Financial Product Design in Decentralized Markets*

Marzena Rostek†
University of Wisconsin-Madison

Ji Hee Yoon‡
University College London

July 15, 2024

Abstract

Decentralized trading motivates financial innovation, making synthetic products like derivatives nonredundant, even when all traders trade all assets. This nonredundancy arises because derivatives affect cross-security inference (information) and, in markets with large traders, equilibrium price impact (liquidity). The efficient securities differ from the underlying assets. While the market index/mutual funds are efficient in decentralized markets with competitive investors, heterogeneous portfolios that balance index tracking with liquidity transformation become efficient in markets with large traders. Efficient securities facilitate the trading of all fundamental risks but generally forgo hedging all contingencies to minimize the price impact costs associated with risk sharing and diversification.

JEL Classification: D47, D53, D82, G11, G12

Keywords: Imperfect competition, Decentralized market, Security design, Market design, Liquidity, Price impact, Efficiency

---

*We thank numerous colleagues and participants at the CEPR-Collegio Carlo Alberto Workshop on "Contracts, Incentives, and Information," the Warwick Icebreaking Workshop, and Women in Microstructure Meeting, and seminar participants at Columbia University, Brown University, CUHK-HKU-HKUST, ECARES, Emory University, Harvard&MIT, University of Michigan, Michigan State University, Northwestern University, NYU, University of Pennsylvania&Wharton, Stony Brook University, and UCL for comments and suggestions. Rostek gratefully acknowledges the financial support from the NSF Grant SES-1851857. Yoon acknowledges the financial support of UCL. This paper was edited by Andrew Atkeson.

†Email: `mrostek@ssc.wisc.edu`.

‡E-mail: `jihee.yoon@ucl.ac.uk`.

1 Introduction

Since their inception, synthetic financial products have been one of the great successes of financial markets.[1] Derivatives and related products like ETFs have assumed a significant role in portfolio strategy and risk management. Apart from their potential tax-related benefits and greater transparency, derivatives are typically traded for one of three primary reasons: to improve the liquidity (lower the price impact) of the underlying assets, or to diversify payoffs for contingencies not covered by the underlying assets (i.e., due to some form of market incompleteness), or because trading the underlying assets is restricted.[2]

Regulators often classify the derivatives market as challenging, in part because it is difficult to assess the impact and the scope for innovating nonredundant derivatives. In the standard equilibrium model, the introduction of securities whose payoffs relate linearly to the payoffs of traded assets is neutral to equilibrium and welfare. Textbook methods of derivative pricing either assume the existence of a replicating portfolio, thereby assuming derivatives are redundant or endow a market with an exogenously given demand for a derivative product. An equilibrium model of nonredundant derivatives would help address the following question: "Given the structure of the underlying assets, which derivatives should be introduced in markets?" This paper contributes to the growing literature on nonredundant derivatives, which we will review below.

We start with the following observation. The standard multi-asset equilibrium model of financial markets assumes that traders submit fully *contingent* schedules: The demand for each asset is contingent on the prices of *all* assets, $q_k^i(p_1, ..., p_K) : \mathbb{R}^K \to \mathbb{R}$. While some providers employ contingent orders in practice,[3] such orders are not prevalent. Orders submitted for one asset typically do not depend on the prices of other traded assets, i.e., demands are *uncontingent*, $q_k^i(p_k) : \mathbb{R} \to \mathbb{R}$. The absence of contemporaneous-round cross-asset conditioning in traders' demands not only aligns more closely with market practice but, as we show, also renders derivatives nonredundant. Namely, the introduction of securities that do not alter the traded assets' span changes the traders' equilibrium payoffs, even if the new securities are in zero supply, the underlying assets are not scarce due to regulatory restrictions such as short-selling constraints or margin requirements, and access to the underlying assets is the same for all traders—the assumptions we maintain throughout our analysis.[4]

[1] In 2022, the gross market value of contracts in the derivatives market reached \$18.3 trillion. The total notional amounts outstanding for contracts during the first half of 2022 were estimated at \$632 trillion; `https://stats.bis.org/statx/toc/DER.html`. The daily average turnover of ETF futures and options was \$9.8 trillion in 2022; `https://www.bis.org/statistics/extderiv.htm`. At year-end 2022, assets in ETFs accounted for 22% of the 28.9 trillion total net assets in investment companies (`https://www.ici.org/faqs`).

[2] For an overview of the derivatives market, see, e.g., Stulz (2004).

[3] Types of contingent orders are available in trading platforms for financial assets, such as Active Trader Pro, Etrade, Street Smart, and Tradehawk, as well as in electricity markets and auctions of spectrum.

[4] It is well understood that the introduction of a security with a non-zero supply would change prices and utilities. However, our analysis does not rely on these effects—even with a zero supply, derivatives are nonredundant. Instead, our results point to the role of limited versus full cross-asset demand conditioning for nonredundancy.

Derivative products come in many varieties. We consider the introduction of new securities that capture the defining property of standard derivative products: Their payoffs are linear combinations of the payoffs of traded assets (e.g., ETFs, ETPs, futures). Our analysis is cast in the double-auction model for $K > 1$ assets and $I < \infty$ strategic traders based on the uniform-price mechanism (e.g., Wilson (1979), Klemperer and Meyer (1989), Kyle (1989), Vives (2011)) in the quadratic-Gaussian setting. Traders have private information about their asset holdings. We build on the recently introduced version of this model with uncontingent demand schedules (Chen and Duffie (2021), Rostek and Yoon (2021), Wittwer (2021)) and extend it to allow designs with derivatives and/or the underlying assets.

In hindsight, accounting for endogenous transaction costs due to price impact is critical to the efficient design and regulation of synthetic products. Many markets for financial assets are dominated by large institutional investors whose behavior impacts prices. In practice, dealing with strategic behavior (price impact) often serves as a primary motivation for establishing alternative trading venues or introducing new financial products.[5] Our equilibrium model with large traders allows an explicit treatment of price impact and directly relates to that motive for innovation. We present four sets of results.

Demand conditioning determines how the market clears. With fully contingent demands, all securities in the market must clear *jointly*. In contrast, with uncontingent demands, the market can clear independently for each security; in this sense, trading is decentralized.[6] We first show that when the (sole) assumption that assets are cleared jointly is dispensed with, derivatives become generally nonredundant even if all traders trade all assets (Theorem 1 and Corollary 6). Although their payoffs are neutral to the underlying assets' span, derivatives are payoff-relevant because they affect traders' cross-security price inference. They thus change traders' information in the *total* demand for the underlying assets. In imperfectly competitive markets, another, potentially countervailing, effect of synthetic products arises: Namely, traders have price impact. We show that derivatives improve the liquidity per unit of the underlying assets provided the covariances between these assets are not too heterogeneous (Lemma 3);

[5] E.g., Biais, Bisière, and Spatt (2010), Knight Capital Group (2010), Angel, Harris, and Spatt (2011); see also Bollen and Whaley (2004), Gârleanu, Pedersen, and Poteshman (2009), Frazzini, Israel, and Moskowitz (2018), Zhang (2022), Bali, Beckmeyer, Moerke, and Weigert (2023), and references there. The global derivatives market is notably concentrated, with the top five largest counterparties, excluding the central counterparty (CCP), accounting for nearly 50% of the outstanding notional amount across all asset classes (EU Derivatives Market Report 2023). Significant price impacts in derivative markets have been observed in various markets, including inflation swaps (Bahaj et al. (2023)), synthetic and cash dollar rates in FX derivatives (Wallen (2022)), and commodity ETFs (Todorov (2023)). Furthermore, concerns about price manipulation, both in derivative and underlying asset markets, are closely linked to the observed price impact (e.g., Zhang (2022)).

[6] If traders were able to choose the type of demands they submit, individual optimization would imply the choice of fully contingent schedules, given that they permit conditioning on the actual asset price realizations and subsequent trades. However, in practice, the available types of schedules are determined by the providers, and fully contingent schedules are not commonly offered. Our analysis indicates that providers generally lack incentives to allow fully contingent demands. The benefits derived from appropriately designed securities primarily stem from improving liquidity and increasing trading volume.

otherwise, derivatives may worsen liquidity.[7] Both effects are present except when *all* asset payoffs are either independent or perfectly correlated, regardless of whether the underlying assets are traded.

The nonredundancy of synthetic products implies that collections of securities that share the same span are not payoff-equivalent—in fact, none of them are. Our second set of results examines the efficient design of securities in markets where securities are cleared independently. If efficiency is the objective, which securities should traders trade? Should they opt for an index (a "market portfolio," i.e., the security whose weight is proportional to the average trading needs) or multiple funds (i.e., securities that hedge the same risk as the underlying assets), or could trading factors be efficient, considering that independent payoffs eliminate inference error across securities due to limited demand conditioning? In the absence of weight regulations like the ETF Rule 2019, derivatives offer flexibility in designing security payoffs and hence designing inference error and price impact.

Per the standard (fully contingent) multi-asset equilibrium model's recommendation, only the number of funds with linearly independent payoffs matters so long as they have the same span as the underlying assets. Instead of trading the K underlying assets, traders can equivalently trade the same or a smaller number of funds. Specifically, the same funds are efficient for any investor class, regardless of their trading needs' size, direction, or heterogeneity, as well as the distribution of asset payoffs.

However, these results do not match institutional portfolio data: Institutional investors tend to have heterogeneous portfolio holdings that deviate from the market portfolio. Institutions often hold several dozen stocks and funds, and they frequently work directly with investment banks and issuers to create custom investments tailored to their desired risk profile. (See, e.g., Koijen and Yogo (2019) and Lettau, Ludvigson, and Manoel (2021)).

Dispensing with the assumption that all assets are cleared jointly qualifies the efficiency of trading mutual funds in the competitive ($I \to \infty$) versus imperfectly competitive ($I < \infty$) markets: The same-funds-for-all recommendation does not apply when traders have price impact. In fact, any set of securities can increase or lower the welfare that can be attained with funds—the characteristics of assets (other than their payoff span) and traders (distributions of traders' asset holdings) (Proposition 1) determine which securities are efficient. Furthermore, the efficient design generally involves more than K yet a limited number of securities—*strictly* between K and $\frac{K(K+1)}{2}$ securities (Corollary 3 and Proposition 1).

If traders have no price impact, the same-funds-for-all recommendation continues to hold with independent market clearing, albeit in a stronger sense: Rather than being equivalent to the underlying assets, either sufficiently many funds (at least $\frac{K(K+1)}{2}$) or factors (i.e., securities with independent payoffs) would be *strictly* more efficient than any other securities, irrespective

[7]Derivatives can introduce potentially complex effects on the equilibrium price impact—of the underlying assets as well as those not included in the new security—whose structural properties change when the assumption of fully contingent demands is relaxed (Appendix A.1).

of the characteristics of assets and traders (Corollaries 1 and 2). In particular, when traders have no impact on prices, efficiency entails that all contingencies be hedged.[8] Our results support the promotion of financial innovation across all levels of market competition while indicating an efficiency-based rationale for institutional and small investors to trade different securities.

Third, in markets with any traders and assets, derivatives can be designed so that the market that clears securities independently can reproduce the equilibrium from the market with the underlying assets alone cleared jointly (in which synthetic products would be redundant). Two results are of note: Market structures with securities independent of trader characteristics implement a welfare guarantee in markets with any number of traders (Corollary 1). Furthermore, in markets with any (imperfectly correlated) underlying assets, synthetic products can be designed to *strictly* improve efficiency relative to both such welfare-guaranteeing securities (Proposition 1) and underlying assets (Theorem 2 and Example 3).

Despite allowing inference error, the efficient set of derivatives decreases the costs associated with the endogenous price impact. A key question is which synthetic products mitigate the inefficiencies due to price impact. The efficient design of securities with independent market clearing can be understood by considering two opposing effects that derivatives have on price impact: specifically, on the trading costs due to own-asset and cross-asset price impact, which affect sharing and diversification, respectively. Derivatives cannot be too strongly correlated since higher correlation increases the own-asset price impact due to cross-asset inference.[9] Consequently, when assets are not cleared jointly, using an index or funds that align with traders' *ex ante* trading needs (or *target portfolios*;[10] Gârleanu and Pedersen (2013), Kyle, Obzihaeva, and Wang (2017)) may be too costly in terms of the own-asset price impact. On the other hand, trading factors, which have independent payoffs, forgoes the benefit of designing cross-asset price impact to enhance diversification: Derivatives can transform price impact avoidance into price impact seeking when it comes to cross-asset price impacts, as appropriate derivative weights can induce the desired sign of per-unit cross-asset price impact given the trading positions across assets.

In the competitive market, financial innovation's efficiency role lies in the elimination of inference errors among the underlying assets to replicate the fully contingent outcome. Our analysis indicates that in markets with large traders where securities are cleared independently,

[8]In a competitive market, if full demand conditioning were restricted and hence nonredundant innovation were allowed, the only result would be the introduction of inference errors, leading to a decrease in overall welfare. However, in markets with large traders, cross-asset price inference modifies the traders' price impacts within and across underlying assets. These changes in liquidity can offset the inefficiency due to inference error, and the efficient design depends on market characteristics.

[9]This finding corresponds with the 2019 revision of the SEC ETF rule, which permits the composition of baskets to diverge from closely mirroring indices or diversified portfolios, a measure intended to boost liquidity (SEC (2019, Section II.C.5)). It also coincides with the strategy of actively managing corporate bond ETFs by adjusting their baskets to deviate from index tracking, thereby reducing transaction costs (Koont, Ma, Pastor, and Zeng (2023)).

[10]Traders' target portfolios coincide with the market portfolio when trading needs are symmetric across assets (i.e., $E[\mathbf{q}_0^i] = \xi_{ij} E[\mathbf{q}_0^j]$ for some $\xi_{ij} \in \mathbb{R}$, for all i and $j \neq i$) but not otherwise.

the role of securities that do not alter the underlying asset span is not to mimic diversification facilitated by these assets but to (1) improve upon it and (2) enhance risk sharing alongside diversification. As a general principle, in contrast to markets where all assets are cleared jointly, efficient securities (i) must correlate security payoffs to enhance portfolio diversification while compromising on own-asset price impact, (ii) correlate differently from the underlying assets, and (iii) imperfectly so, thus letting traders trade in the direction "around" the target rather than perfectly matching it, thus reducing the own-asset price impact; and (iv) the efficient security correlations differ between one- and two-sided markets and depend on the correlations among the underlying assets (Theorem 2).

Taken together, our results show that decentralized trading weakens the role of spanning methods—which represent risk through state-contingent securities. In particular, when markets do not clear all assets jointly, the implied representation of risk depends on the endogenous price impact and is thus not invariant to counterfactual changes in market structure. We discuss some of the implications for asset pricing and financial innovation in Section 7.

Fourth, decentralized trading not only motivates the design of financial innovation within the underlying asset payoff span but also encourages suitable innovation on unspanned risks. To show this, we also consider markets where the payoffs of K traded assets only span a subset of $Z > K$ fundamental risks. In contrast to markets that clear assets jointly, where the introduction of *arbitrary* securities for unspanned risk (i.e., in the span of $Z \setminus K$ assets) always increases welfare with quasilinear utilities, the introduction of such securities may lower welfare when trading is decentralized (Proposition 2).[11] Nonetheless, if the securities on the Z risks can be designed, then opening markets with fewer than Z securities is inefficient; however, listing fewer than $\frac{Z(Z+1)}{2}$ securities is generally efficient when traders face price impact. Thus, our results underscore the derivatives' role in treating the incompleteness associated with unspanned fundamental risks differently than limited information given the span: Allowing trading on *all* fundamental risks while *limiting* the cross-asset information traders' demands can condition on is efficient. The key difference lies in the fact that fundamental risks can be factorized, allowing for the introduction of securities without creating an externality on the traded assets' price impact—but risks due to imperfect information cannot.

Related literature. This paper highlights decentralized trading as a source of synthetic products' nonredundancy: Accounting solely for the fact that assets are not all (or typically) cleared jointly in financial markets implies that derivatives are nonredundant, except under trivial conditions. This perspective complements the growing body of research that investigates the nonredundancy of securities brought by fixed costs of issuing them (Allen and Gale (1988)), differences in margin requirements between derivatives and the underlying assets (Gârleanu and

[11] This result thus differs from the competitive equilibrium price effects in the literature following Hart (1975), which are absent with quasilinear utilities. The effect we investigate, which applies to innovation within and outside the traded assets' span, can be seen as an imperfectly competitive counterpart of these effects: Innovation is not neutral if and only if it alters the relative *price impacts* of the traded assets rather than their price levels.

Pedersen (2011)), derivatives' ability to relax binding short-sale constraints when the underlying security is scarce (Banerjee and Graveline (2014)), heterogeneous beliefs among the traders (Fostel and Geanakoplos (2012), Simsek (2013), Che and Sethi (2014), Oehmke and Zawadowski (2015)), and limited pledgeability (Biais, Hombert, and Weill (2021)).

More generally, our results underscore the role of financial innovation in imperfectly competitive markets, particularly in enhancing risk sharing and diversification for spanned risks. The competitive markets literature has emphasized the value of financial innovation in improving the risk-sharing of unspanned risks; Allen and Gale (1994) and Duffie and Rahi (1995) provide surveys on the spanning motive as a determinant of financial innovation.

There has been significant interest in understanding how market fragmentation influences security design (e.g., Rahi and Zigrand (2009), Babus and Hachem (2021, 2023), Biais, Hombert, and Weill (2021); see also Cabrales, Gale, and Gottardi (2015)). In decentralized market settings, prior literature has examined the impact of derivatives on financial stability (Allen and Carletti (2006)) and on facilitating the hedging of counterparty exposures within financial networks (Zawadowski (2013)). While these studies consider market fragmentation in the sense of limited trader participation, our results instead explore the effects of market fragmentation in the sense of limited demand conditioning. Thus, our analysis most directly applies to exchange-traded products, but analogous results will apply to over-the-counter derivative markets, where both types of effects are present. In this sense, one expects the derivatives' nonredundancy and its implications to be even more robust in other fragmented market structures.

Our paper also contributes to the literature on financial market design with large traders.[12] Several authors have recently challenged the assumption that assets are cleared jointly to explore the potential welfare benefits of market fragmentation (Chen and Duffie (2021), Rostek and Yoon (2021), Wittwer (2021)).[13] In both this paper and Rostek and Yoon (2021), we observe that accounting for independence in market clearing among traded assets is not only more realistic but also makes nonredundant financial innovation that does not alter the underlying asset span, traders' initial asset holdings, or asset supply—spanning no longer holds. Rostek and Yoon (2021) studies the design of market-clearing technology in a model that incorporates contingent demands for subsets of assets traded in different exchanges, without bundling asset payoffs. Changes in market-clearing technology encompass the introduction of new exchanges

[12] E.g., Du and Zhu (2017a,b), Malamud and Rostek (2017), Kyle and Lee (2017, 2022), Babus and Kondor (2018), Baisa and Burkett (2018), Duffie (2018), Yang and Zhu (2020, 2021), Antill and Duffie (2021), Baldauf and Mollner (2021), Rostek and Yoon (2021), Somogyi (2021), Babus and Hachem (2021, 2023), Babus and Parlatore (2022), Cespa and Vives (2022), Chen (2022), Zhang (2022), Allen and Wittwer (2023).

