Purchase history and product personalization*

Laura Doval[†] Vasiliki Skreta[‡]

July 27, 2023

Abstract

Product personalization opens the door to price discrimination. A rich product line allows firms to better tailor products to consumers' tastes, but the mere choice of a product carries valuable information about consumers that can be leveraged for price discrimination. We study this trade-off in an upstream - downstream model, where a consumer buys a good of variable quality upstream, followed by an indivisible good downstream. The downstream firm's use of the consumer's purchase history for price discrimination introduces a novel distortion: The upstream firm offers a subset of the products that it would offer if, instead, it could jointly design its product line and downstream pricing. By controlling the degree of product personalization the upstream firm curbs ratcheting forces that result from the consumer facing downstream price discrimination.

KEYWORDS: *product-line design, product line pruning, price discrimination, privacy, upstream downstream interactions, dynamic mechanism design, information design, limited commitment, information externalities*

JEL CLASSIFICATION: D84, D86, L12, L13, L15

*We thank the Editor, David Myatt, and three anonymous referees for feedback that has greatly improved this paper. We also thank Nageeb Ali, Anthony Dukes, Ran Eilat, Shota Ichihashi, Nikhil Vellodi, and seminar participants in the Symposium on Communication and Persuasion, EARIE (Virtual, 2020), AMETS, Columbia University, Texas A&M, University of Essex, University College London, University of Montreal, Durham University, IO Day 2021 (New York, 2020), CEPR 2022 Workshop on Contracts, Incentives, and Information (Turin, 2022), TOI Workshop 2022 (Santiago, 2022), Workshop in Antitrust Economics (Rochester, 2023), and the 24th Economics and Computation Conference (London, 2023) for helpful comments. We owe special thanks to Maher Said for his insightful discussion at IO Day 2021. Nathan Hancart provided excellent research assistance. This research is supported by grants from the National Science Foundation (Doval: SES-2131706; Skreta: SES-1851729). Vasiliki Skreta is grateful for generous financial support through the ERC consolidator grant 682417 "Frontiers in design."

[†]Columbia University and CEPR. E-mail: laura.doval@columbia.edu

[‡]University of Texas at Austin, University College London, and CEPR. E-mail: vskreta@gmail.com

1 Introduction

"[...] versioning has the benefit of reducing concerns about inequity that arise with personalized pricing, and big data may facilitate versioning strategies based on 'mass customization' " White House report on "Big Data and Differential Pricing"

"[...] there's a positive side to all that tracking companies do, too: it allows them to customize offers that customers do want."

Randall Rothenberg (Interactive Advertising Bureau)

The trade-off between product personalization and price discrimination is at the center of the debate about the use of consumer data. Consumer data in the form of purchase histories is becoming increasingly available to firms[1] and used for personalized pricing.[2] Firms, however, also use consumer data for product personalization, allowing the firm to better meet consumers' needs, which may result in Pareto improvements (Anderson and Dana, 2009). As the opening quotes suggest, policy makers and industry practitioners note that rich product lines may compensate for the the costs of the availability of consumer data in the form of price discrimination.

Absent from this debate is that consumer data in the form of purchase histories is *endogenously* determined. On the one hand, if the consumer is aware of price discriminating practices, the data is *selected* as it reflects the consumer's trade-off between a better product match and the costs of price discrimination.[3] On the other hand, a firm designs the set the consumer chooses from and hence, how informative purchase histories may be about the consumer's preferences, taking into account the environment the firm and the consumer interact in. By selecting which products the consumer can choose from, the firm can control how much information can be gleaned about the consumer from their interaction.

In this article, we shed light on this debate by showing that the endogenous nature of consumer data in the form of purchase histories coupled with ratcheting forces

[1] For instance, Google's Gmail keeps detailed data on consumers' purchases (CNBC, 2019); brick and mortar stores track buying histories (ABCNews, 2013).

[2] Priceline acknowledges it "personalizes search results based on a user's history of clicks and purchases" (Forbes, 2014); Orbitz steers Mac users to pricier hotels (Wall Street Journal, 2012); auto dealerships tailor prices to buyers' willingness to pay, using the way they dress and the car they currently drive (Harvard Business Review, 2017); supermarkets peg prices to purchase histories (ABCNews, 2013).

[3] A report on the impact of big data on differential pricing, Executive Office of the President of the United States (2015), states *"[...] three broad trends suggest that concerns about big data and personalized pricing are not stifling consumer activity on the Internet [...]: (1) the rapid growth of electronic commerce, (2) the proliferation of consumer- empowering technologies, and (3) the slow uptake of privacy tools."*

can lead to narrow product lines. We study a dynamic mechanism design problem in the presence of limited commitment. We consider a canonical yet stylized model of upstream-downstream interaction, in which the upstream firm faces the aforementioned trade-off. On the one hand, a rich product line allows the firm to better tailor the product to the consumer's tastes. On the other hand, a rich product line creates a richer purchase history, which can be exploited by a downstream firm (either the upstream firm itself or a third party) for price discrimination. In the spirit of the ratchet effect, the consumer demands upstream rents to be compensated for downstream rent extraction. Anticipating this, the product line and hence, the consumer's choice may be distorted.

In the model, a consumer interacts with an upstream and a downstream firm over two periods. In the first period, the upstream firm chooses its product line as in Mussa and Rosen (1978): It produces a good of variable quality q_1 at quadratic cost, whereas the buyer has private information about her willingness to pay for quality, indexed by $\theta \sim U[0,1]$. (At the end of the introduction, we discuss the robustness of our results to removing the assumption that θ is uniformly distributed.) In the second period, a downstream firm sells an indivisible good at zero cost for which the buyer has a binary private value $v \in \{v_L, v_H\}$, $v_L < v_H$. Valuations are correlated over time: the consumer's type θ also parametrizes the probability that the buyer's value for the second good is v_H. We assume that the downstream firm maximizes downstream profits, whereas the upstream firm's payoff is the sum of the upstream profits and a percentage $\gamma \in [0,1]$ of downstream profits. When $\gamma \in \{0,1\}$, we span the cases in which the firms are separate or the same entities. As we explain in Section 2, we can interpret the case $\gamma \in (0,1)$ as the upstream firm obtaining a payment from the downstream firm for its use of upstream consumer data.

Preview of results: As a benchmark, we derive the solution for the case in which the upstream firm can design both the upstream product line and the downstream allocation under commitment as would be the case if the firms were (vertically) integrated. We show that the product line is determined *independently* of the downstream allocation. This is intuitive: There is no payoff-relevant link between the upstream and downstream allocations and under commitment, the upstream firm internalizes the information externalities across periods. Thus, the product line coincides with that in the static analysis of Mussa and Rosen (1978): Because of decreasing returns to quality, the firm offers a *complete* product line. Importantly, the product line coincides with that in the first best (Anderson and Celik, 2015). Furthermore, the commitment solution features either no price discrimination or *reverse* price discrimination, that is, the upstream firm offers the consumer a *discount* for quality in the downstream good.

We then characterize the upstream firm's optimal mechanism when the downstream firm observes the consumer's choice out of the product line *before* choosing the period-2 mechanism. We refer to this as *limited commitment* because the downstream mechanism must be optimal given the downstream firm's information. In other words, the downstream firm does not internalize the upstream costs of using the consumer's purchase history for pricing. Anticipating the possibility of downstream price discrimination, the upstream firm offers a narrower product line than that in the commitment solution. By curtailing the range of products it offers to the consumer, the upstream firm obfuscates how much information can be gleaned about the consumer, thereby softening price discrimination downstream.

The distortions to the product line depend on (i) the price the downstream firm would set absent consumer data and (ii) how much the upstream firm cares about the downstream profits. The first determines the *shape* of the product line by determining which products are used to convey the consumer's willingness to pay for the downstream good. Below we say that the downstream market is *premium* or *mass* depending on whether v_H or v_L is the optimal price absent consumer information. The second determines the willingness of the upstream firm to share information with the downstream firm: After all, downstream profits are maximized by tailoring prices to the consumer's information. Figure 1 depicts the consumer's choice out of the product line as a function of her type in the first best (black), second best (blue), and limited commitment (red).

Proposition 3 shows that in a premium market, the upstream firm offers a high-end product line to convey that the consumer has a high willingness to pay for quality and hence a high willingness to pay for the downstream good. In contrast to the commitment solution, low quality products are no longer offered, as Figure 1a illustrates. Moreover, because only the most exclusive products remain, the decision to purchase the good of the lowest quality no longer implies that the consumer's utility for quality is low. This guarantees that the consumer faces no price discrimination downstream, having to pay a price of v_H to obtain the downstream good. Relative to the commitment solution, the consumer faces (weakly) higher downstream prices and higher upstream prices because of the narrow product line. As a result, each consumer type is worse off under limited commitment.

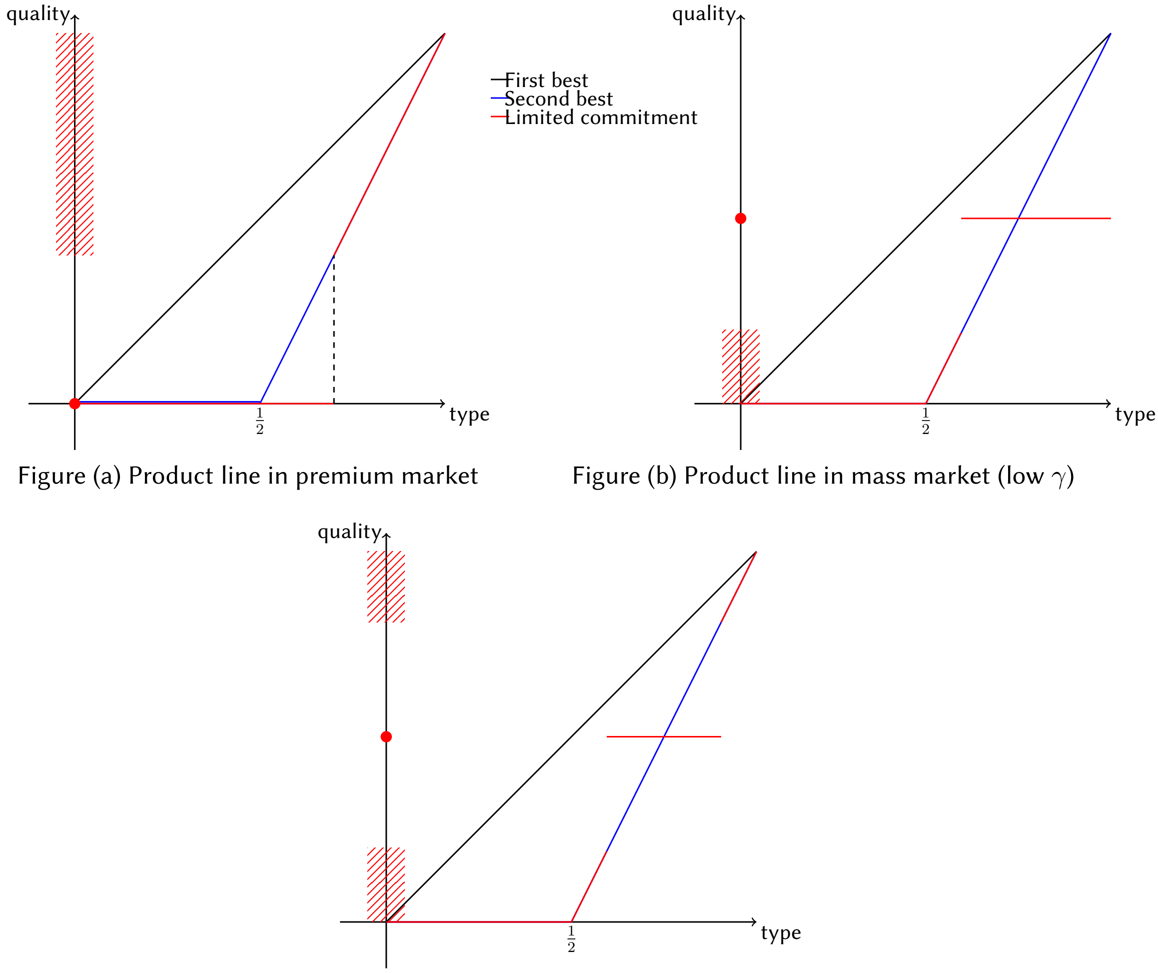

Figure (a) Product line in premium market

Figure (b) Product line in mass market (low γ)

Figure (c) Product line in mass market (high γ)

Figure 1: Product line and product choice. The product line under limited commitment is depicted in red on the y-axis.

Proposition 4 shows that in a mass market the optimal product line depends on how the upstream firm trades off upstream rents and downstream rent extraction. When γ is low, if there is price discrimination, the upstream firm would get a small share of downstream rent extraction, but would pay the costs of the rents the consumer demands. Thus, it is optimal to not allow for downstream price discrimination, which in this case corresponds to inducing a price of v_L with probability 1. To do so, the upstream firm offers either one product or a two-tier product line: A mass downstream market necessitates a product line with mass-market appeal (see Figure 1b). As γ increases, the upstream firm cares more about downstream profits, so that it internalizes the value of sharing consumer information for downstream pricing. The cost of offering a price of v_L is higher the higher is θ, as these consumer types are more likely to have high values in period 2. Thus, the upstream firm is willing to share this information. Because a high type consumer faces a high downstream price, there is no point in distorting their quality purchase and they obtain the same product as under commitment. The upstream firm then offers a three-tier product line with a distinct high-end range intended to compensate high type consumers for the downstream rent extraction, as Figure 1c illustrates. Because in a mass market there are both downward and upward distortions in quality, not all consumer types are worse off under limited commitment.

Privacy as a remedy One interpretation of our results is that absent a privacy policy that prevents the consumer's purchase history from being accessed by the downstream firm, the upstream firm then offers the consumer privacy through a coarser product line. We materialize this intuition in Section 5, where we consider the case in which the upstream firm designs both the product line *and* the information available to the downstream firm. The comparison between this *data intermediation* benchmark and the limited commitment solution separates the frictions introduced by optimal downstream pricing from those introduced by the observability of the purchase history. We show that the upstream firm offers the second best product line and offers the consumer *full* privacy. In other words, a carefully designed privacy policy alleviates product line distortions, increases upstream profits, but its welfare effects depend on whether the downstream market is premium or mass.[4] We consider other remedies to limited commitment in Section 5.

Modeling choices We conclude the introduction by discussing the role of our modeling choices; namely, the linear quadratic framework of Mussa and Rosen (1978), uniformly distributed types, and the linearity in θ of the probability that the consumer's value is v_H. Letting F_1 denote the consumer's type distribution and $p(\theta)$ denote the probability that a consumer of type θ has value v_H, we show that the solutions to the commitment and data intermediation benchmarks are qualitatively the same if we instead considered $p(\theta) = F_1(\theta)$.[5] Under this assumption, the result that both benchmarks lead to complete product lines does not rely on the linear-quadratic framework of Mussa and Rosen (1978), but rather on the log-supermodularity of the upstream social surplus (Anderson and Dana, 2009).

[4]This is consistent with Amazon's recent push to reduce the data available to third-party sellers about consumer's transactions with Amazon (see The growing customer data war).

[5]As we discuss in Section 3.2, this is the assumption behind the optimality of no/reverse price discrimination in the commitment solution.

The assumption that $p = F_1$ and θ is uniformly distributed guarantees that the posterior mean of θ is a sufficient statistic for the upstream mechanism design problem under limited commitment.[6] Indeed, by leveraging the revelation principle under limited commitment and Markov environments in Doval and Skreta (2022b), we construct the firm's optimal product line by marrying elements of mechanism design and information design. On the information design side, we rely on the techniques developed for continuum type spaces to transform the design of the product line, and hence how much information the firm learns about the consumer, into an information design problem (e.g., Gentzkow and Kamenica, 2016; Kolotilin, 2018; Dworczak and Martini, 2019; Arieli et al., 2023). On the mechanism design side, we rely on the first order approach in dynamic mechanism design and dynamic public finance to characterize the solution to a relaxed problem and then provide conditions under which the firm can implement the solution to the relaxed problem (Pavan et al., 2014; Stantcheva, 2020).

Whereas the characterization of the optimal product line under limited commitment under more general parametric specifications remains an open question, Proposition 2 characterizes the product line distortions introduced by ratcheting forces when the upstream firm offers a menu of quality-transfer pairs under the assumption that p is convex. Intuitively, a menu that induces no price discrimination must feature pooling at the bottom (if the downstream price is v_H) or at the top (if the downstream price is v_L). Importantly, we show that any incentive compatible menu that induces price discrimination downstream *must* have a product line gap – an interval of qualities that is not offered – to compensate the consumer for the forgone downstream rents. These results rely on showing that incentive compatibility together with the optimality of downstream prices imply sorting of higher types into higher qualities.

Related Literature: Our work contributes to the literatures on product line design (e.g., Mussa and Rosen, 1978; Itoh, 1983) and (intertemporal) price discrimination (e.g., Armstrong, 2006; Stokey, 1979; Hart and Tirole, 1988; Acquisti and Varian, 2005; Fudenberg and Villas-Boas, 2006), which for the most part have proceeded on separate tracks. An exception is Sun (2014), who studies a repeated version of the model in Mussa and Rosen (1978). Sun (2014) shows that offering a single variety may be optimal with binary values. Furthermore, when the firm chooses from a restricted class of mechanisms, he provides conditions under which a single variety is offered in the first period when types are drawn from a continuum: either it is optimal to offer a single variety under commitment, or the firm is patient so that it sacrifices product

[6] In fact, the main qualitative features of our results extend to more general specifications of the period-1 interaction, as long as the consumer's utility is linear in θ.

personalization today, in lieu of product personalization tomorrow. Whereas our article shares with Sun (2014) the observation that limited commitment limits varieties in the market, the results are not related otherwise: we do not restrict the set of mechanisms the firm offers to the consumer, but there is no product personalization in period 2. For this reason, it can still be optimal to offer fully personalized products in period 1 to consumers on the high-end of the type distribution in our model.

We also contribute to the literatures on multiproduct competition and behavioral product line design (e.g., Johnson and Myatt, 2003; Villas-Boas, 2004; Ellison, 2005; Kamenica, 2008; Johnson and Myatt, 2015, 2018; Xu and Dukes, 2019).[7] Johnson and Myatt (2018) and Ellison (2005) are the most relevant references. Starting from a benchmark of complete product lines absent firm competition, Johnson and Myatt (2018) characterize how cost and technological asymmetries drive product line pruning. Our model starts from the same benchmark and shows that intertemporal competition and the incentive to soften price discrimination lead to product line pruning. Ellison (2005) shows that hidden upgrade prices (add-on pricing) may lead to product proliferation when competing against a horizontally differentiated rival, when absent competition the firm would only offer one product. Whereas in Ellison (2005) there is strategic complementarity in the choice to engage in second-degree price discrimination, in our model, downstream price discrimination makes it costly to engage in second-degree price discrimination upstream.

By interpreting the upstream and downstream firm as different parties, our analysis contributes to the literature on downstream markets (e.g., Calzolari and Pavan, 2006a,b; Argenziano and Bonatti, 2020). Unlike in Calzolari and Pavan (2006a,b) and consistent with the increasing availability of consumer data, we consider the case in which the upstream firm cannot prevent the downstream firm from observing the consumer's purchase. However, our data intermediation benchmark echoes the main result in Calzolari and Pavan (2006b) that the upstream firm would prefer not to share information with the downstream firm. In Argenziano and Bonatti (2020), a downstream firm observes the quantity purchased by a consumer upstream before setting prices. The authors study the effects of this *data link* on upstream pricing decisions and evaluate the welfare implications of various privacy policies. Because there is

[7] Whereas in Johnson and Myatt (2003) product line pruning softens competition by increasing differentiation, Zhang (2011) shows that competition can also soften differentiation when price discrimination is possible. In Zhang (2011), two firms choose a location in a Hotelling line (a product) in the first period, anticipating that in the second period they can make price offers conditional on whether the consumer purchased from the firm or the rival. Zhang (2011) shows that both firms choose the same location in period one, making the decision to purchase uninformative.

no upstream product line design in their model, the ratchet effect only exacerbates downward distortions. Because these articles inherently embed a complex information feedback problem between consumer and firm's choices, like us, these articles rely on assumptions that reduce the dimension of the sufficient statistics needed to solve the problem (binary types and allocations in Calzolari and Pavan, 2006b; linear equilibrium and pricing in Argenziano and Bonatti, 2020).[8]

By showing the welfare implications of a carefully designed privacy policy, we relate to the works that study consumer privacy starting from the classic contributions of Taylor (2004) and Calzolari and Pavan (2006b).[9] A series of recent articles study the (incentive compatible) use of consumer information in static models of second-degree price discrimination (Hidir and Vellodi, 2021; Ichihashi, 2020; Eilat et al., 2021).

