

Applications of Game Theory to Policy-Relevant Questions

Francesca Arduini

8th May 2025

Thesis submitted for degree: Economics PhD from UCL

1 Signed Declaration

I, Francesca Arduini, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

The second chapter of this thesis is joint work with Florine Le Henaff. We jointly developed the idea, and worked together throughout the project. Florine's contribution was more weighted on situating our contribution within the environmental literature, preparing the expenditure and emissions data, and comparing our results to others in the environmental literature. My contribution was more focused on situating our work within the household literature, setting out the structural household model, and estimating the structural parameters of the model.

2 Abstract

Game theory is a versatile and powerful tool that can help us shed light on a variety of important, policy-relevant questions. In this thesis, I use game theory to develop novel methodologies to answer three economic questions of substantial relevance to institutions including governments and courts. The first two chapters are grounded in the static collective household model, a widely used cooperative game theoretic model of family decision-making.

In the first chapter, I develop an innovative approach to estimating how families share resources. This allows me to estimate inequality measures, such as poverty rates, at the individual level. By contrast, standard inequality measures, used to guide and evaluate policy, only take into account inequality between households. This can lead to underestimating levels of inequality, masking important dimensions of heterogeneity, and incorrect conclusions about trends of inequality over time and the impact of specific policies. I apply my methodology to data on UK heterosexual working couples, finding that women are more than 20% more likely than men to be in poverty.

My second chapter draws on the same underlying model, but applies it to a novel application in environmental economics. This is the first paper to extend a household bargaining model to a setting with greenhouse gas emissions. We show women's purchasing behaviour is consistent with women preferring more environmentally-friendly goods and services than men. We estimate that, if the average UK heterosexual working couple transitioned to gender-equal bargaining, this would lower household emissions by more than 2%.

The third chapter tackles a question that has been brought up in several antitrust cases, in a variety of jurisdictions. I use non-cooperative game theory to debunk a common misconception that focal pricing (charging only special prices e.g. ending in 9s) reduces pass-through of input cost changes to consumers.

Word count: 300 words (max 300 words)

3 Impact Statement

Both of the first chapters contribute to a growing body of evidence that bargaining power within households is uneven, and that this has important welfare consequences. In the first chapter, I develop a new approach to estimating inequality within households. This is key to more accurately estimating inequality measures, such as poverty rates, to target and evaluate policies more effectively. This literature has the potential to be very impactful, although further work is needed to fully realise its potential, as discussed in the conclusion. The second chapter adds to a growing body of work analysing patterns in emissions from household consumption, with the end-goal of guiding urgent efforts to combat climate change. This literature is still in a state of relative infancy, and, as discussed in the conclusion, improvements in emissions data, and further methodological work, are both important to provide more precise and robust advice to institutional decision-makers in the coming years. Both of these chapters also have important implications for institutional actors considering the cost-benefit analysis of policies aimed at narrowing gender pay gaps. In addition to standardly acknowledged direct benefits on gender inequality, these policies are likely to have further indirect effects through the redistribution of bargaining power from men to women within households. I show this is likely to reduce poverty rates and greenhouse gas emissions from household consumption.

The third chapter was inspired by a widespread misconception about focal pricing, which is the common practice of only charging ‘special’ prices, e.g. ending in 9. It is often claimed that (i) focal pricing leads to little, or no, pass-through of upstream input cost changes to consumers, and (ii) hence that, simply because an industry is characterised by focal pricing, consumers cannot be harmed by anticompetitive behaviour of upstream producers of inputs to that industry. This misconception has been raised in several high-profile cases, and was a key factor in the Lithium Ion Batteries class action being dismissed in the US. I show that, under game theory models standardly used in competition cases, focal pricing has no impact on the expected pass-through rate. This chapter has attracted discussion by competition consultants and lawyers, and is already being cited in expert economic evidence in a UK class action. In this way, the chapter is likely to have a concrete impact on the decision-making of antitrust courts, contributing to bringing abusive firms to justice and obtaining compensation for overcharged consumers.

In addition to the impact of each specific chapter, the research agenda underlying this thesis has a broader indirect impact of supporting general efforts to draw on economic theory to improve institutional decision-making. For instance, by exploring a novel policy-relevant application of household bargaining models, the second chapter hopes to inspire more work

taking these models to other unexplored, and empirically important, contexts. Moreover, by evidencing the usefulness of game theory to antitrust courts, the third chapter may contribute to a long-standing debate about the value of expert economic evidence to courts, facilitating further use of economics in court settings.

Word count: 499 words (max 500 words)

Contents

1	Signed Declaration	2
2	Abstract	3
3	Impact Statement	4
4	Acknowledgements	13
5	Introduction	14
5.1	Economic Theory at Work	14
5.2	Structural Economics	15
5.3	Game Theory	16
5.4	This Thesis	18
5.4.1	First Chapter: Estimating Intra-Household Sharing From Time-Use Data	18
5.4.2	Second Chapter: Female Empowerment and Household Emissions . .	20
5.4.3	Third Chapter: Focal Pricing and Pass-Through	21
5.5	The Broader Research Agenda	22
I	Estimating Intra-Household Sharing from Time-Use Data	24
6	Introduction	24
6.1	Contributions to the literature	26
6.2	Structure of this paper	30
7	The model	30
7.1	Framework	31
7.2	The household's optimisation problem	33
7.3	A problem in two stages	34
8	Methodology	35
8.1	Assignable good: private leisure	36
8.1.1	Advantages of using private leisure	37

8.1.2	Theoretical consistency of using time-use data to estimate sharing . . .	38
8.1.3	Wage endogeneity	39
8.1.4	Restricting the sample to working couples	39
8.2	Parametric specification: Cobb-Douglas	40
8.3	Assumptions restricting preference heterogeneity	42
8.3.1	Identifying assumption: private leisure C-D SAP	42
8.3.2	Additional restriction for ease of estimation: public good C-D SAP . .	44
8.4	Linear approximation of resource shares	45
8.5	Estimation	46
8.5.1	A structural approach to the error term	46
8.5.2	Estimation procedure	46
8.5.3	Intuition for identification	47
8.6	Data requirements	48
9	Estimating resource shares from UK time-use data	49
9.1	UKTUS	49
9.2	Sharing rule estimates	51
9.2.1	Interpreting regression results and estimating marginal effects	51
9.3	Empirical performance of the model	54
9.3.1	Testing the private leisure C-D SAP assumption	55
9.3.2	Testing the public good C-D SAP assumption	56
9.3.3	Homotheticity tests	57
9.3.4	Internal consistency of estimates	58
9.4	Comparison to other approaches	59
9.4.1	Non-market-work hours as the assignable good	59
9.4.2	Clothing as the assignable good	60
9.4.3	Estimates from other papers with UK applications	61
10	Individual-level consumption for UK couples	62
10.1	Methodology	62
10.2	Findings and policy implications	65
10.2.1	Individual-level full consumption (including time-use)	66

10.2.2	Policy implications in the UK setting	67
10.2.3	Gender health gaps in the UK	68
11	Conclusion	69
II	Female Empowerment and Emissions	71
12	Introduction	71
13	Related literature	74
14	Framework	75
14.1	Model	75
14.2	Singles	77
14.3	Couples	78
14.3.1	Bargaining in the household	78
14.3.2	Household-level expenditure functions for couples	79
14.4	Environmental extension	80
15	Methodology	82
15.1	Estimating bargaining weights	82
15.1.1	Estimating bargaining weights from time-use data	82
15.1.2	Estimating bargaining weights in the expenditure dataset	84
15.2	Estimation of emissions propensities	85
15.2.1	Singles	85
15.2.2	Couples	86
16	Data	87
16.1	UK Time-use Survey	87
16.2	LCF	88
16.3	DEFRA emissions multipliers	91
16.3.1	Bridging LCF and DEFRA data	92

17 Findings	94
17.1 Findings for singles	94
17.2 Findings for couples	105
17.2.1 Bargaining weights	105
17.2.2 Environmental preferences	106
17.2.3 Alternative proxies of bargaining power	108
17.2.4 Comparison to other findings in the literature	111
17.3 Discussion of findings	112
18 Conclusion	113
 III Focal Pricing Constraints and Pass-Through of Input Cost Changes	 115
19 Introduction	115
20 A general framework to assess the impact of focal pricing constraints on pass-through	119
20.1 The framework	119
20.2 The Irrelevance Theorem	123
20.3 The distribution of pass-through	124
21 Examples of specific models	125
21.1 Monopolistic competition	126
21.1.1 Linear demand	126
21.1.2 Non-linear demand	127
21.2 Perfect competition	128
21.3 Differentiated price competition	129
21.4 Other models	131
22 Discussion	132
 IV Conclusion	 134

V	References	137
VI	Appendices to all chapters	144
A	Identification in a more general setting	144
A.1	Assignable good	144
A.2	Linear approximation of resource shares	145
A.3	Parametric forms and estimating equations	146
A.4	Identification	146
A.4.1	Identification result	147
A.4.2	Identifying assumption	150
A.4.3	Identification proof	151
A.4.4	Discussion of the identification result	152
A.5	Almost Ideal Demand System	153
B	Time-use data	155
B.1	Definition of private leisure	163
C	Expenditure data and individual-level consumption estimation	167
C.1	The dataset	167
C.2	Categorisation of expenditures into private and public	168
C.2.1	Excluded categories	168
C.2.2	Private vs public consumption	169
C.2.3	Treatment of durables, including homes and vehicles	169
D	Considerations on choosing the assignable good	172
D.1	Availability and accuracy of assignable good data	172
D.2	Credibility of the assignable good being private	173
D.3	Frequency of purchases and magnitude of budget share	173
D.4	Further considerations	174
E	The issue of non-participation	174
E.1	Why excluding couples with non-participation is unlikely to bias estimates	175

F Cobb-Douglas: additional material	176
F.1 Cobb-Douglas Domestic production functions	176
F.2 Normalising assumptions	179
F.3 Public good C-D SAP	179
F.3.1 Under public good C-D SAP, the resource share and Pareto weight coincide	179
F.3.2 Under public good C-D SAP, second-stage demands do not depend on public expenditure	180
F.4 Structural approach to the error term	180
F.5 Individual-level full consumption when expenditure data is separate from time-use data	181
F.6 Back-of-the-envelope check	182
F.6.1 Three example households	183
F.6.2 General result	186
G Estimation and hypothesis testing	189
H Emission multipliers	191
H.1 UKMRIO methodology	191
H.2 UKMRIO good categories	192
I From Household Emissions to Total UK Emissions	194
J A monopolist facing linear demand	195
K A monopolist facing a logarithmic demand function	196

List of Figures

1	Baseline results, UKTUS.	52
2	Estimated female shares using different assignable goods	59
3	Individual-level consumption	64
4	Average Emission factors, kgCo2eq/£, 2001-2014	93
5	Budget and carbon footprint shares by good categories.	94

6	Gender difference in singles' emissions propensities on overall consumption: Women's - Men's.	95
7	Gender difference in singles' emissions propensities on transport: Women's - Men's.	100
8	Gender difference in singles' emissions propensities on food: Women's - Men's.	100
9	Gender difference in singles' emissions propensities on clothing: Women's - Men's.	103
10	Gender difference in singles' emissions propensities on energy: Women's - Men's.	103
11	Interior solution vs. corner solution for time-use	175

List of Tables

1	Summary statistics: UKTUS and LCF	51
2	Regression results, baseline specification, UKTUS.	53
3	Interpretation of regression coefficients	54
4	Female resource shares vary with wages (left) and assignable good choice (right)	54
5	Summary statistics for LCF main sample	90
6	Summary statistics for LCF regression sample	91
7	Emissions Propensities and Carbon Intensities for singles: overall consumption	97
8	Emissions Propensities and Carbon Intensities for singles: Food	99
9	Emissions Propensities and Carbon Intensities for singles: Transport	101
10	Emissions Propensities and Carbon Intensities for singles: Energy	102
11	Emissions Propensities and Carbon Intensities for singles: Clothing	104
12	Emissions Propensities and Carbon Intensities, couples	105
13	Proxy for bargaining power with gender of the breadwinner	109
14	Proxy for bargaining power with hourly wage ratios	110
15	Proxy for bargaining power with weekly pay ratios	111
18	Expenditures categories included in the UKMRIO (COICOP-level)	192

4 Acknowledgements

I have been very lucky to be supported throughout my life by capable and caring instructors and loving and open-minded family and friends.

I am thankful to all the teachers, professors, colleagues and mentors who have helped me in my academic and professional journey, which started with ISM, then Oxford, Oxera, UCL, IFS and Yale. For my progress in the PhD, I benefited from the generous help and thoughtful advice of many colleagues and mentors at UCL, IFS, where I am affiliated as a PhD Scholar, and Oxford, where I completed my previous studies, and have taught at since. For this thesis, I am particularly indebted to my UCL PhD supervisors, Lorraine Dearden and Alex Bryson, whose unwavering belief in my abilities, and commitment to helping me, has been a much-needed source of strength to overcome the various challenges that I encountered over these past years. I also wish to thank Fabien Postel-Vinay and Valérie Lechene, who were part of my broader supervisory team, and my upgrade examiners, Nikki Shure and Hedvig Horvath, for their insightful advice. My PhD further benefited greatly from my exchange with Yale University, and the continued support of my Yale mentors, Costas Meghir and Cormac O’Dea, even once I returned to London. I thank Richard Blundell and Abi Adams for agreeing to be my PhD examiners, and look forward to discussing my work with them and incorporating any suggestions they may have.

I am deeply grateful to have family and friends who always spontaneously embrace my dreams, help me follow them through, and rejoice in every step forward with me. I would not have been able to complete this PhD in less than four years, while also becoming a mother, had it not been for the incredible support of my husband Tom, his love, and his effortless feminism. I owe an even deeper debt to my parents, Daniela and Ferdinando, and my brother, Stefano, who brought me up to care about things, put work into understanding them, and try with determination to make them better. I dedicate this thesis to my baby, William, who was born during it, and is my joy and inspiration.

Finally, I gratefully acknowledge studentship funding from the Economic and Social Research Council including an Advanced Quantitative Methods Stipend [grant number ES/P000592/1].

5 Introduction

5.1 Economic Theory at Work

It is fashionable in some circles to object to policymakers making use of the insights of experts, particularly economists. However, when we consider the complexity of most policy decisions, it becomes hard to see how these choices could be taken in an informed manner without input from specialist work. At the same time, it is important to acknowledge that the public dislike of expert economists stems partly from real faults, such as the propensity of the profession to quantitatively analyse situations without first attempting to understand them on a qualitative level.

After my MPhil in Economics at Oxford, I stepped out of academia for three years, working as an economic consultant with Oxera, to understand how economic ideas are put to practice on a daily basis outside of academic circles. I saw first-hand how economic analysis is put to use fruitfully in a variety of contexts, such as comparing the potential costs and benefits of potential policies, quantifying the damages from anti-competitive firm behaviour, and regulating prices in key industries. There is no real alternative to using economics in these contexts, so if we are not satisfied with the results we can either try to improve the underlying economic approaches, or work harder to effectively bridge academic economics and its applications. I started my PhD motivated by the fact that there are several aspects of institutional decision-making that could greatly benefit from more policy-relevant academic work, and hoping to contribute to that effort.

An example that motivated me is the misconception which I debunk in the third chapter of this thesis. I saw how huge antitrust cases, worth vast amounts of money and with many professionals working on them, can be dismissed based on an intuitive argument, not based in economic theory or evidence, even though this argument can be shown to be fallacious with reasonably simple economic analysis. It is a failure of academic economics if we are not able to fill in gaps in the literature that have important real-world relevance.

I was also inspired by some big policy questions of our time, such as how to re-design parental leave policies to support gender equality in workplaces and in the family. This question is so complex, and it is so challenging to obtain high-quality data pertaining to it, that it ended up not making it in this thesis, even though I continue to work on that project. In fact, the chapters in this thesis partly constitute building blocks towards that question, since they tackle questions about the decisions of families, and the decisions of firms, separately - one of the important next steps will be to combine those, as well as make further modifications to the models.

5.2 Structural Economics

There is some debate amongst economists about the relative merits of taking a structural, or a reduced-form, approach to economic analysis. A structural approach always starts with a model of the relevant situation, and derives from that model-driven equations which can be taken to the data to estimate model parameters. Once these parameters have been estimated, it is generally possible to use the model for a wide variety of purposes, including estimating the potential impact of yet-unobserved policies as well as of policies that have actually been implemented. These policy counterfactuals are deeply valuable in terms of advising policymakers because they enable us to make predictions without having to test out all the options we are considering. However, these predictions are only as good as the underlying model, and proponents of reduced-form analysis prefer to reduce reliance on specific models and increase focus on flexible, data-driven estimation. Moreover, estimation of structural models is often very complex, making it almost impenetrable to policy-makers, practitioners, and even academics outside of that specific field. At the same time, purely reduced-form approaches can end up being black boxes too. This can happen when extremely flexible functional forms, which are very hard to interpret, are estimated from the data. Moreover, any estimation approach can be mapped back to a set of theoretical models consistent with it, so there is no such thing as a truly model-free approach. In fact, not specifying a model can result in the assumptions being made in a less intentional and transparent manner, and therefore also less sense-checked and empirically tested.

At a first glance, this thesis sits squarely within the structural approach as it aims to take game theory models to applied policy questions. However, in reality the approach taken in each of the chapters attempts to reconcile the structural and reduced form approaches in two senses. Firstly, I make a real effort to avoid relying on overly restrictive assumptions. For instance, relative to similar papers in the literature, my first chapter makes more of an effort to test model assumptions, and the identifying assumption. There are limits to what can be done here - I cannot directly test my identifying assumption with the available data - but I have plans for future work addressing this further. Secondly, while my estimation approaches are grounded in economic theory, I strive to make them as simple as possible. In both of my first chapters, I only run simple linear regressions. In order to do this, even though I start with a very general model of household bargaining in the first chapter, I focus on a very simple version for the empirical application in both the first and second chapter. In the third chapter, I make my argument conceptually, and provide a proof for some simple models of competition between firms. In that case, I don't perform any empirical estimation, because the relevant argument can be made effectively at a conceptual level.

A related issue is a core question from the philosophy of science: how detailed should our models be? There is an inevitable tension between making a model as simple as possible, capturing only the essence of the issue at hand, and making it as accurate as possible, fully taking reality into account. A model that is too far removed from the world is going to yield incorrect predictions. At the same time, a model which is too detailed will not be tractable, and will not serve any practical purpose. Economics, like other sciences, must strive for the right balance. In all three chapters of this thesis, I strive to use the simplest models that capture the key features of the situations of interest. In the case of the first two chapters, additional complexity of the model would potentially be beneficial to accuracy, but would likely detract from the possibility to estimate the model in a simple and transparent way, from widely available data. In the case of the third chapter, since the argument is merely that we cannot in general assume no pass-through when there is focal pricing, a more detailed model would not particularly aid the discussion.

5.3 Game Theory

Game theory is "the methodology of using mathematical tools to model and analyze situations of interactive decision making. These are situations involving several decision makers (called players) with different goals, in which the decision of each affects the outcome for all the decision makers" (Maschler et al. (2013), p.xxiii). A very substantial proportion of the subjects studied in economics, and also in other disciplines, involve strategic interactions. Game theory has been effectively applied from ethics to military issues, and from political science to evolutionary biology.

Several important branches of economics are grounded in game theory. For instance, industrial organisation, the study of firm behaviour and competition, is deeply reliant on game theoretic models. In any real-life market, the profit of a firm depends both on its behaviour and that of its competitors, making it a strategic situation. Insights and empirical tools from industrial organisation play a key role in a variety of contexts, such as the regulation of mergers, the assessments of abusive behaviour and the estimation of follow-on damages in antitrust cases. The third chapter of this thesis draws on models of competition between firms.

Auction theory is a branch of game theory that has particularly clear applications to policy. This includes the allocation of organs to potential recipients (e.g. Agarwal et al. (2021)), of food donations to different food banks (e.g. Prendergast (2022)), and of public housing to those in need (e.g. Waldinger (2021)). Auctions have also been used, sometimes in disastrous ways and sometimes in spectacularly successful ways, to sell government property.

For instance, carefully designed 3G spectrum auctions in the UK raised £22.5 billion for the Exchequer, vastly more than comparable sales in other European countries which were either allocated administratively or had poorly designed auction set-ups (see Klemperer (2004) for a detailed discussion of this case study and a broader discussion of auction design).

Auction design can help ensure that lots are assigned to whoever values them most and hence that allocation is efficient. Moreover, they can raise vastly more money for governments than alternative arrangements, providing substantial benefits to taxpayers. Where auctions are intended to sell spots in a market, as happens with spectrum auctions, careful auction design can also encourage bidding from more players, supporting the formation of more competitive markets thereafter (more providers to choose from). The importance of auction theory has been recognised widely, for example the 2020 Nobel prize in Economics was awarded to Paul Milgrom and Robert Wilson for their contributions to auction theory, which benefited sellers, buyers and taxpayers around the world. Many of the key insights from auction theory, such as the winner's curse (if you win, you probably overpaid, so bid less than you think you should) can be widely applied outside of large formal auction contexts, e.g. when a small number of people make offers to purchase a home through an estate agent.

The game theoretic approach is also widely considered to be the best existing tool to formalise bargaining contexts (see Muthoo (1999)), which are relevant to a range of situations. For example, models of wage bargaining between employers and employees are at the heart of much of labour economics. These models can help us understand how to best achieve important goals such as reducing unemployment, improving the match quality between employers and employees, and increasing the well-being of more vulnerable workers. For instance, recent work by Coviello et al. (2024) finds that female workers benefit less than male workers from a minimum wage increase because they have a weaker bargaining position with employers, due to worse outside options. While men and women receive similar wage increases, women increase their effort levels more than men to reduce the risk of losing their job.

Bargaining theory is also central to household economics, the study of how family units form, dissolve and make decisions. Moreover, it can help us understand how these patterns are influenced by institutional choices, such as divorce law (see e.g. Voena (2015)). While it may seem unnatural to attempt to analyse marriage, divorce and family decisions with economic models, it is important for us to formalise our understanding of multi-person households. While game theory may seem a crude tool in this context, it is the best one available to us, and in fact is a particularly apt tool in settings with few players, where the actions of other players matter substantially to each of them (Fudenberg and Tirole (1991)).

We may worry that game theory, especially when applied in contexts such as household decision-making, makes excessively strong assumptions about the rationality and cognitive

processing power of individuals. It is one thing to assume that large multi-nationals behave rationally, and another to assume that spouses do so. This is an important critique of game theory, and two branches of game theory have developed in response to it. The first is behavioural game theory, which allows for more nuance in preferences and decision-making than standard game theory. For example, the level-k model considers different players as being able to perform k levels of iterated logic when choosing their strategy, rather than the infinite rounds of common knowledge of rationality required to reach Nash equilibrium in some games, including beauty contests.

Evolutionary game theory, which in some ways is closer to biology than other branches of economics, also allows us to depart from full-rationality assumptions, and incorporate cognitive limitations. Evolutionary game theory models are typically characterised by inertia in decision-making, players best responding in a myopic way (to historic behaviour by their ‘neighbours’, rather than forward-looking full optimisation), and some probability of random mutations, or errors. Evolutionary game theory can provide interesting insights in many contexts, such as the evolution of norms, and patterns of residential racial segregation. Interestingly, many classic outcomes of standard game theory can be rationalised in these models through learning over time (for more, see Young (2001)).

5.4 This Thesis

As set out below, in each chapter of this thesis I employ approaches from game theory to tackle policy-relevant questions. As discussed above, I attempt to keep the models as simple as possible, while capturing key features of the situation at hand.

5.4.1 First Chapter: Estimating Intra-Household Sharing From Time-Use Data

An important question for policy-makers is how to measure inequality. These measures are widely used to target and evaluate policies, e.g. to track progress against the UN’s Sustainable Development Goals. However, standard measures, such as the World Bank’s poverty rates, only take into account inequality between households, assuming that resources are shared equally within households. This is because expenditure surveys, which are used to estimate inequality, are taken at the household level. Collecting precise individual-level data is very resource intensive, and hence not an appealing way forward. In order to estimate individual-level inequality from the available household-level data, we require a model to take to that data, to help us make use of the information we have to answer the question of how resources are shared within households.

Because we are trying to understand how households split resources amongst themselves, the most natural model to apply is a bargaining model. The way families make decisions depends on the preferences, and the relative bargaining power, of all members. Bargaining power depends on what would happen in the event of disagreement, e.g. non-cooperation within the household, or dissolution of the household. The implicit assumption behind standard methods is that all household members have equal bargaining power, but this is too unrealistic. A very general, but also tractable, model that is used in the household literature to answer questions of this form is the static collective household model. This does not pin down a specific bargaining solution, and only requires that decisions are made efficiently in the household. The static version of this model, which is much simpler than the dynamic version, takes households to be already formed, with specific bargaining weights (which are a function of household characteristics), and making decisions at a specific point in time without consideration of inter-temporal issues such as savings and durable goods. The static and efficiency assumptions are not innocuous, but they are not too problematic when we are interested in estimating inequality within households at a specific point in time, in a specific country, based on their patterns of consumption. Note that these assumptions would be out of place if we were investigating issues where inter-temporal dynamics are highly relevant, e.g. investments in labour market skills.

In this chapter, I make two main contributions to this literature. Firstly, I develop a new approach which is both grounded in a general model of household decision-making, and simple to estimate by OLS with widely available data. Compared to existing approaches which are similarly simple to implement (Dunbar et al. (2013)), my methodology allows for public, as well as private, consumption, and carefully incorporates time-use as well as material consumption. While both public consumption and time-use create some challenges for identification, they are crucial components of household decision-making, and I argue that incorporating them is important for estimate accuracy.

My second contribution is using time-use data to identify sharing. The literature to date uses household expenditure surveys both to estimate the sharing rule and to apply that to observed household-level expenditure. However, expenditure surveys contain very limited individual-level variation, and it is not always clear that the available variation satisfies identification requirements. Instead, I propose splitting this exercise into two phases, and show that the first step can be conducted using time-use data. This data contains individual-level variation which, I argue, satisfies identification requirements more credibly than alternatives from expenditure surveys used in the literature to date.

I apply my methodology to UK working couples. First, I estimate the household sharing rule from time-use data, finding that women have fewer resources than men on average.

The sharing rule varies at the household level as a function of household characteristics, including the wages and educational attainment of members. These characteristics drive substantial heterogeneity in my data, in a manner which is consistent with bargaining theory, e.g. women with higher wages have higher resource shares. I then apply my estimated sharing rule to a separate but comparable expenditure dataset in order to estimate individual-level consumption. I find that, on average, men consume 8.53% more than women, and that the poverty rate is 20.59% higher for women than men.

These gaps, which would go undetected by standard approaches to measure consumption inequality, provide an important rationale for UK policy-makers to target help to women in low-income households. I suggest potential policy responses, including direct transfers, subsidising childcare, and narrowing gender pay gaps.

The substantial degree of intra-household inequality I estimate for UK households is indicative that we are likely to find even larger gaps in other contexts. The UK is one of the most gender equal countries in the world, and the literature suggests that working couples are characterised by less intra-household inequality than households with cohabiting children and non-participation (e.g. Bargain, Donni and Hentati (2022)). Hence, we can think of my results almost as a lower bound for the magnitude of inequality we are likely to find more generally, reinforcing the broader importance of estimating intra-household inequality to inform policy-making.

5.4.2 Second Chapter: Female Empowerment and Household Emissions

The second chapter of this thesis is joint work with Florine Le Henaff, a PhD candidate at ECARES. She is an environmental economist, and this project draws both on her expertise in that area, and my work on household bargaining from my first chapter.

It is widely acknowledged that we need to find ways of reducing our greenhouse gas emissions, and that household consumption patterns are indirectly responsible for a very substantial component of those emissions. The environmental literature has started exploring how consumption patterns relate to emissions, and whether there are important dimensions of heterogeneity between households. One understudied area is the relationship between gender and emissions. Evidence from surveys and text analysis of social media suggests that women have more pro-environmental attitudes than men (e.g. De Rock and Le Henaff (2023)). However, more work is needed to link this to concrete gender differences in patterns of greenhouse gas emissions.

The second chapter of this thesis contributes to filling that gap in the literature. The main empirical challenge is similar to the one tackled in the first chapter of this thesis: dis-

entangling gender differences in emissions patterns from household-level expenditure data requires a structural approach to understand the degree to which decisions are influenced by the preferences of men and women in the household. While for singles we could take a more reduced-form approach to estimating gender differences, it is important to also develop approaches to separately analyse multi-person households, since there is evidence that preferences are not always stable across household compositions (see e.g. Hubner (2020)). To this end, an appropriate model is the household bargaining model used in the first chapter, with an extension to incorporate differential emissions associated with consumption of different goods. This yields structural equations for emissions as a function of the preference parameters of men and women, and their relative bargaining weights. To our knowledge, this chapter is the first to employ a household bargaining model to answer a question in environmental economics.

Using a similar approach to the one developed in my first chapter, we estimate bargaining weights as a function of household characteristics from time-use data. We then apply the estimated bargaining rule to expenditure data, so we obtain household-specific bargaining weights. We also merge in data with granular estimates of the emissions associated with consumption of different goods. This allows us to analyse emissions patterns for households as a function of the bargaining weight of women, and hence estimate gender differences in emissions propensities. We find that women have lower emissions propensities than men, and therefore greater female empowerment is associated with lower household emissions intensity. If the average UK heterosexual working couple transitioned to gender-equal bargaining, its emissions would fall by 2.1%. Our findings suggest that policies aimed at increasing female bargaining power, for instance by narrowing gender pay gaps, may have the additional benefit of reducing household emissions.

5.4.3 Third Chapter: Focal Pricing and Pass-Through

My third chapter was inspired by an observation I made during my years working as an economic consultant. There is a common misconception that the adoption of focal pricing practices (discrete pricing e.g. with prices ending in 9) leads to little, or no, pass-through of upstream input cost changes. The argument appears intuitive: if firms charge focal prices, then a small change in input costs will not lead to a change in the charged focal price. However, equally intuitively, this fails to recognise the other side of the coin: sometimes a small change in input costs can lead to a big jump from one focal price to the next. Hence, it is not at all evident that the presence of focal pricing generally reduces pass-through.

However, this misconception has been raised in several high-profile cases, and for instance

was a key factor in the Lithium Ion Batteries class action being dismissed by the Court. My third chapter aims to formalise the intuition that focal pricing needn't affect average focal pricing, and therefore that follow-on antitrust damages cases should proceed as usual in the presence of focal pricing, i.e. with industry-specific empirical estimation of pass-through. Firms' pricing decisions are strategic interactions, best modelled using a game theoretic approach. In order to achieve my goal, it seemed appropriate to make use of simple models of competitions that are widely used as benchmarks in court: monopoly and perfect competition (the outcomes of which can also be replicated under oligopolistic competition with collusion, for monopoly, and undifferentiated product price competition, for perfect competition). In this way, my model shows that the assertion that focal pricing leads to no pass-through cannot be taken to be a general conceptual truth. This is not to say that this could not happen under specific circumstances, but case-specific empirical estimation is likely to be the best way of judging the degree of pass-through, regardless of the use of focal pricing.

I show that the adoption of focal pricing has no impact on the expected pass-through rate, but it does generally lead to more heterogeneous distribution of pass-through. I also discuss conceptually some other relevant factors, such as the whether firms are able to adjust non-price product characteristics, such as quality, on a near-continuous spectrum. If so, when we consider hedonic pricing, taking other characteristics into account as well as prices, we no longer have focal pricing. In practice, this is common, and may render the original claim about focal pricing reducing pass-through moot in many contexts.

This chapter has attracted discussion by competition practitioners including economic consultants and lawyers, and is already being drawn upon in antitrust cases in multiple jurisdictions. this chapter may contribute to bringing abusive firms to justice and obtaining compensation for overcharged consumers.

5.5 The Broader Research Agenda

This thesis is the start of a broader research agenda on policy-relevant applications of game theoretic models.

I have started working on two projects extending my first chapter. One paper is an extension to more general household compositions, including children and people who do not do paid work. The other is a project directly testing assumptions of different approaches, and evaluating their predictions against both data collected specifically for this project and simulated data. This will be an important step forward towards the goal outlined above of working towards strengthening the structural approach taken in the first paper with more data-driven

testing of the model, and potentially modifications to be more consistent with the empirics. Amongst other things, we will also test whether game theoretic approaches other than the collective model (e.g. specific, potentially inefficient, bargaining models) are a better fit. This exercise would be an important step towards consolidating the multi-faceted literature on estimating intra-household sharing into something like a mainstream approach with sufficient credibility to be employed by policymakers for practical decision-making, replacing currently standard approaches which assume equal sharing within the household.

Working on the first two chapters of this thesis allowed me to deeply understand the static collective model of household bargaining. Combining this with my background in competition economics, and my understanding of industrial organisation models such as those used in my third chapter, is going to be an important next step in this agenda. Questions such as how to optimally design parental leave require understanding both the relevant family decision-making processes and the relevant aspects of employer-employee relations.

Reforming parental leave has the potential to drive profound change in the intra-household allocation of work and childcare responsibilities (see e.g. Bünning (2015)). This could unlock welfare improvements for households; narrow the gender pay gap; increase fertility rates in ageing societies; and increase workforce participation in tight labour markets, promoting macroeconomic growth. In the recent UK election, most parties mentioned parental leave reform in their manifestos, and the Labour government has committed to reviewing the system. My planned work in this area will contribute to a nascent literature on the impact of parental leave policies on pay gaps (e.g. Johnsen et al. (2024)), and a growing literature on the motivation for employers to provide non-wage benefits, including enhanced parental leave (e.g. Goldin et al. (2020)).

To tackle the question of how to optimally design parental leave, it may also become important to model dynamics and consider carefully whether the efficiency assumption of the collective model is fit-for-purpose or requires some degree of relaxation. Potentially, non-cooperative models of bargaining may be more apt in that context. Going forward, the work I have done for this thesis will form the basis of this continued effort to find the most suitable type of game theoretic approach to apply to different policy-relevant questions and take the model to the data in as simple and transparent a way as possible, while also permitting rich analysis, including evaluating the impact of policy counterfactuals.

Part I

Estimating Intra-Household Sharing from Time-Use Data

6 Introduction

It is empirically challenging to measure individual-level consumption inequality because expenditure data is typically collected at the household level. The simplest solution to this problem, which is still widely used, for instance by the World Bank, is to assume that there is no intra-household inequality, so that household-level expenditure, divided by the number of household members, directly provides estimates of individual-level consumption (the per capita approach). A growing body of evidence suggests that the equal sharing assumption is unrealistic and that it is crucial to account for intra-household, as well as inter-household, inequality (e.g. Lechene et al. (2022)). Not doing so leads to inaccurate, generally downward-biased, estimates of the aggregate level of consumption inequality, such as the poverty rate. It also fails to appropriately capture key dimensions of inequality, such as gender and age, and their intersection (e.g. see Calvi (2020)). Therefore, standard measures may lead to incorrect conclusions on trends of inequality over time, and about the impact of specific policies. Moving towards more accurate individual-level estimates of inequality is important to improve the targeting of public funds to prioritise more effective policies, and policies aimed at groups that are particularly in need.

I develop a new approach to estimating intra-household sharing, which is both grounded in a general collective household model, and simple to implement with widely available data. My approach incorporates both private and public goods, and both material goods and detailed time-use. Estimation proceeds by OLS from a small number of estimating equations. Instead of requiring the whole expenditure system at the individual level, I need only individual-level expenditure data on a single private good (the ‘assignable good’), which is much more commonly available. The responsiveness of assignable good expenditure to household budget is informative about how that budget is shared between members. Together with an identifying assumption restricting preference heterogeneity, this yields point-identification of the sharing rule i.e. household-specific estimates of sharing, as a function of household characteristics. Applying sharing estimates to household expenditure surveys enables estimation of individual-level consumption and inequality.

In the collective household model, the household agrees on time-use and material consumption in the same bargaining problem, and hence the same sharing rule applies both to time-use and material goods. Therefore, sharing can in principle be estimated just as well with a material or a time-use assignable good. This paper is the first to identify sharing from time-use data, and uses a novel assignable good: private leisure. This is time spent on leisure activities without other household members co-present, for instance reading a book alone or having a coffee with a friend. This can be measured accurately from time-use data combining activity and co-presence information. I provide conceptual arguments and empirical tests that suggest that this new source of identifying variation may satisfy the required assumptions more credibly than alternatives, contributing to more accurate estimates of individual-level consumption inequality.

I adapt the identification result of Dunbar et al. (2013) to a setting with time-use, price variation and public goods. The first two extensions are required in order to use my novel assignable good. The third extension is required for theoretical consistency with my proposed approach to applying estimated resource shares to estimating individual-level consumption, where the distinction between private and public goods is important to avoid overestimating inequality. To make these three extensions tractable, my implementation uses Cobb-Douglas expenditure functions, rather than Almost Ideal Demand System Engel curves, as estimating equations.¹ I estimate sharing from time-use data alone, using the UK Time-Use Survey (UKTUS). Women generally command a lower proportion of resources than men,² 45% on average, but this varies substantially between households. Characteristics affect resource shares in a manner consistent with bargaining theory, e.g. women with higher wages, or matched with less educated men, command a higher resource share.

I then apply the sharing rule to a separate but comparable expenditure dataset, the Living Costs and Food Survey (LCF), allowing me to estimate individual-level material consumption. I define this as the monetary market value of an individual’s consumption, i.e. the sum of (i) the household’s expenditure on public goods, and (ii) the individual’s estimated share of the household’s expenditure on private goods. This is a useful metric to calculate objective and policy-relevant measures of individual-level consumption inequality. I find that, on average, men’s material consumption is 8.53% higher than women’s, with a wider gap for poorer households, so that the poverty rate is 20.59% higher for women than for men. These estimates suggest that policymakers should target policies to reduce female poverty in the UK, as further discussed in section 10.2. When I incorporate time-use as well as material

¹It is possible to use my approach with other functional forms, including Almost Ideal Demand System Engel curves, as discussed in online appendix A.

²In this paper I use ‘resource share’ and ‘share of resources’ as a short-hand for ‘conditional resource share’ or ‘share of household private expenditure’.

consumption, I estimate a wider average gender gap in ‘full’ consumption of 10.1%. This supports the view that in order to fully estimate gender inequality we must take into account time-use as well as material expenditure.

These gender gaps are very substantial, especially when considering the context. The UK is one of the most gender equal countries in the world³ and previous literature finds that women command higher resource shares in working couples than in households with children or where the woman does not work (e.g. Bargain, Donni and Hentati (2022)). Therefore, we can think of this as an approximate lower bound for the kind of intra-household inequality we are likely to encounter in other countries and with broader household compositions. This reinforces previous findings in the literature that equal sharing does not hold, and that it is important to estimate intra-household inequality instead of using a per capita approach.

Finally, I note that my results also shed light onto the apparently puzzling fact that women with higher wages tend to have less leisure, even though they have more bargaining power in the household. These facts are squared by noticing that *expenditure* on women’s leisure is higher for women with higher wages, even if their *quantity* of leisure is somewhat lower; and also that expenditure on material consumption increases in a way that compensates for lower leisure. Importantly, we can account for all this with a collective household model, with a single sharing rule which applies both to material goods and time-use, because changes in wages have an income effect (both through the household budget and through bargaining) as well as a price effect (on the price of time).

6.1 Contributions to the literature

This paper builds on a large and growing literature on estimating intra-household sharing. A small number of papers estimate sharing directly from data on individual-level expenditure e.g. Cherchye, De Rock and Vermeulen (2012). However, this type of data is rarely available in practice and most of the literature focuses on estimating individual-level consumption from household-level expenditure. Several papers have shown that this is theoretically possible under different versions of the collective household model, and with different restrictions (e.g. Chiappori (1992), Browning, Bourguignon et al. (1994), Blundell, Chiappori and Meghir (2005), Blundell, Chiappori, Magnac and Meghir (2007) and Chiappori and Ekeland (2009)). A variety of papers have estimated intra-household sharing in different countries, both developing and developed, based on a range of approaches, with different underlying versions of the model, identifying assumptions, and estimation approaches. For example, Cherchye,

³United Nations Development Programme. (2022). Human development report 2021-22. <http://report.hdr.undp.org>

De Rock, Lewbel and Vermeulen (2015) draw on revealed preference techniques to estimate bounds on sharing. Other papers, including Browning, Chiappori and Lewbel (2013) (henceforth BCL), Lewbel and Pendakur (2008), Bargain and Donni (2012) and Brown et al. (2021) point-identify the sharing rule from estimates of Engel curves of multiple goods combined with identifying assumptions restricting preference heterogeneity. Dunbar et al. (2013) (henceforth DLP) shows that, in a setting with non-public material consumption and no price variation, it is possible to identify the sharing rule from individual-level data for a single private good, known as the ‘assignable good’. A similar approach is followed by Bargain, Donni and Hentati (2022) and Lechene et al. (2022) (henceforth LPW). Lise and Seitz (2011) also estimates sharing from data on a single assignable good, but based on a different identifying assumption, restricting the bargaining process so that men and women with equal earnings potential share resources equally. Most papers in the literature find substantial inequality in intra-household sharing, with women and children typically receiving fewer resources than men.

While several approaches in this literature involve complex estimation, recent development have moved towards simpler methods. In particular, LPW shows that, building on the identification result of DLP, it is possible to estimate resource shares by linear regression. I adapt the identification approach of DLP, and the linear estimation framework of LPW, to a setting with time-use, price variation and public goods. In this sense, this paper attempts to combine the tractability of DLP and LPW and the generality of a model with time-use and public goods such as in Lise and Seitz (2011).

Many papers in this literature, including BCL, DLP and LPW, treat all goods as non-public and model household economies of scale through the shareable goods framework. This framework is very tractable, but it does not allow us to impose the restriction which is at the core of public consumption, i.e. the requirement that all household members consume the same quantity of a public good (see Browning, Chiappori and Lewbel (2013) and Chiappori, Meghir and Okuyama (2024)). For this reason, household optimisation under shareable goods and public goods generally yields different solutions. Since households spend a very large proportion of their budgets on public goods, e.g. housing, it is important to accurately capture economies of scale driven by public consumption. This aspect is particularly important when it comes to applying the estimated sharing rule to estimate individual-level consumption and inequality. The sharing rule should apply only to private expenditures since public goods are, by definition, non-rivalrous and consumed by all members in the same quantity. If we incorrectly treat public expenditure as if it were private (as done e.g. by LPW), or exclude it, we substantially over-estimate intra-household inequality. For theoretical consistency, I include public consumption in my model and derive estimating equations

for the sharing rule that are consistent with the presence of public goods. In principle this inclusion can make a substantial difference at the stage of resource share estimation. However, in my application, my use of separable preferences and of a further assumption, restricting preference heterogeneity on the private-public good split, entail that the public good extension only really bites in section 10, where I estimate individual-level consumption by applying the estimated sharing rule only to private expenditure, and assigning full public expenditure to each household member. As discussed in section 10.2, the inclusion of public goods matters substantially, even when it only bites in this stage of the analysis.

Time-use, similarly to public consumption, generates challenges for identification. For this reason, many papers in this literature focus on material consumption. However, as argued by Becker (1965), time-use is a crucial component of household decision-making and well-being. My extension to time-use allows me to estimate ‘full’ individual-level inequality, inclusive of time-use as well as material consumption. This paper is the first to provide such estimates. Additionally, recognising the endogeneity of time-use is important to correctly defining the household budget as the full budget (how much the household could earn if all members worked all the time, in addition to any unearned income), rather than realised earnings, often used in other papers in this literature, and which are endogenous. Finally, incorporating time-use opens up an alternative source of identifying variation to the standardly used expenditure data on clothing.

The literature to date has been constrained in its choice of individual-level variation to identify the sharing rule from. This is because papers to date use household expenditure surveys both to (i) estimate the sharing rule, and (ii) apply the estimated household-specific shares to household expenditure, to obtain individual-level expenditure. Household expenditure data contains very little individual-level variation, and for this reason most papers in the literature have used clothing as the assignable good, as it is often available split into men’s, women’s and children’s clothing. However, as discussed in more detail in section 8.1, clothing has limitations, including a substantial proportion of zeros in the data, which reduces accuracy and theoretically may bias estimates towards equal sharing. As alternatives to clothing, a small number of papers has attempted using individual-level food expenditure or the residual from usual hours worked (non-market-work) as the assignable good. The former is rarely available in practice. The latter can often be obtained from recall data on usual hours worked in expenditure surveys. As discussed in more detail in section 8.1, non-market-work can be thought of as a proxy for leisure, but is conceptually not a private good because it includes domestic work and joint leisure (time spent on leisure activities with other household members co-present),⁴ conceptually biasing estimated towards equal

⁴The distinction between private and joint leisure is an extension of the distinction between individual

sharing.

I propose splitting the estimation of individual-level consumption into two phases. The first, estimating the sharing rule, can be conducted with time-use data. In the second phase, the estimated sharing rule can then be applied to a separate but comparable household expenditure survey, to obtain individual-level consumption. This is advantageous because time-use data (i) is widely collected and recorded at the individual level, and (ii) typically includes sufficiently detailed activity and co-presence information to enable careful definition of the assignable good in a way that is both credibly private and has a large enough budget share, with few or no zeros. I propose using private leisure as the assignable good, and providing conceptual reasons why this may lead to more accurate estimates in some contexts, as appears to be borne out in my application to UK data (see section 9.4).

The main limitation of private leisure is that its price is only known for those who work, for whom it coincides with the observed wage. For non-participants, the price of time is higher than their (unobserved) wage and we would have to estimate that in order to use time-use data to identify sharing. That extension is left for future work: in my application to UK data, I focus on working couples. Focusing on this sub-sample also allows for more credible testing of assumptions using singles data, as discussed in section 8.1.4. In section 8.1.4, I also explain why this sample restriction is not problematic in the UK context.

My paper is the first to use private leisure as the assignable good. This is partly because several papers in this literature, e.g. DLP and LPW, cannot use private leisure as the assignable good because their identification result precludes price variation, and the price of time varies at the individual level. I adapt the identification approach in DLP to a context with price variation. Other papers in the literature have shown identification is possible with price variation, but they either do not allow for point identification of the levels of the sharing rule (e.g. Blundell, Chiappori, Magnac and Meghir (2007), which additionally does not allow for public goods) or require more complex estimation than this paper (e.g. Lise and Seitz (2011)).

My paper is closely related to two papers which also focus on the UK context: Bargain, Donni and Hentati (2022) and Lise and Seitz (2011).⁵ In section 9.4, I compare my results

and spousal leisure in Fong and Zhang (2001).

⁵My paper builds on Lise and Seitz (2011) in a few dimensions. Firstly, I add detailed time-use to the model, use a more credibly private assignable good (private leisure instead of non-market work) and estimate the sharing rule from high-quality time-use diary data, instead of expenditure data with recall question on hours worked. Secondly, I drop the ‘symmetry assumption’ in Lise and Seitz (2011), that men and women with the same hourly wage share household full income equally, and instead use an identifying assumption restricting preference heterogeneity. I also show that it is possible to estimate sharing from time-use data alone. Finally, I adopt the simpler linear estimation approach pioneered by LPW. Relative to this paper, Lise and Seitz (2011) incorporates the complexity of an income taxation system based on household earnings,

to those from these papers: my findings can be reconciled with those from Lise and Seitz (2011), but are qualitatively different from those of Bargain, Donni and Hentati (2022). A systematic comparison of different approaches is left to future work, but in section 9.4 and section 10 I illustrate how my estimates vary in response to varying the assignable good, the identifying assumption, and how public consumption is incorporated.

As well as contributing to the literature on estimating intra-household sharing, I contribute to the literature on measuring individual-level consumption inequality. I define individual-level consumption similarly to Lise and Seitz (2011), but build on existing approaches in two directions. Firstly, I propose a methodology that carefully takes into account (i) durable goods as well as consumables, (ii) goods with prices that vary substantially at a regional level, and (iii) partly public goods, as well as fully private and public goods. This allows me to incorporate large expenditures which are often excluded, such as housing and cars. Additionally, this is the first paper to set out an approach to including time-use as well as material consumption, to consider ‘full’ consumption.

6.2 Structure of this paper

This paper is organised as follows. Section 7 sets out the structural model of the household. In section 8 I propose a new methodology to estimate intra-household sharing. Section 9 discusses the estimation of the sharing rule from UK time-use data, the findings and their interpretation, empirical performance of the model and comparison to other approaches. In section 10, I set out a methodology to estimate individual-level consumption. I implement it by applying my estimated sharing rule to a separate but comparable dataset on UK household material expenditure, and discuss policy implications. I conclude with some reflections for future work in section 11.

7 The model

In this section I set out a general static collective model of the household with both private and public goods and both material goods and time-use. To date, the collective model of the household is both more general, and provides a better empirical fit, than alternatives in the literature (e.g. see Browning, Chiappori and Weiss (2014)). This structural model underpins the methodology proposed in this paper.

which was an important feature of the UK economy in the time-period they studied, but is no longer in use in the UK.