[13] Wittwer (2021) and Rostek and Yoon (2021) focus on the welfare implications of allowing traders to submit orders contingent on cross-exchange prices. These studies show when market fragmentation, understood as a departure from fully contingent demands, can enhance efficiency. Chen and Duffie (2021) consider orders contingent on prices within the same exchange to assess how fragmentation, understood as increasing the number of exchanges, influences allocative efficiency and price informativeness, also incorporating a dynamic market perspective. The authors demonstrate that sufficient fragmentation can lead to allocative efficiency. Cespa (2004) examined competitive uncontingent markets with two assets and noise traders and also characterized how uncontingent trading affects price informativeness.

for underlying assets, mergers of exchanges, and the listings of assets in exchanges where they were not previously traded. This paper instead examines the efficient design of financial products in fragmented markets. While technology innovations enable the joint clearing of some assets by modifying traders' strategies, synthetic products cleared independently change the joint distribution of securities' payoffs. Our paper is the first to draw a connection between the independent clearing of securities and the nonredundancy of derivatives, while also identifying how market fragmentation in this "minimal" sense revises the efficient security design.

Recognizing that assets are not cleared jointly in practice motivates the study of various other types of innovation. In particular, the design of exchanges (contingent demands for asset subsets), the design of financial products (uncontingent demands for asset bundles), and the joint design of both innovations are not equivalent for traders' payoffs and involve synergies and tradeoffs. Essentially, these designs transform the game and, consequently, the equilibrium price impact in different ways, as discussed by Rostek and Yoon (2024a); see also ft. 30. Although motivated by different considerations from market fragmentation, the design proposed by Budish, Cramton, Kyle, Lee, and Malec (2021) serves as another example of innovation that remains neutral to the asset span and is not equivalent to the aforementioned innovations: It involves each trader submitting a demand $\mathbb{R} \to \mathbb{R}$ for a single asset portfolio of their choice rather than demands for individual assets.

2 Model

Our model is based on the uniform-price double auction in the quadratic-Gaussian setting. Unlike the standard multi-asset version of that model, where all assets are cleared jointly, we consider markets where assets are cleared independently.[14]

2.1 Traders, Assets, and Synthetic Products

Consider a market with $I \geq 3$ traders and K risky assets. The payoffs of the K risky assets are jointly normally distributed $\mathbf{r} \equiv (r_k)_k \sim \mathcal{N}(\boldsymbol{\delta}, \boldsymbol{\Sigma})$ with a vector of expected payoffs $\boldsymbol{\delta} \equiv (\delta_k)_k \in \mathbb{R}^K$ and a positive semidefinite covariance matrix $\boldsymbol{\Sigma} \equiv (\sigma_{k\ell})_{k,\ell} \in \mathbb{R}^{K \times K}$. There is also a riskless asset (a numéraire).

Each trader i has a quadratic in the quantity of risky assets (mean-variance) utility:

$$u^i(\mathbf{q}^i) = \boldsymbol{\delta} \cdot (\mathbf{q}^i + \mathbf{q}^i_0) - \frac{\alpha}{2}(\mathbf{q}^i + \mathbf{q}^i_0) \cdot \boldsymbol{\Sigma}(\mathbf{q}^i + \mathbf{q}^i_0), \tag{1}$$

where $\mathbf{q}^i = (q^i_k)_k \in \mathbb{R}^K$ is the vector of trader i's traded quantities of risky assets, $\mathbf{q}^i_0 = (q^i_{0,k})_k \in$

[14] Our main characterization result, Theorem 3 and the corresponding lemmas in Appendix A.1 extend the equilibrium characterization in Rostek and Yoon (2021) to markets where synthetic products are traded alongside or instead of the underlying assets. Example 1 Cont'd in Appendix A.2 compares equilibrium with fully contingent and uncontingent demands.

$\mathbb{R}^K$ represents the units of risky assets with which trader i is initially endowed, and $\alpha \in \mathbb{R}_+$ is the risk aversion parameter. Asset holdings $\mathbf{q}_0^i$ are trader i's private information distributed according to $\mathcal{N}(E[\mathbf{q}_0^i], \sigma_q^2\mathbf{Id})$ and are independent across assets and independent of asset payoffs $\mathbf{r}$. Specifically, trader i is uncertain about other traders' asset holdings $\{\mathbf{q}_0^j\}_{j\neq i}$, and thus, about the per-capita aggregate asset holdings $\overline{\mathbf{q}}_0 = \frac{1}{I}\sum_j \mathbf{q}_0^j$ (equivalently, prices).[15] Given that the traders share the same prior on the asset payoffs $(\boldsymbol{\delta}, \boldsymbol{\Sigma})$ and risk aversion α, gains from trade come from the desire to hedge and diversify risk in asset holdings, which are heterogeneous.

In markets with K assets, we consider the introduction of $D \geq 0$ synthetic products (derivatives) whose payoffs derive from the K underlying assets' payoffs without altering the asset span: The payoff of derivative d is a linear combination of asset payoffs $r_d = \mathbf{w}_d'\mathbf{r} \in span(\{r_k\}_k)$ for some weight vector $\mathbf{w}_d \equiv (w_{dk})_k \in \mathbb{R}^K$, where $w_{dk} \in \mathbb{R}$ for any $k \in K$ and $d \in D$. In particular, weights can be of any sign, i.e., a derivative can be a combination of long and short positions of the assets. The *span* of random variables $\{r_k\}_k$ is the set of random variables that are linear combinations of $\{r_k\}_k$,

$$span(\{r_k\}_k) \equiv \{r = \sum\nolimits_k w_k r_k : w_k \in \mathbb{R} \text{ for each } k\}. \tag{2}$$

The payoffs of $K + D$ *securities* (assets and derivatives) are thus jointly normally distributed according to $\mathcal{N}(\boldsymbol{\delta}^+, \boldsymbol{\Sigma}^+)$ with the moments

$$\boldsymbol{\delta}^+ \equiv \begin{bmatrix} \boldsymbol{\delta} \\ \mathbf{W}_d'\boldsymbol{\delta} \end{bmatrix} \in \mathbb{R}^{K+D} \qquad \text{and} \qquad \boldsymbol{\Sigma}^+ \equiv \begin{bmatrix} \boldsymbol{\Sigma} & \boldsymbol{\Sigma}\mathbf{W}_d \\ \mathbf{W}_d'\boldsymbol{\Sigma} & \mathbf{W}_d'\boldsymbol{\Sigma}\mathbf{W}_d \end{bmatrix} \in \mathbb{R}^{(K+D)\times(K+D)}, \tag{3}$$

where in the weight matrix $\mathbf{W}_d \equiv (\mathbf{w}_1, \cdots, \mathbf{w}_D) \in \mathbb{R}^{K\times D}$, the d^{th} column $\mathbf{w}_d$ corresponds to the d^{th} derivative.

In practice, synthetic products are traded because they help increase the liquidity of the underlying assets or because the underlying assets are harder to trade. Accordingly, our model encompasses two types of security innovations in Eq. (3) whose payoffs lie in the span of the payoffs of the K underlying assets:

- The introduction of D securities to be traded with the K underlying assets; Example 4 in Appendix A.2, i.e., the weight matrix is $\mathbf{W} = \begin{bmatrix} \mathbf{Id} & \mathbf{W}_d \end{bmatrix} \in \mathbb{R}^{K\times(K+D)}$.
- A delisting of some or all K underlying assets while listing new securities whose payoffs have the same span, i.e., $\mathbf{W} \in \mathbb{R}^{K\times N}$ for $N \leq K + D$ securities to be arbitrary; Example 2.

The introduction of derivatives does not change the market primitives: Traders' initial holdings for the derivatives are zero: $q_{0,d}^i = 0$ for all i and d and the derivative supply is zero.

[15] To ensure that $\overline{\mathbf{q}}_0$ is random in the limit large market ($I \to \infty$), we allow traders' asset holdings to be correlated across traders through a common value component (Eq. (19) in Appendix A.1).

We make this assumption to focus the analysis on the core reasons for the nonredundancy, which are not related to the non-zero supply of derivatives. Given a vector of trades of the K assets, $\mathbf{q}_a^i \equiv (q_k^i)_k \in \mathbb{R}^K$, and a vector of trades of D derivatives, $\mathbf{q}_d^i \equiv (q_d^i)_d \in \mathbb{R}^D$, the utility of trader i as a function of the vector of total trades $\mathbf{q}_a^i + \mathbf{W}_d \mathbf{q}_d^i$ is $u^i(\mathbf{q}_a^i + \mathbf{W}_d \mathbf{q}_d^i)$, using (1). By the definition of $\boldsymbol{\delta}^+$ and $\boldsymbol{\Sigma}^+$ in Eq. (3), we can equivalently treat the derivatives in utility (1) as distinct assets:

$$u^i(\mathbf{q}^i) = \boldsymbol{\delta}^+ \cdot (\mathbf{q}^i + \mathbf{q}_0^{i,+}) - \frac{\alpha}{2}(\mathbf{q}^i + \mathbf{q}_0^{i,+}) \cdot \boldsymbol{\Sigma}^+ (\mathbf{q}^i + \mathbf{q}_0^{i,+}), \tag{4}$$

where $\mathbf{q}^i = (\mathbf{q}_a^i, \mathbf{q}_d^i) \in \mathbb{R}^{K+D}$ is the vector of trades for all $K + D$ securities and $\mathbf{q}_0^{i,+} = (\mathbf{q}_0^i, \mathbf{0}) \in \mathbb{R}^{K+D}$ is the asset holdings vector whose elements corresponding to derivatives are zeros. Equilibrium is invariant to any split of initial asset holdings among the underlying assets and synthetic products.

Remark. In the rest of this section and in Section 3, we omit the superscript '+' unless it is helpful.

2.2 Market Structure

Each of the $N \leq K + D$ securities is traded in an *exchange*. All traders participate in all exchanges (i.e., they trade all securities). The securities are cleared independently across N exchanges. Each exchange n is organized as a uniform-price double auction (e.g., Kyle (1989), Vives (2011)) in which traders submit (net) demand schedules. The schedule $q_n^i(\cdot) : \mathbb{R} \to \mathbb{R}$ represents the limit orders (i.e., a menu of (q_n^i, p_n)) a trader submits, each specifying a quantity demanded q_n^i for the price p_n that may occur.

Definition 1 (Double Auction) *Each trader i submits N* uncontingent *demand schedules $\mathbf{q}^i(\cdot) \equiv (q_1^i(p_1), \ldots, q_N^i(p_N))$, with $q_n^i(\cdot) : \mathbb{R} \to \mathbb{R}$ specifying the quantity of security n demanded for any realization of price p_n. If $q_n^i > 0$, the trader is a buyer of security n; if $q_n^i < 0$, the trader is a seller. The market clears* independently *across securities: The market-clearing price p_n is determined by the zero aggregate net demand in exchange n, $\sum_i q_n^i(p_n) = 0$.[16] Trader i trades $\{q_n^i\}_n$, pays $\sum_n p_n q_n^i$, and receives a payoff of $u^i(\mathbf{q}^i) - \mathbf{p} \cdot \mathbf{q}^i$.*

The standard multi-asset equilibrium model is instead based on fully *contingent schedules*: Each trader i submits N demand schedules $\mathbf{q}^{i,c}(\cdot) \equiv (q_1^{i,c}(\mathbf{p}), \ldots, q_N^{i,c}(\mathbf{p}))$, with the demand $q_n^{i,c}(\cdot) : \mathbb{R}^N \to \mathbb{R}$ for each security n specifying the quantity demanded of that security for all price realizations of all assets $\mathbf{p} \equiv (p_1, \ldots, p_N)$. With fully contingent demands, securities must be cleared *jointly*: Market-clearing applies to all assets simultaneously to determine the

[16] As is customary in defining the game, if p_n is such that $\sum_i q_n^i(p_n) = 0$ does not exist or is not unique, the market ends with no trade, i.e., $q_n^i = 0$ for all i.

equilibrium price vector $\sum_i \mathbf{q}^{i,c}(p_1, \cdots, p_N) = \mathbf{0} \in \mathbb{R}^N$.[17] This paper shows that the fact that not all assets are cleared jointly in financial markets implies that synthetic products that would be neutral under joint market clearing become nonredundant.

We study the Bayesian Nash Equilibrium in linear demand schedules.[18] All traders are strategic and consider the impact of their demands on prices.

Definition 2 (Equilibrium) *A profile of (net) demand schedules* $\{\{q_n^i(\cdot)\}_n\}_i$ *is a linear Bayesian Nash equilibrium if, for each* i, $\{q_n^i(\cdot)\}_n$ *maximizes the expected payoff:*

$$\max_{\{q_n^i(\cdot)\}_n} E[\boldsymbol{\delta} \cdot (\mathbf{q}^i + \mathbf{q}_0^i) - \frac{\alpha}{2}(\mathbf{q}^i + \mathbf{q}_0^i) \cdot \boldsymbol{\Sigma}(\mathbf{q}^i + \mathbf{q}_0^i) - \mathbf{p} \cdot \mathbf{q}^i | \mathbf{q}_0^i], \tag{5}$$

given other traders' schedules $\{\{q_n^j(\cdot)\}_n\}_{j \neq i}$ *and market clearing* $\sum_j q_n^j(\cdot) = 0$ *for all* n.

Remark. This paper focuses on derivatives whose payoffs lie in the linear span of underlying asset payoffs (Eq. (2)) in the quadratic-Gaussian setting that supports equilibrium with linear demands and hence tractability of equilibrium, welfare, and design analysis of derivatives. The assumption on the linear combination of asset payoffs also facilitates a clear demonstration of how accounting for market fragmentation in the sense of independent market clearing weakens the spanning methods; see Section 7.1.

ETFs, ETPs, and futures are examples of synthetic products whose payoffs are in the linear span of the underlying assets. Although we do not study exotic derivatives and options whose payoffs are non-linear functions of the underlying asset payoffs, the qualitative effects we report for welfare results and portfolio analysis will apply there; see Section 7.2.

In addition, we assume that the introduction of derivatives does not change the markets' trading mechanism (Definition 1). Consequently, our analysis focuses on exchange-traded products. However, the main effects we explore do not depend on a specific price mechanism and will be present in over-the-counter derivative markets with private information; see Section 7.3.

3 Equilibrium

In this paper, we show that decentralized trading—the independent rather than joint clearing of assets—leads to the nonneutrality of security innovations (Section 4; e.g., Examples 2 and 4), alters the efficient portfolio recommendations (Section 5), and impacts asset pricing properties—linear pricing no longer holds for security innovations (Section 7.1). In this section, we explain

[17]Some electronic trading platforms for financial assets, including Active Trader Pro, Etrade, Street Smart, and Tradehawk, offer traders the ability to express their demands for an asset contingent on the prices of other assets. Still, such conditioning is limited to only a small number of assets and need not be mutual across assets; see also ft. 6. The effects we explore in this paper are applicable as long as not all asset demands are fully contingent).

[18]Equilibrium is *linear* if schedules have the functional form of $\mathbf{q}^i(\cdot) = \boldsymbol{a}^i - \mathbf{B}^i \mathbf{q}_0^i - \mathbf{C}^i \mathbf{p}$.

the mechanisms that motivate security innovation in a market with $N \leq K + D$ securities, where D derivatives can be traded with all or only some of the K underlying assets (Section 3.1). Example 2 illustrates these effects (Section 3.2).

3.1 Equilibrium Effects of Independent Market Clearing

Traders submit their schedules before prices are realized. Nevertheless, as is well known, equilibrium is *ex post* when demands are fully contingent:[19] A trader can express his demand for security n as a function of the actual quantities traded of all other securities that will be realized, and the price vector is one-to-one with the quantity vector. The key implication of independent market clearing is price uncertainty, which affects traders' equilibrium price impacts: A trader's demand for security n is contingent on, and hence measurable only with respect to the price p_n. Given the technological constraint on the schedules traders can submit, as captured by the uncontingent demands, a trader's optimization differs from maximizing the *ex post* payoff.

Optimization problem and cross-security inference. Consider the optimization problem (5) of trader i who submits (net) demand schedules $q^i_n(\cdot) : \mathbb{R} \to \mathbb{R}$ in $N > 1$ exchanges, each for one security n (an asset or a derivative). The first-order conditions of trader i equalize, for each security n, his *expected* marginal utility with *expected* marginal payment pointwise for all $p_n \in \mathbb{R}$: Denoting the n^{th} row of $\mathbf{A}$ by $\mathbf{A}_n$, these conditions can be written as:

$$\delta_n - \alpha \mathbf{\Sigma}_n E[\mathbf{q}^i + \mathbf{q}^i_0 | p_n, \mathbf{q}^i_0] = p_n + \lambda^i_n q^i_n \qquad \forall p_n \in \mathbb{R}, \tag{6}$$

each taking as given the trader's residual market $\{\{q^j_\ell(\cdot)\}_\ell\}_{j \neq i}$ and his own demands for other securities $\{q^i_\ell(\cdot)\}_{\ell \neq n}$. Here, $\lambda^i_n \equiv \frac{dp_n}{dq^i_n} \in \mathbb{R}_+$ represents the *price impact* of trader i (i.e., "Kyle's lambda") for security n. The second-order condition $(-\alpha\mathbf{\Sigma} - \mathbf{\Lambda}^i - (\mathbf{\Lambda}^i)' < 0)$ is satisfied when $\lambda^i_n > 0$ for all i and n, or equivalently, when traders' demands are downward-sloping, i.e., $\frac{\partial q^i_n(\cdot)}{\partial p_n} < 0$ for all i and n (Eq. (7)).

Since schedules are not fully contingent, a trader's demand for security n depends on his *expected trades* $E[q^i_\ell(p_\ell)|p_n, \mathbf{q}^i_0]$ (equivalently *expected prices* $E[p_\ell|p_n, \mathbf{q}^i_0]$) of other securities $\ell \neq n$ (Eq. (6)).[20] Cross-asset inference becomes relevant, except when the payoff Hessian ($\mathbf{\Sigma}$) is separable, i.e., the security payoffs are independent. As a result, even though all securities are cleared independently, the equilibrium outcome is not independent across exchanges. When the market is imperfectly competitive ($I < \infty$), the imperfect cross-security price inference under

[19]Equilibrium is *ex post* if the equilibrium demands $\{q^i_n(\cdot; \mathbf{q}^i_0)\}_n$ are optimal for all i, for all realizations of asset holdings of all traders $\{\mathbf{q}^j_0\}_j$:

$$\{q^i_n(\cdot; \mathbf{q}^i_0)\}_n = \underset{\{q^i_n(\cdot)\}_n}{\operatorname{argmax}} \; E[\boldsymbol{\delta} \cdot (\mathbf{q}^i + \mathbf{q}^i_0) - \frac{\alpha}{2}(\mathbf{q}^i + \mathbf{q}^i_0) \cdot \mathbf{\Sigma}(\mathbf{q}^i + \mathbf{q}^i_0) - \mathbf{p} \cdot \mathbf{q}^i | \{\mathbf{q}^j_0\}_j].$$

[20]Intuitively, with uncontingent schedules, a trader's choices are based on his best estimate of the prices: Price p_n contains information about the aggregate holdings for security n and, due to the nonseparability of $\mathbf{\Sigma}$, also contains information about the aggregate holdings for all other securities and hence their respective prices.

independent market clearing also affects the equilibrium price impacts $\{\lambda_n^i\}_n$ (Eq. (7)).