Finally, we contribute to the literature on dynamic mechanism design. From a methodological perspective, we contribute to the literature on limited commitment with continuum type spaces, which models mechanisms as menus (Skreta, 2006; Deb and Said, 2015; Skreta, 2015). Whereas we rely on a more general class of mechanisms thanks to the revelation principle in Doval and Skreta (2022b) for Markovian environments such as those analyzed in this article, the optimal upstream mechanism can nevertheless be implemented by a menu. Conceptually, we contribute to the literature that studies conditions under which a designer with commitment power would release exogenously available information to an agent about her type, such as Eső and Szentes (2007) and Li and Shi (2017). Instead, we study a designer's incentives under limited commitment to disclose *endogenous* information about a privately informed agent to a third party (or the designer's future self) that may use it for price discrimination.

Organization: The rest of the article proceeds as follows. Section 2 describes the model and notation. Section 3 solves two benchmarks: Section 3.1 characterizes the optimal downstream mechanism as a function of the information gleaned from the upstream interaction, whereas Section 3.2 characterizes the optimal mechanism under commitment. Section 4 derives the optimal mechanism under limited commitment and its welfare implications. Section 5 discusses different remedies to the distortions introduced by limited commitment. Section 6 concludes. All proofs are in Appendix A.

[8] Calzolari and Pavan (2006b) rely on binary types and allocations to characterize the optimal mechanism when the conditions for the optimality of full privacy do not hold and hence, the optimal mechanism must deal with a complex information feedback problem.

[9] Cummings et al. (2015) study the impact of ad targeting in monopoly pricing in a two-period model with a continuum of types in the first period and binary types in the second period.

2 Model

A consumer interacts with an upstream and a downstream firm over two periods $t \in \{1, 2\}$. The consumer is fully patient and thus, does not discount payoffs across periods.

In period 1, the upstream firm (henceforth, U-firm) produces a good of variable quality q_1 at quadratic costs, $c(q_1) = cq_1^2/2$. Upstream allocations are described by $(q_1, x_1) \in [0, Q] \times \mathbb{R}_+ \equiv A_1$, where x_1 denotes the payment from the consumer to the U-firm.

In period 2, the downstream firm (henceforth, D-firm) produces an indivisible good at 0 marginal cost. Period-2 allocations are described by $(q_2, x_2) \in \{0, 1\} \times \mathbb{R}_+ \equiv A_2$, where q_2 denotes whether the period-2 good is assigned to the consumer and x_2 denotes the payment from the consumer to the D-firm.

Profits: Downstream profits are given by the period-2 payments, x_2. Instead, upstream profits depend on the profits from the sale of the good of variable quality and downstream profits. We assume that the U-firm earns a portion $\gamma \in [0, 1]$ of the D-firm's profits. Letting (q_1, x_1) and (q_2, x_2) denote the upstream and downstream allocations, respectively, the U-firm's profits are given by $x_1 - c(q_1) + \gamma x_2$.

Consumer information and payoffs: The consumer's valuation for each of the goods is her private information. In period 1, if the consumer purchases a good of quality q_1 and pays x_1, her flow payoff is $u_1(q_1, x_1, \theta) = \theta q_1 - x_1$, where $\theta \in [\underline{\theta}, \overline{\theta}] \equiv \Theta$ denotes the consumer's type. In period 2, if she purchases the downstream good and pays x_2, her flow payoff is $u_2(q_2, x_2, v) = vq_2 - x_2$, where $v \in \{v_L, v_H\}$, $0 < v_L < v_H$. In what follows, we let Δv denote the difference $v_H - v_L$.

The consumer's period-1 type is distributed according to distribution F_1 on $[\underline{\theta}, \overline{\theta}]$. We assume that F_1 has a density $f_1 > 0$ and is such that the *virtual values*, $\hat{\theta}(F_1) = \theta - (1 - F_1(\theta))/f_1(\theta)$ are increasing in θ. In period 1, the consumer does not know her valuation for the good in period 2. Conditional on the consumer's type in period 1 being θ, her valuation in period 2 is v_H with probability $p(\theta) = F_1(\theta)$. In particular, a consumer who values quality more is more likely to value the downstream good more, that is, (θ, v) are *positively* correlated. The parametrization $p = F_1$ ensures that p is well-defined as a probability.[10] We assume that p, and hence F_1, is Lipschitz continuous.

[10] As we explain in Section 3.2, the parametrization also ensures that when $\gamma = 1$ the optimal dynamic mechanism under commitment features no price discrimination.

Finally, to characterize the optimal upstream mechanism under limited commitment, we assume in Section 4 that F_1 is the uniform distribution. As the analysis that follows makes clear, the main role of this assumption is to enable the application of the existing tools of information design with continuum type spaces, which have been developed exclusively for the case in which the sender and the receiver care only about the posterior mean (e.g., Gentzkow and Kamenica, 2016; Kolotilin, 2018; Dworczak and Martini, 2019; Arieli et al., 2023). However, as we discuss in that section, the economic forces underlying the distortions in the product line are more primitive and, as we discuss in Section 6, we expect that they will arise under more general assumptions.

Model interpretation The above is a canonical model of dynamic consumer-firm interactions and thus admits different interpretations:

One interpretation is that the U-firm and the D-firm are the same firm selling different goods over time to the consumer, as in the case of a car, followed by add-on features; a computer, followed by accessories; a hotel room, followed by amenities.[11] Under this interpretation, the parameter γ can be interpreted as weights on the profits of the different sales $1/(1+\gamma)$ on the upstream sale and $\gamma/(1+\gamma)$ on the downstream sale), or a discount factor. The most natural parametrization would be $\gamma = 1$.

Another interpretation is that the U-firm and the D-firm are different firms, in which case $\gamma = 0$ would be a natural parametrization. In that case, the downstream interaction is merely an externality to the U-firm through the consumer's ratchet effect. There are two ways to interpret the case $\gamma > 0$. First, the D-firm may pay the U-firm a commission for directing the consumer to the D-firm. Second, the D-firm may pay the U-firm a fee for its use of consumer data in the form of the purchase history.[12] The largest price that the D-firm would be willing to pay for the consumer's data is the difference in profits between accessing and not accessing this data.[13] Letting γ denote the U-firm's bargaining power vis-à-vis the D-firm, the U-firm's total profits would then consist of the profits from the sales of the product line plus a percentage γ of the difference in the D-firm's profits. Because the D-firm's profits without access to the purchase history do not depend on the upstream product line, it is without loss of generality to assume that the U-firm's payoffs are as we specified.

[11] Because there are no consumption externalities across periods, our period-2 good is closer to an accessory than a prototypical add-on.

[12] This is similar to the interpretation of the model in Calzolari and Pavan (2006b).

[13] In Section 5, we let the U-firm jointly design the product line and the data – the *cookies* – available to the D-firm. In that setting, the U-firm releases the purchase history only if it is in the U-firm's interest to do so.

2.1 Mechanisms, timing, and solution concept

We follow a mechanism design approach and thus, place no constraints on the mechanisms available to the firms. We consider constraints on the firms' ability to commit. Under *(full) commitment*, we assume that the U-firm can write long-term contracts with the consumer that, among other things, specify the terms of the downstream interaction. In particular, the downstream mechanism *need not* be optimal given the information revealed by the consumer's choice out of the product line. Instead, under *limited commitment*, we assume that the U-firm can only commit to a short-term mechanism with the consumer and conditional on the outcome of that mechanism, the D-firm offers the consumer a downstream mechanism that is optimal given the information revealed by the consumer's interaction with the upstream mechanism. The constraints on the U-firm's ability to commit can be seen as a constraint on the completeness of contracts in this setting: Under full commitment, it is as if the firms are (vertically) integrated. Instead, under limited commitment, the U-firm cannot contract on the downstream allocation and hence designs its mechanism anticipating the effect that it may have on the downstream allocation.

Given a sequence of mechanisms $\mathbb{M}_1, \mathbb{M}_2$ offered by the firms, Figure 2 summarizes the sequence of events that unfold on the consumer's side:[14]

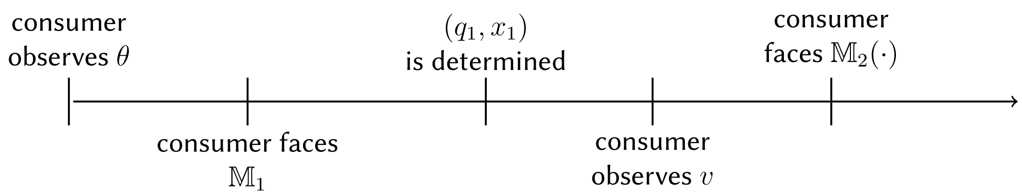

Figure 2: Timeline given a sequence of mechanisms $(\mathbb{M}_1, \mathbb{M}_2(\cdot))$

After observing her type θ, the consumer decides whether to participate in the upstream mechanism $\mathbb{M}_1$. If she does not participate, she makes no payment to the upstream firm $(x_1 = 0)$ and receives $q_1 = 0$. Instead, if she participates in the mechanism, she submits a report, which determines, among other things, the consumer's allocation at the end of period 1. The consumer then learns her value v and faces the downstream mechanism $\mathbb{M}_2$, which can depend, among other things, on the consumer's participation decision and her period-1 allocation. Given the downstream mechanism, the consumer decides whether to participate. If she does not participate, the no trade allocation, $(q_2, x_2) = (0, 0)$, obtains. Instead, if she participates, she submits a report and the downstream mechanism determines the allocation.

[14] Appendix B formally defines the mechanisms and the solution concept.

Given a sequence of mechanisms faced by the consumer, we focus on consumer participation and reporting strategies that are sequentially rational. That is, the consumer's strategy is optimal in each period given her information and the sequence of mechanisms.

3 Two benchmarks

Section 3 characterizes the solutions to two scenarios that help build intuition for the optimal mechanism under limited commitment. Section 3.1 characterizes the optimal downstream mechanism as a function of the D-firm's beliefs about the consumer's period-1 type, which is useful to understand the optimal mechanisms under commitment and limited commitment. Section 3.2 then characterizes the U-firm's optimal mechanism when it can design the upstream and downstream allocations under commitment.

3.1 Downstream pricing without product-line design

Consider the D-firm's mechanism design problem. Standard arguments imply the optimal downstream mechanism is a posted price.[15] Whether this posted price is v_L or v_H depends on the likelihood the firm assigns to the consumer's value being v_H. This likelihood, in turn, depends on the D-firm's beliefs in period 2 about the consumer's type, θ.

Letting F_2 denote the D-firm's belief about θ in period 2, the D-firm assigns probability $\mathbb{E}_{F_2}[p(\theta)] \equiv p_{F_2}$ to the consumer's value being v_H. Then, the optimal downstream price is given by:

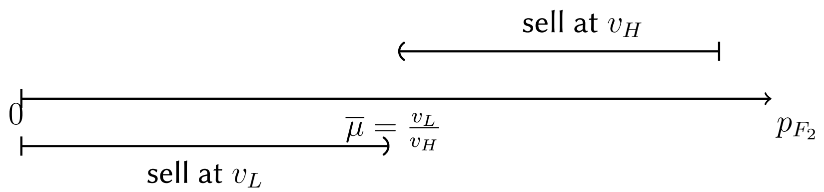

Figure 3: Optimal downstream price as a function of p_{F_2}

where $\overline{\mu} = v_L/v_H$ is the belief about v_H at which the D-firm is indifferent between

[15]This observation is immediate if the D-firm offers a mechanism that elicits v alone. The standard revelation principle, however, implies that the downstream mechanism can elicit, a priori, both θ and v. However, because θ is payoff irrelevant to the consumer, the D-firm cannot elicit it. It follows that the optimal downstream mechanism can only elicit v. A formal proof of (a more general version of) this observation can be found in Doval and Skreta (2022a,b).

selling at a price of v_H (obtaining revenue $p_{F_2}v_H$) and selling at a price of v_L (obtaining revenue of v_L). Figure 3 illustrates two important themes for what follows. First, optimal downstream pricing is sensitive to the information about θ, which gives rise to the possibility of price discrimination. Second, optimal downstream pricing only depends on the posterior mean of $p(\theta)$, p_{F_2}.

Downstream best response For future reference, we define the D-firm's best response correspondence via the probability that the D-firm serves the consumer when her value is v_L. Formally, we let $\mathcal{Q}_2^*(v_L)$ denote the set of mappings $q_2^*(v_L, \cdot) : \Delta(\Theta) \mapsto [0, 1]$ that satisfy the following:

$$q_2^*(v_L, F_2) \begin{cases} = 1 & \text{if } p_{F_2} < \overline{\mu} \\ \in [0,1] & \text{if } p_{F_2} = \overline{\mu} \\ = 0 & \text{if } \overline{\mu} < p_{F_2} \end{cases} . \tag{D-BR}$$

Downstream pricing without consumer data In what follows, the optimal downstream mechanism at the *prior* mean of $p(\cdot)$ plays a role. This describes what the D-firm would do in the absence of information and corresponds to $p_{F_1} = 1/2$. We say that the downstream market is a *premium* market when $\overline{\mu} < p_{F_1} = 1/2$ because the optimal downstream price is v_H. Instead, we say that the downstream market is a *mass* market when $\overline{\mu} > p_{F_1} = 1/2$ because the optimal downstream price is v_L.

3.2 Product-line design under commitment

As our next benchmark, we consider the case in which the U-firm can design the upstream and downstream allocations under full commitment. This allows us to study the distortions introduced to the product line when the U-firm can internalize the information externalities across periods. Because of this, we draw an analogy with models of vertical integration. The contrast between the results in this section and those in Section 4 is reminiscent of *double marginalization*: when the U-firm can internalize the information externalities across periods, the only distortions to the product line that remain are those coming from the consumer's adverse selection constraints.

Our model is a special case of the environments studied in Pavan et al. (2014), so we can rely on the standard revelation principle to characterize the U-firm's optimal mechanism (see, e.g., Myerson, 1986). Without loss of generality, the U-firm chooses a direct revelation mechanism

$$\{(q_1(\theta), x_1(\theta), q_2(\theta, v), x_2(\theta, v)) : (\theta, v) \in \Theta \times \{v_L, v_H\}\},$$

which specifies the allocation that the consumer receives in each period, as a function of the reports in each period. Importantly, when the consumer submits a type report, θ', in period 1, she restricts the menu from which she chooses in period 2, to $(q_2(\theta', \cdot), x_2(\theta', \cdot))$.

A direct revelation mechanism determines the consumer's upstream and downstream payoffs as a function of her private information, (θ, v), and her reports, (θ', v'), as follows. Her downstream payoff is given by

$$u_2(\theta', v') = vq_2(\theta', v') - x_2(\theta', v'), \tag{1}$$

whereas her upstream payoff is given by

$$W_1^C(\theta', \theta) = \theta q_1(\theta') - x_1(\theta') + p(\theta)u_2(\theta', v_H) + (1 - p(\theta))u_2(\theta', v_L). \tag{2}$$

Let $U_1^C(\theta) = W_1^C(\theta, \theta)$ denote the payoff from reporting θ truthfully in period 1.

The upstream-profit maximizing mechanism then solves

$$\max_{q_1, x_1, q_2, x_2} \int_\Theta [x_1(\theta) - c(q_1(\theta)) + \gamma\,(p(\theta)x_2(\theta, v_H) + (1 - p(\theta))x_2(\theta, v_L))]\, F_1(d\theta) \quad \text{(C-OPT)}$$

$$\text{s.t.}(\forall \theta \in \Theta)U_1^C(\theta) \geq 0 \quad \text{(C-PC)}$$

$$(\forall \theta \in \Theta)(\forall \theta' \in \Theta)U_1^C(\theta) \geq W_1^C(\theta', \theta) \quad \text{(C-TT}_1\text{)}$$

$$(\forall \theta \in \Theta)(\forall v, v' \in \{v_L, v_H\})u_2(\theta, v) \geq u_2(\theta, v'). \quad \text{(C-TT}_2\text{)}$$

That is, the U-firm's optimal mechanism must satisfy the following constraints. First, the consumer must find it optimal to participate (C-PC). Second, the consumer must find it optimal to report her type θ truthfully (C-TT$_1$). Finally, for each type report, the consumer must find it optimal to truthfully report her value (C-TT$_2$). Equation C-TT$_2$ implies that the consumer does not benefit from deviating by first misreporting θ and then misreporting v.

Proposition 1 describes the optimal mechanism under commitment:

Proposition 1 (Vertical integration). *The following is the optimal mechanism when the U-firm can design the product line and downstream allocation under commitment:*

1. ***Product line:*** *The product line is given by* $[0, \overline{\theta}/c]$*, with a type* θ*-consumer obtaining quality* $q_1(\theta) = \max\{0, \hat{\theta}(F_1)/c\}$*,*

2. ***Period 2:*** *The period-2 allocation* $q_2(\theta, \cdot)$ *is as follows:*

   (a) There is no distortion at the top, that is, for all $\theta \in [\underline{\theta}, \overline{\theta}]$*,* $q_2(\theta, v_H) = 1$*,*

   (b) Let θ_* *be such that* $F_1(\theta_*) = \max\{0, 1-2\overline{\mu}\}/(2(1-\overline{\mu})-\gamma)$*. Then,* $q_2(\theta, v_L) = \mathbb{1}[\theta \geq \theta_*]$.

   In particular, when $\overline{\mu} \geq 1/2$*, the consumer obtains the downstream good with probability 1.*

See Section A.1 for the proof. Figure 4 describes the upstream product line and the downstream probability of trade with v_L as a function of θ for different parameter values under the assumption that F_1 is the uniform distribution.

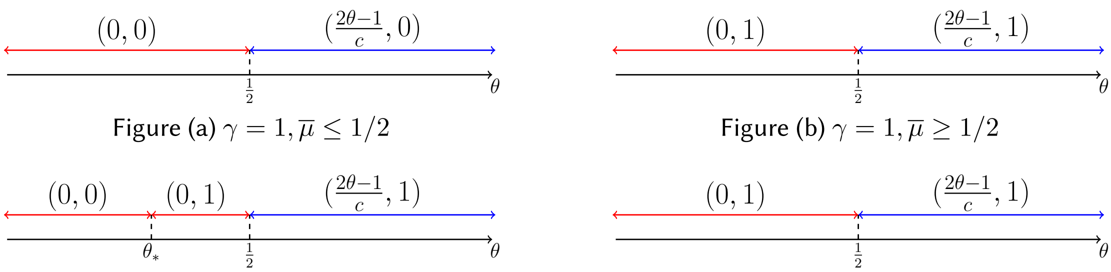

Figure (a) $\gamma = 1, \overline{\mu} \leq 1/2$

Figure (b) $\gamma = 1, \overline{\mu} \geq 1/2$

Figure (c) $\gamma = 0, \overline{\mu} \leq 1/2$

Figure (d) $\gamma = 0, \overline{\mu} \geq 1/2$

Figure 4: Upstream quality and downstream probability of serving v_L as a function of θ under uniform distribution. The top two panels depict the solution when $\gamma = 1$ and the bottom two panels depict the solution when $\gamma = 0$.

Two features of the optimal mechanism under commitment are worth highlighting. First, the product line is determined *independently* of the downstream allocation. This is illustrated in Figure 4, where across all panels the product line is the same. This is intuitive: There is no payoff-relevant link between the upstream and downstream allocations and the commitment solution internalizes the information externalities across periods. Thus, the product line coincides with that in the static analysis of Mussa and Rosen (1978): Because of decreasing returns to quality, the firm offers a *complete* product line.[16] Second, the commitment downstream allocation features either *no* price discrimination or *reverse* price discrimination, that is, consumer types that value quality more receive lower downstream prices. Indeed, in a mass downstream market ($\overline{\mu} > 1/2$), the consumer receives the good with probability 1 *independently*

[16] Johnson and Myatt (2018) refer to this as *cost driven* second-degree price discrimination.

of her upstream type. Instead, in a premium downstream market, a type θ_* exists such that the consumer is excluded downstream when her value is v_L as long as her upstream type is below θ_*. In other words, in a premium market, the consumer receives a discount for quality on the downstream good.[17]

To understand the pattern of downstream prices when the U-firm can jointly design the upstream and downstream allocations, note the following. The downstream incentive constraints (C-TT$_2$) imply the consumer needs to obtain rents so that she truthfully reveals her value. These downstream rents also determine the rents the U-firm must leave the consumer in period 1. Indeed, when a consumer of type θ reports θ', her information rents are given by

$$(\theta - \theta')q_1(\theta') + (p(\theta) - p(\theta'))(u_2(\theta', v_H) - u_2(\theta', v_L)),$$

where the dependence of u_2 on θ' represents the possibility of downstream price discrimination as a function of the reported type in period 1. That is, higher consumer types enjoy rents because they enjoy quality more and they are more likely to accrue downstream rents. Thus, the U-firm has two instruments to minimize the consumer's information rents: downward distortions in quality (familiar from Mussa and Rosen, 1978) and downward distortions in downstream rents. Thus, if downstream price discrimination exists, it is intuitive that it would be *reversed*: By giving consumers who purchase high quality goods a discount on the downstream good, the U-firm makes purchasing low quality goods less attractive. As illustrated in Figure 4c, when $\gamma = 0$ and F_1 is the uniform distribution, the consumer types who are excluded downstream are also those who receive the lowest quality good upstream.