7.1 Framework

An individual i belongs to a household h . Each individual has a person type t (e.g. man, woman, child).⁶ Households are in the same category g (e.g. singles, heterosexual couples...) if they have the same composition, i.e. the same number of household members of each type: $N_{h,g}^t = N_g^t, \forall h \in g$. To keep the notation leaner, I index individual variables only by i and h (even though these have an associated t and g), and household variables only by h (not g). I index type-specific parameters that vary between household categories with both t and g , and household category-specific parameters only with a g . Where summing over multiple individuals, I use the letter s instead of i , and the notation st for the type of person s . A household may have multiple members of the same type. Individuals have a vector of characteristics $\pi_{i,h}$ (e.g. age, educational attainment) and their households have a vector of characteristics ζ_h (e.g. the gender ratio in the region).

The household purchases two types of goods on the market: private material goods and public material goods.⁷ $c_{i,h}$ is the vector of market purchased private goods consumed by individual i . c is the sum of these vectors over all household members. A specific good, e.g. food, is indexed by $j \in \Omega^c$. The private good j has price p^j , and the vector of prices is p . X is the vector of market purchased public goods. Since these are public goods, all household members consume the full amount purchased by the household. A specific good, e.g. housing, is indexed by $j \in \Omega^X$.⁸ The public good j has price r^j , and the vector of prices is r .⁹ I do not index p^j and r^j by h because in many applications they will be constant across the sample, but where they do vary between households everything carries through with the

⁶Depending on the application of interest, types can be defined more granularly to model more heterogeneity. Children may be modelled as decision-makers or as public goods, depending on their age.

⁷This paper focuses only on the private-public framework (which can incorporate partially public goods by appropriately defining goods e.g. splitting car fuel into car fuel used for holidays and car fuel used for work trips) but it readily extends to a model with both shareable and public goods. In this case, actual consumption of each member is obtained by multiplying the vector of market purchases for that member by an economies of scale matrix A . The A matrix depends only on household size. Analysis remains the same, with the addition of this matrix (potentially leading to more complex forms for demands, especially where cross-good economies of scale are allowed). An important strength of the shareable goods framework is that it does not impose the degree of economies of scale of different goods. This can be put to use in papers, such as BCL, which are able to estimate economies of scale parameters in addition to the sharing rule. However, the identification approach in this paper, similarly to DLP and LPW, would not allow recovering these additional parameters.

⁸I do not restrict the types of goods. They can be normal or inferior, goods or bads...

⁹The more commonly used notation in this literature is q and p for the quantity and price of private goods, and Q and P for those of public goods. I deviate from it for two reasons. Firstly, to highlight the fact that c and X are only material private and public goods. This allows me to use Q to refer to the set of all public goods. Secondly, for the very practical reason that it is easy to confuse lower and upper case instances of the same letter, particularly for the letter p . For this reason, I use R , rather than P for the price vector associated with Q .

small addition of h subscripts to the material good prices.

Time is continuous,¹⁰ and each individual has time-endowment normalised to 1, which can be spent in different activities. For exposition, I distinguish between four types of time-use: private leisure, joint leisure, market work and domestic work. Private leisure $\ell_{i,h}$ includes leisure activities enjoyed by an individual without the co-presence of other household members, e.g. reading a book alone or having coffee with a friend. It is therefore a private good, which enters the utility function directly. Joint leisure $jt_{i,h}$ and domestic work $d_{i,h}$ may instead be thought of as inputs to public goods $D = f^D(\mathbf{d})$ and $JT = f^{JT}(\mathbf{j}\mathbf{t})$ enjoyed by the household (or by specific sub-sets of members in the case of joint leisure involving only some household members). These production functions can accommodate heterogeneous productivity by type.¹¹ Market work $m_{i,h}$ does not enter the utility function.¹² In order to use private leisure as the assignable good, the key is to distinguish between it and other types of time-use (a more detailed categorisation is consistent with the model, but is not necessary). In particular, it is important to distinguish between private and public leisure¹³ as well as between leisure and non-leisure activities such as domestic work.

We write Q_h for the vector of public goods including both material and time-use public goods (X_h, JT_h, D_h) and R_h for the associated price vector. R_h is indexed by h because it includes the wages $w_{i,h}$ of each of its members. Each individual commands an exogenous wage $w_{i,h}$ for a unit of market work.¹⁴ We can think of an individual's unobserved skills as determining their hourly pay. Individuals then choose, within the set of jobs available to them given their pre-determined skills, whether to work in a longer-hour, higher-overall-pay

¹⁰By examining UK time-use data used for the application in this paper, this modelling assumption appears realistic.

¹¹Productivity at home, and preference parameters, can in principle depend on characteristics such as age and productivity in labour markets. Hence, the model can account for very general patterns of behaviour, such as those documented in the time-use literature (e.g. see Bastian and Lochner (2020, August)). To simplify exposition, I treat preferences and productivity as constants which vary at the person type level. The theoretical results of this paper extend to the context with heterogeneity.

¹²This model can be extended to accounting for individuals taking some pleasure in their work (and/or domestic work) by modifying the time budget constraint so that an hour spent working reduces leisure time by less than an hour, capturing the fact that part of the time spent working is enjoyed. See Browning, Chiappori and Weiss (2014) for a discussion.

¹³Browning, Donni and Gørtz (2020) finds that these are far from perfect substitutes.

¹⁴A static collective model is incompatible with hourly pay being endogenous (see e.g. Browning, Chiappori and Weiss (2014) section 4.4.2). Since hourly pay generally affects bargaining power in the household, the household problem becomes inefficient if it is a choice variable. In a static framework, we can reconcile (i) modelling wages as exogenous in the household problem, and (ii) wages in actuality depending on time spent working, if individuals are myopic about the impact of their present time-use on future periods. If instead we wish to allow bargaining weights to be endogenous (through time-use affecting wages, or another channel), we must depart from the static collective model either by (i) employing a dynamic collective model with limited commitment, or (ii) remaining in a static framework, but choosing an inefficient household model instead of the static collective model (which is efficient by assumption). I note that this is a general issue for this literature, and is not specific to using private leisure as the assignable good.

job, or a shorter-hour, lower-overall-pay job. For instance, someone with high numerical literacy will likely command a high hourly wage, and might choose between the longer hours and higher overall salary of investment banking, and the shorter hours and lower (though still substantial) overall pay of industry forecasting. While contracted hours cannot be chosen freely in many jobs, in this model m_i are actual hours worked, which can be chosen freely and are more naturally modelled as a continuous choice variable.

Given the time endowment has been normalised to 1, we model each individual as having a labour income endowment equal to $w_{i,h}$. In addition, each member can be endowed with non-labour income $y_{i,h}^{NL}$ (or alternatively the household as a whole can have non-labour income y_h^{NL}). An individual's overall endowment is $y_{i,h} = y_{i,h}^{NL} + w_{i,h}$. The household's endowment (or full income) is $y_h = \sum_{i \in h} y_{i,h}$. I refer to this as the household budget.

Each individual i of type t living in a household category g has utility function $u_{t,g}$. This allows preference heterogeneity across types and household compositions. For instance, a woman living alone can have different preferences from a woman living with a partner, and also different from a man living with a partner. Person types and household categories can be chosen to be arbitrarily granular, allowing additional heterogeneity in preferences. However, for implementation, it will be practical to restrict the number of person types and household categories. Types should be chosen to capture the key likely dimensions of heterogeneity in the context of interest, but without reducing the sample size for each category too drastically. Estimation of the sharing rule must be conducted separately for households of different categories. This is because the form of the bargaining solution of each household depends on the number and types of members, so that the sharing rule for different household categories must be estimated separately.

7.2 The household's optimisation problem

In a household, the constituent individuals bargain over how to divide resources. Depending on the bargaining process, and on the outside options of the individuals, the different individuals will have different bargaining power and the resulting division of resources will be different.

The collective model of the household does not restrict bargaining to any specific solution, and only requires that this process be efficient.¹⁵ Relative bargaining power will in general depend both on (i) market variables such as prices (including wages), and (ii) distribution factors, which enter the household's optimisation problem only indirectly through the distribution of bargaining power, e.g. age and education of members (elements of the vectors of

¹⁵As opposed to models with specific bargaining solutions e.g. McElroy and Horney (1981).

individual and household characteristics $\pi_{i,h}, i \in h$ and ζ_h). I refer to the vector of variables that affect bargaining power as z_h . Note that a specific member's bargaining power will depend not only on their own characteristics, but also on the characteristics of all other household members, hence the household-level subscript.

A key result from the existing literature (see Browning, Chiappori and Weiss (2014)) is that the problem solved by any collective household, regardless of the underlying bargaining process and outside options, can be represented as an optimisation problem where the maximand is the weighted sum of the members' utility functions. Each member's utility function is weighed by their Pareto weight $\mu_{i,h}(z_h)$ normalised so that $\sum_{i \in h} \mu_{i,h}(z_h) = 1$. The higher an individual's Pareto weight, the more weight the collective household gives their utility in determining its choices. The household's optimisation problem is therefore to maximise $\sum_{i \in h} (\mu_{i,h}(z_h) u_{t,g}(c_{i,h}, X_h, l_{i,h}, JT_h, D_h))$.

7.3 A problem in two stages

We can re-cast this as a two-stage problem.¹⁶ This representation is very helpful for identifying individual-level resources. In the first stage, the household chooses expenditure on public goods $R_h Q_h = \sum_{j \in \Omega^X} r^j X_h^j + \sum_{i \in h} w_{i,h}(d_{i,h} + jt_{i,h})$, and how to divide the remaining household budget into individual budgets $\rho_{i,h} = (y_h - R_h Q_h) \eta_{i,h}$ for members. The sharing rule¹⁷ determines the share of household budget net of public good expenditure assigned to each member (the individual's resource share $\eta_{i,h}$), with the shares normalised to sum to one $\sum_{i \in h} \eta_{i,h} = 1$. In the second stage, members decide how to allocate their individual budgets $\rho_{i,h}$ to private expenditure.

First stage: public goods and individual budgets

$\max_{\rho, X, d, jt} \sum_{i \in h} (\mu_{i,h}(z_h) v_{t,g}(\rho_{i,h}, X_h, D_h, JT_h))$ s.t. the following constraints:

- Budget constraint: $\sum_{i \in h} w_{i,h}(d_{i,h} + jt_{i,h}) + \sum_{j \in \Omega^X} r^j X_h^j + \sum_{i \in h} \rho_{i,h} = y_h$
- Time feasibility constraint: $d_{i,h} + jt_{i,h} \leq 1$
- Non-negativity constraints: $\rho_{i,h}, d_{i,h}, X_h, jt_{i,h} \geq 0$
- Domestic and joint leisure production functions: $D = f^D(\mathbf{d}), JT = f^{JT}(\mathbf{jt})$

¹⁶Separability is often assumed when employing the two-stage representation of a collective model with public goods (e.g. Lise and Seitz (2011)), but it is not a necessary assumption. Without separability, second-stage demands generally depend on public good consumption.

¹⁷In a model with private and public goods, what is called the sharing rule here is sometimes referred to as the conditional sharing rule. This is because it only affects the share of private expenditure of members, conditional on the household's choice of public good expenditure.

Second stage: individual optimisation over private good consumption

$\max_{c_{i,h}, l_{i,h}, m_{i,h}} u_{t,g}(c_{i,h}, l_{i,h}, X_h, D_h, JT_h)$ s.t. the following constraints:

- Budget constraint: $\sum_{j \in \Omega^c} p^j c_{i,h}^j + w_{i,h} l_{i,h} = \rho_{i,h}$
- Time feasibility constraint: $l_{i,h} + m_{i,h} = 1 - (d_{i,h} + jt_{i,h})$
- Non-negativity constraints: $c_{i,h}, l_{i,h}, m_{i,h} \geq 0$

In general, second-stage demand depends on the prices of all private goods, including leisure (but not directly on the prices of public goods). Unless preferences are separable in private and public goods, it depends on public good consumption as determined in the first stage. Finally, it depends on the individual's second-stage budget, which is their resource share multiplied by the household budget net of public good expenditure (determined in the first stage). The assignable good c^a is a private good for which we observe individual-level demand (or expenditure) data. Individual i 's (of type t) second-stage demand for the assignable good c^a takes the form: $c_{t,g}^a(p, w_{i,h}, Q_h, (y_h - R_h Q_h) \eta_{i,h})$.

Note that this two-stage representation allows us to think of foregone wages from public time-use as foregone first-stage household expenditure on material public goods and on individual second stage budgets. At the same time, the two-stage representation implies that foregone wages from private leisure are foregone individual second-stage expenditure on private material goods. This is *not* an assumption that is imposed on a model, but rather a feature of the static collective model with time-use. The result is the same regardless of whether the model is solved in its original form or its two-stage form.

8 Methodology

Crucially for identification, individual i 's second-stage demand for any private good, including the assignable good, depends only on observables, i 's own resource share, and i 's own preference parameters. Hence, if we observe individual-level expenditure for one assignable good for each household member, and make an identifying assumption restricting preference heterogeneity, we can in general identify resource shares. The advantage of this approach is that identification requires individual-level demand for one good only, rather than the whole unobservable individual-level expenditure system. The structural model simplifies the mapping from household-level expenditures to individual-level expenditures, so that we only need to estimate a very limited number of parameters (the resource share parameters).

In this paper I propose a new method to estimate the sharing rule under the model set out in section 7, with (i) Cobb-Douglas preferences, (ii) leisure as the assignable good, (iii) an identifying assumption restricting preference heterogeneity, and (iv) an additional restriction on preference heterogeneity allowing estimation of the sharing rule from time-use data alone. In section 9 I employ this very tractable approach to estimate resource shares from UK time-use data. As explained in section 8.1.4, I focus my application only on heterosexual working couples without cohabiting children. The model seems to fit well and the assumptions made do not appear to be strong-arming the data in this context. The methodology set out in this section can straightforwardly be applied to different assignable goods: in section 9.4, I compare my baseline results to findings based on using clothing and non-market-work as the assignable good. Similarly, the methodology can be adapted to different identifying assumptions. In section 9.4 I discuss how the results for clothing change substantially depending on which identifying assumption is used for clothing. While Cobb-Douglas preferences greatly simplify the problem, it is possible to extend the approach in this paper to different parametric specifications. Online appendix A sketches out a general identification result and the required restrictions on functional form of preferences. In Online appendix A, I also outline a detailed example with the Almost Ideal Demand System.

8.1 Assignable good: private leisure

Regardless of our choice of assignable good and exact methodology, the core approach underlying papers including DLP, LPW and this paper is the following. The variation in assignable good expenditure with household budget is taken to be informative about how household budget is split between individual members' second stage budgets. This relies on the strong assumption that men and women have the same preference parameter for the assignable good, or else that we already know the preference parameters of men and women, e.g. from singles data. These types of identifying assumptions can be criticised for any assignable good, and any disputes on their credibility must ultimately be resolved empirically. This paper provides some analysis in this regard, but there is substantial scope for future work to comment on the credibility of the assumptions used in this, and other, papers in the literature.

If the identifying assumption is considered to be credible enough (at least better than the assumption that men and women share resources equally) then the rest of the exercise is rather straightforward. The variation in assignable good expenditure allows for the identification of the sharing rule, i.e. how the budget is allocated to members. This is all that is required to compute estimates of individual-level *overall* consumption. There is no assumption that

expenditure on the assignable good is informative about the split of expenditure on other *specific* goods, and it is maintained that men and women may have heterogeneous and unknown preference parameters on all goods apart from the assignable good. Therefore, when using private leisure as the assignable good, this should not be interpreted as assuming that expenditure on any specific material good is split between members in the same manner as leisure. Rather, the assumption that I make in this paper, and which is supported by UK singles' data, is that men and women have a similar preference parameter for private leisure specifically. This means that, if they have the same wages and the same individual budget, then they will have the same hours of private leisure. Hence we can use data on expenditure on private leisure, together with household budget and wages, to infer how the budget was split between men and women. This does not mean, say, that men and women also spend a similar proportion of their budget on clothing - in fact that is strongly rejected by my data.

8.1.1 Advantages of using private leisure

Clothing has been commonly used to identify sharing because it is often available split into women's, men's and children's clothing in household expenditure surveys. However, it also has limitations, as discussed for instance by LPW. Perhaps most concerning, clothing is a durable good, with a high proportion of 0s in expenditure surveys. In my UK data, 73.20% of households have zero recorded expenditure on at least one of male and female clothing. This is an issue both because it reduces the heterogeneity we can leverage for identification, and because most of the literature treats all observations as interior solutions, and when a large proportion of them are in fact corner solutions this can lead to inaccuracy and also bias towards equality because people with lower resource shares are more likely to hit the corner.

LPW show that they obtain similar estimates from clothing and individual-level food, using data from Bangladesh, and argue this supports the use of clothing as an assignable good because individual-level food expenditure is conceptually a more reliable assignable good. However, individual-level food expenditure cannot generally be used as the assignable good because that type of data is rarely available due to how resource-intensive it is to collect.

Other papers, including Lise and Seitz (2011), have used non-market work as the assignable good. This is obtained as a residual by subtracting usual hours worked (which is often asked as a recall question in expenditure surveys) from the time endowment. The idea is that this is a proxy for leisure which is often available in expenditure data. As acknowledged by authors who have used this assignable good, it is conceptually a good with substantial public components because it includes both domestic work and leisure time spent jointly with other

household members. This is likely to bias estimates towards equal sharing.

I propose using individual time-use, instead of household expenditure, data to identify sharing. This is helpful because it allows us to strip out domestic work from non-market-work, based on activity information. Additionally, time-use data typically includes co-presence information, allowing us to also strip out leisure time spent jointly with other household members. Having removed the public components of non-market-work, we are left with a more credibly private assignable good: private leisure. Another advantage of private leisure is there are no zeros in the data, unlike clothing, and in fact it commands a very large budget share.

For these reasons, we may expect private leisure to yield more accurate estimates of intra-household sharing than commonly available alternative assignable goods, at least in some contexts. In section 9.4, I show that the conceptual arguments outlined above appear to be borne out empirically in my UK data: for instance, the distribution of estimated resource shares obtained using non-market-work is shifted towards equal sharing relative to that obtained from private leisure.

8.1.2 Theoretical consistency of using time-use data to estimate sharing

One may worry that household decision-making over time-use in reality is quite separate from decisions about material consumption, and that we cannot conceptualise a single resource share as applying to both types of goods. This would require a departure from the static efficient collective model which underpins this paper and other papers in this literature. If in reality there was such a divergence in sharing between material and time-use goods, then it may be problematic to use private leisure as the assignable good to estimate a sharing rule to apply to material goods, and vice versa it may be problematic to use a material assignable good to estimate a sharing rule to apply to time-use. Reassuringly, Calvi et al. (2022) find that estimates of women’s resource shares (from material assignable goods) are closely related to those women’s satisfaction with their availability of leisure time, suggesting that the collective model is right to conceptualise the same sharing rule being applied both to time-use and material goods. Moreover, my resource share estimates, obtained from private leisure as the assignable good, are highly correlated to the wage ratio,¹⁸ a commonly used proxy for sharing.¹⁹

Note that the model underpinning this paper is compatible with a woman’s leisure consumption and her material consumption moving in opposite directions when her resource share

¹⁸Defined as the ratio of female wages to the sum of female and male wages.

¹⁹The correlation coefficient is 0.76.

increases. For example her material consumption may increase while her leisure decreases if her increased bargaining power is driven by an increased wage (the price effect may outweigh the income effect).

It is also important to note that private leisure is never ‘imposed’ by other household members in the collective model with participation. Intuitively, while it is true that I cannot take joint leisure with my partner if they are not taking joint leisure with me at that time, I am able to choose whether to undertake private leisure, domestic work, or market work instead. This is similar to what the model implies for public goods: while it is true that I cannot consume more heating if we do not collectively decide to spend more on heating, this does not imply that I must spend any additional money on my private consumption of food, e.g. I could purchase more clothing or spend more time having private leisure. In the same way, if we collectively choose not to spend more of our budget on public time-use, I could spend more budget, and more time, on private leisure, or I could spend more on material goods (and spend more time working). In order for this to be the case, it is important that hours worked can be chosen.

8.1.3 Wage endogeneity

In the static collective household model wages are exogenous. However, we may worry that in reality wages are endogenous. If time-use affects wages, and wages affect bargaining, that cannot be accounted for within the static collective framework, as it would lead to inefficient bargaining. The fact that previous papers have found the collective model to be a good empirical fit (see Browning, Chiappori and Weiss (2014)) suggests that wage endogeneity is not a substantial issue. For instance, it may be that individuals do not take into account the transitive implication of (i) their wage depending on their time-use, and (ii) their bargaining power depending on their wage, when making decisions. Note that if there was a substantial endogeneity issue, while this issue is more salient when using private leisure as the assignable good, it would be problematic regardless of the choice of assignable good, as the estimating equations derived from the static collective model would not be valid.

8.1.4 Restricting the sample to working couples

For people who participate in market work, the price of time is the observed wage. For those who do not, the price of time is an unobserved quantity, which exceeds the (also unobserved) wage, as explained in online appendix E. Therefore, in general to estimate

sharing for households with non-participation we will require either a different approach,²⁰ or additional data, assumptions and estimation steps, to recover this unobserved price. This is not necessary if we choose an assignable good which is separable from private leisure, but it is always relevant when using private leisure itself as the assignable good.

Partly for this reason, incorporating non-participation is left to future work. In this paper, I focus on working couples without cohabiting children. This sample restriction is not unusual in the literature. In particular, other papers which explicitly incorporate time-use generally exclude children or else model expenditure on children as a public good (Almås et al. (2020)). This sub-sample of the population is still of great interest, and in particular can shed light on gender inequalities and some of their drivers.

A further reason why I focus on working couples is that we can more credibly test the identifying assumptions for this sub-sample. As discussed in section 9.4.2, tests from data on working singles appear to be highly informative for this sub-sample, but, as explained in section 9.3.1, may be less convincing for other household compositions.

It is important to note that, in the UK context, selecting only couples where both members work is unlikely to introduce any bias in the estimates. The reason is that non-participation in couples without cohabiting children is similar between genders in the UK and is generally due to reasonably exogenous drivers such as long-term illness or temporary unemployment (see online appendix E). In other countries, with much lower female participation in labour markets, it will be important to incorporate non-participation.

8.2 Parametric specification: Cobb-Douglas

Utilities and domestic production functions are Cobb-Douglas, so that we can write:²¹

$$u_{t,g} = \sum_{j \in \Omega^c} \left(\alpha_{t,g}^{c^j} \ln(c_{i,h}^j) \right) + \sum_{j \in \Omega^X} \left(\alpha_{t,g}^{X^j} \ln(X_h^j) \right) + \alpha_{t,g}^l \ln(l_{i,h}) + \sum_{s \in h} \left(\alpha_{t,g}^{T^{st,g}} \ln(j t_{s,h}) \right) + \sum_{s \in h} \left(\alpha_{t,g}^{D^{st,g}} \ln(d_{s,h}) \right)$$

In the context of resource share estimation, Cobb-Douglas preferences have several advantages. Importantly, they lead to parsimonious estimating equations, even with a rich characteristics vector z_h affecting bargaining power. This enables estimation from realistic, widely available data, which generally involves small sample sizes. Small sample sizes are common in this literature, partly due to limitations of existing data, and partly by construction, be-

²⁰e.g. Blundell, Chiappori, Magnac and Meghir (2007) estimate the sharing rule, up to a constant, using a revealed preference approach.

²¹See online appendix F for a more detailed discussion, including normalising assumptions.

cause estimation must proceed separately for households of different categories, so that even with large data the sample is divided into smaller sub-samples for estimation.

Additionally, Cobb-Douglas preferences, unlike e.g. the Almost Ideal Demand System, have a direct utility representation, which presents several advantages. Firstly, it allows for very clear interpretation of identifying assumptions in terms of restrictions on preferences. Moreover, the tractable direct utility representation of Cobb-Douglas enables estimation of the sharing rule from time-use data alone. Furthermore, under Cobb-Douglas preferences, resource shares are equal to Pareto weights scaled by a composite preference parameter. In my application to UK data, I find a significant effect of wages on the sharing rule, which with Cobb-Douglas preferences implies that Pareto weights vary with wages. This is inconsistent with the unitary model and can be interpreted as evidence for the collective model. This finding is in keeping with several other findings in the household economics literature (see Cherchye, De Rock and Vermeulen (2012)). Additionally, the simple, explicit utility representation of Cobb-Douglas preferences allows one to sense-check estimates. In online appendix F.6, I set out a back-of-the-envelope method to check that the estimated sharing rule parameters are quantitatively, as well as qualitatively, consistent with bargaining theory. The separability properties of Cobb-Douglas are also helpful in reducing the amount of data required for estimation. For instance, assuming that private leisure is separable from various types of public time-use means I do not need to estimate time spent on specific other types of time-use. That exercise can be complex due to the sometimes fine line between joint leisure and domestic work, and would require careful consideration of household production functions.

As explained in online appendix A, while it is possible to identify the sharing rule using more flexible functional forms than Cobb-Douglas, this often requires either focusing on type-specific constant shares or else making some additional linear approximations. For instance, with covariates z , an additional small linear approximation is required for Almost Ideal Demand System Engel curves to be estimated by OLS (similarly to LPW). This limits the degree to which alternative preferences would better capture any non-linearities. Moreover, more flexible parametric forms often lead to additional practical issues for estimation because, in order to allow resource shares to depend on several characteristics z_h , estimation requires a large number of regressors, many of which are highly correlated with each other (see online appendix A.5). For these reasons, implementation will often benefit from using very parsimonious functional form assumptions, such as Cobb-Douglas.

The price of the tractability and ease of application achieved with Cobb-Douglas is a strong parametric assumption about preferences. In particular, Cobb-Douglas preferences impose homotheticity, which may seem to strong an assumption for certain goods. It is important to

note that I do not require the whole utility function to be Cobb-Douglas: more specifically, I require preferences that are Cobb-Douglas over the assignable good and some aggregate good. Therefore, I am imposing homotheticity on the assignable good, but not on other specific goods. As discussed in section 9.3.3, I test homotheticity of private leisure using data on working singles, and cannot reject it (while, for instance, clothing appears to be a necessity good in my singles data). In section 9.3.3, I outline additional tests that suggest that the functional form assumption is not driving my results in my application to UK data. The Cobb-Douglas assumption is less likely to be realistic for other candidate assignable goods such as food and clothing, which are strongly non-homothetic.

8.3 Assumptions restricting preference heterogeneity

8.3.1 Identifying assumption: private leisure C-D SAP

Identification of the sharing rule from individual-level demand data for a single assignable good requires a strong identifying assumption. It is important to note that the chosen assumption has an important impact on estimates, and should be chosen carefully. Assumptions could be made on the bargaining process e.g. Lise and Seitz (2011) assume that women and men with the same potential earnings divide full income equally. Other papers have focused on restricting preference heterogeneity. For example, Lewbel and Pendakur (2008) and Bargain and Donni (2012) rely on the ‘SAT’ assumption (similarity across household types), that preferences for the assignable good are stable across household compositions, allowing for identification of the preference parameter from singles data. However, we may be concerned by findings that, at least in some contexts, preference stability across household composition is rejected empirically (see Hubner (2020)). DLP and LPW use the alternative ‘SAP’ assumption (similarity across people) that women and men in the same household composition have the same preference parameter for the assignable good. Whether this is a good assumption depends on whether we have reason to suspect substantial gender differences for the assignable good.

My data does not allow me to directly test this identifying assumption for private leisure, but I am able to check that it is consistent with singles data, as explained in section 9.3.1, and hence use the SAP identifying assumption. In other contexts, SAP may not be a realistic assumption: as discussed in section 9.4, singles data suggest that women have much stronger preferences than men for clothing. Therefore, if using clothing as the assignable good, it may be better to opt for SAT, or a somewhat weaker version of that assumption, which I term ‘SRAT’ (similarity across household types). I use this novel SRAT assumption in my robustness check with clothing as the assignable good, finding that the results are

much closer to my baseline results from private leisure than when using clothing and the SAP assumption. I also show that, under my baseline sharing estimates, structural estimates of the ratio between men and women’s preference parameters for clothing are consistent with the ratio observed in singles’ clothing expenditure patterns.

I note that the interpretation of the identifying assumptions differs depending on the functional form of preferences and the assignable good of choice. Because I use a SAP-type identifying assumption in conjunction with Cobb-Douglas preferences and private leisure as the assignable good, the specific identifying assumption I am making may be referred to as ‘private leisure C-D SAP’. Everyone, regardless of their type (in a given household category) is assumed to have the same Cobb-Douglas preference parameter for private leisure: $\alpha_{t,g}^l = \alpha_g^l$. Preference heterogeneity across types is maintained for all other preference parameters, and across household compositions.

In the context of heterosexual couples, this assumption means that men and women spend the same share of their individual second-stage budget on leisure, but will split the remaining budget differently between other private goods. This means that I can infer how the budget is split between them by observing their expenditure on private leisure. It *does not* imply that I conclude that expenditure on each other good is split in the same manner as expenditure on private leisure.

Where SAP is not considered to be realistic, we can use alternatives from the literature such as SAT. In the context of private leisure as the assignable good, and Cobb-Douglas preferences, SAT implies that: $\alpha_{t,g}^l = \alpha_t^l$. The preference parameter for private leisure is stable across household composition for people of a specific type. This does not require preferences for other goods to also be stable across household compositions. Alternatively, I propose a novel, milder version of SAT: SRAT. Instead of requiring the level of preference parameters for the assignable good to be stable across household composition, I require only that the ratio between men and women’s preference parameter for the assignable good is stable across household compositions. Private leisure C-D SRAT requires that: $\frac{\alpha_{t,g}^l}{\alpha_{st,g}^l} = \frac{\alpha_t^l}{\alpha_{st}^l}$. Hence, it is implied by (but does not in turn imply) SAT. Where we cannot reject equality between men and women’s preference parameters for the assignable good in singles data, as is the case for private leisure, SAP and SRAT can be considered equivalent. Further work is required to carefully test the validity of different identifying assumptions applied to different assignable goods and different samples.

8.3.2 Additional restriction for ease of estimation: public good C-D SAP

The second-stage demand for assignable goods is conditional on public good expenditure (even when preferences are separable) because the resource share enters multiplicatively with the household budget *net of public good expenditure*. This presents two challenges. Firstly, it entails the substantial data requirement of having high-quality data both on time-use and material expenditure to construct overall expenditure on public goods, including domestic time and joint leisure as well as material goods. Ideally, this data would be in the same dataset. Alternatively, there would be a separate but comparable dataset including similar household characteristics - then the missing data could be estimated based on the separate dataset. The second challenge is correctly constructing a net of public expenditure budget, which in reality has substantial inter-temporal components, within a static framework. Squaring a static full income with observed public expenditure and implied budget for private consumption, requires careful consideration of savings, investments and large durables (especially housing). In section 10 I explain how I tackle these challenges for the purpose of estimating individual-level material consumption (after having recovered the sharing rule). A similar approach could be taken here too, but there is a risk that it might give rise to inaccuracies which could affect the accuracy of sharing rule estimates.

To circumvent the challenges presented by estimating household budget *net of public good expenditure*, one can impose an additional preference restriction that total expenditure on public goods does not vary with the distribution of bargaining power in the household (public good SAP). This assumption *does not* amount to a unitary model in that the distribution of bargaining power still affects the composition of public good (and private good) expenditure; it simply does not affect the *overall* split between private and public expenditure. For example, female empowerment could lead to more spending on housing and less on television subscriptions.

Intuitively, public good SAP may seem too strong assumption because some papers have found gender differences on investment in children, although the evidence is mixed (e.g. see Gitter and Barham (2008)). Regardless of this debate, there is no clear reason why public good SAP would be problematic for the working couples in my UK application. In fact, this assumption appears to be consistent with the data in the context of this paper's application, as discussed in section 9.3.2. Care should be taken extending this assumption to other contexts, especially in households with children if children are modelled as public goods, as gender differences are more likely to emerge in the desired split between private and public expenditure.

With Cobb-Douglas preferences, we can achieve the desired restriction with public good C-D

SAP, as explained in more detail in online appendix [F.3](#). Public good C-D SAP imposes that the sum of all public good preference parameters is the same for all individuals of the same household category g :²²

$$\left(\sum_{j \in \Omega^X} \alpha_{t,g}^j + \sum_{st \in h} \alpha_{t,g}^{Tst,g} + \sum_{st \in h} \alpha_{t,g}^{Dst,g} \right) = a_g^Q$$

This assumption allows heterogeneity in how different types of people would choose to divide the budget between specific private goods, and also between specific public goods, while requiring that they would choose the same overall split between private and public goods. As a result of the public good SAP assumption, it is possible to estimate the sharing rule from time-use data alone, if using private leisure as the assignable good. Alternatively, if using a material assignable good, it is possible to account for time-use as well as material goods, and still estimate the sharing rule from expenditure data alone. Note that this assumption is not required if we instead choose to estimate household budget net of public good expenditure.

I also note that, for the purposes of estimating the sharing rule, this assumption allows us to proceed as if all goods were private, though in a manner that is consistent with a model with public goods. The extension to public goods bites primarily when it comes to using the estimated sharing rule to estimate individual-level consumption. At that stage it makes a substantive difference whether we consider some goods as public, as in this paper, or not, as in LPW (see section [10.2](#)).

8.4 Linear approximation of resource shares

Resource shares will generally depend on household budget, prices (including wages of all household members) and Pareto weights of household members: $\eta_{i,h}(y_h, p, r, w_h, \mu_h)$. Recall that Pareto weights μ_h are an unknown function of a vector of variables z_h that determine relative bargaining power in the household, and that the collective model does not pin down a specific bargaining solution. Since z_h in general includes the other variables which affect $\eta_{i,h}$, we can write $\eta_{i,h}(z_h)$. To estimate resource shares by linear regression, we linearly approximate them as:

$$\eta_{i,h} = \eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})$$

- $\eta_{t,g}^0$ is the average resource share of a type (living in a specific household category). This is the resource share evaluated at the average characteristics in the sample. By definition, $\sum_{t \in h} \eta_{t,g}^0 = 1$. In the context of heterosexual couples, the average resource share of men and women sum to one.

²²Equivalently, the sum of all private good parameters is also homogeneous within category: $\left(\sum_{j \in \Omega^c} \alpha_{t,g}^j + \alpha_{t,g}^l \right) = a_g^{c,l}$.

- $\eta_{t,g}^z$ captures the impact on sharing of a household's characteristic z_h deviating from the sample average ($z_h - \bar{z}$). For instance, a higher-than-average wage for the woman might increase the woman's resource share, so that she would have a higher-than-average-for-women resource share. Since resource shares must sum to one within the household, this implies her partner must have a correspondingly lower-than-average-for-men resource share: $\sum_{t \in h} \eta_{t,g}^z = 0$. We can interpret $\eta_{t,g}^z$ as the marginal impact of characteristic z_h on the resource share.

A linear approximation does not guarantee estimates of resource shares which fall within the unit interval. However, by construction, only estimates in this range are theory-consistent. This provides a useful test of model fit. Reassuringly, in my application to UK data, my baseline resource share estimates from time-use data are all within the unit interval.

8.5 Estimation

Under the framework set out above, I structurally derive a linear and parsimonious system of estimating equations. This comprises an equation for each type of person in the form:

$$w_{i,h} l_{i,h} = \alpha_{t,g}^l y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z}))$$

8.5.1 A structural approach to the error term

To proceed to estimation, we must consider the source of any error terms. In this case, there are three likely sources of error in our estimating equations. The first is approximation error from linearly approximating the resource share based on the characteristics vector z_h . The second is household optimisation error at the first stage of the household problem. The third is individual optimisation error at the second stage of the household problem. As explained in online appendix F.4 we can write $w_{i,h} l_{i,h} = \alpha_{t,g}^l y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})) + \epsilon_{i,h}$ where $\epsilon_{i,h}$ is mean-zero. If z_h includes all key characteristics that affect bargaining (or at least all those that are correlated with regressors) then estimates should be unbiased. The errors are negatively correlated within household (with a correlation weaker in magnitude than -1), so that a SURE estimation approach is recommended.

8.5.2 Estimation procedure

For estimation, I focus on heterosexual working couples without cohabiting children²³ for the reasons discussed above. The estimating equations are:

²³Where there are multiple household categories, the approach set out below should be carried out separately for each household category.

1. Equation for women: $w_{i,h}l_{i,h} = \beta_{f,g}^0 y_h + \sum_z \beta_{f,g}^z y_h (z_h - \bar{z})$
2. Equation for men: $w_{i,h}l_{i,h} = \beta_{m,g}^0 y_h + \sum_z \beta_{m,g}^z y_h (z_h - \bar{z})$

where $\beta_t^0 = \alpha_{t,g}^l \eta_t^0$ and $\beta_t^z = \alpha_{t,g}^l \eta_t^z$

Estimation follows these steps:

1. Run linear SURE regressions of leisure expenditure, one for men and one for women, with the restriction that $\beta_f^z + \beta_m^z = 0$. This is because $\beta_t^z = \alpha_t^l \eta_t^z$ and from private leisure C-D SAP and the definition of the linear approximation parameters, $\sum_{t \in h} \beta_t^z = \alpha_t^l \sum_{t \in h} \eta_t^z = 0$. Note that there are other restrictions imposed by the model which I do not impose during estimation, and that can be used to test the fit of the model as discussed in section 9.3.
2. Estimate each type's average resource share as $\hat{\eta}_t^0 = \frac{\hat{\beta}_t^0}{\hat{\beta}_m^0 + \hat{\beta}_f^0}$. To see why, first note that $\frac{\beta_t^0}{\sum_{st \in h} \beta_{st}^0} = \frac{\alpha_t^l \eta_t^0}{\sum_{st \in h} \alpha_{st}^l \eta_{st}^0}$ and by the leisure SAP assumption the α parameters cancel out so that $\frac{\beta_t^0}{\sum_{st \in h} \beta_{st}^0} = \frac{\eta_t^0}{\sum_{st \in h} \eta_{st}^0} = \eta_t^0$ since resource shares sum to one by definition.
3. Estimate the marginal impact of different characteristics as follows. First, estimate $\hat{\alpha}^l = \frac{\hat{\beta}_t^0}{\hat{\eta}_t^0}$ since $\beta_t^0 = \alpha^l \eta_t^0$. Then, estimate $\hat{\eta}_t^z = \frac{\hat{\beta}_t^z}{\hat{\alpha}^l}$ since $\beta_t^z = \alpha^l \eta_t^z$.
4. The estimated parameters yield the estimated sharing rule $\hat{\eta}_{i,h}$

Armed with the sharing rule, one can estimate individual-specific resource shares in the dataset used for estimation (the time-use dataset in this case): $\hat{\eta}_{i,h} = \hat{\eta}_t^0 + \sum_z \hat{\eta}_t^z (z_h - \bar{z})$. Moreover, one can take the sharing rule across to other comparable datasets (e.g. expenditure data) and estimate household-specific resource shares there by applying the sharing rule.²⁴ This can be helpful to proceed to further applications of sharing rule estimation, as exemplified in section 10.

8.5.3 Intuition for identification

The identification problem we face is that the man's expenditure on private leisure may be more or less responsive than the woman's to changes in the household budget for one of two reasons, or a combination of them. The first possibility is that the man receives a larger proportion of the household budget (net of public good expenditure), i.e. he has a higher resource share than the woman. The second possibility is that the man's preference

²⁴To do so, I calculate how household characteristics in the expenditure dataset deviate from the averages in the time-use data, and substitute these deviations ($z_h - \bar{z}$) in the estimated sharing rule $\hat{\eta}_g$.

for private leisure is stronger relative to the woman's, so that he spends a larger proportion of his individual budget on private leisure. In order to disentangle these two channels, I shut down the preference channel with the identifying assumption, allowing me to identify the sharing rule. Under CD-SAP, private leisure expenditure for individuals of any type (within a given household category) $l_i w_i = \alpha^l y_h \eta_{i,h}$ responds in the same way to the same increase in individual budget $y_h \eta_{i,h}$, so differences in responsiveness to changes in y_h identify differences in sharing.

Empirically, identification of the levels of resource shares, and the marginal effects of different characteristics on them, is driven by different sources of variation in the data. For example, *ceteris paribus*, a higher wage for the man will reduce his leisure demand through a price effect, but increase it through a twofold income effect: (i) increased overall household budget, and (ii) increased own-resource share.²⁵ Therefore, the leisure patterns of two similar couples that differ in the man's wage help identify both the effect of male wage on sharing, and, by increasing the overall household budget, the levels of resource shares. Similarly, couples which are identical in characteristics apart from the age gap, and have different leisure patterns, help identify the effect of intra-couple age gaps on sharing.

8.6 Data requirements

In order to identify the sharing rule we require:

- A cross-section household dataset with both household-level information and individual-level data for all members.
- Individual-level data on demand / expenditure on an assignable good for all members. Sometimes assignable good demand or expenditure data will be directly available in the dataset. In other cases, it may require constructing. For instance, when using private leisure as the assignable good, it may be necessary to classify time-use data based on activity and co-presence information.
- Data on full household income. Full income will generally have to be constructed from information on wages, together with any non-labour income of members (in this model, the household budget is not equal to realised earnings nor to expenditure on material goods).

In order to identify household-specific shares (instead of type-specific constants) we additionally require:

²⁵Browning and Gørtz (2012) find that the 'unitary effect', that leisure demand falls as its price increases, dominates the 'collective effect' of the individual's bargaining power increasing in own-wage.

- Data on key characteristics which are likely to affect bargaining e.g. sex, age, educational attainment and wages of all members. Wages may need to be estimated from earnings and hours worked.

Depending on the choice of assignable good and parametric specification, we may additionally require:

- Price data for any goods for which (i) their price enters the estimating equations, and (ii) their price varies substantially in the sample. This includes wage data if private leisure is the assignable good. No price data is required if assignable good preferences are separable from preferences of all goods with price heterogeneity.
- Expenditure data on specific public goods, if there are any which enter the estimating equations (none if the assignable good is separable from public goods).
- Total public expenditure (on both material goods and time-use). As discussed in section 8.3.2, this requirement can be challenging to meet and can be avoided by making an additional assumption restricting preference heterogeneity.

9 Estimating resource shares from UK time-use data

I apply the novel methodology set out in section 8 to UK time-use data. In section 9.1, I describe the data (the UK Time-Use Survey). In section 9.2, I set out my findings and how to interpret them. Women command substantially lower resources than men, with substantial variation between households. The directions of marginal effects and their magnitudes are consistent with bargaining theory. In section 9.3, I discuss several tests of the model and assumptions, which appear realistic in this context. In section 9.4, I illustrate how my findings vary with different assignable goods, and compare my findings to those from two other papers on household sharing in the UK.

9.1 UKTUS

The UK Time-Use Survey (UKTUS 2000, 2014),²⁶ is a high-quality time-use survey that has been used in the economics literature (e.g. Kalenkoski et al. (2005)) but never before with the goal of estimating intra-household inequality. It is a national household-based

²⁶Office for National Statistics. (2019b). United Kingdom Time Use Survey [data series] 2nd Release. <https://doi.org/http://doi.org/10.5255/UKDA-Series-2000054>

study composed of: (i) a household questionnaire, (ii) an individual questionnaire, and (iii) individual time-diaries. A single household representative answered the household questionnaire, including questions on household characteristics such as composition, dwelling type, and location. The other components were answered by the individual in question. This is likely to substantially increase the quality of the data relative to datasets where a single member answers on behalf of all individuals. The individual questionnaire asks about individual characteristics including age, educational attainment and earnings. Each member completed a weekday and weekend time diary identifying primary and secondary activities for each 10-minute interval over the two days. The time-use data includes very detailed activities, location, and co-presence of others (distinguishing between household and non-household members). This enables me to define private leisure very precisely, as time spent doing leisure activities without other household members co-present. For a list of activities, see online appendix B. The time-diaries are constructed carefully to minimise measurement error, for instance with the possibility of writing a simple sign to signify the same activity for multiple time intervals. The quality of the data is very high, with approximately 95% of observations having more than 5 distinct activities recorded in a day, and less than 90 minutes of unrecorded time.

As discussed in section 8, I focus on heterosexual working couples without cohabiting children.²⁷ After cleaning the data, the final pooled (2000 and 2014) sample comprises 711 households (1,422 individuals). While the sample size is not particularly large, it is of a good size relative to this literature.²⁸

Summary statistics for some of the key variables are reported in table 1. Hourly pay is obtained by dividing labour income by actual hours worked, rather than contractual hours.²⁹ Women on average command lower hourly wages than men. Couples generally form with an older, higher-earning man and a younger, lower-earning woman. Household budget is full income: the sum of the labour endowment of members (hourly wage multiplied by 24 hours).³⁰ Women on average are more qualified.

²⁷The sample includes both couples who never had children and couples who have children who no longer cohabit with them. I am not able to run the analysis separately for these two types of household because the data does not contain information on non-cohabiting children over the age of 15 (none of the households in my sample have children under 15 who do not cohabit with them). Therefore, the age effect in the resource share (older couples assign lower resource shares to women) could be partly driven by older couples being more likely to have non-cohabiting children.

²⁸e.g. Cherchye, De Rock and Vermeulen (2012) estimate a closely related model on a sample of 212 Dutch households. Small sample sizes are to some extent inevitable in applications where different household compositions are analysed separately.

²⁹This, together with self-employed labour, explains the lower end of hourly wages (which are sometimes lower than the official minimum hourly wage).

³⁰Accurate non-labour income data is not available in UKTUS.

³¹Highest qualification obtained, simplified into three categories. Category 2 is equivalent to an under-

		UKTUS sample					LCF sample				
		mean	sd	p25	p50	p75	mean	sd	p25	p50	p75
Hourly pay (2020 GBP)	female	9.87	5.58	6.6	8.61	11.6	10.5	5.39	7.51	9.37	12.2
	male	11.6	8.74	7.51	9.78	13.1	12.2	7.61	7.79	10.4	14.3
Age (years)	female	41.9	13	30	43	53	42.1	13.6	29	43	55
	male	44	13.1	31	46	55	44	13.6	30	44	57
Qualification (levels 0,1,2) ³¹	female	.864	.824	0	1	2	1.04	.871	0	1	2
	male	.816	.813	0	1	2	.956	.857	0	1	2

Table 1: Summary statistics: UKTUS and LCF

9.2 Sharing rule estimates

Using the UKTUS pooled sample, I estimate the resource sharing rule for UK working couples. I find that households do not share resources equally: the mean resource share for women is 0.45 and 0.55 for men (the median is 0.44 and 0.56). As shown in table 4, and explained in more detail in online appendix G, I comfortably reject equal sharing. Household-specific resource shares vary substantially (as can be seen in figure 1). This heterogeneity is driven both by market variables (e.g. wages) and distribution factors (e.g. the gap between the age of the man and the woman). These findings confirm the importance of accounting for intra-household unequal sharing, and for household-specific variation in sharing as a function of characteristics. Moreover, these findings reinforce the prevalence of systematic inequality in sharing by gender, with the distribution of female resource shares to the left of that for men, as illustrated in figure 1.

9.2.1 Interpreting regression results and estimating marginal effects

Sharing rule estimates are obtained by running seemingly unrelated regressions of leisure expenditure equations for men and women. I allow resource shares to depend on the following characteristics: the hourly pay and educational attainment of both members, the age gap and average age of the couple³² and regional wealth in the household’s region. The results are reported in table 2.

The mean female resource share is calculated by dividing the coefficient on ‘Budget’ in the female regression by the sum of the ‘Budget’ coefficients across the two equations: $\hat{\eta}_f^0 = 0.45$. The preference coefficient is estimated by dividing the coefficient on ‘Budget’ in the female regression by the estimated average female resource share: $\hat{\alpha}^l = 0.44$. This number is

graduate degree or higher. Category 1 is equivalent to end-of-school diplomas e.g. A levels, IBDP, or equivalent technical qualifications. Category 0 is anything less than that, including lower technical qualifications and school diplomas obtained before the end of school e.g. GCSEs.

³²The age gap is the difference in years between the man and the woman




Figure 1: Baseline results, UKTUS.

The distribution of male resource shares is a reflection of the distribution of female resource shares in the equal sharing (dotted) line. This is because each household's shares must sum to one.

consistent with individuals choosing to take private leisure for a substantial proportion of their daily time endowment i.e. spending a large proportion of their budget on private leisure.³³

To interpret the regression coefficients, consider an example household h . The starting point for their resource shares are the averages $\hat{\eta}_f^0 = 0.45, \hat{\eta}_m^0 = 0.55$. We then adjust for any deviations from the sample mean for determinants of sharing. For instance, consider a deviation on female hourly pay. The effect of a unit deviation from the mean of the female wage (w_f) is calculated by dividing the coefficient on the 'Budget * dev. fem. hourly pay' term by the preference coefficient: $\eta_f^{wf} = 0.01$. We estimate the resource share for a couple where the woman earns £15 an hour, instead of the mean of £9.87 as: $\hat{\eta}_{f,h} = \hat{\eta}_f^0 + \eta_f^{wf} * (15 - 9.87)$, which is $\hat{\eta}_{f,h} = 0.51$ and hence $\hat{\eta}_{m,h} = 1 - 0.51 = 0.49$. Note that the impact of deviations from the mean female wage are very substantial. The change in predicted female resource share for changes in different determinants of sharing is summarised in table 3.