Residual market and price impact. Trader i's price impact λ_n^i in the exchange for security n is determined by the slope of the inverse residual supply:

$$\lambda_n^i = -(\sum_{j\neq i} \frac{\partial q_n^j(\cdot)}{\partial p_n})^{-1}. \tag{7}$$

Since the market clears independently across securities, the cross-asset price impacts $\{\lambda_{n\ell}^i \equiv \frac{dp_\ell}{dq_n^i}\}_{n,\ell\neq n}$ of traders are zero. Therefore, the price impact matrices of all traders are diagonal: $\mathbf{\Lambda}^i \equiv diag(\lambda_n^i)_n \in \mathbb{R}^{N\times N}$ for all i. Here, $diag(x_n)_n = diag(x_1,\cdots,x_N)$ denotes a diagonal matrix in $\mathbb{R}^{N\times N}$, where the n^{th} diagonal element is x_n and all off-diagonal elements are zero. The interdependence among traders' demands for different securities manifests itself through the price impacts *within* the exchanges, which, in contrast to joint clearing (see Example 1) depend on cross-security price inference.

Example 1 (Market That Clears All Securities Jointly) Suppose that, in the same market, traders submit fully contingent (net) demand schedules for N securities. The familiar counterparts of optimization and price impact conditions (6) and (7) are: for all i,

$$\delta_n - \alpha\mathbf{\Sigma}_n(\mathbf{q}^i + \mathbf{q}_0^i) = p_n + \mathbf{\Lambda}_n^i\mathbf{q}^i \qquad \forall \mathbf{p} \in \mathbb{R}^N, \tag{8}$$

where $\mathbf{\Lambda}^i \equiv (\frac{d\mathbf{p}}{d\mathbf{q}^i})' \in \mathbb{R}^{N\times N}$ is the *price impact* of trader i and

$$\mathbf{\Lambda}^i = -((\sum_{j\neq i} \frac{\partial \mathbf{q}^j(\cdot)}{\partial \mathbf{p}})^{-1})' = \frac{\alpha}{I-2}\mathbf{\Sigma}. \tag{9}$$

□

Independent market clearing transforms the relationship between traders' incentives (the equilibrium price impact $\mathbf{\Lambda}^i$) and the fundamental risk (covariance $\mathbf{\Sigma}$) in two ways: Due to the imperfect cross-asset price inference, the price impact (i) depends not only on the fundamental asset covariance but also on the distribution of traders' privately known initial holdings, and (ii) is no longer separable across assets (i.e., the price impact for any pair of assets depends on the covariance of *all* asset payoffs). Mathematically, the transformation of risk is captured by the lack of proportionality between $\mathbf{\Lambda}^i$ and $\mathbf{\Sigma}$, which holds in the case of joint clearing (Eq. (9)).

Equilibrium. As is common in demand submission games, equilibrium is characterized by two conditions for all traders: (i) a trader's optimization (the first-order condition (6)), given his residual supply; and (ii) the requirement that trader's residual supply is correct (Eq. (7)). Our main characterization result, Theorem 3 in Appendix A.1, shows that the profile of all traders' price impact matrices $\{\mathbf{\Lambda}^i\}_i$ is sufficient for equilibrium. In particular, the inference's effect on equilibrium is accounted for by a profile of traders' price impacts.[21] We will highlight the crucial role that price impact plays in security design with independent market clearing.

Definition 3 (Competitive Market, Competitive Equilibrium) *Letting* $\{\mathbf{q}^{i,I}(\cdot)\}_i$ *be the equilibrium in the market with* $I < \infty$ *traders, the* competitive equilibrium $\{\mathbf{q}^{i}(\cdot)\}_i$ *is the limit of equilibria* $\{\mathbf{q}^{i,I}(\cdot)\}_i$ *as* $I \to \infty$*:* $\mathbf{q}^{i}(\cdot) = \lim_{I\to\infty} \mathbf{q}^{i,I}(\cdot)$ *for all* i*.*

In the competitive market, $\mathbf{\Lambda} \to \mathbf{0}$ as $I \to \infty$ (Lemma 1 in Appendix A.1). While traders' price inference is imperfect with independent market clearing irrespective of the market size (see Eq. (6) with $\lambda^i_n = 0$ for all i and n), security innovation plays a limited role in efficient design when traders are price-takers.[22]

Proposition 5 in Appendix A.4 shows that an equilibrium exists and is unique in markets with K symmetric assets and D derivatives with symmetric weights (Definition 5).[23]

3.2 Motivating Example

Example 2 elaborates on the equilibrium effects of traders' cross-security inference and price impact to illustrate a key result: Independence in market clearing renders financial innovation that is neutral to security payoff span nonneutral to traders' equilibrium payoffs (Section 4).

Example 2 (Assets vs. Factors) Consider $K = 2$ imperfectly correlated assets with covariance $\mathbf{\Sigma} = \begin{bmatrix} \sigma^2 & \sigma^2\rho \\ \sigma^2\rho & \sigma^2 \end{bmatrix}$. Suppose that instead of trading the underlying assets, traders engage in trading two *factors* with independent payoffs $r^+_1 = w_{11}r_1 + w_{12}r_2$ and $r^+_2 = w_{21}r_1 + w_{22}r_2$ whose span coincides with that of the underlying assets' payoffs $\{r_1, r_2\}$. The factors are constructed with the weight matrix $\mathbf{W} \in \mathbb{R}^{K\times K}$ that orthogonalizes the asset covariance $\mathbf{\Sigma}$: The joint distribution of the K securities is $\mathcal{N}(\boldsymbol{\delta}^+, \mathbf{\Sigma}^+)$, where $\mathbf{\Sigma}^+ \equiv \mathbf{W}'\mathbf{\Sigma}\mathbf{W} = \sigma^2\mathbf{Id}$ is a diagonal matrix and $\boldsymbol{\delta}^+ = \mathbf{W}'\boldsymbol{\delta}$; $\mathbf{q}^{i,+}_0 \equiv \mathbf{W}^{-1}\mathbf{q}^i_0$ for all i are each trader's initial holdings $\mathbf{q}^i_0$ in units of the traded securities.

We show that with independent market clearing (part (i)), but not with joint market clearing (part (ii)), securities whose prices are imperfectly correlated give rise to cross-asset inference effects and thus alter traders' price impacts and outcomes.

[21] This result holds even though, in contrast to markets with fully contingent demands, a trader's price impact is not sufficient for his best-response problem due to imperfect cross-security price inference.

[22] In competitive markets, equilibrium is characterized by a fixed point in the demand Jacobian (substitution) matrix $\mathbf{C} \equiv diag(\frac{\partial q^i_n(\cdot)}{\partial p_n})_n$, thus with zero off-diagonal elements (vs. $\mathbf{C}^c = (\alpha\mathbf{\Sigma})^{-1}$ with fully contingent demands). Similarly, in imperfectly competitive markets ($I < \infty$, Theorem 3 in Appendix A.1), the equilibrium fixed point for $\mathbf{\Lambda}$ can be written as a fixed point for $\mathbf{C}$, given the one-to-one map between $\mathbf{\Lambda}$ and $\mathbf{C}$ (Eq. (7)).

[23] In extensive simulations that consider asymmetric assets or weights, the numerical iteration that solves the equilibrium fixed point equation (Theorem 3 in Appendix A.1) converges to the same equilibrium, allowing for random initial values, using different forms of the fixed point equation, and considering fixed points defined with respect to different variables.

(i) (Independent market clearing) Suppose the market clears securities independently. The first-order condition (6) for the underlying assets with payoffs $\{r_1, r_2\}$ is

$$(\alpha\sigma^2 + \lambda_1)q_1^i = \delta_1 - p_1 - \alpha\sigma^2 q_{0,1}^i - \alpha\sigma^2\rho q_{0,2}^i - \alpha\sigma^2\rho E[q_2^i|p_1, \mathbf{q}_0^i], \tag{10}$$

and the first-order condition for the factors with payoffs $\{r_1^+, r_2^+\}$ is

$$(\alpha\sigma^2 + \lambda_1^+)q_1^{i,+} = \delta_1^+ - p_1^+ - \alpha\sigma^2 q_{0,1}^{i,+} - \underbrace{\alpha\sigma^2\rho^+ E[q_2^{i,+}|p_1^+, \mathbf{q}_0^i]}_{=0}.$$

The change in security payoffs via $\mathbf{W}$ is not neutral to the first-order condition. This is because the change from the correlated contingent variables in traders' demands $\{p_1, p_2\}$ to independent variables $\{p_1^+, p_2^+\}$ eliminates cross-asset inference (i.e., $E[q_2^i|p_1, \mathbf{q}_0^i]$). With no cross-asset inference, the price impact coincides with its contingent-market counterpart, $\mathbf{\Lambda}^+ = \mathbf{\Lambda}^{c,+}$. The *total* equilibrium trade of underlying asset 1 in both exchanges is the same as with fully contingent demands, i.e., $w_{11}q_1^{i,+} + w_{21}q_2^{i,+} = q_1^{i,c}$, where $\mathbf{q}^{i,c} = (q_1^{i,c}, q_2^{i,c})$ is the equilibrium outcome for the underlying assets with joint market clearing.

This equilibrium outcome fails to satisfy the first-order condition (10) in the market for the underlying assets unless $\rho = 0$ (i.e., the underlying assets are independent, then $\lambda_1 = \lambda_1^c$) or $|\rho| = 1$ (the underlying assets are perfectly correlated, then, $\lambda_1 = 2\lambda_1^c$ and $E[q_2^i|p_1, \mathbf{q}_0^i] = q_2^i$). The changes to price impact and expected trades do not generally offset each other since the fixed point between inference and price impact is non-linear (Eq. (7)).

(ii) (Joint market clearing) If the market clears all securities jointly, the standard result applies: Trading factors is payoff-equivalent to trading the underlying assets. This equivalence can be seen from the fact that the first-order condition (8) is invariant to changes in security payoffs through $\mathbf{W}$, as neither cross-asset inference nor the price impact per unit of the underlying assets are affected. Specifically, changes to security payoffs via $\mathbf{W}$ preserve the *ex post* property of the equilibrium with fully contingent demands. Furthermore, due to the proportional relationship between the equilibrium price impact $\mathbf{\Lambda}$ and the covariance $\mathbf{\Sigma}$ (Eq. (9)), traders' marginal utility and marginal payment per unit of the underlying assets are invariant to such simultaneous transformations. See Example 2 Cont'd in Appendix A.2 for further details. □

Example 2 shows that allowing traders to trade securities with independent payoffs instead of the underlying assets is not neutral to welfare. Theorem 1 establishes that no set of K securities with the same span as the underlying assets is payoff-equivalent. This nonredundancy of security innovation under independent market clearing holds in both competitive and imperfectly competitive markets. The redundancy of these securities with fully contingent demands applies irrespective of market size as well: Then, changes to the traded securities do not affect traders' inference and price impact, and equilibrium is *ex post*.

Example 4 in Appendix A.2 shows that allowing traders to trade additional securities whose payoffs lie within the span of the underlying assets is also non-neutral. In fact, no such security is redundant, except for one that duplicates a traded asset (Corollary 6(ii) in Appendix A.3).

4 Security Innovation

In the standard multi-asset equilibrium model with fully contingent demands, the introduction of securities whose payoffs lie in the span of the traded assets, such as either type of innovation in Examples 2 and 4, is *redundant*, i.e., does not change traders' equilibrium payoffs.[24] This result holds in imperfectly competitive ($I < \infty$) and competitive ($I \to \infty$) markets. However, when assets are cleared independently, the introduction of such synthetic products, without altering any traders' asset holdings or creating a non-zero supply of the securities, is generally not neutral. It changes traders' cross-asset price inference and their price impact. Notably, synthetic products remain payoff-relevant in markets with large traders even when there is no inference error (i.e., the uncertainty in aggregate asset holdings or prices vanishes, $\sigma_q^2 \to 0$),[25] as they still modify the equilibrium price impact. Table 1 provides a summary of the nonredundancy results.

Table 1: Nonredundancy of Security Innovation

	Joint Market Clearing	Independent Market Clearing
$I \to \infty$	X	✓ Information
$I < \infty$	X	✓ Information and Liquidity

Notes. When securities are cleared independently, security innovation that preserves the span of the traded assets is not neutral to traders' equilibrium payoffs, as it changes cross-security inference (information) and/or the equilibrium price impact (liquidity).

Taking as an objective the traders' *ex ante* total welfare, $\sum_i E[u^i(\mathbf{q}^i) - \mathbf{p} \cdot \mathbf{q}^i]$ (Eq. (13)), we present two main results: A market with a sufficiently large number of independently clearing derivatives can reproduce the equilibrium from the market with only the underlying assets, cleared jointly, where synthetic products would be redundant (Section 4.2). However, when traders have price impact, securities that do not reproduce the contingent-market outcome are more efficient, except in trivial cases (Section 5).

[24]New securities are redundant in the sense of equilibrium outcomes (Corollary 5 in Appendix A.1) and payoffs (Eq. (13)), not equilibrium demands $\{\mathbf{q}^{i,+}(\cdot)\}_i$.

[25]In contrast, with limited demand conditioning, new securities are nonredundant in the competitive market ($I \to \infty$) only if there is inference error (i.e., $\sigma_q^2 > 0$). While our main interest is in the incomplete information environment, the complete information scenario (i.e., $\sigma_q^2 \to 0$) is useful for isolating the effects of price impact; see Eq. (13) in Section 5.1. Defining the complete information benchmark as the limit equilibrium sidesteps the equilibrium multiplicity issue with $\sigma_q^2 = 0$ (e.g., Klemperer and Meyer (1989)).

4.1 Synthetic Products Are Nonredundant

At its core, security design is payoff-relevant because the variables it introduces in traders' expected trades affect inference, regardless of whether the underlying assets are traded directly. Our main results in this section show that security design is nonredundant when it modifies cross-asset inference or, equivalently (with imperfect competition), when it alters price impact across assets (Theorem 1), and derivatives are nonredundant except under special conditions on the underlying asset and derivative payoffs (Corollary 6 in Appendix A.3).

Assessing the nonredundancy of innovation necessitates comparing equilibrium payoffs across different market structures. In each, the fixed point among traders' schedules is equivalent to the fixed point among their price impact matrices (Theorem 3 in Appendix A.1). However, price impact matrices $\mathbf{\Lambda}$ cannot be compared directly across market structures, as their dimension and the securities for which they are defined may differ. Example 2 motivates the key analytic tool, which also provides insights into the analysis of nonredundancy.

Definition 4 (Per-Unit Price Impact) *In a market for N securities whose payoffs are defined for K underlying assets, let $\widehat{\mathbf{q}}^i \equiv (\widehat{q}^i_k)_k = \mathbf{W}\mathbf{q}^{i,+} \in \mathbb{R}^K$ be trader i's* total equilibrium trade *vector of all assets traded across N exchanges for all securities. The* per-unit price impact *$\widehat{\mathbf{\Lambda}} \in \mathbb{R}^{K\times K}$ is a positive semi-definite matrix, such that for all i and any* ex ante *initial holdings $\{E[\mathbf{q}^i_0]\}_i \in \mathbb{R}^{IK}$,*

$$E[\widehat{\mathbf{q}}^i] \equiv E[\mathbf{W}\mathbf{q}^{i,+}] = (\alpha\mathbf{\Sigma} + \widehat{\mathbf{\Lambda}})^{-1}\alpha\mathbf{\Sigma}(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}^i_0]); \tag{11}$$

i.e., if the price impact in a market structure with a single *exchange for K assets were $\widehat{\mathbf{\Lambda}}$, then the expected trade of each asset $k \in K$ in the counterfactual exchange would be equal to its expected total equilibrium trade in the market with N securities.*[26]

The proof of Theorem 1 (Lemma 2) shows that by examining the *single-exchange* counterfactual for the K underlying assets that determines $\widehat{\mathbf{\Lambda}}$, one can, first, compare equilibrium outcomes across markets with arbitrary securities and, second, identify nonredundant innovation with the change in $\widehat{\mathbf{\Lambda}}$ without the need to consider changes in expected trades.[27]

[26] $\widehat{\mathbf{\Lambda}}$ is not defined as an equilibrium variable in a single-exchange game. The proof of Lemma 2 in Appendix A.1 demonstrates the existence and uniqueness of $\widehat{\mathbf{\Lambda}}$ when $\mathbf{\Sigma}$ is non-singular. With $N = K + D$ securities, the inverse of the per-unit price impact matrix is characterized as (Eq. (28) in Lemma 2):

$$\widehat{\mathbf{\Lambda}}^{-1} = \mathbf{\Lambda}_a^{-1} + \underbrace{\underbrace{\boldsymbol{W}_d \mathbf{\Lambda}_d^{-1} \boldsymbol{W}_d'}_{\substack{\text{Basis change from } D \text{ derivatives} \\ \text{to } K \text{ underlying assets through } \boldsymbol{W}_d}}}_{\text{Projection of derivatives to the underlying asset}}. \tag{12}$$

Eq. (12) shows that the per-unit price impact $\widehat{\mathbf{\Lambda}}$ projects the liquidity risk from derivative trades' units to the liquidity risk per unit of *total* trades of the K underlying assets from all exchanges, thus transforming the basis of liquidity risk from derivative trades' units to the underlying assets.

[27] In competitive markets ($I \to \infty$), $\widehat{\mathbf{\Lambda}}$ becomes a zero matrix for any securities N and nonredundant innovation

Theorem 1 (Nonredundancy of Derivatives) Let $I < \infty$, $K > 1$, and suppose that $\mathbf{\Sigma}$ is non-singular. Consider two markets: One for $N \geq 1$ securities with payoffs $\mathbf{r}_N \equiv (r_n)_{n \in N} = \mathbf{W}'_N \mathbf{r}$ and another for $N' \geq 1$ securities with payoffs $\mathbf{r}_{N'} \equiv (r_{n'})_{n' \in N'} = \mathbf{W}'_{N'} \mathbf{r}$. The security innovation is redundant, i.e., $u^{i,N} = u^{i,N'}$ for all i, if and only if one of the following conditions holds for these two markets:

(i) *(Price impact)* The per-unit price impacts for the underlying assets coincide: $\widehat{\mathbf{\Lambda}}^N = \widehat{\mathbf{\Lambda}}^{N'}$.

(ii) *(Inference)* The expected total trades $E[\widehat{q}^i_k | p_\ell, \mathbf{q}^i_0]$, $\ell \neq k$, are the same in both markets for each k.

(iii) *(Market structure)* The span of the contingent variables in the total demand of each asset is the same: $span(\{p_n | w_{nk} \neq 0\}_{n \in N}) = span(\{p_{n'} | w_{n'k} \neq 0\}_{n' \in N'})$ for each k.

In which markets should one expect derivatives not to be neutral? Corollary 6 in Appendix A.3 establishes that no nonredundant derivatives can be created if and only if the payoffs of all underlying assets are either independent or perfectly correlated; see also Example 2(i). However, in all other markets, synthetic products alter the correlation of traded asset prices and, as a result, affect traders' price impacts. Notably, only derivatives that correlate *perfectly* with some traded assets or assets they replace are neutral. See Example 4 in Appendix A.2.

4.2 Bound on Security Innovation

Are there limits to introducing nonredundant derivatives? Theorem 1 implies a bound on the number of derivatives that can be introduced and remain nonredundant. Corollary 1 shows that the market structure with derivatives and/or the underlying assets that are cleared independently can reproduce the equilibrium in the fully contingent market for the underlying assets state by state (i.e., for all realizations of initial asset holdings); then, equilibrium becomes *ex post* and additional derivatives would be redundant.