Whereas the above explains why price discrimination, if any, should be reverse, it does not explain (i) why it may be optimal to not price discriminate, (ii) the comparison between $\overline{\mu}$ and $1/2$ in determining the optimal mechanism, and (iii) the role of γ. To understand (i), note the following. Because the U-firm contracts with the consumer in period 1, it can actually recoup part of the consumer's downstream rents. Indeed, from the perspective of period 1, what matters for rents is (i) how θ determines the willingness to pay for quality and (ii) how informative θ is about v (portion of downstream rents that go to the consumer). In particular, the U-firm can extract any portion of the consumer's downstream payoffs that is *independent* of θ. Indeed, the U-firm can always extract $u_2(\theta, v_L)$ if there is no price discrimination downstream.

[17] Airlines are a good example: A business ticket is more expensive than an economy one and provides access to free drinks and amenities that are available for a price to economy travelers.

Now, if the downstream allocation is independent of θ, then the best the firm can do is to select the allocation as a function of the prior mean of $p(\theta)$, which is $1/2$. Indeed, the *ex ante* downstream profit maximizing mechanism excludes the consumer when her value is v_L in a premium market and serves the consumer with probability 1 in a mass one. This helps explain why the comparison between $\overline{\mu}$ and the ex-ante mean of $p(\theta)$ is one of the determinants of the downstream allocation.

Finally, consider the role of γ. When γ is smaller than 1, the U-firm does not enjoy downstream profits in their entirety, but pays the costs of those profits through the consumer's period-1 rents. Thus when γ is smaller than 1, the mechanism trades off downstream profit maximization with downstream rent minimization. This effect comes through more forcefully when $\overline{\mu} < 1/2$: The incentive to diminish consumer rents together with not internalizing downstream profit losses leads the U-firm to offer downstream discounts for the purchase of high quality goods in period 1.

No price discrimination when $\gamma = 1$ When $\gamma = 1$, we obtain the stark result that there is *no* price discrimination regardless of the value of $\overline{\mu}$ (see Figures 4a and 4b). In particular, when $\gamma = 1$ and $\overline{\mu} < 1/2$, we have that $\theta_* = 1$, so in a premium market all consumer types face a price of v_H. This is a consequence of the assumption that $p(\theta) = F_1(\theta)$. Under this assumption, the *dynamic* virtual value of a consumer of type θ and value v_L is *independent* of θ when $\gamma = 1$. Thus, the decision of whether to serve v_L is independent of θ.

Commitment vs limited commitment It is immediate to see that the downstream allocation in the commitment solution is not downstream optimal given the information revealed about the consumer's type by her quality purchase upstream. This is most easily seen when F_1 is the uniform distribution, where for consumer types above $1/2$ the quality provided at the end of period 1 fully reveals the consumer's private information. For instance, consider the case in which $\overline{\mu} > 1/2$ and $\theta > \overline{\mu}$. For such a consumer, the optimal downstream price is v_H, whereas the commitment solution allocates the good to the consumer also when her value is v_L. In other words, when the U-firm can jointly design the product line and the downstream allocation under commitment, it can ignore the information revealed by the product line when choosing the downstream price.

4 Product-line design with limited commitment

Section 4 studies the properties of the U-firm-profit maximizing mechanism when the downstream mechanism is chosen based on the information provided by the consumer's choice out of the U-firm's product line. Section 4.1 applies the revelation principle in Doval and Skreta (2022b) to derive a constrained optimization problem, L-OPT, the solution to which characterizes the upstream-profit maximizing mechanism under limited commitment. We use this program to illustrate the forces that lead the U-firm to offer a reduced product line relative to the commitment solution (see Proposition 2). Section 4.2 characterizes the solution to the *relaxed* version of L-OPT under the assumption that F_1 is the uniform distribution. When the solution to the relaxed program solves L-OPT, we show that it can be implemented with the consumer choosing from a menu of quality-transfer pairs.

4.1 The upstream mechanism design problem

We set up in this section a constrained optimization problem, L-OPT, the solution to which characterizes the U-firm's optimal mechanism, and in particular, product line, under limited commitment.

Direct Blackwell mechanisms Our model is a special case of the environment studied in Doval and Skreta (2022b), so we can rely on the revelation principle in that article to characterize the U-firm's maximum payoff under limited commitment. Without loss of generality, the U-firm chooses a *direct Blackwell mechanism*, which consists of two mappings

$$\beta : \Theta \mapsto \Delta\left(\Delta(\Theta)\right),\ \alpha : \Delta(\Theta) \mapsto \Delta(A_1),$$

that specify for each consumer type θ, a Blackwell experiment $\beta(\cdot|\theta) \in \Delta\left(\Delta(\Theta)\right)$ and for each realized posterior, F_2, a distribution over allocations $\alpha(\cdot|F_2)$. Conditional on participating in a direct Blackwell mechanism, the consumer *privately* submits a type report, θ'. This determines the distribution $\beta(\cdot|\theta')$ from which a posterior F_2 is drawn. In turn, this determines the distribution $\alpha(\cdot|F_2)$ from which an allocation $(q_1, x_1) \in A_1$ is drawn. Whereas the consumer's type report is unobserved to the firms, the realized posterior and allocation are publicly observed.

Finally, Lemma 1 shows that it is without loss of generality to consider direct Blackwell mechanisms in which each posterior F_2 is mapped to one quality-transfer pair,

$(q_1(F_2), x_1(F_2))$.

Downstream information and pricing The Blackwell experiment β summarizes the information that the D-firm obtains from observing the consumer participate in the mechanism and her choice out of the upstream product line. Indeed, by the revelation principle in Doval and Skreta (2022b), it is without loss of generality to assume that when the mechanism outputs F_2, then F_2 is the downstream belief about the consumer's type. Section 3.1 characterized the optimal downstream mechanism as a function of the downstream information about the consumer's type, which is summarized by the D-firm's best response correspondence, $\mathcal{Q}_2^*(v_L)$.

Fix a D-firm's best response, $q_2^*(v_L, \cdot) \in \mathcal{Q}_2^*(v_L)$, and a posterior F_2. The D-firm's profits as a function of θ and the downstream belief F_2 are given by

$$\Pi_D^L(\theta, F_2 \mid q_2^*) = q_2^*(v_L, F_2)v_L + (1 - q_2^*(v_L, F_2))p(\theta)v_H. \tag{3}$$

Upstream mechanism design A direct Blackwell mechanism together with the downstream best response determine the consumer's upstream payoff as a function of her private information θ and her report θ' as follows:

$$W_1^L(\theta', \theta) = \int_{\Delta(\Theta)} [\theta q_1(F_2) - x_1(F_2) + p(\theta)q_2^*(v_L, F_2)\Delta v]\, \beta(dF_2|\theta'), \tag{4}$$

where $q_2^*(v_L, F_2)$ is defined in Equation D-BR. To see how Equation 4 obtains, note that in period 2, the consumer makes a positive payoff only when her valuation is v_H and the D-firm sells the good at a price of v_L, in which case, she earns $v_H - v_L \equiv \Delta v$. Let $U_1^L(\theta) = W_1^L(\theta, \theta)$ denote the payoff from truthfully reporting θ.

Theorem 1 in Doval and Skreta (2022b) implies that the U-firm's optimal mechanism solves the following constrained optimization problem:

$$\max_{\beta, q_1, x_1, q_2^* \in \mathcal{Q}_2^*(v_L)} \int_\Theta \int_{\Delta(\Theta)} \left[x_1(F_2) - c(q_1(F_2)) + \gamma \Pi_D^L(\theta, F_2 \mid q_2^*)\right] \beta(dF_2|\theta) F_1(d\theta) \tag{L-OPT}$$

$$\text{s.t. } (\forall \theta \in \Theta) U_1^L(\theta) \geq 0 \tag{L-PC}$$

$$(\forall \theta \in \Theta)(\forall \theta' \in \Theta) U_1^L(\theta) \geq W_1^L(\theta', \theta) \tag{L-TT$_1$}$$

$$(\forall \widetilde{\Theta} \subseteq \Theta)(\forall \widetilde{U} \subseteq \Delta(\Theta)) \int_{\widetilde{\Theta}} \beta(\widetilde{U}|\theta) F_1(d\theta) = \int_\Theta \int_{\widetilde{U}} F_2(\widetilde{\Theta}) \beta(dF_2|\theta) F_1(d\theta). \tag{BP}$$

That is, it is without loss of generality to restrict attention to upstream mechanisms such that the consumer participates with probability 1 (L-PC), truthfully reports her type (L-TT_1), and the Blackwell experiment β satisfies a Bayes' plausibility constraint (BP), which we explain below. To understand the right hand side of Equation L-PC, note the following. Because without loss of generality the consumer participates in the mechanism, then non-participation is an off-path event. Thus, Bayes' rule does not pin down the D-firm's beliefs about θ conditional on not participating. In particular, the D-firm can assign probability 1 to the consumer's type being $\theta = 1$ upon non-participation. Thus, offering a price of v_H in period 2 is optimal. It follows that if the consumer does not participate in the period-1 mechanism, the consumer's payoff is 0. Finally, the Bayes' plausibility constraint BP states that whenever the mechanism outputs a belief F_2, this is the belief the D-firm has about the consumer's type.

Comparison with C-OPT: We can always interpret the program L-OPT as one in which the U-firm chooses the upstream and downstream allocations as the U-firm does in C-OPT, but subject to additional constraints. As we explain next, both constraints push in the direction of product line pruning.

Sequential rationality. The first constraint is that the downstream allocation must satisfy the downstream *sequential rationality* constraint, summarized by $q_2^*(v_L, \cdot)$. The sequential rationality constraint captures that in order to maximize profits the D-firm leverages the data revealed from the upstream purchase for price discrimination. In other words, relative to the commitment solution, the U-firm chooses the product line taking into account how the information revealed through the product line affects the D-firm's optimal price in period 2. As the analysis in Section 4.2 illustrates, this constraint alone may induce product line pruning. In other words, product line pruning may obtain precisely to *avoid* price discrimination downstream (as is the case in Propositions 3 and 4, item 1). By offering a coarser product line than would be optimal, the U-firm induces more pooling of consumer types, which in turn obfuscates the information available downstream for price discrimination.

Truthtelling. Contrary to C-OPT, the consumer's truthtelling constraint, L-TT_1, is now intertwined with the downstream-sequential rationality constraints: Consumer data available to the D-firm (and hence, downstream prices) must be consistent with the consumer's choices by the Bayes' plausibility constraint, BP. At the same time, the consumer's willingness to report truthfully depends on the downstream prices by L-TT_1. Thus, the upstream product line design can no longer be separated from that of the data that is generated about the consumer downstream. If this data is used for price discrimination, the consumer will demand to be compensated upstream for the lost downstream rents. As the analysis in Section 4.2 illustrates, this compensation takes the form of a product line *gap* (see Proposition 4, item 2). Whereas the analysis in Section 4.2 relies on the assumption that θ is uniformly distributed, Proposition 2 below shows that the property that price discrimination induces a product line gap is more fundamental and holds whenever the U-firm's optimal mechanism takes the form of a *menu*.

Menu mechanisms Looking ahead, the U-firm's optimal mechanism in Section 4.2 can be implemented as a menu. That is, each consumer type is mapped to one posterior belief, F_2 and hence, each consumer type is mapped to one pair $(q_1(\theta), x_1(\theta))$. Blackwell mechanisms are richer than menus, as they allow the U-firm to obfuscate downstream information by assigning a consumer type θ to different posterior beliefs. Despite this, the optimal mechanism in Section 4.2 can be implemented as a menu, so that the analysis that follows allows us to interpret the results in that section without assuming that F_1 is the uniform distribution.

To a menu mechanism $\mathcal{M} = \{(q_1(\theta), x_1(\theta)) : \theta \in \Theta\}$ we can associate a family of posterior beliefs $\{F_2^\theta \in \Delta(\Theta) : \theta \in \Theta\}$ such that F_2^θ is the D-firm's belief about the consumer's period-1 type conditional on observing the upstream allocation $(q_1(\theta), x_1(\theta))$. This, in turn, determines the D-firm's best response $q_2^*(v_L, F_2^\theta)$. We say that the menu $\mathcal{M}$ is *incentive compatible* if given the D-firm's best response, the upstream truthtelling constraints, L-TT$_1$ hold. We say that an incentive compatible menu *induces price discrimination* if two types, θ and θ', exist such that $q_2^*(v_L, F_2^\theta) \neq q_2^*(v_L, F_2^{\theta'})$. Finally, we say that there is a *product line gap* if $\{q_1(\theta) : \theta \in \Theta\}$ is not an interval.

Proposition 2 summarizes the properties of incentive compatible menus:

Proposition 2 (Price discrimination induces a product line gap)**.** *Suppose $\mathcal{M}$ is incentive compatible. Furthermore, suppose $p'(\theta)$ is (weakly) increasing and bounded below by $G > 0$. Then, the following hold:*

1. *Downstream beliefs track qualities, that is, downstream prices depend on period-1 qualities but not on period-1 prices,*

2. *Downstream prices are monotone in the quality purchased,*

3. *If $\mathcal{M}$ induces price discrimination, then there is a product line gap. Formally, if q_1 is followed by v_L and q_1' is followed by v_H, then $q_1 + G\Delta v \leq q_1'$.*

The proof is in Section A.2.2. Whereas Proposition 2 does not depend on the definition of $p = F_1$, note that when $p = F_1$ and F_1 is the uniform distribution the assumption on p' automatically holds with $G = 1$.

Proposition 2 illustrates how the consumer's incentives lead to product line pruning: An incentive compatible menu that induces price discrimination *must* have a product line gap to compensate the consumer for the forgone downstream rents. Key to this result is the observation that higher types sort into higher qualities in an incentive compatible menu. Thus, downstream prices are increasing in the quality purchased (item 2) and in fact, independent of the upstream prices (item 1). As a consequence, if q_1 is followed by v_L and q_1' is followed by v_H, then $q_1 \leq q_1'$. Item 3 qualifies this statement by describing the extent to which q_1' must differ from q_1 if they are followed by different prices: the quality difference must account for the forgone downstream rents after the purchase of q_1'. Finally, note that Proposition 2 also has implications for the product line absent downstream price discrimination: Indeed, item 2 implies that there must be pooling at the bottom (top) whenever the downstream price is v_H (v_L).

Having understood the forces that lead to product line pruning, we now turn to the characterization of the U-firm-profit maximizing mechanism in Section 4.2.

4.2 Product line design as data design

Section 4.2 characterizes the U-firm-profit maximizing mechanism under limited commitment when F_1 is the uniform distribution relying on the first order approach in dynamic mechanism design and public finance (Pavan et al., 2014; Stantcheva, 2020). We do so by a combination of mechanism design and information design tools, which reflects the underlying theme of the article: When designing its product line – a typical mechanism design problem– the U-firm is also designing the data on the basis of which downstream prices are determined – a typical information design problem.

Relaxed program As is standard in the literature in dynamic mechanism design, we achieve this characterization by studying the solution to a relaxed version of L-OPT, which only involves the consumer's downward looking incentive constraints. In Section A.2, we show how to obtain an envelope representation of the consumer's utility, $U_1^L(\theta)$, which we use to replace the transfers out of the U-firm's profits. This representation, in turn, allows us to express the U-firm's expected profit in terms of virtual values and to reduce the upstream mechanism design problem to the problem of choosing two objects: the data available downstream, in the form of a Bayes' plausible distribution over posteriors, $\tau \in \Delta(\Delta(\Theta))$, and for each posterior, a quality level, $q_1(F_2)$. The U-firm chooses these objects to maximize the virtual surplus subject to a *monotonicity* constraint, requiring that the consumer's marginal utility

$$U_1^{L\prime}(\theta) = \int_{\Delta(\Theta)} \left[q_1(F_2) + p'(\theta)\Delta v q_2^*(v_L, F_2)\right] \beta(dF_2|\theta), \text{ is increasing in } \theta. \quad \text{(MON)}$$

The *relaxed* program corresponds to maximizing the virtual surplus with respect to the posterior distribution τ and the product line $q_1(\cdot)$ ignoring the monotonicity constraint.

To understand the role of the constraint MON and the implications of ignoring it in the analysis that follows, consider again a menu mechanism, where each consumer type θ is mapped to one posterior distribution, F_2^θ. In that case, MON is equivalent to the requirement that $q_1(F_2^\theta) + p'(\theta)\Delta v q_2^*(v_L, F_2^\theta)$ is increasing in θ. As we showed in Proposition 2, the posterior mean of p, p_{F_2}, is increasing in θ, and thus $q_2^*(v_L, F_2^\theta)$ has a *cutoff* structure, so that MON can be written as[18]

$$U_1^{L\prime}(\theta) = \begin{cases} q_1(F_2^\theta) + p'(\theta)\Delta v & \text{if } \theta \leq \hat{\theta} \\ q_1(F_2^\theta) & \text{otherwise} \end{cases}, \text{ is increasing in } \theta. \quad (5)$$

In other words, the monotonicity constraint is precisely the product line gap constraint in Proposition 2. Thus, the solution to the relaxed program may lead to a product line that is pruned less than what the monotonicity constraint would dictate (see the discussion before Proposition 4). Despite this, as the results that follow highlight, the U-firm nevertheless offers a coarser product line because the relaxed program still captures the downstream sequential rationality constraints and the consumer's downward-looking truthtelling constraints, all of which push towards product line pruning relative to the commitment solution.

Posterior means In the relaxed problem, the assumption that F_1 is the uniform distribution yields the result that the virtual surplus is a function of a low-dimensional sufficient statistic of the consumer data: the posterior mean of θ. Thus, we can solve this problem with the information design tools for continuum type spaces, which deal exclusively with the case in which the receiver's action and the sender's payoff are a function of the posterior mean (e.g., Gentzkow and Kamenica, 2016; Kolotilin, 2018;

[18] If $\hat{\theta} = \underline{\theta}$, then the downstream price is always v_H. Instead, if $\hat{\theta} = \overline{\theta}$, then the downstream price is always v_L.

Dworczak and Martini, 2019).

In what follows, we describe the solution to the relaxed problem (Propositions 3 and 4) and show that under a wide range of parameter configurations, the solution to the relaxed problem is a solution to L-OPT (Corollaries 1 and 3). The optimal mechanism turns out to be a compromise between the U-firm's desire to offer a rich product line and at the same time discipline downstream price discrimination. Except when $\overline{\mu} \leq 1/4$, the optimal mechanism distorts quality provision to discipline the revelation of information about θ across periods. How the product line is distorted and whether price discrimination arises depends on $\overline{\mu}$, which determines the downstream price in the absence of upstream information, and on γ, which determines how much the U-firm cares about downstream profits (γ) and downstream consumer payoffs $(1 - \gamma)$. The sections that follow describe the U-firm's optimal mechanism depending on whether it faces a premium or mass downstream market.

4.2.1 Low-end bundling: Premium downstream market begets a premium product line

Absent further consumer data, the D-firm would serve the consumer only when her value is v_H in a premium downstream market. Instead, except when $\gamma = 1$, the U-firm prefers to give the consumer a discount for quality, setting a downstream price equal to v_H for $\theta \leq \theta_*$ and v_L, otherwise. Unfortunately, there is no way to provide consumer data that would make it optimal to offer the discount for quality downstream.[19] As a consequence, the U-firm designs the product line so as to sustain high downstream prices whenever possible. Thus, in a premium downstream market, both the product line and the downstream allocation may be distorted relative to the commitment solution.

