Risk-Sharing, Investment, and Asset Prices According to Cournot and Arrow-Debreu*

Daniel Neuhann† Michael Sockin‡

November 2020

Abstract

We study the macroeconomic consequences of financial market concentration in a complete markets economy with production. We propose a theory in which differences in preferences, productivity, and risk exposure generate gains from trade, but these gains are not fully realized because some large agents internalize their impact on asset prices. In equilibrium, risk-sharing is incomplete and rents from strategic trading feed back into the real economy by distorting the marginal value of production. Along with an increase in valuations, the model can generate a joint decline in investment, productivity, risk-free rates and the market risk premium. Welfare losses from strategic trading can be measured using asset market data, and are largest in the deepest markets.

*First version: January 2019. This version: November 2020. We thank Gian Luca Clementi (discussant), Nuno Coimbra (discussant), Winston Dou (discussant), Wenxin Du, Selman Erol, Emmanuel Farhi, William Fuchs, John Hatfield, Guido Lorenzoni, Yiming Ma, Semyon Malamud (discussant), Andrey Malenko, Gregor Matvos, Aaron Pancost, Alberto Teguia (discussant), Robert Townsend, Xavier Vives, Neng Wang, Ariel Zetlin-Jones, and seminar participants at 2nd HKUST-Jinai Macroeconomics Workshop, NASMES 2019, SED 2019, EFA 2019, 10th Tepper-LAEF Advances in Macro-Finance conference, UBC Winter Finance Conference 2020, MFA 2020, Boston College, Carnegie Mellon, and UT Austin for helpful comments. We thank Mahyar Sefidgaran for excellent research assistance.

†UT Austin. daniel.neuhann@mccombs.utexas.edu

‡UT Austin. michael.sockin@mccombs.utexas.edu

1 Introduction

There is mounting evidence that many financial markets, including centralized trading venues, are not perfectly competitive and that market participants must take into account that their trading behavior impacts asset prices (e.g. Koijen and Yogo (2019)).[1] The aggregate consequences of financial market power, however, are still not well-understood. We introduce financial price impact into a canonical general equilibrium model with production and complete markets, and use the model to examine the following questions: How does price impact interact with the distribution of productivity, risk, and preferences to distort risk sharing and asset prices in different states of the world? Can financial market concentration generate outcomes in line with recent macroeconomic trends, such as the joint decline in investment, risk-free rates, and the market risk premium?[2] What are the welfare implications of price impact, and can we use asset market data to measure its welfare costs?

In our theory, agents may differ in preferences, average productivity, and exposure to aggregate and idiosyncratic risk. This heterogeneity creates gains from trade that would ordinarily be realized in financial markets. The key friction is that some large agents, such as large financial institutions or firms, take into account that their trades affect equilibrium asset prices. As a result, strategic agents trade fewer assets than they would otherwise. While the model is similar in some respects to Cournot's approach to product markets, an important aspect of our analysis is that marginal costs and valuations are determined in general equilibrium, and that there is sorting of strategic agents into both sides of every financial market.

Agents in the economy buy claims against states with low private income for in-

[1]Koijen and Yogo (2019) show that many financial institutions have substantial price impact even in relatively liquid U.S. equity markets, with price elasticities around 3 on average. In its 2017Q4 quarterly report, the OCC estimates that over 90% of the notional amount of interest-rate swaps is accounted for by four banks, while over 95% of the CDS market is accounted for by three banks. Corbae and Levine (2018) shows that the five largest U.S. banks held 47% of total U.S. bank assets in 2015; in the U.K., France, Germany, Italy, and Canada, the range is from 71% to 84%. De Loecker, Eeckout, and Unger (2018) estimate that markups in the financial industry have increased since the 1980s. Hau, Hoffmann, Langfield, and Timmer (2017) document non-competitive pricing in derivatives markets.

[2]Gutierrez and Philippon (2017b), Fernald and Watson (2017), and Alexander and Eberly (2018) find a secular decline in investment since the early 2000s, while Fernald and Watson (2017) and Covarrubias (2019) provide evidence of lower productivity growth over this period. Laubach and Williams (2016) and Del Negro and Tambalotti (2018) find a secular decline in real risk-free rates over the past few decades; Bianchi, Lettau, and Ludvigson (2020) find a concurrent decline in risk premia.

surance and sell claims against states with high private income to fund investment. The basic mechanism is that strategic concerns prevent the realization of all gains from trade: regardless of trading motive, whether it be differences in average productivity, idiosyncratic risk exposure, or preferences, sellers restrict supply to prevent prices from falling while buyers reduce demand to prevent prices from rising.

We provide a sharp characterization of the unrealized gains from trade in every state by showing that the consumption allocation (and thus the degree of risk sharing) in our complete-markets economy with price impact is observationally equivalent to that in an economy with perfect competition but *exogenously incomplete* markets. Interestingly, this implied market incompleteness can take familiar forms. We provide examples of economies for which an outside observer would infer that agents have access only to a risk-free bond or a market index, respectively, even though markets are complete. In contrast to models with exogenously incomplete markets, however, the extent of incompleteness is endogenously determined, and is such that asset payoffs are orthogonal to residual gains from trade. As a consequence, we can the link degree of incomplete risk sharing to fundamentals such as the cross-sectional distribution of productivity risk.