Corollary 1 (Redundancy of Derivatives and Equivalence with Joint Market Clearing) Suppose that $I < \infty$ and consider a market with $N \geq 1$ securities defined for $K > 1$ underlying assets. Assume that $\mathbf{\Sigma}$ is non-singular. All derivatives traded with the N securities are redundant if and only if either (i) the N securities orthogonalize the payoffs of K assets or (ii) in the total trade of each asset k, the number of linearly independent contingent variables in expected trades of other assets $\ell \neq k$ is at least K: for each k, $span(\{r_n | w_{nk} \neq 0\}_{n \in N}) = span(\{r_\ell\}_{\ell \in K})$.

A couple of implications of this result are worth highlighting. First, Corollary 1 characterizes the nonredundant derivatives that can be introduced in the market. With $K = 2$ assets,

can be identified by the change in the Jacobian matrix of total demands $\widehat{\mathbf{C}} \equiv \mathbf{WCW}' \in \mathbb{R}^{K \times K}$; see also ft. 22. The conditions (ii) and (iii) in Theorem 1 apply to both competitive and imperfectly competitive markets.

the introduction of just one nonreplicating derivative is sufficient for the market where the derivative and the underlying assets are cleared independently to function as the market where the underlying assets are cleared jointly—equilibrium is *as if* traders could condition their demands on the price vector. For markets with $K \geq 2$ assets, the result places a bound on the number of such derivatives at $\frac{K(K-1)}{2}$, which applies for any underlying assets' payoffs and traders' trading needs; the only requirement for such derivatives' weights is that their joint payoffs with the underlying assets are linearly independent. Intuitively, for the traders' inference about expected trades to be perfect, it must be that for every pair of assets, some derivative weighs their assets' payoffs. Once $\frac{K(K-1)}{2}$ derivatives are introduced, any additional derivatives are redundant.

Second, derivatives allow the implementation of the fully contingent outcome with lower-dimensional schedules ($\mathbb{R}^1 \to \mathbb{R}^1$). In a market where such a *complete set of securities* is cleared independently, the cross-exchange price inference mimics the cross-asset price impact from a single exchange that clears the underlying assets jointly;[28] then, the per-unit price impact satisfies $\widehat{\mathbf{\Lambda}} = \mathbf{\Lambda}^c$. Intuitively, in the counterfactual that defines the price impact of trader i, other traders react to the additional demand of trader i *as if* their inference was perfect.

With $K \geq 2$ assets, Corollary 1 shows that the complete set of securities comprises either the K factors in place of the underlying assets or $\frac{K(K+1)}{2}$ imperfectly correlated securities (whether or not the underlying assets are included). However, in Section 5, we show that a market with a limited (incomplete) set of derivatives cleared independently generally improves upon the welfare bound of the joint market clearing when traders have price impact. Example 4 in Appendix A.2 illustrates Corollary 1.

Rostek and Yoon (2021) showed that it is possible to achieve the fully contingent outcome with an alternative design that features exchanges for multiple assets and no synthetic products, where demands are contingent on the prices of the assets traded within each exchange but uncontingent across exchanges. However, this design requires contingent schedules for every pair of assets. On the other hand, Corollary 1 shows that by introducing synthetic products instead of relying on a technology that clears multiple underlying assets jointly, a simpler design where traders' demands are all *uncontingent* can mimic the fully contingent outcome.

In fact, the result of Corollary 1 extends to subsets of the underlying assets $L \subset K$: The number of nonredundant derivatives whose payoffs underlie L assets is bounded by $\frac{L(L-1)}{2}$—the set of L securities is *locally complete*. Furthermore, in markets with K underlying assets, if the correlation among some assets is sufficiently similar, it is possible to achieve a welfare outcome close to the fully contingent outcome using fewer derivatives than the bound $\frac{K(K-1)}{2}$. For example, when the correlation between any pair of assets is $\rho_{k\ell} = \rho < 0$ for all k and $\ell \neq k$, a single derivative with the payoff $r_d = \frac{1}{K}\sum_k r_k$ can implement the outcome of joint market

[28] This result holds whether or not the underlying assets are traded alongside derivatives. Example 4 in Appendix A.2 illustrates the former case by introducing a sufficient number of derivatives with $K = 2$ underlying assets, while Example 2 demonstrates the latter by replacing the underlying assets with factors.

clearing. The general result is provided by Lemma 9 in Appendix A.4. These results do not have counterparts with multi-asset exchanges.[29,30]

5 Welfare Impact of Derivatives

We have shown that in markets that clear securities independently, no set of K securities with the same span as the underlying assets can achieve the same equilibrium payoffs as those assets, except in trivial cases (Corollary 6 in Appendix A.3). Taking efficiency as the objective, which securities should traders then trade? Should they trade an index (a "market portfolio" with weights proportional to average trading needs) or multiple funds? Or, might trading factors be efficient, given that independent payoffs would eliminate the cross-security inference effects in price impact, thereby reproducing the fully contingent outcome (Corollary 1)?

The standard model, which assumes joint clearing of assets, suggests that traders can achieve equivalent outcomes by trading the same or a smaller number of funds instead of individual assets. Moreover, the same funds are efficient and recommended for all investor classes, regardless of their trading needs or the joint distribution of assets. Only the number of funds with linearly independent payoffs matters as long as they have the same span.

However, we show that in markets where assets are not cleared jointly, the recommendation of using the same funds for all traders (analogous to the Mutual Fund Theorem) is only applicable under specific conditions. Specifically, it holds in competitive markets (Theorem 2 and Proposition 1) or when traders' motives for trade are purely speculative (i.e., they have zero trading needs $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i] = \mathbf{0}$ for all i; see Eq. (13)). When traders have price impact and nontrivial trading needs, the mutual-fund-type results do not hold, even in a weaker sense.

[29]With multi-asset exchanges, the outcome that coincides with the fully contingent market outcome must involve at least K contingent variables in each trader's total demand for each asset. This requirement holds regardless of whether or not the market is symmetric.

[30]Furthermore, a comparative analysis by Rostek and Yoon (2024a) shows that the equilibrium payoffs from designs involving exchanges for multiple assets (as investigated in Rostek and Yoon (2021)) cannot reproduce or be reproduced by synthetic products studied in this paper on a state-by-state basis. For example, the introduction of $\frac{L(L-1)}{2}$ derivatives with uncontingent demands does not generally enable perfect inference among these assets, in contrast to the effect of creating new exchanges for $L \subset K$ assets. More generally, these two types of innovations have different effects on the per-unit price impact. Independently cleared derivatives offer the flexibility to adjust the weights bundling asset payoffs, thereby achieving the desired own- and cross-asset per-unit price impacts. On the other hand, the joint clearing of multiple assets in a single venue generally leads to asymmetries in cross-asset price impacts, which are absent when clearing is done independently (or jointly) for all assets. Essentially, only with independent or joint clearing of all assets does the demand for any asset allow for symmetric inference regarding the prices of other assets. Otherwise, even if every pair of assets in L is jointly traded in some exchange, introducing new exchanges can still be nonredundant unless $L = K$ (i.e., the fully contingent outcome).

5.1 Welfare Effects of Derivatives

A key observation so far is that synthetic products are generally nonneutral because they affect traders' price impact and inference. By utilizing the welfare decomposition in Eq. (13), we can further attribute the welfare impact of derivatives to three effects. Namely, we can express the expected equilibrium payoff of trader i as a function of $\widehat{\mathbf{\Lambda}}$:

$$E[u^i(\mathbf{q}^i) - \mathbf{p}\cdot\mathbf{q}^i] \;=\; \underbrace{E[\boldsymbol{\delta}\cdot\mathbf{q}_0^i - \frac{1}{2}\mathbf{q}_0^i\cdot\alpha\mathbf{\Sigma}\mathbf{q}_0^i]}_{\text{Payoff without trade}} + \underbrace{(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i])\cdot\mathbf{\Upsilon}(\widehat{\mathbf{\Lambda}})(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i])}_{\text{Equilibrium surplus from trade}} \tag{13}$$

$$+ \underbrace{\frac{1}{2}\frac{I-2}{I-1}\sigma_q^2 tr\big(\alpha\mathbf{\Sigma}\big)}_{\text{Payoff term due to } Var[\overline{\mathbf{q}}_0|\mathbf{q}_0^i]>0} - \underbrace{\frac{I-1}{I}\sigma_q^2 tr\Big((\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}))'\alpha\mathbf{\Sigma}(\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}})) + \frac{2}{I-1}\alpha\mathbf{\Sigma}(\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}))\Big)}_{\text{Inference error}},$$

where the equilibrium surplus matrix coefficient $\mathbf{\Upsilon}(\widehat{\mathbf{\Lambda}}) = \frac{1}{2}\alpha\mathbf{\Sigma} - \mathbf{\Theta}(\widehat{\mathbf{\Lambda}})$ captures how liquidity risk $\mathbf{\Theta}(\widehat{\mathbf{\Lambda}})$ (see ft. 31) modifies the fundamental risk $\frac{1}{2}\alpha\mathbf{\Sigma}$, $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]$ is the vector of *ex ante* trading needs (i.e., the *target portfolio*; cf. Eq. (26)), and $\widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}})$ and $\mathbf{B}^c$ are the matrix coefficients on $\mathbf{q}_0^i$ in the uncontingent and the fully contingent demands, respectively.[31] The second line pertains to the variance of realized equilibrium surplus, where the term $Var[\overline{\mathbf{q}}_0|\mathbf{q}_0^i]$ is due to the uncertain aggregate asset holdings and the second term is due to the inference error resulting from limited demand conditioning.

From Eq. (13), security design affects welfare via $\widehat{\mathbf{\Lambda}}$ through[32]

(i) *Own-asset price impact*: $\widehat{\lambda}_k$ captures the marginal cost component of risk sharing for a specific asset;

(ii) *Cross-asset price impact*: Derivatives introduce non-zero off-diagonal elements of $\widehat{\mathbf{\Lambda}}$ (see, e.g., Example 2 Cont'd below); $\widehat{\lambda}_{k\ell} \neq 0$ captures the marginal cost or benefit components of diversification across assets;

(iii) *Inference error.*

To seed the intuition for the welfare effects of derivatives, it is helpful to consider a competitive market (i.e., $I \to \infty$) where assets are cleared independently. There, security design involving at least K securities affects the inference error alone, which can be minimized with the complete set of securities as characterized by Corollary 1, ensuring zero inference error.[33]

[31] Lemma 2 in Appendix A.1 characterizes $\mathbf{\Theta}(\widehat{\mathbf{\Lambda}}) = \frac{1}{2}\widehat{\mathbf{\Lambda}}(\alpha\mathbf{\Sigma} + \widehat{\mathbf{\Lambda}})^{-1}\alpha\mathbf{\Sigma}(\alpha\mathbf{\Sigma} + \widehat{\mathbf{\Lambda}})^{-1}\widehat{\mathbf{\Lambda}}$, $\widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}) = \mathbf{W}\mathbf{B} = ((1-\sigma_0)\alpha\mathbf{\Sigma} + (1+(I-2)\sigma_0)\widehat{\mathbf{\Lambda}})^{-1}\alpha\mathbf{\Sigma}$, and $\mathbf{B}^c = \frac{I-2}{I-1}\mathbf{Id}$. In the fully contingent market, the inference error in Eq. (13) is zero since $\widehat{\mathbf{B}}(\mathbf{\Lambda}^c) = \mathbf{B}^c$, i.e., equilibrium is *ex post*.

[32] In light of Theorem 1, the selection of a more efficient design involves choosing among the price impacts induced by the designs. Indeed, a design modifies the inference error in Eq. (13) if and only if it changes the per-unit price impact ($I < \infty$), since $\widehat{\mathbf{B}}$ is a function of $\widehat{\mathbf{\Lambda}}$.

[33] Fewer than K securities are not efficient. This can be observed in Eq. (13), where with zero price impact

Thus, in competitive markets where assets are cleared independently, an analogous principle to the Mutual Fund Theorem holds, implying the suitability of using the same funds for all traders.

Corollary 2 (Competitive Markets: Mutual Funds) Suppose $I \to \infty$. The market structure with the complete set of securities is efficient for all distributions of asset payoffs and traders' asset holdings.

However, we will show that in imperfectly competitive markets where assets are cleared independently, one can generally improve upon the efficiency of the complete set of securities despite allowing inference error. This can be accomplished by designing derivatives that mitigate the welfare costs associated with price impact (effects (i) and (ii) above). These derivatives effectively transform traders' equilibrium liquidity risk $\mathbf{\Theta}(\widehat{\mathbf{\Lambda}})$ and offset the welfare loss caused by imperfect inference. Therefore, unlike in markets that clear assets jointly, the role of securities is not merely to mimic the diversification achieved by the assets, as in the traditional Mutual Fund Theorem. Instead, their purpose is (1) to improve upon that diversification and (2) to facilitate more effective risk sharing.

To focus on the effects of price impact on equilibrium surplus, we subsequently assume that the inference error is zero, i.e., $\sigma_q^2 \to 0$. Under this assumption, both variance terms in Eq. (13) become zero. The means of initial holdings can be arbitrary; we will show that the heterogeneity in asset holdings plays a crucial role in determining the efficient security design.

5.2 Price Impact Benefits of Derivatives

Example 2 illustrates the significance of price impact in security design and highlights the main result of this section: Which securities should traders trade? Notably, the own-asset and cross-asset price impact components encourage different types of security design: Trading the independent factors (Example 2 in Section 3.2) minimizes the own-asset price impact while allowing an index (Eq. (17) below) can induce the cross-asset price impact beneficial for the trader's equilibrium payoff.

Example 2 Cont'd (Welfare Improvement with Securities). In the setting of Example 2, the design with $N = 2$ factors with payoffs defined by the weight vector $\mathbf{W}$ that orthogonalizes the underlying assets' payoffs would be efficient in the competitive market ($I \to \infty$; Corollary 2). However, when traders have price impact, changing the security weight from $\mathbf{W}$ to $\mathbf{W}^+$ that correlates security payoffs while preserving their span can increase traders' payoffs. This welfare change results from the interplay of two countervailing effects.

$\widehat{\mathbf{\Lambda}} = \mathbf{0}$, both the equilibrium surplus and the first variance term depend solely on the asset span and not the market structure. Both terms are maximized and remain constant with any K or more securities.

(Cross-asset price impact: diversification) When the security payoffs are independent (i.e., $\sigma_{12}^{+} = 0$), an increase in demand for one security does not affect the prices of other securities:

$$p_1^{+} + \lambda_{11}^{+} q_1^{i,+} + 0 \; q_2^{i,+} \quad \text{and} \quad p_1^{c,+} + \lambda_{11}^{c,+} q_1^{i,c,+} + 0 \; q_2^{i,c,+}.$$

Consider a change in one security's weight vector from $\mathbf{w}_1$ to $\mathbf{w}_1^{+}$, resulting in a new security payoff of $r_1^{++} = w_{11}^{+} r_1^{+} + w_{12}^{+} r_2^{+}$, given the factors' payoffs $\{r_1^{+}, r_2^{+}\}$. Relative to factors, the new weights *induce non-zero cross-security per-unit price impact* $\widehat{\lambda}_{12}^{+} \neq 0$: in the unit of factors,

$$p_1^{+} + \widehat{\lambda}_{11}^{+} q_1^{i,+} + \widehat{\lambda}_{12}^{+} q_2^{i,+}, \tag{14}$$

which may (or may not) reduce the marginal payment (i.e., the RHS of Eq. (14)) and be beneficial to a trader's equilibrium payoff. The desired condition for the cross-security per-unit price impact is[34]

$$sign(\widehat{\lambda}_{12}^{+}) = -sign(\frac{E[\overline{q}_{0,1}^{+}] - E[q_{0,1}^{i,+}]}{E[\overline{q}_{0,2}^{+}] - E[q_{0,2}^{i,+}]}). \tag{15}$$

Recall that with joint market clearing, $sign(\widehat{\lambda}_{12}^{+}) = sign(\rho_{12}^{+})$ always holds (Example 1), so condition (15) cannot be satisfied for all trading needs. However, with independent market clearing, the sign of the off-diagonal per-unit price impact induced by security weight $\mathbf{w}_1^{+}$ does not necessarily match the sign of the covariance. In fact, the security weight vector $\mathbf{w}_1^{+}$ can always be chosen to reduce the trading cost of diversification across securities. Using Eq. (28),

$$\widehat{\mathbf{\Lambda}}^{+} = \left(\begin{bmatrix} 1 & w_{11}^{+} \\ 0 & w_{12}^{+} \end{bmatrix} \begin{bmatrix} \lambda_1^{+} & 0 \\ 0 & \lambda_2^{+} \end{bmatrix}^{-1} \begin{bmatrix} 1 & 0 \\ w_{11}^{+} & w_{12}^{+} \end{bmatrix} \right)^{-1} = \frac{1}{(w_{12}^{+})^2} \begin{bmatrix} (w_{12}^{+})^2 \lambda_1^{+} & -w_{11}^{+} w_{12}^{+} \lambda_1^{+} \\ -w_{11}^{+} w_{12}^{+} \lambda_1^{+} & (w_{11}^{+})^2 \lambda_1^{+} + \lambda_2^{+} \end{bmatrix}, \tag{16}$$

we have that $sign(\widehat{\lambda}_{12}^{+}) = -sign(w_{11}^{+}/w_{12}^{+})$. For example, choosing security weight $\mathbf{w}_1^{+} = (w_{11}^{+}, w_{12}^{+})$ to be proportional to a trader's expected trading needs $E[\overline{\mathbf{q}}_0^{+}] - E[\mathbf{q}_0^{i,+}]$, i.e.,

$$\mathbf{w}_1^{+} = \xi(E[\overline{\mathbf{q}}_0^{+}] - E[\mathbf{q}_0^{i,+}]) \tag{17}$$

for a proportionality coefficient $\xi \in \mathbb{R}$, ensures that Eq. (15) holds for any $E[\overline{\mathbf{q}}_0^{+}] - E[\mathbf{q}_0^{i,+}]$ irrespective of $\mathbf{\Sigma}^{+}$. Thus, *derivatives can transform price impact avoidance into price impact seeking* when it comes to cross-security per-unit price impact. Lemmas 3 and 5 in Appendix A.1 provide additional insights into the effects of derivatives on the per-unit price impact, effects that cannot be achieved with contingent demands.

(Own-asset price impact: risk sharing) Nevertheless, the securities in Eq. (17), which aim to

[34] That is, the cross-security per-unit price impact is beneficial when $\widehat{\lambda}_{12}^{+} < 0$ and the trader takes the same positions (buying or selling) in the securities, or when $\widehat{\lambda}_{12}^{+} > 0$ and the trader takes opposite positions across securities.

induce the desired cross-security price impact by aligning with traders' target portfolios, are typically not welfare-maximizing, because they fail to account for the change in the securities' own price impacts $\{\widehat{\lambda}_k^+\}_k$ brought by the inference effects they induce. Specifically, when a security with weight $\mathbf{w}_1^+$ is traded instead of one of the factors, the non-zero covariance increases the securities' price impact $\widehat{\lambda}_{11}^+$ (Lemma 4 in Appendix A.1). □

The effects of security innovation on the own- and cross-asset price impact $\widehat{\lambda}_{11}^+$ and $\widehat{\lambda}_{12}^+$ in Eq. (14) give rise to a trade-off that shapes the efficient design characterized by Theorem 2.