Figure 5 describes the product line and downstream pricing distortions in the case of a premium downstream market. In a premium market, the solution to the relaxed problem differs in whether it prevents (Figure 5a) or allows for (Figure 5b) for price discrimination. However, the product line has the same qualitative features in both cases: The U-firm provides a high-end product line to convey to the D-firm that the consumer's valuation for the downstream good is high. That is, the U-firm bundles a series of low quality upgrades into a minimum quality good that is of high enough

[19] Formally, reverse price discrimination in the commitment solution obtains from the constraint that v_L does not imitate v_H binding in the commitment solution for high consumer types. The optimality of the D-firm's mechanism implies that only the constraint that v_H does not imitate v_L binds, except when both v_H and v_L receive the same allocation.

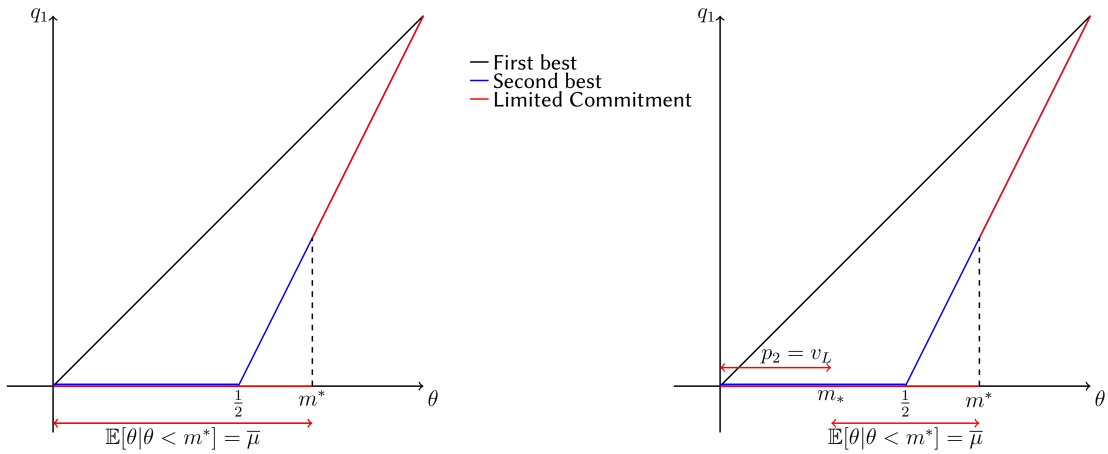

Figure (a) No price discrimination

Figure (b) Price discrimination in the low-end

Figure 5: Product choice in premium downstream market in first best (black), commitment (blue), and limited commitment (red)

quality that the consumer's purchase history does not necessarily convey that the consumer's value for quality is low.

Figure 5a illustrates the case in which the U-firm does not allow for downstream price discrimination. Thus, the minimum quality good (the one that corresponds to type m^* in the figure) is sufficiently high that all consumer types face a price of v_H downstream, regardless of their purchase history. The solution in Figure 5a obtains either when $\overline{\mu} \leq 1/4$ or when $\overline{\mu} \geq 1/4$ and the cost of personalization is high. In the first case, the U-firm offers a complete product line: By observing that the consumer buys the lowest quality good, the D-firm assigns probability $\mathbb{E}[\theta|\theta \leq 1/2] = 1/4 > \overline{\mu}$ that the consumer's value is v_H, and hence is willing to offer a high price.[20] In the second case, the high costs of personalization imply that the U-firm does not find it worth it to pay the cost of downstream rents as well. As $\overline{\mu}$ goes above $1/4$, the product line is pruned at the bottom to guarantee no price discrimination.

However, the closer $\overline{\mu}$ is to $1/2$ the more high-end the product line needs to be to prevent price discrimination in period 2. Thus, when $\overline{\mu}$ is high or the cost of product personalization is low (see Proposition 3 below), the U-firm would prefer to disclose more than just whether the consumer purchased a good of the lowest quality in period 1. Figure 5b illustrates the solution to the relaxed problem in this case: The U-firm separates the lowest consumer types that buy $q_1 = 0$ (i.e., those in $[0, m_*)$) from the

[20] However, limited commitment shapes the way information is disclosed relative to the commitment solution. Indeed, when $\overline{\mu} \leq 1/4$ there is a sense in which the commitment solution reveals "too much" information: consumer types below $1/2$ reveal θ to the mechanism which is used neither for upstream product personalization nor for downstream price discrimination.

low-to-middle consumer types that buy $q_1 = 0$, (i.e., those in $[m_*, m^*]$). The D-firm then offers the former a price of v_L and the latter a price of v_H. In turn, this allows the U-firm not to sacrifice product personalization for the high consumer types in period 1 (i.e., those above m^*), because the U-firm no longer needs to pool them with the lowest types to keep a price of v_H in period 2. Proposition 2 implies that in this case the solution to the relaxed program is not a solution to L-OPT: After all, the solution to the relaxed program is a menu mechanism that induces price discrimination *and* there is no product line gap to compensate consumer types in $[m_*, m^*]$ for the forgone rents.

Proposition 3 describes the solution to the relaxed program in a premium market and Corollary 1 describes when the solution to the relaxed program solves L-OPT. In what follows, to simplify notation we denote the product cv_H by $\tilde{c}$.

Proposition 3 (Premium downstream market). *Suppose F_1 is the uniform distribution. Let*

$$l_0(\overline{\mu}, \gamma, \tilde{c}) = (1 - \overline{\mu}) - \gamma\overline{\mu} - \frac{(4\overline{\mu} - 1)^2}{2\tilde{c}}.$$

The solution to the relaxed program is as follows:

1. *If $\overline{\mu} \leq 1/4$ or $\overline{\mu} \in (1/4, 1/2]$ and $l_0(\overline{\mu}, \gamma, \tilde{c}) \geq 0$*

   (a) ***Product line:*** *There is product line pruning at the bottom, that is, the product line is given by $[\max\{0, {}^{4\overline{\mu}-1}/_{c}\}, {}^{1}/_{c}]$, with a type θ-consumer being assigned quality equal to ${}^{2\theta-1}/_{c}$ whenever $\theta \geq \max\{2\overline{\mu}, 1/2\}$ and 0 otherwise,*

   (b) ***Period-2 pricing:*** *There is no price discrimination downstream: the price is v_H for all consumer types.*

   In particular, when $\overline{\mu} \leq 1/4$, the product line coincides with that in Proposition 1.

2. *If $\overline{\mu} \in [1/4, 1/2]$ and $l_0(\overline{\mu}, \gamma, \tilde{c}) \leq 0$, let m_*, m^* be such that $\overline{\mu} = \mathbb{E}[\theta | \theta \in [m_*, m^*]]$ and Equation 24 in Section A.2.3 holds.*[21] *Then, we have the following*

   (a) ***Product line:*** *There is product line pruning at the bottom, that is, the product line is given by $[{}^{2m^*-1}/_{c}, {}^{1}/_{c}]$, with a type θ-consumer being assigned quality equal to ${}^{2\theta-1}/_{c}$ whenever $\theta \geq m^*$ and 0 otherwise,*

   (b) ***Period-2 pricing:*** *There is price discrimination downstream: Consumer types below m_* face a price of v_L and consumer types above m_* face a price*

[21] Among the pairs m_*, m^* that satisfy $\overline{\mu} = \mathbb{E}[\theta | \theta \in [m_*, m^*]]$, Equation 26 identifies the one that is part of the solution to the relaxed program.

of v_H.

Corollary 1. *Under the assumptions of Proposition 3, the solution to the relaxed program solves L-OPT in case 1 of Proposition 3.*

Note that in case 1 there is product line pruning and no price discrimination. As anticipated in Section 4.1, the downstream sequential rationality constraint is enough to introduce distortions in the product line. In this case, the product line is pruned at the bottom relative to the commitment solution. By doing this, the U-firm forces low θ consumers to pool and hence, sustain high downstream prices.

Under the conditions of case 1 in Proposition 3, both the U-firm and the consumer lose from the U-firm's inability to control downstream actions. Relative to the commitment solution, the consumer faces either the same downstream price ($\theta \leq \theta_*$) or a higher price ($\theta > \theta_*$). Moreover, in period 1, a consumer with type below m^* receives the lowest quality good, whereas a consumer with type above m^* faces higher prices (see Figure 5a). Indeed, by pruning products from the product line, the U-firm gives the consumer fewer opportunities to self-select in period 1. Therefore, the U-firm needs to leave less rents to the consumer in period 1, and hence charges higher prices. In other words, in premium markets, limited commitment exacerbates downward distortions. Despite the possibility of charging higher prices, the U-firm is clearly worse off because it cannot implement the commitment solution. Corollary 2 summarizes this discussion:

Corollary 2 (Consumer welfare and social surplus in premium market). *Suppose the conditions in case 1 in Proposition 3 hold. Then, all consumer types are worse off under limited commitment, that is, $U_1^C(\theta) \geq U_1^L(\theta)$ for all $\theta \in [0,1]$, and upstream profits are lower. Thus, upstream social surplus is lower under limited commitment.*

4.2.2 Bundling and pruning: Mass downstream market begets a product line gap

Recall that in a mass downstream market, the D-firm would serve the consumer with probability 1 absent upstream consumer data, which is also optimal in the commitment solution. One way in which the U-firm can induce low downstream prices is to produce a *mass* product line, consisting of low to middle-range quality products. The absence of high-end products forces high θ consumers to buy mid-range products, thereby convincing the D-firm that it is facing a mass market downstream. Low downstream prices, in turn, allow the U-firm to recoup part of the downstream rents, while at the same time avoiding ratcheting forces. This, however, comes at the cost of not offering the consumer personalized products when her type is high. As we describe below, the solution in a mass downstream market depends on the weight the U-firm attaches to the cost of revealing consumer data $(1-\gamma)$.

Mid-range bundling and high-end pruning Figure 6 illustrates the optimal product line for low values of γ (see Proposition 4). The possibility of downstream price discrimination implies the consumer demands rents up front from the U-firm. When γ is low, this is very costly to the U-firm, so that the optimal product line induces low prices downstream. This is accomplished in two ways. First, the U-firm offers no high-end products (see Figure 6): The perception of a mass downstream market is sustained by a mass product line. Second, when the U-firm provides a range of low-quality products as in Figure 6b, it creates a distinct middle-quality product, by bundling a bunch of intermediate-level upgrades to just one upgrade. This allows the U-firm to engage in some second-degree price discrimination, while at the same time preventing downstream price discrimination. Because there is no price discrimination, the solution to the relaxed program is a solution to L-OPT when γ is low.

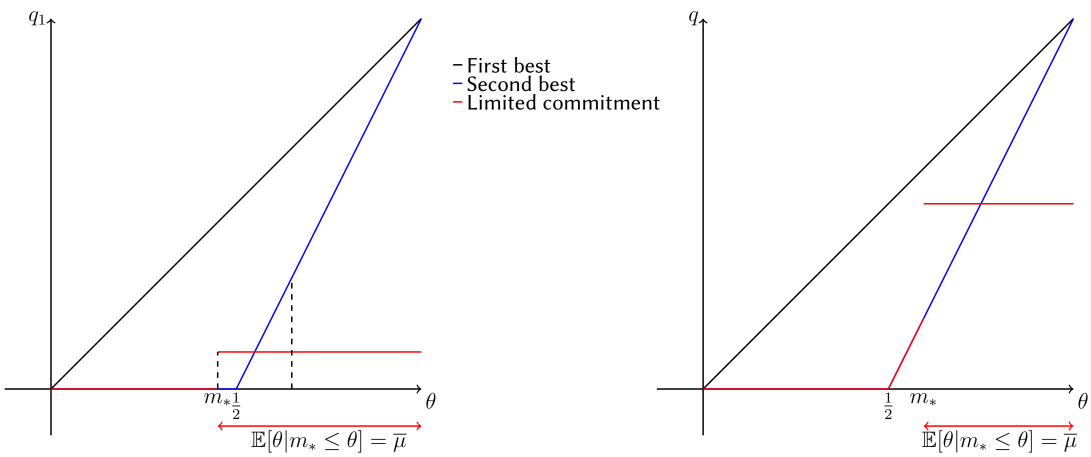

Figure (a) $\overline{\mu} \in [1/2, 3/4)$ Figure (b) $\overline{\mu} \geq 3/4$

Figure 6: Product choice in a mass market and "low" γ in first best (black), commitment (blue), and limited commitment (red)

In this case, the limited commitment solution always features a distortion at the top, whereas sometimes it does not feature distortions at the bottom (Figure 6b).

Two upgrades: mid-range bundling and high-end price discrimination As γ increases, the incentive cost induced by downstream rent extraction is compensated by downstream profits, which can be more effectively maximized with access to more detailed consumer data. The U-firm then offers a three-tier product line, with a range of low and high-end quality products, and a mid-range quality product as illustrated in Figure 7.[22] The consumer now faces downstream price discrimination: Whereas consumer types who upgrade from low to mid-quality products face a price of v_L, those that upgrade from the middle to high-quality products face a price of v_H. This has two effects on the product line: First, because consumer types who buy high-end products face downstream price discrimination, the U-firm then offers them personalized products, which fully reveal their data to the D-firm. Second, as anticipated in Proposition 2, there is a product line gap between middle and high quality products to compensate for the forgone downstream rents.

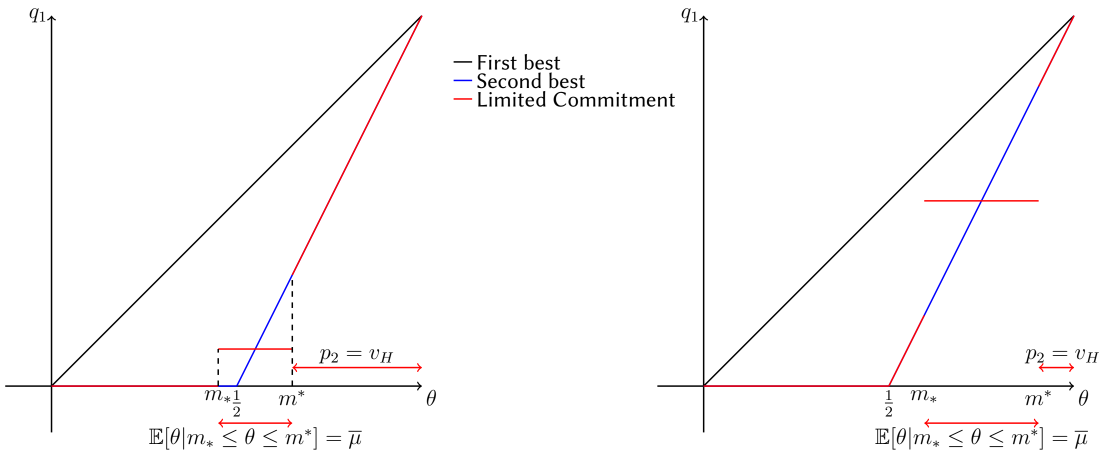

Figure (a) $\overline{\mu} \in [1/2, 3/4)$ Figure (b) $\overline{\mu} \geq 3/4$

Figure 7: Product choice in a mass market, "high" γ in first best (black), commitment (blue), and limited commitment (red)

Proposition 4 summarizes the above discussion:

Proposition 4 (Mass downstream market). *Assume F_1 is the uniform distribution and $\overline{\mu} \geq 1/2$. Let*

$$l_1(\overline{\mu}, \gamma, \tilde{c}) = (1 - \overline{\mu})(1 - \gamma) + \frac{(2\overline{\mu} - 1)^2}{\tilde{c}} - \frac{1}{2\tilde{c}},$$
$$l_2(\overline{\mu}, \gamma, \tilde{c}) = 1 - \gamma - \frac{4}{\tilde{c}}(1 - \overline{\mu}).$$

The solution to the relaxed program is as follows:

[22] Three-tier product lines are ubiquitous, with the easiest example being economy, business, first in airlines, or silver, gold, platinum in loyalty programs. For instance, Netflix has basic, standard, and premium subscriptions; Mailchimp has new business, growing business, and pro marketer; Slickplan has basic, premium, and unlimited.

1. *If either (i) $\overline{\mu} \leq 3/4$ and $l_1 \geq 0$, or (ii) $\overline{\mu} \geq 3/4$ and $l_2 \geq 0$, then*

   (a) ***Product line:*** *The upstream product line is pruned at the top and is given by*

   $$\left[0, \frac{2\min\{2\overline{\mu}-1, 1/2\}-1}{2c}\right] \cup \left\{\frac{2\overline{\mu}-1}{c}\right\},$$

   so that the U-firm offers only one quality whenever $\overline{\mu} \leq 3/4$,

   (b) ***Period-2 pricing:*** *There is no price discrimination downstream: all consumer types face a price of v_L.*

2. *Instead, if either (i) $\overline{\mu} \leq 3/4$ and $l_1 \leq 0$, or (ii) $\overline{\mu} \geq 3/4$ and $l_2 \leq 0$, then let m_*, m^* be such that $\overline{\mu} = \mathbb{E}[\theta|\theta \in [m_*, m^*]]$ and Equation 26 in Section A.2.3 holds.*[23] *We have the following:*

   (a) ***Product line:*** *There is a gap in the upstream product line, which is given by*

   $$\left[0, \frac{\max\{2m_*-1, 0\}}{c}\right] \cup \left\{\frac{2\overline{\mu}-1}{c}\right\} \cup \left[\frac{2m^*-1}{c}, \frac{1}{c}\right],$$

   (b) ***Period-2 pricing:*** *There is price discrimination at the high end of the product line: Consumer types below m^* face a price of v_L, and above m^* face a price of v_H.*

The solution to the relaxed program features price discrimination in case 2 in Proposition 4. Because the relaxed program ignores the monotonicity constraint, the product line gap due to downstream price discrimination is not necessarily enough to compensate the consumer for the forgone downstream rents, $\theta\Delta v \propto \theta(1-\overline{\mu})$. When $\overline{\mu}$ is small, these upfront rents are "tempting" for low consumer types in period 1, who may now wish to report that they value quality in period 1 more than they actually do.[24] In other words, the solution to the relaxed problem may fail to satisfy the monotonicity constraint for low values of $\overline{\mu}$.

Corollary 3 provides conditions under which the solution to the relaxed problem satisfies the monotonicity constraints in case 2 in Proposition 4. Whereas we provide the full set of conditions in the appendix, we state them only for the cases of $\gamma = 0$ and

[23] Among the pairs m_*, m^* that satisfy $\overline{\mu} = \mathbb{E}[\theta|\theta \in [m_*, m^*]]$, Equation 26 identifies the one that is part of the solution to the relaxed program.

[24] The above logic is reminiscent of the "take the money and run" strategy in Laffont and Tirole (1988).

$\gamma = 1$ for clarity.

Corollary 3. *Under the assumptions of Proposition 4, the solution to the relaxed problem is a solution to L-OPT in case 1 and whenever the following holds in case 2: Either $\gamma = 0$, or $\gamma = 1$ and $\overline{\mu} \leq 1/2 + \tilde{c}/4$.*

Figure 8 in the appendix illustrates the tuples $(\overline{\mu}, \tilde{c})$ for which the monotonicity constraint holds for $\gamma = 0$ and $\gamma = 1$ across premium and mass downstream markets.

In contrast to a premium downstream market, the implications of limited commitment for consumer welfare and social surplus are more nuanced in a mass market. The reason is that in a mass market there are both upward and downward distortions in quality and less exclusion than in the commitment solution. Serving lower consumer types has the benefit of making purchase histories less informative, at the cost of giving more consumer rents upstream in the form of more quality provision than in the commitment solution. That is, some consumer types benefit from receiving higher quality products than in the commitment solution (as is the case for types $[m_*, \overline{\mu}]$ in Figure 7), whereas those consumer types who receive lower quality products face lower prices than in the commitment solution (as is the case for types $[\overline{\mu}, m^*]$ in Figure 7). It follows that not all consumer types are worse off under limited commitment: Intuitively, high consumer types are those that may prefer the commitment solution because they face either the largest quality distortions or price discrimination. Under our parametric assumptions, we obtain that average consumer welfare may be higher under limited commitment, but there are also instances in which each consumer type may be (weakly) better off under limited commitment.

Corollary 4 (Consumer welfare in mass market)**.** *In a mass market, limited commitment has heterogeneous impact on consumer welfare. Indeed, the following hold:*

1. *If $\overline{\mu} \geq 3/4$ and $l_2 \geq 0$, then $U_1^L(\theta) \geq U_1^C(\theta)$ for all types, with the inequality being strict for $\theta > m_*$,*
2. *In all other cases, there is an intermediate range of consumer types who strictly prefer the allocation under limited commitment to that under commitment, whereas high consumer types have the opposite preference.*

The proof of Corollary 4 provides more detail about the welfare comparison across the commitment and limited commitment allocations.