The complete markets economy with imperfect competition behaves very differently from its competitive counterpart. In competitive complete-markets economies, a pure increase in cross-sectional dispersion leads to larger gross asset positions without affecting real allocations or prices. With price impact, larger gross asset positions create stronger incentives distort portfolios because there are more infra-marginal gains from trade to be exploited. Since cross-sectional dispersion has been rising over time relative to aggregate volatility (Campbell, Lettau, Malkiel, and Xu (2001)), our model suggests the distortions from price impact may have risen over time even absent an independent increase in financial market concentration.

Given these distortions to risk-sharing, we explore the asset pricing implications of financial market power. Price impact introduces an endogenous wedge between the marginal valuations (*state prices*) of strategic agents and asset prices in every state of the world. Although this typically invalidates the no-arbitrage approach to asset pricing, no arbitrage continues to hold because the model includes a competitive fringe of price-taking agents, such as households and small firms, who trade in all markets while taking prices as given. As a result, assets can be priced with the unique stochastic discount factor

(SDF) of the competitive fringe, where the effects of large agents' price impact manifest in a distorted fringe consumption process. If the fringe has convex marginal utility, price impact decreases non-linearly in fringe consumption.

In every market, buyers act strategically to lower asset prices while sellers attempt to raise them. Although this unambiguously reduces trading volumes, the net effect on prices depends on whether the demand of buyers is more elastic than the supply of sellers. Taking investment policies as given, we identify three strategic forces that favor upward pressure on asset prices, which we refer to as the *sorting*, *coordination*, and *selection* channels. The first two represent inflationary forces within a given security market, while the third characterizes relative distortions across states.

The *sorting* channel arises because agents *choose* whether to supply or demand a security. Sellers sell a particular state-contingent claim because they have high income in that state; while buyers are buyers because they have low income in the same state. Since the slope of marginal utility measures the cost of trading away from a preferred portfolio, for sufficiently similar preferences, sellers are more willing to distort their portfolios than buyers, which puts upward pressure on prices. This intuition further generalizes to the case of heterogenous risk aversion or wealth, with more risk-tolerant and wealthier agents responding more aggressively to price impact.

The *coordination* channel stems from the non-linearity of price impact under convex marginal utility. If a buyer reduces demand in a given state, the competitive fringe consumes more and price impact falls, which reduces the incentives of other buyers to reduce demand. In contrast, if a seller reduces supply, price impact increases and other sellers have an incentive to distort more. Consequently, while demand reductions are strategic substitutes, supply reductions are strategic complements, which favors price increases.

The *selection* channel provides a ranking of the distortions of price impact across security markets. Sellers most heavily distort supply in states of the world that are most dear to buyers because there are more rents to extract when demand is relatively more inelastic. In contrast to models with exogenously incomplete markets, distortions appear in states of the world that have outsized implications for asset prices and welfare.

Price distortions at the security level have aggregate implications for risk premia and the risk-free rate. Since price impact introduces dispersion in state prices, we consider a valid, consensus SDF that we recover by averaging the Euler Equations across all agents for each Arrow security and taking its implied average state price, which we refer to as the state-price implied SDF. We find that the associated risk-free rate is decreasing in price impact under natural conditions. We also consider direct measures of returns recovered from equilibrium prices alone. Fixing investment, the market-implied risk-free rate is decreasing in market power because strategic concerns raise asset prices state by state.

The behavior of the market risk premium is more subtle because there is an intricate relation between aggregate and idiosyncratic risk in the presence of strategic concerns. While trading volume responds to any shock, asset prices only respond when there is sufficient ex post asymmetry in marginal rates of substitution. The states of the world that see the highest change in asset prices are consequently those where one side of the market is particularly elastic and the other is inelastic. If these are high aggregate output states, financial market concentration compresses the market risk premium. We show a calibrated example where this is the case because there is sufficient heterogeneity in private returns even when returns are high on average.

With endogenous investment, the return distribution is governed by privately optimal capital allocation decisions. There is a tight link between price impact and investment because price impact distorts marginal valuations of state-contingent income. We formalize this intuition by defining a q-theory of investment that incorporates distortions from strategic trading. Distortions that reduce state prices (as in the case of sellers) make it less attractive to invest in technologies that yield returns in the same state, while rent extraction by other agents makes it more costly to reallocate capital through financial markets. In equilibrium, price impact reduces average productivity because agents begin to invest in dominated technologies in lieu of engaging in costly trades on financial markets. In the aggregate, moreover, inefficient risk sharing reduces total risky investment. Importantly, low investment despite low risk-free rates and a compressed risk premium is *not* a puzzle; all three observations are symptoms of the same underlying cause.

We also find that asset prices and capital may respond highly non-linearly to changes in market concentration. For low levels of price impact, increasing market concentration is primarily reflected in asset prices through a declining risk-free rate. Once agents prefer to self-insure using an inefficient storage technology, however, risky investment declines while the risk-free rate remains constant; changes in market concentration are now primarily reflected in quantities. We illustrate these non-monotonicities through a numerical example calibrated to empirical measures of price impact from Koijen and Yogo (2019). The distortion to investment can be sizable; in the benchmark calibration, a doubling of price impact leads to a decline in risky investment of about 15%.