5.3 Which Derivatives Are Efficient?

For markets with $K \geq 2$ underlying assets, Theorem 2 characterizes the $N = K$ efficient securities in the span of the underlying assets. In particular, neither factors (or the complete set; Corollary 1) nor the index (Eq. (17)) are generally efficient when traders have price impact.

Theorem 2 (Efficient Security Design) Let $I < \infty, K > 1$, and assume no inference error, i.e., $\sigma_q^2 \to 0$. Suppose that the payoffs of the K underlying assets are symmetric (i.e., $\sigma_{kk} = \sigma^2$ and $\sigma_{k\ell} = \sigma^2\rho$ for all k and $\ell \neq k$) and traders' *ex ante* trading needs are symmetric across assets and traders (i.e., $E[\overline{q}_{0,k}] - E[q_{0,k}^i] = E[\overline{q}_{0,\ell}] - E[q_{0,\ell}^i]$ for all k, ℓ, and i). For all market structures with K symmetrically correlated securities (i.e., $\sigma_{kk}^+ = \sigma^2$ and $\sigma_{k\ell}^+ = \sigma^2\rho^+$ for all k and $\ell \neq k$), the following results hold:

(a) *(Design with underlying assets is inefficient)* For any I, K, and ρ, there exists a unique security correlation $\rho^+(\rho; I, K)$ that maximizes *ex ante* welfare. Furthermore, there exists a unique cutoff $\overline{\rho}$ such that $\rho^+(\overline{\rho}; I, K) = \overline{\rho}$, and correlation $\rho^+(\rho; I, K)$ is higher than the underlying asset correlation ρ if and only if the latter is below the cutoff $\overline{\rho}$:

$$\rho^+(\rho; I, K) > \rho \quad \text{if and only if} \quad \rho < \overline{\rho}.$$

The following weight matrix characterizes the securities with the optimal correlation ρ^+:

$$\mathbf{W} = \frac{\sqrt{1-\rho^+}}{\sqrt{1-\rho}}\mathbf{Id} - \frac{\sqrt{1-\rho^+}}{K\sqrt{1-\rho}}\mathbf{1}\mathbf{1}' + \frac{\sqrt{1+(K-1)\rho^+}}{K\sqrt{1+(K-1)\rho}}\mathbf{1}\mathbf{1}'. \tag{18}$$

(b) *(Correlating securities is efficient)* The *ex ante* welfare is maximized with the securities whose payoffs orthogonalize those of the underlying assets (i.e., $\rho^+(\rho; I, K) = 0$) if and only if the market is competitive (i.e., $I \to \infty$).

(c) *(Index is not efficient)* The *ex ante* welfare is maximized with the securities whose weights mimic the target portfolio (i.e., $\mathbf{w} = \mathbf{1}$, equivalently $\rho^+(\rho; I, K) = 1$) if and only if $\rho = 1$ or the market is competitive (i.e., $I \to \infty$).

With price impact, design matters. Theorem 2 underscores several implications of price impact. First, synthetic products can be designed to strictly increase *ex ante* welfare compared to both the underlying assets and the market with the complete set of securities. In Section 5.5, we show that allowing more than K securities can further increase welfare.

Second, in the presence of price impact, the efficient security design is no longer invariant to trader or asset characteristics other than their span. Asset payoff correlation ρ and traders' target portfolios $\{E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]\}_i$, which are determined by their initial holdings, become relevant. As a result, the efficient set of derivatives differs between one-sided markets (where some traders buy and others sell all securities, such as the primary market for Treasury securities) and two-sided markets (where traders buy and sell different securities, such as inter-dealer markets).[35] See Table 2 below.

The results in the presence of price impact contrast sharply with those in competitive markets, where K factors or any other complete set of securities are weakly efficient for any $\mathbf{\Sigma}$ and $\{E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]\}_i$, as we discussed in Section 5.1. Neither the Mutual Fund Theorem (Corollary 2) nor the Separation Theorem (Corollary 3) hold when traders have price impact.

Next, we highlight several implications of Theorem 2 for security design.

How to design efficient derivatives? Example 2 anticipates the design recommendation of Theorem 2: The optimal level of security correlations $\rho^+(\rho; I, K)$ balances two countervailing effects on the price impact cost components associated with diversification and risk sharing, given the symmetry of trading needs across traders and assets. Each effect favors different securities.

- *Cross-asset price impact favors an index:* Holding fixed the own-asset price impacts $\{\widehat{\lambda}_k\}_k$, the welfare benefit due to cross-asset price impact (condition (15)) is maximal when $\rho^+ = 1$. The proportional to traders' target portfolios $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]$ weight $\mathbf{w}$ for all i (Eq. (17)), as implied by $\rho^+ = 1$, indicates that there is a single efficient security: the index (Eq. (18)).[36]

- *Own-asset price impact favors factors:* Holding fixed the cross-asset price impacts $\{\widehat{\lambda}_{k\ell}\}_{k,\ell\neq k}$, $\rho^+ = 0$ minimizes the inference errors and the own-asset price impacts λ_k^+ for all k (Lemma 4 in Appendix A.1). Thus, factors would be efficient (Example 2).

It follows that an index is too costly in terms of the own-asset price impacts, while trading factors ignores the benefit of the cross-asset price impacts for diversification (Fig. 1(A); the

[35] In one-sided markets, the optimal asset covariance for trading the underlying assets is positive (i.e., $\overline{\rho} > 0$). However, in two-sided markets, it is *negative* (i.e., $\overline{\rho} < 0$), given the joint symmetry of $\mathbf{\Sigma}$ and $\{E[\overline{q}_{0,k}] - E[q_{0,k}^i]\}_{i,k}$. The proof of Theorem 2 encompasses both one- and two-sided markets, but for simplicity of exposition, we present the result for one-sided markets here.

[36] Similarly, in the symmetric two-sided market, securities correlation $\rho^+ = -\frac{1}{K-1}$ (i.e., the bound imposed by the positive semidefiniteness of $\mathbf{\Sigma}^+$) minimizes the cost of diversification (see Fig. 5 and the proof of Theorem 2 in Appendix A.4).

red lines in Fig. 1(B)). Therefore, the efficient design of securities (i) correlates security payoffs (Theorem 2(b)) (ii) differently from the underlying assets (except when $\rho = \overline{\rho}$, Theorem 2(a)), but (iii) not perfectly (except when $\rho = 1$, Theorem 2(c)).

Figure 1: Optimal Security Correlations in Theorem 2

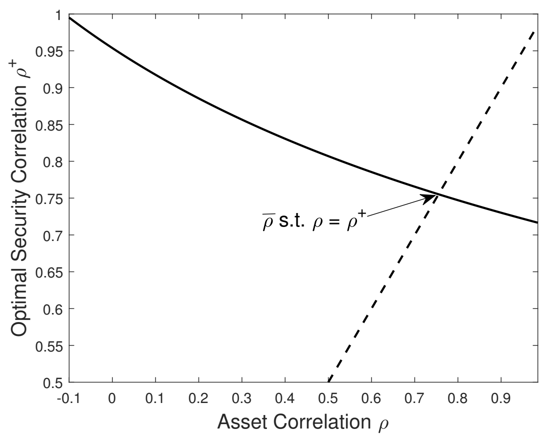

(A)

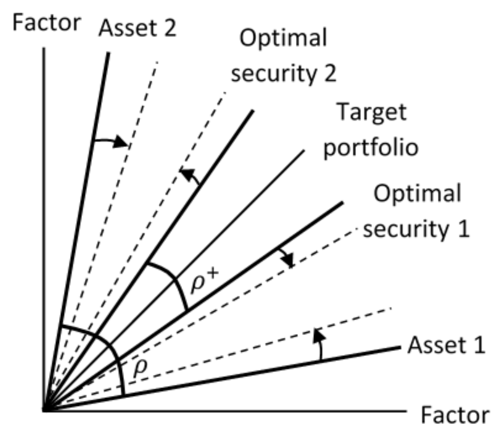

(B)

Notes. Panel A: The figure plots the optimal security correlation ρ^+ as a function of the underlying asset correlation ρ. The underlying assets are the optimal securities (i.e., $\rho = \rho^+$) only if $\rho = \overline{\rho}$ (Theorem 2(a)). The dashed line is the 45° degree line, i.e., $\rho = \rho^+$.

Panel B: Cast in terms of trading towards a trader's target portfolio (Gârleanu and Pedersen (2013), Kyle, Obzihaeva, and Wang (2017)), the following principle summarizes the efficient securities. The efficient securities align with the traders' target portfolio (which favors an index, $\rho^+ = 1$; the 45° line) subject to trading costs $\{\lambda_k\}_k$ (which favor factors, $\rho^+ = 0$; the axes). If the underlying asset correlation ρ is small (the solid lines), the efficient securities act as closer substitutes than the underlying assets; if the underlying asset correlation ρ is large (the dashed lines), the efficient securities are weaker substitutes than the underlying assets. The index, which matches the traders' target, is not welfare-maximizing except in the competitive market with no inference error (Corollary 2).

Moreover, the trade-off between two components of price impact determines whether the efficient securities are more or less strongly correlated relative to the underlying assets: Because λ_k^+ is convex in ρ^+ (Lemma 4 in Appendix A.1), the securities' effect on risk sharing dominates when the underlying asset payoffs are strongly correlated. Fig. 1 summarizes this result.

- When the underlying asset payoffs are strongly correlated (the dashed lines in Fig. 1(B)), efficient securities lower the correlation to decrease the cost of risk sharing while sacrificing some level of diversification: $\rho^+ < \rho$.
- When the payoffs of the underlying assets are close to independent (the solid lines in Fig. 1(B)), efficient securities increase the correlation to improve risk diversification while limiting the benefit of risk sharing: $\rho^+ > \rho$.

Table 2 summarizes the efficient securities' weights, characterized in Eq. (18).

Theorem 2 focuses on markets with symmetric asset covariance and trading needs. When either is heterogeneous, the own- and cross-asset price impacts are not necessarily monotone in

Table 2: Efficient Design Correlates Securities

	$\rho < \overline{\rho}_- < 0$	$\overline{\rho}_- < \rho < \overline{\rho}_+$	$0 < \overline{\rho}_+ < \rho$
One-sided	positive weights	positive weights	positive and negative weights
Two-sided	positive weights	positive and negative weights	positive and negative weights

Notes. The cutoffs for the underlying asset correlations $\overline{\rho}_+$ and $\overline{\rho}_-$ in one- and two-sided markets, respectively, are characterized in the proof of Theorem 2. In one-sided markets, increasing security correlations relative to the underlying asset correlations (i.e., positive weights; Eq. (18)) is welfare-improving when the underlying asset correlations are not too large so the inference effects in λ_k^+ are not too strong ($\rho < \overline{\rho}_+$; Theorem 2(a)). In two-sided markets, increasing security correlations (i.e., positive weights; Eq. (18)) is welfare-improving when the underlying asset correlations are sufficiently large so the inference effects in λ_k^+ are strong ($\rho < \overline{\rho}_-$).

securities correlation ρ^+. Nevertheless, the key implications continue to apply: Trading factors or the underlying assets is generally suboptimal and the optimal security weight is determined by the trade-off between own- and cross-asset price impacts (see Example 3).

5.4 Index Trading?

Theorem 2 and Example 2 illuminate the role of price impact for the decentralized market's analog of the Mutual-Fund-Theorem. This result continues to apply when traders are price-takers—then, K or fewer (linearly independent) funds, which are the same for all traders, will be efficient as long as a sufficient number of funds are available; only the number of funds matters (Corollary 2)—with one difference. Namely, a stronger (viz. fully-contingent-demand theory) recommendation applies with independent market clearing: Instead of being merely payoff-equivalent, funds are *strictly* more efficient than the underlying assets.

However, this result no longer applies when traders have price impact. To appreciate the extent to which the mutual-fund-style results fail once one accounts for the fact that not all assets are cleared jointly, consider markets where traders are *ex ante* symmetric, i.e., $E[\mathbf{q}_0^i] = \xi_{ij} E[\mathbf{q}_0^j]$ for some $\xi_{ij} \in \mathbb{R}$ and for all i and $j \neq i$. If traders are price-takers and assuming zero inference error, a single fund—an index that matches the trading needs $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]$ for all i (Eq. (17))—maximizes *ex ante* welfare. All traders hold a combination of the index (or "market portfolio") $\mathbf{w}_m$ and the risk-free asset (numéraire).[37] However, in an imperfectly competitive market, an index is not efficient, *even when traders are* ex ante *symmetric and there is zero information loss.* In fact, a stronger result emerges: Trading fewer than K securities is suboptimal.

[37] When the vector of asset trading needs is asymmetric across traders, the two-fund separation theorem fails to hold, irrespective of whether the market is contingent or uncontingent, competitive or imperfectly competitive. In such cases, it is not possible for any single security to replicate the equilibrium outcome attained with the underlying assets. In Eq. (17), the derivative is tailored to the traders' individual *ex ante* trading needs $E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}]$. It follows from Theorem 2, in imperfectly competitive decentralized markets, the securities implied by the Mutual Fund Theorem (Corollary 2) do not correspond to the efficient design even with symmetric trading needs.

Corollary 3 (Two-Fund Separation Does Not Hold) When a market with K securities is imperfectly competitive (i.e., $I < \infty$), fewer than K securities give a lower welfare compared to the maximum welfare with K securities for any asset covariance $\mathbf{\Sigma}$ and initial holdings $F(\mathbf{q}_0^i)_i$.

By Theorem 2, the optimal security correlation with price impact satisfies $|\rho^+| < 1$, except in trivial cases (see also Lemma 7 in Appendix A.3). Therefore, efficiency requires traders to trade more securities than just an index. Furthermore efficiency dictates that the security design must be tailored to the traders' trading needs $(F(\mathbf{q}_0^i)_i)$ and the distribution of assets $(\mathbf{\Sigma})$ (Theorem 2 and Proposition 1 in the next section), even if the *ex ante* trading needs are symmetric across assets.

These findings are in line with the rationale for the SEC's 2019 revision of ETF regulations. The updated rule eased the previous mandate that ETFs' basket compositions must closely align with indices or diversified portfolios, due to concerns that "ETFs without basket flexibility [...] could result in wider bid-ask spreads and potentially less efficient arbitrage." The agency emphasized that the flexibility to use custom ETF baskets can be in the best interests of investors and shareholders. See SEC (2019, Section II.C.5). Koont, Ma, Pastor, and Zeng (2023) find that corporate bond ETFs adjust their baskets to lower the transaction costs, which constrains their index-tracking capacity. The study indicates that ETFs are active—create issuer-specific portfolios rather than passively following market indices—driven by considerations for both index tracking and liquidity transformation. It's this active management of baskets that enables ETFs to effectively balance these objectives.

Our analysis implies that even a weaker version of the mutual-fund result, which states that "sufficiently many funds are efficient irrespective of trader characteristics," does not hold when traders have price impact. Regarding the number of securities, the complete set of securities is generally inefficient (Theorem 2 and Proposition 1). In fact, as we will show next, efficient design involves strictly more than K securities, yet still a limited number (Proposition 1 and Example 3).

5.5 Additional Derivatives?

So far, we established that a design with K securities with the same span as the underlying assets generally increases *ex ante* welfare compared to the underlying assets (Theorem 2). This section shows that whether more than K securities can further enhance welfare depends on the heterogeneity in security covariances $\mathbf{\Sigma}^+$ and trading needs *across securities* $\{E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}]\}_i$. It turns out that when both are symmetric, as in Theorem 2, the introduction of any nonredundant securities in a market with K correlated ones *decreases* welfare. However, when either is heterogeneous, a market structure with *more* than K securities, suitably designed, can always improve *ex ante* welfare relative to designs with any K securities. The additional flexibility provided by the new securities helps to induce asymmetries in trading costs across assets, which can be beneficial precisely when there are heterogeneities in either the covariances or trading needs across K securities. However, too many securities are generally inefficient (Proposition 1).

Example 3 (Derivatives and Welfare) Let $I < \infty$ and $\sigma_q^2 \to 0$. Consider a market with $K = 3$ securities, two of which (securities 2 and 3) have symmetric covariances and *ex ante* trading needs:[38]

$$\mathbf{\Sigma}^+ = \begin{bmatrix} 1 & \rho_k^+ & \rho_k^+ \\ \rho_k^+ & 1 & \rho_\ell^+ \\ \rho_k^+ & \rho_\ell^+ & 1 \end{bmatrix} \quad \text{and} \quad E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}] = \begin{bmatrix} E[\overline{q}_{0,k}^+] - E[q_{0,k}^+] \\ E[\overline{q}_{0,\ell}^+] - E[q_{0,\ell}^+] \\ E[\overline{q}_{0,\ell}^+] - E[q_{0,\ell}^+] \end{bmatrix}.$$

Suppose an additional derivative with payoff $r_d = w_k r_1^+ + w_\ell r_2^+ + w_\ell r_3^+$ is introduced. We compare the *ex ante* welfare in the market with just the K securities and the market with the additional security, considering an arbitrary weight vector (w_k, w_ℓ, w_ℓ). Fig. 2 illustrates the welfare change with the welfare-maximizing derivative as a function of relative security correlations ρ_k^+/ρ_ℓ^+ and relative trading needs $(E[\overline{q}_{0,k}^+] - E[q_{0,k}^{i,+}])/(E[\overline{q}_{0,\ell}^+] - E[q_{0,\ell}^{i,+}])$.

(a) *(Symmetric markets)* When the security covariances in $\mathbf{\Sigma}^+$ and trading needs $E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}]$ are both symmetric across securities (i.e., $\rho_k^+/\rho_\ell^+ = 1$ and $(E[\overline{q}_{0,k}^+] - E[q_{0,k}^{i,+}])/(E[\overline{q}_{0,\ell}^+] - E[q_{0,\ell}^{i,+}]) = 1$ in Fig. 2(A)), no additional derivative d increases the *ex ante* welfare.

(b) *(Heterogeneity matters)* When either security covariances $\mathbf{\Sigma}^+$ or trading needs $E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}]$ are heterogeneous across securities, there exists a derivative d that increases *ex ante* welfare relative to markets with the K securities.

(c) *(Welfare-maximizing derivatives)* Derivatives help lower the per-unit price impacts of securities strongly correlated with other securities or associated with larger trading needs by placing a larger weight on such securities. It follows from the comparative statics of per-unit price impact (Lemmas 3 and 5 in Appendix A.1) that such derivatives reduce the per-unit price impact $\widehat{\lambda}_k$ as well as the relative per-unit price impact $\widehat{\lambda}_k/\widehat{\lambda}_\ell$ for those securities. See Fig. 2(B). □

Example 3 shows that, compared to the market with K securities, additional derivatives can generally be designed to increase *ex ante* welfare even when these traded securities cannot be delisted, for any security covariances $\mathbf{\Sigma}^+$ and trading needs $E[\overline{\mathbf{q}}_0^+] - E[\mathbf{q}_0^{i,+}]$. Clearly, when

[38] The example extends to markets with either K underlying assets or K welfare-maximizing securities. Given the symmetry between assets 2 and 3, the set of K welfare-maximizing securities includes two symmetric securities.