The direction of estimated marginal effects is consistent both with bargaining theory and

³³My baseline definition of private leisure includes sleep. haring rule estimates are similar when sleep is excluded, though the preference parameter falls substantially as expected.

Dependent variable	leisure expenditure	
Equation	male	female
Budget	0.243*** (0.00262)	0.198*** (0.00236)
Budget * dev. fem. hourly pay	-0.00559*** (0.000198)	0.00559*** (0.000198)
Budget * dev. mal. hourly pay	0.00215*** (6.87e-05)	-0.00215*** (6.87e-05)
Budget * dev. fem. qualification	-0.00295 (0.00253)	0.00295 (0.00253)
Budget * dev. mal. qualification	0.0154*** (0.00237)	-0.0154*** (0.00237)
Budget * dev. average age	0.000905*** (0.000151)	-0.000905*** (0.000151)
Budget * dev. age gap	0.000481 (0.000338)	-0.000481 (0.000338)
Budget * dev. regional wealth p.c.	-5.44e-07** (2.19e-07)	5.44e-07** (2.19e-07)
Observations (households)	711	711
R-squared	0.937	0.931

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 2: Regression results, baseline specification, UKTUS.

The symmetry of coefficients on the interaction terms is imposed as a constraint during SUR estimation. Budget * dev. variable = budget * (household-specific value of variable - sample average of variable)

with previous findings in the literature. As shown in table 4, women's resource shares are higher in households with high female hourly pay and low male hourly pay. The impact of the former outweighs the latter, so that women in households with higher household budgets³⁴ have higher resource shares. Older couples (either because they are older or belong to more traditional generations) and couples with a larger age gap (older man relative to the woman) are characterised by a higher fraction of resources going to the man, although the latter effect is not statistically significant. Women have higher resource shares in regions with higher wealth per capita, potentially because of better outside options for women in wealthier regions, or due to wealthier regions proxying more gender progressive regions. More educated men have higher resources shares, and the same goes for females, although for the

³⁴The budget here is the full daily budget, i.e. the labour endowment of the couple (sum of male and female hourly pay, multiplied by 24 hours)

variable		mean	sd	impact of..	on fem. share
Hourly pay (2020 GBP)	female	9.87	5.58	↑ 1 s.d.	0.0706
				↑ £1	0.0127
	male	11.63	8.74	↑ 1 s.d.	-0.0426
				↑ £1	-0.0049
Age (years)	average	42.93	12.82	↑ 1 s.d.	-0.0263
				↑ 10 years	-0.0205
	gap	2.06	4.72	↑ 1 s.d.	-0.0051
				↑ 1 year	-0.0011
Regional wealth p.c. (2020 GBP)		30,446	6,780	↑ 1 s.d.	0.0084
				↑ £5k	0.0062
Qualification (levels 0,1,2)	female	0.86	0.82	↑ 1 s.d.	0.0055
				↑ 1 level	0.0067
	male	0.82	0.81	↑ 1 s.d.	-0.0284
				↑ 1 level	-0.0349

Table 3: Interpretation of regression coefficients

latter the effect is not statistically significant. A possible explanation is that, in the UK, female qualifications have trended upwards more strongly than male qualifications over time and generations, substantially reducing the matching market returns of female education.³⁵

		male hourly pay		mean	95% confidence interval		
		< median	> median		Lower bound	Upper bound	
female	< median	0.43	0.39	Private leisure	0.45	0.44	0.46
	> median	0.50	0.47	Non-market-work	0.49	0.48	0.49
hourly pay				Clothing (adjusted)	0.46	0.39	0.53

Table 4: Female resource shares vary with wages (left) and assignable good choice (right)

9.3 Empirical performance of the model

The model appears to fit the data well. Results are stable, and change in the theoretically consistent direction, when performing various robustness checks, including:

- Different definitions of private leisure, e.g. defining leisure in a much more narrow way, excluding time spent sleeping, eating, and work breaks.
- Estimation on different samples, e.g. excluding outliers with particularly low or high wages, and estimating the sharing rule separately on 2000 and 2014 data.

³⁵i.e. how valuable an asset female education is to women in finding a good partner and being able to secure a strong bargaining position in that relationship.

- Running alternative regression specifications, e.g. substituting the gap between male and female qualification for the levels of male and female qualification, and excluding the age gap and regional wealth, as determinants of bargaining power.

Moreover, the model imposes several testable restrictions on the sign and magnitude both of the regression coefficients and the structural parameters recovered from them. These restrictions are not imposed by the estimation method, and testing them suggests the model fits the data well. Importantly, all resource shares fall within the unit interval.³⁶ Additionally, in conformity to the way noise was incorporated in the model (see section 8.5.1), residuals are negatively correlated within the household, with a correlation coefficient of -0.09.³⁷

The only restrictions that were imposed during estimation required that all coefficients apart from that on household budget sum to zero across equations. I run unconstrained SURE regressions and find the magnitudes of coefficients are similar across the two equations, and the signs are opposites of each other. This suggests that imposing the restriction that these coefficients sum to zero across the equations is not far off correct, and is not strong-arming the results (although some of the estimated coefficients are statistically significantly different from each other).

9.3.1 Testing the private leisure C-D SAP assumption

To test the validity of the private leisure C-D SAP assumption, I use data on working singles using UKTUS pooled data. I focus on working singles to ensure comparability in the way the budget is calculated, and also since non-participation leads to different interpretations of time-use choices, including involuntary leisure for the unemployed.

For singles, $w_i l_i = \alpha_t^l y_i$. I run this regression for men and women, finding the coefficients are approximately equal: $\hat{\alpha}_m^l \simeq \hat{\alpha}_f^l \simeq 0.6$.³⁸ I test the null hypothesis that $\alpha_m^l = \alpha_f^l$ and cannot reject the private leisure CD-SAP assumption at any of the usual significance levels. These findings suggest single working men and women satisfy private leisure C-D SAP, and hence that private leisure C-D SAP is plausible for working couples without cohabiting children. It is unclear why the preferences of men and women would change differently in this specific respect between being single (and working) and being in a couple (and still working). As

³⁶I note that this is very much not mechanical: estimates using clothing and non-market-work fall outside of the unit interval at the extremes.

³⁷A statistical test of cross-equation independence rejects at the 5% significance level

³⁸Not only are these preference parameters similar to each other but, comfortingly, they are higher than the estimated preference parameter for private leisure for working couples. We would expect that to be the case because joint leisure is not an option for singles, so we would expect singles to have a higher relative preference (higher preference coefficient, under the normalising assumption that preference coefficients sum to 1) for private leisure.

explained in section 9.4.2, this type of test appears to be highly informative for working couples. We may not be as sanguine about the validity of the test in other contexts, e.g. for couples with cohabiting children, since there are substantial gender differences in norms around childcare.

As an additional check, I also look at high-level patterns in the data on UK couples without cohabiting children. If private leisure C-D SAP holds then the proportion of household budget spent on overall private leisure (the sum of expenditure on female private leisure and expenditure on male private leisure) should not vary with the distribution of bargaining power. This is consistent with UKTUS data: the budget share spent on private leisure is mildly, and not statistically significantly, correlated with my estimated resource shares. I repeat the same check with the wage ratio, a frequently used proxy for bargaining power, yielding similar results. This provides additional reassurance that private leisure C-D SAP is a reasonable approximation in the context of this paper’s application.

9.3.2 Testing the public good C-D SAP assumption

Recall that, in order to estimate sharing from time-use data alone, I also assumed that preference parameters over all public goods sum to the same quantity for men and women. Equivalently, I can test whether the preference parameters over all private goods sum to the same quantity for men and women.³⁹

Having already tested that private leisure preferences are similar across types, I only need to test that preferences for aggregate material private good consumption are the same for men and women. I test this with data from a widely used UK expenditure survey (Living Costs and Food Survey, LCF, see section 10) for working singles.⁴⁰ For singles, $\sum_{j \in \Omega^c} p^j c_{i,h} = \left(\sum_{j \in \Omega^c} \alpha_t^j \right) y_{i,h} = \alpha_t^c y_{i,h}$. The test therefore requires regressing aggregate expenditure on material private goods on the budget, and checking whether the coefficients are the same for men and women. The coefficients are very similar in magnitude ($\hat{\alpha}_m^c \simeq \alpha_m^c \simeq 0.07$) and a test of equality cannot reject at any usual significance level. This test is consistent with the public good C-D SAP assumption.

It is unclear why the preferences of single men and women would change differently with regard to private v public expenditure when forming a couple without cohabiting children. Therefore we may feel reasonably confident that the singles test provides good evidence in

³⁹The latter implication is preferable for testing because data on private good consumption is less lumpy, and collected more accurately, in most expenditure surveys, including the one used here.

⁴⁰For singles there isn’t a distinction between private and public goods in practice, but we can still distinguish theoretically between goods which are non-rivalrous and those that aren’t, categorising goods in the same way for singles and couples.

favour of public good C-D SAP in the context of this paper’s application. As explained in section 9.4.2, this type of test appears to be highly informative for working couples.

I also perform additional high-level checks using data on couples, to provide additional evidence that public good C-D SAP does hold for UK couples. If public good C-D SAP holds then the share of household budget devoted to private goods should be independent of the distribution of bargaining power. Because private leisure C-D SAP entails that the share of household budget spent on private leisure is independent of the distribution of bargaining power, these assumptions together entail that the share of household budget devoted to private material goods is independent of the distribution of bargaining power. This is consistent with correlations in the LCF expenditure data. The proportion of budget spent on private material goods is mildly, and not statistically significantly, correlated with my estimated resource shares. This suggests internal consistency of my approach. As a further check, I also repeat this exercise with the ratio of female-to-male wages as a commonly used proxy for household sharing. Again, the correlation is mild and not significant, suggesting that public good C-D SAP is a reasonable simplification in the context of this paper’s application.

As discussed for private leisure C-D SAP, while public C-D SAP appears to be a reasonable approximation for couples without cohabiting children in the UK, it is important to consider the context when evaluating this assumption. For instance, we may be less willing to make this assumption in contexts with cohabiting children as gender differences are more likely to emerge in that context.

9.3.3 Homotheticity tests

Cobb-Douglas is clearly a strong simplification of underlying preferences, and we may be particularly concerned about imposing homotheticity. I note that I am not imposing homotheticity on all goods, but only at a high level between private leisure and other consumption (and between overall private and public consumption if relying on public good SAP too).

I test homotheticity of private leisure by estimating singles’ private leisure expenditure separately for higher and lower income singles. I cannot reject equality of coefficients at any of the usual significance levels. Similarly, for singles, the correlation between the budget share of private leisure and their budget is mild and not significant. As a further check, I estimate resource shares separately for higher and lower budget couples. Parameter estimates are broadly similar to the baseline, and the resulting estimated resource shares are highly cor-

related (correlation of 0.85)⁴¹ with the baseline estimates, suggesting that the homotheticity assumption is not strong-arming results. Finally, I relax the homotheticity assumption by estimating the sharing rule under the assumption of Stone-Geary preferences.⁴² Comparing the results to the baseline estimates, the direction of marginal effects is the same, and the estimated resource shares are very highly correlated (correlation of 0.97) with baseline estimates. Overall, it appears that, in this context, assuming homotheticity is not strong-arming results.

9.3.4 Internal consistency of estimates

An additional advantage of Cobb-Douglas preferences is that they have a simple, explicit utility representation which can be used to sense-check estimates. In Appendix F.6, I set out a back-of-the-envelope method to check that the estimated sharing rule parameters are quantitatively, as well as qualitatively, consistent with bargaining theory. Consider the example household discussed in section 9.2.1, characterised by sample-averages for all characteristics apart from the female wage, which is £15 instead of £9.87 an hour, and hence by a female share of 0.51 as opposed to the average female share of 0.45. We wish to check that the implied utility is higher for the higher-earning female than for the average female, as she has a better outside option. This is confirmed by a back-of-the-envelope check: the higher-earning female has slightly lower-than-average hours of leisure, and more material consumption, with the latter effect outweighing the former.⁴³ The fact that leisure is lower for her means the price effect outweighs the income effect (which here is a double effect through the household budget and through the sharing rule). Similar findings hold for men, and for a reasonable range of possible wage changes.

A further check of internal consistency is discussed in section 9.4, where I am able to replicate the same gendered pattern in preferences for clothing observed in singles data by structurally estimating clothing preference parameters for couples, using my sharing estimates.

⁴¹This is for the sample excluding outliers. The correlation coefficient including outliers is still high but falls to 0.75 as each of the two sub-samples only contains outliers in one direction.

⁴²In order to avoid colinearity issues due to the lack of accurate non-labour income data, this version of the estimates assumes all goods are private and that subsistence levels for private leisure are equal to zero (while they can be non-zero for other goods).

⁴³Different papers find different results on the relationship between wages and leisure, see Browning, Donni and Gørtz (2020). Those authors find that leisure, both private and public, is broadly stable across the wage distribution of both members of the couple. This is consistent with high-level correlations in UKTUS, though it is important to remember that there is significant assortative matching in the data and hence that these high-level correlations conflate the effects of female and male wages.

9.4 Comparison to other approaches

Reassuringly, my estimated resource shares are positively (correlation coefficient of 0.76) and very significantly correlated with the ratio of female to male wages,⁴⁴ which is often used as a proxy for sharing in the household. As further checks, I re-estimate resource shares using two alternative assignable goods: (i) non-market-labour time, and (ii) clothing, and compare the findings. Finally, I consider how my findings relate to existing estimates in the literature on UK data.

9.4.1 Non-market-work hours as the assignable good

I repeat my analysis using non-market-work as the assignable good. I construct it from a recall question on usual hours worked from the LCF expenditure data (see section 10).




Figure 2: Estimated female shares using different assignable goods

As can be seen in table 1 and further discussed in section 10, while the expenditure and time-use data are comparable, the LCF sample is characterised by somewhat higher wages and education. Therefore, to construct the figures comparing the distribution of resource shares for different assignable goods, and to estimate correlation coefficients between different sharing estimates, I compare estimates of sharing in the LCF obtained from different assignable goods. Non-market-work and clothing estimates are constructed directly in the expenditure survey, and for my baseline estimates I use my baseline sharing rule applied to LCF (baseline female shares are marginally higher than in UKTUS).

My baseline results and estimates using non-market-work as the assignable good are positively correlated, with a correlation coefficient of 0.76. As can be seen in figure 2, the whole

⁴⁴Defined as the ratio of female wages to the sum of female and male wages.

distribution is shifted towards equal sharing relative to my baseline estimates, and the mean female share is estimated at 0.49, substantially higher than the 0.45 from my baseline results. As discussed in more detail in online appendix G, and shown in table 4, the difference between this estimate and my baseline estimate is statistically significant. This difference is consistent with the conceptual reasons set out in section 8.1 why non-market-work may lead to estimates biased towards equal sharing. This bias may be even greater in other applications, with cohabiting children and non-participation, where we may expect that women spend substantially more time than men on childcare and domestic work.

9.4.2 Clothing as the assignable good

I repeat my analysis using clothing as the assignable good, using the LCF expenditure data. Clothing is of interest because it has been used as the assignable good almost universally in this literature. I find a mean female resource share of 0.63, substantially higher than my baseline finding of 0.45. As can be seen in figure 2, the whole distribution of female resource shares is shifted upwards, and the extremes of the distribution violate the testable restriction that resource shares lie in the unit interval. Using singles data, I find that women’s preference parameter for clothing is approximately twice as large as men’s, and the identifying assumption (clothing CD-SAP) is rejected at the 5% significance level. As a further check, I also estimate clothing preference parameters *for couples*, by running structural regressions of clothing expenditure on estimated individual budget computed using my baseline sharing estimates. This also yields a female preference parameter for clothing which is approximately twice the male parameter, as for singles. Therefore, I adopt an adjusted identifying assumptions, assuming the female-to-male ratio of preference parameters estimated in singles data holds for working couples (clothing C-D SRAT). The resulting distribution of adjusted clothing resource shares is much closer to the baseline, as can be seen from figure 2, and the mean resource share for women falls to 0.46. This is encouraging, in that it suggests that tests of identifying assumptions using singles data are highly informative for working couples.

As discussed in more detail in online appendix G, and shown in table 4, the confidence intervals for the clothing estimates are very wide, and partially overlap with the confidence intervals for my baseline estimates and non-market-work estimates. Leaving aside the noisiness of the clothing estimates, the fact that the estimated female share from clothing is higher than that the one from private leisure, even after the adjustment to the identifying assumption, may be explained by the conceptual reasons set out in section 8.1. However, a full comparison is left to future work - in particular, I note that my implementation with

Cobb-Douglas preferences is not a good fit for clothing, which is not a homothetic good. In singles' LCF data I find that there is a significantly negative correlation between budget share spent on clothing and household budget, suggesting clothing is a necessity good. Potentially for this reason (or perhaps for other reasons, such as the high proportion of clothing zeros in my data), using clothing, the estimated marginal effects of household characteristics on sharing are at odds with bargaining theory (female hourly pay decreases female resource shares, and male hourly pay decreases male resource shares), the estimates are noisier, and the estimated shares are negatively correlated with the wage ratio (commonly used as a proxy for sharing).

9.4.3 Estimates from other papers with UK applications

Lise and Seitz (2011) use non-market-work as the assignable good, and their framework differs from this paper in several respects, including parametric specification, identifying assumption,⁴⁵ and estimation approach.

Their estimated share for the last cohort in their data, born in the 1960s, has a female resource share of 44.2% on average. Their estimates include couples with non-participation, so are not directly comparable: on average, we would expect women to have lower resource shares in households with non-participation. We might expect a partial cancelling out of the opposite effects of (i) sample differences, and (ii) non-market-work leading to some degree of bias towards equal sharing. Hence, the similar magnitude of the estimate is reassuring as it suggests that the results can be reconciled based on sample differences and differences in the assignable good used, notwithstanding the other substantial methodological differences.

However, the results of Bargain, Donni and Hentati (2022), using clothing as the assignable good are substantially different. They find the average resource share for women in heterosexual couples without children (including non-participants) in 1978-2007 is 51.7%. This is an estimate for the pooled 1978-2007 sample, with an upward trend over time, implying a higher estimate for 2000-14. Therefore, it appears that the approach in Bargain, Donni and Hentati (2022) (which differs from this paper in several respects, including parametric specification and estimation approach) leads to much higher estimates than the adjusted clothing approach discussed above. Moreover, it differs substantially from my baseline estimates using private leisure as the assignable good. Qualitatively, my baseline result that, on average, women have a lower resource share than men is reversed. It would be surprising

⁴⁵Lise and Seitz (2011) make the 'symmetry' identifying assumption that women and men with the same potential earnings have the same resource shares. Empirically, this assumption is not supported by my findings. A woman with an hourly wage equal to the average male wage in the sample is estimated to have only a 47% resource share, substantially below the average male resource share of 55%.

if women had higher resource shares than men on average in a society which, while comparatively gender progressive, still has a patriarchal tradition and norms. More systematic and thorough investigation of how different approaches compare is a priority for future work, to enable additional progress in this field.

10 Individual-level consumption for UK couples

Having estimated the sharing rule, it can be applied to other comparable datasets to investigate additional questions. A natural goal is to estimate individual-level consumption inequality. To do so, I apply my estimated sharing rule to a UK expenditure survey (LCF). Both my time-use and expenditure datasets are representative of the UK population, and they were collected at a similar time. Even so, as shown in table 1, the LCF sample of heterosexual working couples is characterised by somewhat higher wages and education than the UKTUS sample. This is not a problem, since when the sharing rule estimate is applied to the expenditure data the process explicitly takes into account differences from the sample averages in the time-use data. As expected, the distribution of my baseline sharing estimates are similar in UKTUS and LCF, with very slightly higher female shares in LCF. Applying the sharing estimates to recorded household expenditures, I find substantial gender gaps between men’s and women’s material consumption. I also estimate gender gaps in ‘full’ consumption, and discuss policy implications.

10.1 Methodology

I measure consumption inequality based on estimates of individual-level consumption, defined as the monetary market value of an individual’s consumption.⁴⁶ Similarly to Lise and Seitz

⁴⁶This terminology is somewhat imprecise since I am not referring to the consumption bundles of different individuals. However, it captures the essence of this measure, which compares the value of different consumption bundles at their market prices. I avoid the term individual-level expenditure because this could be confused with a different approach, taking into account individual Lindahl prices for public goods. For instance, Chiappori, Meghir and Okuyama (2024) define the Money Metric Welfare Index (MMWI) as the expenditure that would be necessary for an individual to achieve the same level of utility they currently enjoy (i) without access to household economies of scale for public consumption, (ii) with the possibility of varying their chosen bundle of goods relative to what they are factually consuming, but (iii) keeping their preferences unchanged. We can conceptualise this approximately as the expenditure that would be necessary for an individual to achieve, as a single, the same level of utility they currently enjoy in a multi-person household (the caveat is that preferences are not adjusted to counterfactual single-hood preferences, and SAT is not assumed).

Similarly to this paper, Chiappori, Meghir and Okuyama (2024) work within a collective household model with Cobb-Douglas preferences. Unlike this paper, they do not estimate intra-household sharing but directly observe individual-level material consumption in the data. Depending on the available data, and the assump-

(2011), individual-level consumption is calculated as the sum of household material public expenditure and the individual’s share of the household’s material private expenditure. This metric allows me to compare the objective (preference-independent) value of material consumption of different individuals, and is a clearly policy-relevant measure.

Relative to existing approaches, I make four contributions. Firstly, I propose an approach to including large durables, such as housing and cars, by imputing rents for them to reconcile their dynamic nature with the static household model. Several papers in the literature exclude these goods from measures of individual consumption, but these are the largest expenditures in my UK data, making it desirable to take them into account. Secondly, I suggest drawing on usage data to estimate the degree to which partly public goods, such as cars, are private vs public. Thirdly, I propose standardising prices of goods across regions for goods with substantial geographic variation in prices, such as housing. Finally, this paper is the first to also provide estimates of ‘full’ consumption, inclusive of time-use as well as material consumption.

Individual-level consumption, defined as the monetary market value of an individual’s material consumption, is the correct metric to measure consumption inequality only when prices of goods are reasonably similar for people in the sample. For material goods, individuals in the same country generally face similar market prices for the same goods (although the composition of consumption varies between individuals).⁴⁷ The exception is housing, which in a country like the UK is priced very differently across regions. That price heterogeneity is relevant for the purpose of estimating wealth inequality, but should be abstracted from in estimating consumption inequality. For instance, someone living in London may receive a higher salary than someone living in Manchester, spend a larger proportion of that salary on a home of the same underlying quality, and spend the same amount on other consumption. Then the person living in London would have higher wealth (if they own the home) but the same level of consumption as the person living in Manchester. Similarly, when thinking about ‘full’ consumption, we may wish to standardise the price of time by using the average wage, rather than individual wages.

Below, I summarise my proposed methodology to estimate individual-level material consumption, and how I implemented it using UK data. Further detail is provided in online appendix C. In section 10.2.1, I extend this to ‘full’ consumption, inclusive of time-use.

1. A cross-section household expenditure dataset is required.

tions made to estimate the sharing rule, it may be possible to use the MMWI instead of, or in addition to, individual-level consumption measures. In the context of this paper, I focus on individual-level consumption because, being preference-agnostic, it is a more natural concept to measure individual-level poverty.

⁴⁷For instance, see DellaVigna and Gentzkow (2019) on uniform pricing across the U.S.




Figure 3: Individual-level consumption

- I use the LCF (Living Costs and Food Survey)⁴⁸; a high-quality, large-scale survey that is used to estimate official government statistics.
2. If the expenditure data is contained in a separate dataset from the estimating dataset (e.g. a time-use dataset) apply the estimated sharing rule to the expenditure dataset. It is important that the dataset used to estimate the sharing rule and the expenditure dataset be comparable, e.g. nationally representative data for the same country in the same year, although they needn't have exactly the same distribution of relevant variables, as any such differences are explicitly accounted for in the methodology. The data should include any household characteristics and member's characteristics which influence the sharing rule (e.g. members' wages, ages, etc.)
 - The LCF, like UKTUS, is nationally representative for the UK. I use data for the year 2014,⁴⁹ and focus on the sub-sample of working heterosexual couples without cohabiting children to ensure comparability with the UKTUS data used for resource share estimation. The LCF contains information on individual and household characteristics, as well as detailed income data, which allows me to estimate household-specific resource shares using my estimated sharing rule. The distribution of resource shares in the LCF is similar to that in the UKTUS, and the mean female resource share remains 45%.
 3. In the expenditure dataset, divide household expenditure into public and private expenditures. Moreover, some outflows should be excluded if they are incompatible with the static nature of the model (e.g. saving, insurance and investments). To enable

⁴⁸Office for National Statistics. (2019a). Living Costs and Food Survey [data series] 3rd Release. <http://doi.org/10.5255/UKDA-Series-2000028>

⁴⁹The sharing rule was estimated using data from 2000 and 2014. I do not additionally analyse LCF data from 2000 as there were substantial changes in the dataset which undermine the feasibility of accurately pooling the 2000 and 2014 LCF datasets.

accurate classification of expenditures, the dataset should record household expenditure by sufficiently granular categories. Ideally, the data would allow inclusion both of durable and non-durable expenditure. Additionally, it is helpful if the data contains information on vehicles owned and the characteristics of the home the family lives in to enable imputing rent and lease prices consistently across the sample. This is important because expenditures on purchasing or renting a home, and purchasing vehicles, are extremely large in relative terms (especially in some contexts, including the UK).

- The LCF contains household expenditure divided into very granular categories (COICOP codes). In addition, the LCF records a two-week expenditure data, recall questions on infrequent expenditures, data on the number of vehicles owned, and the characteristics of the home the family lives in.
4. In the expenditure dataset, estimate individual-level material consumption for each household member by summing (i) public good expenditure and (ii) private good expenditure weighted by the individual’s resource share: $C\hat{C}_{i,h} = \eta_{i,h} \left(\sum_{j \in \Omega^c} p^j c_h^j \right) + \left(\sum_{j \in \Omega^x} r^j X_h^j \right)$ This yields a distribution of individual-level material consumption, which can then be used to estimate different measures of consumption inequality.

10.2 Findings and policy implications

For heterosexual working couples without cohabiting children, the average gender gap in consumption is 8.53%.⁵⁰ Note that this gap is substantial, and would have gone completely undetected by the standard ‘per capita’ approach to estimating consumption inequality because it is fully driven by intra-household inequality. The reason this gap is smaller than the average gender gap in resource shares is that the sharing rule only applies to private expenditures, while consumption is also inclusive of public expenditure. I also estimate two common measures of inequality: the Gini coefficient is 0.21, and 6.43% of the group live in relative poverty (consuming less than 60% of the median individual-level consumption). My baseline estimates of inequality are low relative to standard estimates of inequality, partly because I focus on the sub-sample of working couples without cohabiting children. This excludes sources of inequality such as non-participation. Moreover, Bargain, Donni and Hentati (2022) suggests that women’s resource shares in the UK are lower in couples with children than without. Both of these factors suggest that my estimates are underestimates of inequality for the broader UK population.

⁵⁰This gap is defined as the difference between mean male consumption and mean female consumption, divided by mean male consumption.

A further driver of my inequality estimates being lower than some of the available estimates is that I attempt to fully account for public expenditure, and assign all public expenditure to each household member. Focusing only on private consumption, or treating voices of public expenditure as private, substantially overestimates inequality (especially absolute measures of poverty, as these mechanically increase if some expenditures are excluded). Excluding public expenditure, the Gini coefficient increases to 0.30 and the poverty rate to 15.95%. When treating public expenditure as private (similarly to the approach in LPW), the Gini coefficient increases to 0.24 and the poverty rate to 9.01%. This suggests it is important for us to properly take public expenditure into account to avoid over-estimating inequality.

Since women have lower resource shares than men, especially in lower income households, the poverty rate for working heterosexual couples is higher for women (7.03%) than for men (5.83%); a gender gap of 20.59%.⁵¹ As previously noted, this is a lower-bound estimate for the population-wide gender gap in poverty, and an even greater underestimate of the level of intra-household inequality we are likely to find in other countries.

10.2.1 Individual-level full consumption (including time-use)

We may be interested in the monetary value of individual-level consumption including time-use. Following Becker (1965) and Lise and Seitz (2011) I term this ‘full’ consumption.

When we include private leisure expenditure, full individual-level consumption is given by:⁵²

$$T\hat{C}_{i,h} = w_{i,h}l_{i,h} + \hat{\eta}_{i,h} \left(\sum_{j \in \Omega^c} p^j c_h^j \right) + \sum_{j \in \Omega^x} r^j X_h^j$$

When time-use and expenditure datasets are separate, it may still be possible to estimate full individual-level consumption, as discussed in online appendix F.5 for the case of Cobb-Douglas preferences. I apply this methodology to the LCF data and find that the average gender gap in individual-level consumption is higher (11.3%) when taking into account time-use. While with material goods the prices of goods are generally the same for all households (even though the composition of consumption varies), the price of time varies at the individual level, making time-inclusive measures of individual-level consumption harder to interpret at face value. I re-estimate individual-level consumption using the average wage, instead of own-wage, as the price of time. This lowers the average gender gap inclusive of time because men’s wages are higher than women’s on average. However, the age gap

⁵¹Defined as the female poverty rate minus the male poverty rate, divided by the male poverty rate

⁵²Ideally, we would also include expenditure on joint leisure and public time $T\hat{C}_{i,h} = w_{i,h}l_{i,h} + \hat{\eta}_{i,h} \left(\sum_{j \in \Omega^c} p^j c_h^j \right) + \sum_{j \in \Omega^x} r^j X_h^j + \sum_i w_{i,h} (d_{i,h} + jt_{i,h})$. This exercise is left to future work as it requires careful analysis to disentangle time spent on joint leisure and time spent on domestic activities, which is beyond the scope of this paper.

estimated in this manner (10.1%) is still higher than the measure excluding time-use.⁵³

These findings confirm the importance of including time-use to fully account for gender gaps in consumption. Further work is needed to carefully incorporate time-use into inequality measures, also considering public time-use and potentially accounting for complex issues such as any dis-utility from domestic work.

10.2.2 Policy implications in the UK setting

These findings suggest that policymakers should take into account intra-household inequality in targeting and evaluating policies. More accurate estimation of consumption inequality is relevant to a range of policy decisions. For instance, one of the United Nations' Sustainable Development Goals for 2030 is the eradication of extreme poverty. A recent enquiry in the UK concluded that foreign aid must be targeted to more closely align with this goal, and that as part of this it is important for the impact of aid on poverty be measured and kept track of. Accurate estimates of poverty rates in this context can affect where foreign aid is targeted, and even whether the foreign aid budget is cut if it is seen as being ineffective. Moreover, as part of its commitment to the goal of eradicating poverty, the UK aims to halve its poverty rate by 2030. Accurate measurements of poverty would help both track progress with this goal, and evaluate which policies may be most impactful, for instance by targeting women specifically.

To reduce female poverty rates, policymakers may consider targeting income transfers to women specifically, especially in deprived households. Another way of targeting help to females is to expand policies subsidising products which are likely to disproportionately help women. This includes providing free or subsidised childcare for young children, especially for lower-earning households; something the UK has been moving towards, with the introduction of 30 hours of free childcare for children aged 9 months and above, from September 2025. Since women often take primary custody of children after divorce, free childcare improves the outside option for women relative to men, reducing asymmetries in bargaining power underlying unequal sharing. Free childcare can also play a role in narrowing gender wage gaps and can also equalise sharing through this channel. Another initiative that can help reduce female poverty is lowering the cost of female sanitary products. In the UK, the VAT rate on these products was dropped to zero in 2021, and as part of the 'Period Product Scheme', schools provide them for free since 2020. Expanding this policy to free provision for all women would be in line with the current UK policy of providing contraceptives, including

⁵³Note that this does not contradict the identifying assumption that men and women spend the same share of their individual budget on leisure. We detect more inequality when taking leisure into account because leisure is a private, rather than public, good, so that it is shared unequally between members.

the pill, free of charge to all. Moreover, policymakers may consider that an added advantage of pursuing policies to narrow gender pay gaps is their equalising effect on intra-household sharing (this paper, in line with the literature, finds that sharing is more equal in households with higher female wages and/or lower male wages).

10.2.3 Gender health gaps in the UK

An additional reason why this paper’s findings are relevant to UK policymakers is the impact of poverty on health problems, and the fact that women suffer from worse health outcomes than men in the UK as measured by metrics such as the morbid obesity rate. Decreasing intra-household gender-based inequalities in sharing may also help reduce health inequalities at a national level. For instance, women living in poverty are more likely to suffer from anxiety disorders (Remes et al. (2017)) and to be obese (Griffith (2022)). As noted by Griffith (2022), there are several reasons why living in deprivation can lead to obesity, and in order to tackle high obesity rates it is important for policies to target the right people. Women in poverty are likely to have little time, little money, and live in ‘food deserts’ with lower availability of healthy foods. Hence they are likely to consume unhealthy, cheap, store-bought foods, and not be able to exercise regularly.

Obesity is a common target for policymakers because of its wide-ranging impacts on physical and mental health, as well as mortality. For women specifically, obesity is also linked to gestational problems and adverse outcomes for children’s health. Interestingly, in the UK, women are more likely than men to be morbidly obese. This is consistent with this paper’s finding that women are more likely than men to be living in poverty. Moreover, in the UK, regional deprivation and household income are stronger predictors of female obesity than male obesity (The King’s Fund (2021)). This is consistent with this paper’s findings that (i) men have higher resource shares than women, and (ii) this is more marked in poorer households in poorer regions, so that household income and regional deprivation are better predictors of individual-level poverty for women than for men. Beyond obesity rates, while in the UK women have a longer life expectancy than men, the Women’s Health Strategy for England policy paper notes that women “spend a significantly greater proportion of their lives in ill health and disability when compared with men” (Department of Health & Social Care (2022)). Targeting cash transfers to women in poor households, and labelling them as meant for expenditure on healthy foods and exercise for women may be particularly helpful in tackling obesity as well as poverty.⁵⁴

⁵⁴See Beatty et al. (2014) for an example of the effectiveness of labelling cash transfers in the UK.

11 Conclusion

My findings add to the growing literature on the importance of estimating intra-household inequality and the importance of considering the dimensions, such as gender, as well as the levels, of inequality. The approach set out in this paper is grounded in the well-established collective model of the household, has a clear source of identification, and is straight-forward to implement with widely available data. Moreover, while the approach is simple, tests of model fit in my application to UK data suggest that it is empirically sound. This approach is also well-suited to easily estimating a variety of measures of individual-level well-being, including individual-level material and ‘full’ consumption, allowing for important policy applications. These characteristics may contribute to the goal of facilitating adoption of individual-level inequality measures outside of academia. The empirical finding of substantial gender inequality within households, even in the context of the UK, one of the most gender equal countries in the world, strongly suggests that the widely used equal sharing assumption is not fit for purpose, and that it is empirically important for policy to be guided by measures of consumption inequality that take into account unequal sharing within households.

To further aid this goal, additional work is required to investigate the empirical performance of the different approaches that have been proposed in the literature. Bargain, Lacroix and Tiberti (2021) test a specific approach to estimating intra-household sharing against reported individual-level consumption. Future work might systematically compare the performance of different approaches using simulated and empirical data, as well as comparing their more general advantages in terms of applicability. This exercise would be very valuable in a literature which has been growing in several different directions, and may benefit from consolidation into something like a mainstream approach. In order to focus data collection efforts, it would be particularly useful to understand the partial discrepancy between estimates obtained from different assignable goods.

Another important direction for future work is to extend the approach in this paper to contexts with non-participation. Identifying the sharing rule from time-use data for non-participants requires estimating the unobserved price of their time (which is higher than their, unobserved, wage). This would allow one to appropriately take into account households with children and people who do not supply market work. For those households, the tests of identifying assumptions provided in this paper (which appear to be highly informative for working couples) may not be sufficiently convincing. As illustrated with clothing-based estimates of sharing in the UK application, the choice between different identifying assumptions can have a large impact on the magnitude of estimates and even on qualitative conclusions. Hence, it may also be important to work toward reduced reliance on identifying assumptions, or new

approaches to testing them credibly for a wider range of household compositions.

Finally, this paper is the first to provide estimates of ‘full’ individual-level consumption inequality. These are supportive of the widely held view that gender inequality may be underestimated when only considering material consumption. Further work is required to consider how to best measure ‘full’ inequality, carefully overcoming the additional difficulty relative to material consumption that the price of time varies at an individual level.

Part II

Female Empowerment and Emissions

12 Introduction

Households play a crucial role in global greenhouse gas (GHG) emissions. The direct and indirect emissions associated with household consumption account for more than 60% of global GHG emissions (Ivanova, Stadler et al. 2016), and almost three quarters of the global carbon emissions (Druckman and Jackson 2016). Policies aimed at mitigating emissions due to household demand could reduce global GHG emissions in end-use sectors by up to 70% relative to the Intergovernmental Panel on Climate Change (IPCC)’s baseline scenarios (IPCC 2022). Therefore, measuring households’ emissions, and understanding its determinants, are crucial to identifying behavior-changing policies to meet urgent climate objectives. Henceforth, we refer interchangeably to “carbon footprint” and “GHG emissions”, since our paper follows the standard approach in the literature of converting greenhouse gas emissions into CO₂-equivalent emissions using equivalence scales for different gases.

There is a growing literature suggesting that women express greater environmental concern and are more likely to engage in pro-environmental behaviors than men (e.g., Baraldi and Fosco 2025; De Rock and Le Henaff 2023; Kennedy and Kmec 2018; May et al. 2021; Xiao and McCright 2015). This high-level pattern holds across a range of settings relating to stated attitudes, infrequent decisions, and political choices. However, there are only a few papers that have investigated the degree to which actual individual daily consumption patterns vary by gender (Carlsson Kanyama et al. (2021), Osorio et al. (2024), Toro et al. (2024) and Büchs and Schnepf (2013)). To disentangle gender effects, those papers either focus on singles or perform reduced-form analysis with simple, discrete proxies of whether observed consumption and emissions patterns mostly reflect the preferences of men or women in the household, i.e. whether the household has a female or male “leader” or “head”. Even though most of the global population lives in multi-person households, where decisions are taken collectively, we are not aware of any study that has explored the relationship between the distribution of bargaining power within the household and households’ carbon footprints using a structural approach. Such an approach is required to carefully estimating the degree to which collective decisions represent the preferences of men or women.

This paper proposes to fill these gaps. Using a structural method, we investigate the relationship between gender and greenhouse gas emissions both by (i) analysing singles data

by gender, and by (ii) estimating the relationship between female bargaining power within heterosexual couples, and the household’s carbon footprint. To do this, we build on recent developments in household economics, extending them to a setting with emissions. To our knowledge, this is the first paper to use a household bargaining model to answer a question in environmental economics. Our novel approach allows us to contribute to the discussion of gender differences in pro-environmental behaviours, and also to the debate on the potential benefits of policies aimed at empowering women within households. We also contribute to the household literature, by demonstrating the potential for household bargaining models to be applied much more widely than they have to date, including to environmental questions. We model household bargaining through the static collective household model, which is widely used in family economics, and enhance it with a framework to consider the greenhouse gas emissions associated with consumption choices. To estimate bargaining power, we propose an approach which is closely related to that developed in the first chapter of this thesis, building on contributions by Dunbar et al. (2013) and Lechene et al. (2022). Using individual-level variation in time-use data, from the UK Time Use Survey (UKTUS), we estimate bargaining power as a function of household characteristics for UK heterosexual working couples without cohabiting children, as a function of household characteristics. The intuition behind identification is the following. We observe that variation in the household budget is associated with more variation in men’s expenditure on private leisure (time spent on leisure activities without other household members co-present) than women’s. Combined with an identifying assumption that men and women have similarly strong preferences for private leisure,⁵⁵ this implies that men have more of a say than women on how to spend household budget, i.e. that they have more bargaining power.

With estimates of bargaining power at hand, our structural model yields estimating equations which we can use to recover men’s and women’s preference parameters on different categories of goods, with different associated emissions. To do this, we combine (i) granular household expenditure data for the UK over the period 2001-2014, from the Living Costs and Food Survey (LCF), with (ii) detailed emissions data for different categories of goods in different years published by the Department for Environment, Food & Rural Affairs (DEFRA). The DEFRA emissions data was estimated using Environmentally-Extended Input-Output tables, and is the best available source of information of the total direct and indirect emissions associated with purchasing goods in the UK. Combining expenditure surveys with data on emissions multipliers is the current standard approach in the environmental literature, and is the approach taken by other papers that have attempted to estimate gender differences in emissions propensities (Carlsson Kanyama et al. (2021), Osorio et al. (2024), Toro

⁵⁵See the first chapter of this thesis for a detailed discussion of why this assumption is valid

et al. (2024) and Büchs and Schnepf (2013)).

Taking our estimated bargaining rule (the relationship between household characteristics and bargaining power) from the UKTUS to our merged LCF and DEFRA data, we structurally estimate gender-specific preference parameters for different goods, with different associated emissions. For both singles and couples, we find that women have lower emissions propensities than men. By this, we mean that women’s preferences are conducive to lower greenhouse gas emissions from consumption than those of men, per GBP of budget.

We find that if the average UK household transitioned from our estimate of current bargaining to equal bargaining, this would reduce emissions from UK couples by 2.1%. On a per-household basis, this is a sizeable effect: it is in a similar order of magnitude to estimates of the impact of the average North American household switching to a vegetarian diet (Ivanova, Barrett et al. 2020). Re-weighting our estimate by the relevant proportion of UK households, and the proportion of UK greenhouse gas emissions which are (directly and indirectly) attributable to households, we estimate that transitioning to equal bargaining could, alone, reduce the UK’s overall carbon footprint by 1.1%. This suggests that policies aimed at empowering women within households, such as narrowing gender pay gaps, may have substantial additional benefits in terms of reduced greenhouse gas emissions.

Our more granular findings for specific consumption categories partly differ across household compositions. This is in line with other evidence from the household literature, e.g. Hubner (2020), that preferences are not stable across household compositions. This reinforces the importance of our approach to assessing gender differences in couples, as opposed to inferring these directly from singles data. It is also a strength of our approach relative to papers in the environmental literature, which typically jointly analyse a variety of household compositions.

As a robustness check, we repeat our couples analysis using alternative measures of bargaining power, including simple proxies of whether the household has a female or male lead, following papers in the environmental literature. We still find some partial evidence of women having more environmentally friendly preferences than men, but the findings are weaker, often not statistically significant. This suggests the value of drawing on structural approaches from the household bargaining literature, instead of simple proxies of bargaining.

The remainder of the paper is structured as follows. The next section provides a brief overview of the related literature and underlines our contributions. In section 14, we set out our model of household decision-making. Section 15 outlines our methodology to estimate bargaining power and gender differences in preferences. In section 16, we describe the UK data used in this paper. In section 17, we report and discuss the results of our analysis. Section section 18 concludes.

13 Related literature

To date, there are very few papers answering the question of whether actual individual daily emissions patterns from consumption vary by gender. One paper, focusing only on Swedish singles data, finds that women have lower total emissions than men (Carlsson Kanyama et al. 2021). Other papers have considered broader household compositions, using simple proxies for the degree to which a household’s behaviour should be taken as indicative of the preferences of men or women.

Using Spanish data, Osorio et al. (2024) use as a proxy the female share in households. They define this as the proportion of members older than 13 who are women. This amounts to assuming that bargaining power is divided equally between household members, which has been strongly rejected in the household literature (e.g. Lechene et al. (2022)). Controlling for potential socio-economic confounders, they find that households with a higher female share have (i) lower total household carbon footprints, (ii) lower carbon intensity, (iii) higher emissions on housing and food products, and (iv) lower emissions on restaurant and transport services. Toro et al. (2024) also investigate gender differences in greenhouse gas emissions in Spain, but using a different proxy for female bargaining power. They define the breadwinner in a household as the individual who contributes the most to the household income. They find that female breadwinner households have lower total GHG emissions, primarily driven by lower emissions from private transportation.

Büchs and Schnepf (2013), using UK data, follow a similar approach to proxying for bargaining power: they define the head of the household as the individual financially responsible for accommodation or, in cases of equal financial contribution, the higher earner. Similarly to other studies, they find that female-headed households are less likely to have high emissions associated with their consumption. A more complex picture emerges when they condition on different variables: they then find that female headed households have conditionally higher overall CO₂ emissions, but lower emissions for transport, primarily due to reduced motor fuel consumption. As discussed in section 17.2.4, the partial differences between our findings and those of Büchs and Schnepf (2013) are due to differences in our methodological approaches and the emissions data used. Our analysis benefits both from our more accurate estimation of bargaining power and from more recent, higher quality, emissions data.

Our findings support the broad consensus in this literature that women have lower emissions propensities than men. Our main contribution relative to previous approaches is that we structurally estimate gender differences using an approach grounded in a general household bargaining model. Instead of using simple, discrete, wage-based proxies, our methodology combines a general theoretical household model, data on observed patterns of time-use,

and data on detailed characteristics of household members, including not only wages, but also, for instance, educational attainment. In this way, we provide a more reliable, granular (continuous rather than discrete), theory-consistent, measure of bargaining in the household. This allows us to estimate gender heterogeneity in emissions propensity more accurately.

Our work also contributes to the literature on estimating intra-household bargaining. Firstly, by extending the approach in Dunbar et al. (2013), Lechene et al. (2022) and the first chapter of this thesis, we show that it is possible to estimate intra-household bargaining in a manner which is both grounded in a general household model and simple to estimate from widely available data. Secondly, estimates of intra-household bargaining have mostly been applied to estimating individual-level consumption, inequality and poverty measures. For example, see Browning, Chiappori and Lewbel (2013), Dunbar et al. (2013), Lechene et al. (2022), Lise and Seitz (2011), Bargain, Donni and Hentati (2022) and Calvi (2020). There has been little work on how female empowerment relates to different policy-relevant consumption patterns. This has been limited to two applications to date: investment in children’s development, e.g. Blundell, Chiappori and Meghir (2005) and Cherchye, De Rock and Vermeulen (2012) and household portfolio composition (Thörnqvist and Vardardottir 2014). This paper explores a new question: how does female empowerment relate to the emissions of households? Apart from being interesting in itself, this provides an illustration of how methods to estimate intra-household bargaining may have wider applications than have so far been explored in the literature.

14 Framework

We set out a structural model of household decision-making and augment it with good-specific emissions parameters to formalise how gender differences in preferences translate into different emissions patterns. For singles, household consumption choices take the form of individual utility maximisation. For larger households, we need to additionally consider intra-household bargaining to understand how the preferences of different individual members are mapped onto observed household-level consumptions and emissions patterns.

14.1 Model

This paper is grounded in a static collective model of the household with both private and public goods and both material goods and time-use. The collective model of the household is both more general, and provides a better empirical fit, than alternatives in the literature

(e.g. see Browning, Chiappori and Weiss (2014)). For a more in-depth discussion of this model, and of a more general version of it, see the first chapter of this thesis.

We focus on two household compositions $g = \{S, C\}$.⁵⁶

- **Working singles (S).** Each single person is categorised as either male or female $t = \{f, m\}$ and is indexed by a household number belonging to the set of single households $h \in H^S$.
- **Heterosexual working couples (C).** Each of these household has two members, indexed by their type $t = \{f, m\}$ for female and male. Each couple is indexed by $h \in H^C$.

The household purchases both rivalrous (e.g. bread) and non-rivalrous (e.g. heating) material goods on the market. We refer to the conceptually rivalrous goods as private and the conceptually non-rivalrous goods as public, even for single households, where in practice there is no distinction between these types of goods. Additionally, individuals enjoy (i) private leisure, i.e. time spent on leisure activities without other household members co-present (for singles, all leisure time is private leisure), (ii) joint leisure, i.e. time spent on leisure activities with other household members co-present (this is not available to singles), and (iii) a domestic good produced by household members' domestic work. Individuals also spend time on paid work, but, as is standard, this is assumed not to directly enter the utility function.

Denote the consumption vector of a person of type t in household h of composition g by: $q_{t,h,g}^j \in \Omega^Q$. The price of each good j is $p_{t,h,g}^j$. For most goods, the price does not vary within the sample, so we often write simply p^j and omit the indices. However, some prices vary at the household or individual level, notably for time-use. The price of an individual's time is their wage $w_{t,h,g}$.⁵⁷

Time is continuous,⁵⁸ and each individual has time-endowment normalised to 1, so each individual has labour endowment equal to $w_{t,h}$. In addition, each member can be endowed with non-labour income $y_{t,h}^{NL}$, or alternatively the household as a whole can have non-labour income y_h^{NL} . The household's endowment (or full income) is $y_h = y_h^{NL} + \sum_{t \in h} w_{t,h}$. We refer to this as the household budget.

⁵⁶Note that we do not claim that these two groups are “comparable”, e.g. we do not focus only on heterosexual singles (it is not possible in the data) nor do we require people to have stable preferences across household compositions: this is why we conduct the analysis separately for the two groups.