We close Section 4 by discussing the implementation of the optimal mechanism and the difficulties in incorporating the monotonicity constraint:

Menu implementation: The results in Propositions 3 and 4 show that the U-firm distorts its product line relative to the commitment solution in an attempt to obfuscate how much the consumer's choice out of the product line reveals information about her preferences for quality. Indeed, as described above, the optimal product line sorts the different consumer types in (sometimes multiple) separation and pooling intervals. Despite this, whenever the monotonicity constraint MON holds, the U-firm's optimal mechanism has a *simple* implementation. Indeed, the U-firm can offer the consumer a menu of qualities and payments, such that what the D-firm learns from observing the consumer's choice from the menu coincides with the information that is induced by the optimal mechanism.

Incorporating the monotonicity constraint As discussed, the monotonicity condition MON ensures that the quality upgrade between products that induce low prices and those that induce high prices compensates the consumer for the forgone rents. However, incorporating the monotonicity constraint into the upstream mechanism design problem is not without difficulty. First, we would turn the product line design problem into a constrained information design one, as in Doval and Skreta (Forthcoming). Unfortunately, the tools developed for these kinds of problems do not readily extend to continuum type spaces. Second, we can no longer rely on the existing tools for continuum type spaces as we would lose the property that the virtual surplus depends only on the posterior mean of θ.

5 Policy implications

We briefly discuss different remedies to the product line distortions introduced by limited commitment, whose feasibility may depend on the context:

Bundling One interpretation of the commitment solution is that the sale of the upstream good is *bundled* with that of the downstream good. From this perspective, bundling the upstream and downstream allocation and pricing decisions restores commitment and eliminates the distortions in the upstream product line and downstream pricing.

Downstream competition As long as there are no exclusivity clauses or compatibility requirements with the upstream good, price competition in the downstream market may also help eliminate product line distortions. To see this, suppose that at least two firms can offer the period-2 good at 0 marginal cost and they compete in prices. In this case, there is no price discrimination in the downstream market *regardless* of how informative purchase histories are. This implies that the U-firm can implement the commitment product line. This enhances consumer welfare in a premium downstream market and has ambiguous welfare effects in a mass downstream market. Having said this, exclusivity clauses or compatibility with the upstream good may do away with the benefits of downstream competition.

Privacy design Recall that one possible interpretation of the U-firm earning downstream profits is that the downstream firm pays a fee for its use of the consumer data stemming from the purchase history. Given the distortions documented thus far, it is natural to ask whether the upstream firm wants to share information with the downstream firm and if so, whether it would be in the form of the consumer's purchase history.

Motivated by this we characterize the upstream profit-maximizing mechanism when the U-firm can design the product line together with the data available to the D-firm.[25] Importantly, the D-firm's only source of consumer data is that from the U-firm. That is, the D-firm no longer observes the consumer's purchase out of the product line. Thus, the U-firm acts as a *data intermediary* between the consumer and the D-firm.

In this data intermediation benchmark, the U-firm's optimal mechanism features no product line distortions, full consumer privacy, and no price discrimination. Proposition 5 states this formally:

Proposition 5 (Upstream data design). *Under upstream product and data design, the upstream profit maximizing mechanism is as follows:*

1. ***Product line:*** *The product line is given by* $[0, \overline{\theta}/c]$, *with a type* θ*-consumer obtaining quality* $q_1(\theta) = \max\{0, \hat{\theta}(F_1)/c\}$,
2. ***Period 2:*** *There is no downstream price discrimination:*
   (a) *In a premium market, all consumer types face a price of* v_H,
   (b) *In a mass market, all consumer types face a price of* v_L.

The proof is in Section A.3. When the U-firm designs the data the D-firm has access to, it preserves full consumer privacy. Thus, the D-firm cannot engage in behavior-based price discrimination. Because the U-firm can conceal consumer choices, it offers a

[25] This is analogous to the setting in Calzolari and Pavan (2006b), but without perfectly persistent types.

complete product line. In other words, the U-firm utilizes the consumer data to engage in second-degree price discrimination and does not share this data downstream to be used for price discrimination.

We note the following three features of the optimal mechanism in Proposition 5. First, relative to the commitment solution, the U-firm loses the ability to engage in reverse price discrimination. This explains why this feature is absent under limited commitment as well. Second, recall the interpretation of the U-firm's profits as arising from the D-firm paying for access to the consumer's purchase history. Because no value of γ exists such that the U-firm shares consumer data with the D-firm, it follows that the U-firm would not sell consumer data to the D-firm, even if the U-firm could extract the entirety of the downstream profits.[26] Finally, relative to the limited commitment solution, the U-firm benefits from a carefully designed privacy policy, but whether this benefits the consumer depends on whether the downstream market is premium or mass. In particular, by implementing the full product line, the U-firm excludes more consumers relative to the limited commitment solution in a mass market. This illustrates the nuanced effects of privacy policies: They may neither harm firms, nor benefit consumers. Moreover, because the welfare effects of limited commitment are not uniform across consumer types, it is not immediate that giving the consumer the choice to reveal her purchase history to the D-firm would be a solution.

6 Concluding remarks

Our results provide an additional perspective on the availability of consumer information and its use for price discrimination (Bergemann and Ottaviani, 2021). Economists usually highlight that the consumer data available to the firms is endogenous because it is the result of consumer choices *given* the choice sets offered by the firm. Our results complement this view by highlighting another way in which this data is endogenous: the choice set the consumer faces is *chosen* by the firm. It is precisely because we study the *joint* determination of the firm's product line together with the consumer's behavior within the set of products offered by the firm (and its informational impact

[26] Proposition 5 echoes the main theorem in Calzolari and Pavan (2006b). In a persistent-type setting, they provide sufficient conditions under which the upstream firm may choose not to share consumer information downstream. Whereas we consider non-perfectly persistent types, the conditions under which privacy is optimal in their setting also hold in our model. Namely, the consumer's and U-firm's payoffs are additively separable in the period-1 and period-2 allocations, and there is positive correlation in the valuations. In Calzolari and Pavan (2006b), positive correlation is not in the statistical sense, but rather the mechanism design sense: It means that the set of binding incentive constraints is the same in the upstream and downstream interaction.

on period-2 pricing) that we uncover a distortion in the product line on top of the one driven by rent extraction. This perspective puts at the forefront a concern in empirical work about the endogeneity of the set of options consumers face (e.g., Ivaldi and Martimort, 1994; Miravete, 2002; Luo et al., 2018), making salient the possibility that this endogeneity may be pervasive in settings where firms and consumers interact repeatedly over time.

Our work opens up several avenues for future research. First, the strength of the ratchet effect, and hence the extent to which the U-firm may prune its product line, depends on how forward looking consumers are. Considering the effects of having consumers of different levels of sophistication in the U-firm's optimal mechanism would be interesting. Second, one interpretation of our results is that the U-firm optimally responds to ratcheting forces by choosing a mechanism that is less "history dependent." Thus, our framework and results could be used as a first step to understand the costs and benefits of making procurement contracts sensitive to past performance. In a setting with moral hazard, Decarolis et al. (2016) illustrate the benefits of taking into account past performance in awarding procurement contracts. Instead, our results suggest that in settings with adverse selection procurement contracts should be less sensitive to past performance.

Finally, whereas our analysis relies on a number of parametric assumptions whose only role is to allow us to apply the existing methodology of information design for continuum type spaces, the economic force underlying the optimal product line is likely to extend to more general settings. Indeed, an interpretation of our model is that, faced with the dynamic inconsistency of the commitment solution, the U-firm prefers to acquire less information about the consumer as a (self-)disciplining device, as in Carrillo and Mariotti (2000). Even if it is natural to conjecture that this economic force extends to more general settings, showing this formally requires extending the existing toolkit of information design with continuum type spaces. Because it will enable a deeper exploration of the issues raised in this article and open the analysis of new problems, we see this extension as a fruitful avenue for further research.

A Omitted proofs

Remark 1. *Throughout, we make the following technical assumptions. Unless noted otherwise, all spaces are Polish spaces; we endow them with their Borel σ-algebra. Second, product spaces are endowed with their product σ-algebra. Third, for a Polish space X, we let $\Delta(X)$ denote the set of Borel probability measures over X, endowed with the* *weak* topology. Thus, $\Delta(X)$ is also Polish (Aliprantis and Border, 2013). Finally, for any two measurable spaces X and Y, a mapping $\varphi : X \mapsto \Delta(Y)$ is a* transition probability *from X to Y if, for any measurable $C \subseteq Y$, $\varphi(C|x) \equiv \varphi(x)(C)$ is a measurable real valued function of $x \in X$.*

A.1 Proofs of Section 3.2

Proof of Proposition 1. Let $\Delta_v u_2(\theta) = u_2(\theta, v_H) - u_2(\theta, v_L)$. Equation C-TT$_2$ implies a monotonicity condition for $\Delta_v u_2(\theta)$, that we use later on:

$$\Delta v q_2(\theta, v_L) \leq \Delta_v u_2(\theta) \leq \Delta v q_2(\theta, v_H). \tag{6}$$

Furthermore, because $p(\cdot)$ is differentiable, we can apply the envelope theorem in Milgrom and Segal (2002) to obtain the following envelope condition from Equation C-TT$_1$:

$$U_1^{C\prime}(\theta) = q_1(\theta) + p'(\theta)\Delta_v u_2(\theta), \tag{C-E}$$

Equation C-E delivers the following expression for the period-1 transfers

$$x_1(\theta) = \theta q_1(\theta) + p(\theta)\Delta_v u_2(\theta) + u_2(\theta, v_L) - \int_{\underline{\theta}}^{\theta} \left(q_1(s) + p'(s)\Delta_v u(s)\right) ds.$$

Replacing this in Equation C-OPT, we obtain:

$$\Pi_U^C = \int_\Theta \left[q_1(\theta)\hat{\theta}(F_1) - c(q_1(\theta)) + \gamma \Pi_D^C(\theta) + (1-\gamma)W_D^C(\theta)\right] F_1(d\theta), \tag{C-VS}$$

where $\Pi_D^C(\theta)$ denotes the downstream profits adjusted by the dynamic information rents

$$\Pi_D^C(\theta) = p(\theta)v_H q_2(\theta, v_H) + (1 - p(\theta))q_2(\theta, v_L) - p'(\theta)\frac{(1 - F_1(\theta))}{f_1(\theta)}\Delta_v u_2(\theta), \tag{7}$$

whereas $W_D^C(\theta)$ is a measure of the consumer's downstream payoffs in terms of the dynamic rents:

$$W_D^C(\theta) = u_2(\theta, v_L) + \left(p(\theta) - \frac{1 - F_1(\theta)}{f_1(\theta)}p'(\theta)\right)\Delta_v u_2(\theta). \tag{8}$$

We now show that the mechanism described in Proposition 1 maximizes the objective in C-VS. Because it satisfies all constraints, it follows that it is the optimal mechanism.

Pointwise maximization of the objective in C-VS with respect to $q_1(\theta)$ delivers that the product line described in Proposition 1 is optimal.

We now show the properties of the period-2 allocation. Note that the term multiplying $\Delta_v u_2(\theta)$ in C-VS is given by

$$\left[(1-\gamma)p(\theta) - p'(\theta)\frac{1-F_1(\theta)}{f_1(\theta)}\right] = (2-\gamma)F_1(\theta) - 1, \tag{9}$$

and note that this is increasing in θ, negative for $\theta = \underline{\theta}$ and positive for $\theta = \overline{\theta}$. Define θ_γ to be such that $F_1(\theta_\gamma) = 1/(2-\gamma)$. Applying Equation 6, we obtain the following implication on $\Delta_v u_2(\theta)$:

$$\Delta_v u_2(\theta) = \begin{cases} \Delta v q_2(\theta, v_L) & \text{if } \theta \leq \theta_\gamma \\ \Delta v q_2(\theta, v_H) & \text{if } \theta > \theta_\gamma \end{cases}, \tag{10}$$

which we can replace in Equation C-VS. For $\theta \leq \theta_\gamma$, we obtain that

$$\begin{aligned} &\gamma \Pi_D^C(\theta) + (1-\gamma)W_D^C(\theta) \\ &= \gamma p(\theta) v_H q_2(\theta, v_H) + q_2(\theta, v_L) v_H \left(2\overline{\mu} - 1 + F_1(\theta)(2 - 2\overline{\mu} - \gamma)\right) - (1-\gamma)x_2(\theta, v_L). \end{aligned} \tag{11}$$

Clearly, $q_2(\theta, v_H) = 1$ and $x_2(\theta, v_L) = 0$. When $\overline{\mu} < 1/2$, a threshold type $\theta_* \leq \theta_\gamma$ exists such that $F_1(\theta_*) = (1-2\overline{\mu})/(2(1-\overline{\mu})-\gamma)$ and it is optimal to set $q_2(\theta, v_L) = \mathbb{1}[\theta \geq \theta_*]$.

Instead, when $\overline{\mu} \geq 1/2$, the term multiplying $q_2(\theta, v_L)$ is (weakly) positive for all $\theta \leq \theta_\gamma$, in which case, $q_2(\theta, v_L) = \mathbb{1}[\theta \geq 0]$.

Consider now $\theta > \theta_\gamma$. We have that

$$\begin{aligned} \gamma \Pi_D^C(\theta) + (1-\gamma)W_D^C(\theta) &= q_2(\theta, v_L) v_L \left(\gamma(1 - F_1(\theta)) + 1 - \gamma\right) - (1-\gamma)x_2(\theta, v_L) \\ &= q_2(\theta, v_H) v_H \left[F_1(\theta)(\gamma(2-\overline{\mu}) + 2(1-\gamma)(1-\overline{\mu})) - (1-\overline{\mu})\right]. \end{aligned} \tag{12}$$

In this case, $q_2(\theta, v_L) = 1$ and $x_2(\theta, v_L) = 0$ for all $\theta > \theta_\gamma$. Furthermore, the term multiplying $q_2(\theta, v_H)$ is increasing in θ and positive for $\theta \geq \theta_\gamma$ so that $q_2(\theta, v_H) = 1$ for all $\theta > \theta_\gamma$. □

Observation 1 (Implementability). *The solution to the relaxed program can be implemented. Indeed, it is possible to show that the monotonicity constraint ($U_1^{C'}(\theta)$ increasing in θ) is satisfied.*

A.2 Proofs of Section 4

A.2.1 Proofs of Section 4.1

Given a mechanism $\langle\beta,\alpha,x_1\rangle$, denote the average quality conditional on F_2 by $q_1(F_2)=\int_0^Q q_1\alpha(dq_1|F_2)$. Lemma 1 shows that it is always payoff-improving for the U-firm (and payoff irrelevant for the consumer) to have the mechanism induce $q_1(F_2)$ with probability 1:

Lemma 1 (Deterministic q_1 conditional on F_2). *Let $\mathbb{M}=\langle\beta,\alpha,x_1\rangle$ denote a mechanism that satisfies L-PC, L-TT$_1$, BP. Then, $\mathbb{M}'=\langle\beta,\alpha',x_1\rangle$ such that $\alpha'(\cdot|F_2)$ assigns probability 1 to $q_1(F_2)$ satisfies L-PC, L-TT$_1$, BP and is preferred by the U-firm.*

Proof of Lemma 1. The proof follows from the expressions for upstream profits in L-OPT and consumer payoffs in Equation 4. Because downstream information does not depend on the allocation rule α, $\mathbb{M}'$ induces the same downstream response as $\mathbb{M}$. Note that the U-firm's profits from $\mathbb{M}$ are given by:

$$\begin{aligned}&\int_\Theta\int_{\Delta(\Theta)}\int_0^Q\left[x_1(F_2)-c(q_1)+\gamma\left(q_2^*(v_L,F_2)v_L+(1-q_2^*(v_L,F_2))p(\theta)\Delta v\right)\right]\alpha(dq_1|F_2)\beta(dF_2|\theta)F_1(d\theta)\\&=\int_\Theta\int_{\Delta(\Theta)}\left[x_1(F_2)-\int_0^Q c(q_1)\alpha(dq_1|F_2)+\gamma\left(q_2^*(v_L,F_2)v_L+(1-q_2^*(v_L,F_2))p(\theta)\Delta v\right)\right]\beta(dF_2|\theta)F_1(d\theta)\\&\leq\int_\Theta\int_{\Delta(\Theta)}\left[x_1(F_2)-c(q_1(F_2))+\gamma\left(q_2^*(v_L,F_2)v_L+(1-q_2^*(v_L,F_2))p(\theta)\Delta v\right)\right]\beta(dF_2|\theta)F_1(d\theta),\end{aligned}$$

where the latter are the profits from $\mathbb{M}'$. The second equality uses that conditional on F_2 the only element that depends on q_1 are the costs and the inequality in the third line follows from convexity of the costs (and it is strict whenever $\operatorname{supp}\,\alpha(\cdot|F_2)$ has at least two elements. The consumer's payoffs from reporting θ' when her type is θ under $\mathbb{M}$ are given by:

$$\begin{aligned}&\int_{\Delta(\Theta)}\int_0^Q\left[\theta q_1-x(F_2)+p(\theta)\Delta v q_2^*(v_L,F_2)\right]\alpha(dq_1|F_2)\beta(dF_2|\theta')\\&=\int_{\Delta(\Theta)}\left[\theta q_1(F_2)-x(F_2)+p(\theta)\Delta v q_2^*(v_L,F_2)\right]\beta(dF_2|\theta'),\end{aligned}$$

where the latter are the payoffs from reporting θ' under $\mathbb{M}'$. The result follows. □

A.2.2 Proof of Proposition 2

The proof of Proposition 2 follows from three lemmas, which we state and prove below. Given the menu $\{(q_1(\theta),x_1(\theta)):\theta\in\Theta\}$, recall that $F_2^\theta\equiv F_2^{(q_1(\theta),x_1(\theta))}$ denotes the downstream belief when $(q_1(\theta), x_1(\theta))$ is the choice out of the menu.

Lemma 2 (Downstream beliefs track qualities). *Let $\mathcal{M}$ be incentive compatible. Let $\theta < \theta'$ and suppose that $q_1(\theta) = q_1(\theta')$. Then, $q_2^*(v_L, F_2^{\theta}) = q_2^*(v_L, F_2^{\theta'})$.*

Proof of Lemma 2. Towards a contradiction, suppose that $q_2^*(v_L, F_2^{(q_1(\theta),x_1(\theta))}) \neq q_2^*(v_L, F_2^{(q_1(\theta'),x_1(\theta'))})$. Then, it must be that case that $x_1(\theta) \neq x_1(\theta')$ as downstream prices are different after the allocations of θ and θ'. Because the menu is incentive compatible, the following holds:

$$(p(\theta') - p(\theta))\Delta v q_2^*(v_L, F_2^{\theta'}) \geq x_1(\theta') - x_1(\theta) \geq (p(\theta') - p(\theta))\Delta v q_2^*(v_L, F_2^{\theta}),$$

which means that $q_2^*(v_L, F_2^{\theta'}) = 1 > 0 = q_2^*(v_L, F_2^{\theta})$. Therefore, we must have

$$\mathbb{E}\left[p|F_2^{\theta}\right] \geq \overline{\mu} \geq \mathbb{E}\left[p|F_2^{\theta'}\right]. \tag{13}$$

There are two possibilities:

1. $(\exists \theta_+ < \theta')$ such that $(q_1(\theta_+), x_1(\theta_+)) = (q_1(\theta'), x_1(\theta'))$, or
2. $(\exists \theta^+ > \theta)$ such that $(q_1(\theta^+), x_1(\theta^+)) = (q_1(\theta), x_1(\theta))$.

Otherwise, simultaneous violation of both of the above implies that $\mathbb{E}[p|F_2^{\theta}] < \mathbb{E}[p|F_2^{\theta'}]$, a contradiction.

Suppose that item 1 holds. Incentive compatibility implies that *any* such θ_+ must satisfy that $\theta_+ > \theta$. Hence, $\mathbb{E}[p|F_2^{\theta'}] \geq p(\theta)$. It follows that item 2 must also hold. However, incentive compatibility implies that *any* such θ^+ must satisfy that $\theta^+ < \theta'$. In fact, any such θ^+ must be less than any $\theta_+ < \theta'$ that receives θ''s allocation. It follows that Equation 13 cannot hold. □

Because of Lemma 2, below we write $\mathbb{E}[p|q_1(\theta)]$ instead of $\mathbb{E}[p|F_2^{\theta}]$.