We then explore the normative and empirical implications of financial market power. Not only is the equilibrium constrained inefficient but, perhaps surprisingly, it is the asset markets with the *highest* trading volume and *smallest* price discounts that suffer the largest distortions from price impact in that they have the largest unrealized gains from trade. In a special case, the aggregate welfare cost can be measured as *average price impact* $\times$ *price of a risk-free bond* $\times$ *average risk-free trading volumes of strategic agents*. We discuss these and other empirical predictions in detail.

We conclude with two applications. First, we show our model can account for recent evidence on persistent arbitrage opportunities in financial markets if the competitive fringe cannot trade simultaneously in all markets. Second, we explore the implications of price impact when there is a market for capital and control rights.

1.1 Related Literature

Our paper is primarily related to the literature on asset pricing and investment under imperfect risk-sharing. It is well-understood that exogenous market incompleteness can depress the risk-free rate and raise the equity premium when income shocks are sufficiently persistent (Kocherlakota (1996)). Since there are significant drawbacks to imposing exogenous barriers to trade, the literature has explored a number of microfoundations for imperfect risk sharing in endowment economies. Alvarez and Jermann (2000) and Kehoe and Levine (1993) explore solvency constraints due to limited commitment; Alvarez and Jermann (2001) show that these constraints can lower the risk-free rate and raise the equity premium. Golosov and Tsyvinski (2003) consider private information about income risk and show that asset pricing predictions are based on Inverse Euler Equations; Kocherlakota and Pistaferri (2009) show that this can lower the risk-free rate and raise the equity premium when the SDF is derived from cross-sectional consumption.

Similar to this literature, we find that price impact hampers risk sharing and lowers the risk-free rate. This is perhaps surprising because we consider an entirely unrelated trading friction. There are also a number of important differences. One is that price impact can *lower* the equity premium despite inefficient risk sharing, which is consistent with the aforementioned evidence on risk compression.[3] Another is that risk-sharing is constrained *inefficient* rather than constrained efficient in our model, which leads to the equivalence between our model and the one with exogenously, missing markets. We also provide sharp predictions for capital allocation.

Gutierrez and Philippon (2017a) and Jones and Philippon (2016) argue that imperfect competition in product markets can contribute to low investment and low interest rates, respectively. We focus on imperfect competition in financial markets rather than product markets. Given the role of financial intermediaries in capital allocation, financial market power can potentially affect investment across multiple industries.

Strategic trading under price impact has previously been investigated in the literature on liquidity provision and market microstructure following Kyle (1985) and Kyle (1989). Typically, this literature focuses on a small number of assets under CARA preferences and Gaussian payoffs. Pritsker (2005), for instance, considers large traders with price impact who can front-run distressed sales. Recently, Carvajal and Weretka (2012) and Rostek and Weretka (2015) investigate the role of constant price impact for risk-sharing in thin financial markets. We differ by considering a general equilibrium model with unrestricted payoffs and non-CARA utility in which risk-free rates are determined in general equilibrium. We also introduce a competitive fringe to pin down price impact. Malamud and Rostek (2017) study risk-sharing among agents with heterogenous coefficients of absolute risk aversion under different market structures. They also emphasize that sellers shade supply and buyers shade demand, but they do not consider production or aggregate asset pricing consequences. Glebkin, Malamud, and Teguia (2020) explore nonlinear, symmetric equilibria with large traders and non-normal payoffs without a competitive fringe, but do not allow for production or asymmetric equilibria.

A number of papers study the behavior of imperfectly competitive financial intermediaries by focusing on double-margining between banks, firms, and depositors. Examples include Drechsler, Savov, and Schnabl (2018), Corbae and Levine (2018), and Wang, Whited, Wu, and Xiao (2018). Eisenbach and Phelan (2020) investigate how price impact interacts with pecuniary externalities in incomplete markets. We consider multi-sided strategic behavior among large financial institutions in complete security markets.

[3] Although borrowing constraints also inflate asset prices, they inflate prices more for high than low marginal utility states, which increases the equity premium.

2 Model

Time is discrete and there are two dates, $t = \{1, 2\}$. The economy is populated by a finite number of strategic agents with price impact, such as large financial institutions or firms, and a unit mass of non-strategic agents who take prices as given, such as households. Uncertainty is represented by a discrete set of states of the world $\mathcal{Z}$ with cardinality $Z = |\mathcal{Z}|$. State $z \in \mathcal{Z}$ is realized at date 2 with probability $\pi(z) \in (0, 1)$. There is a single good that serves as both the consumption good and the capital good for real investment.

Since we are interested in endogenous barriers to trade, we assume that security markets are complete. There is a full set of Arrow-Debreu securities that can be traded in centralized markets at date 1. Securities can be thought of as financial claims or trade credit extended for future deliverables. The claim on state z promises one unit of the numeraire in state z and zero in all other states, and its market price is $q(z)$.

Competitive Fringe. The competitive fringe represents small firms or households who participate in financial markets but do not have market power. The presence of some price-taking agents ensures the Law of One Price and makes the framework amenable to an analysis of asset prices in general equilibrium.

The fringe has quasi-linear preferences. Utility is linear in consumption at date 1, but consumption at date 2 is evaluated using a strictly concave utility function $u_f(c_{f2}(z))$ with weakly convex marginal utility $u'_f(c_{f2}(z))$. Common utility functions satisfying these restrictions are power utility (constant relative risk aversion), exponential utility (constant absolute risk aversion), and quadratic utility (mean variance). Quasi-linearity has several specific advantages in our setting that we discuss below.