⁵⁷For non-participants in the labour market, the price of their time is higher than their (unobserved) wage, and would require estimating. For this reason, we restrict our attention to working couples without cohabiting children.

⁵⁸By examining UK time-use data used for the application in this paper, this modelling assumption appears realistic.

Preferences are heterogenous in two dimensions: sex and household composition. Single men and women have different preferences from each other. Men and women in couples have different preferences from each other, and also have different preferences from singles. This allows us to capture the key dimension of preference heterogeneity of interest to us, i.e. the difference between men and women, without assuming preference stability across household composition, which has been rejected empirically by other papers e.g. by Hubner (2020). Therefore an individual has preferences of one of four types: $u_{f,S}$, $u_{m,S}$, $u_{f,C}$, or $u_{m,C}$.

All preferences and domestic production functions are modelled as Cobb-Douglas, so that an individual's utility takes the form:

$$u_{t,g} = \sum_{j \in \Omega^Q} (\alpha_{t,g}^j \ln(q_{t,h,g}^j)) \quad (1)$$

This widely used, tractable, specification has several advantages in this context. Firstly, it enables estimation of bargaining weights even from data with small sample sizes, which are prevalent in the household bargaining literature, including in the application to environmental preferences in this paper. Secondly, the parsimonious direct utility representation and demand equations lend themselves to the environmental extension in this paper and enable transparent interpretation of estimates. The price of these advantages is that Cobb-Douglas imposes homotheticity and separability, which are strong assumptions.

As explained in the methodology section, to estimate bargaining weights we focus only on data on private leisure. Hence, in effect, we are only imposing the restrictions of Cobb-Douglas on this specific good. As discussed in the first chapter of this thesis, in this context the Cobb-Douglas assumption does not seem problematic. For instance, UK time-use data on singles shows homothetic patterns for expenditure on private leisure.

However, for our environmental analysis we need to draw more heavily on parametric form of preferences. This introduces complexities because some goods, such as food and clothing, exhibit strong non-homotheticity. To keep our approach simple and easy to interpret, we use Cobb-Douglas utility functions, but only make comparisons between people with similar budgets. This is akin to taking a linear approximation of expenditure functions which is local to a specific budget.

14.2 Singles

For singles, Cobb-Douglas utility functions yield very simple expenditure functions:

$$E_{t,h}^j = q_{t,h}^j p^j = \alpha_{t,S}^j y_h$$

Expenditure on a subset of goods Ω^X is:

$$E_{t,h}^X = \sum_{j \in \Omega^X} (q_{t,h}^j p^j) = \left(\sum_{j \in \Omega^X} \alpha_{t,S}^j \right) y_h$$

These expenditure functions yield simple structural estimating equations that allow us to directly recover the preference parameters of singles.

14.3 Couples

14.3.1 Bargaining in the household

In a household with multiple members, the constituent individuals bargain over how to divide resources. In the context of this paper this only applies to couples, so we drop the household category subscripts g . Depending on the bargaining process, and on the outside options of the individuals, the different individuals will have different bargaining power and the resulting household-level consumption patterns will be a closer reflection of the preferences of one or the other member. Individuals have a vector of characteristics $\pi_{t,h}$ (e.g. age, educational attainment) and their households have a vector of characteristics ζ_h (e.g. the gender ratio in the region).

The collective model of the household does not restrict bargaining to any specific solution, and only requires that this process be Pareto efficient.⁵⁹ Relative bargaining power will in general depend on a variety of variables, including prices (e.g. hourly wages of all members), individual characteristics (e.g. age and educational attainment of all members) and household characteristics (e.g. local area the household is located in). We refer to the vector of variables that affect bargaining power as z_h . Note that a specific member's bargaining power will depend not only on their own characteristics, but also on the characteristics of all other household members, hence the household-level subscript.

A key result from the existing literature (see Browning, Chiappori and Weiss (2014)) is that the problem solved by any collective household, regardless of the underlying bargaining process and outside options, can be represented as an optimisation problem where the maximand is the weighted sum of the members' utility functions. Each member's utility function is weighed by their Pareto weight $\mu_{t,h}(z_h)$ normalised so that $\sum_{i \in h} \mu_{t,h}(z_h) = 1$. The higher an individual's Pareto weight, the more weight the collective household gives their utility

⁵⁹As opposed to models with specific bargaining solutions e.g. McElroy and Horney (1981), for which there is little consensus in the literature.

in determining its choices. The household's optimisation problem is therefore to maximise $\sum_{i \in h} (\mu_{t,h}(z_h) u_{t,g}(q_{t,h}))$.

Because the pareto weights could take any form, we linearly approximate them as:

$$\mu_{t,h} = \mu_t^0 + \sum_z \mu_t^z (z_h - \bar{z})$$

- μ_f^0 is the average pareto weight of women and μ_m^0 is the average pareto weight of men in the sample. This is the pareto weight evaluated at the average characteristics in the sample. By definition, the average resource share of men and women sum to one: $\mu_f^0 + \mu_m^0 = 1$.
- $\mu_f^z (\mu_m^z)$ captures the impact on sharing of a household's characteristic z_h deviating from the sample average $(z_h - \bar{z})$ on women's (men's) bargaining power. For instance, a higher-than-average wage for the woman might increase the woman's pareto weight, so that she would have a higher-than-average-for-women bargaining power. Since pareto weights must sum to one within the household, this implies her partner must have a correspondingly lower-than-average-for-men pareto weight: $\mu_f^z + \mu_m^z = 0$. We can interpret μ_t^z as the marginal impact of characteristic z_h on the resource share on people of type t.

Given the properties of Pareto weights, in the rest of the paper we write $\mu_{f,h}$ for women's bargaining power and $(1 - \mu_{f,h})$ for the man's.

14.3.2 Household-level expenditure functions for couples

In this context, the household-level expenditure functions for couples are a weighted sum of the preferences of members, where the weights represent their relative bargaining power:

$$E_h^j = \mu_{m,h} \alpha_{m,C}^j y_h + \mu_{f,h} \alpha_{f,C}^j y_h$$

Using the fact that bargaining weights sum to one, we can re-write this as:

$$E_h^j = \alpha_{m,C}^j y_h + (\alpha_{f,C}^j - \alpha_{m,C}^j) \mu_{f,h} y_h$$

Note that here the interpretation of expenditure on a private good is that part of it is consumed by the man and part of it by the woman, while for public expenditure, the whole consumption is enjoyed by both household members. In expenditure data, we only observe household expenditure on a good such as food E_h^{food} , but that is an aggregate of the expenditure on food for the woman $\mu_{f,h} \alpha_{f,C}^{food} y_h$ and food for the man $\mu_{m,h} \alpha_{m,C}^{food} y_h$.

Expenditure on a subset of goods Ω^X is:

$$E_{t,h}^X = \sum_{j \in \Omega^X} E_h^j = \left(\sum_{j \in \Omega^X} \alpha_{m,C}^j \right) y_h + \left(\sum_{j \in \Omega^X} (\alpha_{f,C}^j - \alpha_{m,C}^j) \right) \mu_{f,h} y_h$$

We first estimate bargaining weights, and then substitute them into these structural equations, allowing us to directly estimate these equations to recover the preference parameters of men and women in couples.

14.4 Environmental extension

For consistency with emissions data, we model emissions per GBP of expenditure as a good-specific constant ϕ^j , called conversion factor or emission multiplier.

A household's expenditure on a specific good therefore has associated emissions:

$$\epsilon_h^j = \phi^j E_h^j$$

Similarly, a household's expenditure on a subset of goods X has associated emissions:

$$\epsilon_h^X = \sum_{j \in \Omega^X} \phi^j E_h^j$$

The total emissions (carbon footprint) of a household are:

$$\epsilon_h^Q = \sum_{j \in \Omega^Q} \phi^j E_h^j$$

We would expect higher budget households to emit more than lower budget households in absolute terms (because most consumption has associated emissions), so rather than looking at absolute emissions, we focus on emissions relative to budget, which depends on how money is spent. For single households, we define their **emissions propensity** as their emissions-to-budget ratio:

$$\pi_{h,S}^Q = \frac{\epsilon_h^Q}{y_h} = \sum_j \alpha_{t,S}^j \phi^j$$

or, for a subset of goods:

$$\pi_{h,S}^X = \frac{\epsilon_h^X}{y_h} = \sum_{j \in \Omega^X} \alpha_{t,S}^j \phi^j$$

For couples, we define emissions propensity of a member as the emissions-to-budget ratio

that the household would have under a dictatorship of that individual.⁶⁰

$$\pi_{h,C}^Q = \sum_j \alpha_{t,C}^j \phi^j$$

And for a subset of goods:

$$\pi_{h,C}^X = \sum_{j \in \Omega^X} \alpha_{t,C}^j \phi^j$$

We are interested in whether women or men have a higher emissions propensity, i.e. comparing $\pi_{f,g}^X \leq \pi_{m,g}^X$. Note that lower emissions propensity on a subset of goods X is driven by two factors:

1. Consuming less of X: spending a lower proportion of budget on goods in subset X and spending a larger proportion on some other subset of goods.
2. Consuming more environmentally friendly products within X: keeping the proportion of budget spent on X constant, a higher proportion is spent on goods with lower associated emissions.

Lower emissions propensity overall indicate exclusively the second effect.

To disentangle between these two effects, we also consider the **carbon intensity of choices**. For singles, this is defined as the emissions propensity on category X, divided by the share of the budget that is spent on category X. For a member of a couple, this is defined as that person's emissions propensity on category X, divided by the share of the budget that would be spent on that category under the dictatorship of that member:

Overall:

$$CI_{t,C} = \frac{\sum_j \alpha_{t,g}^j \phi^j}{\sum_j \alpha_{t,g}^j}$$

For a subset of goods X:

$$CI_{t,C}^X = \frac{\sum_{j \in \Omega^X} \alpha_{t,g}^j \phi^j}{\sum_{j \in \Omega^X} \alpha_{t,g}^j}$$

⁶⁰We define male dictatorship as a household where decisions are fully shaped by the male preference parameter. In this case, the female bargaining weight is equal to zero. Similarly, under female dictatorship, the female bargaining weight is equal to one. In our analysis, we also consider an equal household, where each members' preferences are weighted equally.

15 Methodology

We begin by estimating household-specific bargaining weights from UK time-use data, based on observable household characteristics and our bargaining rule.⁶¹ These estimates are then applied to expenditure data (specifically the LCF) to calculate household-specific bargaining weights within that dataset. To analyse emissions patterns, we integrate emissions data from DEFRA into the LCF, creating an emissions-augmented expenditure dataset. Within this dataset, we compare the consumption and emissions behaviours of single men and women with similar budgets. We then examine couples, using the estimated bargaining weights, again focusing on comparisons between households with similar financial means.

15.1 Estimating bargaining weights

15.1.1 Estimating bargaining weights from time-use data

Estimating structural parameters for multi-person households is complicated by substantial data limitations. In general, we only observe aggregate household expenditures, and cannot directly observe how these are split between members, or what process the household underwent to arrive to this choice. In order to get under the hood of household decision-making, and uncover differences in bargaining power and in preferences between members, we require some identifying variation in the data - something that varies at the individual, rather than household, level, and which we observe for all household members.

Dunbar et al. (2013) shows that we can recover resource sharing (the way that expenditures are split between household members overall, not on each specific good) in quite a general setting, as long as we observe individual-level expenditure on a single private good, the ‘assignable good’, for all members of a household. Lechene et al. (2022) shows that the approach proposed by Dunbar et al. (2013) can be estimated simply by OLS. Building on those results, the first chapter of this thesis proposes a new approach to estimating resource sharing from time-use data.

We use a similar approach, but (i) estimate bargaining weights instead of resource sharing, and (ii) apply our structural estimates to explore gender differences in terms of emissions, rather than individual-level consumption. We set out our approach below.

First, we use individual time diaries (collected for all members of a household) to obtain individual-level expenditure on private leisure ℓ for both men and women in working heterosexual couples. Private leisure is time spent on leisure activities without other household

⁶¹Note that the bargaining rule may well vary between countries and over time. In this paper, we focus only on the UK in the period 2001-14, and hence do not allow the bargaining rule to vary itself.

members co-present. For instance, one of the members may be reading a book on their own, or having a coffee with a friend but without their partner. Expenditure on private leisure is the amount of time spent on private leisure multiplied by the price of time, which for our sample is the observed wage from paid market work. From our structural model we obtain the equations:

$$E_{f,h}^\ell = w_{f,h} \ell_{f,h} = \alpha_{f,C}^\ell \mu_{f,h} y_h$$

$$E_{m,h}^\ell = w_{m,h} \ell_{m,h} = \alpha_{m,C}^\ell (1 - \mu_{f,h}) y_h$$

Substituting in the approximation of the Pareto weight, we obtain the following structural estimating equations:

$$E_{f,h}^\ell = \alpha_{f,C}^\ell \mu_f^0 y_h + \alpha_{f,C}^\ell \sum_z \mu_f^z (z_h - \bar{z})$$

$$E_{m,h}^\ell = \alpha_{m,C}^\ell (1 - \mu_f^0) y_h - \alpha_{f,C}^\ell \sum_z \mu_f^z (z_h - \bar{z})$$

Re-writing this in terms of observables and regression coefficients:

$$E_{f,h} = \beta_f^0 y_h + \sum_z \beta_f^z y_h (z_h - \bar{z})$$

$$E_{m,h} = \beta_m^0 y_h + \sum_z \beta_m^z y_h (z_h - \bar{z})$$

where:

- $\beta_f^0 = \alpha_{f,C}^l \mu_f^0$
- $\beta_f^z = \alpha_{f,C}^l \mu_f^z$
- $\beta_m^0 = \alpha_{m,C}^l (1 - \mu_f^0)$
- $\beta_m^z = -\alpha_{m,C}^l \mu_f^z$

To identify the bargaining weight parameters, we:

1. Run linear regressions⁶² of leisure expenditure for men and women with the restriction that $\beta_f^z + \beta_m^z = 0$.
2. We make the identifying assumption that men and women have similar preferences for private leisure so that $\alpha_{t,C}^l = \alpha_C^l$.
3. Estimate each type's average bargaining weight as $\hat{\mu}_{t,C}^0 = \frac{\beta_{t,C}^0}{\beta_{f,C}^0 + \beta_{m,C}^0}$.
4. Estimate the marginal impact of different characteristics as follows. First, estimate $\hat{\alpha}^l = \frac{\hat{\beta}_t^0}{\hat{\mu}_t^0}$. Then, we estimate $\hat{\mu}_t^z = \frac{\hat{\beta}_t^z}{\hat{\alpha}^l}$.

The intuition for the identification result is as follows. The man's expenditure on private leisure may be more or less responsive to changes in the household budget for one of two reasons, or a combination of them. The first possibility is that the man has more bargaining power, and hence more of a say on how household budget is spent. In this case, more of the budget will be assigned to his own private consumption (including his private leisure). This is the channel we wish to estimate. The second possibility is that the man's preferences for private leisure are stronger relative to the woman's, so that even if they have equal say in how additional budget is used, the man may choose to spend additional budget on his private leisure while the woman may choose to spend additional budget on her clothing. This is a potentially confounding factor.

In order to disentangle these two channels, we shut down the preference channel through the identifying assumption, allowing us to identify bargaining power. Under our identifying assumption, men and women have the same preferences for their own private leisure, so differences in responsiveness to changes in y_h identify differences in bargaining weights. This identifying assumption is consistent with UK singles' time-use data, in which we observe similar patterns of expenditure on private leisure between men and women.⁶³

15.1.2 Estimating bargaining weights in the expenditure dataset

Having estimated the bargaining rule, we can apply our estimates to a separate but comparable dataset, the LCF. Note that both the UKTUS and LCF are representative of the UK population and we use data for the same time period. This allows us to estimate bargain-

⁶²We run the two regressions jointly, as seemingly unrelated regressions, because the error terms are likely correlated across equations.

⁶³While this does not guarantee that men and women *in couples* will also have similar preferences for leisure, we are not able to test our identifying assumption directly. We note that this is also the case for similar papers in the household literature.

ing weights for the households in our expenditure data as a function of those households' characteristics.⁶⁴

15.2 Estimation of emissions propensities

15.2.1 Singles

The first part of our analysis focuses on singles, to directly elicit gender differences from observed consumption patterns. For singles, the structural estimating equations for greenhouse gas emissions from category X are:

$$\begin{aligned}\epsilon_{f,h}^X &= \left(\sum_{j \in \Omega^X} \alpha_{f,S}^j \phi^j \right) y_h \\ \epsilon_{m,h}^X &= \left(\sum_{j \in \Omega^X} \alpha_{m,S}^j \phi^j \right) y_h\end{aligned}$$

Therefore, we linearly regress:

$$\epsilon_{t,h,S}^X = \beta_{f,S}^X y_h I_f + \beta_{m,S}^X y_h I_m \quad (2)$$

Where I_t is an indicator function which takes value 1 for individuals of type t and 0 for individuals who are not of type t. We conduct the singles analysis separately for households in different budget groups to ensure we are comparing single men and women with similar budgets.⁶⁵

Our estimated coefficients are: $\beta_{t,S}^X = \pi_{t,S}^X$ and $\beta_{t,S} = \pi_{t,S}$. Hence, if $\beta_{f,S}^X < \beta_{m,S}^X$ we conclude that women have a lower emissions propensity than men for category of goods X. Then, to obtain the carbon intensity of choices, we run an additional regression on expenditures, rather than emissions:

$$E_{t,h,S}^X = \gamma_{f,S}^X y_h I_f + \gamma_{m,S}^X y_h I_m \quad (3)$$

Here the estimated coefficients are: $\gamma_{t,S}^X = \sum_{j \in \Omega^X} \alpha_{t,S}^j$, so we recover the share of budget spent on category X from these coefficients directly. We then obtain the carbon intensity of

⁶⁴To do so, we calculate how household characteristics in the expenditure dataset deviate from the averages in the time-use data, and use these deviations ($z_h - \bar{z}$) to estimate household-specific bargaining.

⁶⁵We divide singles into categories where the budget varies by no more than 100GBP per week, and conduct our analysis on buckets which contain at least 200 households.

choices for single men and women as:

$$CI_{t,s} = \frac{\beta_{t,s}}{\gamma_{t,s}}$$

15.2.2 Couples

For couples, the structural estimating equation for greenhouse gas emissions for category X is:

$$\epsilon_h^X = \left(\sum_{j \in \Omega^X} \alpha_{m,C}^j \phi^j \right) y_h + \left(\sum_{j \in \Omega^X} ((\alpha_{f,C}^j - \alpha_{m,C}^j) \phi^j) \right) \mu_{f,h} y_h$$

Therefore, we linearly regress:

$$\epsilon_h^X = \beta_{m,C}^X y_h + \beta_{\Delta,C}^X \mu_{f,h} y_h \quad (4)$$

For couples, we also conduct our analysis separately for people with similar budgets. To do so, we focus our analysis on households which (i) are in the middle 50% of the distribution of total household budget, and (ii) have bargaining weights between 0.4 and 0.6. This means that, between households, we are comparing households with budgets that are not too dissimilar and, within households, we are comparing men and women who have a reasonably similar split of their household budget at their disposal.

For men, we estimate the emissions propensity as:

$$\pi_{m,C}^X = \beta_{m,C}^X$$

And for women, we estimate it as:

$$\pi_{f,C}^X = \beta_{m,C}^X + \beta_{\Delta,C}^X$$

To estimate the carbon intensity of choices, we additionally run a regression on expenditure, rather than emissions, to recover further structural parameters:

$$E_h^X = \gamma_{m,C}^X y_h + \gamma_{\Delta,C}^X \mu_{f,h} y_h \quad (5)$$

Using these estimates we obtain shares of the budget that would be spent on category X

under dictatorship of men:

$$\sum_{j \in \Omega^X} \alpha_{m,C}^j = \gamma_{m,C}^X$$

And for women:

$$\sum_{j \in \Omega^X} \alpha_{f,C}^j = \gamma_{m,C}^X + \gamma_{\Delta,C}^X$$

Hence, we estimate men's carbon intensity as:

$$CI_{m,g}^X = \frac{\beta_{m,C}^X}{\gamma_{m,C}^X}$$

And for women:

$$CI_{f,g}^X = \frac{\beta_{m,C}^X + \beta_{\Delta,C}^X}{\gamma_{m,C}^X + \gamma_{\Delta,C}^X}$$

16 Data

In this section, we present the data sources used in our analysis. We begin by using the UK Time-Use Survey to estimate bargaining power among working couples in the UK, based on household characteristics (the bargaining rule). These estimates are then applied to the Living Costs and Food Survey (LCF), a comparable expenditure dataset, to compute household-specific bargaining. We enrich the expenditure data with emissions multipliers from Environmentally-Extended Input-Output tables, enabling us to estimate the emissions associated with the spending patterns of different households.

16.1 UK Time-use Survey

The UK Time-Use Survey (UKTUS 2001, 2014)⁶⁶, is a high-quality time-use survey, representative of the UK population. It is a national household-based study composed of: (i) a household questionnaire, (ii) an individual questionnaire, and (iii) individual time diaries. A single household representative answered the household questionnaire, including questions on household characteristics such as composition, dwelling type, and location. The other components were answered by the individual in question. This is likely to substantially increase the quality of the data relative to datasets where a single member answers on behalf of all individuals. The individual questionnaire asks about individual characteristics including age, educational attainment and earnings. Each member completed a weekday and weekend

⁶⁶Office for National Statistics. (2019b). United Kingdom Time Use Survey [data series] 2nd Release. <https://doi.org/http://doi.org/10.5255/UKDA-Series-2000054>.

time diary identifying primary and secondary activities for each 10-minute interval over the two days.

The time-use data is extremely detailed, including very granular activities, location, and co-presence of others (distinguishing between household and non-household members). A full list of activities can be found on the UKTUS webpage (it is not included due to its length), but examples of activities include “shopping for and ordering food via the internet”, “repairing dwelling”, “ironing”, “cleaning yard”, and “working”.

To the end of identifying bargaining weights, we require information on how much time people spend doing leisure activities without other household members co-present (private leisure). This variation is used to identify bargaining weights, as explained in Section ?? . The activities included in our definition of private leisure include, for instance: “reading books”, “walking and hiking”, “watching a film on TV”, “listening to sport on the radio”, “visiting and receiving visitors”, and “sleeping”.⁶⁷

The time diaries are constructed carefully to minimise measurement error, for instance with the possibility of writing a simple sign to signify the same activity for multiple time intervals. The quality of the data is very high, with approximately 95% of observations having more than 5 distinct activities recorded in a day, and less than 90 minutes of unrecorded time.

16.2 LCF

The Living Costs and Food Survey (previously FES) is a nationally representative UK expenditure survey. It is a high-quality, large-scale survey that is used to estimate official government statistics, and has been widely used in academic papers. The LCF is a repeated cross-section available yearly since 1978. We focus on the years 2001-2014 for comparability with the time-use data used to estimate bargaining within couples.⁶⁸ The survey has three components:

1. A household survey recording household characteristics and retrospective questions on infrequent expenses (rent, vehicles, house repairs...). These infrequent expenditures are transformed to an equivalent weekly value to make them comparable to other categories. The household survey is answered by the reference person (potentially jointly with other household members).

⁶⁷The full list of activities is chosen to be consistent with the approach outlined in the first chapter of this thesis. As discussed there, results are robust to alternative definitions of leisure activities

⁶⁸The expenditure survey changes from 2000 to 2001 and is available in its current form only from 2001, hence we exclude 2000. For the years 2004 and 2006, coding problems led to the detailed diary data not being reliable. The data files with more aggregate expenditure categories for those years are high quality but are not sufficient for our purposes.

2. An individual questionnaire, answered by each member, with information including educational attainment, hours worked and sources of income.⁶⁹
3. A detailed two-week expenditure diary, completed individually by each member. For our purposes, who purchased a good is not relevant, so our analysis is based on overall household expenditure, summed across different members' diaries.

After restricting the pooled LCF data (2001-14) to working singles and working heterosexual couples, and cleaning the data, our sample ('main' sample) comprises 13,913 households, of whom 7,448 are couples and 6,465 are working singles. Table 5 provides summary statistics for the 'main' sample, which we use to understand high-level patterns and to interpret our findings relative to the population. As explained in section 14, for our regression analysis, we further restrict our sample to compare only people with similar budgets. This substantially reduces our sample size, but is important to weaken the constraint imposed by our homotheticity assumption.⁷⁰ Table 6 provides summary statistics for the regression sample, which is composed of 4,141 singles and 3,149 couples.

Hourly pay is obtained by dividing labour income by actual hours worked, rather than contractual hours.⁷¹ Household budget is full income: the sum of the labour endowment of members (hourly wage multiplied by 24 hours, multiplied by 7 days).⁷² On average, despite having higher educational attainment,⁷³ women command lower hourly wages than men. Single women are slightly older than single men, and the opposite is the case for couples.

⁶⁹In rare instances, income is top-coded. We adjust top-coded values using data on after-tax income percentiles from HM Revenue & Customs (HM Revenue & Customs 2023).

⁷⁰The homotheticity assumption was reasonable in my first chapter, applied just to private leisure, but is too unrealistic for the expenditure system as a whole. Since we are restricting our sample to the middle of the distribution, in interpreting our results we focus on the effect of changes for the average couple.

⁷¹This, together with self-employed labour, explains the lower end of hourly wages (which are sometimes lower than the official minimum hourly wage).

⁷²Accurate non-labour income data is not available in UKTUS and hence is also excluded in LCF for comparability.

⁷³We divide qualifications into three categories: 2 for undergraduate degrees, equivalent or higher; 1 for school-leaving qualifications, such as A-levels, or equivalent; and 0 for any lower qualifications.

Table 5: Summary statistics for LCF main sample

(a) Singles			
Variable	(1) Female	(2) Male	(3) Difference
Household Budg.	1,938.387 (1,154.741)	2,133.766 (1,442.175)	195.380*** (32.700)
Hourly pay	11.538 (6.873)	12.701 (8.584)	1.163*** (0.195)
Age	46.616 (12.857)	44.586 (11.917)	-2.030*** (0.308)
Qualification score	0.929 (0.834)	0.830 (0.841)	-0.099*** (0.021)
Observations	3,079	3,386	6,465

(b) Couples			
Variable	(1) Female	(2) Male	(3) Difference
Household Budg.	3,982.508 (1,942.905)	3,982.508 (1,942.905)	0.000 (31.838)
Hourly pay	10.933 (5.758)	12.773 (8.178)	1.840*** (0.116)
Age	43.066 (12.956)	45.086 (13.004)	2.020*** (0.213)
Qualification score	0.922 (0.847)	0.809 (0.848)	-0.113*** (0.014)
Observations	7,448	7,448	14,896

Notes: Standard deviation in parenthesis in Columns (1) and (2), Standard errors of two-sample t-test in parenthesis in Column (3). Column (3) provides the difference between single females and males for selected variables in Sub-table (a), and the difference between female and male members of two-person households in the main sample for selected variables in Sub-table (b).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Summary statistics for LCF regression sample

(a) Singles			
Variable	(1) Female	(2) Male	(3) Difference
Household Budg.	1,535.653 (327.737)	1,555.116 (324.261)	19.463* (10.134)
Hourly pay	9.141 (1.951)	9.257 (1.930)	0.116* (0.060)
Age	46.560 (13.121)	44.339 (12.177)	-2.221*** (0.394)
Qualification score	0.801 (0.800)	0.654 (0.783)	-0.148*** (0.025)
Observations	2,099	2,042	4,141

(b) Couples			
Variable	(1) Female	(2) Male	(3) Difference
Household Budg.	3,540.823 (468.824)	3,540.823 (468.824)	0.000 (11.815)
Hourly pay	10.376 (2.388)	10.700 (2.512)	0.324*** (0.062)
Age	40.963 (12.627)	42.745 (12.606)	1.782*** (0.318)
Qualification score	0.965 (0.836)	0.696 (0.809)	-0.269*** (0.021)
Observations	3,149	3,149	6,298

Notes: Standard deviation in parenthesis in Columns (1) and (2), Standard errors of two-sample t-test in parenthesis in Column (3). Column (3) provides the difference between single females and males for selected variables in Sub-table (a), and the difference between female and male members of two-person households in our regression sample for selected variables in Sub-table (b).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

16.3 DEFRA emissions multipliers

To translate GBP expenditures into GHG emissions, we make use of emissions multipliers provided by the Department for Environment Food and Rural Affairs (DEFRA). These quantify the GHG emissions associated with one GBP spent in a given year on a given good category, in the UK. Emissions of different GHGs are converted into CO_2 equivalent to provide a single emissions multiplier per GBP expenditure on a category of goods.

These multipliers are derived from the UK Multi-Regional Input-Output model (UKMRIO

model) (DEFRA 2024), which is an environmentally-extended multi-regional input-output database. This approach takes into account all direct and indirect emissions associated with expenditures by any UK end-consumer (households, national and local government, and charities) regardless of where those emissions actually happened. For instance, household emissions include direct emissions from burning wood in a fireplace, indirect emissions from a local restaurant using energy to cook food purchased by the household, and indirect emissions abroad involved in a foreign firm producing a technological device purchased by the household.

The DEFRA emission multipliers are available at the COICOP level (Classification of Individual Consumption by Purpose). This includes 108 product categories - see appendix H for a full list, as well as additional information about the DEFRA multipliers. Figure 4 depicts the average GHG emissions for one GBP spent on selected goods, in the period 2001 - 2014. Two opposite forces are driving emissions multipliers: the price of a good drives the multiplier down, while its associated emissions drive it up.

Emissions multipliers from environmentally-extended multi-regional input-output databases are widely used and viewed as the best approach to estimate consumption-based emissions (e.g. Christis et al. (2019), Kilian et al. (2023), Osorio et al. (2024) and Owen and Büchs (2024)). It is worth noting that they suffer from two main limitations. Firstly, at the time of writing, they are available only at a higher level of aggregation than would be ideal. For instance, in the DEFRA data, while we can distinguish between meat and vegetable products, we cannot distinguish between red and white meat. Secondly, being based on expenditure data, carbon footprint estimates inherit some limitations of expenditure surveys. For instance, there is no way of accounting for whether a product was purchased at a discounted price. If there were gendered patterns in purchasing discounted products, this could affect our estimates.

16.3.1 Bridging LCF and DEFRA data

We merge the DEFRA data into LCF by COICOP category. This gives us category-specific emissions per GBP expenditure alongside the information on how much our households spend on those categories. Our analysis focuses on the following categories of expenditure: food, energy, transport, clothing, and an overall expenditure measure including these categories and other expenditures (e.g. on toiletries). Our selection of consumption categories aligns with our objective of analyzing daily individual behavior and follows common practices in

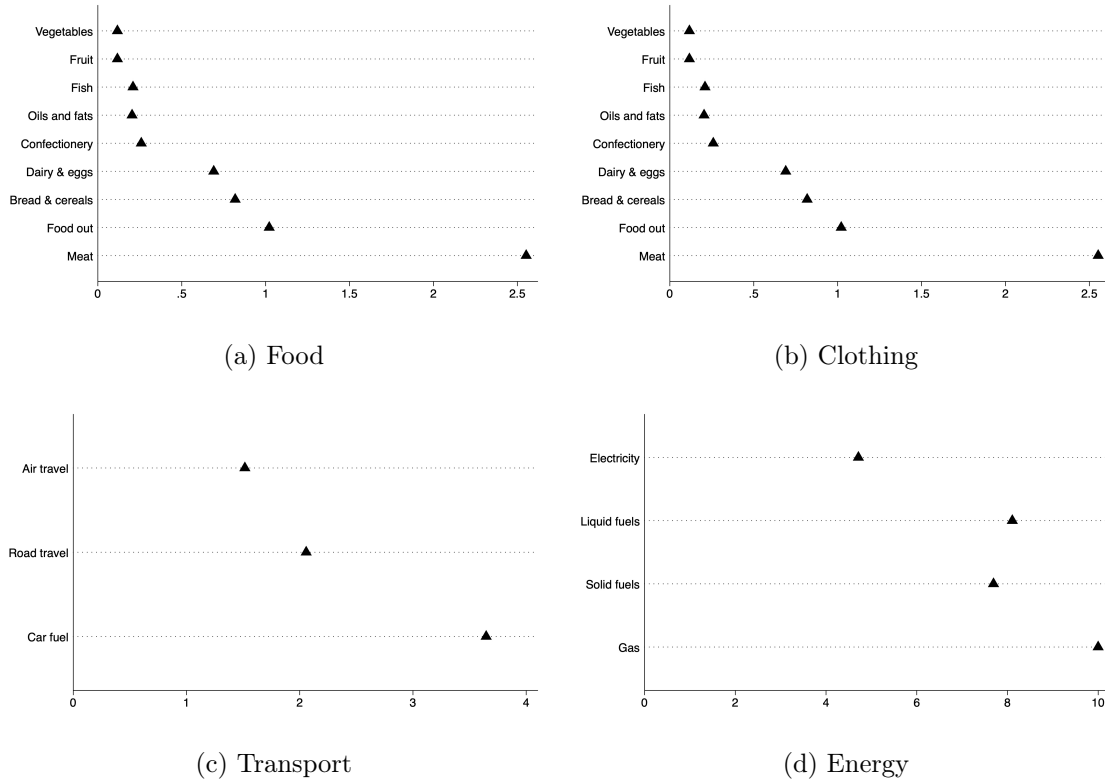


Figure 4: Average Emission factors, kgCo2eq/£, 2001-2014

Notes: Graphical depiction of the average GhG multipliers in kgCo2eq/£ from 2001 to 2014 for selected consumption categories. Yearly multipliers provided by DEFRA (2024).

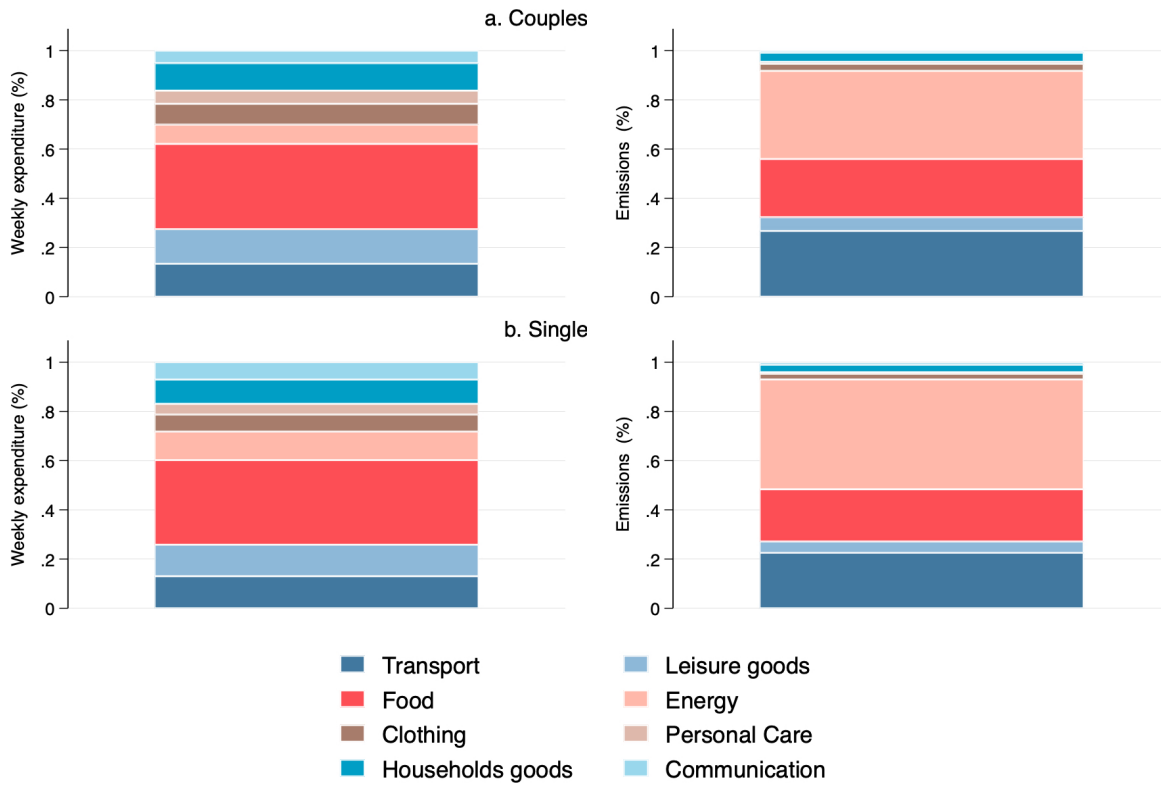
the literature.⁷⁴ We exclude from the analysis expenditure categories that have almost no immediate consumption value and cannot easily be squared with the static model underlying the methodology in this paper. This includes: savings, insurance, investments, purchase and lease of vehicles, major house works including renovation (minor repairs are included), financial gifts, bets, gambling, and expenditure on education (the latter is minor for working couples without cohabiting children). We exclude expenditure on accommodation because it is too noisy a measure of emissions associated with building that accommodation in the first place. Similarly, we exclude medical expenditures, which map very inaccurately onto emissions in the UK context because of the complex pricing structure for treatment and medication in the UK. The average contribution of each consumption category to spending and emissions for households in the ‘main’ sample is illustrated in figure 4.

Our overall expenditure measure captures a very large proportion of household expenditure. It represents on average 90.54% of the measure of total consumption in the LCF data (the

⁷⁴For instance, Carlsson Kanyama et al. (2021) examine only food, furnishing, and holiday-related expenses, while Osorio et al. (2024) incorporate both consumption and non-consumption expenditures in their carbon footprint calculations.

ONS total consumption measure). The remaining proportion of expenditure can be ascribed to the excluded categories, particularly rental and mortgage costs. We note that our overall expenditure measure is a much smaller proportion of household budget. This is because household budget in this paper is defined as full income, rather than realised income or disposable income, i.e. it includes income that would have been earned if household members worked 24 hours a day, 7 days a week. People spend a substantial proportion of their time, and hence of their full income, on leisure and domestic activities.

Figure 5: Budget and carbon footprint shares by good categories.



Notes: The Figure depicts the share of weekly household budget dedicated to each good category (left bars) and the share of emissions out of the total weekly emissions these goods represent (right bar) for the singles and couples in our sample.

17 Findings

17.1 Findings for singles

Table 7 outlines the regression results and structural estimates for overall consumption. Each column refers to a particular household budget bucket: Column (1) to households with a




Figure 6: Gender difference in singles' emissions propensities on overall consumption: Women's - Men's.

Notes: The graph depicts the differences between single women's and men's emission propensities, estimated via Equation 2, and their confidence intervals at the 90% confidence level.

budget between 1000 and 1100£ per week, column (2) to households with 1100 to 1200£ per week, etc. The first two rows of the table provide the coefficients estimated from Equation 4, i.e. the regression on household emissions. Together with results from the regression on household expenditure, this allows us to recover the estimated emissions propensities (π_t), Cobb-Douglas preference parameters (α_t) i.e. the share of budget spent on those goods, and carbon intensity of choices (CI_t) for single men and women separately. We can then disentangle the budget share (α_t) from the carbon intensity (CI_t) channels, to explain gender differences in propensities. These are reported in the rows below. We additionally report for each estimate whether the difference between genders is statistically significant (significant diff.). Tables 8, 9, 10, and 11 are built similarly, but for specific categories of consumption. For each category of consumption (overall, food, transport, energy and clothing), we graph the difference between single men's and women' emission propensities together with its confidence interval in Figures 6, 7, 8, 9 and 10.

Overall environmental preferences. In our singles regression sample, the average weekly greenhouse gas emissions of single women are $330kgCo_2eq$ per week, and $346kgCo_2eq$ per week for single men, i.e. 4.85% lower for women than men. Moreover, as detailed in Figure

6 the overall pattern across budget categories suggests that single women have somewhat lower emissions propensities π than single men. These differences are mostly driven by men preferring goods that are more carbon intensive within the basket of goods included in our analysis, and are partly counteracted by women wishing to spend more of their budget on the categories of goods included in our analysis. Computing the weighted average for the regression sample, women's Cobb-Douglas preference parameter α on overall consumption is 4.3% larger than men's, while the carbon intensity (CI) of women's choices is 7.4% lower than men's.

Household budget groups(in 1000£)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-1.1	1.1-1.2	1.2-1.3	1.3-1.4	1.4-1.5	1.5-1.6	1.6-1.7	1.7-1.8	1.8-1.9	1.9-2	2-1.2	2.1-2.2
Fem. Budg.	0.269*** (0.013)	0.251*** (0.011)	0.250*** (0.008)	0.219*** (0.008)	0.229*** (0.009)	0.208*** (0.009)	0.207*** (0.008)	0.222*** (0.008)	0.185*** (0.007)	0.201*** (0.010)	0.173*** (0.014)	0.180*** (0.010)
Mal. Budg.	0.306*** (0.015)	0.288*** (0.012)	0.231*** (0.009)	0.241*** (0.008)	0.221*** (0.009)	0.232*** (0.008)	0.217*** (0.008)	0.196*** (0.008)	0.198*** (0.008)	0.200*** (0.009)	0.200*** (0.014)	0.186*** (0.010)
π_m	0.306	0.288	0.231	0.241	0.221	0.232	0.217	0.196	0.198	0.200	0.200	0.186
π_f	0.269	0.251	0.250	0.219	0.229	0.208	0.207	0.222	0.185	0.201	0.173	0.180
Significant diff.	yes	yes	no	yes	no	yes	no	yes	no	no	no	no
α_m	0.165	0.159	0.130	0.142	0.129	0.143	0.133	0.119	0.113	0.124	0.133	0.111
α_f	0.159	0.151	0.147	0.140	0.140	0.141	0.138	0.149	0.121	0.130	0.121	0.127
Significant diff.	no	no	yes	no	no	no	no	yes	no	no	no	yes
CI_m	1.849	1.813	1.780	1.694	1.712	1.623	1.629	1.641	1.743	1.615	1.505	1.668
CI_f	1.694	1.659	1.700	1.558	1.629	1.478	1.500	1.492	1.535	1.550	1.434	1.415
Significant diff.	yes	no	no	yes	no	no	no	yes	yes	no	no	yes
N	310	403	447	450	435	381	340	314	309	256	272	224
R-squared	0.736	0.731	0.781	0.773	0.750	0.786	0.797	0.821	0.805	0.771	0.568	0.765

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Emissions Propensities and Carbon Intensities for singles: overall consumption

Notes: The sample is divided into 12 subsamples of at least 200 households, whose budgets vary by no more than 100 pounds per week within the subsample. Each column refers to a distinct subsample: column 1 to the subsample including households whose weekly budget is between 1000 and 1100 pounds per week, column 2 to households whose weekly budget is between 1100 and 1200 pounds, etc. Fem. Budg. and Mal. Budg are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 90% confidence level.

Environmental preferences for food, transport, energy, and clothing. We then go on to look at gender differences in emission propensities and carbon intensity of consumption for specific good categories. We can see from Figure 8 that single women’s food preferences exhibit a lower emission propensity than single men’s. These gender differences are driven by both lower shares of budget dedicated to food consumption for single women, and less carbon-intensive food choices. Similar conclusions can be drawn for transportation, although the gap between single men’s and women’s emission propensities is lower, and rarely statistically significant.

As shown in Figure 10, we observe the opposite trend for energy. Although the differences in emission propensity are not statistically significant, the point estimates suggest that women’s emission propensities may be higher than men’s. Table 10 shows that single women tend to allocate a larger share of their budget to energy, in line with previous findings (notably Büchs and Schnepf (2013)). The results for carbon intensity are mixed, consistent with the fact that individuals have limited control over the type of energy used in their accommodation.

For clothing, as depicted in Figure 9, single women have substantially higher emissions propensity than single men. As illustrated in Table 11, this is driven almost fully by the substantially higher budget share that single women assign to clothing. The carbon intensity of clothing choices is similar for single men and single women - this is unsurprising since clothing is divided into a small number of, not very granular, categories in DEFRA emissions data.

Household budget groups(in 1000£)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-1.1	1.1-1.2	1.2-1.3	1.3-1.4	1.4-1.5	1.5-1.6	1.6-1.7	1.7-1.8	1.8-1.9	1.9-2	2-1.2	2.1-2.2
Fem. Budg.	0.045*** (0.003)	0.041*** (0.002)	0.042*** (0.002)	0.037*** (0.002)	0.038*** (0.002)	0.033*** (0.002)	0.035*** (0.002)	0.036*** (0.002)	0.030*** (0.002)	0.032*** (0.002)	0.032*** (0.002)	0.025*** (0.002)
Mal. Budg.	0.064*** (0.003)	0.053*** (0.002)	0.046*** (0.002)	0.051*** (0.002)	0.042*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.039*** (0.002)	0.044*** (0.003)	0.042*** (0.002)	0.042*** (0.002)	0.038*** (0.002)
π_m	0.064	0.053	0.046	0.051	0.042	0.047	0.047	0.039	0.044	0.042	0.042	0.038
π_f	0.045	0.041	0.042	0.037	0.038	0.033	0.035	0.036	0.030	0.032	0.032	0.025
Significant diff.	yes	yes	no	yes	no	yes	yes	no	yes	yes	yes	yes
α_m	0.053	0.046	0.041	0.045	0.037	0.040	0.040	0.034	0.037	0.038	0.036	0.033
α_f	0.045	0.040	0.039	0.037	0.036	0.033	0.034	0.036	0.030	0.031	0.031	0.027
Significant diff.	yes	yes	no	yes	no	yes	yes	no	yes	yes	yes	yes
CI_m	1.209	1.152	1.128	1.147	1.125	1.173	1.166	1.123	1.215	1.108	1.147	1.128
CI_f	0.996	1.028	1.083	1.011	1.044	0.994	1.018	0.999	0.980	1.047	1.019	0.921
Significant diff.	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	yes
N	310	403	447	450	435	381	340	314	309	256	272	224
R-squared	0.642	0.665	0.665	0.629	0.635	0.638	0.684	0.647	0.594	0.711	0.629	0.681

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Emissions Propensities and Carbon Intensities for singles: Food

Notes: The sample is divided into 12 subsamples of at least 200 households, whose budgets vary by no more than 100 pounds per week within the subsample. Each column refers to a distinct subsample: column 1 to the subsample including households whose weekly budget is between 1000 and 1100 pounds per week, column 2 to households whose weekly budget is between 1100 and 1200 pounds, etc. Fem. Budg. and Mal. Budg are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 90% confidence level.




Figure 7: Gender difference in singles' emissions propensities on transport: Women's - Men's.
Notes: The graph depicts the differences between single women's and men's emission propensities, estimated via Equation 2, and their confidence intervals at the 90% confidence level.




Figure 8: Gender difference in singles' emissions propensities on food: Women's - Men's.
Notes: The graph depicts the differences between single women's and men's emission propensities, estimated via Equation 2, and their confidence intervals at the 90% confidence level.

Household budget groups(in 1000£)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-1.1	1.1-1.2	1.2-1.3	1.3-1.4	1.4-1.5	1.5-1.6	1.6-1.7	1.7-1.8	1.8-1.9	1.9-2	2-1.2	2.1-2.2
Fem. Budg.	0.043*** (0.007)	0.045*** (0.004)	0.047*** (0.004)	0.043*** (0.003)	0.050*** (0.003)	0.047*** (0.004)	0.047*** (0.004)	0.049*** (0.004)	0.040*** (0.003)	0.046*** (0.004)	0.040*** (0.004)	0.044*** (0.004)
Mal. Budg.	0.070*** (0.008)	0.060*** (0.005)	0.049*** (0.004)	0.054*** (0.003)	0.054*** (0.003)	0.057*** (0.004)	0.050*** (0.004)	0.051*** (0.004)	0.047*** (0.004)	0.050*** (0.004)	0.047*** (0.004)	0.049*** (0.004)
π_m	0.070	0.060	0.049	0.054	0.054	0.057	0.050	0.051	0.047	0.050	0.047	0.049
π_f	0.043	0.045	0.047	0.043	0.050	0.047	0.047	0.049	0.040	0.046	0.040	0.044
Significant diff.	yes	yes	no	yes	no	yes	no	no	no	no	no	no
α_m	0.022	0.018	0.015	0.016	0.016	0.016	0.015	0.014	0.014	0.015	0.015	0.015
α_f	0.013	0.015	0.014	0.013	0.015	0.014	0.014	0.014	0.012	0.013	0.012	0.015
Significant diff.	yes	no	no	yes	no	no	no	no	no	no	no	no
CI_m	3.197	3.401	3.351	3.323	3.406	3.476	3.378	3.570	3.407	3.429	3.206	3.313
CI_f	3.209	3.019	3.359	3.288	3.264	3.352	3.298	3.399	3.367	3.437	3.211	2.987
Significant diff.	no	yes	no	no	no	no	no	no	no	no	no	yes
N	310	403	447	450	435	381	340	314	309	256	272	224
R-squared	0.290	0.408	0.446	0.500	0.515	0.499	0.481	0.547	0.503	0.526	0.529	0.538

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Emissions Propensities and Carbon Intensities for singles: Transport

Notes: The sample is divided into 12 subsamples of at least 200 households, whose budgets vary by no more than 100 pounds per week within the subsample. Each column refers to a distinct subsample: column 1 to the subsample including households whose weekly budget is between 1000 and 1100 pounds per week, column 2 to households whose weekly budget is between 1100 and 1200 pounds, etc. Fem. Budg. and Mal. Budg are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 90% confidence level.