Lemma 3 (Prices are monotone in qualities). *An incentive compatible menu cannot have two qualities $q_1 < q_1'$ such that $q_2^*(v_L, q_1) = 0 < q_2^*(v_L, q_1') = 1$.*

Proof of Lemma 3. Towards a contradiction, suppose this is the case. Let θ' denote the consumer type that chooses q_1' and let θ denote the consumer type that chooses q_1. Note that we must have that $\mathbb{E}[p|q_1'] \leq \overline{\mu} \leq \mathbb{E}[p|q_1]$. Now, incentive compatibility implies that

$$(\theta' - \theta)(q_1' - q_1) + (p(\theta') - p(\theta))\Delta v \geq 0. \tag{14}$$

Note that we must have that $\theta' > \theta$. Otherwise, because $q_1' > q_1$ and p is increasing, Equation L-E cannot hold. Then, as in the proof of Lemma 2, we must have that either

1. $(\exists \theta_+ < \theta')$ such that $q_1(\theta_+) = q_1'$,
2. $(\exists \theta^+ > \theta)$ such that $q_1(\theta^+) = q_1$.

Suppose that item 1 holds. Note that for any such θ_+, incentive compatibility implies that $\theta < \theta_+$. Thus, we have that $\mathbb{E}[p|q_1'] \geq p(\theta)$ and hence $\mathbb{E}[p|q_1] \geq p(\theta)$.

There are two possibilities. If θ is the only type that purchases q_1, then we have $p(\theta) = \mathbb{E}[p|q_1] \geq \overline{\mu} \geq \mathbb{E}[p|q_1'] \geq p(\theta)$, contradicting that $q_2^*(v_L, q_1) \neq q_2^*(v_L, q_1')$. Otherwise, there exists $\theta^+ > \theta$ that satisfies item 2. Incentive compatibility then implies that $\theta^+ < \theta'$. As before, it must be that $\theta^+ < \theta_+$ for any θ_+ that satisfies item 1. It follows that we cannot have that $\mathbb{E}[p|q_1'] \leq \mathbb{E}[p|q_1]$. □

We have the following corollary:

Corollary 5 (Cutoff). *Let (q_1, x_1, q_2^*) be an incentive compatible menu. Then, there exists a cutoff quality $q_1^\star$ such that*

1. $q_2^*(v_L, q_1) = 1$ *if* $q_1 < q_1^\star$,
2. $q_2^*(v_L, q_1) = 0$ *if* $q_1 > q_1^\star$.

Lemma 4 (Price discrimination induces gap). *If q_1 and q_1' are part of an incentive compatible menu and $q_1 < q_1'$ and $q_2^*(v_L, q_1) = 1 > q_2^*(v_L, q_1') = 0$, then $q_1' \geq q_1 + G\Delta v$.*

Proof of Lemma 4. Towards a contradiction, suppose that $q_1' > q_1$, $q_2^*(v_L, q_1') = 0 < q_2^*(v_L, q_1) = 1$, but $q_1' < q_1 + \Delta v$. Let θ' denote the type that receives q_1' and θ denote the type that receives q_1. Suppose that $q_1' < q_1 + G\Delta v$. Incentive compatibility implies that:

$$(\theta' - \theta)(q_1' - q_1) - (p(\theta') - p(\theta))\Delta v \geq 0. \tag{15}$$

We want to show that it cannot be that $\theta' > \theta$. Indeed, suppose that $\theta' > \theta$, then

$$
\begin{aligned}
(\theta' - \theta)(q_1' - q_1) - (p(\theta') - p(\theta))\Delta v &< (\theta' - \theta)G\Delta v - (p(\theta') - p(\theta))\Delta v \\
&= (\theta' - \theta)\Delta v \left(G - \frac{p(\theta') - p(\theta)}{\theta' - \theta} \right) \leq 0, \qquad (16)
\end{aligned}
$$

where the first inequality follows from $q_1' < q_1 + G\Delta v$ and the last one follows from convexity of p and the assumption that $p' \geq G$. Thus, if Equation 15 holds, then it must be that $\theta' < \theta$.

To complete the proof of Lemma 4, we need to show that all types $\tilde{\theta} < \theta'$ prefer $(q_1', q_2^*(v_L, q_1'))$ to $(q_1, q_2^*(v_L, q_1))$ and all types $\tilde{\theta} > \theta$ have the opposite preference. Let $\tilde{\theta} < \theta'$. $\tilde{\theta}$ prefers $(q_1', q_2^*(v_L, q_1'))$ over $(q_1, q_2^*(v_L, q_1))$ if the following holds:

$$\tilde{\theta}(q_1' - q_1) - p(\tilde{\theta})\Delta v \geq x_1' - x_1,$$

where x_1 is the payment associated to q_1 and x_1' is the payment associated to q_1'. Note that we have that

$$\theta'(q_1' - q_1) - p(\theta')\Delta v \geq x_1' - x_1,$$

so that it would suffice that we show

$$\tilde{\theta}(q_1' - q_1) - p(\tilde{\theta})\Delta v \geq \theta'(q_1' - q_1) - p(\theta')\Delta v \Leftrightarrow (\theta' - \tilde{\theta})(q_1' - q_1) - (p(\theta') - p(\tilde{\theta}))\Delta v \leq 0.$$

Note that this follows from $q_1' < q_1 + G\Delta v$ as in Equation 16. Similar steps show that $\tilde{\theta} > \theta$ prefers $(q_1, q_2^*(v_L, q_1))$ to $(q_1', q_2^*(v_L, q_1'))$.

Because all types $\tilde{\theta} < \theta'$ prefer q_1' (and the implied period-2 price) to q_1 and all types $\tilde{\theta} > \theta$ prefer q_1 (and the implied period-2 price) to q_1', it follows that in order for the condition on the period-2 prices to hold, we must have that either there is $\theta^+ > \theta'$ that chooses q_1' or $\theta_+ < \theta$ that chooses q_1. It must be that $\theta' < \theta^+ < \theta_+ < \theta$ for any such types. Indeed, for $\tilde{\theta} \in [\theta', \theta]$ write

$$u(\tilde{\theta}) = \tilde{\theta}q_1' - x_1' + \max\{0, \tilde{\theta}(q_1 - q_1') + p(\tilde{\theta})\Delta v + x_1 - x_1'\}.$$

Note that the second term is continuous and increasing in $\tilde{\theta}$. By assumption it is positive for $\tilde{\theta} = \theta$ and it is zero or negative when $\tilde{\theta} = \theta'$. Thus, there is $\theta^\star \in [\theta', \theta]$ such that for θ above $\theta^\star$, q_1 is the preferred quality and for $\theta < \theta^\star$, q_1' is the preferred quality. We conclude again that it cannot be the case that $\mathbb{E}[p|q_1'] \geq \overline{\mu} \geq \mathbb{E}[p|q_1]$. $\square$

A.2.3 Proofs of Section 4.2

Envelope representation of payoffs: We first argue that the consumer's period-1 payoff defined in Section 4.1 is Lipschitz continuous, and hence almost everywhere differentiable. To see this, consider the payoff from the following deviation: The consumer with type θ reports θ' and then follows the strategy of θ' in period 2. Her payoff would then be given by $W_1^L(\theta', \theta)$ as defined in Equation L-TT$_1$.

The optimality of truthtelling implies

$$U_1^L(\theta) = \max_{\theta' \in \Theta} W_1^L(\theta', \theta).$$

We now establish that the family $\{W_1^L(\theta', \cdot) : \theta' \in \Theta\}$ is equi-Lipschitz continuous. Let θ and $\tilde{\theta}$ be such that $\theta \neq \tilde{\theta}$, and consider

$$|W_1^L(\theta', \theta) - W_1^L(\theta', \tilde{\theta})| = \left| \int_{\Delta(\Theta)} ((\theta - \tilde{\theta}) q_1(F_2) + (p(\theta) - p(\tilde{\theta}))\Delta v q_2^*(v_L, F_2))\beta(dF_2|\theta') \right|$$
$$\leq |\theta - \tilde{\theta}| \int_{\Delta(\Theta)} \left[q_1(F_2) + \left| \frac{p(\theta) - p(\tilde{\theta})}{\theta - \tilde{\theta}} \right| \Delta v q_2^*(v_L, F_2) \right] \beta(dF_2|\theta') \leq |\theta - \tilde{\theta}|(Q + K\Delta v),$$

where K is the Lipschitz constant of p and Q is the bound on quality. We conclude that U_1^L is Lipschitz continuous because it is the max over a family of equi-Lipschitz continuous functions. Moreover, at any point of differentiability of $U_1^L(\cdot)$, we have

$$U_1^{L\prime}(\theta) = \int_{\Delta(\Theta)} [q_1(F_2) + p'(\theta)\Delta v q_2^*(v_L, F_2)] \, \beta(dF_2|\theta). \tag{L-E}$$

Incentive compatibility implies $U_1^{L\prime}(\theta)$ is nondecreasing. Equation L-E implies

$$\int_\Theta \int_{\Delta(\Theta)} x_1(F_2)\beta(dF_2|\theta)F_1(d\theta) = \int_\Theta \int_{\Delta(\Theta)} [\theta q_1(F_2) + p(\theta)\Delta v q_2^*(v_L, F_2)] \, \beta(dF_2|\theta)F_1(d\theta)$$
$$- \int_\Theta \int_0^\theta \left(\int_{\Delta(\Theta)} [q_1(F_2) + p'(\theta)\Delta v q_2^*(v_L, F_2)] \, \beta(dF_2|u) \right) du F_1(d\theta). \tag{17}$$

Virtual surplus: Replacing Equation 17 in L-OPT and integrating by parts, we obtain the virtual surplus representation of the U-firm's profits. Denote by τ the distribution on $\Theta \times \Delta(\Theta)$ defined as $P(\widetilde{\Theta} \times \widetilde{U}) = \int_{\widetilde{\Theta}} \beta(\widetilde{U}|\theta)F_1(d\theta)$, for all measurable subsets $\widetilde{\Theta}, \widetilde{U}$ of Θ and $\Delta(\Theta)$. Letting τ denote its marginal on $\Delta(\Theta)$, Proposition 3.6

in Crauel (2002) delivers the virtual surplus representation of the U-firm's problem:

$$\max_{q_1(\cdot),\tau} \int_{\Delta(\Theta)} \int_{\Theta} \left[\hat{\theta}(F_1)q_1(F_2) - c(q_1(F_2)) + \gamma \Pi_D^L(\theta, F_2) + (1-\gamma)W_D^L(\theta, F_2)\right] F_2(d\theta)\tau(dF_2) \tag{NC-VS}$$

$$\text{s.t. } \tau \text{ is Bayes' plausible given } F_1 \tag{BP}$$

$$U_1^{L'}(\theta) = \int_{\Delta(\Theta)} \left[q_1(F_2) + p'(\theta)\Delta v q_2^*(v_L, F_2)\right] \beta(dF_2|\theta) \text{ is increasing in } \theta. \tag{MON}$$

In the above expression, Π_D^L are the downstream profits adjusted by *dynamic* rents:

$$\Pi_D^L(\theta, F_2) = p(\theta)v_H(1 - q_2^*(v_L, F_2)) \\ + q_2^*(v_L, F_2)\left[p(\theta)v_H + (1-p(\theta))\left(v_L - \frac{1-F_1(\theta)}{f_1(\theta)}\frac{p'(\theta)}{1-p(\theta)}\Delta v\right)\right],$$

and W_D^L is a measure of downstream consumer welfare, adjusted by *dynamic* rents:

$$W_D^L(\theta, F_2) = \Delta v q_2^*(v_L, F_2)\left(p(\theta) - \frac{1-F_1(\theta)}{f_1(\theta)}p'(\theta)\right).$$

Relaxed problem: Ignoring the monotonicity constraint, we can choose $q_1(F_2)$ to maximize pointwise the integrand in Equation NC-VS, in which case, we obtain

$$q_1(F_2) = \max\left\{0, \frac{2\mu_{F_2} - 1}{c}\right\} \equiv q_1(\mu_{F_2}). \tag{18}$$

Replacing Equation 18 in the objective in NC-VS leads to the following expression for the objective function in NC-VS:

$$\int_{\Delta(\Theta)} \frac{(\max\{2\mu_{F_2} - 1, 0\})^2}{2c}\tau(dF_2) + (1-\gamma)\int_{\Delta(\Theta)} q_2^*(v_L, F_2)\Delta v(2\mu_{F_2} - 1)\tau(dF_2) \tag{19}$$
$$+ \gamma \int_{\Delta(\Theta)} \left[(1 - q_2^*(v_L, F_2))v_H\mu_{F_2} + q_2^*(v_L, F_2)\left(v_H\mu_{F_2} + (v_L - \Delta v)(1 - \mu_{F_2})\right)\right]\tau(dF_2).$$

Thus, the virtual upstream profits are a function of the posterior mean of F_2, and we can pool together all distributions F_2 which induce the same posterior mean.

Product line design as information design: Thus, letting m denote the posterior mean, the firm's payoff can be written as:

$$\int_0^1 \Pi(m)G(dm),$$

where G is a distribution on Θ that is dominated by F_1 in the convex order, and in a slight abuse of notation, the function Π corresponds to the virtual surplus as a function of the posterior mean m and is defined as follows. When $\overline{\mu} < 1/2$, we have

$$\Pi(m) = \begin{cases} v_H\left((2m-1)(1-\overline{\mu}) + \gamma\overline{\mu}\right) & \text{if } m < \overline{\mu} \\ v_H \gamma m & \text{if } \overline{\mu} \leq m < 1/2 \\ v_H\left(\gamma m + \frac{(2m-1)^2}{2cv_H}\right) & \text{if } m \geq 1/2 \end{cases}. \tag{20}$$

Instead, when $\overline{\mu} > 1/2$, we have

$$\Pi(m) = \begin{cases} v_H\left((2m-1)(1-\overline{\mu}) + \gamma\overline{\mu}\right) & \text{if } m < 1/2 \\ v_H\left((2m-1)(1-\overline{\mu}) + \gamma\overline{\mu} + \frac{(2m-1)^2}{2cv_H}\right) & \text{if } 1/2 \leq m \leq \overline{\mu} \\ v_H\left(\gamma m + \frac{(2m-1)^2}{2cv_H}\right) & \text{if } m > \overline{\mu} \end{cases}. \tag{21}$$

The above expressions show that the value of v_H does not matter, so in what follows we write $\tilde{c} \equiv cv_H$. The relaxed problem then reduces to

$$\max_{G: F_1 \succ_{cx} G} \int_0^1 \Pi(m) G(dm). \tag{22}$$

Because Π satisfies the conditions of Dworczak and Martini (2019), their results imply that G^* is a solution to Equation 22 if and only if a convex function P exists such that (i) $\mathbb{E}_{G^*} P(m) = \mathbb{E}_{F_1} P(m)$, (ii) $P \geq \Pi$, and (iii) supp $G^* \subseteq \{m : P(m) = \Pi(m)\}$. The proof of Propositions 3 and 4 proceeds as follows. In each case, we prove that the induced distribution over posterior means is optimal by constructing a convex function P and verifying that the above conditions hold. We then use Equation 17 to construct the transfers and provide conditions under which the solution satisfies that $U'(\theta)$ is nondecreasing.

Proof of Proposition 3. The solution when $\overline{\mu} \leq 1/4$ follows from the proof of Proposition 5. Suppose that $\overline{\mu} > 1/4$. Let $m^* = 2\overline{\mu}$ and note that $\overline{\mu} = \mathbb{E}[\theta | \theta \leq m^*]$. When $l_1(\cdot) \geq 0$, the distribution over posteriors is as follows:

$$G(m) = \begin{cases} 0 & \text{if } m < \overline{\mu} \\ F_1(m^*) & \text{if } \overline{\mu} \leq m < m^* \\ F_1(m) & \text{otherwise} \end{cases}.$$

The price function that supports this as a solution is

$$P(m) = \begin{cases} \Pi(\overline{\mu}) + \frac{\Pi(m^*)-\Pi(\overline{\mu})}{m^*-\overline{\mu}}(m-\overline{\mu}) & \text{if } m \leq m^* \\ \Pi(m) & \text{otherwise} \end{cases}.$$

Note that we only need to check that $P(0) \geq \Pi(0)$. This is the case if and only if

$$P(0) = 2\Pi(\overline{\mu}) - \Pi(m^*) = v_H\left(-\frac{(4\overline{\mu}-1)^2}{2\tilde{c}}\right) \geq \Pi(0) = v_H(\gamma\overline{\mu} - (1-\overline{\mu}))$$
$$\Leftrightarrow \frac{(4\overline{\mu}-1)^2}{2\tilde{c}} \leq (1-\overline{\mu}) - \gamma\overline{\mu},$$

which corresponds to $l_0 \geq 0$.

Instead, when $l_0 \leq 0$, let m_*, m^* be such that $\mathbb{E}[\theta|\theta \in [m_*, m^*]] = \overline{\mu}$. In this case, the distribution over posteriors is as follows:

$$G(m) = \begin{cases} F_1(m) & \text{if } 0 \leq m < m_* \text{ or } m^* < m \leq 1 \\ F_1(m_*) & \text{if } m_* \leq m < \overline{\mu} \\ F_1(m^*) & \text{if } \overline{\mu} \leq m \leq m^* \end{cases}.$$

The price function that supports this as a solution is

$$P(m) = \begin{cases} \Pi(m_*) + \frac{\Pi(m^*)-\Pi(m_*)}{m^*-m_*}(m-m_*) & \text{if } m_* \leq m \leq m^* \\ \Pi(m) & \text{otherwise} \end{cases}.$$

Assume first that m_*, m^* exist. We construct a convex function $P(m)$ that satisfies the necessary properties. Define $P : [0,1] \mapsto \mathbb{R}$ as follows:

$$P(m) = \begin{cases} \Pi(m) & \text{if } m \notin [m_*, m^*] \\ \Pi(m_*) + \frac{\Pi(m^*)-\Pi(m_*)}{m^*-m_*}(m-m_*) & \text{otherwise} \end{cases}. \tag{23}$$

Clearly, P is convex and $P(m) \geq \Pi(m)$. To see that $\mathbb{E}_G[P] = \mathbb{E}_{F_1}[P]$, note that

$$\mathbb{E}_{F_1}[P(m)] - \mathbb{E}_G[P(m)] = \int_{m_*}^{m^*} P(m)dm - (m^* - m_*)P(\overline{\mu}).$$

Now, because P is linear on $[m_*, m^*]$, we have that

$$\int_{m_*}^{m^*} P(m)dm = (m^* - m_*)P\left(\frac{m_* + m^*}{2}\right) = (m^* - m_*)P(\overline{\mu}).$$

To finish the proof, we show m_*, m^* exist so that $P(\overline{\mu}) = \Pi(\overline{\mu})$. The goal is to find

m_*, m^* such that

$$\frac{m_* + m^*}{2} = \overline{\mu} \qquad \Pi(\overline{\mu}) = \Pi(m_*) + \frac{\Pi(m^*) - \Pi(m_*)}{m^* - m_*}(\overline{\mu} - m_*). \tag{24}$$

Note that the second condition is equivalent to

$$\Pi(m^*) = \Pi(m_*) + \frac{\Pi(\overline{\mu}) - \Pi(m_*)}{\overline{\mu} - m_*}(m^* - m_*). \tag{25}$$

Verifying that m_*, m^* exist that satisfy the above equations is equivalent to showing that the quadratic equation

$$h_0(x) \equiv \frac{2}{\tilde{c}}x^2 - x\left(\frac{2}{\tilde{c}} + 2(1-\overline{\mu}) - \gamma\right) + \frac{1}{2\tilde{c}} + 4\overline{\mu}(1-\overline{\mu}) - (1-\overline{\mu}) - \gamma\overline{\mu} = 0,$$

has a solution $m^* \in [1/2, 1]$. It is easy to verify that $h_0(x)$ achieves its minimum at $\frac{1}{2} + (1-\gamma+1-2\overline{\mu})(\tilde{c}/4) \geq 1/2$. So we need to show that $h_0(2\overline{\mu}) > 0$. This ensures that $m_* = 2\overline{\mu} - m^* > 0$. The latter condition is equivalent to $l_0 \leq 0$, which completes the proof. □