Figure 2: Security Design and Asset Heterogeneity

(*A*) (*B*)

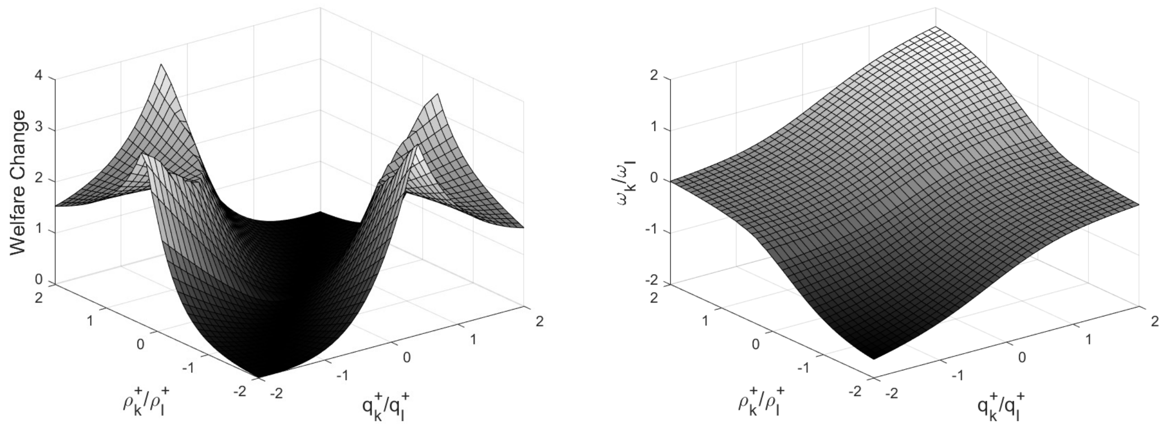

Notes. Panel A: For any K securities, *ex ante* welfare weakly increases with additional derivatives optimized to maximize welfare, relative to the uncontingent and fully contingent markets (Example 3(a) and (b)). The figure illustrates this increase relative to the maximum welfare achieved by the two benchmarks.

Panel B: The optimal weights that characterize the additional derivative are larger for securities that exhibit stronger correlation or larger trading needs (e.g., $|w_k| > |w_\ell|$ when $|\rho_k^+| > |\rho_\ell^+|$ or $|q_k^+| > |q_\ell^+|$). Table 2 gives the weight signs that induce the joint substitution; see also Example 3(c).

The payoff variance of the derivative is normalized to match the security payoff variances; $\rho_\ell^+ = -0.1, E[\overline{q}_{0,\ell}^+] - E[q_{0,\ell}^{i,+}] = 10$, and $I = 10$.

combined with delisting K securities (Theorem 2 and Proposition 4 in Appendix A.3), the design of the number and payoffs (weights) of securities can further improve the *ex ante* welfare.

Too many derivatives are inefficient. Proposition 1 shows that, unlike in the competitive market (Corollary 2), markets with large traders generally achieve higher welfare with fewer derivatives than required to complete the market. In fact, *any* set of securities that does not mimic the equilibrium of the fully contingent market can result in higher or lower welfare for some distributions of asset holdings (see also Example 3). The proof of Proposition 1 encompasses markets where the underlying assets can be traded as well as those where they cannot.

Proposition 1 (Welfare with Securities and Underlying Assets) Let $I < \infty$ and $K > 1$. Assume that $\mathbf{\Sigma}$ is non-singular and there is no inference error, i.e., $\sigma_q^2 \to 0$. Suppose an *arbitrary* set of D derivatives is introduced and the equilibrium with $K + D$ securities is not *ex post*. The *ex ante* welfare is strictly larger than with the complete set of securities for some distribution of asset holdings $F((\mathbf{q}_0^i)_i)$.

The key observation underlying Proposition 1 is as follows (Lemma 8 in Appendix A.3): The per-unit price impact $\mathbf{\Lambda}^c$ for the fully contingent design (which matches the covariance, up to a scaling factor) is not unambiguously ranked with $\widehat{\mathbf{\Lambda}}$ in the positive-semidefinite sense in any uncontingent design with $K+D$ securities. Consequently, the liquidity risk matrices $\mathbf{\Theta}(\mathbf{\Lambda}^c)$ and

$\boldsymbol{\Theta}(\widehat{\boldsymbol{\Lambda}})$ in Eq. (13) are also not unambiguously ranked. Hence, unlike the competitive-market case, where the complete set of securities is efficient irrespective of the traders' trading needs and the asset distribution, the efficient design responds to market characteristics. Proposition 1 shows that heterogeneity in either traders' trading needs across the underlying assets or asset covariances further amends the efficient weights of Theorem 2. This captures that the welfare-maximizing weights balance the securities' effects on the trading costs of diversification and risk sharing while taking into account the assets' relative importance to traders' trading needs and price impact.

Table 3 provides a summary of the efficient number of securities.

Table 3: EFFICIENT NUMBER OF SECURITIES

(A) *Ex ante* heterogeneous traders

	Contingent	Uncontingent
$I \to \infty$	K	A complete set
$I < \infty$	K	Strictly between K and $\frac{K(K+1)}{2}^*$

(B) *Ex ante* symmetric traders

	Contingent	Uncontingent
$I \to \infty$	1	1
$I < \infty$	1	Strictly between K and $\frac{K(K+1)}{2}^*$

Notes. This table provides information on the efficient number of securities. Panel B assumes no inference error (i.e., $\sigma_q^2 \to 0$). The symbol "*" denotes markets where the complete set of securities, such as K factors, is efficient. In imperfectly competitive uncontingent markets, the efficient number of securities is generally strictly between K and $\frac{K(K+1)}{2}$.
If there is a non-zero inference error, Table 3(A) applies regardless of whether traders are *ex ante* symmetric or heterogeneous. The first variance term in Eq. (13) increases in the securities' span, so it is minimized with only 1 security and maximized with K linearly independent securities.

6 Innovating Securities Outside the Span

Thus far, our analysis has focused on introducing securities in the span of the K underlying assets. However, independent market clearing also impacts the welfare effects of securities outside the K assets' span. Suppose that each trader i holds $\mathbf{q}_0^i = (q_{0,z}^i)_z \in \mathbb{R}^Z$ units of Z assets, but only a subset $K \subset Z$ of the assets is traded in the market. Theorem 3 and Lemma 2 in Appendix A.1, as well as Theorem 1, can be extended to apply to such markets.[39]

It is worth noting that in a competitive contingent market, the introduction of *any* securities for unspanned risk—i.e., whose payoffs weigh the $Z \setminus K$ underlying assets—always leads to a weak Pareto improvement with quasilinear utilities. This result extends to imperfectly competitive markets as well (Corollary 8 in Appendix A.3). However, when the market clears securities independently, Proposition 2 shows that the introduction of any securities for unspanned risk can lower welfare even with quasilinear utilities; see also Fig. 3.

[39] Duffie and Jackson (1989, 1990) investigate the impact of introducing securities outside the traded asset span in a dynamic model with exogenous transaction costs.

Proposition 2 (Welfare with Securities Outside the Span of Traded Assets) Let $I < \infty$ and $K < Z$, and assume no inference error, i.e., $\sigma_q^2 \to 0$. Suppose we introduce $D \geq 1$ securities outside the span of the K traded assets; for each $d \in D$, $r_d = \sum_{k \in K} w_{dlk} r_k + \sum_{z \in Z \setminus K} w_{dz} r_z$ with $w_{dz} \neq 0$ for some $z \in Z \setminus K$. The *ex ante* welfare with K assets and any such D securities is strictly lower than in the market with only K assets for some distribution of initial holdings $F((\mathbf{q}_0^i)_i)$.

The introduction of securities outside the span of the K traded assets enables risk sharing for the $Z \setminus K$ securities (i.e., $\{E[\overline{q}_{0,z}] - E[q_{0,z}^i]\}_{z \in Z \setminus K}$; Eq. (56)) in any market structure, whether it is competitive or imperfectly competitive, contingent or uncontingent. However, the price inference among the K assets and $Z \setminus K$ securities changes the price impacts of the K traded assets, unless the security payoffs are independent of those of the K assets. Consequently, despite the welfare benefits of risk sharing, the $Z \setminus K$ securities can decrease welfare. This occurs when these securities either significantly increase the price impact on the traded assets K (e.g., when $|\rho|$ and $|\rho^+|$ are large in Fig. 3), or fail to align with the target portfolio (i.e., Eq. (17) applied to security weights $\mathbf{w} \in \mathbb{R}^Z$ and trading needs for all risks $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i] \in \mathbb{R}^Z$; see Fig. 3).

Figure 3: Welfare Change with Securities Outside the Underlying Asset Span

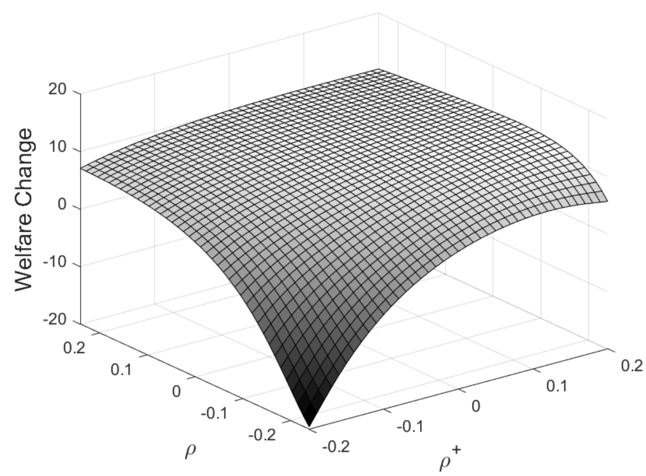

Notes. $K = D = 2$. For any assets with $\sigma_{k\ell}^+ = \sigma^2 \rho$, $k \in Z$, and $\ell \neq k$, one can design $D = Z - K$ securities that strictly increase the *ex ante* welfare compared to the market with only the K assets (Corollary 7 in Appendix A.3): The payoffs of these D securities are correlated with the payoffs of the K traded assets ($\sigma_{kd}^+ = \sigma^2 \rho^+$, $k \in K$, and $d \in D$) and are independent among the D securities ($\sigma_{dd'}^+ = 0$, $d, d' \in D$, and $d' \neq d$). However, D securities that are strongly correlated with the K assets (i.e., a large $|\rho^+|$) can lower the *ex ante* welfare (Proposition 2). In the figure, traders' *ex ante* trading needs, $E[\overline{q}_{0,z}] - E[q_{0,z}^i] > 0$ for all $z \in Z$, are symmetric among assets K and among assets $Z - K$, and $I = 10$.

Suppose that all the additional $Z - K$ securities can be designed. In light of Proposition 1, one might wonder whether the efficient number of securities could be lower than Z. It is not. Instead, the efficient set of securities allows for trading *all* Z fundamental risks (Corollary 7,

Appendix A.3) but generally *limits* the cross-asset information upon which traders' demands can condition (Theorem 2, Proposition 1). The efficient number of securities can be characterized using an argument analogous to that of Corollary 3 and Table 3, but with Z replacing K.

In conclusion, efficient design treats market incompleteness resulting from unspanned risk differently from limited conditioning. (The dimension of fundamental risk in the aggregate initial holdings is Z, but the minimum number of securities required to hedge all contingencies and eliminate inference error is $\frac{Z(Z+1)}{2}$.) While the fundamental risk can be factorized into independent components up to Z factors, it cannot be factorized with more than Z securities.

7 Conclusion and Discussion

By accounting for the independence in market clearing, we may gain insights into why there are active markets for many types of financial products that would be neutral in the standard theory, and we can determine which markets should be active, given asset and trader characteristics and the level of imperfect competition. The *equilibrium* model of nonredundant derivatives implied by the independence in market clearing in an otherwise standard framework offers rich opportunities for empirical research.

Our main focus has been on the welfare-improving effects of derivatives, though it is clear that derivatives can also lower welfare. Arguments akin to those in the proof of Theorem 2 can be applied to show that derivatives that enhance welfare in certain markets can decrease welfare in others. Specifically, derivatives that induce correlations opposite to those characterized by these results lead to lower welfare (Corollary 9 in Appendix A.4). In essence, derivatives detract from welfare if they create inadequate correlations among underlying assets, whether by introducing excessive or insufficient correlation, thus failing to mitigate the trading costs associated with price impact.

We conclude with a discussion of the additional implications of independent market clearing.

7.1 Pricing Derivatives

The classical techniques based on spanning have laid the foundation for the analysis of equilibrium and asset pricing. These techniques involve the representation of uncertainty through the implied state space over which contingent securities are defined. However, these techniques and the implied representation of risk are not generally applicable when markets clear securities independently.

With fully contingent trading, any security whose payoff lies in the span of the traded assets can be priced via a linear combination of the traded assets' prices. It follows from our analysis (Corollary 5 in Appendix A.1 and Corollary 6 in Appendix A.3) that linear pricing does not generally hold for new securities whose payoffs are synthetically created in markets that clear the underlying assets independently.

Corollary 4 (Linear Pricing Does Not Hold) Consider a market with K traded assets that are cleared independently. Assume that $\mathbf{\Sigma}$ is non-singular. For any non-traded security with payoff $r_d = \mathbf{w}_d'\mathbf{r}$, linear pricing $p_d^+ = \mathbf{w}_d'\mathbf{p}$ holds for any $\mathbf{w}_d \in \mathbb{R}^K$ if and only if equilibrium with K assets is *ex post*.

This result holds in competitive and imperfectly competitive markets. It implies that factor prices are not useful in the same way in markets that clear securities independently, except when the security payoffs are all independent, or there are at least $\frac{K(K+1)}{2}$ securities with linearly independent payoffs traded (Corollary 1). The latter result can be seen as the counterpart of the spanning number (Duffie (1986)) for markets that clear securities independently.[40] In more general market structures like the efficient ones, the implied representation of risk is not solely determined by the fundamental covariance but also depends on the endogenous price impact.

The approach to characterizing Bayesian Nash equilibrium in demand schedules as a fixed point in price impacts could be useful for the counterfactual pricing and valuation of financial innovation, including new assets, exchanges, or contracts. For instance, recent work by Gabaix and Koijen (2022) and Koijen and Yogo (2019) introduces methods for estimating demand elasticities across securities. Combined with these methods, the techniques based on the per-unit price impact matrix would allow counterfactual pricing and analysis of security and market design—the first-order conditions in our model can be estimated for general utility functions (see Wittwer (2021, Appendix F)). These techniques can be extended to enable comparative design analysis of different instruments, utilizing the per-unit price impacts derived from the matrix representations of the corresponding fixed points.

7.2 Other Security Innovations

We do not examine options and exotic derivatives whose payoffs are non-linear functions of the underlying asset payoffs. While we should acknowledge that additional effects would arise in a nonlinear equilibrium, amending the welfare analysis accordingly, the main findings of this paper concerning independent market clearing remain applicable: the nonredundancy of financial innovation, the liquidity-information welfare trade-off, and the alternative implementation of fully contingent outcomes. These results do not rely on linearity, only the imperfect inference among independently cleared securities. Regardless of linearity, introducing derivatives affects traders' inference across different securities, thereby influencing price impacts; see also Chabakauri, Yuan, and Zachariadis (2021) and Keller and Tseng (2023). In settings with

[40] In a continuous-time model with fully contingent demands and K diffusion state variables, Duffie (1986) introduced the concept of a spanning number, which represents the minimum number of security markets necessary for market completeness: With K securities (in addition to the riskless asset), any additional security in the span of K state variables can be priced linearly using the prices of these K securities. In the uncontingent market with K state variables $\overline{\mathbf{q}}_0 \in \mathbb{R}^K$, Corollary 1 implies that the spanning number is $\frac{K(K+1)}{2}$. Unlike with the K securities that are cleared jointly, however, the $\frac{K(K+1)}{2}$ securities' payoffs cannot be independent.

non-linear demands, our model's first-order conditions can be written analogously to those in Wittwer (2021, Appendix F) and Glebkin, Malamud, and Teguia (2023).

We have assumed (correlated) private values to isolate the core effects that underlie the nonredundancy of financial innovation. These effects relate to price inference across assets and do not rely on inference among traders related to common values. In decentralized markets with common values, there are additional reasons to innovate derivatives.

7.3 Over-the-Counter Derivatives

The focus on exchange-traded derivatives has allowed us to isolate independence in market clearing as a source of nonredundancy. For financial products traded over the counter, transactions tend to occur through protocols that resemble bargaining rather than a price mechanism (see, e.g., Atkeson, Eisfeldt, and Weill (2015)). Nevertheless, the key effects we explore are not reliant on a particular price mechanism such as the uniform-price auction, and apply to other games in contracts. Rather, the essence of our results is rooted in the presence of inefficiency resulting from two-sided (buyer and seller) private information among strategic traders in any (budget-balanced) mechanism, and how changes in the available financial products affect that inefficiency.

On the other hand, an over-the-counter market structure modifies the information on which traders can condition their demands with different counterparties and introduces counterparty heterogeneity. The design of over-the-counter derivatives motivated by the joint effects merits its own study.

7.4 Dynamic Trading

Dynamic trading strategies can reproduce the outcomes of static joint clearing in markets where trading rounds are frequent relative to shocks (to information or liquidity) that renew the gains from trade (Lyu, Rostek, and Yoon (2022a)). More generally, the ability to submit demands for different assets contingent on contemporaneously determined prices has distinct implications from conditioning demands on past prices in dynamic markets. Conditioning on past prices allows demands—of any type—to at least partially incorporate the information from past shocks. On the other hand, conditioning demands on contemporaneous-round prices affects how current-round shocks impact behavior. If shocks renewing the gains from trade occur as frequently as trading rounds, or frequently enough, the effects of independent market clearing that we examine will be present even in the continuous trading limit.

In such dynamic markets, demand conditioning on past prices—potentially of all securities—induces intertemporal inference effects in price impact, which interact with the per-period inference vs. price impact tradeoff. Lyu, Rostek, and Yoon (2022b) show that the innovation of dynamic securities (whose payoffs and prices are determined in different rounds, e.g., futures and repo contracts) can further improve efficiency. A study of how static and dynamic security design interacts in markets with large traders would be worthwhile.