The fringe does not invest and receives endowment w_1 at date 1 and $w_2(z)$ in state z. Fringe demand for the state z claim is denoted by $D(z)$. Hence the decision problem is

$$\begin{aligned} U_f = \max_{\{D(z)\}_{z \in \mathcal{Z}}} \quad & c_{f1} + \sum_z \pi(z) u_f(c_{f2}(z)) \\ \text{s.t.} \quad & c_{f1} = w_1 - \sum_z q(z) D(z), \\ & c_{f2} = w_2(z) + D(z). \end{aligned} \tag{1}$$

The constraints ensure budget balance at date 1 and in state z at date 2, respectively.

Strategic Agents. We now describe the demographics of strategic agents with price impact. There are N *types* of strategic agents indexed by $i \in \{1, ..., N\}$. Types may differ along three dimensions: their initial endowment, their preferences, and their production technologies. These differences represent sources of gains from trade.

The endowment received at date 1 by type i is $e_i > 0$ and preferences over consumption are given by $u_1^i(c_{i1}) + u_2^i(c_{i2}(z))$. The second-period utility function $u_2^i(\cdot)$ is concave and strictly increasing, continuously differentiable and satisfies the Inada condition. Moreover, marginal utility $u_2^{i\prime}(\cdot)$ is homothetic of degree $\gamma > 0$ and strictly convex. Common preferences that satisfy these restrictions are power utility (constant relative risk aversion) and exponential utility (constant absolute risk aversion). Unless specified otherwise, the first-period utility function $u_1^i(\cdot)$ satisfies the same properties. The curvature of $u_2^i(\cdot)$ can be viewed as reflecting the preferences of a risk-averse manager or owner of a firm, and curvature over initial consumption $u_1^i(\cdot)$ as an effective budget constraint.

To isolate the effects of price impact and nest the benchmark economy with perfect competition, we assume that there are $\frac{1}{\mu_i}$ strategic agents of type i, each of which has mass $\mu_i \in [0, 1]$. If $\mu_i = 0$, we say that there is a continuum of infinitesimal agents who takes asset prices as given. Perfect competition obtains if $\mu_i = 0$ for all i. Conversely, agents have price impact for all $\mu_i > 0$.

In addition to trading in financial markets, strategic agents can transfer resources across time using an agent-specific real production technology with constant returns to scale. In the case of real firms, the technology simply represents the firm's production frontier. In the case of financial institutions, production technologies represent lending opportunities to the real economy, where productivity differences represent differences in average returns and risk available to different banks as a result of specialization.

Our specification is flexible and does not impose any particular restrictions on the distribution of returns. Type i's technology transforms an investment k_i into $y_i(z)k_i$ units of consumption good in state z. Payoffs are bounded almost surely and may be subject to both idiosyncratic and aggregate risk. All agents also have access to a storage technology that can transfer resources $s_i \geq 0$ elastically from date 1 to date 2 at a riskless rate of transformation $R > 0$. Storage is less productive than the risky technology on average, $E[y_i(z)] > R$ for all i. One can think of storage as cash, where cash holdings cannot be negative to ensure that borrowing must occur in financial markets. We introduce this technology as an inferior means of self-insurance that may be used if trading frictions in financial markets are sufficiently severe. To isolate the effects of financial market concentration, we assume that there is no price impact or market power in the real sector.

Strategic agents choose risky investment, storage, and an asset portfolio to maximize expected utility over consumption at dates 1 and 2. The key difference to the decision problem of the fringe is that strategic agents must take into account their price impact. Since all agents of a given type are symmetric, we will restrict attention to a *representative agent* of type i going forward, and refer to it simply as agent i.

Let $a_i(z)$ denote the quantity of claim z held by an agent i where negative quantities indicate sales. Let $A(z) = \sum_i a_i(z)$ be the aggregate net demand for claim z by strategic agents. By a slight abuse of notation, in the description of the decision problem we denote the perceived equilibrium pricing functional by $\tilde{q}(A(z); z)$, where we anticipate that the pricing function can be written solely as a function of the aggregate demand of strategic agents. Later we simply write $q(z)$ to economize on notation and highlight that the perceived pricing function corresponds to the correct pricing function in equilibrium. Strategic agent $i's$ decision problem can then be stated as

$$
\begin{aligned}
U_i = \max_{\{c_{i1}, k_i, s_i, \vec{a}_i\}} \quad & u_1^i\left(c_{i1}\right) + \sum_{z \in \mathcal{Z}} \pi\left(z\right) u_2^i\left(c_{i2}\left(z\right)\right) \\
\text{s.t.} \quad & c_{1i} = e_i - k_i - s_i - \sum_{z \in \mathcal{Z}} \tilde{q}(A(z); z) a_i\left(z\right), \\
& c_{i2}\left(z\right) = y_i\left(z\right) k_i + a_i\left(z\right) + R s_i.
\end{aligned} \tag{2}
$$

The constraints ensure budget balance at date 1 and in state z at date 2, respectively.

Equilibrium Concept. We use as our equilibrium concept Cournot-Nash Equilibrium in Demand Schedules. In such an equilibrium, strategic agents take beliefs, investments, demand schedules of other strategic agents, and the pricing functional for financial claims as given. We refer to equilibrium as the *market equilibrium*.