Proof of Corollary 1. The mechanism described in Proposition 3 when $\overline{\mu} \leq 1/4$ or $\overline{\mu} \geq 1/4$ and $l_0 \geq 0$ implies that in the solution to the relaxed problem types below m^* are excluded in period 1, whereas types above m^* receive $(2\theta - 1)/c$ and pay $\frac{\theta^2 - m^*(1-m^*)}{c}$. It is immediate to show $U_1^{L\prime}(\theta)$ is monotone. Instead, when $\overline{\mu} \geq 1/4$ and $l_0 \leq 0$, $U_1^{L\prime}(\theta)$ is not monotone, so the solution to the relaxed program cannot be implemented. □

Proof of Corollary 2. Relying on the envelope representation of the consumer's payoffs for the commitment (C-E) and limited commitment (L-E) solutions, we obtain the consumer's payoffs under the assumptions of Proposition 3, case 1. Consumer payoffs under limited commitment are

$$U_1^L(\theta) = \begin{cases} 0 & \text{if } \theta \leq \min\{2\overline{\mu}, 1/2\} \\ \int_{\max\{1/2, 2\overline{\mu}\}}^{\theta} \frac{2x-1}{c}dx & \text{otherwise} \end{cases}.$$

Instead, consumer payoffs under commitment depend on whether $\theta_* \geq 1/2$ (top) or

$\theta_* \leq 1/2$ (bottom):

$$U_1^C(\theta) = \begin{cases} 0 & \text{if } \theta \leq 1/2 \\ \int_{1/2}^{\theta} \frac{2x-1}{c} dx & \text{if } \theta \in [1/2, \theta_*] \\ \int_{1/2}^{\theta} \frac{2x-1}{c} dx + (\theta - \theta_*)\Delta v & \text{if } \theta \geq \theta_* \end{cases},$$

$$U_1^C(\theta) = \begin{cases} 0 & \text{if } \theta \leq \theta_* \\ (\theta - \theta_*)\Delta v & \text{if } \theta \in [\theta_*, 1/2] \\ \int_{1/2}^{\theta} \frac{2x-1}{c} dx + (\theta - \theta_*)\Delta v & \text{if } \theta \geq 1/2 \end{cases}.$$

Tedious, but straightforward algebra, verifies that $U_1^C(\theta) \geq U_1^L(\theta)$, strictly so whenever (i) $\gamma < 1$ and $\theta > 0$, or (ii) $\gamma = 1$, $\overline{\mu} > 1/4$, and $\theta > 1/2$. □

Proof of Proposition 4, item 1. Define m_* to be such that $\overline{\mu} = \mathbb{E}[\theta|\theta \geq m_*]$. That is, $m_* = 2\overline{\mu} - 1$. Note that whether m_* is above or below $1/2$ depends on whether $\overline{\mu}$ is above or below $3/4$. In both cases, the posterior distribution is the same:

$$G(m) = \begin{cases} F_1(m) & \text{if } 0 \leq m < m_* \\ F_1(m_*) & \text{if } m_* \leq m \leq \overline{\mu} \\ 1 & \text{otherwise} \end{cases},$$

and it is supported by the following convex function:

$$P(m) = \begin{cases} \Pi(m) & \text{if } 0 \leq m \leq m_* \\ \Pi(m_*) + \frac{\Pi(\overline{\mu}) - \Pi(m_*)}{\overline{\mu} - m_*}(m - m_*) & \text{if } m_* \leq m \leq 1 \end{cases}.$$

The only remaining step is to check that $P(m) \geq \Pi(m)$ for all m. Indeed, it suffices to check that this happens at $m = 1$.

Suppose first that $1/2 \leq \overline{\mu} \leq 3/4$ so that $m_* = 2\overline{\mu} - 1 \leq 1/2$. Then,

$$P(1) = 2\Pi(\overline{\mu}) - \Pi(m_*) = (1 - \overline{\mu}) + \gamma\overline{\mu} + \frac{(2\overline{\mu} - 1)^2}{2\tilde{c}}$$
$$\Pi(1) = \gamma + \frac{1}{2\tilde{c}}.$$

Then, $P(1) \geq \Pi(1)$ if and only if

$$(1 - \overline{\mu})(1 - \gamma) + \left[\frac{(2\overline{\mu} - 1)^2}{2\tilde{c}} - \frac{1}{2\tilde{c}}\right] \geq 0,$$

yielding the condition that $l_2 \geq 0$.

Consider now $\overline{\mu} \in [3/4, 1]$, in which case $m_* \geq 1/2$. Then,

$$P(1) = 2\Pi(\overline{\mu}) - \Pi(m_*) = (1-\overline{\mu}) + \gamma\overline{\mu} + \frac{(2\overline{\mu}-1)^2}{2\tilde{c}} - \frac{(2(2\overline{\mu}-1)-1)^2}{2\tilde{c}}$$
$$\Pi(1) = \gamma + \frac{1}{2\tilde{c}}.$$

Then, $P(1) \geq \Pi(1)$ if and only if

$$(1-\overline{\mu})\left[1-\gamma-\frac{4}{\tilde{c}}(1-\overline{\mu})\right] \geq 0,$$

yielding the condition that $l_3 \geq 0$. □

Proof of Proposition 4, item 2. Let m_*, m^* be such that $\mathbb{E}[\theta|\theta \in [m_*, m^*]] = \overline{\mu}$. In this case, the distribution over posteriors is as follows:

$$G(m) = \begin{cases} F_1(m) & \text{if } 0 \leq m < m_* \text{ or } m^* < m \leq 1 \\ F_1(m_*) & \text{if } m_* \leq m < \overline{\mu} \\ F_1(m^*) & \text{if } \overline{\mu} \leq m \leq m^* \end{cases}.$$

The price function that supports this as a solution is

$$P(m) = \begin{cases} \Pi(m_*) + \frac{\Pi(m^*)-\Pi(m_*)}{m^*-m_*}(m-m_*) & \text{if } m_* \leq m \leq m^* \\ \Pi(m) & \text{otherwise} \end{cases}.$$

Similar to the steps in the proof of part 2 of Proposition 3, it suffices to check that m_*, m^* can be chosen so that $P(\overline{\mu}) = \Pi(\overline{\mu})$. The goal is to find m_*, m^* such that

$$\frac{m_* + m^*}{2} = \overline{\mu} \qquad \Pi(\overline{\mu}) = \Pi(m_*) + \frac{\Pi(m^*)-\Pi(m_*)}{m^*-m_*}(\overline{\mu}-m_*). \tag{26}$$

As in the proof of item 1 in Proposition 4, we consider two cases. Suppose first that $m_* \leq 0.5$. Then, finding a solution to Equation 26 is equivalent to finding a solution $m^* \in [\overline{\mu}, 1]$ to the following quadratic equation:

$$h_1(x) = \frac{2}{\tilde{c}}x^2 + x\left(\gamma - 2(1-\overline{\mu}) - \frac{2}{\tilde{c}}\right) - \left(\frac{4\overline{\mu}^2 - 4\overline{\mu}}{\tilde{c}} + \frac{1}{2\tilde{c}} + \overline{\mu}(1+\gamma) - 1\right).$$

Now, because

$$h_1(\overline{\mu}) = -\frac{(2\overline{\mu}-1)^2}{2\tilde{c}} + (1-\overline{\mu})(1-2\overline{\mu}) < 0, \tag{27}$$

we need to check that $h_1(2\overline{\mu} - 1/2) \leq 0$ and $h_1(1) \geq 0$. The first ensures that $m_* = 2\overline{\mu} - m^* \leq 1/2$ and the latter ensures that $m^* \leq 1$. Note that because $m^* \in [2\overline{\mu} - 1/2, 1]$, it must be the case that $\overline{\mu} \leq 3/4$.

Now, $h_1(1) \geq 0$ is equivalent to the condition that

$$l_1(\cdot) = (1-\gamma)(1-\overline{\mu}) + \frac{(2\overline{\mu}-1)^2}{\tilde{c}} - \frac{1}{2\tilde{c}} \leq 0.$$

Instead, $h_1(2\overline{\mu} - 1/2) \leq 0$ is equivalent to

$$\overline{\mu} \leq \frac{2 + (4-\gamma)\tilde{c}}{4(1+\tilde{c})} \in \left[\frac{1}{2}, \frac{4-\gamma}{4}\right].$$

Suppose now that $m_* \in [0.5, \overline{\mu})$. Then, finding a solution to Equation 26 is equivalent to finding a solution $m_* \in [1/2, \overline{\mu}]$ to the following quadratic equation:

$$h_2(x) = \frac{4}{\tilde{c}}x^2 - x\left(\frac{8}{\tilde{c}}\overline{\mu} + \gamma - 2(1-\overline{\mu})\right) + \overline{\mu}\left(\frac{4\overline{\mu}}{\tilde{c}} + \gamma\right) - (1-\overline{\mu})(4\overline{\mu}-1).$$

Now, because

$$h_2(\overline{\mu}) = (1-\overline{\mu})(1-2\overline{\mu}) < 0,$$

we need to check that $h_2(1/2) \geq 0$ and $h_2(2\overline{\mu} - 1) \geq 0$. The first ensures that $1/2 \leq m_*$ and the second ensures that $m^* = 2\overline{\mu} - m_* \leq 1$. Now, when $\overline{\mu} \leq 3/4$, $2\overline{\mu} - 1 \leq 1/2$ so that $h_2(1/2) \geq 0$ implies $h_2(2\overline{\mu} - 1) \geq 0$. The condition $h_2(1/2) \geq 0$ requires that

$$\overline{\mu} \geq \frac{2 + (4-\gamma)\tilde{c}}{4(1+\tilde{c})}.$$

Thus, we have that

$$\frac{2 + (4-\gamma)\tilde{c}}{4(1+\tilde{c})} \leq \overline{\mu} \leq \frac{3}{4}.$$

This implies that $(1-\gamma)\tilde{c} \leq 1$. As a consequence, we obtain that

$$2(1-\overline{\mu})(1-\gamma)\tilde{c} + 2(2\overline{\mu}-1)^2 - 1 \leq 2(1-\overline{\mu}) + 2(2\overline{\mu}-1)^2 - 1 = (2\overline{\mu}-1)(4\overline{\mu}-3) \leq 0,$$

which is the remaining condition in item 2 in Proposition 4 when $\overline{\mu} \leq 3/4$. Instead, when $\overline{\mu} > 3/4$, $h_2(2\overline{\mu} - 1) \geq 0$ requires that

$$(1-\overline{\mu})\left[\frac{4}{\tilde{c}}(1-\overline{\mu}) - (1-\gamma)\right] \geq 0.$$

□

Proof of Corollary 3. Standard results imply that if $U_1^{L'}$ is monotone, then transfers exist that implement the allocation that solves the relaxed program. This immediately holds under the conditions of case 1 in Proposition 4. Consider then the solution to the relaxed problem under the conditions of case 2. Note that when $m_* \leq 1/2$, we have that

$$U_1^{L'}(\theta) = \begin{cases} \Delta v & \text{if } \theta < m_* \\ \frac{2\overline{\mu}-1}{c} + \Delta v & \text{if } m_* \leq \theta \leq m^* \\ \frac{2\theta-1}{c} & \text{otherwise} \end{cases},$$

whereas when $m_* \geq 1/2$, we have that

$$U_1^{L'}(\theta) = \begin{cases} \Delta v & \text{if } \theta \leq 1/2 \\ \frac{2\theta-1}{c} + \Delta v & \text{if } 1/2 \leq \theta < m_* \\ \frac{2\overline{\mu}-1}{c} + \Delta v & \text{if } m_* \leq \theta \leq m^* \\ \frac{2\theta-1}{c} & \text{otherwise} \end{cases}.$$

Thus, in both cases monotonicity is satisfied if and only if

$$m^* \geq \overline{\mu} + \frac{(1-\overline{\mu})\tilde{c}}{2}. \tag{28}$$

Case 1: $m_* \leq 1/2$ Recall that in this case the conditions of Proposition 4 boil down to $\overline{\mu} \leq \min\{3/4, \frac{2+(4-\gamma)\tilde{c}}{4(1+\tilde{c})}\}$ and $l_1(\overline{\mu}, \gamma, \tilde{c}) \leq 0$.

Equation 28 holds if $h_1(\overline{\mu} + (1-\overline{\mu})\tilde{c}/2) \leq 0$ and $\overline{\mu} + (1-\overline{\mu})\tilde{c}/2 \leq 1$. The latter holds if $\tilde{c} \leq 2$. This former holds if and only if:

$$\overline{\mu}(1-\overline{\mu}) \leq \frac{1+(1-\gamma)\tilde{c}^2}{4+\tilde{c}^2}.$$

When $\gamma = 1$, the above equation is inconsistent with $\overline{\mu} \leq \frac{2+(4-\gamma)\tilde{c}}{4(1+\tilde{c})}$. Instead, when $\gamma = 0$, the above equation is always satisfied.

Case 2: $m_* \geq 1/2$ Because $h_2(\cdot)$ is expressed in terms of m_*, Equation 28 is equivalent to $m_* \leq \overline{\mu} - (1-\overline{\mu})\tilde{c}/2$. Thus, in this case we need that

$$\overline{\mu} - \frac{\tilde{c}(1-\overline{\mu})}{2} \geq \frac{1}{2} \Leftrightarrow 1 - 2\overline{\mu} + \tilde{c}(1-\overline{\mu}) \leq 0$$
$$h_2(\overline{\mu} - (1-\overline{\mu})\tilde{c}/2) \leq 0 \Leftrightarrow (1-\overline{\mu})\left(1 - 2\overline{\mu} + \frac{\gamma\tilde{c}}{2}\right) \leq 0.$$

Recall that when $m_* \geq 1/2$, the conditions of Proposition 4 can be split into two groups:

Case 2a: In this case, $\overline{\mu} \in [\frac{2+(4-\gamma)\tilde{c}}{4(1+\tilde{c})}, 3/4]$ (As we showed before, this implies that $l_1(\overline{\mu}, \gamma, \tilde{c}) \leq 0$). We thus have the following three equations:

$$1 - 2\overline{\mu} + 2\tilde{c}(1-\overline{\mu}) \leq 0, \qquad 1 - 2\overline{\mu} + (1-\overline{\mu})\tilde{c} \leq 0, \qquad 1 - 2\overline{\mu} + \tilde{c}\gamma/2 \leq 0,$$

where the first is the condition that defines Case 2a. We thus obtain that when $\gamma = 0$, monotonicity always holds. Instead, when $\gamma = 1$, the condition $\overline{\mu} \in [\frac{2+(4-\gamma)\tilde{c}}{4(1+\tilde{c})}, 3/4]$ implies $\overline{\mu} = 3/4$, which is part of the next case.

Case 2b: In this case, $\overline{\mu} \geq 3/4$ and $(4/\tilde{c})(1-\overline{\mu}) - (1-\gamma) \geq 0$ and the monotonicity constraint implies that

$$1 - 2\overline{\mu} + (1-\overline{\mu})\tilde{c} \leq 0, 1 - 2\overline{\mu} + \tilde{c}\gamma/2 \leq 0.$$

When $\gamma = 0$, the conditions are implied by $\tilde{c} \leq 4(1-\overline{\mu})$, so that monotonicity always holds. Instead, when $\gamma = 1$, we obtain the condition that $1 - 2\overline{\mu} + \tilde{c}/2 \leq 0$. □

Figure 8 shows the configurations of $(\overline{\mu}, \tilde{c})$ such that the monotonicity constraint holds.

Proof of Corollary 4. In a mass downstream market, consumer's payoffs under commitment are as follows:

$$U_1^C(\theta) = \theta\Delta v + \mathbb{1}[\theta \geq 1/2] \int_{1/2}^{\theta} \frac{2x-1}{c} dx,$$

whereas average consumer welfare is given by

$$AW_C = \int_0^1 U_1^C(\theta) d\theta = v_H \left(\frac{(1-\overline{\mu})}{2} + \frac{1}{24\tilde{c}}\right).$$

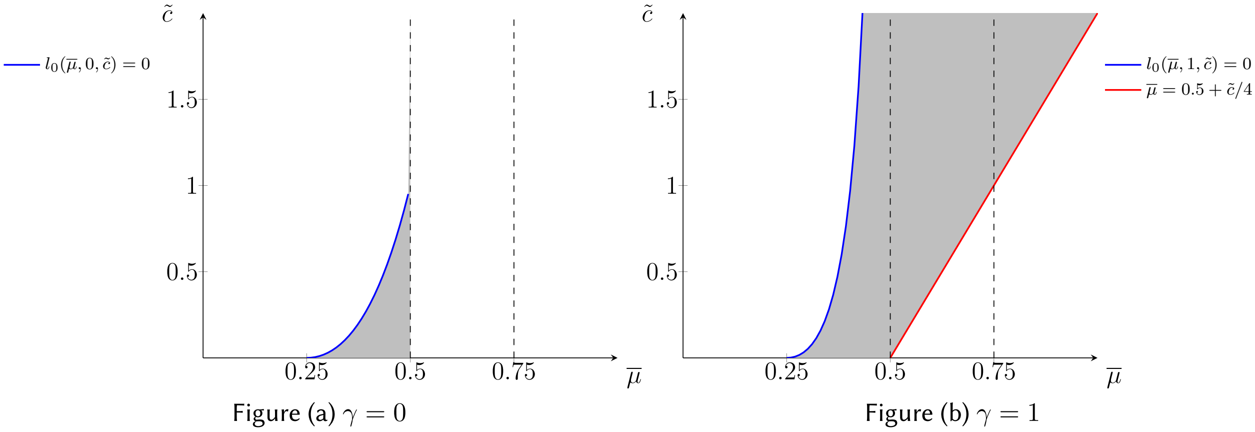

Figure 8: Shaded gray area shows parameter values for which the solution to the relaxed problem does not satisfy monotonicity

Instead, the payoffs under limited commitment depend on $\overline{\mu}$ and γ. For ease of exposition, we refer to the cases by the figures that illustrate them.

Case 1: Product line is as in Figure 6a Relative to the commitment solution, there is less exclusion (all types above $2\overline{\mu}-1 \leq 1/2$ are served) and only one quality besides the outside option one is provided ($(2\overline{\mu}-1)/c$). Like in the commitment solution, there is no downstream price discrimination. Consumer payoffs are given by:

$$U_1^L(\theta) = \theta\Delta v + \mathbb{1}[\theta \geq 2\overline{\mu} - 1](\theta - (2\overline{\mu} - 1))\frac{2\overline{\mu} - 1}{c}.$$

It is easy to see that $\theta \leq 2\overline{\mu} - 1$ are indifferent and $\theta \in [2\overline{\mu} - 1, 1/2]$ prefer limited commitment to the commitment solution. Consider now $\theta \geq 1/2$, then

$$U_1^C(\theta) - U_1^L(\theta) = -\left(\frac{1}{2} - (2\overline{\mu} - 1)\right)\frac{2\overline{\mu} - 1}{c} + \frac{1}{c}\left(\theta - \frac{1}{2}\right)\left(\theta + \frac{1}{2} - 2\overline{\mu}\right).$$

This difference is increasing in θ for $\theta \geq \overline{\mu}$, negative for $\theta = \overline{\mu}$, and positive for $\theta = 1$. Instead, average consumer welfare under limited commitment is given by

$$AW_L = \Delta v + \frac{2\overline{\mu} - 1}{c}\int_{2\overline{\mu}-1}^{1}\left[\theta - (2\overline{\mu} - 1)\right]d\theta = \Delta v + \frac{2\overline{\mu} - 1}{c}2(1 - \overline{\mu})^2.$$

Figure 9 illustrates the difference $\tilde{c}\,(AW_C - AW_L)$ as a function of $\overline{\mu}$. When $\overline{\mu}$ is close to $1/2$, the limited commitment solution is close to full exclusion – even if it serves a larger number of types. As $\overline{\mu}$ grows, the upward distortion in quality kicks in so that average consumer welfare is larger under limited commitment.

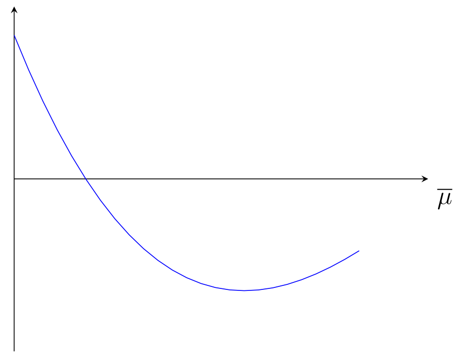

Figure 9: Difference in average welfare under commitment and limited commitment as a function of $\overline{\mu}$ for $\overline{\mu} \in [0.5, 0.75]$.