A Appendix

Appendix A.1. Equilibrium Characterization

Appendix A.2. Comparison with Contingent Demands

Appendix A.3. Results for General Markets

Appendix A.4. Results for Symmetric Markets

A.1 Equilibrium Characterization

We consider markets for an arbitrary set of $N \leq K + D$ securities, with K underlying assets and D derivatives. Matrix $\mathbf{W} \equiv (\mathbf{w}_1, \cdots, \mathbf{w}_N) \equiv (w_{nk})_{k,n} \in \mathbb{R}^{K\times N}$ represents the security weights, and the distribution of security payoffs is jointly Normal, $\mathcal{N}(\boldsymbol{\delta}^+, \boldsymbol{\Sigma}^+)$, where $\boldsymbol{\delta}^+ = \mathbf{W}'\boldsymbol{\delta}$ and $\boldsymbol{\Sigma}^+ = \mathbf{W}'\boldsymbol{\Sigma}\mathbf{W}$. We allow asset holdings to be imperfectly correlated across the underlying assets, $\boldsymbol{\Omega} \equiv (Cov[q^i_{0,k}, q^i_{0,\ell}])_{k,\ell} \in \mathbb{R}^{K\times K}$ for all i. In particular, to ensure that the per-capita aggregate asset holdings (equivalently, price) are random in the limit large market ($I \to \infty$), we allow for the common value component $\mathbf{q}^{cv}_0 = (q^{cv}_{0,k})_k \in \mathbb{R}^K$ in traders' asset holdings. For each asset k, privately known asset holdings $\{q^i_{0,k}\}_i$ are correlated among traders through $q^{cv}_{0,k} \sim \mathcal{N}(E[q^{cv}_{0,k}], \sigma^2_{cv})$: for each i,

$$q^i_{0,k} = q^{cv}_{0,k} + q^{i,pv}_{0,k} \qquad \text{and} \qquad q^{i,pv}_{0,k} \sim \mathcal{N}(E[q^{i,pv}_{0,k}], \sigma^2_{pv}), \tag{19}$$

where $q^{i,pv}_{0,k}$ are independent across i and k.[41] Trader i knows his asset holdings $\mathbf{q}^i_0$ but not its components $\mathbf{q}^{cv}_0$ or $\mathbf{q}^{i,pv}_0 = (q^{i,pv}_{0,k})_k \in \mathbb{R}^K$. The asset holdings $\{q^i_{0,k}\}_i$ and the common value $q^{cv}_{0,k}$ are independent across assets k.

Equilibrium characterization (Theorem 3, Corollary 5, Lemma 2) adapts the steps from Rostek and Yoon (2021) to extend their Theorem 1 to a model with D derivatives, whose asset holdings are zero, traded along with K underlying assets; additionally, $\mathbf{W}$ is a weight matrix rather than an indicator matrix.

Theorem 3 shows that a fixed point in demand schedules is equivalent to one in price impacts. We parameterize the best response demand (6) of trader i for security n as follows:

$$q^i_n(p_n) \equiv a^i_n - \mathbf{b}^i_n \mathbf{q}^i_0 - c^i_n p_n \qquad \forall p_n \in \mathbb{R} \quad \forall \mathbf{q}^i_0 \in \mathbb{R}^K, \tag{20}$$

with the demand intercept $a^i_n \in \mathbb{R}$, the demand coefficients $\mathbf{b}^i_n \in \mathbb{R}^{1\times K}$, and the demand slope $c^i_n \in \mathbb{R}_+$. To write the fixed point problem in matrix form, we denote the matrix demand coefficient on $\mathbf{q}^i_0$ by $\mathbf{B}^i \equiv (\mathbf{b}^i_n)_n \in \mathbb{R}^{N\times K}$, and denote the slope of a trader's demand for N securities by $\mathbf{C}^i \equiv diag(c^i_n)_n \in \mathbb{R}^{N\times N}$. The matrix operator $[\cdot]_d : \mathbb{R}^{N\times N} \to \mathbb{R}^{N\times N}$ is defined such that for any matrix $\mathbf{A} \in \mathbb{R}^{N\times N}$, $[\mathbf{A}]_d \equiv diag(a_{nn})_n$ is a matrix whose diagonal elements coincide with those of $\mathbf{A}$ and off-diagonal elements equal zero.

[41] The presence of a common value component in $\{\mathbf{q}^i_0\}_i$ does not qualitatively affect any results and only impacts the magnitude of inference coefficients. For simplicity, we have assumed that asset holdings have a symmetric variance across traders and are independent across assets. The results hold qualitatively without these assumptions. The equilibrium characterization allows for correlation in asset holdings, which changes the magnitude of price impacts but not the nonredundancy result.

Theorem 3 (Equilibrium: Fixed Point in Demand Schedules) The equilibrium (net) demand schedules, defined by matrix coefficients $\{\boldsymbol{a}^i\}_i, \mathbf{B}$, and $\mathbf{C}$, and price impact $\mathbf{\Lambda}$ are characterized by the following conditions: for each trader i,

(i) (*Demand coefficients, given price impact*) Given price impact matrix $\mathbf{\Lambda}$, the coefficients of (net) demands $\boldsymbol{a}^i, \mathbf{B}$, and $\mathbf{C}$ are characterized by

$$\boldsymbol{a}^i = \mathbf{C}\underbrace{\left(\boldsymbol{\delta}^+ - (\alpha\mathbf{W}'\mathbf{\Sigma} - \mathbf{C}^{-1}\mathbf{B})E[\overline{\mathbf{q}}_0]\right)}_{=\mathbf{p}-\mathbf{C}^{-1}\mathbf{B}\overline{\mathbf{q}}_0} + \underbrace{\left((\alpha\mathbf{\Sigma}^+ + \mathbf{\Lambda})^{-1}\alpha\mathbf{W}'\mathbf{\Sigma} - \mathbf{B}\right)(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i])}_{\text{Adjustment due to cross-asset inference}}, \tag{21}$$

$$\mathbf{B} = \Big((1-\sigma_0)(\alpha\mathbf{\Sigma}^+ + \mathbf{\Lambda}) + \underbrace{\sigma_0\mathbf{C}^{-1}}_{\substack{\text{Adjustment due to}\\ \text{cross-asset inference}}}\Big)^{-1}\alpha\mathbf{W}'\mathbf{\Sigma}, \tag{22}$$

$$\mathbf{C} = \Big[(\alpha\mathbf{\Sigma}^+ + \mathbf{\Lambda})\underbrace{(\mathbf{B}Var[\overline{\mathbf{q}}_0]\mathbf{B}')[\mathbf{B}Var[\overline{\mathbf{q}}_0]\mathbf{B}']_d^{-1}}_{\substack{\text{Inference coefficient}\\ Var[\mathbf{s}^{-i}|\mathbf{q}_0^i][Var[\mathbf{s}^{-i}|\mathbf{q}_0^i]]_d^{-1}}}\Big]_d^{-1}, \tag{23}$$

where $\overline{\mathbf{q}}_0 \equiv \frac{1}{I}\sum_j \mathbf{q}_0^j \in \mathbb{R}^K$ is the aggregate asset holdings and $\sigma_0 \equiv \frac{\sigma_{cv}^2 + \frac{1}{I}\sigma_{pv}^2}{\sigma_{cv}^2 + \sigma_{pv}^2} \in \mathbb{R}_+$.

(ii) (*Correct price impact*) Price impact $\mathbf{\Lambda}$ equals the transpose of the Jacobian of the trader's inverse residual supply function:

$$\mathbf{\Lambda} = \frac{1}{I-1}(\mathbf{C}^{-1})'. \tag{24}$$

Corollary 5 (Equilibrium Prices and Allocations) Given the equilibrium demand coefficients $\{\boldsymbol{a}^i\}_i, \mathbf{B}, \mathbf{C}$, and price impact $\mathbf{\Lambda}$ in Theorem 3, the equilibrium prices and trades are

$$\mathbf{p} = \boldsymbol{\delta}^+ - (\alpha\mathbf{W}'\mathbf{\Sigma} - \mathbf{C}^{-1}\mathbf{B})E[\overline{\mathbf{q}}_0] - \mathbf{C}^{-1}\mathbf{B}\overline{\mathbf{q}}_0, \tag{25}$$

$$\mathbf{q}^i = \left((\alpha\mathbf{\Sigma}^+ + \mathbf{\Lambda})^{-1}\alpha\mathbf{W}'\mathbf{\Sigma} - \mathbf{B}\right)(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]) + \mathbf{B}(\overline{\mathbf{q}}_0 - \mathbf{q}_0^i). \tag{26}$$

A trader's equilibrium trade depends on their *ex ante* trading needs $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}_0^i]$ and *ex post* trading needs $\overline{\mathbf{q}}_0 - \mathbf{q}_0^i$.

Lemma 1 states that the equilibrium price impact $\mathbf{\Lambda}$ converges to $\mathbf{0}$ in the competitive market (i.e., $I \to \infty$). Its proof mimics that of Lemma 3 from Rostek and Yoon (2021).

Lemma 1 (Price Impact in Competitive Markets) In a market for $N \geq 1$ securities, let α^I be the traders' risk aversion, and suppose $\mathbf{\Lambda}^I$ is the equilibrium price impacts for all $I < \infty$ and in the limit market as $I \to \infty$. The equilibrium price impact becomes zero as $I \to \infty$ if $\alpha^I = \alpha\gamma^I \in \mathbb{R}_+$ increases slower than linearly with $\gamma^I \sim o(I^{1-\varepsilon})$ for some $\varepsilon > 0$: $\mathbf{\Lambda} = \lim_{I\to\infty}\mathbf{\Lambda}^I = \mathbf{0}$.

Lemma 2 shows that the per-unit price impact $\widehat{\mathbf{\Lambda}}$ (Eq. (11)) is a sufficient statistic for a trader's *ex ante* equilibrium payoff. The *cross-asset inference* $\widehat{\mathbf{B}}$ is the coefficient on the privately known asset holdings $\mathbf{q}_0^i$ in a trader's total demand that matches the variance of the total equilibrium trade (cf. Eq. (22)): for all i,

$$Var[\widehat{\mathbf{q}}^i] \equiv Var\left[\mathbf{W}\mathbf{q}^{i,+}\right] = \widehat{\mathbf{B}}Var[\overline{\mathbf{q}}_0 - \mathbf{q}_0^i]\widehat{\mathbf{B}}' = \frac{I-1}{I}\sigma_{pv}^2\widehat{\mathbf{B}}\widehat{\mathbf{B}}'; \tag{27}$$

$(\widehat{\mathbf{B}}\widehat{\mathbf{B}}')_{k\ell}(\widehat{\mathbf{B}}\widehat{\mathbf{B}}')^{-1}_{kk}$ is the cross-asset inference coefficient in the expected total trade $E[\widehat{q}^i_\ell|\widehat{q}^i_k, \mathbf{q}^i_0]$ (see Eq. (26)). Appendix A.3 provides the proof of Lemma 2.

Lemma 2 (Sufficient Statistic for Equilibrium Payoffs) Consider a market with $I < \infty$ and N securities defined by $\mathbf{W} \equiv (w_{nk})_{k,n} \in \mathbb{R}^{K\times N}$. Assume that $\mathbf{\Sigma}$ is non-singular.[42]

(1) *(Per-unit price impact)* The per-unit price impact $\widehat{\mathbf{\Lambda}} \in \mathbb{R}^{K\times K}$, defined by condition (11), is characterized as

$$\widehat{\mathbf{\Lambda}} = \left(\mathbf{W}\mathbf{\Lambda}^{-1}\mathbf{W}'\right)^{-1}, \tag{28}$$

and cross-asset inference $\widehat{\mathbf{B}} \in \mathbb{R}^{K\times K}$ is characterized as a function of $\widehat{\mathbf{\Lambda}}$:

$$\widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}) \equiv \left((1-\sigma_0)(\alpha\mathbf{\Sigma} + \widehat{\mathbf{\Lambda}}) + \sigma_0(I-1)\widehat{\mathbf{\Lambda}}\right)^{-1}\alpha\mathbf{\Sigma}. \tag{29}$$

(2) *(Expected payoff)* The expected equilibrium payoff of trader i as a function of $\widehat{\mathbf{\Lambda}}$ is:

$$E[u^i(\mathbf{q}^i) - \mathbf{p}\cdot\mathbf{q}^i] = \underbrace{E[\boldsymbol{\delta}\cdot\mathbf{q}^i_0 - \frac{1}{2}\mathbf{q}^i_0\cdot\alpha\mathbf{\Sigma}\mathbf{q}^i_0]}_{\text{Payoff without trade}} + \underbrace{(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}^i_0])\cdot\mathbf{\Upsilon}(\widehat{\mathbf{\Lambda}})(E[\overline{\mathbf{q}}_0] - E[\mathbf{q}^i_0])}_{\text{Equilibrium surplus from trade}} \tag{30}$$
$$+ \underbrace{\frac{1}{2}\frac{I-2}{I-1}\sigma^2_{pv}tr\left(\alpha\mathbf{\Sigma}\right)}_{\text{Payoff term due to } Var[\overline{\mathbf{q}}_0|\mathbf{q}^i_0]>0} - \underbrace{\frac{I-1}{I}\sigma^2_{pv}tr\left((\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}))'\alpha\mathbf{\Sigma}(\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}})) + \frac{2}{I-1}\alpha\mathbf{\Sigma}(\mathbf{B}^c - \widehat{\mathbf{B}}(\widehat{\mathbf{\Lambda}}))\right)}_{\text{Inference error}},$$

where $\mathbf{B}^c = \frac{I-2}{I-1}\mathbf{Id}$ is the coefficient of the contingent demand on $\mathbf{q}^i_0$ and

$$\mathbf{\Upsilon}(\widehat{\mathbf{\Lambda}}) \equiv \frac{1}{2}\mathbf{\Sigma} - \mathbf{\Theta}(\widehat{\mathbf{\Lambda}}) \equiv \frac{1}{2}\alpha\mathbf{\Sigma}(\alpha\mathbf{\Sigma}+\widehat{\mathbf{\Lambda}})^{-1}(\alpha\mathbf{\Sigma}+2\widehat{\mathbf{\Lambda}})(\alpha\mathbf{\Sigma}+\widehat{\mathbf{\Lambda}})^{-1}\alpha\mathbf{\Sigma} \in \mathbb{R}^{K\times K}$$

represents the marginal payoff per unit of *ex ante* trading needs $E[\overline{\mathbf{q}}_0] - E[\mathbf{q}^i_0]$.

In the contingent market, the inference error in Eq. (30) is zero (i.e., equilibrium is *ex post*).

New securities and equilibrium price impact. From Theorem 3 and Lemma 2 (Eqs. (23) and (28)), new securities affect price impact in two ways: They

(i) provide more exchanges to trade units of the underlying assets when $N > K$ and

(ii) change inference across securities.

See also Eq. (31) for (i) and Eq. (32) for (ii). Because of the joint effect of (i) and (ii), the introduction of a derivative may increase or lower the per-unit price impact. We present three results:

First, the introduction of derivatives that cover (weigh) all assets indeed lowers the per-unit price impact of the underlying assets when asset correlations $\mathbf{\Sigma}$ and derivative weights $\mathbf{W}_d$ are not too heterogeneous (Lemma 3(i)). When covariances $\mathbf{\Sigma}^+$ are heterogeneous—either because $\mathbf{\Sigma}$ or $\mathbf{W}_d$ are—the same derivative can increase or lower the per-unit price impact $\widehat{\lambda}_k$ on an asset included in the derivative (Fig. 4).

Second, the introduction of a derivative alters the liquidity of the assets not included in the new security (Lemma 5 and Fig. 4).

[42] The proof allows a singular covariance matrix $\mathbf{\Sigma}$. Then, the uniqueness of $\widehat{\mathbf{\Lambda}}$ holds up to payoff equivalence.

Third, even if derivatives cover all assets, fewer derivatives that each cover more assets lower the price impact more than a large number of derivatives that each cover fewer assets (Lemma 3(ii)). This occurs because the cross-security inference effect is weaker with the former design.

These results have no analogs with contingent demands.

Lemma 3 (Equilibrium Price Impact: Symmetric Markets) Let $I < \infty$, $K = DL > 1$, $L \geq 1$, $D \geq 1$, and suppose that asset covariances are symmetric ($\sigma_{kk} = \sigma^2$ for all k and $\sigma_{k\ell} = \sigma^2\rho < 0$ for all k and $\ell \neq k$). Apart from the K assets, there are $D = \frac{K}{L}$ derivatives whose payoffs are each defined by the unweighted average of L asset payoffs; the assets underlying derivatives are disjoint, $r_d = \frac{1}{L}\sum_{m=1}^{L} r_{L(d-1)+m}$ for each $d \in D$. The equilibrium per-unit price impact $\widehat{\mathbf{\Lambda}}$ satisfies the following properties:

(i) The per-unit price impacts are lower than the price impacts in the uncontingent market (and higher than those in the contingent market) for only K assets: $\lambda^c_{kk} \leq \widehat{\lambda}_k \leq \lambda^u_k$. The first equality holds when $L = K$ or $\rho = 0$, and the second equality holds when $L = 1$ or $\rho = 0$.

(ii) The per-unit price impacts decrease as L increases: $\frac{\partial \widehat{\lambda}_k}{\partial L} < 0$ and $\frac{\partial \widehat{\lambda}_{k\ell}}{\partial L} < 0$.

Lemma 4 (Comparative Statics of Equilibrium Price Impact: Symmetric Markets) In the setting of Lemma 3, the equilibrium per-unit price impact $\widehat{\mathbf{\Lambda}}$ satisfies the following properties:

(i) The per-unit price impacts increase as $|\rho|$ increases: $\frac{\partial(\widehat{\lambda}_k - \lambda^c_{kk})}{\partial|\rho|} > 0$ and $\frac{\partial(\widehat{\lambda}_{k\ell} - \lambda^c_{k\ell})}{\partial|\rho|} > 0$.

(ii) The per-unit price impacts decrease as I increases: $\frac{\partial(\widehat{\lambda}_k - \lambda^c_{kk})}{\partial I} < 0$ and $\frac{\partial(\widehat{\lambda}_{k\ell} - \lambda^c_{k\ell})}{\partial I} < 0$.

Lemma 5 (Equilibrium Price Impact: Derivatives for an Asset Subset) Let $I < \infty$ and $K > 1$, and suppose that the covariances are symmetric for all assets (i.e., $\sigma_{kk} = \sigma^2$ for all k; $\sigma_{k\ell} = \sigma^2\rho$ for all k and $\ell \neq k$). We introduce a derivative whose payoff is an unweighted average of a strict asset subset $K_1 = K/2$, $r_d = \frac{1}{K_1}\sum_{m \in K_1} r_m$. The equilibrium per-unit price impact $\widehat{\mathbf{\Lambda}} = diag((\widehat{\lambda}_{k1} - \widehat{\lambda}_{k\ell})\mathbf{Id} + \widehat{\lambda}_{k\ell}\mathbf{11'}, \widehat{\lambda}_{k2}\mathbf{Id})$ satisfies

$$\lambda^c_k \leq \widehat{\lambda}_{k1} \leq \lambda^u_k \quad \forall k \in K_1 \qquad \text{and} \qquad \lambda^c_k \leq \lambda^u_k \leq \widehat{\lambda}_{k2} \quad \forall k \in K \setminus K_1,$$

where $\mathbf{\Lambda} = \lambda^u_k \mathbf{Id}$ and $\mathbf{\Lambda}^c = \frac{\alpha}{I-2}\mathbf{\Sigma}$ are the equilibrium price impacts in the uncontingent and contingent markets, respectively, before the derivative is introduced.

Lemma 3 and Lemma 5 give sufficient conditions on the primitives for a derivative to lower the per-unit price impact of the assets that underly the derivative: When the covariances $\mathbf{\Sigma}^+$ are sufficiently symmetric across securities, the introduction of derivatives that weigh assets symmetrically lowers the per-unit price impact $\widehat{\lambda}_k$ for all assets k underlying the derivative, with effect (i) dominating effect (ii), as described above. The proofs of the lemmas are presented in Appendix A.4. Here we elaborate on the two effects underlying these lemmas.