Definition 1 (Market equilibrium) *A market equilibrium is a Cournot-Nash Equilibrium in Demand Schedules consisting of strategy profiles* $\sigma_i = (c_{i1}, k_i, s_i, \{a(z), c_{i2}(z)\}_{z \in \mathcal{Z}})$ *for each strategic agent i, demand functions* $\{D(z)\}_{z \in \mathcal{Z}}$ *for the competitive fringe, and a set of pricing functions* $\{q\left(z\right)\}_{z \in \mathcal{Z}}$ *such that:*

1. *Fringe demand functions solve decision problem* (1).

2. *For each i,* σ_i *solves decision problem* (2) *given* σ_{-1} *and* $\{q\left(z\right)\}_{z \in \mathcal{Z}}$.

3. *Each market clears with zero excess demand:* $A(z) + D(z) = 0$ *for all* $z \in \mathcal{Z}$.

4. *All agents have rational expectations of their and others' strategies.*

2.1 Discussion of Assumptions

Our model features large strategic agents who have access to real investment opportunities and a competitive fringe which participates in all asset markets. The competitive fringe represents small end-investors in the financial sector, such as households, who participate in financial markets either directly or indirectly. Collectively, these end-investors are deep-pocketed. A simple way of capturing this feature is to assume that the fringe has quasi-linear preferences. A benefit of this assumption is that asset prices are separable across states *conditional on consumption*, and price impact can be recovered in closed form. This allows us to incorporate heterogeneity in preferences, productivity risk, and wealth, and to consider asymmetric equilibria, all of which has proven challenging for the literature following Kyle (1985) and Kyle (1989). The presence of the fringe also ensures the Law of One Price holds, which is crucial for an equilibrium theory of asset pricing.[4]

State-separability in the pricing functional does *not* imply that there are no cross-state linkages in optimal portfolios. Since we do not impose quasi-linear preferences on strategic agents, these agents *do* take into account price impact in all states when choosing investments and portfolios. A state-separable pricing functional only ensures that a change in strategic demand in state z does not *directly* affect price *impact* in state $z' \neq z$; it does, however, indirectly affect it through the portfolio rebalancing of strategic agents.

3 Characterization of Market Equilibrium

We now characterize several basic properties of the equilibrium and clarify the main strategic distortions. In the sequel, we derive the real and financial implications of these distortions.

[4]Section 5.1 shows this not the case if the fringe does not participate in all asset markets.

3.1 Equilibrium Pricing Functional and Price Impact

We first characterize the pricing functional and equilibrium price impact. These two objects represent the sources of strategic interaction among the large agents. We then prove the existence of equilibrium and discuss the directions of financial market distortions as a result of price impact. Finally, this will allow us to construct a counterfactual economy with exogenously incomplete markets that delivers the same degree of risk sharing as our complete-markets economy with price impact.

The equilibrium pricing functional can be derived directly from the decision problem of the competitive fringe. In particular, the first-order conditions for portfolio optimality require that market prices are equal to the fringe's marginal utility. Under quasilinearity, this delivers a closed-form solution for the pricing functional. By market-clearing, each strategic agent can then infer how much the fringe's consumption will move when the agent demands more or less of a given security, holding other agents' portfolios fixed. This observation delivers a closed-form expression for price impact. Since each agent's influence on aggregate consumption scales with her mass μ_i, so does her price impact. As a consequence, agents do not have price impact in the competitive equilibrium benchmark in which $\mu_i = 0$. We summarize this discussion in the following proposition.

Proposition 1 *The Law of One Price holds. All available assets are traded, but investment, consumption, and prices are invariant to the introduction of redundant assets. Asset prices satisfy*

$$q(z) = \pi(z) u_f'\left(c_{f2}(z)\right) \qquad \textit{where} \qquad c_{f2}(z) = w_2(z) - A(z). \tag{3}$$

Price impact is symmetric across agents and satisfies

$$\frac{\partial q(z)}{\partial a_i(z)} = \mu_i q'(z) \qquad \textit{where} \qquad q'(z) \equiv \frac{\partial q(z)}{\partial A(z)} = -\pi(z) u_f''\left(c_{f2}(z)\right) > 0. \tag{4}$$

Our model retains two important properties for asset pricing: the Law of One Price holds, and equilibrium allocations are invariant to the introduction of redundant assets. The first result follows because the fringe takes prices as given and participates in all asset markets; as such, it would trade to eliminate any mispricing between assets. The neutrality result follows because prices depend on the fringe's consumption. Any combination of strategic agents' portfolios that delivers the same consumption to the fringe consequently induces the same asset prices, and, given these asset prices, the same consumption allocation is optimal for strategic agents. In contrast to perfectly competitive models, however, all redundant assets *must be traded* in equilibrium even though their presence does not affect real allocations. This is because untraded assets can be traded at an infinitesimal price impact in a neighborhood around zero quantities.

To guide economic intuition, it is useful relate price impact to standard measures of risk aversion for the competitive fringe. We define $\epsilon_f(z) \equiv q'(z)/q(z)$ to be the semi-elasticity of $q(z)$, and $\tilde{\epsilon}_f \equiv q''(z)/q'(z)$ to be semi-elasticity of price impact. They are equal to the coefficient of absolute risk aversion and the coefficient of prudence of the fringe, respectively. Whenever fringe risk aversion is non-linear in consumption, the model gives rise to price impact that is non-linear price in the fringe's consumption. This differentiates our model from the literature following Kyle (1985), in which price impact is linear, and will be relevant when discussing positive implications for the risk-free rate and risk premia.