Household budget groups(in 1000£)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-1.1	1.1-1.2	1.2-1.3	1.3-1.4	1.4-1.5	1.5-1.6	1.6-1.7	1.7-1.8	1.8-1.9	1.9-2	2-1.2	2.1-2.2
Fem. Budg.	0.138*** (0.007)	0.120*** (0.006)	0.116*** (0.005)	0.097*** (0.004)	0.097*** (0.004)	0.089*** (0.005)	0.088*** (0.004)	0.088*** (0.004)	0.079*** (0.004)	0.082*** (0.005)	0.065*** (0.004)	0.067*** (0.005)
Mal. Budg.	0.131*** (0.008)	0.130*** (0.007)	0.104*** (0.005)	0.094*** (0.004)	0.086*** (0.004)	0.088*** (0.004)	0.082*** (0.004)	0.077*** (0.004)	0.078*** (0.004)	0.071*** (0.005)	0.062*** (0.004)	0.066*** (0.005)
π_m	0.131	0.130	0.104	0.094	0.086	0.088	0.082	0.077	0.078	0.071	0.062	0.066
π_f	0.138	0.120	0.116	0.097	0.097	0.089	0.088	0.088	0.079	0.082	0.065	0.067
Significant diff.	no	no	yes	no	yes	no	no	yes	no	no	no	no
α_m	0.016	0.015	0.013	0.012	0.011	0.010	0.010	0.010	0.009	0.008	0.008	0.008
α_f	0.017	0.015	0.014	0.012	0.012	0.011	0.011	0.010	0.010	0.009	0.008	0.009
Significant diff.	no	no	no	no	yes	no	yes	no	no	yes	no	no
CI_m	8.168	8.413	7.877	8.022	7.864	8.385	8.253	7.878	8.216	8.589	7.668	8.359
CI_f	8.140	7.743	8.443	7.826	7.948	7.858	7.886	8.381	7.839	8.618	8.170	7.553
Significant diff.	no	yes	no	no	no	no	no	no	no	no	no	yes
N	310	403	447	450	435	381	340	314	309	256	272	224
R-squared	0.666	0.636	0.680	0.681	0.709	0.673	0.718	0.759	0.745	0.649	0.689	0.616

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Emissions Propensities and Carbon Intensities for singles: Energy

Notes: The sample is divided into 12 subsamples of at least 200 households, whose budgets vary by no more than 100 pounds per week within the subsample. Each column refers to a distinct subsample: column 1 to the subsample including households whose weekly budget is between 1000 and 1100 pounds per week, column 2 to households whose weekly budget is between 1100 and 1200 pounds, etc. Fem. Budg. and Mal. Budg. are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 90% confidence level.




Figure 9: Gender difference in singles' emissions propensities on clothing: Women's - Men's.
Notes: The graph depicts the differences between single women's and men's emission propensities, estimated via Equation 2, and their confidence intervals at the 90% confidence level.




Figure 10: Gender difference in singles' emissions propensities on energy: Women's - Men's.
Notes: The graph depicts the differences between single women's and men's emission propensities, estimated via Equation 2, and their confidence intervals at the 90% confidence level.

Household budget groups(in 1000£)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-1.1	1.1-1.2	1.2-1.3	1.3-1.4	1.4-1.5	1.5-1.6	1.6-1.7	1.7-1.8	1.8-1.9	1.9-2	2-1.2	2.1-2.2
Fem. Budg.	0.007*** (0.001)	0.007*** (0.001)	0.006*** (0.001)	0.006*** (0.001)	0.006*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.007*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Mal. Budg.	0.005*** (0.001)	0.005*** (0.001)	0.002*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.003*** (0.001)	0.003*** (0.001)	0.003*** (0.001)	0.002*** (0.001)	0.002*** (0.001)	0.003*** (0.001)	0.003*** (0.001)
π_m	0.005	0.005	0.002	0.004	0.004	0.003	0.003	0.003	0.002	0.002	0.003	0.003
π_f	0.007	0.007	0.006	0.006	0.006	0.005	0.005	0.007	0.005	0.005	0.005	0.005
Significant diff.	no	no	yes	yes	no	yes	yes	yes	yes	yes	yes	yes
α_m	0.010	0.010	0.006	0.009	0.007	0.006	0.008	0.006	0.005	0.005	0.007	0.006
α_f	0.014	0.015	0.013	0.013	0.013	0.012	0.012	0.016	0.012	0.011	0.011	0.011
Significant diff.	no	no	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
CI_m	0.466	0.472	0.404	0.467	0.559	0.454	0.444	0.411	0.451	0.431	0.435	0.521
CI_f	0.478	0.460	0.455	0.486	0.423	0.452	0.449	0.438	0.451	0.465	0.453	0.480
Significant diff.	no	no	no	no	no	no	no	no	no	no	no	no
N	310	403	447	450	435	381	340	314	309	256	272	224
R-squared	0.161	0.093	0.261	0.143	0.127	0.187	0.227	0.282	0.254	0.290	0.277	0.227

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Emissions Propensities and Carbon Intensities for singles: Clothing

Notes: The sample is divided into 12 subsamples of at least 200 households, whose budgets vary by no more than 100 pounds per week within the subsample. Each column refers to a distinct subsample: column 1 to the subsample including households whose weekly budget is between 1000 and 1100 pounds per week, column 2 to households whose weekly budget is between 1100 and 1200 pounds, etc. Fem. Budg. and Mal. Budg are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 90% confidence level.

17.2 Findings for couples

Table 12: Emissions Propensities and Carbon Intensities, couples

	(1) Overall	(2) Food	(3) Transport	(4) Energy	(5) Clothing
HH Budget	0.209*** (0.015)	0.034*** (0.004)	0.045*** (0.007)	0.086*** (0.007)	0.003** (0.001)
Female share	-0.084*** (0.032)	0.004 (0.008)	0.000 (0.015)	-0.072*** (0.015)	0.003 (0.003)
π_m	0.209	0.034	0.045	0.086	0.003
π_f	0.125	0.038	0.045	0.014	0.006
Significant diff.	yes	no	no	yes	no
α_m	0.125	0.037	0.013	0.012	0.007
α_f	0.109	0.029	0.014	0.001	0.014
Significant diff.	no	no	no	yes	no
CI_m	1.667	0.917	3.397	7.455	0.487
CI_f	1.144	1.304	3.274	24.661	0.442
Significant diff.	yes	yes	no	no	no
N	3,149	3,149	3,149	3,149	3,149
R-squared	0.825	0.764	0.612	0.688	0.316

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is restricted to households belonging to the middle 50% of the household budget distribution, and in which inequality between members is limited (the female resource share is constrained between 40 and 60% of the household's budget) HH Budg and Female share are the estimated coefficients of Equation 4 π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 5, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 95% confidence level.

17.2.1 Bargaining weights

Our bargaining regressions are identical to those run in the first chapter, reported in Table 2. By our identifying assumption, and the simple steps discussed in the methodology section, we are able to recover the structural parameters of the bargaining rule. Again, these are identical to those estimated in the first chapter: see section 9.2 for a discussion of those findings. Applying these estimates to households in the pooled LCF data, we find that

women have lower bargaining power than men on average: the mean bargaining weight for women is 0.46 and 0.54 for men in our main sample (the median is 0.45 and 0.55).⁷⁵

17.2.2 Environmental preferences

For couples, our findings confirm the high-level conclusion that women have more environmentally friendly preferences than men. As shown in Table 7, for overall consumption, men have a statistically significantly higher emission propensity than women: the total emissions of a household with a male dictator would be 67% higher than those of a household with a female dictator, and the same household budget; a very substantial impact.⁷⁶ The gender difference in carbon intensities is also significantly estimated: under male dictatorship, the carbon intensity of the basket of chosen goods is 45% higher than under female dictatorship; a stark difference.⁷⁷

To understand some more concrete implications of these estimates, let us consider a household with a weekly household budget of £3,983, which is the mean in the main sample⁷⁸. If this household had a female bargaining weight of 0.46, which is the main sample mean, then we would predict the household's overall level of emissions to be 679 *kgCo₂eq* per week.⁷⁹ We can repeat the same calculation for an equal household, where the bargaining weights are both set to 0.5, and then we obtain a predicted level of overall emissions of 665 *kgCo₂eq* per week. This is a 2.1% fall in household consumption emissions.⁸⁰ For comparison, Ivanova, Barrett et al. (2020) find that transitioning to a vegetarian diet could reduce the average North American citizen's annual per capita carbon footprint by 3.7%. To consider the impact relative to total GHG emissions, we weight our estimate by the proportion of relevant households in the UK, and the proportion of UK GHG emissions associated with households as opposed to other end-consumers (the government and non-profit organisations). As explained in appendix I, we estimate that transitioning to gender equal bargaining would reduce total direct and indirect GHG emissions by UK end-consumers by 1.1%.

Within narrower categories of goods, a somewhat different picture emerges relative to what

⁷⁵These estimated average female share is marginally higher in this pooled LCF sample than in the data used in the first chapter.

⁷⁶This number is obtained as the gap between male and female emissions propensities, divided by the female emissions propensity

⁷⁷Obtained as the difference between the carbon intensity for men and women, divided by the carbon intensity for women.

⁷⁸Recall that this is full income rather than disposable income, i.e. it is what the household would earn if both members worked 24 hours a day, 7 days a week.

⁷⁹This is obtained by multiplying the household budget by the weighted average of the male and female emissions propensities

⁸⁰Calculated as the difference between the emissions of the average household and the counterfactual equal household, divided by the emissions of the average household.

we saw for singles. For food, while women continue having a lower preference parameter than men on the overall category, women in couples prefer a basket of goods with a statistically significantly higher carbon intensity than men in couples (while for singles, women also had lower carbon intensity for food). These two opposite effects lead to no statistically significant difference in emissions propensity on food for men and women in couples.

For transport, men and women in couples have similar emissions propensities. This appears to be due to the canceling out of two opposite forces: women preferring to spend slightly more than men on this category (the reverse of the finding for singles), but choosing statistically less carbon intensive goods within this basket (similarly to singles).

For energy, women in couples have a much lower emissions propensity than men, driven by a much lower preferred share of budget assigned to this category of consumption. At the same time, for couples, women’s preferred basket of goods within energy is much more carbon intensive than the preferred basket of goods for men, although the latter is not statistically significant. This is a reversal of each of the findings for singles on energy. We note that it is particularly challenging to interpret findings for energy because a substantial component of observed emissions patterns from energy may be attributable to heterogeneity between the available energy sources in different UK homes during the period being considered. Part of the reversal in findings between singles and couples on energy may also be related to the fact that singles and couples often consider different sets of potential homes from each other, which may have different profiles in terms of energy systems and the possibility of upgrading or not.

For clothing, we find similar emissions propensities for men and women, with noisy measurement.

The partial divergence between findings for singles and couples is consistent with previous findings in the literature, e.g. Hubner (2020), that preferences are generally not stable across household compositions. This makes it particularly important to develop approaches, such as the one proposed in this paper, to uncover gender differences in preferences of men and women in couples directly from couples data, combined with estimates of bargaining power, rather than relying on findings from singles data. There are numerous reasons why preferences may differ between household compositions, including sample selection effects and a genuine transformative effect of becoming part of a couple. Future work may extend our analysis to other household compositions.⁸¹

⁸¹we have not attempted this here due to substantial additional complications when dealing with members who don’t do paid work, and children.

17.2.3 Alternative proxies of bargaining power

We compare our baseline findings to what would happen if we were to use alternative estimates of bargaining power. We consider three options:

1. **Breadwinner approach**, as sometimes used in the environmental literature. The breadwinner (the household member with the highest weekly pay) is assumed to have full control over all the consumption decisions of the household (i.e. is a dictator).
2. **Wage ratio**, i.e. the hourly wage of the woman divided by the sum of the hourly wages of the woman and the man in the household. This is a commonly used proxy for bargaining power in the household economics literature because it captures an important aspect of the difference in earnings potential of men and women if they were to separate and live on their own individual income.
3. **Income ratio**, i.e. the actual income of the woman divided by the sum of her income and the man's. This is a half-way approach between the breadwinner proxy from the environmental literature and the wage ratio proxy from the household literature. Like the wage ratio, it is a continuous measure rather than a binary, but like the breadwinner approach, it focuses on earned income rather than potential earnings.

Using the breadwinner approach, for couples we do not find any statistically significant gender differences in terms of emissions propensity, as reported in Table 13. This is consistent with these proxies being noisier estimates than our structural bargaining estimates, since they only incorporate a limited signal of bargaining (relative income or wages) rather than also drawing on information on other characteristics, e.g. education, and responsiveness of time-use to household budget. With the breadwinner approach, the direction of effects is suggestive of overall somewhat lower emissions propensities for women on all categories apart from clothing. Using this proxy, women in couples are estimated to have significantly lower carbon intensity than men in couples for all categories of expenditure but energy.

Using the wage ratio proxy, we find similar results on emissions propensity: nothing is statistically significant, but the directions indicate lower emissions propensity for women on all categories except clothing. As shown in Table 14, this is driven by both lower budget share spent on these categories, and lower carbon intensity of choices, except for transport and clothing, although none of these gender differences are statistically significant. The latter finding for transport is in contrast with most evidence from the literature.

Using the income ratio, as outlined in Table 15, we find that women have statistically significantly lower emissions propensity overall, but not for any specific category of goods.

Carbon intensity of choices is never statistically significantly lower for women than men. The overall findings are similar to our baseline ones qualitatively, but using this proxy, we estimate a substantially lower average bargaining power for women than in our baselines: 0.43 instead of 0.46. Driven by this, the impact of moving to equal bargaining (equal earnings here) is estimated to be a 3.2% decrease in emissions, which is higher than our baseline estimate of a 2.1% reduction.

Table 13: Proxy for bargaining power with gender of the breadwinner

	(1) Overall	(2) Food	(3) Transport	(4) Energy	(5) Clothing
Fem. Budg.	0.167*** (0.002)	0.035*** (0.001)	0.045*** (0.001)	0.052*** (0.001)	0.005*** (0.000)
Mal. Budg.	0.171*** (0.002)	0.036*** (0.000)	0.045*** (0.001)	0.054*** (0.001)	0.004*** (0.000)
π_m	0.171	0.036	0.045	0.054	0.004
π_f	0.167	0.035	0.045	0.052	0.005
Significant diff.	no	no	no	no	no
α_m	0.117	0.033	0.013	0.006	0.010
α_f	0.119	0.033	0.014	0.007	0.010
Significant diff.	no	no	no	no	no
CI_m	1.467	1.082	3.352	8.239	0.458
CI_f	1.408	1.060	3.265	7.909	0.457
Significant diff.	yes	yes	no	yes	no
N	3,724	3,724	3,724	3,724	3,724
R-squared	0.815	0.750	0.596	0.688	0.298

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is restricted to households belonging to the middle 50% of the household budget distribution. Fem. Budg. and Mal. Budg are the regression estimates of Equation 2. π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 3, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 95% confidence level.

Table 14: Proxy for bargaining power with hourly wage ratios

	(1) Overall	(2) Food	(3) Transport	(4) Energy	(5) Clothing
HH Budget	0.187*** (0.015)	0.041*** (0.004)	0.045*** (0.007)	0.054*** (0.007)	0.004*** (0.001)
Female share	-0.034 (0.031)	-0.011 (0.008)	-0.001 (0.014)	-0.003 (0.014)	0.000 (0.003)
π_m	0.187	0.041	0.045	0.054	0.004
π_f	0.152	0.030	0.044	0.052	0.005
Significant diff.	no	no	no	no	no
α_m	0.124	0.037	0.014	0.006	0.010
α_f	0.111	0.030	0.013	0.007	0.010
Significant diff.	no	no	no	no	no
CI_m	1.504	1.110	3.242	8.910	0.433
CI_f	1.371	1.019	3.373	7.417	0.486
Significant diff.	no	no	no	no	no
N	2,321	2,321	2,321	2,321	2,321
R-squared	0.826	0.763	0.604	0.704	0.297

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is restricted to households belonging to the middle 50% of the household budget distribution, and in which inequality between members is limited (the female hourly wage ratio is constrained between 40 and 60%) HH Budg and Female share are the estimated coefficients of Equation 4 π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 5, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 95% confidence level.

Table 15: Proxy for bargaining power with weekly pay ratios

	(1) Overall	(2) Food	(3) Transport	(4) Energy	(5) Clothing
HH Budget	0.208*** (0.018)	0.040*** (0.004)	0.048*** (0.008)	0.062*** (0.008)	0.003** (0.002)
Female share	-0.077** (0.036)	-0.010 (0.009)	-0.004 (0.016)	-0.021 (0.017)	0.003 (0.003)
π_m	0.208	0.040	0.048	0.062	0.003
π_f	0.131	0.031	0.044	0.041	0.006
Significant diff.	yes	no	no	no	no
α_m	0.134	0.034	0.013	0.007	0.007
α_f	0.103	0.033	0.016	0.006	0.014
Significant diff.	no	no	no	no	no
CI_m	1.549	1.182	3.834	9.383	0.480
CI_f	1.276	0.941	2.818	6.759	0.444
Significant diff.	no	no	no	no	no
N	1,860	1,860	1,860	1,860	1,860
R-squared	0.812	0.764	0.604	0.647	0.304

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is restricted to households belonging to the middle 50% of the household budget distribution, and in which inequality between members is limited (the female weekly pay ratio is constrained between 40 and 60%.) HH Budg and Female share are the estimated coefficients of Equation 4 π_m and π_f are the emission propensities (retrieved from the regression estimates), α_m and α_f are Cobb Douglas preferences parameters (retrieved by estimating Equation 5, CI_m and CI_f give the carbon intensities of consumption. Significant diff. indicates whether differences between the estimates for men and women are statistically significant at the 95% confidence level.

17.2.4 Comparison to other findings in the literature

The closest study to ours is Büchs and Schnepf (2013), which also uses UK data, but finds that female-led households have higher emissions propensities than male-led households, conditional on other household characteristics. Our results are more in line with other findings in literature, including from Sweden and Spain. The differences between our paper and Büchs and Schnepf (2013) are likely due to differences in our methodological approaches and the emissions data used.

Firstly, we employ a structural estimate of household bargaining rather than a binary variable for whether the house is female-headed or male-headed. This explains part of the difference in our results: when we repeat our analysis using (i) the same sample as for our baseline analysis,

and (ii) a similar proxy to theirs, we find no statistically significant different emissions propensities for female and male headed households. This is consistent with this proxy being a noisier estimate than our structural bargaining estimates, which draw more fully on the rich characteristics and time-use data we have. Even so, in our robustness analysis, using a variety of proxies of bargaining, we never find that women have overall higher emissions propensities (as in Büchs and Schnepf (2013)).

The remaining difference in our findings is likely due to improvements in data over time. Similarly to us, Büchs and Schnepf (2013) combine LCF data with DEFRA emissions multipliers. However, while we use DEFRA’s 2023 emissions data release, that was not yet available at the time of writing of Büchs and Schnepf (2013), who instead had to rely on the 2011 release, which has been revised repeatedly in more recent releases. It is also useful to keep in mind that, while both of our papers contain applications to the UK, we focus on different samples, including different time periods and household compositions, so that the results should be compared with care.⁸²

17.3 Discussion of findings

Policy implications Policy-makers can intervene in a variety of ways to empower women within households, for instance by pursuing policies to narrow gender pay gaps and providing free childcare to all children. This paper suggests an additional, novel, rationale for such policies: they are likely to substantially lower household consumption emissions, and thus overall GHG emissions at a national level.

Limitations It is important to recall that the data used in this paper has some limitations, including the limited degree of granularity of categories for which we can obtain emissions multipliers. Moreover, our analysis does not encompass some consumption goods which are difficult to analyse with our expenditure and emissions data, such as the emissions associated with building and renovating homes, and pharmaceuticals. When more granular and extensive emissions data becomes available in future years, it would be important to repeat a similar analysis using it. Moreover, this paper focuses on the period 2001-2014. It would be interesting to check whether gender differences in emissions propensities have remained stable, or become more or less pronounced, over time, with growing overall environmental

⁸²Firstly, we cover a longer time period (they focus only on 2006-9). Secondly, we focus only on working singles and heterosexual working couples, while they consider all different household compositions in aggregate.

consciousness.⁸³

We also note that further work is needed to rule out potential confounding factors that are not addressed in this paper. In particular, we do not allow for preference heterogeneity within genders. In reality, this is likely to exist, for instance by age. Moreover, it is possible that couples form in a way that is assortatively matched so that couples tend to either involve two environmentally friendly people or two individuals who are not concerned about the environment (even though on average women may be more environmentally friendly than men). This in itself is not an issue for our analysis, unless environmental preferences are correlated with factors that also affect the distribution of bargaining power in the household. If that were the case, that could introduce bias into our results. For instance, it may be that some of our estimated gender differences are in fact driven by some couples being more left-wing and/or younger, and these couples being characterised by both higher female empowerment and greater environmental consciousness of both members of the couple. We do not allow for more heterogeneity in preference parameters in our approach because the data that we use to estimate bargaining power is too small to conduct the exercise separately for granular groups (e.g. younger people in left-wing regions vs older people in right-wing regions), and we leave it to future work to investigate these possibilities.

Finally, we note that findings are sensitive to the chosen measure of bargaining power, and that we obtain somewhat different conclusion in our baseline, using structural bargaining estimates, and using each of three alternative proxies of bargaining. This suggests the importance of carefully modeling and estimating bargaining power, rather than relying on a proxy. The approach we use itself makes substantial simplifications and assumptions, and it will be important for future work to use whichever methods emerge from the household economics literature as being more reliable. Currently, there is a gap in that literature when it comes to systematically assessing the predictive success of different approaches, but we hope that gap will be filled in the coming years, allowing applications to a range of contexts, including the environmental literature, to rely on best-performing approaches.

18 Conclusion

This paper contributes to the environmental literature by investigating gender as a determinant of carbon footprints of consumption both for singles and couples by using a revealed preference approach, rather than relying on reported attitudes. Moreover, this is the first

⁸³We do not have appropriate data to perform this analysis - in particular we cannot estimate bargaining power reliably for years after 2014.

paper to structurally estimate intra-household bargaining power to examine gender differences in emissions in two-person households. In so doing, this paper also contributes to the household economics literature, by exploring a new application of intra-household bargaining power estimates.

Consistent with the literature to date, our findings suggest that women in general have more environmentally friendly preferences than men. For single women, this is driven by a lower carbon intensity of the chosen basket of goods. For women in couples, this is driven both by a lower carbon intensity of preferred baskets of goods, and a preference to spend a lower proportion of the budget on goods considered in this analysis.

We do not attempt to explore the psychological mechanisms driving gender differences in preferences. However, our results, combined with existing evidence in the literature that women have stronger pro-environmental attitudes than men, are suggestive that awareness of environmental issues and concerns about them translates into concrete differences in emissions. Further work is needed to draw this link more strongly and explore potential avenues for policy to leverage this link to combat climate change.

Part III

Focal Pricing Constraints and Pass-Through of Input Cost Changes

19 Introduction

A firm sells goods to consumers, and incurs costs in the form of input costs and sales taxes. The degree to which changes in these costs are passed-through to changes in the prices paid by consumers is relevant in a variety of contexts, including tax incidence, merger control, and antitrust cases. Governments considering a hike in sales taxes are interested in the degree to which it will be absorbed by firms, and to what degree by consumers. Competition authorities deciding whether to allow a merger consider whether any potential efficiencies⁸⁴ from the proposed merger will in fact be passed through to consumers in the form of lower prices, and therefore may offset detrimental impacts from the merger, including reduced competitive pressure. Courts awarding antitrust damages rely on estimates of pass-through to quantify fair compensation. Because focal pricing has been raised as an issue in several class actions, the discussion in this chapter focuses on the case of indirect purchaser class actions. In these cases, it is alleged that the claimants (end-consumers) paid inflated prices for end-products because input costs of production were inflated due to abusive conduct by one or more upstream firms (producers of the input). In order to estimate appropriate damages, it is necessary to estimate the extent to which inflated input costs were passed through to consumers, instead of being absorbed by downstream firms (purchasers of the input and producers of the end-product).⁸⁵

If a £2 increase in input costs leads to a £2 increase in prices, then there is ‘complete’ pass-through, or an 100% pass-through rate. Depending on the context, pass-through rates can be below, above, or at 100%. They can be affected by numerous industry-specific characteristics, including the nature of competition between firms, the curvature of demand, returns to scale, and any frictions in setting prices. For instance, Bulow and Pfleiderer (1983) discuss the theoretical impact of different demand curvatures on pass-through rates in the context of tax incidence. This chapter aims to answer a question that has received very little academic

⁸⁴Merger efficiencies are cost reductions resulting from the merger.

⁸⁵In some cases there may be additional layers in the industry, e.g. the input may be sold by an upstream firm to an intermediate firm to produce an output which is then sold by a retailer. The logic remains the same, though with a chain of pass-through.

attention to date: the impact of focal pricing constraints on pass-through rates. To my knowledge, this question has been addressed only by Alexandrov (2013), who sets out a high-level argument why focal pricing will typically not have a substantial impact on pass-through. I formally prove that the expected pass-through rate is unchanged by the presence of focal pricing constraints, under a simple but general framework. I refer to this as the Irrelevance Theorem. The Irrelevance Theorem holds exactly under several standard models of competition between firms, and different demand functions, and as an approximation in other cases. I also discuss the distribution of pass-through, which is generally characterised by more dispersion under focal pricing constraints.

Focal pricing is a widely observed phenomenon, consisting in firms only charging prices with specific characteristics. These are often prices with 9s in the last digits, as noted e.g. by Levy et al. (2011), and Snir et al. (2017). This can be explained if firms face constraints in the form of consumer inattention to certain digits, so that demand is unchanged within non-singleton sets of prices. This phenomenon in turn could be formalised with behavioural models (see e.g. Strulov-Shlain (2023)) or as rational inattention (see e.g. Basu (1997)). Focal pricing constraints may also arise for other reasons. For instance, Knotek (2008) and Knotek (2011) show that ‘convenient’ prices (multiples of cash denominations) are widespread for frequently purchased goods paid in cash in high-traffic transactions. Moreover, virtually all firms employ focal pricing to some degree, because money denominations constrain most prices to being multiples of pence. We can think of these phenomena in terms of firms facing exogenous constraints on the set of prices they are able to choose from. Alternatively, we might think of a step-wise demand function, where several prices have equivalent effects on quantity demanded. Regardless of the exact nature of these constraints, they make it optimal for firms to charge focal prices, rather than non-focal alternatives.

The focal pricing constraints I study are distinct from two other closely related phenomena. The first are endogenous focal prices, which may emerge for instance as a mechanism to promote tacit collusion in an oligopolistic setting (see e.g. Scherr (1981)). This is a very different context from the exogenous focal pricing studied here, since firms are in fact able to charge non-focal prices, and in some cases might profit from doing so instead of charging a focal price, at least in the short-run. The other related, but distinct, phenomenon is menu costs. Menu costs involve an actual cost to changing prices, which is not present with focal pricing. Moreover, under menu costs, firms are able to price continuously rather than discretely, unlike with focal pricing. While they are distinct phenomena, they are easily conflated because they both cause price rigidity (stickiness), i.e. prices being changed infrequently, and typically not in very small increments. The impact of price rigidities on pass-through has been explored, with a focus on menu costs. For instance, Nakamura and

Zerom (2010) investigate different potential explanations for why there is incomplete (sub-100%) pass-through of exchange rate fluctuations. Among other candidates, they include price rigidity (which in their model is explained through menu costs); finding that they introduce a time lag in pass-through, but have a negligible impact on long-run pass-through rates.⁸⁶ In this chapter, I focus only on the impact of exogenous focal pricing constraints and not on other related phenomena.

This chapter is motivated by debunking the misconception that the presence of focal pricing constraints in general reduces pass-through rates. In multiple high-profile antitrust cases, it has been claimed that if an industry adopts focal pricing that in itself implies there will be very little, if any, pass-through of input cost changes. The argument is that if firms round to certain special prices, they are unlikely to adjust their prices in response to small input cost changes. In *re Lithium Ion Batteries antitrust litigation*, the Court struck out the case, amongst other reasons, because the defendants argued that the plaintiffs’ expert analysis of pass-through did not take into account focal pricing. The defendants alleged that “focal point pricing is prevalent in the pricing of products within the class definition, and will result in no pass-through when a small cost change—such as the estimated \$2.16 overcharge for a notebook computer battery here—in presence of focal points that are wider apart than the cost difference itself.”⁸⁷ A similar argument has been made by Qualcomm in antitrust litigation in several jurisdictions.⁸⁸

While this argument has intuitive appeal, it fails to recognise the other side of the coin: in the presence of focal pricing constraints, some prices will be over-adjusted if the input cost change leads to a jump from one focal price to another. Consider the example of laptops, and let their prices be constrained to end in 9. Then it is true that an overcharge on batteries of \$2.16 could result in no impact on the focal price charged, but it could also lead to a \$9 jump (or more). This is confirmed empirically, e.g. Levy et al. (2011) find price changes are less frequent, but bigger in magnitude, with focal pricing. It is also easy to empirically disprove the claim that focal pricing implies very low pass-through rates: for instance, see Conlon and Rao (2020) for examples of high pass-through in industries which adopt focal pricing.

While there is empirical evidence that pass-through can still be high in contexts with focal

⁸⁶The authors find that exposure to local costs plays an important role in dampening the impact of exchange rate fluctuations, and that curvature of demand also explains the observed pass-through patterns.

⁸⁷United States District Court Northern District of California *In re Lithium Ion Batteries Antitrust Litigation*, Case No.: 13-MD-2420 YGR (2018, March).

⁸⁸For example, see United States District Court Northern District of California San Jose Division *In re: Qualcomm Antitrust Litigation, Defendant Qualcomm Incorporated’s Opposition to Plaintiff’s Motion for Class Action*, Case No. 5:17-md-02773-LHK-NMC (2018, September).

pricing, little attention has been devoted to theoretical modelling of this issue. This may be due to the fact that many models of competition become highly complex when pricing is discrete. Filling this theoretical gap is important to complement the existing empirical evidence. Under a simple but general framework, I show that focal price constraints have no impact on expected pass-through rates. The claim is not that focal pricing *cannot* impact pass-through rates, in fact it often will; but rather that there is *no general reason* to expect pass-through rates to be lower in the presence of focal pricing constraints. If we knew exactly how an industry worked, in full detail, we might be able to say that for a specific input cost overcharge none of it was passed through to downstream consumers, or conversely that it was passed-through at a higher than 100% rate. However, it is almost impossible for this sort of information to be available in detailed and reliable form at the outset of a court case, even though some of this information may be obtained through a long process of disclosure and economic analysis. Hence, it is not safe to dismiss a court case based on the assumption that the adoption of focal pricing will reduce pass-through relative to a similar context without focal pricing. Instead, as usual, it is best to estimate pass-through rates specifically for the context of interest. Note that the misconception that focal prices lead to little, or no, pass-through, has substantial economic consequences. For instance, the dismissal of class actions covering millions of claimants, and very large total damage claims, has a direct impact on the claimants, who do not receive fair damages, and an indirect impact on firms, who are more likely to engage in abusive behaviour if they are likely to avoid paying for damages.⁸⁹

I also discuss the distribution of pass-through across different purchases. I show that focal pricing generally increases the dispersion in the pass-through rate. This is particularly relevant in class actions for two reasons. Firstly, the class certification stage often requires demonstrating sufficient homogeneity of the class. This involves different requirements in different jurisdictions, but broadly covers the idea that claimants must have been injured in a sufficiently similar manner, and that similar methods may be used to estimate the damages to be awarded to different claimants. Secondly, the quantification of damages may need to be performed separately for sub-classes which were harmed to different extents. Where claimants purchased few, or a single, product(s), it is likely that different pass-through rates applied to different claimants. With multiple purchases at different points in time, the average pass-through rate is a reasonable estimate of the pass-through rate that applies to

⁸⁹Note that, if abusive upstream firms successfully argue that there was no pass-through in consumer class actions, then that may increase the risk of them having to pay damages to downstream firms (who would have absorbed the inflated cost in the absence of pass-through). However, for that to happen, the downstream firms would have to start their own litigation with the upstream firms. Downstream firms may be unwilling to enter into litigation with their input producers, especially if those producers have market power and have already engaged in abusive practices. For instance, they may fear retribution in the form of disruption to their input supply.

all claimants. In the former case, it may be best to estimate pass-through rates separately for different groups of products. In the latter case, it can be legitimate to not explicitly consider focal pricing in the economic analysis.

This chapter is set out as follows. First, in section 20.1, I set out a general model to assess the impact of focal pricing constraints on pass-through. Then, in section 20.2, I prove the Irrelevance Theorem under a set of assumption. I discuss each of these assumptions in turn, arguing that they are general. To further substantiate that discussion, in section 21, I analyse different models of competition, and different curvatures of the demand function, relating them back to the assumptions required for the Irrelevance Theorem to hold. In section 22, I conclude with a brief discussion of my findings and their relevance to antitrust damages class actions.

20 A general framework to assess the impact of focal pricing constraints on pass-through

20.1 The framework

Consider a firm facing constant marginal costs c and choosing prices to maximise profits. Here, I remain agnostic about the number of firms in the market, the type of competition, and curvature of demand.⁹⁰ I explore specific examples in section 21. To understand how focal pricing constraints affect pass-through, I compare pass-through rates under two scenarios: an unconstrained optimisation problem and a constrained optimisation problem.

In the unconstrained optimisation problem, the firm has a strategy $p^u(c)$ mapping any possible cost level to an optimal price.

Now consider the case where there are focal pricing constraints in the market. In this case, the quantity demanded of the firm's product will no longer depend on the charged price, but on the corresponding focal price f . Prices within a certain interval will be associated with a specific focal price: $f = f_i \quad \tau_{i-1} < p \leq \tau_i$. For instance, we can imagine that consumers are inattentive to some digits, so quantity demanded is a function of price rounded up to

⁹⁰Where there are n symmetric firms, the single firm we consider here can be taken to be a representative firm, so that if pass-through rates are not affected by focal pricing for this one firm, they are not affected at an industry level. Where there are multiple asymmetric firms, with potentially different pass-through rates, we can in principle repeat the same analysis for each of them - this chapter focuses on symmetric settings, and asymmetric ones are left to future work.

the next focal price.⁹¹ In this case, $\tau_i = f_i$ and $f = f_i$ $f_{i-1} < p \leq f_i$.⁹² In general, it will be optimal for the firm to charge only focal prices.⁹³ Therefore, while in the unconstrained problem the firm's strategy was $p^u(c)$, now it takes the form $p^c = f_i$ $t_{i-1} < c \leq t_i$.

It is useful to consider for what value of costs a given focal price would be unconstrained optimal: $f_i = p^u(\chi_i)$. Then we can define the gap between costs at which consecutive focal prices are unconstrained optimal as: $\chi_{i+1} - \chi_i$.

Let the firm's input cost increase from c to $c + \Delta$. The pass-through rate of an increase in marginal cost of size Δ in the unconstrained case is:

$$\frac{p^u(c + \Delta) - p^u(c)}{\Delta}$$

The pass-through rate of the same increase in marginal cost of size Δ in the constrained case is:

$$\frac{p^c(c + \Delta) - p^c(c)}{\Delta}$$

The question at hand is how the unconstrained pass-through rate compares to the pass-through rate under focal pricing constraints. I answer this question formally under the following set of assumptions.

Assumption Set A

1. **Regularity Condition** for focal prices. The distance between each consecutive focal price is consistent: $f_{i+1} - f_i = G$.
2. **Equal Spacing condition** for cost thresholds. The distance between consecutive cost thresholds is constant, and is the same as the distance between costs at which consecutive focal prices are unconstrained optimal. We can therefore write $t_{i+1} - t_i = \chi_{i+1} - \chi_i = \theta G \quad \forall i$. This assumption implies that the unconstrained pass-through rate is constant (see the proof of the Irrelevance Theorem).
3. **Uniformity Condition.** If we observe a firm charging the focal price f_i we can infer that $c \sim U(t_i, t_{i+1})$.

⁹¹The concept of a latent unconstrained demand function can be thought of as representing the demand function that consumers tend towards as frictions decrease.

⁹²Results readily extend to rounding up, or to the nearest, focal price.

⁹³In certain oligopolistic settings, it may be possible to have multiple equilibria, some of them involving charging non-focal prices, but we can select equilibria where focal prices are charged to match our empirical observations.

Discussion of Assumption Set A

The Regularity Condition is likely to cover the vast majority of real-life cases which, as discussed above, involve consistently rounding to prices ending in specific digits, generally 9s. Moreover, the discretisation of prices due to currency limitations is also an example of regular focal pricing. Regular spacing of focal prices is a convenient modelling assumption, but it can be dropped with the Irrelevance Theorem still holding, at least approximately. If we wanted to drop the Regularity Condition, we would also relax the Equal Spacing condition for cost thresholds, to require that the (variable) gaps between cost thresholds are approximately the same as the (variable) gaps between costs at which consecutive price thresholds are unconstrained optimal.

The Equal Spacing Condition is key for the Irrelevance Theorem to hold exactly. As discussed in section 21, this assumption holds for monopolistic, perfect, and differentiated Bertrand competition with linear demand. Under perfect competition, the Equal Spacing assumption holds with any demand function. In the monopolistic case the assumption holds exactly for demand functions other than linear demand, including the logarithmic demand function. Under monopolistic competition, the Equal Spacing Condition holds approximately for any curvature of demand (including demand functions for which the unconstrained pass-through rate is not constant over the cost interval).

The Uniformity Condition is necessary for the Irrelevance Theorem to hold – if it does not then the result needn't hold even approximately. This assumption is likely to hold under any realistic context where the Irrelevance Theorem may be usefully deployed, i.e. in situations where we have limited information about the context of interest. This would certainly be the case in the early stages of a class action, when funding has not yet been obtained to carry out in-depth analysis, and disclosure has not been granted. We are also likely to face substantial lack of information at later, intermediate, stages of court cases. In such situations, the Uniformity Condition represents our uncertainty prior to a full analysis. By definition, if we observe a firm charging the focal price f_i we know that $t_i < c < t_{i+1}$. Moreover, in general we do not have sufficiently detailed and reliable information (prior to detailed empirical investigation) whether the marginal cost is close to the upper bound of this interval or not, which is relevant to whether focal pricing constraints will lead to reduced pass-through, or increased pass-through. It is difficult enough to obtain accurate cost estimates, even with court mandated disclosure, which can be challenging to obtain.

Consider the example of an antitrust damages class action. The claimants are end-consumers of a product which included an input which allegedly had an overcharge on it. The defendants are firms who sold the input to downstream firms which in turn produced and sold a product,

with this input, to the consumers (for simplicity, I refer to them as the retailers). The claimants, and the defendants, do not possess detailed information about the cost structure and pricing strategy in the relevant retail industry. Moreover, the retail firms generally have no incentive to voluntarily disclose this sensitive information unless a court requires them to do so. Since they are not the ones accused of wrong-doing, obtaining disclosure is particularly hard. It is possible in some cases to request disclosure of certain documents, but it can be a lengthy and challenging process, including difficulties such as the redaction of sensitive information which can be crucial to obtaining a full picture. Even after obtaining access to documents, these are typically so numerous and lengthy that it is very resource-intensive to obtain a clear picture of firm strategy from their documents. It might also be possible to interview people who work in the sector, but it can be challenging. This is partly because the abusive firm(s) are likely to command substantial influence in the sector, so that it is challenging to find reputable sources willing to go against them in court.

Even with full disclosure, we are very unlikely to be able to obtain reliable information on the full contingent pricing strategy of a firm, as it may well not exist in written form, or be subject to frequent discussion and alteration. Therefore, we are unlikely to know where the marginal cost was located relative to cost thresholds prior to the alleged overcharge. Hence, prior to detailed analysis (which is warranted by the Irrelevance Theorem) it is reasonable to assume the Uniformity Condition (which allows us to prove the Irrelevance Theorem). It is hard to construct examples where this is not the case in a manner which undermines the Irrelevance Theorem.

The first way in which the Uniformity Condition may not hold is if we have already conducted detailed empirical analysis. In this case, it is no surprise that the theorem is no longer valid: having performed detailed analysis, we will have information on whether in the specific case focal pricing increased, decreased, or did not affect pass-through. The Irrelevance Theorem is about the expected value of pass-through prior to detailed analysis. The point is that, prior to detailed analysis, we cannot assume that focal pricing led to low pass-through, and hence that empirical analysis is required. The second way in which the Uniformity Condition may not hold is if we have reason to believe that focal pricing is not an exogenous constraint (e.g. due to consumer inattention to digits), but an industry practice which is endogenous to firm optimisation (e.g. focal pricing as a tacit collusive strategy). In this case, costs will be located in relation to thresholds in a profit-maximising manner. This is a different phenomenon from the one addressed in this chapter. It is potentially of interest too, but challenging to analyse because it is best conceptualised in a dynamic oligopolistic setting, with the ensuing multiplicity of equilibria; for this reason, it is left to future work. The third way in which the Uniformity Condition may not hold is if firms can easily adjust

non-price characteristics (such as quality, pack-size, components included in a bundle..). Then we might expect that firms would adjust their marginal cost, by adjusting these other characteristics, so that the focal price being charged is as close to unconstrained optimal as possible. But then the reason to suspect focal pricing may dampen pass-through is moot: holistic pricing, taking into account non-price characteristics, is (almost) unconstrained, so that we would expect the holistic pass-through, capturing characteristic-adjusted prices, to be the same regardless of the presence of focal pricing.

20.2 The Irrelevance Theorem

Irrelevance Theorem *Under Assumption Set A, the expected pass-through rate is the same in the unconstrained setting, and in the setting with focal pricing constraints.*

Proof First consider the unconstrained case. By the definition of χ_i , we can write the pass-through rate of an input cost increase from χ_i to χ_{i+1} as:

$$\frac{f_{i+1} - f_i}{\chi_{i+1} - \chi_i}$$

By the Regularity Condition and Equal Spacing condition, we can write:

$$\frac{f_{i+1} - f_i}{\chi_{i+1} - \chi_i} = \frac{G}{\theta G} = \frac{1}{\theta} \quad \forall i$$

Since we can write this for any value of the focal prices, it follows that the unconstrained pass-through rate is constant. Let the firm's input cost increase from c to $c + \Delta$, then:

$$\frac{p^u(c + \Delta) - p^u(c)}{\Delta} = \frac{1}{\theta}$$

$\theta = 1$ corresponds to complete pass-through of 100%, $\theta > 1$ to incomplete pass-through, and $\theta < 1$ to pass-through above 100%.

In the absence of focal pricing constraints, the charged price increases by:

$$p^u(c + x) - p^u(c) = \frac{\Delta}{\theta}$$

Now consider the constrained problem. By the Equal Spacing assumption, the price charged will jump by $\lfloor \frac{\Delta}{\theta G} \rfloor$ focal prices with certainty where $\lfloor \cdot \rfloor$ is the floor operator. If the change in input cost is weakly greater than the gap between cost thresholds, $\Delta \geq \theta G$, then the charged price is sure to increase. By the Regularity Assumption, each of these jumps entails a price

change of G . Therefore, the charged price will increase by $\lfloor \frac{\Delta}{\theta G} \rfloor G$ with certainty. Where $\Delta < \theta G$, we do not have certainty about any price jumps.

Additionally, by the Uniformity Assumption and Equal Spacing Assumption, there is a $\frac{\Delta \text{mod} \theta G}{\theta G}$ probability of a further jump in focal price (where *mod* is the modulo function). By the Regularity Assumption, this further jump, if it occurs, would lead to an additional increase of G in the charged price.

In expectation, the input price increase therefore leads to the following change in the constrained optimal price charged:

$$E[p^c(c + \Delta) - p^c(c)] = \lfloor \frac{\Delta}{\theta G} \rfloor G + \frac{\Delta \text{mod} \theta G}{\theta G} G = G(\lfloor \frac{\Delta}{\theta G} \rfloor + \frac{\Delta \text{mod} \theta G}{\theta G})$$

By definition of the modulo and floor operators:

$$E[p^c(c + \Delta) - p^c(c)] = G \frac{\Delta}{\theta G} = \frac{\Delta}{\theta} = p^u(c + x) - p^u(c)$$

The expected pass-through rate is $\frac{1}{\theta}$ regardless of the presence of focal pricing constraints. **QED.**

As I discuss in section 21.4, while it is conceivable that in some specific settings the Irrelevance Theorem may not hold, it is hard to find realistic examples of this. Therefore, in general it is not safe to assume that focal pricing reduces pass-through.

20.3 The distribution of pass-through

In general, the presence of focal pricing increases the dispersion of pass-through. Consider an industry where, in the absence of focal pricing constraints, all firms have a pass-through rate of $\frac{1}{\theta}$. Then the price change associated with a cost increase of size Δ ⁹⁴ is always $\frac{\Delta}{\theta}$, and there is 0 variance in pass-through.

The introduction of focal pricing constraints increases the dispersion in pass-through rates. Let $i\theta G < \Delta < (i + 1)\theta G$, where i is a weakly positive integer. Then the impact on the price of a specific product at a specific time is iG with probability $1 - \frac{\Delta \text{mod} \theta G}{\theta G}$ and $(i + 1)G$ with probability $\frac{\Delta \text{mod} \theta G}{\theta G}$. Therefore, in the presence of focal pricing constraints, pass-through may be higher or lower than in the unconstrained problem. The probability with which it is higher (or lower) depends on the size of the cost change relative to the gap between cost thresholds. For instance, if the cost change is much smaller than the gap between focal prices,

⁹⁴e.g. because of an overcharge of Δ on an input.

it will result in no pass-through with higher probability, and extremely high pass-through with lower probability. The variance is now:

$$\begin{aligned} E [(p^c(c + \Delta) - p^c(c))^2] - E [p^c(c + \Delta) - p^c(c)]^2 \\ = \left(1 - \frac{\Delta_{mod}\theta G}{\theta G}\right) (nG)^2 + \frac{\Delta_{mod}\theta G}{\theta G} ((n + 1)G)^2 - \left(\frac{\Delta}{\theta}\right)^2 \end{aligned}$$

By Jensen's inequality, the variance is weakly positive under focal pricing constraints, and hence higher than in the unconstrained case.

In the context of a class action, where each member of the class purchased a single product, focal pricing constraints may lead to substantial heterogeneity in the pass-through rates which apply to each claimant,⁹⁵ so that it may be best to estimate pass-through rates specific to sub-classes. Where each member purchased multiple products, at different times, by the law of large numbers, we might expect the average pass-through rate experienced by each claimant to be similar to the expected pass-through rate. In this case, it may be sensible to estimate the pass-through rate as one would in the absence of focal pricing constraints, and consider a single number to be an appropriate estimate for all claimants.

21 Examples of specific models

In this section, I show that the Equal Spacing assumption holds in a variety of standard models, and hence that the Irrelevance Theorem is widely applicable. Given this chapter's focus on exogenous constraints on pricing, it is more natural to consider firms as price-setting, rather than quantity-setting. The rest of this discussion is therefore grounded in models of price competition.⁹⁶

In turns, I consider three standard competitive frameworks: monopolistic competition, perfect competition, and differentiated Bertrand competition. In the context of monopolistic competition and perfect competition, I also discuss the impact of different curvatures of demand.⁹⁷ Throughout this discussion, I maintain a number of basic assumptions. In par-

⁹⁵Other input costs change over time, so even in a fully symmetric industry there may be heterogeneity in pass-through due to other costs impacting the degree of pass-through of the overcharge Δ .

⁹⁶It is left to future work to consider how we might incorporate focal pricing constraints in a quantity-setting context, and whether in specific models of quantity competition the Irrelevance Theorem does not hold.

⁹⁷It is left to future work to consider different curvatures of demand with differentiated Bertrand competition.

ticular, I focus on static models, in which all firms face the same constant marginal cost c (so when there is a change in that cost, it is faced by all firms). In section 21.4 I briefly discuss what may happen when relaxing these assumptions.

21.1 Monopolistic competition

21.1.1 Linear demand

I start by considering the simplest possible model, with a single monopolist selling a single good, facing a linear demand function and marginal cost c . The monopolist can only change the price of the good, not any non-price characteristics. As discussed earlier, when firms can alter non-price characteristics, focal pricing constraints are likely to affect pass-through rates, so that I am focusing on the most clear case where focal pricing could matter to pass-through.