(Part (i): Per-unit price impact of the underlying assets) For (i), trading the same asset in more than one exchange through derivatives can reduce the per-unit price impact. Suppose D

Figure 4: Derivative-Induced Changes in the Per-Unit Price Impacts $\widehat{\lambda}_k$ and $\widehat{\lambda}_n$

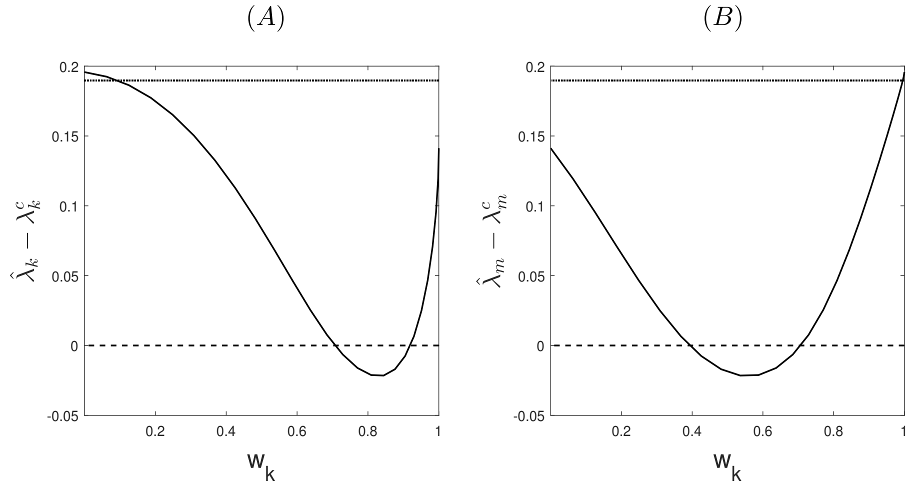

Notes. With heterogeneous derivative weights, the per-unit price impacts with the derivative can be higher than in the uncontingent market (the dotted lines) or lower than in the contingent market (the dashed lines). The plots are based on a market with $K = 4$ assets in which a derivative is introduced with heterogeneous weights on two assets $r_d = w_k \frac{1}{K_1} \sum_{k \in K_1} r_k + w_m \frac{1}{K_2} \sum_{m \in K_2} r_m$. $I = 10, \alpha = 10, \sigma_0 = 0.1, \rho = -0.25$, and w_m and w_k are such that $Var[r_d] = 1$.

Lemma 5 shows that when a derivative weighs a subset of assets $K_1 \subset K$ (i.e., $w_k = 1$ in Panels A and B), the per-unit price impacts satisfy $\lambda_k^c \leq \widehat{\lambda}_k \leq \lambda_k^u$ for assets $k \in K_1$, and $\lambda_m^c \leq \lambda_m^u \leq \widehat{\lambda}_m$ holds for $m \in K_2$. Moreover, when asset covariances and derivative weights are sufficiently symmetric, the per-unit diagonal price impact $\widehat{\lambda}_k$ is lower with derivatives for the assets whose payoffs are assigned higher weights: $\widehat{\lambda}_k < \lambda_k^c = \lambda_m^c$(the dashed lines) $< \widehat{\lambda}_m$ if $w_k > 0.7$.

derivatives $\mathbf{r}_d = \mathbf{W}_d' \mathbf{r}$ for some $\mathbf{W}_d \in \mathbb{R}^{K \times D}$ are traded along with the existing securities. Eq. (28) can be written as follows:

$$\widehat{\mathbf{\Lambda}} \quad = \quad \underbrace{\mathbf{\Lambda}_a}_{\substack{\text{Price impact of } K \text{ underlying assets} \\ \text{changes with new securities}}} \quad - \quad \underbrace{(\mathbf{\Lambda}_a \mathbf{W}_d (\mathbf{\Lambda}_d + \mathbf{W}_d' \mathbf{\Lambda}_a \mathbf{W}_d)^{-1} \mathbf{W}_d' \mathbf{\Lambda}_a)}_{\substack{\text{Per-unit price impact decreases} \\ \text{due to projection}}}, \tag{31}$$

where the equilibrium price impact $\mathbf{\Lambda}^+ = diag(\mathbf{\Lambda}_a, \mathbf{\Lambda}_d)$ is a block diagonal matrix, with blocks corresponding to the K underlying asset and D derivatives. The reduction of per-unit price impact relative to the price impact in exchanges for existing securities $(\mathbf{\Lambda}_a - \widehat{\mathbf{\Lambda}})$ occurs only for the assets underlying the derivatives.

(Part (ii): Introduction of securities affects inference across securities) For (ii), the introduction of a new security—with or without delisting an existing security—can increase or lower the price impact by altering inference effects $\frac{\partial q_n^i}{\partial p_\ell} \frac{\partial E[p_\ell | p_n, \mathbf{q}_0^i]}{\partial p_n}$ among the existing securities and give rise to inference effects $\frac{\partial q_n^i}{\partial p_d} \frac{\partial E[p_d | p_n, \mathbf{q}_0^i]}{\partial p_n}$ between the price of each traded security n and the expected price of each derivative d:[43]

$$\lambda_n^i = -\Big(\underbrace{\sum_{j\neq i} \frac{\partial q_n^j}{\partial p_n}}_{\text{Direct effect}} + \underbrace{\sum_{\ell\neq n} \frac{\partial q_n^j}{\partial p_\ell}\frac{\partial E[p_\ell|p_n,\mathbf{q}_0^i]}{\partial p_n}}_{\substack{\text{Inference effect among}\\ \text{existing securities}}} + \underbrace{\sum_{d} \frac{\partial q_n^i}{\partial p_d}\frac{\partial E[p_d|p_n,\mathbf{q}_0^i]}{\partial p_n}}_{\substack{\text{Inference effect between}\\ \text{security } n \text{ and derivatives } d}}\Big)^{-1}. \tag{32}$$

The inference effects depend on the substitutabilities and complementarities in all security payoffs, which are arbitrary (subject to $\mathbf{\Sigma}$ being positive semidefinite). Derivative weights affect security covariances, and hence, the cross-security inference effects in price impact. Derivatives do not change the direct effects in price impact.

When security correlations and derivative weights on all underlying assets are sufficiently symmetric (i.e., $\sigma_{n\ell} = \sigma^2\rho$ for all n and $\ell \neq n$ and $w_{dn} = w_{d\ell}$ for all underlying n and $\ell \neq n$, for all derivatives d), the introduction of new derivatives—or replacing a security by another that is more strongly correlated with other securities—increases price impact λ_m (Lemma 3). Larger security covariances increase inference effects in Eq. (32).

When either security correlations or derivative weights are *heterogeneous*, however, a security innovation can lower the price impact λ_n for some securities. This result is due to the nonseparability of the price impact for any security pair in all security payoffs. For example, prices of complementary securities ($\sigma_{\ell n} < 0$) can be positively correlated ($Cov[p_\ell, p_n] > 0$).

[43] Differentiating the first-order condition (6) of trader $j \neq i$ for asset n with respect to p_n gives the price elasticity $\frac{\partial q_n^j(\cdot)}{\partial p_n}$: for each n and $j \neq i$,

$$-\alpha\sigma_{nn}\frac{\partial q_n^j(\cdot)}{\partial p_n} - \sum_{\ell\neq n}\sigma_{n\ell}(-c_\ell^j)\frac{\partial E[p_\ell|p_n,\mathbf{q}_0^j]}{\partial p_n} - \sum_{d}\sigma_{nd}(-c_d^j)\frac{\partial E[p_d|p_n,\mathbf{q}_0^j]}{\partial p_n} = 1 + \lambda_n^j\frac{\partial q_n^j(\cdot)}{\partial p_n}.$$

The partials $\frac{\partial q_n^j}{\partial p_n} \equiv \frac{1}{\alpha\sigma_{nn}+\lambda_n^j}$ and $\frac{\partial q_n^j}{\partial p_\ell} \equiv \frac{\sigma_{n\ell}c_\ell^j}{\alpha\sigma_{nn}+\lambda_n^j}$ characterize the direct and inference effects in λ_n^i (Eq. (32)).

Data Availability

Code replicating the figures in this article can be found in Rostek and Yoon (2024b) in the Harvard Dataverse, `https://doi.org/10.7910/DVN/IBGXQU`.

References

Allen, F. and E. Carletti. 2006. "Credit Risk Transfer and Contagion." *Journal of Monetary Economics*, 53, 1: 89-111.

Allen, F. and D. Gale. 1988. "Optimal Security Design." *Review of Financial Studies*, 1, 3: 229-263.

Allen, F. and D. Gale. 1994. "Limited Market Participation and Volatility of Asset Prices." *American Economic Review*, 84, 4: 933-955.

Allen, J. and M. Wittwer. 2023. "Centralizing Over-the-Counter Markets?" *Journal of Political Economy*, forthcoming.

Angel, J., L. Harris, and C. Spatt. 2011. "Equity Trading in the 21st Century." *Quarterly Journal of Finance*, 1, 1: 1-53.

Antill, S. and D. Duffie. 2021. "Augmenting Markets with Mechanisms." *Review of Economic Studies*, 88, 4: 1665-1719.

Atkeson, A., A. Eisfeldt, and P-O. Weill. 2015. "Entry and Exit in OTC Derivatives Markets." *Econometrica*, 83, 6: 2231-2292.

Babus, A. and K. Hachem. 2021. "Regulation and Security Design in Concentrated Markets." *Journal of Monetary Economics*, 121, 1: 139-151.

Babus, A. and K. Hachem. 2023. "Markets for Financial Innovation." *Journal of Economic Theory*, 208: 105-615.

Babus, A. and P. Kondor. 2018. "Trading and Information Diffusion in Over-the-Counter Markets." *Econometrica*, 86: 1727-1769.

Babus, A. and C. Parlatore. 2022. "Strategic Fragmented Markets." *Journal of Financial Economics*, 145: 876-908.

Bahaj, S., R. Czech, S. Ding, and R. Reis. 2023. "The Market for Inflation Risk." Working Paper.

Baisa, B. and J. Burkett . 2018. "Large Multi-Unit Auctions with a Large Bidder." *Journal of Economic Theory*, 174: 1-15.

Baldauf, M. and J. Mollner. 2021. "Trading in Fragmented Markets." *Journal of Financial and Quantitative Analysis*, 56, 1: 93-121.

Bali, T. G., H. Beckmeyer, M. Moerke, and F. Weigert. 2023. "Option Return Predictability with Machine Learning and Big Data." *Review of Financial Studies*, 36, 9: 3548-3602.

Banerjee, S. and J. Graveline. 2014. "Trading in Derivatives when the Underlying is Scarce." *Journal of Financial Economics*, 111, 3: 589-608.

Biais, B., C. Bisiére, and C. Spatt. 2010. "Imperfect Competition in Financial Markets: An Empirical Study of Island and Nasdaq." *Management Science*, 56, 12: 2237-2250.

Biais, B., J. Hombert, and P-O. Weill. 2021. "Incentive Constrained Risk Sharing, Segmentation, and Asset Pricing." *American Economic Review*, 111, 11: 3575-3610.

Bollen, N. and R. Whaleyi. 2004. "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?" *Journal of Finance*, 59, 2: 711-753.

Budish, E., P. Cramton, A. S. Kyle, J. Lee, and D. Malec. 2021. "Flow Trading." Working Paper.

Cabrales, A., D. Gale, and P. Gottardi. 2015. "Financial Contagion in Networks." The Oxford Handbook of the Economics of Networks (Y. Bramoullé, A. Galeotti and B. Rogers, eds.).

Cespa, G. 2004. "A Comparison of Stock Market Mechanisms." *RAND Journal of Economics*, 35, 4: 803-823.

Cespa, G. and X. Vives. 2022. "Exchange Competition, Entry, and Welfare." *Review of Financial Studies*, 35, 5: 2570-2624.

Chabakauri, G., K. Yuan, and K. Zachariadis. 2021. "Multi-Asset Noisy Rational Expectations Equilibrium with Contingent Claims." *Review of Economic Studies*, 89, 5: 2445-2490.

Che, Y-K. and R. Sethi. 2014. "Credit Market Speculation and the Cost of Capital." *American Economic Journal: Microeconomics*, 6, 4: 1-34.

Chen, D. 2022. "Optimal Design of a Financial Exchange." Working Paper, Stanford GSB.

Chen, D. and D. Duffie. 2021. "Market Fragmentation." *American Economic Review*, 111, 7: 2247-2274.

Du, S. and H. Zhu. 2017a. "Are CDS Auctions Biased and Inefficient?" *Journal of Finance*, 72: 2589-2628.

Du, S. and H. Zhu. 2017b. "What is the Optimal Trading Frequency in Financial Markets?" *Review of Economic Studies*, 84, 4: 1606-1651.

Duffie, D. 1986. "Stochastic Equilibria: Existence, Spanning Number, and the 'No Expected Financial Gain From Trade' Hypothesis." *Econometrica*, 54, 5: 1161-1183.

Duffie, D. 2018. "Notes on LIBOR Conversion." Working Paper, Stanford GSB.

Duffie, D. and M. Jackson. 1989. "Optimal Innovation of Futures Contracts." *Review of Financial Studies*, 2, 3: 275-296.

Duffie, D. and M. Jackson. 1990. "Optimal Hedging and Equilibrium in Dynamic Futures Market." *Journal of Economic Dynamics and Control*, 14, 1: 21-33.

Duffie, D. and R. Rahi. 1995. "Financial Market Innovation and Security Design: An Introduction." *Journal of Economic Theory*, 65, 1: 1-42.

EU Derivatives Market Report. 2023. European Securities and Markets Authority, ESMA50-524821-2930, https://www.esma.europa.eu/sites/default/files/2023-12/ESMA50-524821-2930_EU_Derivatives_Markets_2023.pdf.

Frazzini, A., R. Israel, and T. Moskowitz. 2018. "Trading Costs." Working Paper.

Fostel, A. and J. Geanakoplos. 2012. "Tranching, CDS, and Asset Prices: How Financial Innovation Can Cause Bubbles and Crashes." *American Economic Journal: Macroeconomics*, 4, 1: 190-225.

Gabaix, X. and R. Koijen. 2022. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." Working Paper.

Gârleanu, N. and L. Pedersen. 2011. "Margin-Based Asset Pricing and Deviations from the Law of One Price." *Review of Financial Studies*, 24, 6: 1980-2022.

Gârleanu, N. and L. Pedersen. 2013. "Dynamic Trading with Predictable Returns and Transaction Costs." *Journal of Finance*, 68, 6: 2309-2340.

Gârleanu, N., L. Pedersen, and A. Poteshman. 2009. "Demand-Based Option Pricing." *Review of Financial Studies*, 22, 10: 4259-4299.

Glebkin, S., S. Malamud, and A. Teguia. 2023. "Illiquidity and Higher Cumulants." *Review of Financial Studies*, 36, 5: 2131-2173.

Hart, O. 1975. "On the Optimality of Equilibrium When the Market Structure is Incomplete." *Journal of Economic Theory*, 11, 3: 418-443.

Keller, C. and M. Tseng. 2023. "Price Discovery for Derivatives." Working paper.

Klemperer, P. and M. Meyer. 1989. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, 57, 6: 1243-1277.

Knight Capital Group, eds. 2010. "Current Perspectives on Modern Equity Markets: A Collection of Essays by Financial Industry Experts." Washington, DC: Knight Capital Group.

KOIJEN, R. AND M. YOGO. 2019. “A Demand System Approach to Asset Pricing.” *Journal of Political Economy*, 127, 4, 1475-1515.

KOONT, N., Y. MA, L. PASTOR, AND Y. ZENG. 2023. “Steering a Ship in Illiquid Waters: Active Management of Passive Funds.” Working paper.

KYLE, A. S. 1989. “Informed Speculation and Imperfect Competition.” *Review of Economic Studies*, 56, 3: 517-556.

KYLE, A. S. AND J. LEE. 2017. “Toward a Fully Continuous Exchange.” *Oxford Review of Economic Policy*, 33, 4: 650-675.

KYLE, A. S. AND J. LEE. 2022. “When Are Financial Markets Perfectly Competitive?” Working Paper.

KYLE, A. S., A. A. OBIZHAEVA, AND Y. WANG. 2017. “Smooth Trading with Overconfidence and Market Power.” *Review of Economic Studies*, 85, 1: 611-662.

LETTAU, M., S. LUDVIGSON, AND P. MANOEL. 2021. “Characteristics of Mutual Fund Portfolios: Where are the Value Funds?” Working paper.

LYU, C., M. ROSTEK, AND J. H. YOON. 2022a. “Innovation in Market-Clearing Technology: Design of Multi-Asset Trading Mechanisms.” Working Paper.

LYU, C., M. ROSTEK, AND J. H. YOON. 2022b. “Futures Contracts and Imperfect Competition.” Working Paper.

MALAMUD, S. AND M. ROSTEK. 2017. “Decentralized Exchange.” *American Economic Review*, 107, 11: 3320-3362.

OEHMKE, M. AND A. ZAWADOWSKI. 2015. “Synthetic or Real? The Equilibrium Effects of Credit Default Swaps on Bond Markets.” *Review of Financial Studies*, 28, 12: 3303-3337.

RAHI, R. AND J.-P. ZIGRAND. 2009. “Strategic Financial Innovation in Segmented Markets.” *Review of Financial Studies*, 22, 8: 2941-2971.

ROSTEK, M. AND J. H. YOON. 2021. “Exchange Design and Efficiency.” *Econometrica*, 89, 6: 2887-2928.

ROSTEK, M. AND J. H. YOON. 2024a. “Innovation in Decentralized Markets: Synthetic Products vs. Trading Technology.” *American Economic Journal: Microeconomics*, 16, 1: 63-109.

ROSTEK, M. AND J. H. YOON. 2024b. “Replication Data for: ‘Financial Product Design in Decentralized Markets’,” *Harvard Dataverse*, `https://doi.org/10.7910/DVN/IBGXQU`.

SEC. 2019. Exchange-Traded Funds. `https://www.sec.gov/files/rules/final/2019/33-10695.pdf`.

SIMSEK, A. 2013. “Speculation and Risk Sharing with New Financial Assets.” *Quarterly Journal of Economics*, 128, 3: 1365-1396.

SOMOGYI, F. 2021. “Dollar Dominance in FX Trading.” Working Paper.

STULZ, R. M. 2004. “Should We Fear Derivatives?” *Journal of Economic Perspectives*, 18, 3: 173-192.

TODOROV, K. 2023. “When Passive Funds Affect Prices: Evidence from Volatility and Commodity ETFs.” Working Paper.

VIVES, X. 2011. “Strategic Supply Function Competition with Private Information.” *Econometrica*, 79, 6: 1919-66.

WALLEN, J. 2022. “Markups to Financial Intermediation in Foreign Exchange Markets.” Working Paper.

WILSON, R. 1979. “Auctions of Shares.” *Quarterly Journal of Economics*, 93, 4: 675-689.

WITTWER, M. 2021. “Connecting Disconnected Financial Markets?” *American Economic Journal: Microeconomics*, 13, 1: 252-282.

YANG, L. AND H. ZHU. 2020. “Back-Running: Seeking and Hiding Fundamental Information in Order Flows.” *Review of Financial Studies*, 33, 4: 1484-1533.

YANG, L. AND H. ZHU. 2021. “Strategic Trading when Central Bank Intervention is Predictable.” *Review of Asset Pricing Studies*, 11, 4: 735-761.

ZAWADOWSKI, A. 2013. “Entangled Financial Systems.” *Review of Financial Studies*, 26, 5: 1291-1323.

Zhang, A. 2022. "Competition and Manipulation in Derivative Contract Markets." *Journal of Financial Economics*, 144, 2: 396-413.