3.2 Optimal Portfolios and Existence

We now investigate how strategic agents chose portfolios and real investment with price impact. Given that our economy features imperfect competition, it is not obvious that an equilibrium exists. We show that this is the case, and characterize the first-order conditions for optimality. To summarize these conditions, it is convenient to define a *state price* $\Lambda_i(z)$ for agent i and state z as the ratio of expected marginal utility in state z and marginal utility at date 1,

$$\Lambda_i\left(z\right) \equiv \frac{\pi\left(z\right) u_2^{i\prime}\left(c_{i2}\left(z\right)\right)}{u_1^{i\prime}\left(c_{i1}\right)}. \tag{5}$$

The next proposition proves the existence of an equilibrium in which optimal policies are homogenous of degree 1 and therefore scale in initial wealth. This is a useful property for our model in which changes in an agent's mass are reflected in her price impact.

Proposition 2 *There exists an equilibrium in which the optimal policies of agent i are homoge-*

neous of degree 1 in e_i. Optimal policies for c_{i1}, k_i, s_i, and $a_i(z)$ satisfy the optimality conditions

$$\begin{aligned}
k_i: &\qquad \sum_{z\in\mathcal{Z}} \Lambda_i(z)\, y_i(z) \leq 1 \;\; (\textit{and} = \textit{if } k_i > 0), \\
s_i: &\qquad \sum_{z\in\mathcal{Z}} \Lambda_i(z) \leq 1, \;\; (\textit{and} = \textit{if } s_i > 0), \\
a_i(z): &\qquad \Lambda_i(z) = q(z) + \mu_i q'(z) a_i(z).
\end{aligned}$$

The optimality conditions for risky investment and storage appear standard since agents equate the state-price weighted expected return to the marginal cost of investing. Price impact, however, feeds back to investment by distorting marginal rates of substitution in consumption. We return to this issue in detail in Section 4.3.

Since large agents internalize price impact, their first-order conditions for optimal asset holdings include an endogenous wedge between Arrow prices and state prices. This is because they distort trading volumes on the margin to tilt Arrow prices in their favor. In addition, asset positions enter multiplicatively because there are more infra-marginal benefits to distorting prices when an agent has a larger position.

We next discuss several basic properties of the market equilibrium by distinguishing the strategic incentives of buyers and sellers. Gains from trade can arise from differences in income risk and/or preferences, and we use the dispersion in state prices across buyers and sellers to measure the unrealized gains from trade. Since all gains from trade are exhausted in the benchmark with perfect competition (denoted by superscript CE), we can attribute any unrealized gains from trade as a distortion from price impact. Similarly, we can measure the existence of gains from trade by asking whether trading volumes would be strictly positive under perfect competition.

The next result characterizes these distortions, and shows that state price dispersion increases when there is more dispersion in income risk and/or preferences. To isolate the role of financial market distortions, in what follows we hold investment policies fixed.

Proposition 3 *In the market equilibrium, distortions from price impact have the following properties:*

(i) The market equilibrium is equivalent to the competitive equilibrium if and only if there are no gains from trade in the competitive equilibrium, $a_i^{CE}(z) = 0$ for all i and z.

(ii) If there are gains from trade in the competitive equilibrium, then in the market equilibrium sellers supply less and buyers demand less than in the competitive.

(iii) State prices are dispersed in the market equilibrium. If j is a seller of claim z and i is a buyer, then the seller has a strictly lower state price than the buyer, $\Lambda_i(z) > \Lambda_j(z)$.

(iv) State price dispersion is increasing in the cross-sectional dispersion of productivity or risk aversion.

All results follow from the first-order condition for financial assets in which price impact enters as a multiple of $a_i(z)$, and price impact is zero if μ_i. An example in which trading volumes are zero because there are no gains from trade is pure systematic risk and homogeneous preferences. More important is the converse: distortions from price impact distorts financial trading whenever there are gains from trade. This is because it is strictly optimal for sellers to marginal shade down supply and for buyers to shade down demand in order to capture infra-marginal rents. State prices therefore reveal that there is inefficient income inequality ex post: sellers retain too much state-contingent wealth and buyers attain too little.

The last statement describes an important interaction between price impact and fundamental gains from trade, such as dispersion in risk and/or preferences. Fundamental dispersion creates larger *potential* gains from trade. In a competitive model, this leads to higher trading volumes without affecting equilibrium allocations. With price impact, however, higher trading volumes raise the infra-marginal benefits of distorting prices. Agents respond by distorting more and *realizing fewer* gains from trade. As a result, there are allocational and asset pricing consequences to any shock to idiosyncratic dispersion even when the aggregate income frontier is unchanged.

3.3 Endogenous Market Incompleteness

The previous section shows that price impact leads to state price dispersion between any pair of buyers and sellers whenever there are gains from trade, but did not address the general equilibrium consequences of this result. State price dispersion is the hallmark of incomplete-markets models with heterogeneous agents in which exogenous barriers to trade (such as missing markets) prevent the alignment of marginal rates of substitution.