The monopolist's optimisation problem is:

$$\max (\alpha - \beta p) (p - c) \quad \alpha, \beta > 0$$

The optimal price is $p^u = \frac{\alpha}{2\beta} + \frac{c}{2}$ and we obtain the standard result of 50% pass-through under monopolistic competition. In this case, the profit function is symmetric around the optimal price, so that:

$$(\alpha - \beta (p^u + x)) ((p^u + x) - c) = (\alpha - \beta (p^u - x)) ((p^u - x) - c)$$

Now consider that the monopolistic faces focal pricing constraints: consumers are inattentive to some digits, so that demand is a function of price rounded up to the next focal price $f = f_i$ $f_{i-1} < p \leq f_i$.⁹⁸ Then the monopolist's constrained problem is $\max (\alpha - \beta f_i) (p - c)$ $\alpha, \beta > 0$. The profit maximising price in this constrained problem must be a focal price, because for any other price it would be possible to increase the price up to the next focal price without decreasing demand. Moreover, because of the symmetry of the profit function around the optimal unconstrained price, it is always constrained optimal for the monopolist to charge the focal price nearest to the optimal unconstrained price.

In appendix J, I show that the Equal Spacing Condition holds, with $\theta = 2$. With uncertainty about the full details of the context of interest (the Uniformity Condition), the conditions for the Irrelevance Theorem hold: the expected pass-through rate is 50% regardless of the presence of focal pricing constraints.

⁹⁸Results readily extend to rounding up, or to the nearest, focal price.

21.1.2 Non-linear demand

In practice, demand may well not be linear, and the degree of curvature of demand is an important determinant of the rate of pass-through of cost changes, as discussed by Bulow and Pflaederer (1983).⁹⁹ In the monopolistic case, the Equal Spacing assumption holds exactly under broader demand functions than linear demand, for instance with the logarithmic demand function $p = \alpha - \beta \ln q$ $\alpha, \beta > 0, 0 < q < e^{\alpha/\beta}$.¹⁰⁰ In this case, the unconstrained monopolist chooses $p^u = c + \beta$, so there is a constant mark-up, and ‘complete’ pass-through of 100%. The monopolist constrained by focal prices (spaced out at regular intervals G) chooses which focal price to charge based on a cut-off rule. As shown in appendix K, these cut-offs are evenly spaced at the same regular intervals G . Since the gap between costs at which consecutive focal prices are unconstrained optimal is also G , the Equal Spacing assumption holds. In the presence of focal pricing constraints, pass-through remains 100%.

With other demand functions, we could have two sources of failure of the Equal Spacing assumption. The first is non-constant pass-through rates. Then the requirement that the gaps between cost thresholds is constant will not hold. This is not a fundamental issue for the Irrelevance Theorem, as long as the (variable) gaps between thresholds are approximately the same as the (variable) gaps between costs at which consecutive price thresholds are unconstrained optimal. If this is the case, then we can adjust the proof of the Irrelevance Theorem to account for non-constant pass-through, and will obtain that focal pricing leaves local pass-through approximately unchanged.

Even where unconstrained pass-through rates are constant, as in the case for the class of demand functions discussed in Bulow and Pflaederer (1983), for some curvatures of demand, the Irrelevance Theorem will hold approximately, rather than exactly. This will happen if the Equal Spacing condition does not hold locally everywhere, in which case we can see that it will hold as a global approximation. The key intuition here is the following. It is possible, locally, for cost thresholds to be more closely (or widely) spaced out than the costs at which focal prices are unconstrained optimal. Then, locally, focal pricing reduces or increases pass-through because we cannot perform the cancelling out operation that leads us to the Irrelevance Theorem. However, it cannot consistently be the case for cost thresholds to be non-negligibly more closely (or widely) spaced out than the costs at which focal prices are unconstrained optimal. This is because the cost at which a specific focal price is unconstrained optimal must be in the interval between the cost thresholds where this specific

⁹⁹There are other important determinants, such as returns to scale, which are left for future work to consider.

¹⁰⁰It is left to future work to determine whether there is a specific class of demand functions for which this is the case.

focal price is charged in the constrained problem: $t_{i-1} \leq \chi_i \leq t_i$. If we had consistently non-negligibly narrower (or wider) gaps between cost thresholds $t_{i+1} - t_i$ than between the costs at which focal prices are unconstrained optimal $\chi_{i+1} - \chi_i$, then at some point we would find that a cost at which a specific focal price is optimal is in fact outside of the interval when that focal price is charged in the constrained problem:

$$t_{i-1} \geq \chi_i \quad \text{or} \quad \chi_i \geq t_i$$

This is a contradiction. Hence, focal pricing may increase (or lower) expected pass-through over specific segments of the cost interval. However, these effects are local, and are typically balanced out by opposite effects on other segments of the cost interval.¹⁰¹ Therefore, without very detailed information about the industry, of a variety which is very unlikely to be readily available (even subject to court disclosure), we cannot know a priori whether we are in a cost segment where the presence of focal pricing locally increases or decreases expected pass-through. Moreover, we know that, globally, expected pass-through will be approximately unaffected by the presence of focal pricing. Therefore, the Irrelevance Theorem may be considered to be approximately true regardless of the curvature of demand, in a monopolistic setting.

21.2 Perfect competition

Consider a market characterised by perfect competition (for instance, we might conceptualise this in terms of undifferentiated Bertrand competition). The aggregate demand function Q can take any form, and depends only on the lowest price offered by any firm j in the market $\min_j \{p_j\}$. There are N firms in the market, and the demand function faced by a specific firm n takes the following form:

$$q_n = \begin{cases} \frac{Q}{\sum_j I(p_j = p_n)} & \forall j \quad p_n \leq p_j \\ 0 & \exists j \neq n \quad p_n > p_j \end{cases}$$

All firms face the same marginal cost c . The unique equilibrium involves all firms charging $p = c$, and hence each firm facing demand $q_n = \frac{Q}{N}$, and making zero profits. In this context,

¹⁰¹I note that it is possible for pass-through to be mildly higher (or lower) in the presence of focal pricing over some of the region, without being cancelled out by the opposite effect, as long as pass-through is almost the same as without focal pricing. For instance, it is possible to construct examples with the constant elasticity demand curve where pass-through with focal pricing constraints is slightly lower than in the unconstrained problem for low value of input costs, and tends towards the same pass-through rate as costs increase. In this case, expected pass-through is only very slightly reduced, so that the Irrelevance Theorem can be said to hold approximately.

there is complete pass-through of input cost changes. Therefore, a Δ overcharge will result in a Δ increase in prices.

We now introduce focal pricing constraints for the firms. For instance, consider a case of consumer inattention to certain digits, so that demand is a function of price p_n rounded up to the nearest focal price $f_n = f_i$ $f_{i-1} < p_n \leq f_i$.¹⁰²

$$q_n = \begin{cases} \frac{Q}{\sum_j I(f_j = f_n)} & \forall j \quad f_n \leq f_j \\ 0 & \exists j \neq n \quad f_n > f_j \end{cases}$$

Once again there is a unique equilibrium, and each firm faces demand $q_n = \frac{Q}{N}$, but now all firms charge $p = f_c$ where f_c is the marginal cost c rounded up to the next focal price.¹⁰³ Therefore, all firms now make weakly positive profit. We can easily see that the cost thresholds at which firms adjust prices are identical to the focal prices themselves. Similarly, we know that the costs for which focal prices are unconstrained optimal are also identical to the focal prices themselves. Therefore, the Equal Spacing condition holds under perfect competition, for any demand function. In this context, the expected pass-through rate is still 100% even when we introduce focal pricing constraints.

21.3 Differentiated price competition

Having discussed the two extreme cases of monopoly and perfect competition (which are also informative about collusion and undifferentiated price competition) I now turn to an intermediate case: differentiated price competition.

With continuous pricing, it can be shown under general conditions that there exists a unique equilibrium for differentiated price competition (see Mizuno (2003)). This result no longer holds with discrete pricing, and in general there is a multiplicity of equilibria. With indeterminacy of equilibria it is harder to draw conclusions about the impact of focal pricing, since it is conceptually possible for it to lead to higher or lower pass-through. Here, I focus on a standard setting, and propose a simple equilibrium strategy for firms facing focal pricing constraints. I show that in this case the Equal Spacing assumption holds.

Consider a market with N firms producing a differentiated product, and simultaneously competing on prices in a one-shot game. Demand for firm n 's product satisfies standard conditions for differentiated price competition (see Mizuno (2003)). Here, we take it to be:

¹⁰²Results straightforwardly extend for rounding down, or rounding to the nearest focal price.

¹⁰³There is no profitable deviation to a lower focal price because it would entail negative profits, nor any lower non-focal price because it would entail the same firm-specific demand, but at a lower price. There is also no profitable deviation to a higher price because it would entail zero firm-specific demand.

$$q_n = \begin{cases} Q - p_n + p_{-n}^- & Q > p_n - p_{-n}^- \\ 0 & Q \leq p_n - p_{-n}^- \end{cases}$$

where p_{-n}^- is the average price set by other firms in the market, and Q is some positive constant. Take the symmetric input case, in which all firms face the same marginal cost $c_j = c$, $\forall j$ and are all exposed to the same changes in input costs (including the one of interest).

In the unconstrained case, the profit function for firm n is $\pi_n = (p_n - c)(Q - p_n + p_{-n}^-)$. Best responses are linear and symmetric, and there is a unique symmetric equilibrium where all firms charge $p_n = p = Q + c$, $\forall n$. Therefore, this industry is characterised by 100% pass-through, i.e. a Δ increase in the input cost faced by all firms will result in a Δ increase in the price charged by all firms.

Now consider the case of exogenously determined focal pricing, e.g. due to consumer inattention to some digits, spaced out with gaps G . Let there be a focal price point at Q , which we refer to as the low price p^L , a medium focal price at $p^M = Q + G$, and a high focal price at $p^H = Q + 2G$ (we focus on this interval, but the logic generalises to the whole interval of possible prices). We can see that the low price is unconstrained optimal when $c = 0$, the medium price is unconstrained optimal when $c = G$, and the high price is unconstrained optimal when $c = 2G$. Therefore, the spacing between costs at which consecutive focal prices are unconstrained optimal is G .

With focal pricing, the demand function can be written as:

$$q_n = \begin{cases} Q - f_n + f_{-n}^- & Q > f_n + f_{-n}^- \\ 0 & Q \leq f_n + f_{-n}^- \end{cases}$$

where f_n is p_n rounded up to the next focal price, and $f_{-n}^- = \frac{\sum_{j \neq n} f_j}{N-1}$.

In this case, the conditions for uniqueness of equilibrium no longer hold. Focusing just on symmetric equilibria, we can divide the support of the marginal cost into segments for each of which there are two possible symmetric equilibria. For $0 \leq c \leq G$ there is an all-low-price equilibrium, and an all-medium-price equilibrium. For $G \leq c \leq 2G$ there is an all-medium-price equilibrium, and an all-high-price equilibrium. Therefore, we can construct an equilibrium where each firm's strategy is to charge medium prices if $0 \leq c \leq G$ and high prices if $G \leq c \leq 2G$ (we could easily extend this to higher costs and prices). Hence the spacing between cost thresholds is G , which is the same as the spacing between the gaps between the costs at which focal prices are unconstrained optimal. Hence, the Equal Spacing assumption holds, and in the presence of uncertainty (the Uniformity assumption) the Irrelevance Theorem holds.

Let us say that prior to the cost increase, the industry was characterised by medium prices, meaning that marginal costs were in the range $0 \leq c \leq G$. Then if costs increase by Δ , there is a $\frac{\Delta}{G}$ chance of costs increasing to the range $G \leq c \leq 2G$, in which case all firms charge high prices. The expected change in prices is therefore $\frac{\Delta}{G}G = \Delta$, i.e. the industry is still characterised by 100% pass-through, as it was in the absence of focal pricing.

21.4 Other models

There are a multiplicity of modelling assumptions that influence pass-through, and its relation to focal pricing. These include the number of firms, whether they compete in a one-shot game or repeatedly, whether there are asymmetries between firms, the degree of product heterogeneity, what characteristics of the product can be chosen by the firm (price, quantity, quality...), etc. Therefore, it is hard to rule out the possibility that there exist models where the Equal Spacing assumption does not apply, and hence focal pricing reduces, or increases pass-through in expectation. However, it is challenging to find examples where this is the case. For instance, consider relaxing the assumption of symmetric costs in the context of undifferentiated price competition. In this case, with focal pricing constraints, if the increase in input costs was sustained only by a subset of firms, then the affected firms would be forced to exit the market. There would be no direct impact on prices,¹⁰⁴ and hence zero direct price pass-through. However, this is exactly what would happen in the unconstrained setting too, so that the pass-through rate is 0% regardless of the presence of focal pricing constraints.

Even in models with multiplicity of equilibria it is hard to see how the Irrelevance Theorem could not hold at least as a global approximation. For instance, consider the differentiated price competition setting discussed above. Consider an alternative equilibrium, where each firm's strategy is to charge low prices if $0 \leq c \leq G$ and high prices if $G \leq c \leq 2G$. Then the expected pass-through is twice as high as without focal pricing. However, note that we could also construct the opposite example, where firms are in the medium price equilibrium for the whole set of possible costs, and hence there is no pass-through. Similarly to the discussion of different curvatures of demand in section 21.1.2, locally higher (or lower) pass-through rates will in general be 'cancelled out' by other locally lower (or higher) pass-through rates; or at least it cannot be the case that focal pricing systematically increases or decreases pass-through along the cost interval. Therefore, the Irrelevance Theorem still holds as a global approximation when we do not possess detailed knowledge about whether we are in a segment where focal pricing locally increases or decreases pass-through rates. Continuing

¹⁰⁴It is worth noting that the exit of a subset of firms might lead to a large indirect impact on prices through increased market concentration and higher chance of collusion.

the differentiated Bertrand example, if we extended our analysis to further along the cost interval, if we were in an all-medium equilibrium at $G \leq c \leq 2G$, that would not be feasible for $2G \leq c \leq 3G$, where there must be either an all-high equilibrium or an equilibrium with the next focal price up: it is not sustainable for there to not be pass-through along more than a local segment of the cost interval.

The conceptual indeterminacy of the impact of focal pricing on pass-through, combined with the fact that the Irrelevance Theorem holds exactly in many standard settings, and approximately in many others, provides a clear rationale for conducting context-specific empirical analysis to estimate pass-through, rather than assuming that the presence of focal pricing will lead to little, or no, pass-through.

22 Discussion

In the context of antitrust damages class actions, the Irrelevance Theorem provides a strong rationale not to accept arguments that cases should be dismissed because there will be low, or zero, pass-through with focal pricing constraints. However, we should take seriously the possibility that focal pricing may increase heterogeneity in the distribution of pass-through. Depending on the specifics of the contexts, this might mean that the majority of the class has suffered no harm, while a minority has suffered substantial harm. In other cases, all class members suffered damages, perhaps to similar, or perhaps to different, degrees. It is also possible that no class members suffered any damages at all. In the, unusual, case of all consumer having purchased the same product at the same time, and overcharge being very small relative to the gaps between focal prices, it is likely that there was no pass-through. However, there is also a non-negligible possibility that there was pass-through, and that it was very high.

In some jurisdictions, the increased heterogeneity in the distribution of pass-through introduced by focal pricing may be perceived as a challenge to the homogeneity requirement that class members should all have suffered damages in a similar way. However, it is not clear why heterogeneity in damages arising from focal pricing should be treated any differently from heterogeneity arising from class members having purchased slightly different products at different times, as is standardly the case in class actions. In these cases too, class members may well have suffered from different damages to begin with, and then accumulated differential interest on damages over time. It also seems remarkable that abusive firms can avoid paying damages to consumers simply because those consumers were affected to different degrees by the anti-competitive behaviour. From an economic viewpoint, the way to address this issue

is to perform more granular analysis, and obtain estimates of damages specific to members of different sub-classes. Where this kind of detailed expert analysis is considered too expensive by the court, using average estimates for the whole class seems like a more sensible approach than dismissing the case as a whole.

Part IV

Conclusion

This thesis contributes to a large and multi-faceted literature showing how powerful and versatile game theory is. The first and third chapter are ‘classic’ applications of game theory, respectively of cooperative models to household economics and non-cooperative models to industrial organisation settings. The second chapter explores a novel application of cooperative game theory in environmental economics.

The first chapter of this thesis contributes to the literature on estimating inequality at the individual, rather than household level. This literature has the potential to improve the targeting and evaluation of policies aimed at achieving greater equity, including commitments to cut poverty rates in line with the UN’s sustainable development goals. The standard measures that policy-makers currently rely on, such as the World Bank’s poverty rates, assume equal sharing of resources within households. However, there is an increasing body of evidence that this assumption is unrealistic (e.g. Lechene et al. (2022)); a finding which is also supported in my first chapter. This means that standard measures may lead to an inaccurate understanding of levels of poverty (typically underestimates), dimensions of heterogeneity (such as gender), and inaccurate conclusions about trends over time and the impact of specific policies. Approaches such as those developed in the recent household economics literature, and in the first chapter of this thesis, aim to replace standard measures of inequality with more accurate ones. In order for this to happen, two important steps are required.

Firstly, the proposed approaches must be feasible with the kind of data that is available at a large scale, across many countries, at reasonably frequent time intervals. Moreover, estimation must be sufficiently transparent and simple that it can be replicated easily and understood outside of a narrow circle of academic specialists. The literature has made great progress in this direction, in particular with the work of Dunbar et al. (2013) and Lechene et al. (2022), who show that we can estimate intra-household sharing by linear estimation of a small number of equations, with data on just one assignable good. My chapter further contributes in this direction by proposing the use of a novel source of identifying variation: time-use data, which is widely available and feasible to collect at scale. However, the price of simplicity tends to be a loss of generality, and in my first chapter I only focus on working heterosexual couples without cohabiting children. I plan to work on an extension of my methodology that can be applied to more general household compositions, including children and people who do not supply paid market work. Understanding intra-household inequality

for these household compositions is particularly important for policy-makers, as we have reason to think they may be characterised by more inequality than working heterosexual couples without cohabiting children (e.g. see Bargain, Donni and Hentati (2022)).

Secondly, there must be strong evidence that the proposed approaches perform better than standard measures of inequality between households. Because we need to provide estimates of inequality at frequent intervals for many countries, we cannot rely on methods that exploit plausibly exogenous variations in bargaining power within the household, because these are typically available only for a specific time and place. Instead, we must rely reasonably heavily on structural models, and their assumptions. This makes it particularly important to supply clear evidence that the proposed models fit the data well, and the estimates are more accurate than standard ones. However, the literature has a long way to go in this direction. Different approaches don't always yield consistent answers about inequality within households (e.g. see Bargain, Donni and Hentati (2022) and Lise and Seitz (2011)), and in order for policymakers to be able to replace standard approaches with new ones it is essential to systematically evaluate different methods and understand which methods are more successful in practice. In my first chapter, I attempt to test assumptions, including my identifying assumption, as extensively as possible with my data. However, I cannot directly test my identifying assumption, and can only rely on indirect evidence. I also provide high-level comparisons of my approach and some alternatives from the literature. A more detailed comparison was too substantial a project in its own right to include as part of that chapter, but I am working on it as part of a longer-term follow-up project. With my co-authors, we plan to investigate the assumptions underlying different approaches, and compare estimates to direct measures of the objects of interest. To do so, we plan to draw on a combination of novel data and simulations. We hope this will help this literature to move towards something like a cohesive mainstream approach that can be relied upon by policymakers. Understanding what works best in this literature will also be beneficial to inform data collection efforts.

Going forward, the work that has been done in the household literature, mostly with the goal of estimating individual-level inequality, could also be fruitfully taken to a variety of other applications. The second chapter of this thesis exemplifies this, by extending the model in my first chapter to an environmental context, and using it to estimate gender differences in emissions propensities. Understanding emissions patterns and their determinants is a top priority for governments across the world. This chapter suggests that policies aimed at increasing female empowerment within households may have the important added benefit of helping us reach our climate goals. We contribute to the environmental literature by proposing a new, structural, bargaining approach to estimating gender differences in environmental preferences. Our work is the first to attempt in the literature to do this, and therefore suffers

from many of the common limitations of an early paper in a new field, including three main ones.

The first is that our model is too simple for the purpose: the homotheticity assumption was reasonable in my first chapter, applied just to private leisure, but is too unrealistic for the expenditure system as a whole. We try to limit the empirical ramifications of this by splitting the sample into budget categories and treating the functional form of preferences as a local linear approximation. However, it would be best to retain a larger sample size and employ a more complex functional form. This would require careful consideration of which functional form best fits the data while also being consistent with the approach to estimating bargaining weights. The second is that we restrict heterogeneity more than we ideally would. Because of the small sample used to estimate bargaining, we do not allow preference parameters to vary within men and women, to avoid splitting our sample. However, this opens up our analysis to a potential confounder: we may see that female empowerment is associated with fewer emissions because of preference heterogeneity in some dimension which is correlated with bargaining. Thirdly, our emissions data is not sufficiently granular to yield definitive answers to our questions. We use gold-standard existing data, but this is still at a reasonably aggregated level (e.g. meat, rather than local organic chicken vs imported beef), and could hide important variation. We hope future papers will be able to incorporate steps forward on all of these dimensions.

The third chapter is relatively simple, but this is intentional, and should be a strength, rather than a weakness, to achieve its purpose of influencing the decision-making of courts across the world. However, from an academic point of view, it would certainly be of great interest to consider under which general conditions the irrelevance theorem continues to hold, and under which it does not. Illustrations of more realistic and complex models of oligopolistic dynamic competition with imperfect substitutes may provide further insight into whether there are specific industry characteristics under which the use of focal pricing may be taken as reasonable evidence of low pass-through, and others where instead the opposite may be the case. However, the complexity of the problem is substantial, and future work addressing it may find that it is challenging to tackle it in a manner that combines additional sophistication with sufficient generality to be practically relevant to antitrust courts.

In all three chapters, game theory provides a tractable mathematical formalisation of complex real-life problems. This enables us to see exactly what assumptions are being made, where they might not hold, and what findings follow if we agree on those assumptions. This is powerful. As outlined in this thesis, it can help governments work towards important goals to reduce poverty and emissions, and it can help courts make evidence-based decisions without being swayed by clever, but inaccurate, rhetoric.

Part V

References

References

- Agarwal, N., Ashlagi, I., Rees, M. A., Somaini, P., & Waldinger, D. (2021). Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica*, 89(1), 37–76. <https://doi.org/10.3982/ECTA17017>
- Alexandrov, A. (2013). Pass-through rates in the real world: The effect of price points and menu costs. *Antitrust Law Journal*, 79(1), 349–360. <https://www.jstor.org/stable/43486960>
- Almås, I., Ringdal, C., & Sjørusen, I. H. (2020). Understanding inequality within households. In K. F. Zimmermann (Ed.), *Handbook of labor, human resources and population economics* (pp. 1–27). Springer International Publishing. https://doi.org/10.1007/978-3-319-57365-6_222-1
- Apps, P., & Savage, E. (1989). Labour supply, welfare rankings and the measurement of inequality. *Journal of Public Economics*, 39(3), 335–364. [https://doi.org/10.1016/0047-2727\(89\)90033-9](https://doi.org/10.1016/0047-2727(89)90033-9)
- Baraldi, A. L., & Fosco, G. (2025). Clearing the air: Women in politics and air pollution. *Journal of Environmental Economics and Management*, 103106.
- Bargain, O., & Donni, O. (2012). Expenditure on children: A rothbarth-type method consistent with scale economies and parents’ bargaining. *European Economic Review*, 56(4), 792–813. <https://doi.org/10.1016/j.euroecorev.2012.02.003>
- Bargain, O., Donni, O., & Hentati, I. (2022). Resource sharing in households with children: A generalized model and empirical evidence from the UK. *Journal of the European Economic Association*, 20(6), 2468–2496. <https://doi.org/10.1093/jeea/jvac033>
- Bargain, O., Lacroix, G., & Tiberti, L. (2021). Intrahousehold Resource Allocation and Individual Poverty: Assessing Collective Model Predictions using Direct Evidence on Sharing. *The Economic Journal*, 132(643), 865–905. <https://doi.org/10.1093/ej/ueab085>
- Bastian, J., & Lochner, L. (2020, August). *The EITC and maternal time use: More time working and less time with kids?* (Working Paper No. 27717). National Bureau of Economic Research. <https://doi.org/10.3386/w27717>

- Basu, K. (1997). Why are so many goods priced to end in nine? and why this practice hurts the producers. *Economics Letters*, 54(1), 41–44. [https://doi.org/10.1016/S0165-1765\(97\)00009-8](https://doi.org/10.1016/S0165-1765(97)00009-8)
- Beatty, T. K. M., Blow, L., Crossley, T. F., & O’Dea, C. (2014). Cash by any other name? evidence on labeling from the UK winter fuel payment. *Journal of Public Economics*, 118, 86–96. <https://doi.org/10.1016/j.jpubeco.2014.06.007>
- Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal*, 75(299), 493–517. <https://doi.org/10.2307/2228949>
- Blundell, R., Chiappori, P.-A., Magnac, T., & Meghir, C. (2007). Collective labour supply: Heterogeneity and non-participation. *The Review of Economic Studies*, 74(2), 417–445. <http://www.jstor.org/stable/4626146>
- Blundell, R., Chiappori, P.-A., & Meghir, C. (2005). Collective labor supply with children. *Journal of Political Economy*, 113(6), 1277–1306. <https://doi.org/10.1086/491589>
- Brown, C., Calvi, R., & Penglase, J. (2021). Sharing the pie: An analysis of undernutrition and individual consumption in bangladesh. *Journal of Public Economics*, 200, 104460. <https://doi.org/10.1016/j.jpubeco.2021.104460>
- Browning, M., Bourguignon, F., Chiappori, P.-A., & Lechene, V. (1994). Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political Economy*, 102(6), 1067–1096. <https://www.jstor.org/stable/2138780>
- Browning, M., Chiappori, P.-A., & Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, 80(4), 1267–1303. <https://doi.org/10.1093/restud/rdt019>
- Browning, M., Chiappori, P.-A., & Weiss, Y. (2014). *Economics of the family*. Cambridge University Press.
- Browning, M., Donni, O., & Gørtz, M. (2020). Do you have time to take a walk together? Private and joint time within the household. *The Economic Journal*, 131(635), 1051–1080. <https://doi.org/10.1093/ej/ueaa118>
- Browning, M., & Gørtz, M. (2012). Spending time and money within the household. *The Scandinavian Journal of Economics*, 114(3), 681–704. <http://www.jstor.org/stable/41679527>
- Büchs, M., & Schnepf, S. V. (2013). Who emits most? associations between socio-economic factors and uk households’ home energy, transport, indirect and total co2 emissions. *Ecological Economics*, 90, 114–123.
- Bulow, J. I., & Pfleiderer, P. (1983). A note on the effect of cost changes on prices. *Journal of Political Economy*, 91(1), 182–185. <https://www.jstor.org/stable/1840437>

- Bünning, M. (2015). What happens after the ‘daddy months’? fathers’ involvement in paid work, childcare, and housework after taking parental leave in germany. *European Sociological Review*, 31(6), 738–748. <https://doi.org/10.1093/esr/jcv072>
- Calvi, R. (2020). Why are older women missing in india? The age profile of bargaining power and poverty. *Journal of Political Economy*, 128(7), 2453–2501. <https://doi.org/10.1086/706983>
- Calvi, R., Penglase, J., & Tommasi, D. (2022). Measuring women’s empowerment in collective households. *AEA Papers and Proceedings*, 112, 556–560. <https://doi.org/10.1257/pandp.20221054>
- Carlsson Kanyama, A., Nässén, J., & Benders, R. (2021). Shifting expenditure on food, holidays, and furnishings could lower greenhouse gas emissions by almost 40%. *Journal of Industrial Ecology*, 25(6), 1602–1616.
- Cherchye, L., De Rock, B., Lewbel, A., & Vermeulen, F. (2015). Sharing rule identification for general collective consumption models. *Econometrica*, 83(5), 2001–2041. <http://www.jstor.org/stable/43616995>
- Cherchye, L., De Rock, B., & Vermeulen, F. (2012). Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information. *American Economic Review*, 102(7), 3377–3405. <https://doi.org/10.1257/aer.102.7.3377>
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy*, 100(3), 437–467. <https://doi.org/10.1086/261825>
- Chiappori, P.-A., & Ekeland, I. (2009). The microeconomics of efficient group behavior: Identification. *Econometrica*, 77(3), 763–799. <https://doi.org/10.3982/ECTA5929>
- Chiappori, P.-A., Meghir, C., & Okuyama, Y. (2024). Intrahousehold Welfare: Theory and Application to Japanese Data. *NBER Working Paper Series*. <https://doi.org/10.3386/w32645>
- Christis, M., Breemersch, K., Vercalsteren, A., & Dils, E. (2019). A detailed household carbon footprint analysis using expenditure accounts—case of Flanders (Belgium). *Journal of Cleaner Production*, 228, 1167–1175.
- Conlon, C. T., & Rao, N. L. (2020). Discrete prices and the incidence and efficiency of excise taxes. *American Economic Journal: Economic Policy*, 12(4), 111–143. <https://doi.org/10.1257/pol.20160391>
- Coviello, D., Deserranno, E., & Persico, N. (2024). *Gender disparities in the welfare effect of the minimum wage* (Working Paper). https://www.dropbox.com/scl/fi/f0glruckb0gr4dgzs60uz/Draft%5C_MinWWomen.pdf?rlkey=h3lh5cr9zsh0j9r55z7h2oaaw%5C&dl=0

- De Rock, B., & Le Henaff, F. (2023). Walk the talk: Measuring green preferences with social media data. *IFS Working paper*. <https://doi.org/10.1920/wp.ifs.2023.3423>
- Deaton, A., & Muellbauer, J. (1980). An almost ideal demand system. *The American Economic Review*, 70(3), 312–326. <http://www.jstor.org/stable/1805222>
- DEFRA. (2024). UK full dataset 1990 - 2021, including conversion factors by SIC code. <https://www.gov.uk/government/statistics/uks-carbon-footprint>
- DellaVigna, S., & Gentzkow, M. (2019). Uniform pricing in U.S. retail chains. *The Quarterly Journal of Economics*, 134(4), 2011–2084. <https://doi.org/10.1093/qje/qjz019>
- Department for Transport. (2021). National Travel Survey: 2020. <https://www.gov.uk/government/statistics/national-travel-survey-2020/national-travel-survey-2020#trends-in-car-trips>
- Department for Transport. (2022). Car or van occupancy and lone driver rate by trip purpose, England: 2002 onwards. <https://www.gov.uk/government/statistical-data-sets/nts09-vehicle-mileage-and-occupancy>
- Department of Health & Social Care. (2022). Women’s health strategy for england [Policy paper].
- Druckman, A., & Jackson, T. (2016). Understanding households as drivers of carbon emissions. *Taking Stock of Industrial Ecology*, 181–203.
- Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children’s resources in collective households: Identification, estimation, and an application to child poverty in Malawi. *American Economic Review*, 103(1), 438–471. <https://doi.org/10.1257/aer.103.1.438>
- Fong, Y., & Zhang, J. (2001). The identification of unobservable independent and spousal leisure. *Journal of Political Economy*, 109(1), 191–202. <https://doi.org/10.1086/318599>
- Fudenberg, D., & Tirole, J. (1991). *Game theory*. The MIT Press. <https://doi.org/10.4236/jamp.2023.113043>
- Geary, R. C. (1950). A note on "a constant-utility index of the cost of living". *The Review of Economic Studies*, 18(1), 65–66. <http://www.jstor.org/stable/2296107>
- Gitter, S. R., & Barham, B. L. (2008). Women’s power, conditional cash transfers, and schooling in nicaragua. *The World Bank Economic Review*, 22(2), 271–290. <https://doi.org/10.1093/wber/lhn006>
- Goldin, C., Kerr, S. P., & Olivetti, C. (2020). Why firms offer paid parental leave: An exploratory study. *Paid Leave for Caregiving: Issues and Answers*, 66–92.
- Griffith, R. (2022). Obesity, poverty and public policy. *The Economic Journal*, 132(644), 1235–1258. <https://doi.org/10.1093/ej/ueac013>

- Griffith, R., Jin, W., & Lechene, V. (2022). The decline of home-cooked food. *Fiscal Studies*, 43(2), 105–120. <https://doi.org/https://doi.org/10.1111/1475-5890.12298>
- HM Revenue & Customs. (2023). Percentile points from 1 to 99 for total income before and after tax. <https://www.gov.uk/government/statistics/percentile-points-from-1-to-99-for-total-income-before-and-after-tax>
- Hubner, S. (2020). It's complicated: A Non-parametric Test of Preference Stability between Singles and Couples. *Bristol Economics Discussion Papers* 20/735. <https://research-information.bris.ac.uk/en/publications/its-complicated-a-nonparametric-test-of-preference-stability-betw>
- IPCC. (2022). Summary for policymakers [In Press]. In H. O. Pörtner, D. C. Roberts, M. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem & B. Rama (Eds.), *Climate change 2022: Impacts, adaptation, and vulnerability. contribution of working group ii to the sixth assessment report of the intergovernmental panel on climate change* (In Press). Cambridge University Press.
- Ivanova, D., Barrett, J., Wiedenhofer, D., Macura, B., Callaghan, M., & Creutzig, F. (2020). Quantifying the potential for climate change mitigation of consumption options. *Environmental Research Letters*, 15(9), 093001.
- Ivanova, D., Stadler, K., SteenâOlsen, K., Wood, R., Vita, G., Tukker, A., & Hertwich, E. G. (2016). Environmental impact assessment of household consumption. *Journal of Industrial Ecology*, 20.
- Johnsen, J. V., Ku, H., & Salvanes, K. G. (2024). Competition and career advancement. *The Review of Economic Studies*, 91(5), 2954–2980. <https://doi.org/10.1093/restud/rdad099>
- Kalenkoski, C. M., Ribar, D. C., & Stratton, L. S. (2005). Parental child care in single-parent, cohabiting, and married-couple families: Time-diary evidence from the united kingdom. *American Economic Review*, 95(2), 194–198. <https://doi.org/10.1257/000282805774670176>
- Kennedy, E. H., & Kmec, J. (2018). Reinterpreting the gender gap in household pro-environmental behaviour. *Environmental Sociology*, 4(3), 299–310.
- Kilian, L., Owen, A., Newing, A., & Ivanova, D. (2023). Achieving emission reductions without furthering social inequality: Lessons from the 2007 economic crisis and the covid-19 pandemic. *Energy Research & Social Science*, 105, 103286.
- Klemperer, P. (2004). *Auctions: Theory and practice*. Princeton University Press. <https://www.jstor.org/stable/j.ctv3hh4pm>

- Knotek, E. S. (2008). Convenient prices, currency, and nominal rigidity: Theory with evidence from newspaper prices. *Journal of Monetary Economics*, 55(7), 1303–1316. <https://doi.org/10.1016/j.jmoneco.2008.07.009>
- Knotek, E. S. (2011). Convenient prices and price rigidity: Cross-sectional evidence. *The Review of Economics and Statistics*, 93(3), 1076–1086. https://doi.org/10.1162/REST_a_00124
- Lechene, V., Pendakur, K., & Wolf, A. (2022). Ordinary Least Squares estimation of the intrahousehold distribution of expenditure. *Journal of Political Economy*, 130(3), 681–731. <https://doi.org/10.1086/717892>
- Levy, D., Lee, D., Chen, H. (, Kauffman, R. J., & Bergen, M. (2011). Price points and price rigidity. *The Review of Economics and Statistics*, 93(4), 1417–1431. https://doi.org/10.1162/REST_a_00178
- Lewbel, A., & Pendakur, K. (2008). Estimation of collective household models with engel curves. *Journal of Econometrics*, 147(2), 350–358. <https://doi.org/10.1016/j.jeconom.2008.09.012>
- Lise, J., & Seitz, S. (2011). Consumption inequality and intra-household allocations. *The Review of Economic Studies*, 78(1), 328–355. <https://www.jstor.org/stable/23015857>
- Maschler, M., Solan, E., & Zamir, S. (2013). *Game theory*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511794216>
- May, A. M., McGarvey, M. G., Gustafson, C. R., & Mieno, T. (2021). Gender, environmental issues and policy: An examination of the views of male and female economists. *Ecological Economics*, 182, 106877.
- McElroy, M. B., & Horney, M. J. (1981). Nash-bargained household decisions: Toward a generalization of the theory of demand. *International Economic Review*, 22(2), 333–349. <https://doi.org/10.2307/2526280>
- Mizuno, T. (2003). On the existence of a unique price equilibrium for models of product differentiation. *International Journal of Industrial Organization*, 21(6), 761–793. [https://doi.org/10.1016/S0167-7187\(03\)00017-1](https://doi.org/10.1016/S0167-7187(03)00017-1)
- Muellbauer, J. (1976). Community preferences and the representative consumer. *Econometrica*, 44(5), 979–999. <https://doi.org/10.2307/1911540>
- Muthoo, A. (1999). *Bargaining theory with applications*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511607950>
- Nakamura, E., & Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, 77(3), 1192–1230. <https://doi.org/10.1111/j.1467-937X.2009.589.x>

x

- Office for National Statistics. (2019a). Living Costs and Food Survey [data series] 3rd Release. <http://doi.org/10.5255/UKDA-Series-2000028>
- Office for National Statistics. (2019b). United Kingdom Time Use Survey [data series] 2nd Release. <https://doi.org/http://doi.org/10.5255/UKDA-Series-2000054>
- ONS. (2024). Families and households. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/families/datasets/familiesandhouseholds>
- Osorio, P., Tobarra, M.-Á., & Tomás, M. (2024). Are there gender differences in household carbon footprints? evidence from Spain. *Ecological Economics*, 219, 108130.
- Owen, A., & Büchs, M. (2024). Examining changes in household carbon footprints across generations in the UK using decomposition analysis. *Journal of Industrial Ecology*.
- Prendergast, C. (2022). The allocation of food to food banks. *Journal of Political Economy*, 130(8), 1993–2017. <https://doi.org/10.1086/720332>
- Reiss, P. C., & Wolak, F. A. (2007). Chapter 64 Structural Econometric Modeling: Rationales and Examples from Industrial Organization. In J. J. Heckman & E. E. Leamer (Eds.), *Handbook of Econometrics* (pp. 4277–4415, Vol. 6). Elsevier. [https://doi.org/10.1016/S1573-4412\(07\)06064-3](https://doi.org/10.1016/S1573-4412(07)06064-3)
- Remes, O., Wainwright, N., Surtees, P., LaFortune, L., Khaw, K.-T., & Brayne, C. (2017). Sex differences in the association between area deprivation and generalised anxiety disorder: British population study. *BMJ Open*, 7, e013590. <https://doi.org/10.1136/bmjopen-2016-013590>
- Scherr, F. C. (1981). Focal points and pricing behavior: Results in an experimental oligopoly. *Journal of Behavioral Economics*, 10(1), 47–65. [https://doi.org/10.1016/S0090-5720\(81\)80005-8](https://doi.org/10.1016/S0090-5720(81)80005-8)
- Snir, A., Levy, D., & Chen, H. (2017). End of 9-endings, price recall, and price perceptions. *Economics Letters*, 155, 157–163. <https://doi.org/10.1016/j.econlet.2017.04.001>
- Stone, R. (1954). Linear expenditure systems and demand analysis: An application to the pattern of British demand. *The Economic Journal*, 64(255), 511–527. <http://www.jstor.org/stable/2227743>
- Strulov-Shlain, A. (2023). More than a penny’s worth: Left-digit bias and firm pricing. *The Review of Economic Studies*, 90(5), 2612–2645. <https://doi.org/10.1093/restud/rdac082>
- The King’s Fund. (2021). Tackling obesity.
- Thörnqvist, T., & Vardardottir, A. (2014). Bargaining over risk: The impact of decision power on household portfolios.

- Toro, F., Fernández-Vázquez, E., & Serrano, M. (2024). Who brings emissions home? comparing female and male breadwinner households by matching techniques. *Energy Policy*, 190, 114144. <https://doi.org/https://doi.org/10.1016/j.enpol.2024.114144>
- United Nations Development Programme. (2022). Human development report 2021-22. <http://report.hdr.undp.org>
- United States District Court Northern District of California *In re Lithium Ion Batteries Antitrust Litigation*, Case No.: 13-MD-2420 YGR (2018, March).
- United States District Court Northern District of California San Jose Division *In re: Qualcomm Antitrust Litigation, Defendant Qualcomm Incorporated's Opposition to Plaintiff's Motion for Class Action*, Case No. 5:17-md-02773-LHK-NMC (2018, September).
- Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8), 2295–2332. <https://doi.org/10.1257/aer.20120234>
- Waldinger, D. (2021). Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review*, 111(8), 2660–2696. <https://doi.org/10.1257/aer.20190516>
- Xiao, C., & McCright, A. M. (2015). Gender differences in environmental concern: Revisiting the institutional trust hypothesis in the usa. *Environment and Behavior*, 47(1), 17–37.
- Young, P. (2001). *Individual strategy and social structure*. Princeton University Press.

Part VI

Appendices to all chapters

A Identification in a more general setting

It is possible to extend the approach in this paper to a more general setting. This appendix provides a discussion of how we might attempt identification and linear estimation under different functional form assumptions.

A.1 Assignable good

The assignable good $c_{i,h}^a$ may be clothing, as in most of the existing, private leisure, as in this paper, or some other assignable good. See online appendix [D](#) for a more detailed discussion

of the considerations involved in choosing the assignable good.¹⁰⁵

A.2 Linear approximation of resource shares

Resource shares will generally depend on household budget, prices (including wages of all household members) and Pareto weights of household members: $\eta_{i,h}(y_h, p, r, w_h, \mu_h)$. The form of an individual's resource share will depend on their type and household category so we can write $\eta_{t,g}(y_h, p, r, w_h, \mu_h)$.¹⁰⁶ Recall that Pareto weights are an unknown function of a vector of variables z_h that determine relative bargaining power in the household. Since z_h in general includes y_h, p, r, w_h we can write $\eta_{t,g}(z_h)$. The collective model only restricts bargaining to be efficient, so to estimate resource shares we must approximate the Pareto weights, or directly approximate the resource shares. For different household categories g , resource shares (and Pareto weights) will be different functions of household characteristics z_h . For this reason, the sharing rule $\eta_g(z_h)$, which assigns a resource share $\eta_{i,h}$ to each member of households h of category g , is estimated separately for households of different categories.¹⁰⁷

To linearly estimate resource shares, we linearly approximate them as:

$$\eta_{i,h} = \eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})$$

- $\eta_{t,g}^0$ is the average resource share of a type (living in a specific household category). This is the resource share evaluated at the average characteristics in the sample. By definition, $\sum_{t \in h} \eta_{t,g}^0 = 1$. In the context of heterosexual couples, the average resource share of men and women sum to one.
- $\eta_{t,g}^z$ captures the impact on sharing of a household's characteristic z_h deviating from the sample average ($z_h - \bar{z}$). For instance, a higher-than-average wage for the woman

¹⁰⁵In practice, estimation is likely to proceed from individual-level demand for the same assignable good for everyone in the household. For this reason, the notation in the identification proof implies that the same assignable good c^a is observed for all members. However, it is worth noting that identification can also proceed from different assignable goods for different types of people, as long as the same assignable good is used for each person of the same type in the same household category. In this case, the identification result is unchanged, apart from indexing the assignable good for different members differently depending on their type.

¹⁰⁶Where there are multiple individuals of the same type in a household, the notation must be amended to acknowledge the fact that an individual's share depends both on their type and on their own individual characteristics / Pareto weight. Everything goes through in the same way, so for simplicity of notation I avoid this extension of notation.

¹⁰⁷The approximation is more effective if person types and household categories are defined sufficiently granularly that the bargaining process, as a function of characteristics, would not differ too substantially within each household category sub-sample.

might increase the woman's resource share, so that she would have a higher-than-average-for-women resource share. Since resource shares must sum to one within the household, this implies her partner must have a correspondingly lower-than-average-for-men resource share: $\sum_{t \in h} \eta_{t,g}^z = 0$. We can interpret $\eta_{t,g}^z$ as the marginal impact of characteristic z_h on the resource share.

A linear approximation does not guarantee estimates of resource shares which fall within the unit interval. However, by construction, only estimates in this range are theory-consistent. This provides a useful test of model fit. Reassuringly, in my application to UK data, my baseline resource share estimates from time-use data are all within the unit interval.

A.3 Parametric forms and estimating equations

The identification result holds for several families of preferences. Once we have chosen a parametric form for implementation we can derive the functional form of second-stage demands. We may directly use assignable good demands as estimating equations, or else choose some function f of demands: $f(c_{t,g}^a(p, w_{i,h}, Q_h, (y_h - R_h Q_h) \eta_{i,h}))$. This could be for instance: demand for the assignable good, expenditure on the assignable good, the budget share spent on the assignable good, or expenditure on the assignable good relative to the expenditure on some sub-set of goods.¹⁰⁸ We may choose the most appropriate function depending on the functional form choices made for implementation, or depending on the data available.

A.4 Identification

Previous literature has shown that resource shares are generally identified up to a constant in a context with the underlying model discussed in section 7 and the sort of data set out in section 8.6. This paper is concerned with a somewhat different goal: point-identifying the linear approximation of the sharing rule from coefficients estimated by linear regression.

The linear approximation of the sharing rule is semi-parametrically point-identified under one of three possible identifying assumptions restricting preference heterogeneity. In short, identification proceeds in two steps. First, a composite coefficient is identified by linear regression. This step relies on standard assumptions for linear regression. Then, the sharing rule parameters are identified from the composite coefficient by imposing an identifying

¹⁰⁸Different assignable goods or functions of demand could be used for different types of people, but the same must be used for people of the same type in the same household category. The generalised notation is $f_{t,g}(c_{t,g}^{a,t,g}(p, w_{i,h}, Q_h, (y_h - R_h Q_h) \eta_{i,h}))$ and the same result applies with minimal changes.

assumption restricting preference heterogeneity. In section A.4.1 I set out the result and the conditions required for identification to hold. In section A.4.2 I discuss the three candidate identifying options. In section A.4.3 I prove the identification result with an argument by construction. In section A.4.4 I discuss the intuition behind the identification result. In section A.5, I exemplify the broader identification result with the Almost Ideal Demand System.

A.4.1 Identification result

Theorem A.1. *The linear approximation of the sharing rule $\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})$ is fully identified for households of category g if (i) for all N_g^t household members we observe a function $f(c_{i,h}^a) = f(c_{t,g}^a(p, w_{i,h}, Q_h, (y_h - R_h Q_h) \eta_{i,h}))$ of demand for assignable good $c_{i,h}^a$, and (ii) the following Assumptions hold.*

Assumption A.1.1. Linearity in resource shares.

$f(c_{i,h}^a)$ can be decomposed into a component $f_{t,g}^0 \eta_{i,h}$ which is linear in the resource share:

$$f(c_{i,h}^a) = \sum_{\varphi} f_{t,g}^{\varphi}(p, w_{i,h}, Q_h, (y_h - R_h Q_h) \eta_{i,h}) + f_{t,g}^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h)) \eta_{i,h} \quad (6)$$

This assumption imposes that the assignable good demand functions (or some function of demand) have a *component* which is linear in the individual's second-stage budget. Roy's identity can be used to easily check whether a specific indirect utility function is consistent or not with this requirement. Some functional forms that are consistent with this assumption are: (i) the Linear Expenditure System (Cobb-Douglas or Stone-Geary, see Stone (1954) and Geary (1950)), (ii) Price-Independent Generalised Linear (PIGL) (indirect utility functions that are a function of $\rho_{i,h}^k, k > 0$, see Muellbauer (1976)) and (iii) Price-Independent Generalised Logarithmic (PIGLOG) (indirect utility functions that are a function of $\ln \rho_{i,h}$ e.g. the Almost Ideal Demand System, see Muellbauer (1976) and Deaton and Muellbauer (1980)).¹⁰⁹ This assumption relaxes similar assumptions made in the related literature, allowing for broader families of preferences (e.g. DLP assumes budget share functions are linear in functions of expenditure, and LPW assumes that Engel curves be of the Almost Ideal Demand System, both of which are examples of functions of demand that have a component which is linear in the resource share).

¹⁰⁹Weaker parametric assumptions are also consistent with this assumption. For instance, preferences may be represented as Cobb-Douglas over the assignable good, an aggregate private and an aggregate public good (instead of Cobb-Douglas over each granular good). Alternatively, demand for the assignable good may be modelled as being linear in the budget without restricting the functional form of demands for other goods, and only requiring them to be jointly consistent with rational preferences.

Assumption A.1.2. Linearity in unknowns.

Part A. The vectors of unknown resource share parameters $\eta_{t,g}$ and preference parameters $\alpha_{t,g}$ ¹¹⁰ enter $f(c_{i,h}^a)$ linearly in the sense that:

$$f(c_{i,h}^a) = a_{t,g} + \sum_{\varphi} b_{t,g}^{\varphi} g^{\varphi}(p, w_{i,h}, Q_h, (y_h - R_h Q_h), (y_h - R_h Q_h)(z_h - \bar{z})) \\ + \gamma_{t,g} \left(\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z}) \right) g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h)) \quad (7)$$

Where: (i) $a_{t,g}(\alpha_{t,g}, \eta_{t,g})$ and $b_{t,g}^{\varphi}(\alpha_{t,g}, \eta_{t,g})$ are unknown constants which may be (potentially composite) preference parameters, or composite parameters composed of preference parameters and unknown resource share parameters, (ii) $\gamma_{t,g}(\alpha_{t,g})$ is a potentially composite preference parameter and (iii) g^{φ} and g^0 are observed in the data.¹¹¹

Part B. The functions $a_{t,g}(\alpha_{t,g}, \eta_{t,g})$ and $b_{t,g}^{\varphi}(\alpha_{t,g}, \eta_{t,g})$ are linear in the vector of resource share parameters $\eta_{t,g}$.