Case 2: Product line is as in Figure 6b: The product line coincides with that in the commitment solution for consumer types in $[1/2, 2\overline{\mu}-1]$. Consumer payoffs under limited commitment are given by

$$U_1^L(\theta) = \theta\Delta v + \mathbb{1}[\theta \geq 1/2] \int_{1/2}^{\min\{2\overline{\mu}-1,\theta\}} \frac{2x-1}{c} dx + \mathbb{1}[\theta \geq 2\overline{\mu}-1](\theta - (2\overline{\mu}-1))\frac{2\overline{\mu}-1}{c}.$$

It is immediate that $U_1^C(\theta) = U_1^L(\theta)$ whenever $\theta \leq 2\overline{\mu}-1$. Instead, for $\theta \geq 2\overline{\mu}-1$, we have that θ is worse off under the commitment solution. Indeed,

$$c\left(U_1^C(\theta) - U_1^L(\theta)\right) = \int_{2\overline{\mu}-1}^{\theta} 2(x-\overline{\mu})dx = (\theta - (2\overline{\mu}-1))(\theta - 1) \leq 0,$$

where the difference is strict whenever $\theta \in (2\overline{\mu}-1, 1)$. There are two reasons for this comparison: First, consumer types in $[2\overline{\mu}-1, \overline{\mu}]$ obtain a good of a higher quality than in the commitment solution. Second, whereas consumer types in $[\overline{\mu}, 1]$ obtain a good of lower quality, they also pay less for it, as the most the U-firm can extract is what the consumer with type $2\overline{\mu}-1$ pays for the good.

Case 3: Product line is as in Figure 7a: Relative to the commitment solution, there is less exclusion (all types above $m_* \leq 1/2$ are served) and there is price discrimination for types in the high-end of the product line ($\theta \geq m^*$). Consumer payoffs are given by:

$$U_1^L(\theta) = \max\{\theta, m^*\}\Delta v + \mathbb{1}[\theta \geq m_*](\min\{m^*, \theta\} - m_*)\frac{2\overline{\mu}-1}{c} + \mathbb{1}[\theta \geq m^*] \int_{m^*}^{\theta} \frac{2x-1}{c} dx.$$

Clearly, consumer types below m_* are indifferent and consumer types in $[m_*, 1/2]$ are better off under limited commitment. Consider now a consumer with type in $[m_*, m^*]$.

In this case,

$$U_1^C(\theta) - U_1^L(\theta) = -\left(\frac{1}{2} - m_*\right)\left(\frac{2\overline{\mu}-1}{c}\right) + \frac{2}{c}\left(\theta - \frac{1}{2}\right)\left(\frac{1}{2}\left(\theta + \frac{1}{2}\right) - \overline{\mu}\right).$$

Instead, for $\theta \geq m^*$, we have that:

$$U_1^C(\theta) - U_1^L(\theta) = (\theta - m^*)\Delta v + \frac{1}{c}\left(m^* - \frac{1}{2}\right)\left(m^* + \frac{1}{2} - 2\overline{\mu}\right) - \left(\frac{1}{2} - m_*\right)\left(\frac{2\overline{\mu}-1}{c}\right),$$

which is increasing in θ and positive for $\theta = 1$: $U_1(1) = (1-m^*)\Delta v + (1/c)(m^* + 1/2 - 2\overline{\mu})^2 > 0$.

Case 4: Product line is as in Figure 7b: Exclusion is the same as in the commitment solution, but there is price discrimination on the high-end of the product line ($\theta \geq m^*$). Consumer payoffs are given by

$$U_1^L(\theta) = \max\{\theta, m^*\}\Delta v + \mathbb{1}[\theta \geq m_*](\min\{m^*, \theta\} - m_*)\frac{2\overline{\mu}-1}{c} + \mathbb{1}[\theta \geq 1/2]\int_{1/2}^{\min\{\theta, m_*\}} \frac{2x-1}{c}dx + \mathbb{1}[\theta \geq m^*]\int_{m^*}^{\theta} \frac{2x-1}{c}dx.$$

It is immediate that consumer types $\theta \leq m_*$ are indifferent between the commitment and limited commitment allocations. Consider now $\theta \in [m_*, m^*]$. We have that

$$c\left(U_1^C(\theta) - U_1^L(\theta)\right) = (\theta - m_*)(\theta - m^*) \leq 0.$$

As in case 2, this obtains because types below $\overline{\mu}$ obtain higher qualities and types above $\overline{\mu}$ pay less for the lower-quality good that they obtain. Finally, consider now $\theta \geq m^*$. In this case,

$$U_1^C(\theta) - U_1^L(\theta) = (\theta - m^*)\Delta v \geq 0.$$

Relative to the commitment solution, these types receive the same (average) rents upstream, but now they face price discrimination downstream.

□

A.3 Proof of Proposition 5

When the U-firm can design both the product line and the information available to the D-firm it is without loss of generality to restrict attention to mechanisms $\varphi : \Theta \mapsto \Delta(A_1 \times \{v_L, v_H\})$ that assign to each type θ a lottery over allocations

$(q_1, x_1) \in A_1$ and price recommendations for the D-firm, $x_2 \in \{v_L, v_H\}$ (Myerson, 1982).[27] Furthermore, it is without loss to restrict attention to mechanisms such that the consumer participates and truthfully reports her type and the D-firm *obeys* the received recommendation. Thus, when the consumer's type is θ and the D-firm receives recommendation x_2, downstream profits are given by:

$$\Pi_D^M(\theta, x_2) = \mathbb{1}[x_2 = v_L]v_L + (1 - \mathbb{1}[x_2 = v_L])p(\theta)v_H. \tag{29}$$

Furthermore, the consumer's payoff when her type is θ and she reports θ' is

$$W_1^M(\theta', \theta) = \int_{A_1 \times \{v_L, v_H\}} [\theta q_1(F_2) - x_1(F_2) + p(\theta)\mathbb{1}[x_2 = v_L]\Delta v] \, \varphi(d(q_1, x_1, x_2)|\theta').$$

Let $U_1^M(\theta) \equiv W_1^M(\theta, \theta)$ denote the payoff from truthtelling. The U-firm-profit maximizing mechanism solves

$$\max_{\varphi:\Theta \mapsto \Delta(A_1 \times \{v_L, v_H\})} \int_\Theta \int_{A_1 \times \{v_L, v_H\}} \left[x_1(\theta) - c(q_1(\theta)) + \gamma \Pi_D^M(\theta, x_2)\right] \varphi(d(q_1, x_1, x_2)|\theta) F_1(d\theta) \tag{M-OPT}$$

$$\text{s.t.} (\forall \theta \in \Theta) U_1^M(\theta) \geq 0 \tag{M-PC}$$

$$(\forall \theta \in \Theta)(\forall \theta' \in \Theta) U_1^M(\theta) \geq W_1^M(\theta', \theta) \tag{M-TT$_1$}$$

$$\int_{\Theta \times A_1 \times \{v_L\}} (v_L - p(\theta)v_H) \, \varphi(d(q_1, x_1, x_2)|\theta) F_1(d\theta) \geq 0. \tag{M-OB$_{v_L}$}$$

$$\int_{\Theta \times A_1 \times \{v_H\}} (p(\theta)v_H - v_L) \, \varphi(d(q_1, x_1, x_2)|\theta) F_1(d\theta) \geq 0. \tag{M-OB$_{v_H}$}$$

The last two constraints are the D-firm's obedience constraints. It is possible to show that M-OPT is equivalent to maximizing

$$\max_\varphi \int \left[\hat{\theta}(F_1)q_1 - c(q_1) + \gamma \Pi_D^M(\theta, x_2) + (1 - \gamma) W_D^M(\theta, x_2)\right] \varphi(d(q_1, x_2)|\theta) F_1(d\theta) \tag{M-VS}$$

s.t. M-OB$_{v_L}$, M-OB$_{v_H}$, and

$$U_1^{M\prime}(\theta) = \int_{[0,Q] \times \{v_L, v_H\}} \left[q_1 + p'(\theta)\Delta v \mathbb{1}[x_2 = v_L]\right] \varphi(d(q_1, x_2)|\theta) \text{ is increasing in } \theta.$$

[27] These are the only recommendations D-firm would find optimal to follow.

In the above expression, Π_D^M are the downstream profits adjusted by *dynamic* rents:

$$\Pi_D^M(\theta, x_2) = p(\theta)v_H(1 - \mathbb{1}[x_2 = v_L]) \\ + \mathbb{1}[x_2 = v_L]\left[p(\theta)v_H + (1 - p(\theta))\left(v_L - \frac{1 - F_1(\theta)}{f_1(\theta)}\frac{p'(\theta)}{1 - p(\theta)}\Delta v\right)\right],$$

and W_D^M is a measure of downstream consumer welfare, adjusted by *dynamic* rents:

$$W_D^M(\theta, x_2) = \Delta v \mathbb{1}[x_2 = v_L]\left(p(\theta) - \frac{1 - F_1(\theta)}{f_1(\theta)}p'(\theta)\right).$$

Note that the objective in M-VS would be maximized by placing probability 1 on $q_1 = \hat{\theta}(F_1)/c$ conditional on any recommendation to the D-firm. This can be achieved without affecting the obedience constraints M-OB$_{v_L}$ and M-OB$_{v_H}$.

Thus, the maximization problem M-VS boils down to finding a mapping $\chi : \Theta \mapsto \Delta(\{v_L, v_H\})$ to maximize:

$$\int_\Theta \sum_{x_2 \in \{v_L, v_H\}} \left((1 - \gamma)W_D^M(\theta, x_2) + \gamma\Pi_D^M(\theta, x_2)\right)\chi(x_2|\theta)F_1(d\theta), \tag{30}$$

subject to M-OB$_{v_L}$ and M-OB$_{v_H}$. Standard results imply that this problem is equivalent to finding a Bayes' plausible posterior distribution $\tau \in \Delta\Delta(\Theta)$ and a selection q_2^* from the D-firm's best response correspondence to solve the above problem. Replacing our parametric assumptions shows that the objective function only depends on the posterior mean of $p(\cdot)$. Thus, we need to solve:

$$\max_{G: U \succ_{cx} G} \int_0^1 \left[(1 - q_2^*(v_L, p))\gamma p v_H + q_2^*(v_L, p)\left(2(1 - \overline{\mu})p + \overline{\mu}(1 + \gamma) - 1\right)\right] G(dp), \tag{31}$$

where $U \succ_{cx} G$ states that the uniform distribution (which is the prior distribution of $p = F_1$) dominates G in the convex order. The result that there is no price discrimination in the solution to M-VS follows from establishing that the concavification of the integrand in Equation 31 evaluated at the prior coincides with the integrand evaluated at the prior. In what follows, denote the integrand in Equation 31 by $\Pi_U^M(p)$. Note that it is piecewise linear in p.

When $\overline{\mu} < 1/2$, Π_U^M satisfies the following: (i) There is a jump up at $p = \overline{\mu}$, (ii) the concavification of $\Pi_U^M(\cdot)$ coincides with $\gamma p v_H$ for $p \geq \overline{\mu}$. The reason for the latter is that $\gamma p v_H \geq 2(1 - \overline{\mu})p + \overline{\mu}(1 + \gamma) - 1$ for all $p \in [0, \overline{\mu}]$. Instead, when $\overline{\mu} > 1/2$, Π_U^M satisfies the following: (i) There is a jump down at $p = \overline{\mu}$, (ii) the concavification of $\Pi_U^M(\cdot)$ coincides with $2p(1 - \overline{\mu}) + \overline{\mu}(1 + \gamma) - 1$ for $p \leq \overline{\mu}$. The reason for the latter is that $\gamma p v_H \leq 2(1-\overline{\mu})p + \overline{\mu}(1+\gamma) - 1$ for all $p \in [\overline{\mu}, 1]$. In both cases, the results in Arieli et al. (2023) imply that no information revelation is optimal.

B Mechanism-selection game and solution concept

We describe here the mechanisms available to the U-firm and D-firm in the data intermediation and limited commitment settings. We refer the reader to Doval and Skreta (2022a) for a discussion of mechanism-selection games with a continuum of types and the appropriate solution concept. Here, we take this discussion as given and merely specify the firms' action sets and the sequential rationality requirements.

Mechanisms: The firm which operates in period $t \in \{1, 2\}$ can choose a mechanism $\mathbb{M}_t$, defined as follows. The mechanism $\mathbb{M}_t$ is a tuple (M_t, S_t, φ_t), where M_t is a set of input messages, S_t is a set of output messages, and $\varphi : M_t \mapsto \Delta(S_t \times A_t)$ is a mapping associating to each input message $m \in M_t$ a lottery over output messages, $s_t \in S_t$ and allocations, $a_t \in A_t$. Recall from Section 2 that $A_1 = [0, Q] \times \mathbb{R}_+$ and $A_2 = \{0, 1\} \times \mathbb{R}_+$.

Given a mechanism $\mathbb{M}_t$, the consumer chooses whether to participate. If she does not participate, then the allocation is no trade. If she participates, then she *privately* sends an input message $m \in M_t$ into the mechanism, which determines the distribution $\varphi_t(\cdot|m)$ from which the output message and the allocation are drawn. The firm designing the mechanism and the consumer observe the output message and the allocation.

Information available to D-firm and consumer: The D-firm observes the output message from the U-firm's upstream mechanism and the consumer's participation decision both in the data intermediation and the limited commitment case.[28] In the case of limited commitment, the D-firm also observes the allocation that the consumer obtains in the upstream interaction.

It follows that in the data intermediation setting the D-firm chooses a mechanism for each output message, $\mathbb{M}_2(s_1)$, whereas in the limited commitment setting the D-firm chooses a mechanism for each output message s_1 and each allocation (q_1, x_1).

[28] Calzolari and Pavan (2006b) make the same assumption in their sequential common agency problem.

D-firm sequential rationality: Given an upstream mechanism, $\mathbb{M}_1$, and a family of downstream mechanisms $\mathbb{M}_2(\cdot)$, the consumer faces an extensive form game. The consumer's participation and reporting strategy together with the upstream mechanism determine the D-firm's beliefs as a function of what the D-firm observes via Bayes' rule. We assume that the D-firm chooses $\mathbb{M}_2$ to maximize revenue conditional on its beliefs.

U-firm optimality: Given the consumer's strategy and the downstream mechanism(s), we assume that the upstream mechanism maximizes upstream profits.

References

Acquisti, A. and H. R. Varian (2005): "Conditioning prices on purchase history," *Marketing Science*, 24, 367–381.

Aliprantis, C. and K. Border (2013): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer-Verlag Berlin and Heidelberg GmbH & Company KG.

Anderson, E. T. and J. D. Dana (2009): "When is price discrimination profitable?" *Management Science*, 55, 980–989.

Anderson, S. P. and L. Celik (2015): "Product line design," *Journal of Economic Theory*, 157, 517–526.

Argenziano, R. and A. Bonatti (2020): "Information Revelation and Privacy Protection," *Working paper.*

Arieli, I., Y. Babichenko, R. Smorodinsky, and T. Yamashita (2023): "Optimal persuasion via bi-pooling," *Theoretical Economics*, 18, 15–36.

Armstrong, M. (2006): *Recent Developments in the Economics of Price Discrimination*, Cambridge University Press, vol. 2 of *Econometric Society Monographs*, 97–141.

Bergemann, D. and M. Ottaviani (2021): "Information markets and nonmarkets," in *Handbook of Industrial Organization*, Elsevier, vol. 4, 593–672.

Calzolari, G. and A. Pavan (2006a): "Monopoly with resale," *The RAND Journal of Economics*, 37, 362–375.

——— (2006b): "On the optimality of privacy in sequential contracting," *Journal of Economic theory*, 130, 168–204.

Carrillo, J. D. and T. Mariotti (2000): "Strategic ignorance as a self-disciplining device," *The Review of Economic Studies*, 67, 529–544.

Crauel, H. (2002): *Random probability measures on Polish spaces*, vol. 11, CRC press.

Cummings, R., K. Ligett, M. M. Pai, and A. Roth (2015): "The strange case of privacy in equilibrium models," *arXiv preprint arXiv:1508.03080*.

Deb, R. and M. Said (2015): "Dynamic screening with limited commitment," *Journal of Economic Theory*, 159, 891–928.

Decarolis, F., G. Spagnolo, and R. Pacini (2016): "Past performance and procurement outcomes," Tech. rep., National Bureau of Economic Research.

Doval, L. and V. Skreta (2022a): "Mechanism design with limited commitment," *Econometrica*, 90, 1463–1500.

——— (2022b): "Mechanism Design with Limited Commitment: Markov Environments," Click here.

——— (Forthcoming): "Constrained information design," *Mathematics of Operations research*.

Dworczak, P. and G. Martini (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

Eilat, R., K. Eliaz, and X. Mu (2021): "Bayesian privacy," *Theoretical Economics*, 16, 1557–1603.

Ellison, G. (2005): "A model of add-on pricing," *The Quarterly Journal of Economics*, 120, 585–637.

Eső, P. and B. Szentes (2007): "Optimal information disclosure in auctions and the handicap auction," *The Review of Economic Studies*, 74, 705–731.

Executive Office of the President of the United States (2015): "Big Data and Differential Pricing," Click here.

Fudenberg, D. and J. M. Villas-Boas (2006): "Behavior-based price discrimination and customer recognition," *Handbook on economics and information systems*, 1, 377–436.

Gentzkow, M. and E. Kamenica (2016): "A Rothschild-Stiglitz approach to Bayesian persuasion," *American Economic Review*, 106, 597–601.

Hart, O. D. and J. Tirole (1988): "Contract renegotiation and Coasian dynamics," *The Review of Economic Studies*, 55, 509–540.

Hidir, S. and N. Vellodi (2021): "Privacy, personalization, and price discrimination," *Journal of the European Economic Association*, 19, 1342–1363.

Ichihashi, S. (2020): "Online privacy and information disclosure by consumers," *American Economic Review*, 110, 569–95.

Itoh, M. (1983): "Monopoly, product differentiation and economic welfare," *Journal of Economic Theory*, 31, 88–104.

Ivaldi, M. and D. Martimort (1994): "Competition under nonlinear pricing," *Annales d'Economie et de Statistique*, 71–114.

Johnson, J. P. and D. P. Myatt (2003): "Multiproduct quality competition: Fighting brands and product line pruning," *American Economic Review*, 93, 748–774.

——— (2015): "The properties of product line prices," *International Journal of Industrial Organization*, 43, 182–188.

——— (2018): "The determinants of product lines," *The RAND Journal of Economics*, 49, 541–573.

Kamenica, E. (2008): "Contextual inference in markets: On the informational content of product lines," *American Economic Review*, 98, 2127–49.

Kolotilin, A. (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–635.

Laffont, J.-J. and J. Tirole (1988): "The dynamics of incentive contracts," *Econometrica*, 1153–1175.

Li, H. and X. Shi (2017): "Discriminatory information disclosure," *American Economic Review*, 107, 3363–85.

Luo, Y., I. Perrigne, and Q. Vuong (2018): "Structural analysis of nonlinear pricing," *Journal of Political Economy*, 126, 2523–2568.

Milgrom, P. and I. Segal (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70, 583–601.

Miravete, E. J. (2002): "Estimating demand for local telephone service with asymmetric information and optional calling plans," *The Review of Economic Studies*, 69, 943–971.

Mussa, M. and S. Rosen (1978): "Monopoly and product quality," *Journal of Economic theory*, 18, 301–317.

Myerson, R. B. (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 10, 67–81.

——— (1986): "Multistage games with communication," *Econometrica: Journal of the Econometric Society*, 323–358.

Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic mechanism design: A myersonian approach," *Econometrica*, 82, 601–653.

Skreta, V. (2006): "Sequentially optimal mechanisms," *The Review of Economic Studies*, 73, 1085–1111.

——— (2015): "Optimal auction design under non-commitment," *Journal of Economic Theory*, 159, 854–890.

Stantcheva, S. (2020): "Dynamic Taxation," *Annual Review of Economics*, 12, 801–831.

Stokey, N. L. (1979): "Intertemporal price discrimination," *The Quarterly Journal of Economics*, 355–371.

Sun, C.-J. (2014): "Dynamic price discrimination with customer recognition," *The BE Journal of Theoretical Economics*, 14, 217–250.

Taylor, C. R. (2004): "Consumer privacy and the market for customer information," *RAND Journal of Economics*, 631–650.

Villas-Boas, J. M. (2004): "Communication strategies and product line design," *Marketing Science*, 23, 304–316.

Xu, Z. and A. Dukes (2019): "Product line design under preference uncertainty using aggregate consumer data," *Marketing Science*, 38, 669–689.

Zhang, J. (2011): "The perils of behavior-based personalization," *Marketing Science*, 30, 170–186.