This suggests an equivalence between our model and a competitive counterfactual with appropriately chosen market incompleteness.

The next proposition constructs a mapping between our complete-markets model with price impact and an incomplete market structure with perfect competition. Specifically, we define a mapping to a competitive economy with the equivalent asset span, and show that the asset span is incomplete. This result clarifies that price impact leads to misaligned marginal rates of substitution in equilibrium. A useful interpretation of the result is that an outside econometrician who observed preferences and the consumption allocation of the market equilibrium (but not initial endowments) would infer that markets must be incomplete, subject to the maintained hypothesis that markets are perfectly competitive.

A market structure is defined by asset prices and payoffs. Recall that there are Z states of the world and N types of agents. We summarize the market structure by the number of assets M and a $M \times Z$ dividend-yield matrix d where each row indexes the payoff of an asset across states of the world. Recall that agent i's state price for state z is

$$\Lambda_i(z) = q(z) + \mu_i q'(z) a_i(z) \;\; \forall \;\; (i, z),$$

and let the *market-implied state price* be the cross-sectional average of agents' state prices,

$$\Lambda^*(z) \equiv \mathbb{E}_i[\Lambda_i(z)].$$

Next, define $\vec{q}_i' = [\mu_i q'(1), \mu_i q'(2), \ldots, \mu_i q'(Z)]^T$ to be the vector of price impacts for agent i for all states, $\vec{a}_i = [a_i(1), a_i(2), \ldots, a_i(Z)]^T$ the vector of agent i's asset positions for all states, and $\vec{\Lambda}_i = [\Lambda_i(1), \Lambda_i(2), \ldots, \Lambda_i(Z)]^T$ the vector of agent i's state prices for all states, where superscript T denotes the transpose. Lastly, let ι_M be the $M \times 1$ vector of ones.

Proposition 4 *There exists a market structure with $M < Z$ assets, indexed by a counterfactual dividend-yield matrix d. that replicates the asset span in Proposition 2. Matrix d satisfies*

$$d^T \vec{\Lambda}_i = \iota_M \qquad \text{and} \qquad d^T \left(\vec{q}_i' \odot (\vec{a}_i - \vec{a}_j) \right) = 0_{M \times 1} \qquad \text{for all } i \text{ and } j$$

The extent of implied market incompleteness as measured by unrealized gains from trade satisfies:

$$Cov\left(\Lambda^*(z), \Lambda_i(z) - \Lambda_j(z)\right) = 0,$$

The first statement reveals that intentional mispricing arising from imperfect competition is isomorphic to forced disagreements about state prices stemming from missing markets with perfect competition. As this is reflected in prices, one can infer the degree of market concentration from the implied market incompleteness. The second statement reveals that the incomplete market span is such that any potential gains from trade are unpriced. That is, any unrealized gains from trade are interpreted as untradeable under the maintained hypothesis of perfect competition.

The following examples illustrate the economic content of the proposition by considering two economies in which an outside econometrician would infer that there is only a risk-free bond or a levered market index is available to trade, respectively, despite markets being complete. Hence price impact can have stark implications for agents' ability to realize gains from trade.

Example 1 (Pure Idiosyncratic Risk) *There are two types of agents, $i \in \{1, 2\}$, and two equally likely states, $z \in \{1, 2\}$ with $\pi(z) = \frac{1}{2}$. Each type invests some capital in the risky technology, $k_i > 0$. All risk is idiosyncratic, $y_i(i) = y_h$ and $y_i(-i) = y_l < y_h$. Hence in every state of the world one type has a high return and the other has a low return. Assume further that preferences and endowments are symmetric for all i. Then agents are ex-ante symmetric, and there exist two distinct state prices Λ^- and $\Lambda^+ > \Lambda^-$ such that $\Lambda_i(i) = \Lambda^-$ and $\Lambda_i(-i) = \Lambda^+$. Every agent assigns a low marginal value of consumption to the state with high private returns, and a high marginal value of consumption to the state with low private returns.*

Since there are two states, the equivalent incomplete market structure has a single asset. The construction from Proposition 4 shows that the dividend-yield $\{d_z\}_{z=1,2}$ satisfies

$$\begin{bmatrix} \Lambda^h & \Lambda^l \\ \Lambda^l & \Lambda^h \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Solving this equation gives

$$d_1 = d_2 = \frac{1}{\Lambda^h + \Lambda^l} = r_f^*.$$

This is a risk-free bond whose return is the inverse of sum of state prices. By definition, this is the state-price implied risk-free rate. Since distortions from price impact are increasing in cross-sectional dispersion, a mean-preserving spread of $y_h - y_l$ leads to a lower risk-free rate.

With two types facing purely idiosyncratic risk, consumption allocations are such that an outside observer would infer that only a risk-free bond is available to trade even though markets are complete. The reason is that the risk-free bond only allows agents to shift resources across time, whereas there are gains from trade in shifting resources across states. Since these gains are not realized, the outside observer interprets this as prima facie evidence that they are not tradeable in the market, given a maintained hypothesis of perfect competition. (Some risk is shared even with price impact, which is why we require that endowments are unobservable.) Note also that the implied interest rate is the state-price implied risk-free rate, since this is the return at which individuals are indifferent toward buying more bonds.