The assumption that unknowns enter linearly is required so that the estimating equations can be estimated by linear regression.

If the resource share is modelled as a type-specific constant $\eta_{t,g}^0$ (the vector of characteristics affecting sharing, z_h , is zero) then Part A of this assumption is sufficient. Then, we have:

$$f(c_{i,h}^a) = a_{t,g}(\alpha_{t,g}, \eta_{t,g}^0) + \sum_{\varphi} b_{t,g}^{\varphi}(\alpha_{t,g}, \eta_{t,g}^0) g^{\varphi}(p, w_{i,h}, Q_h, (y_h - R_h Q_h)) \\ + \gamma_{t,g} \eta_{t,g}^0 g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h))$$

In this case, the identification result holds exactly for broader families of preferences, such as the Almost Ideal Demand System.

If instead the resource share depends on a non-zero vector of characteristics, so that resource shares vary at the household level, we additionally require Part B of this assumption. In

¹¹⁰If modelling shareable goods, then read ‘preference or shareability parameters’ for ‘preference parameters’

¹¹¹Note that, as long as we assume the same parametric form for preferences of all individuals (of all types, and in all household categories) these fully observable objects have the same functional form for everyone. Identification also works if we assume different parametric forms for preferences of different types of individuals, or individuals in different household categories, and in that case we write $g_{t,g}^{\varphi}$ and $g_{t,g}^0$. Here I focus on the case where the same parametric form of preferences is assumed for all individuals as this is generally the case for implementation, and it leads to much more natural and realistic interpretations of the identifying assumptions.

this case, the identification result holds exactly for narrower families of preferences, such as the Linear Expenditure System. Where Part B does not hold, it is possible to relax it by linearising whichever component(s) of $f(c_{i,h}^a)$ is problematic. Similarly to LPW, the identification result holds approximately for the Almost Ideal Demand System with covariates (see online appendix A.5).

Assumption A.1.3. Rank condition.

No component of $g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h))$ or $(z_h - \bar{z}) g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h))$ is co-linear with any other component of $f(c_{i,h}^a)$

Substituting $\beta_{t,g}^0 = \gamma_{t,g} \eta_{t,g}^0$ and $\beta_{t,g}^z = \gamma_{t,g} \eta_{t,g}^z$ into equation (7), we obtain the estimating equations:

$$f(c_{i,h}^a) = a_{t,g} + \sum_{\varphi} b_{t,g}^{\varphi} g^{\varphi}(p, w_{i,h}, Q_h, (y_h - R_h Q_h), (y_h - R_h Q_h)(z_h - \bar{z})) + \beta_{t,g}^0 g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h)) + \sum_z \beta_{t,g}^z (z_h - \bar{z}) g^0(p, w_{i,h}, Q_h, (y_h - R_h Q_h)) \quad (8)$$

We can see that the rank condition is required so that coefficients $\beta_{t,g}^0$ and $\beta_{t,g}^z$ can be identified separately from other coefficients in $f(c_{i,h}^a)$. This is necessary because we will identify $\eta_{t,g}^0$ and $\eta_{t,g}^z$ from $\beta_{t,g}^0$ and $\beta_{t,g}^z$. Depending on the functional form of choice, this assumption can imply restrictions on which variables z_h influence the resource share. A related, but less general, assumption which is often made in the literature is that resource shares are invariant to expenditure (see DLP for a discussion of why this assumption is reasonably non-restrictive both theoretically and empirically).

Assumption A.1.4. Invertibility condition.

$$\gamma_{t,g} \neq 0, \forall t, g$$

Depending on the chosen functional form specification, parts of the general form of $f(c_{i,h}^a)$ may be absent, e.g. there may be no constant term $a_{t,g}$. However, the component which is linear in the resource share must be present for identification to work. The interpretation of this assumption depends on the functional form used in implementation. With Cobb-Douglas preferences, we can think of this as a requirement that the assignable good does not command a very small budget share. If using the Almost Ideal Demand System, then this assumption is equivalent to the requirement in LPW that the slope of the Engel curves for the assignable good isn't too close to zero.

Assumption A.1.5. Regression.

Any noise in the observed data is uncorrelated with the right-hand-side variables. The variables have non-zero variance.

In order to run linear regressions we must introduce noise to our model, for example as outlined in section 8.5.1. In order to identify the coefficients of interest by linear regression, this noise must be uncorrelated with right-hand-side variables. Similarly, non-zero-variance is required in order to identify from linear regression.

Assumption A.1.6. *SAP, SAT or SRAT.*

A suitable restriction on $\gamma_{t,g}$ heterogeneity applies.

The candidate preference heterogeneity restrictions are discussed in section A.4.2, and are in line with the identifying restrictions made in the literature.

A.4.2 Identifying assumption

Identification of the sharing rule from individual-level demand data for a single assignable good requires a strong identifying assumption. It is important to note that the chosen assumption has an important impact on estimates, and should be chosen carefully. Assumptions could be made on the bargaining process e.g. Lise and Seitz (2011) assume that women and men with the same potential earnings divide full income equally. Other papers have focused on restricting preference heterogeneity. For example, Lewbel and Pendakur (2008) and Bargain and Donni (2012) rely on the ‘SAT’ assumption (similarity across household types), that preferences for the assignable good are stable across household compositions, allowing for identification of the preference parameter from singles data. However, we may be concerned by findings that, at least in some contexts, preference stability across household composition is rejected empirically (see Hubner (2020)). DLP and LPW use the alternative ‘SAP’ assumption (similarity across people) that women and men in the same household composition have the same preference parameter for the assignable good. Whether this is a good assumption depends on whether we have reason to suspect substantial gender differences for the assignable good. For example, as discussed in section 9.4, singles data suggest that women have much stronger preferences than men for clothing. In this case, we may choose to rely instead on a novel identifying assumption: SRAT.

Assignable good SRAT (similarity of ratios across types): the ratio of the preference parameter $\gamma_{t,g}$ between different types of people t remains stable across household compositions g : $\frac{\gamma_{t,g}}{\gamma_{st,g}} = \frac{\gamma_{t,single}}{\gamma_{st,single}}, \forall t, st \in h$. This assumption implies that preferences on the assignable good change in a similar way for different types of people in the transition from singlehood

to other household categories. This means that the ratio of men's to women's preference parameters is identified from singles data. The SRAT assumption allows preference levels to vary between different people and different household categories, while requiring that the ratio of the parameters is constant across household compositions. For instance, single women's clothing preference parameter may be $\alpha_{w,single}^c = 0.2$ and men's $\alpha_{m,single}^c = 0.1$, while in a working couple it may be that $\alpha_{w,couple}^c = 0.16$ and $\alpha_{m,couple}^c = 0.08$. Note that the rest of the utility function is allowed to vary differently

While SRAT is not an innocuous assumption, it does not seem extraordinarily restrictive for working couples without cohabiting children. There is no particular reason to suspect women's and men's preferences change in a different manner when transitioning from living alone to living in a couple. The assumption is stronger for households with children, since the idea that single men and women's preferences change similarly when they have children is less realistic.

A.4.3 Identification proof

Proof. Firstly, by assumptions 3.1.1 - 3.1.5 and the usual arguments for identification from linear regression, the constants $a_{t,g}$ and $\beta_{t,g}^0$, and the vectors $b_{t,g}^\varphi$ and $\beta_{t,g}^z$ can be estimated by OLS (one for each person type and household category combination) and hence are identified. Secondly, resource share parameters $\eta_{t,g}$ are identified from the β coefficients by assumption 3.1.6. Intuitively, the resource share parameters $\eta_{t,g}$ enter into the β coefficients multiplicatively with the (composite) preference parameter $\gamma_{t,g}$, so we require an identifying assumption to separately identify the resource share parameters. The proof is outlined separately for each of the three candidate identifying assumptions:

- Under SAT, the preference parameter $\gamma_{t,g}$ is identified from singles data, so that $\gamma_{t,g} = \gamma_{t,single} = \gamma_t$ is known. Therefore, we compute: $\eta_{t,g}^0 = \frac{\beta_{t,g}^0}{\gamma_t}$ and $\eta_{t,g}^z = \frac{\beta_{t,g}^z}{\gamma_t}$.
- Under SAP, the average resource share of each type is identified by $\eta_{t,g}^0 = \frac{\beta_{t,g}^0}{\sum_{st \in h} \beta_{st,g}^0}$. This is because under SAP $\frac{\beta_{t,g}^0}{\sum_{st \in h} \beta_{st,g}^0} = \frac{\gamma_g \eta_{t,g}^0}{\sum_{st \in h} \gamma_{st,g} \eta_{st,g}^0} = \frac{\eta_{t,g}^0}{\sum_{st \in h} \eta_{st,g}^0} = \eta_{t,g}^0$ since, by definition, resource shares sum to one within household. The marginal impact of characteristic z is identified by $\eta_{t,g}^z = \frac{\beta_{t,g}^z}{\beta_{t,g}^0} \eta_{t,g}^0$. This is because under SAP $\frac{\beta_{t,g}^z}{\beta_{t,g}^0} \eta_{t,g}^0 = \frac{\gamma_g \eta_{t,g}^z}{\gamma_g \eta_{t,g}^0} \eta_{t,g}^0 = \eta_{t,g}^z$.
- Under SRAT, the proof is similar to SAP with the difference that we identify $\frac{\gamma_{t,g}}{\gamma_{st,g}} = \frac{\gamma_{t,single}}{\gamma_{st,single}} = \Lambda_{t,st}$ from singles data and then identify the average resource share of each

type as $\eta_{t,g}^0 = \frac{\beta_{t,g}^0}{\sum_{st \in h} \Lambda_{st,g} \beta_{st,g}^0}$. The marginal impact of characteristic z is identified by $\frac{\beta_{t,g}^z}{\beta_{t,g}^0} \eta_{t,g}^0 = \frac{\gamma_{t,g} \eta_{t,g}^z}{\gamma_{t,g} \eta_{t,g}^0} \eta_{t,g}^0 = \eta_{t,g}^z$.

□

A.4.4 Discussion of the identification result

This identification result builds on the identification result in DLP and the estimation framework of LPW, extending them to a more general underlying structural model.

The extension to public goods¹¹² leads to a challenge for identification because recovering the sharing rule requires either data on public expenditure (both material and time-use) or an additional preference restriction. I discuss the latter in section 8, where I discuss an implementation with time-use data only.

I also note that the extension to contexts with price-variation (including individual-level wages) can complicate estimation. Firstly, it can greatly increase the number of regressors in the estimating equations, depending on how many prices are heterogenous within the sample and on the functional form of preferences. Moreover, additional data (on non-constant prices and wages) is required for estimation. Furthermore, where there is poor data on non-labour income, the wages of household members become colinear with the household budget, which can complicate or impede estimation under some functional form assumptions. However, allowing for price variation is necessary both to incorporating time-use in the underlying model and to enabling using private leisure as the assignable good. Moreover, estimation with price variation is likely to yield more accurate estimates because it provides additional sources of identifying variation. Finally, as this identification result does not hinge on price variation, the result in this paper includes as a special case a model where all prices are constant.

In practice, I suggest that this identification result be implemented with Cobb-Douglas preferences where these don't strong-arm the data. They have substantial strengths in this context, as discussed in section 8. I also illustrate the identification result with the Almost Ideal Demand System in online appendix A.5. The intuition behind this general result is discussed for a specific implementation in section 8.

¹¹²While all goods are non-public in the LPW setting, economies of scale from household size are modelled by letting the actual consumption vector of each member be equal to the vector of market purchases for that member multiplied by a matrix A . While I do not explicitly incorporate the sharing framework, the identification result in this section readily applies to a model with both public and shareable goods. The most notable effect of this inclusion, if there are cross-good economies of scale, is that estimating equations are likely to take on more complex functional forms and be less tractable.

A.5 Almost Ideal Demand System

Consider preferences which are separable in public and private goods, and where the private good sub-utility corresponds to the Almost Ideal Demand System (Deaton and Muellbauer (1980)).¹¹³ Then Engel curves for the assignable good take the form:

$$\frac{p_{i,h}^a c_{i,h}^a}{\rho_{i,h}} = \alpha_{t,g} + \sum_j \gamma_{t,g}^j \ln p_{i,h}^j + \beta_t \ln \rho_{i,h} - \beta_{t,g} \left(\alpha_{t,g}^0 + \sum_j \alpha_{t,g}^j \ln p_{i,h}^j + \frac{1}{2} \sum_j \sum_k \gamma_{kj} \ln p_{i,h}^k \ln p_{i,h}^j \right)$$

Here the notation for material private goods is used to also include private leisure, so that $p_{i,h}^l = w_{i,h}$. Hence, the i subscript (some other material good prices may vary at the household level). Similarly, the prices of public goods are indexed by h : R_h .

We re-write the Engel curves in terms of observables and the object of interest, remembering that $\rho_{i,h} = \eta_{i,h}(y_h - R_h Q_h)$ and multiplying through by the resource share:

$$\begin{aligned} \frac{p_{i,h}^a c_{i,h}^a}{y_h - R_h Q_h} &= \eta_{i,h} (\alpha_{t,g} - \beta_{t,g} \alpha_{t,g}^0) + \eta_{i,h} \sum_j ((\gamma_{t,g}^j - \beta_{t,g} \alpha_{t,g}^j) \ln p_{i,h}^j) \\ &\quad - \eta_{i,h} \beta_{t,g} \frac{1}{2} \sum_j \sum_k \gamma_{kj} \ln p_{i,h}^k \ln p_{i,h}^j + \eta_{i,h} \beta_{t,g} \ln \eta_{i,h} + \eta_{i,h} \beta_{t,g} \ln (y_h - R_h Q_h) \end{aligned}$$

If we model the resource share as a type-specific constant (rather than varying at the household-level based on covariates z_h) then the identification result holds exactly. However, if we are interested in household-specific resource shares varying with z_h we need to linearly approximate the $\eta_{i,h} \beta_{t,g} \ln \eta_{i,h}$ term (similarly to LPW):

- As required by assumption 3.1.1, the estimating equation has a component which is linear in the resource share. We can write $\frac{p_{i,h}^a c_{i,h}^a}{y_h - R_h Q_h} = \sum_{\varphi} f_{t,g}^{\varphi} + f_{t,g}^0 \eta_{i,h}$, where $f_{t,g}^0 = \beta_{t,g} \ln (y_h - R_h Q_h)$ and

$$\begin{aligned} \sum_{\varphi} f_{t,g}^{\varphi} &= \eta_{i,h} (\alpha_{t,g} - \beta_{t,g} \alpha_{t,g}^0) + \eta_{i,h} \sum_j ((\gamma_{t,g}^j - \beta_{t,g} \alpha_{t,g}^j) \ln p_{i,h}^j) \\ &\quad - \eta_{i,h} \beta_{t,g} \frac{1}{2} \sum_j \sum_k \gamma_{kj} \ln p_{i,h}^k \ln p_{i,h}^j + \eta_{i,h} \beta_{t,g} \ln \eta_{i,h} \end{aligned}$$

¹¹³This is an extension of the set-up in LPW to public goods. If also wishing to incorporate the shareable goods framework, this can be done as in LPW, with their assumption that A is block-diagonal so that there are no cross-good complementarities relating to the assignable good $A_g = \begin{pmatrix} A^X & 0 \\ 0 & A^c \end{pmatrix}$. This assumption ensures that the assignable good Engel curves have the usual form, instead of the more complex form which would result from cross-good complementarities with other goods.

- As required by part A of assumption 3.1.2, the estimating equation is linear in unknowns if we consider resource shares to be type-specific constants. We can write $\frac{p_{i,h}^a c_{i,h}^a}{y_h - R_h Q_h} = a_{t,g} + \sum_{\varphi} b_{t,g}^{\varphi} g^{\varphi} + \gamma_{t,g} \eta_{t,g}^0 g^0$ where (with some slight abuse of notation):

- † $a_{t,g} = \eta_{t,g}^0 (\alpha_{t,g} - \beta_{t,g} \alpha_{t,g}^0) + \eta_{t,g}^0 \beta_{t,g} \ln \eta_{t,g}^0$ is a constant which is a function of unknown preference and resource share parameters
- † $\sum_{\varphi} b_{t,g}^{\varphi} g^{\varphi} = \eta_{t,g}^0 \sum_j ((\gamma_{t,g}^j - \beta_{t,g} \alpha_{t,g}^j) \ln p_{i,h}^j) - \eta_{t,g}^0 \beta_{t,g} \frac{1}{2} \sum_j \sum_k \gamma_{kj} \ln p_{i,h}^k \ln p_{i,h}^j$
 - * $b_{t,g}^1 = \eta_{t,g}^0 \sum_j ((\gamma_{t,g}^j - \beta_{t,g} \alpha_{t,g}^j))$ is a constant which is a function of unknown preference and resource share parameters
 - * $b_{t,g}^2 = \eta_{t,g}^0 \beta_{t,g} \frac{1}{2} \sum_j \sum_k \gamma_{kj}$ is a constant which is a function of unknown preference and resource share parameters
 - * $g^1 = \sum_j (\ln p_{i,h}^j)$ is observed in the data (price data is required for this)
 - * $g^2 = \sum_j \sum_k \ln p_{i,h}^k \ln p_{i,h}^j$ is observed in the data (price data is required for this)
 - * $g^0 = \ln(y_h - R_h Q_h)$ is observed in the data (requires expenditure on all public goods including both time-use and material)
 - * $\gamma_{t,g} = \beta_{t,g}$ and hence is a preference parameter

- However, part B of the assumption is not satisfied because of the non-linear $\eta_{i,h} \beta_{t,g} \ln \eta_{i,h}$ term.

- † We can linearly approximate $\beta_{t,g} \ln \eta_{i,h}$ as $\kappa_{t,g}^0 + \sum_{\omega} \kappa_{t,g}^{\omega} \omega_h$, where ω_h could coincide with the characteristics z_h used in the resource share approximation, or be transformations of z_h . Substituting this in and rearranging we obtain:

$$\begin{aligned} \frac{p_{i,h}^a c_{i,h}^a}{y_h - R_h Q_h} &= \eta_{i,h} (\alpha_{t,g} - \beta_{t,g} \alpha_{t,g}^0 + \kappa_{t,g}^0) + \eta_{i,h} \sum_j ((\gamma_{t,g}^j - \beta_{t,g} \alpha_{t,g}^j) \ln p_{i,h}^j) \\ &\quad - \eta_{i,h} \beta_{t,g} \frac{1}{2} \sum_j \sum_k \gamma_{kj} \ln p_{i,h}^k \ln p_{i,h}^j + \eta_{i,h} \left(\sum_{\omega} \kappa_{t,g}^{\omega} \omega_h \right) + \eta_{i,h} \beta_{t,g} \ln(y_h - R_h Q_h) \end{aligned}$$

- The rank condition (assumption 3.1.3) is satisfied if

- † We are interested in estimating resource shares as type-specific constants. In this case we just need to check that $\ln(y_h - R_h Q_h)$ is not colinear with the other observables i.e. $g^1 = \sum_j (\ln p_{i,h}^j)$ and $g^2 = \sum_j \sum_k \ln p_{i,h}^k \ln p_{i,h}^j$, which it is not.
- † We are interested in estimating household-specific resource shares (with covariates), and $(y_h - R_h Q_h)$ is not an element of ω_h . A related assumption in LPW is that resource shares don't depend on the household budget.

- The invertibility assumption 3.1.4 requires that assignable good Engel curves are sufficiently responsive to household budget net of public expenditure. This has to be assessed in a specific context for application. Assumptions 3.1.5 and 3.1.6 should also be examined in a specific context for application. The identifying assumption restricting preference heterogeneity (SAP, SAT or SRAT) is made on parameter $\beta_{t,g}$ and hence is to be interpreted with regard to the responsiveness of assignable good Engel curves to household budget net of public expenditure.

The Almost Ideal Demand System Engel curves are not very tractable if all prices vary in the sample. This is both because it becomes highly multi-dimensional and because estimating it requires high-quality, detailed price data. As a result, DLP and LPW assume no price variation and drop the price index terms into the constant. This was possible because time-use was not modelled in those papers. In a model with time-use, there will be some price variation in the form of individual-level wages, even if we can assume all other prices to be constant in the sample.¹¹⁴ If this is the case, we can simplify the estimating equations to:

$$\frac{p_{i,h}^a c_{i,h}^a}{y_h - R_h Q_h} = a_{t,g}^0 \eta_{i,h} + a_{t,g}^1 \eta_{i,h} \ln w_i + a_{t,g}^2 \eta_{i,h} (\ln w_i)^2 + \sum_{\omega} a_{t,g}^{\omega} \omega_h \eta_{i,h} + \eta_{i,h} \beta_{t,g} \ln (y_h - R_h Q_h)$$

Even so, as soon as we add covariates affecting intra-household bargaining this exponentially increases the number of regressors (because of the interaction between ω and η , which both depend on these covariates). Without yet more approximation, this is not a tractable specification, especially for estimation with small samples. This is problematic since small samples are almost inevitable in this literature since estimation must be conducted separately for households of different categories. Yet, more approximation is also unpalatable, as it undermines the link between the structural model and the estimated regression. A more restrictive, but more parsimonious functional specification for preferences may be a preferable approach.

B Time-use data

Activity list from UKTUS 2014

Unspecified personal care

Visiting and receiving visitors

Correspondence

Unspecified organisational work

¹¹⁴To avoid dealing with this, if the assignable good is material, we could assume preferences that are separable in material goods and time-use.

Activity list from UKTUS 2014

Other specified ball games
Gymnastics
Visiting a wildlife site
Celebrations
Work for an organisation
Fitness
Unspecified TV video or DVD watching
Watching a film on TV
Telephone conversation
Unspecified childcare as help to other households
Other specified or unspecified arts and hobbies
Volunteer work through an organisation
Unspecified water sports
Other specified TV watching
Unspecified food management
Watching a film on video
Watching sport on video
Other specified water sports
Other specified video watching
Food preparation and baking
Disposal of waste
Other specified physical exercise
Physical care and supervision of own child as help to other household
Study: Unspecified activities related to school or university
Unspecified productive exercise
Dish washing
Study: Classes and lectures
Information searching on the internet
Hunting and fishing
Preserving
Other specified social life
Other specified information by computing
Study: Homework
Communication on the internet
Main job: unspecified main job

Activity list from UKTUS 2014

Unspecified entertainment and culture
Other specified communication by computing
Unspecified other computing
Skype or other video call
Travel escorting to/ from education
Main job: Working time in main job
Cinema
Other specified computing
Second job: Working time in second job
Travel related to household care
Main job: Coffee and other breaks in main job
Unspecified theatre or concerts
Plays musicals or pantomimes
Opera operetta or light opera
Concerts or other performances of classical music
Unspecified informal help to other households
Dance performances
Unspecified listening to radio and music
Other specified theatre or concerts
Sleep
Sleep: In bed not asleep
Food management as help to other households
Other specified food management
Listening to music on the radio
Sleep: Sick in bed
Borrowing books records audiotapes videotapes CDs VDs etc. from a library
Reference to books and other library materials within a library
Using internet in the library
Household upkeep as help to other households
Reading newspapers in a library
Other specified radio listening
Unspecified household upkeep
Other specified library activities
Sports events
Unspecified games

Activity list from UKTUS 2014

Gardening and pet care as help to other households
Cleaning dwelling
Travel related to voluntary work and meetings
Study: other specified activities related to school or university
Construction and repairs as help to other households
Other specified productive exercise
Cleaning yard
Unspecified games and play with others
Billiards pool snooker or petanque
Shopping and services as help to other households
Heating and water
Other specified parlour games and play
Free time study
Swimming
Travel escorting a child other than education
Unspecified sports related activities
Activities related to sports
Arranging household goods and materials
Unspecified video watching
Other unspecified entertainment and culture
Visiting a historical site
Gambling
Visiting a botanical site
Visiting a leisure park
Physical care and supervision of child as help to other household
Teaching non-coresident child
Reading playing & talking to non-coresident child
Accompanying non-coresident child
Other or unspecified entertainment or culture
Reading playing & talking to own non-coresident child
Accompanying own non-coresident child
Other specified childcare as help to other household
Travel related to organisational work
Physical care and supervision of an adult as help to another household
Accompanying an adult as help to another household

Activity list from UKTUS 2014

Other specified help to an adult member of another household
Unspecified help to an adult of another household
Resting - Time out
Other specified informal help to another household
Other specified informal help
Second job: Coffee and other breaks in second job
Outdoor team games
Travel related to informal help to other households
Unspecified participatory activities
Eating
Meetings
Other or unspecified household upkeep
Other specified games
Religious activities
Unspecified making and care for textiles
Laundry
Ironing
Handicraft and producing textiles
Travel related to religious activities
Activities related to employment: Unspecified activities related to employment
Travel to visit friends/relatives in their homes not respondents household
Activities related to employment: Lunch break
Travel related to other social activities
Other personal care: Unspecified other personal care
Visiting an urban park playground designated play area
Travel related to entertainment and culture
Other personal care: Wash and dress
Computing - programming
Other specified making and care for textiles
Travel related to participatory activities other than religious activities
Gardening
Tending domestic animals
Caring for pets
Activities related to employment: Other specified activities related to employment
Activities related to employment: Activities related to job seeking

Activity list from UKTUS 2014

Walking the dog
Unspecified information by computing
Activities related to employment: Other unspecified activities related to employment
Travel related to other leisure
Other personal care: Other specified personal care
Travel related to physical exercise
Travel escorting an adult other than education
Travel related to hunting & fishing
Travel related to productive exercise other than hunting & fishing
Other specified gardening and pet care
Unspecified construction and repairs
Unspecified communication by computer
House construction and renovation
Repairs of dwelling
Travel to work from home and back only
Picking berries mushroom and herbs
Making repairing and maintaining equipment
Woodcraft metalcraft sculpture and pottery
Other specified making repairing and maintaining equipment
Vehicle maintenance
Travel related to gambling
Travel related to hobbies other than gambling
Other specified construction and repairs
Unspecified shopping and services
Unspecified shopping
Shopping mainly for food
Shopping mainly for clothing
Shopping mainly related to accommodation
Shopping or browsing at car boot sales or antique fairs
Window shopping or other shopping as leisure
Other specified shopping
Commercial and administrative services
Personal services
Other specified organisational work
Travel related to changing locality

Activity list from UKTUS 2014

Unspecified household and family care
Travel to holiday base
Travel for day trip/just walk
Other specified shopping and services
Live music other than classical concerts opera and musicals
Household management not using the internet
Shopping for and ordering clothing via the internet
Unspecified household management using the internet
Shopping for and ordering unspecified goods and services via the internet
Shopping for and ordering food via the internet
Shopping for and ordering goods and services related to accommodation via the internet
Shopping for and ordering mass media via the internet
Shopping for and ordering entertainment via the internet
Banking and bill paying via the internet
Other specified household management using the internet
Art exhibitions and museums
Other specified travel
Unspecified radio listening
Unspecified library
Punctuating activity
Unknown: at home
Unspecified childcare
Filling in the time use diary
Study: Unspecified study school or university
Unspecified physical care & supervision of a child
Feeding the child
Other specified participatory activities
No main activity no idea what it might be
Using computers in the library other than internet use
Other and unspecified physical care & supervision of a child
Teaching the child
No main activity some idea what it might be
Reading playing and talking with child
Illegible activity
Accompanying child

Activity list from UKTUS 2014

Listening to recordings
Unspecified time use
Queryable
Other specified performing arts
Listening to sport on the radio
Travel related to unspecified time use
Other or unspecified childcare
Travel related to personal business
Travel related to services
Unspecified mass media
Unspecified help to a non-dependent eg injured adult household member
Physical care of a non-dependent e.g. injured adult household member
Accompanying a non-dependent adult household member e.g. to hospital
Other specified help to a non-dependent adult household member
Unspecified help to a dependent adult household member
Physical care of a dependent adult household member e.g. Alzheimic parent
Accompanying a dependent adult household member e.g. Alzheimic
Solo games and play
Unspecified hobbies games and computing
Other specified help to a dependent adult household member
Watching sport on TV
Travel related to shopping
Unspecified sports and outdoor activities
Unspecified employment
Unspecified social life and entertainment
Travel to/from work
Travel in the course of work
Chess and bridge
Unspecified volunteer work and meetings
Unspecified reading
Travel to work from a place other than home
Reading periodicals
Reading books
Unspecified arts
Unspecified visual arts

Activity list from UKTUS 2014

Painting drawing or other graphic arts
Making videos taking photographs or related photographic activities
Not applicable
Computer games
Other specified visual arts
Unspecified performing arts
Singing or other musical activities
Unspecified physical exercise
Help to other households in employment and farming
Literary arts
Walking and hiking
Taking a walk or hike that lasts at least miles or 1 hour
Other specified arts
Other walk or hike
Jogging and running
Unspecified social life
Unspecified hobbies
Activities related to productive exercise
Biking skiing and skating
Biking
Skiing or skating
Indoor team games
Socialising with family
Collecting
Travel related to education
Unspecified ball games
Indoor pairs or doubles games
Other specified reading
Outdoor pairs or doubles games

B.1 Definition of private leisure

In the 2014 data, I select the following activities in my baseline definition of private leisure. I also conduct robustness analyses removing some of the activities which are borderline between leisure and domestic work, such as gardening, and activities which, like sleep, may

be more likely to have externalities on other family members. Both in my baseline and robustness checks, I focus only on leisure time which was spent ‘NOT with other household members’ (there are several binary variables coding co-presence information), with the exception of sleep, where I allow other household members to be co-present.

Leisure activities

Unspecified personal care
Visiting and receiving visitors
Other specified ball games
Gymnastics
Visiting a wildlife site
Celebrations
Fitness
Unspecified TV video or DVD watching
Watching a film on TV
Other specified or unspecified arts and hobbies
Unspecified water sports
Other specified TV watching
Watching a film on video
Watching sport on video
Other specified water sports
Other specified video watching
Other specified physical exercise
Other specified social life
Communication on the internet
Unspecified entertainment and culture
Other specified communication by computing
Skype or other video call
Cinema
Main job: Coffee and other breaks in main job
Unspecified theatre or concerts
Plays musicals or pantomimes
Opera operetta or light opera
Concerts or other performances of classical music
Dance performances
Unspecified listening to radio and music

Leisure activities

Other specified theatre or concerts

Sleep

Sleep: In bed not asleep

Listening to music on the radio

Reading newspapers in a library

Other specified radio listening

Sports events

Unspecified games

Unspecified games and play with others

Billiards pool snooker or petanque

Other specified parlour games and play

Swimming

Unspecified sports related activities

Activities related to sports

Unspecified video watching

Other unspecified entertainment and culture

Visiting a historical site

Gambling

Visiting a botanical site

Visiting a leisure park

Other or unspecified entertainment or culture

Resting - Time out

Second job: Coffee and other breaks in second job

Outdoor team games

Unspecified participatory activities

Eating

Other specified games

Activities related to employment: Lunch break

Other personal care: Unspecified other personal care

Other personal care: Wash and dress

Gardening

Other personal care: Other specified personal care

Unspecified communication by computer

Woodcraft metalcraft sculpture and pottery

Window shopping or other shopping as leisure

Leisure activities

Travel for day trip/just walk

Live music other than classical concerts opera and musicals

Art exhibitions and museums

Unspecified radio listening

Other specified participatory activities

Other specified performing arts

Listening to sport on the radio

Unspecified mass media

Solo games and play

Unspecified hobbies games and computing

Watching sport on TV

Unspecified sports and outdoor activities

Unspecified social life and entertainment

Chess and bridge

Unspecified reading

Reading periodicals

Reading books

Unspecified arts

Unspecified visual arts

Painting drawing or other graphic arts

Making videos taking photographs or related photographic activities

Computer games

Other specified visual arts

Unspecified performing arts

Singing or other musical activities

Unspecified physical exercise

Literary arts

Walking and hiking

Taking a walk or hike that lasts at least miles or 1 hour

Other specified arts

Other walk or hike

Jogging and running

Unspecified social life

Unspecified hobbies

Biking skiing and skating

Leisure activities

Biking

Skiing or skating

Indoor team games

Collecting

Unspecified ball games

Indoor pairs or doubles games

Other specified reading

Outdoor pairs or doubles games

Activities are defined in a similar manner in UKTUS 2000.

C Expenditure data and individual-level consumption estimation

C.1 The dataset

The LCF (Living Costs and Food Survey) (previously FES) is a UK survey containing information on individual and household characteristics, individual labour supply, detailed income data, and very detailed expenditure data. It is a repeated cross-section available yearly since 1978.

It is a high-quality, nationally representative, large-scale survey that is used to estimate official government statistics. The FES/LCF data has been widely used both for academic and policy applications due both to its scale and high quality. For instance, Bargain, Donni and Hentati (2022) and Lise and Seitz (2011) use this data.

The survey has multiple components: (i) a household survey recording household characteristics and retrospective questions on irregular expenses (rent, vehicles...); (ii) an individual questionnaire with individual characteristics, including demographic characteristics, hours worked and sources of income;¹¹⁵ (iii) a detailed two-week expenditure diary for all members older than 7 (simplified diary for people aged 7-15, full diary for people aged 16 or above). The household questionnaire is answered by the reference person either alone, or together with other household members. Individual surveys, and expenditure diaries, are answered

¹¹⁵In rare instances, income is top-coded. I adjust top-coded values using data on after-tax income percentiles from HM Revenue & Customs. HM Revenue & Customs. (2023). Percentile points from 1 to 99 for total income before and after tax. <https://www.gov.uk/government/statistics/percentile-points-from-1-to-99-for-total-income-before-and-after-tax>

by the relevant person. The expenditure diaries are kept for two weeks by all household members.

The expenditure diary records the type of good in detail, and receipts are attached. Clothing and footwear is divided into male, female, and children.¹¹⁶ Household expenditure is obtained by summing expenditure over all members. Additional information on expenditure on large infrequent expenses, such as house repairs, and regular expenses, such as rent, is obtained during the household survey. These expenditures are transformed to an equivalent weekly value to make them comparable to other categories.

After restricting the LCF 2014 sample to heterosexual working couples and cleaning the data, the final sample comprises 583 households (i.e. 1,166 individuals). Key variables are summarised in table 1.

C.2 Categorisation of expenditures into private and public

The next step towards estimating individual-level inequality is to divide the LCF household-level expenditure data into (i) private expenditure, (ii) public expenditure, and (iii) expenditure to be excluded from consideration. In deciding how to do so, it is important to consider the goal of the exercise. In the case of this paper, the aim is to compare the material standard of living of different individuals in the UK.

C.2.1 Excluded categories

I exclude expenditure categories which have almost no immediate consumption value and cannot easily be squared with the static model underlying the methodology in this paper. This includes: savings, insurance, investments, major house works including renovation (minor repairs are included), financial gifts, bets, gambling and expenditure on education (the latter is minor for working couples without cohabiting children). Future work considering dynamic aspects would enable incorporating these categories into overall resources in a theory-consistent manner. Again with in mind the goal of measuring material well-being, I focus on expenditures gross of any government refund or subsidy . Sometimes expenditures are partly funded by the government and this is visible in the expenditure data for some goods. Where this is the case, I do not detract any subsidies or refunds from expenditure, since these still contribute to material well-being.¹¹⁷

¹¹⁶Note that any particular good may be purchased by a household member but consumed by any combination of them and/or other members, so personal expenditures do not measure personal consumption.

¹¹⁷For other applications, e.g. to estimate what proportion of the population has access to a minimal standard of living, we might not consider the consumption of ‘bads’ to contribute to that standard of living, and hence also exclude those.

C.2.2 Private vs public consumption

Based on the detailed COICOP plus codes into which expenditure is divided in the LCF, I divide expenditure between private and public. While it would be possible to categorise expenditure based on less granular data, the granularity aids accuracy. For instance, most house-cleaning products are categorised as public (as they contribute to the public cleanliness of the house) but washing powders are categorised as private as more is needed to wash the clothes of more members.

Discretion is needed in categorising goods, as most goods have at least some public element, including externalities of consumption on other household members. In some cases this categorisation could be made conditionally on household characteristics. For instance, holiday accommodation is a public good for couples, who will generally stay in the same room. For families with older children, one might consider holiday accommodation as partly private if teenagers are likely to e.g. have a separate hotel room.

In some cases, it might be possible to use other data (e.g. data on car occupancy) to estimate to what degree a good (e.g. cars) is public and to what degree it is private. We can think of car-related expenditures (insurance, fuel, etc.) as separate goods depending on the type of trip, where some of them are private and some of them are public. A car and related expenses could be purely public if everyone in the household only used it together (e.g. to drive to a holiday home). Alternatively it could be purely private if only one household member used it to drive to work. As discussed in more detail below, I approximate the likely economies of scale of car-related expenditures using UK car occupancy data.

C.2.3 Treatment of durables, including homes and vehicles

While some papers in this literature focus on non-durable consumables, I suggest it is important to also consider durables. In particular, housing is a very durable and infrequently purchased, but it is important to take it into account since it is such a large expenditure for many households, and since it is a very important component of public consumption and household economies of scale.

For less expensive durables such as clothing and phones, I use the LCF expenditure data without adjustments. While for any specific category this is likely to lead either to over-estimates or under-estimates, the overall expenditure across categories is likely to be a reasonably accurate estimate of usual expenditure. For very large expenses (buying a home or a vehicle), a different approach is needed, since the magnitude of these expenditures dwarves regular weekly purchases. Moreover, these expenditures are often diluted over some periods of time, with a mortgage or loan, with repayments depending not only on the quality of

the good being purchased but also on factors which are not directly relevant to material well-being, such as macroeconomic conditions at the time of purchase and individual credit score. Keeping in mind the goal of measuring material well-being, we wish to estimate the value of the good being consumed (e.g. a home in a certain area with a certain number of rooms) while abstracting from extraneous considerations (e.g. whether the home is owned or rented). Of course the household may financially benefit from owning outright instead of renting, through decreased monthly expenditures on housing. However, this will translate into increased expenditures on other goods, and hence will still be taken into account insofar as it affects material well-being.

Imputed rents for homes Housing is a particularly complex good from the viewpoint of measuring inequality in the UK. The prices of homes with similar attributes in different regions differs vastly, but households generally have limited choice as to their location, especially in the short-run (due to jobs, relationships, and coordination between members). While it is approximately true that goods like food are similarly priced across the UK, assuming price homogeneity for homes across the UK would be too unrealistic (e.g. homes in the London area command a very large premium). Another source of difficulty is that homes are sometimes owned outright, sometimes they are purchased with a mortgage (the interest on which varies vastly by year and credit score of the purchasing party), sometimes they are rented privately, and sometimes publicly (e.g. council housing provided as a benefit). To avoid both the issue of infrequency and the problem of comparability between renters, outright owners, and owners paying back a mortgage, I impute a weekly standardised rental price. This price increases in the quality of the home (as measured by the number of bedrooms, centrality of location, etc.) but is standardised across different purchasing conditions, as well as for regional house price differences (the latter being mostly something that affects the investment value of the property rather than reflecting the underlying quality of the home).

I suggest taking the following approach to balance between the two opposing goals of comparability between households and granularity of household-specific expenditure:

- Using data on private renters, I regress rent expenditure on
 1. Number of bedrooms. This is the best proxy for home size available in the data.
 2. Area type. The 2011 Area Classification for Output Areas (OAC) categorises postcodes into types e.g. areas dominated by ‘urban professional and families’. Using granular OAC codes, I define four aggregated area types: rural, more desirable urban, less desirable urban, and suburban.

3. Council tax band. Council tax bands are available for homes in England, Wales, and Scotland (I impute them for Northern Ireland). Council tax bands are based on legacy valuations of homes, and provide a good signal of the quality of the home.
 4. Region (Wales, Scotland, Northern Ireland, and the nine regions of England). There are substantial regional disparities in the prices of homes in the UK.
- I impute the rent for all households in the data (including those who own, rather than rent) based on these characteristics.
 - Next, I standardise imputed rents across regions by indexing them to the region with the lowest median imputed rent (the North East of England). I deflate the imputed rents for homes in other regions by the ratio of the median imputed rent in their region to the median imputed rent in the North East of England. This procedure can be thought of as uncovering the fundamental quality of the consumption obtained from the house, disentangling it from other considerations like investment value of property. We can also motivate this choice by noting that households often cannot move region (at least in the short-run), but can choose the specific location of their home (with different associated OAC code), its size (proxied by the number of bedrooms) and broader quality (proxied by council tax band).
 - For each household I record the standardised imputed rent.
 - A small number of households have a second home, but the data contains no information on the second home. Where this is the case, I double the standardised imputed rent of the main home, since the value of the main home is the best available signal in the data of the likely value of the second home.
 - I add the standardised imputed rents to public household expenses.

Imputed lease price for vehicles Car purchases are the second largest expenditure items after homes for many UK households. For this reason, it seems important not to exclude this expenditure. Following a principle similar to the one outlined above for housing, I propose a rental approach to vehicles (cars, vans and motorcycles):

- I calculate the median weekly lease price paid by households that lease a vehicle.
- For each household, I observe how many vehicles they own, and estimate their vehicle lease expenditure by multiplying the average lease price by the number of vehicles owned.

- I add 54% of this imputed lease to private expenditures and the remaining 46% to public expenditures. This approximation is based on data from the National Travel Survey¹¹⁸ on the proportion of car trips by purpose, and data on the occupancy rate of cars by trip purpose.¹¹⁹ This method aims to capture the likely economies of scale of cars, although of course further granularity could be achieved if the data contained information on how the household uses cars. We can think of this as dividing car-related expenditures into multiple goods, some of which are private (e.g. car trips for solo work trips) and some of which are public (e.g. car trips for family holidays).¹²⁰

D Considerations on choosing the assignable good

D.1 Availability and accuracy of assignable good data

Ideally we would have accurate, widely available data for our assignable good.

For clothing, the breakdown between men, women and children is often (though not always) available. Where it is available, the accuracy of clothing expenditure data may vary depending on whether it was derived from expenditure diaries or recall questions. Where expenditure by men, women and children is not available, some papers e.g. Calvi (2020) estimate these expenditure categories from expenditure on specific types of clothing e.g. assigning pyjamas to men, which may add measurement error. Narrowing down to specific clothing types such as skirts and ties is not a viable alternative as these purchases are very infrequent and command a very small budget share (clothing itself commands a low budget share).

Similarly, time-use data may be more or less accurate depending on whether it is based on recall questions or on a detailed time-diary. High-quality time-use data is available for many countries, and generally includes information both on very detailed activities and who was co-present during the activity (e.g. the UK Time Use Data used in the application in this paper). Using the detailed activity information it is possible to construct accurate measures of time spent on leisure activities, and using the co-presence information it is possible to exclude leisure time which was joint rather than private. Where this is possible, private leisure is likely measured accurately, especially where the time diary was taken for a longer

¹¹⁸Department for Transport. (2021). National Travel Survey: 2020. <https://www.gov.uk/government/statistics/national-travel-survey-2020/national-travel-survey-2020#trends-in-car-trips>

¹¹⁹Department for Transport. (2022). Car or van occupancy and lone driver rate by trip purpose, England: 2002 onwards. <https://www.gov.uk/government/statistical-data-sets/nts09-vehicle-mileage-and-occupancy>

¹²⁰I do not attempt to adjust this split based on the numbers of vehicles owned as I am not aware of data that would enable such an adjustment.

period of time. Measurement error may be higher for proxies of leisure obtained from recall questions on hours of leisure in a typical week (non-market-work).

Individual-level food expenditure is a good candidate assignable good, but it is resource-intensive to collect accurate data on it, and hence it is rarely available in practice.

D.2 Credibility of the assignable good being private

Identification relies on the assignable good being private. In general, using an assignable good with a substantial public element (externalities) will bias resource share estimates towards equality. It is likely that, in reality, any assignable good will have some externalities. Ideally we would choose a good which is close to being fully private.

Conceptually, clothing is shareable, even between different types of people. Especially in poorer households, we may worry that this can introduce a substantial public element of clothing. Moreover, even when each person wears their own clothing, people often have strong views about the way their family members, especially their partner, should dress in public.

Non-market work includes a substantial public element due to its inclusion of domestic work and joint leisure. Where high-quality time-use data is available, private leisure is more credibly private. Of course, we may still worry that people care about what their family members do in their spare time, but it is hard to imagine a good for which we would have weaker concerns about externalities than for private leisure (individual-level food expenditure also has similarly small conceptual externalities).

D.3 Frequency of purchases and magnitude of budget share

Ideally, our assignable good will command a high budget share and there will be a very low proportion of people hitting a corner solution on it.

Clothing is an infrequent purchase, and the proportion of zeros is very high in many datasets used in the literature (most expenditure surveys are taken over short periods of time, e.g. two weeks). In the UK expenditure data used in the application in this paper, 73.20% of households have zero recorded expenditure on at least one of male and female clothing.¹²¹ This is problematic because identification relies on demand functions which treat observations as interior solutions. Therefore, using an assignable good with a high proportion of zeros adds substantial inaccuracy to the estimation approach, as corner solutions are treated

¹²¹Some surveys may ask recall questions on clothing expenditure to alleviate this problem, but the accuracy of responses to such questions is more doubtful.

as interior solutions. Moreover, since zeros are not merely driven by random infrequency but are also systematically related to low resource shares, sharing rule estimates are likely to be biased towards equality.

By contrast, households spend a large proportion of their resources on individual-level food expenditure, non-market-work and private leisure, with virtually no zeros in the data, so using it as the assignable good.

D.4 Further considerations

Clothing durability. Clothing is highly durable and often passed down through generations (e.g. older to younger cousins), especially for small children. This means that clothing consumption can be substantially different from clothing purchases. This may lead to additional inaccuracies arising when using clothing as the assignable good. It might potentially also contribute to biasing estimates if e.g. clothes hand-me-downs are more frequently enjoyed by certain sub-groups, such as small children or young women.

Clothing expenditure is not recorded at the individual level. Clothing expenditure is almost never measured at the individual level in expenditure data, and is only available in broad break-downs such as men's, women's and children's clothing. In contexts with larger households and multiple members of each type, common especially in developing countries, this reduces the granularity of estimates to type-level, rather than individual-level, estimates of resources.

E The issue of non-participation

The difference between the case of participation and non-participation is illustrated in figure 11. The left panel shows the case of an interior solution with market participation, while the right panel exemplifies non-participation. To simplify exposition, I group domestic work and joint leisure into a single public time category. This is to be understood as the optimal combination of domestic work and of joint leisure. In the interior solution case, there is an internal market for the individual's time, so that the price of time is the observed market wage. In the case of market non-participation, the (unobserved) market wage is lower than the returns from optimal non-market-work time-use. This substantially complicates the problem. The price of public time-use is not the (unobserved) market wage, so we would need to estimate this unobserved price to use private leisure as the assignable good (and more generally for other assignable goods if separability is not assumed). This is one of two

reasons why this paper focuses only on households where all members participate in market work. The second reason is that we can more credibly test the identifying assumptions for the sub-sample of working couples. As discussed in section 9.4.2, tests from data on working singles appear to be highly informative for this sub-sample, but may be less convincing for other household compositions.

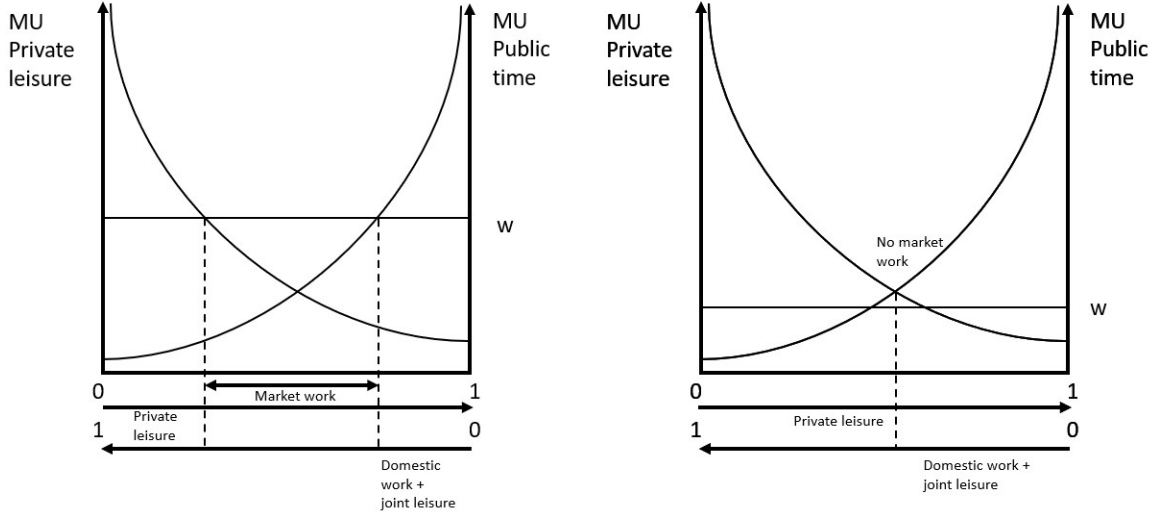


Figure 11: Interior solution vs. corner solution for time-use

E.1 Why excluding couples with non-participation is unlikely to bias estimates

In my application to UK data, I focus only on couples without cohabiting children. In this setting, I restrict my sample only to couples where both members supply some market work. It is important to note that this is unlikely to bias the estimates, for the following reason. In the UK, for couples without dependent children, men and women have very similar, and high, participation rates. Non-participation is due to reasonably exogenous drivers such as long-term illness, disability or temporary unemployment. Importantly, non-participation is not driven by a member having a particularly high, or low, resource share. Hence, excluding non-participating couples should not bias resource share estimates for participating couples. For couples with dependent children, it would be much more problematic to exclude non-participating couples. This is because, in the UK, there is a substantial gap in participation between men and women with dependent (especially very young) children. It is realistic to think that part of this phenomenon is driven by women with very low potential wages, and low resource shares, taking on a full-time childcare and domestic work commitment, leaving

them too little leisure time to additionally take on market work. In my application, I do not incur the latter difficulty because I focus only on couples without cohabiting children. While the having of children may itself be endogenous, it would require a very different, dynamic, model to take this into account appropriately, and difficulties would arise in modelling the preference changes associated with changes in household composition.

Of course it is important to remember that my estimates are not likely to be externally valid to other household compositions, which is why I do not attempt to use my estimated sharing rule for couples with children.

Finally, I note that, if using clothing as the assignable good, excluding households with zero expenditure on clothing may bias results because zero expenditure on clothing may be driven either by infrequency of purchase or, problematically, by actual zeros due to low resource shares.

F Cobb-Douglas: additional material

F.1 Cobb-Douglas Domestic production functions

Household economics, despite its emphasis on households, frequently does not involve modelling domestic production because it introduces identification issues. However, it is clear that domestic production does play an important role in understanding individual-level resources (see for instance Apps and Savage (1989)). One way to model domestic production within a collective model with relative ease is to assume that the domestic good is marketable - it can be bought and sold. This assumption has the advantage of leading to separability between the production and consumption functions of the household as the price of the domestic good is exogenously determined by the market, and not endogenously within the household (see e.g. Browning, Chiappori and Weiss (2014) section 4.6.2).

However, this assumption cannot easily be reconciled with empirical facts. If market work has a constant return of $w_{i,h}$ and domestic work has a constant return $w_{i,h}^{dm}$, then each individual would either supply market work or domestic work, and not both. This is in contrast with empirical evidence that a very substantial proportion of the population do both.

One way of trying to reconcile marketability and this empirical fact is if market work is constrained in terms of hours, a fact which is consistent with some of the literature on elasticity of labour supply. This assumption is sensible in some contexts, but is less likely realistic for low skilled work, where one is more likely to be able to ask to do overtime shifts, or have multiple jobs alongside each other to make ends meet. Moreover, this assumption

would still not be reconcilable with the empirical fact that many people work part-time and also do domestic work. These people are clearly not constrained in terms of number of hours doing market work (especially as part-time hours are quite heterogeneous) and yet they do both types of work.

This is suggestive that we should not think of domestic work as having constant returns. If we wish to maintain the assumption of constant returns to market work, the most coherent way forward is to avoid modelling domestic work as marketable, as it would then become hard to justify decreasing returns to one type of marketable work and constant returns to another. Moreover, it is simply not realistic to model domestic work as being marketable. We can substitute between it and market purchased goods to some degree, but they are not the same good. Cleaning is perhaps the example where the boundary is most blurred, but making a home-cooked meal while looking after one's own child are clearly not the same as buying take-out food and sending the child to a nursery. Empirically, even people with high wages spend some time on domestic work - the reason is that it cannot be purchased, and it is desirable.

Therefore, we are left with the problem of modelling domestic production. Modelling domestic production is a complex task, and one that has historically not received nearly as much attention as aspects relating to market work. The literature which does exist has used a variety of production functions, depending on the aims of the analysis and the key features from the data of interest. For instance, Griffith et al. (2022) employ a Leontief production function for home cooked food where the inputs are market purchased ingredients and domestic time. They further assume constant returns to domestic time, and perfect substitution between different household members' time (time spent cooking is just the sum of individual time spent cooking, and household leisure is just the sum of individual time spent on leisure).

For the purpose of this paper, it is instead important to model decreasing returns to domestic work to avoid the model being irreconcilable with the empirical facts described above. It is also important to model the domestic time of different members as not being perfect substitutes, again to explain patterns of behaviour with market and domestic work. I propose using a simple Cobb-Douglas production function. The productivity of each type of household member (and of the same types in different household categories) is heterogeneous. The concavity of the production function ensures that everyone in the household will do some domestic work.

An advantage of this approach is that the same functional form also seems appropriate to model non-private leisure. We want everyone in the household to have at least some time together (e.g. many households insist on a dinner with everyone, even those who are very

busy), and there are higher returns to public leisure if others are able to join in. The Cobb-Douglas production function captures this concavity, and also allows us to model possible heterogeneity in the returns to joint leisure (e.g. the quality of joint leisure might be higher if children are involved, as it avoids the negative externalities of children being away from their parents). The model lends itself to greater granularity where relevant, for instance separating out leisure that is joint between strict subsets of household members.

However, for the purpose of this paper, consider the following production functions:

$$D_h = \delta_g \prod_{i \in h} (d_{i,h})^{\delta_{t,g}}$$

$$T_h = \phi_g \prod_{i \in h} (jt_{i,h})^{\phi_{t,g}}$$

The Cobb-Douglas production functions, substituted into the Cobb-Douglas utility functions, yield a Cobb-Douglas functional form, preserving the advantages of the Cobb-Douglas preferences:

$$u_{t,g} = \sum_{j \in \Omega^c} (\alpha_{t,g}^{cj} \ln(c_{i,h}^j)) + \sum_{j \in \Omega^X} (\alpha_{t,g}^{Xj} \ln(X_h^j)) + \alpha_{t,g}^l \ln(l_{i,h}) + \alpha_{t,g}^T \ln(T_h) + \alpha_{t,g}^D \ln(D_h)$$

$$\begin{aligned} u_{t,g} = \sum_{j \in \Omega^c} (\alpha_{t,g}^{cj} \ln(c_{i,h}^j)) + \sum_{j \in \Omega^X} (\alpha_{t,g}^{Xj} \ln(X_h^j)) + \alpha_{t,g}^l \ln(l_{i,h}) \\ + \alpha_{t,g}^T \ln(\phi_g \prod_{s \in h} (jt_{s,h})^{\phi_{st,g}}) + \alpha_{t,g}^D \ln(\delta_g \prod_{s \in h} (d_{s,h})^{\delta_{st,g}}) \end{aligned}$$

The above can be re-written as a standard Cobb-Douglas utility function as a direct function of time-use:

$$\begin{aligned} u_{t,g} = \sum_{j \in \Omega^c} (\alpha_{t,g}^{cj} \ln(c_{i,h}^j)) + \sum_{j \in \Omega^X} (\alpha_{t,g}^{Xj} \ln(X_h^j)) + \alpha_{t,g}^l \ln(l_{i,h}) \\ + (\alpha_{t,g}^T \ln(\phi_g) + \alpha_{t,g}^D \ln(\delta_g)) + \alpha_{t,g}^T \sum_{s \in h} \phi_{st,g} \ln jt_{s,h} + \alpha_{t,g}^D \sum_{s \in h} \delta_{st,g} \ln d_{s,h} \end{aligned}$$

We can work with this utility function and drop the production function constraints. We may wish to re-normalise the utility functions by setting the constant terms $(\alpha_{t,g}^T \ln(\phi_g) + \alpha_{t,g}^D \ln(\delta_g))$ to 0, and define some new notation to simplify the coefficients on domestic time and joint leisure:

$$u_{t,g} = \sum_{j \in \Omega^c} (\alpha_{t,g}^{cj} \ln(c_{i,h}^j)) + \sum_{j \in \Omega^X} (\alpha_{t,g}^{Xj} \ln(X_h^j)) + \alpha_{t,g}^l \ln(l_{i,h}) + \sum_{s \in h} (\alpha_{t,g}^{Tst,g} \ln jt_{s,h}) + \sum_{s \in h} (\alpha_{t,g}^{Dst,g} \ln d_{s,h})$$

F.2 Normalising assumptions

Observationally, household behaviour is equivalent up to (i) positive affine transformations of individual utility functions, and (ii) any positive monotonic function of the sum of the individual utilities weighted by their respective Pareto weights.¹²² I normalise the model as follows:

- the constant term in the utility functions is set to zero (omitted in the equations above)
- the sum of each person's preference parameters over all goods (material and time-use, public and private) is set to one $\sum_j \alpha_{t,g}^j = 1$
- the Pareto weights are set to sum to one: $\sum_{i \in h} \mu_{i,h} = 1$

Heterogeneity extension

Note that these normalisations are carried out under the assumption that productivity in domestic work and joint leisure, like preference parameters, are type-specific constants. If types are men and women, this means we are assuming that productivity (as well as preferences) are homogeneous within women and within men. To extend the model to contexts where preferences and/or productivity vary based on characteristics like age or wages, two approaches can be taken. The first is simply to define types more granularly and then conduct estimation separately for different household categories (households with different combinations of types as members). Alternatively, drop the normalising assumptions and explicitly define preference and productivity parameters as functions of characteristics.

F.3 Public good C-D SAP

F.3.1 Under public good C-D SAP, the resource share and Pareto weight coincide

Note that $\rho_{i,h} = y_h \mu_{i,h} \left(\sum_{j \in \Omega^c} \alpha_{t,g}^j + \alpha_{t,g}^l \right)$. Also recall that:

$$\eta_{i,h} = \frac{\rho_{i,h}}{\sum_{s \in h} \rho_{s,h}} = \frac{y_h \mu_{i,h} \left(\sum_{j \in \Omega^c} \alpha_{t,g}^j + \alpha_{t,g}^l \right)}{y_h \sum_{s \in h} \mu_{s,h} \left(\sum_{j \in \Omega^c} \alpha_{st,g}^j + \alpha_{st,g}^l \right)} = \frac{\mu_{i,h} \left(\sum_{j \in \Omega^c} \alpha_{t,g}^j + \alpha_{t,g}^l \right)}{\sum_{s \in h} \mu_{s,h} \left(\sum_{j \in \Omega^c} \alpha_{st,g}^j + \alpha_{st,g}^l \right)}.$$

Now assume public good SAP so that $\left(\sum_{j \in \Omega^x} \alpha_{t,g}^j + \sum_{s \in h} \alpha_{t,g}^{Tst,g} + \sum_{s \in h} \alpha_{t,g}^{Dst,g} \right) = a_g^Q$

Equivalently, $\left(\sum_{j \in \Omega^c} \alpha_{t,g}^j + \alpha_{t,g}^l \right) = a_g^{c,l}$

¹²²Note that there is no uncertainty in this model so that the overall optimisation problem is unchanged by a positive monotonic transformation. The same cannot be said for the individual utilities because the household's optimisation problem is not to maximise a single utility but the weighted sum of all of the utilities, so that each individual utility can only be transformed up to positive affine transformations.

Having made this assumption, we can write $\eta_{i,h} = \frac{\mu_{i,h} a_g^{c,l}}{a_g^{c,l} \sum_{s \in h} \mu_{s,h}}$. By definition, $\sum_{s \in h} \mu_{s,h} = 1$ and the $a_g^{c,l}$ terms cancel out, so that $\eta_{i,h} = \mu_{i,h}$.

In this case, finding that resource shares vary with market variables implies that Pareto weights do too, and hence is evidence in favour of the collective model over the unitary model.

F.3.2 Under public good C-D SAP, second-stage demands do not depend on public expenditure

Note that $l_{i,h} w_{i,h} = \frac{\mu_{i,h} \alpha_{t,g}^l}{\sum_{s \in h} \mu_{s,h} \sum_j \alpha_{st,g}^j} y_h$, where $\sum_j \alpha_{st,g}^j$ is the sum of preference coefficients over all different goods, and was normalised to 1, and $\sum_{s \in h} \mu_{s,h} = 1$. Therefore, $l_{i,h} w_{i,h} = \alpha_{t,g}^l \mu_{i,h} y_{i,h}$. Without assuming public good C-D SAP, re-writing this equation in terms of resource shares, rather than Pareto weights, (this is necessary to be able to estimate the resource shares, which are our object of interest) requires including public expenditure $R_h Q_h$ as a term (or writing the latter in terms of bargaining fundamentals).

With public good C-D SAP, as discussed above, $\mu_{i,h} = \eta_{i,h}$ and hence we can write $l_{i,h} w_{i,h} = \alpha_{t,g}^l \eta_{i,h} y_h$. In this case, we are able to estimate resource shares from data that does not contain public good expenditure. Intuitively, this is driven by the assumption that changing bargaining power may affect how the household divides expenditure on specific public goods, but not on the aggregate public budget, so that we can ignore public good expenditure for the purpose of estimating resource shares.

F.4 Structural approach to the error term

It is good practice to include sources of noise in the structural model (see for instance Reiss and Wolak (2007) for an excellent discussion of this point). I suggest doing this in two ways here.

Firstly, it is unlikely that we can control for all characteristics that affect bargaining within the household. Even if we have many characteristics in our data, there are likely unobserved factors. Moreover, recall that we are using a linear approximation of the Pareto weight, and hence there is likely an approximation error. We can write:

$$w_{i,h} l_{i,h} = \alpha_{t,g}^l y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z}) + e_{i,h})$$

Where by definition $e_{i,h}$ is mean-zero. As long as we assume $e_{i,h}$ is independent of y_h and z_h (or, more weakly, that it is uncorrelated with $y_h, y_h^2, z_h, y_h z_h$) then $\alpha_{t,g}^l y_h e_{i,h}$ is uncorrelated with the other terms. Note that, by definition, the shocks sum to zero within

each household: $\sum_{i \in h} e_{i,h} = 0$. Hence in a two-person household this implies a correlation of -1 between the errors of members of a couple.

Secondly, there may be some optimisation error $u_{i,h}, u_{i,h}^l$ so that:

$$w_{i,h} l_{i,h} = \alpha_{t,g}^l (y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})) + e_{i,h}) + u_{i,h} + u_{i,h}^l$$

We can think of it in these terms:

1. In the first stage of the household problem, the household makes optimisation errors in the division of resources between public goods and individual budgets. This means that a member's individual budget will deviate by $u_{i,h}$ relative to the optimal budget. These errors needn't sum to zero across household members' individual budgets since there is also scope for error in public good expenditure. Note that $\sum u_{i,h} + u_{X,h} = 0$ so that the optimisation errors in the first stage must sum to zero. The $u_{i,h}$ errors will be negatively correlated across household members, with a correlation weakly smaller in magnitude than -1.
2. In the second stage of the household problem, each individual member makes optimisation errors in the division of their budget between private goods for their personal consumption. This means that each private expenditure may deviate from optimal by $u_{i,h}^j$ where $\sum_{j \in \Omega_c} u_{i,h}^j = 0$. Specifically, leisure expenditure may deviate from optimum by $u_{i,h}^l$. These errors could potentially be correlated across members, but there is no clear reason why they would be.

For both sources of optimisation error, as is standard, assume they are mean-zero and uncorrelated with each other and all other variables.

Taking these three sources of error, the final model can be written as:

$$w_{i,h} l_{i,h} = \alpha_{t,g}^l y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})) + (\alpha_{t,g}^l y_h e_{i,h} + a_{t,g}^l u_{i,h} + u_{i,h}^l)$$

Define $\epsilon_{i,h} = \alpha_{t,g}^l y_h e_{i,h} + a_{t,g}^l u_{i,h} + u_{i,h}^l$ so that $w_{i,h} l_{i,h} = \alpha_{t,g}^l y_h (\eta_{t,g}^0 + \sum_z \eta_{t,g}^z (z_h - \bar{z})) + \epsilon_{i,h}$ where $\epsilon_{i,h}$ is mean-zero and uncorrelated with all other regressors. The errors are negatively correlated within household (with a correlation that is negative but smaller in magnitude than -1), so that a SURE estimation approach is recommended.

F.5 Individual-level full consumption when expenditure data is separate from time-use data

Under Cobb-Douglas preferences, we can estimate time-inclusive individual-level consumption even when expenditure and time-use data are separate, as long as we observe hours

worked $m_{i,h}$ and hourly wage $w_{i,h}$ in the expenditure data.¹²³ This is the case in many expenditure datasets which do not contain any other information relating to time-use. The procedure is simple:

- Recall that leisure expenditure is $w_{i,h}l_{i,h} = a^l y_h \eta_{i,h}$ and that we have estimated $\hat{\eta}_{i,h}$ and $\hat{\alpha}^l$ from the time-use data
- Hence we can estimate $l_{i,h}$ in the expenditure data as $\hat{l}_{i,h} = \frac{\hat{\alpha}^l y_h \hat{\eta}_{i,h}}{w_{i,h}}$
- We therefore obtain expenditure on private leisure $w_{i,h}\hat{l}_{i,h}$

F.6 Back-of-the-envelope check

An advantage of working with preferences which have a simple direct utility representation is that we can substitute parameter estimates into indirect utility functions to check that their magnitudes (as well as their signs) are consistent with bargaining theory. Specifically, we can check whether an increase in own-wage implies an increase in own-utility.

First, I illustrate three example households: one characterised by average characteristics, one which deviates in terms of the female wage (to £15 an hour), and one which deviates in terms of the male wage (to £15 an hour). Comparison between the first and, in turn, the second and third, allows us to think about the impact of a change in female wage on female utility and a change in male wage on male utility. This impact is a combination of multiple channels: (i) a price effect on own private leisure, (ii) an income effect through overall household budget, (iii) an income effect through sharing, (iv) a price effect on joint leisure and domestic work, and (v) increased weight on own-preferences in the household's choices of the composition of the budget share devoted to public goods. Because I do not estimate the whole model, I cannot compute exact utility changes. Instead, I conduct a back-of-the-envelope check where I assume away the last two effects (which have opposite effects on utilities). Private leisure falls (the price effect dominates), but overall utility increases through increased material private consumption and increased expenditure on public goods.

Then I consider the general case and show that estimates are consistent with bargaining theory for reasonable variations in wages (I do not consider more extreme variations in wages since the linear approximation of the sharing rule entails that for very extreme wages things break down).

¹²³If the data has income from work $m_{i,h}w_{i,h}$ and hours worked $m_{i,h}$ then we can estimate $w_{i,h} = \frac{m_{i,h}w_{i,h}}{m_{i,h}}$. Where expenditure data only has higher level information, such as part-time or full-time work, employment status and sector, we can estimate hours worked based on those variables.

F.6.1 Three example households

The average household

- The daily budget is $y_h = (9.87 + 11.6) * 24 = 515.28$
- Under C-D SAP, a fixed proportion of household budget a^Q is spent on (material and time-use) public goods and enjoyed by both household members
 - We can conceptualise this expenditure in terms of a composite public good. The composition of this good varies with the distribution of bargaining power as more weight is assigned to goods preferred by the member with more bargaining power. Moreover, the price of public time (domestic work and joint leisure) depends on individual wages. Hence, we can think of the price of the aggregate good as $R(w_f, w_m)$, which for the average household would be $R(9.87, 11.6)$ ¹²⁴
 - To illustrate, consider the example of $a^Q = 0.2$ (as explained further below, the results hold under almost any theory-consistent value of a^Q)
 - Then expenditure on public goods is $R(9.87, 11.6)Q_h = 0.2 * 515.28 = 103.06$
- Under the estimated preference parameter on private leisure (which is the same for men and women), the household spends a budget share equal to $\alpha^l = 0.44$ on private leisure.
 - Hence the household spends a total on private leisure of: $w_m l_m + w_f l_f = 0.44 * 515.28 = 226.72$
 - Because this is a private expenditure, it is assigned to each member according to the sharing rule
 - For the average household, we estimate that the woman has a share of 0.44 and the man has a share of 0.55
 - Therefore, expenditure on the woman's private leisure is: $w_f l_f = 0.44 * 226.72 = 102.03$
 - Evaluated at the mean female wage, $l_f = \frac{102.03}{9.87} = 10.34$
 - Expenditure on the man's private leisure: $w_m l_m = 0.55 * 226.72 = 124.70$
 - Evaluated at the mean male wage, $l_m = \frac{124.70}{11.6} = 10.75$

¹²⁴R depends on wages both as prices and as determinants of sharing. I omit other factors affecting sharing because they are kept constant at sample averages throughout this exercise.

- Total material private expenditure falls out of the budget, private leisure expenditure and public expenditure: $p_f c_f + p_m c_m = (1 - 0.2 - 0.44) * 515.28 = 185.50$
 - Similarly to the composite public good, we can think of a composite material private good for women and a composite material private good for men. Unlike public goods, the composition of each of these material private composite goods does not vary with wages because of the properties of Cobb-Douglas preferences, and hence the price is constant for the purpose of this exercise (the price of the woman's composite material private good may differ from the man's because they may have different preferences on different private material goods).
 - We assign total material private expenditure between members according to the sharing rule
 - Therefore, $p_f c_f = 185.50 * 0.45 = 83.48$
 - And $p_m c_m = 185.50 * 0.55 = 102.03$
 - Hence, $c_f = \frac{83.48}{p_f}$
 - And $c_m = \frac{102.03}{p_m}$
- Back-of-the-envelope simplification:
 - When a member's wage increases, this leads to three effects on their utility from public expenditure:
 - * A price effect (the price of public time has increased) which lowers utility
 - * A composition effect (the composition of the public good becomes more aligned with their preferences on public goods) which increases utility
 - * An income effect (household budget increases) which increases utility
 - Since we do not estimate the whole model, we cannot quantify the first two effects. As they are opposite in direction, as a back-of-the-envelope simplification we assume them away (we can think of the price of the composite public good as a constant R , and of the man and woman having the same preference parameter for it).
- Substituting all of the above into Cobb-Douglas utilities over private leisure, an aggregate private material good, and an aggregate public (time-use and material) good, yields the following utilities:
 - The woman's utility is: $u_f^0 = 0.44 \ln(10.34) + 0.2 \ln(\frac{103.06}{R}) + 0.36 \ln(\frac{83.48}{p_f})$

- And the man's is: $u_m^0 = 0.44 \ln(10.75) + 0.2 \ln(\frac{103.06}{R}) + 0.36 \ln(\frac{102.03}{p_m})$

The household characterised by average characteristics excepting that the woman earns £15 an hour instead of £9.87

To check consistency with bargaining theory, we evaluate the woman's utility in this scenario and compare it to her utility in the average household. We should find it is higher for the woman that earns a higher-than-average wage.

- Household budget: $y_h = (15 + 11.6) * 24 = 638.40$
- Public expenditure: $RQ_h = 0.2 * 638.40 = 127.68$
- Sharing rule: $\eta_f = 0.51$ (see section 9.2.1 for how to calculate the sharing rule based on household characteristics)
- Expenditure on total private leisure: $w_m l_m + w_f l_f = 0.44 * 638.40 = 280.90$
 - Expenditure on the woman's leisure is: $w_f l_f = 0.51 * 280.90 = 143.26$
 - Her private leisure is: $l_f = \frac{143.26}{15} = 9.56$
- Expenditure on material private goods: $p_f c_f + p_m c_m = (1 - 0.2 - 0.44) * 638.40 = 229.82$
 - Expenditure on the woman's private material goods is: $c_f = \frac{0.51 * 229.82}{p_f} = \frac{117.21}{p_f}$
- The woman's utility now is: $u_f^{HF} = 0.44 \ln(9.56) + 0.2 \ln(\frac{127.68}{R}) + 0.36 \ln(\frac{117.21}{p_f})$
- Then $\Delta u_f = u_f^{HF} - u_f^0 > 0$
- Consistently with bargaining theory, the woman's wage increase has increased her utility

The household characterised by average characteristics excepting that the man earns £15 an hour instead of £11.6

Similarly, to check consistency with bargaining theory, we evaluate the man's utility in this scenario and compare it to his utility in the average household. We should find it is higher for the man that earns a higher-than-average wage.

- Household budget: $y_h = (9.87 + 15) * 24 = 596.88$
- Public expenditure: $RQ_h = 0.2 * 596.88 = 119.38$
- Sharing rule: $\eta_m = 0.57$

- Expenditure on total private leisure: $w_m l_m + w_f l_f = 0.44 * 596.88 = 262.63$
 - Expenditure on the man's leisure is: $w_m l_m = 0.57 * 262.63 = 149.70$
 - Her private leisure is: $l_m = \frac{149.70}{15} = 9.98$
- Expenditure on material private goods: $p_f c_f + p_m c_m = (1 - 0.2 - 0.44) * 596.88 = 214.88$
 - Expenditure on the man's private material goods is: $c_m = \frac{0.57 * 214.88}{p_m} = \frac{122.48}{p_m}$
- The man's utility now is: $u_m^{HM} = 0.44 \ln(9.98) + 0.2 \ln(\frac{119.38}{R}) + 0.36 \ln(\frac{122.48}{p_m})$
- Then $\Delta u_m = u_m^{HM} - u_m^0 > 0$
- Consistently with bargaining theory, the man's wage increase has increased his utility

F.6.2 General result

I now generalise the examples above to general changes in female and male wages under general values of a^Q

General result for women's utility

- $u_f = 0.44 \ln(\frac{0.44 \eta_f (w_f + w_m)}{w_f}) + a^Q \ln(\frac{a^Q (w_f + w_m)}{R}) + (1 - a^Q - 0.44) \ln(\frac{\eta_f (1 - a^Q - 0.44) (w_f + w_m)}{p_f})$
 - $\eta_f = 0.45 + 0.01(w_f - 9.87)$
 - $\frac{d\eta_f}{dw_f} = 0.01$
- We can re-write this as:

$$u_f = \ln(w_f + w_m) + (1 - a^Q) \ln(\eta_f) - 0.44 \ln(w_f) + \left(0.44 \ln(0.44) + a^Q \ln(\frac{a^Q}{R}) + (1 - a^Q - 0.44) \ln(\frac{(1 - a^Q - 0.44)}{p_f}) \right)$$

- $\frac{du_f}{dw_f} = \frac{1}{w_f + w_m} + \frac{1 - a^Q}{\eta_f} \frac{d\eta_f}{dw_f} - \frac{0.44}{w_f}$
- substituting in $\frac{d\eta_f}{dw_f}$ and η_f ,
- $\frac{du_f}{dw_f} = \frac{1}{w_f + w_m} + 0.01 \frac{1 - a^Q}{0.45 + 0.01(w_f - 9.87)} - \frac{0.44}{w_f}$
- The condition we are interested in checking is: $\frac{1}{w_f + w_m} + 0.01 \frac{1 - a^Q}{0.45 + 0.01(w_f - 9.87)} - \frac{0.44}{w_f} \leq 0$

- Recall we are thinking of a household with all other characteristics at average, including the male wage of $w_m = 11.6$
- $\frac{1}{w_f + 11.6} + 0.01 \frac{1-a^Q}{0.45 + 0.01(w_f - 9.87)} - \frac{0.44}{w_f} \leq 0$
- We evaluate this for three options of w_f to find which values of a^Q are consistent with bargaining theory in each case:
 - the lower quartile of the female wage (£6.6)
 - the mean female wage (£9.87)
 - the upper quartile of the female wage (£11.6)
- For a woman with a lower quartile wage:
 - $\frac{1}{6.6 + 11.6} + 0.01 \frac{1-a^Q}{0.45 + 0.01(6.6 - 9.87)} - \frac{0.44}{6.6} = \frac{1}{18.2} + \frac{1-\beta}{41.73} - \frac{2.2}{33}$
 - This value is consistent with bargaining if it is positive, which occurs for values of a^Q such that $\frac{1}{18.2} + \frac{1-a^Q}{41.73} - \frac{2.2}{33} > 0$
 - $a^Q < 0.51$
 - This parameter range almost certainly contains the true value of a^Q . Since the estimated share of household budget spent on private leisure is 0.44, if the household were to spend more than 51% of its budget on (time-use and material) public goods, then it would have less than 5% of the budget available to spend on material private goods, which is unrealistically low.
- For a woman with the mean wage of £9.87
 - $\frac{1}{9.87 + 11.6} + 0.01 \frac{1-a^Q}{0.45} - \frac{0.44}{9.87} \leq 0$
 - this is consistent with bargaining theory if
 - $\frac{1}{21.47} + \frac{1-a^Q}{45} - \frac{0.44}{9.87} > 0$
 - $a^Q < 1.09$
 - This is definitionally the case
- For a woman with an upper quartile wage (£11.6):
 - $\frac{1}{11.6 + 11.6} + 0.01 \frac{1-a^Q}{0.45 + 0.01(11.6 - 9.87)} - \frac{0.44}{11.6} \leq 0$
 - This is consistent with bargaining theory if:
 - $\frac{1}{23.2} + \frac{1-a^Q}{46.73} - \frac{0.44}{11.6} > 0$

- $a^Q < 1.24$
- This is definitionally the case

General result for men's utility

- $u_m = 0.44 \ln\left(\frac{0.44\eta_m(w_f+w_m)}{w_m}\right) + a^Q \ln\left(\frac{a^Q(w_f+w_m)}{R}\right) + (1-a^Q-0.44) \ln\left(\frac{\eta_m(1-a^Q-0.44)(w_f+w_m)}{p_m}\right)$
- $\eta_m = 0.55 + 0.005(w_m - 11.6)$
- $\frac{d\eta_f}{dw_f} = 0.005$
- simplifying assumption that $\frac{dR}{dw_m} = 0$
- $\frac{du_m}{dw_m} = \frac{1}{w_f+w_m} + \frac{1-a^Q}{\eta_m} \frac{d\eta_m}{dw_m} - \frac{0.44}{w_m}$
- substituting in $\frac{d\eta_m}{dw_m}$ and η_m ,
- $\frac{du_m}{dw_m} = \frac{1}{w_f+w_m} + 0.005 \frac{1-a^Q}{0.55+0.005(w_m-11.6)} - \frac{0.44}{w_m}$
- The condition we are interested in checking is: $\frac{1}{w_f+w_m} + 0.005 \frac{1-a^Q}{0.55+0.005(w_m-11.6)} - \frac{0.44}{w_m} \leq 0$
- Recall we are thinking of a household with all other characteristics at average, including the female wage of $w_f = 9.87$
- $\frac{1}{9.87+w_m} + 0.005 \frac{1-a^Q}{0.55+0.005(w_m-11.6)} - \frac{0.44}{w_m} \leq 0$
- As for women, we also check this for the lower quartile, mean, and upper quartile values of the male wage
- Lower quartile wage for men: £7.51
 - $\frac{1}{9.87+7.51} + 0.005 \frac{1-a^Q}{0.55+0.005(7.51-11.6)} - \frac{0.44}{7.51} \leq 0$
 - To be consistent with bargaining theory, we require:
 - $\frac{1}{17.38} + \frac{1-a^Q}{105.91} - \frac{0.44}{7.51} > 0$
 - $a^Q < 0.89$
 - This is definitionally the case since the estimated budget share spent on private leisure is 0.44 and $0.89 + 0.44 > 1$
- Men's mean wage: £11.6
 - $\frac{1}{9.87+11.6} + 0.005 \frac{1-a^Q}{0.55} - \frac{0.44}{11.6} \leq 0$
 - To be consistent with bargaining theory, we require:

- $\frac{1}{21.47} + \frac{1-a^Q}{110} - \frac{0.44}{11.6} \leq 0$
- $a^Q < 1.95$
- This is definitionally the case
- Upper quartile wage for men: £13.1
 - $\frac{1}{9.87+13.1} + 0.005 \frac{1-a^Q}{0.55+0.005(13.1-11.6)} - \frac{0.44}{13.1} \leq 0$
 - To be consistent with bargaining theory, we require:
 - $\frac{1}{22.97} + \frac{1-a^Q}{111.5} - \frac{0.44}{13.1} \leq 0$
 - $a^Q < 2.11$
 - This is definitionally the case

G Estimation and hypothesis testing

Under the usual assumptions, our regression coefficients (α_f^0, α_m^0) are consistent estimators of their population counterparts, i.e.

$$N \rightarrow \infty : (\alpha_f^0, \alpha_m^0) \rightarrow (\mu_{\alpha 0f}, \mu_{\alpha 0m})$$

The average female share is estimated as a function of these objects:

$$\eta_f^0(\mu_{\alpha 0f}, \mu_{\alpha 0m}) = \frac{\mu_{\alpha 0f}}{\mu_{\alpha 0f} + \mu_{\alpha 0m}}$$

By the Continuous Mapping Theorem:

$$N \rightarrow \infty : \eta_f^0(\alpha_f^0, \alpha_m^0) \rightarrow \eta_f^0(\mu_{\alpha 0f}, \mu_{\alpha 0m})$$

The asymptotic distribution of the regression coefficients (α_f^0, α_m^0) is found by the Central Limit Theorem:

$$\sqrt{N} \begin{bmatrix} \alpha_f^0 - \mu_{\alpha 0f} \\ \alpha_m^0 - \mu_{\alpha 0m} \end{bmatrix} \xrightarrow{D} \mathcal{N}(0, \Sigma)$$

Where N is the sample size and $\Sigma = \begin{bmatrix} \sigma_{\alpha 0f}^2 & \sigma_{\alpha 0fm}^2 \\ \sigma_{\alpha 0fm}^2 & \sigma_{\alpha 0m}^2 \end{bmatrix}$

As long as the first two moments of this distribution are finite, and that the Jacobian of the resource share as a function of the regression coefficients exists and is non-zero valued, we can apply the delta method to find the asymptotic distribution of the estimated resource share:

$$\eta_f^0(\alpha_f^0, \alpha_m^0) \xrightarrow{D} \mathcal{N}\left(\eta_f^0(\mu_{\alpha 0f}, \mu_{\alpha 0m}), \frac{\delta \Sigma \delta'}{N}\right)$$

Where:

$$\delta = \begin{bmatrix} \frac{\delta \eta_f^0(x,y)}{\delta x} \Big|_{x=\mu_{\alpha 0f}, y=\mu_{\alpha 0m}} & \frac{\delta \eta_f^0(x,y)}{\delta y} \Big|_{x=\mu_{\alpha 0f}, y=\mu_{\alpha 0m}} \end{bmatrix} = \begin{bmatrix} \frac{\mu_{\alpha 0m}}{(\mu_{\alpha 0f} + \mu_{\alpha 0m})^2} & \frac{-\mu_{\alpha 0f}}{(\mu_{\alpha 0f} + \mu_{\alpha 0m})^2} \end{bmatrix}$$

We can construct the confidence interval as:

$$\left[\eta_f^0(\mu_{\alpha 0f}, \mu_{\alpha 0m}) \pm c \sqrt{\frac{\delta \Sigma \delta'}{N}} \right]$$

Since we do not know the population parameters of the distributions, we substitute consistent estimators for the parameters in the δ vector and in the variance-covariance matrix. Hence, we use the following estimates to construct the confidence interval:

$$\hat{\delta} = \begin{bmatrix} \frac{\delta \eta_f^0(x,y)}{\delta x} \Big|_{x=\alpha_f^0, y=\alpha_m^0} & \frac{\delta \eta_f^0(x,y)}{\delta y} \Big|_{x=\alpha_f^0, y=\alpha_m^0} \end{bmatrix} = \begin{bmatrix} \frac{\alpha_m^0}{(\alpha_f^0 + \alpha_m^0)^2} & \frac{-\alpha_f^0}{(\alpha_f^0 + \alpha_m^0)^2} \end{bmatrix}$$

And we use the variance-covariance matrix for our regression coefficients:

$$\frac{\hat{\Sigma}}{N} = \begin{bmatrix} \hat{\sigma}_{\alpha 0f}^2 & \hat{\sigma}_{\alpha 0fm}^2 \\ \hat{\sigma}_{\alpha 0fm}^2 & \hat{\sigma}_{\alpha 0m}^2 \end{bmatrix}$$

We compute the 95% confidence interval for the average female resource share as follows:

$$\left[0.4487294 \pm 1.96 \sqrt{1.971e - 05} \right] = \left[0.4400278 \quad 0.4574310 \right]$$

To estimate confidence intervals for resource shares in Stata it is possible to use the command

runcom as an alternative to coding the confidence intervals based on the methodology shown in this appendix.

H Emission multipliers

H.1 UKMRIO methodology

Take Z to be a transaction matrix, where the row $z_{i.}$ reveals output sold by industry i , and column $z_{.j}$ reveals input bought by industry i . Total output produced x_i produced by industry i can thus be expressed as

$$x_i = z_{i1} + z_{i2} + \dots + z_{ij} + y_i$$

where y_i is final output demanded by consumers, and $z_{i.}$ output sold to other industries. One can reformulate the final output of industry i as

$$x_i = a_{i1}x_1 + a_{i2}x_2 + \dots + a_{ij}x_j + y_i$$

where a_{ij} is the requirement of inputs from industry i for industry j to produce 1 pound of output. Take A to be the matrix of coefficients a_{ij} , X to be the vector of total output, y to be the vector of final consumption, and I to be the identity matrix. We can rewrite in matrix notation: $X = AX + y$. Solving for x , we get

$$X = (I - A)^{-1}y \tag{9}$$

$L = (I - A)^{-1}$ is known as the Leontief Inverse: it provides the inter-industry requirements for a given sector to produce one unit of output to final demand. Call F the vector of greenhouse gas emissions generated by each sector in a given year. The emission intensity vector e is thus:

$$e = FX^{-1}$$

and represents the emissions associated with the production of one unit of output x . Multiplying both sides of Equation 9 by the emission intensity vector e thus gives the GhG emission matrix, which provides the emissions embodied in the final consumption for each sector. Conversion multipliers used to convert expenditures from the LCF to carbon footprints are obtained as $M = eL$: this vector is a conversion factor for indirect emissions on the whole supply chain.

H.2 UKMRIO good categories

Table 18: Expenditures categories included in the UKMRIO (COICOP-level)

COICOP	Expenditure Category
1.1.1	Bread and cereals
1.1.2	Meat
1.1.3	Fish and seafood
1.1.4	Milk, cheese and eggs
1.1.5	Oils and fats
1.1.6	Fruit
1.1.7	Vegetables
1.1.8	Sugar, jam, honey, chocolate and confectionery
1.1.9	Food products n.e.c.
1.2.1	Coffee, tea and cocoa
1.2.2	Mineral waters, soft drinks, fruit and vegetable juices
10.1.1	Pre-primary and primary education
10.2.1	Secondary education
10.3.1	Post-secondary non-tertiary education
10.4.1	Tertiary education
10.5.1	Education not definable by level
11.1.1	Restaurants, cafes and the like
11.1.2	Canteens
11.2.1	Accommodation services
12.1.1	Hairdressing salons and personal grooming establishments
12.1.2	Electrical appliances for personal care
12.1.3	Other appliances, articles and products for personal care
12.3.1	Jewellery, clocks and watches
12.3.2	Other personal effects
12.4.1	Social protection
12.5	Insurance
12.6.2	Other financial services n.e.c.
12.7.1	Other services n.e.c.
13	Non-profit instns serving households
14	Central_x000D_ government
15	Local_x000D_ Authorities
16	Gross fixed_x000D_ capital_x000D_ formation
17	Valuables
18	Changes in inventories
2.1.1	Spirits
2.1.2	Wine
2.1.3	Beer
2.2.1	Tobacco

Table 18 continued from previous page

COICOP	Expenditure category
3.1.4	Cleaning, repair and hire of clothing
3.1.1	Clothing materials
3.1.2	Garments
3.1.3	Other articles of clothing and clothing accessories
3.2.1	Shoes and other footwear
3.2.2	Repair and hire of footwear
4.1.1	Actual rentals paid by tenants
4.1.2	Other actual rentals
4.2.1	Imputed rentals of owner occupiers
4.3.1	Materials for the mainenance and repair of the dwelling
4.3.2	Other services for the maintenance and repair of the dwelling
4.4.1	Water supply
4.4.2	Refuse collection
4.4.3	Sewage collection
4.4.4	Other services relating to the dwelling n.e.c.
4.5.1	Electricity
4.5.2	Gas
4.5.3	Liquid fuels
4.5.4	Solid fuels
4.5.5	Heat energy
5.1.1	Furniture and furnishings
5.1.2	Carpets and other floor coverings
5.1.3	Repair of furniture, furnishings and floor coverings
5.2.1	Household textiles
5.3.1	Major household appliances whether electric or not
5.3.2	Small electric household appliances
5.3.3	Repair of household applicances
5.4.1	Glassware,tableware and household utensils
5.5.1	Major tools and equipment
5.5.2	Small tools and miscellaneous accessories
5.6.1	Non-durable household goods
5.6.2	Domestic services and household services
6.1.1	Pharmaceutical products
6.1.2	Other medical products
6.1.3	Therapeutic appliances and equipment
6.2.1	Medical services
6.2.2	Dental services
6.2.3	Paramedical services
6.3.1	Hospital services
7.1.1	Motor cars
7.1.2	Motor cycles

Table 18 continued from previous page

COICOP	Expenditure category
7.1.3	Bicycles
7.1.4	Animal drawn vehicles
7.2.1	Spare parts and accessories for personal transport equipment
7.2.2	Fuels and lubricants for personal transport equipment
7.2.3	Maintenance and repair of personal transport equipment
7.2.4	Other services in respect of personal transport equipment
7.3.1	Passenger transport by railway
7.3.2	Passenger transport by road
7.3.3	Passenger transport by air
7.3.4	Other transport services
8.1.1	Postal services
8.2.1	Telephone and telefax equipment
8.3.1	Telephone and telefax services
9.1.1	Equipment for the reception, recording and reproduction of sound and pictures
9.1.2	Photographic and cinematographic equipment
9.1.3	Information processing equipment
9.1.4	Recording media
9.1.5	Repair of audio-visual, photographic and information processing equipment
9.2.1	Major durables for outdoor recreation
9.2.2	Musical instruments and major durables for indoor recreation
9.2.3	Maintenance and repair of other durables for recreation and culture
9.3.1	Games, toys and hobbies
9.3.2	Equipment for sport, camping and open-air recreation
9.3.3	Gardens, plants and flowers
9.3.4	Pets and related products
9.3.5	Veterinary and other services for pets
9.4.1	Recreational and sporting services
9.4.2	Cultural services
9.4.3	Games of chance
9.5.1	Books
9.5.2	Newspapers and periodicals
9.5.3	Miscellaneous printed matter
9.5.4	Stationery and drawing materials

I From Household Emissions to Total UK Emissions

Structural estimation of gender-based emission propensities within heterosexual couples indicates that achieving equal bargaining power would reduce household consumption emissions by approximately 2.1%. To assess how this impact translates to total UK emissions,

two key adjustments must be considered.

First, households are not the sole contributors to final demand emissions. Other sectors, such as non-profit institutions, central and local governments, and investment-related components, also play a role. Data from DEFRA, which provides annual breakdowns of UK emissions by final demand category, allow us to quantify the household sector's relative contribution (DEFRA 2024). On average, between 2001 and 2014, households accounted for 73.3% of total GHG direct and indirect emissions due to UK demand.

Second, the equal bargaining counterfactual is only relevant for multi-person households in which decisions are made collectively by couples. It does not apply to single-person households or households without couples. To adjust for this, we use yearly data from the Office for National Statistics (ONS) on family and household composition (ONS 2024). These figures show that, on average, 72.6% of individuals lived in multi-person households with couples over the same 2001-2014 period.

Although we only consider childless couples in the analysis, we do not exclude household with children in the re-weighting of our estimated effect. There is evidence from the household literature that couples with children tend to have lower female bargaining than those without cohabiting children (e.g. see Bargain, Donni and Hentati (2022)). Therefore, here we apply our estimates to families with children, considering this is likely to yield a conservative, lower bound, estimate.

Taking both adjustments into account, we re-weight the initial 2.1% estimated reduction in household emissions. The resulting estimate of the impact on total UK emissions is:

$$2.1\% \times 73.3\% \times 72.6\% = 1.1\%$$

This suggests that equalizing bargaining power within couples would lead to an estimated 1.12% reduction in total UK emissions, once both the sectoral scope of household demand and the relevant household types are taken into account.

J A monopolist facing linear demand

Consider the monopolist's problem in the absence of focal pricing constraints. The optimal price is $p^u = \frac{\alpha}{2\beta} + \frac{c}{2}$. Inverting this, we find that the cost associated with a certain optimal price is: $c = \frac{2\beta p^u - \alpha}{\beta}$. Therefore, the cost for which a focal price is unconstrained optimal is $\chi_i = \frac{2\beta f_i - \alpha}{\beta}$. Using the fact that focal prices are regularly spaced out at intervals G (the Regularity Condition), we substitute in $f_i = G * i$: $\chi_i = \frac{2\beta(G*i) - \alpha}{\beta}$.

Therefore, the gap between costs for which two consecutive focal prices are unconstrained optimal is:

$$\chi_{i+1} - \chi_i = \frac{2\beta(G*(i+1)) - \alpha - 2\beta(G*i) + \alpha}{\beta}$$

$$\chi_{i+1} - \chi_i = 2G$$

Now consider the monopolist's problem when constrained by focal pricing. To find the cost thresholds at which the monopolist switches from one focal price to another, we consider the following indifference condition: $\pi(f_i, t_i) = \pi(f_{i+1}, t_i)$, which can be written as

$$(\alpha - \beta f_i)(f_i - t_i) = (\alpha - \beta f_{i+1})(f_{i+1} - t_i)$$

Using the fact that focal prices are regularly spaced out at intervals G , we substitute in $f_i = G * i$:

$$(\alpha - \beta(G * i))((G * i) - t_i) = (\alpha - \beta(G * (i + 1)))((G * (i + 1)) - t_i)$$

$$\alpha = -\beta(G * i) - \beta((G * (i + 1)) - t_i)$$

$$t_i = \frac{\alpha + \beta(G * i) + \beta(G * (i + 1))}{\beta}$$

So the gap between two consecutive thresholds is:

$$t_{i+1} - t_i = \frac{\alpha + \beta(G * (i + 1)) + \beta(G * (i + 2)) - \alpha - \beta(G * i) - \beta(G * (i + 1))}{\beta}$$

$$t_{i+1} - t_i = \frac{\beta(G * (i + 2)) - \beta(G * i)}{\beta}$$

$$t_{i+1} - t_i = 2G$$

This is the same gap as the gap between costs for which consecutive focal prices are unconstrained optimal. Because there is a $2G$ gap between cost thresholds for focal prices that are G apart, pass-through is 50%, as it was in the unconstrained case.

K A monopolist facing a logarithmic demand function

A monopolist faces the following demand function:

$$p = \alpha - \beta \ln q \quad \alpha, \beta > 0, 0 < q < e^{\alpha/\beta}$$

With constant marginal cost c , the result of unconstrained optimisation is $p^u = c + \beta$, i.e. this market is characterised by a constant mark-up β and 'complete' pass-through of input costs.

Now consider the same monopolist, constrained by focal pricing. Focal price f_i is charged if $t_{i-1} < c \leq t_i$.

The thresholds are found as the points of indifference for the monopolist, i.e. $\pi(f_i, t_i) = \pi(f_{i+1}, t_i)$.

Inverting the demand function, we can write $\pi(f_i, t_i) = (f_i - t_i) e^{\frac{\alpha - f_i}{\beta}}$

We find the threshold t_i as $(f_i - t_i) e^{\frac{\alpha - f_i}{\beta}} = (f_{i+1} - t_i) e^{\frac{\alpha - f_{i+1}}{\beta}}$

Using the fact that focal prices are regularly spaced out at intervals G , we substitute in $f_i = G * i$:

$$(G * i - t_i) e^{\frac{\alpha - G * i}{\beta}} = (G * (i + 1) - t_i) e^{\frac{\alpha - G * (i + 1)}{\beta}}$$

$$(G * i - t_i) = (G * (i + 1) - t_i) e^{\frac{-G}{\beta}}$$

$$t_i \left(e^{\frac{-G}{\beta}} - 1 \right) = G \left((i + 1) e^{\frac{-G}{\beta}} - i \right)$$

$$t_i = G \frac{(i+1)e^{\frac{-G}{\beta}} - i}{e^{\frac{-G}{\beta}} - 1}$$

Hence the interval between any two consecutive cost thresholds is:

$$t_{i-1} - t_i = G \frac{(i+2)e^{\frac{-G}{\beta}} - i - 1}{e^{\frac{-G}{\beta}} - 1} - G \frac{(i+1)e^{\frac{-G}{\beta}} - i}{e^{\frac{-G}{\beta}} - 1}$$

Simplifying, we obtain: $t_{i-1} - t_i = G$

Hence, thresholds are regularly spaced at the same interval as focal prices, and the Irrelevance Theorem holds.