In the previous example, gains from trade stemmed from idiosyncratic risk exposure under symmetric preferences. We now consider the polar opposite case with heterogenous risk aversion but symmetric risk exposure.

Example 2 (Pure Aggregate Risk) *Let there be two types of agents, $i \in \{1, 2\}$, and two states of the world, $z \in \{l, h\}$. The measure of each agent type is identical and strictly positive, $\mu_1 = \mu_2 = \mu > 0$. Assume that each type invests some capital in the risky technology, $k_i > 0$. There is only aggregate risk, $y_i(h) = y_h$ and $y_i(l) = y_l$ for all i. Risk attitudes are heterogeneous: Type 1 is strictly risk-averse and Type 2 is risk-neutral.*

By risk-neutrality, Type 2 has a constant state price $\Lambda_2(z) = \Lambda^* > 0$. For Type 1, risk aversion implies that there are two distinct state prices satisfying $\Lambda_1^h < \Lambda^*$ and $\Lambda_1^l > \Lambda^*$. This is because risk sharing is imperfect since $\mu > 0$. Since there are two states, the implied incomplete market structure has one asset. The construction from Proposition 4 shows that the dividend-yield $\{d_z\}_{z=l,h}$ satisfies

$$\begin{bmatrix} \Lambda^* & \Lambda^* \\ \Lambda^h & \Lambda^l \end{bmatrix} \begin{bmatrix} d_h \\ d_l \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Solving this equation gives

$$d_h = \frac{\Lambda^l - \Lambda^*}{\Lambda^* \left(\Lambda^l - \Lambda^h\right)} \qquad \text{and} \qquad d_l = \frac{\Lambda^* - \Lambda^h}{\Lambda^* \left(\Lambda^l - \Lambda^h\right)}.$$

The expected return is the risk-neutral agent's required return $\frac{1}{\Lambda^*}$, and the asset carries exposure to aggregate risk since $d_h > d_l$ by Jensen's inequality. Hence it is a levered market index.

Under perfect competition, the risk-neutral agent would fully insure the risk-averse agent by taking a long position in the high aggregate state and an offsetting short position in the low state. Since agents do not align their state prices state-by-state, however, an outsider would conclude that they inefficiently shared aggregate risk by trading a market index with a volatility determined by the differences in state prices. Again, not all gains from trade are realized because the market index is tilted in such a manner such that perfect insurance is infeasible.

There are important differences between our setting and canonical models of endogenous incompleteness such as Alvarez and Jermann (2000). There, markets are complete and risk sharing is constrained by limited commitment. Despite this friction, state prices are still aligned state-by-state because agents realize all feasible gains from trade. As such, an outside observer would infer that markets are indeed complete. Price impact therefore offers a distinct form of endogenous incompleteness because state prices are *misaligned*, and welfare losses and asset price distortions are of first-order importance.

4 Real and Financial Implications

So far we established misalignment of state prices as the key force shaping equilibrium outcomes. We now turn to a substantive analysis of asset prices, real investment, and welfare given these trading distortions. We then illustrate our results using a calibrated numerical example and discuss empirical predictions.

4.1 Pricing and Strategic Interaction at the Security Level

We begin with an analysis of asset prices at the security level and explain how they are influenced by strategic interactions. To isolate purely financial distortions, we first hold investment policies fixed and ask how price impact affects asset prices for a given quantity of risk. In Section 4.3, we consider endogenous changes in the quantity of risk. Since asset prices are fully determined by the consumption process of the competitive fringe, the analysis boils down to asking how strategic interactions shape the distribution of consumption across states and, specifically, whether strategic buyers or sellers adjust consumption more aggressively in response to a change in price impact.

It is useful to recall why an agent chooses to buy or sell a particular claim. Given any investment policy, an agent choose to be a seller if her state price in state z is below the market price $q(z)$ of the associated security, and a buyer otherwise. Fixing an asset position, state prices are determined by preferences and income process alone, and state prices are strictly decreasing in state-contingent income and asset holdings.

Strategic incentives to distort are influenced by two main considerations: how costly it is to move away from a preferred asset position, and how elastic prices are with respect to a change in an agent's asset position. The former is measured by the curvature of marginal utility (or, equivalently, state prices) and the latter by price impact. Both are affected by equilibrium considerations. We separately consider each side of the trade-off and find that equilibrium considerations favor price *increases*.

Consider first the costs of distorting. When preferences are homogenous, agents sort into the buy and sell side of a given market based on the distribution of income only. Specifically, agents with high state-specific income choose to be sellers and those with low state-specific income choose to be buyers. Under symmetric preferences, sellers must then have a strictly flatter marginal utility curve for *any* asset position, and so it finds it less costly to distort quantities. Hence sellers react more elastically to an increase in price impact, where we define the elasticity of demand or supply for strategic agent i as

$$\epsilon_i(z) \equiv -\frac{\partial a_i(z)}{\partial \mu_i}\frac{\mu}{a_i(z)}. \tag{6}$$

Since sellers have an interest in preventing prices from falling, this effect leads to upward pressure on asset prices state by state. We refer to this mechanism as the *sorting channel* of price impact. It is summarized in the following proposition, where we focus on quasi-linear preferences for simplicity. This allows us to treat each state separately and removes cross-state spillovers in optimal distortions. The proof of the proposition considers the general case with cross-state spillovers and shows that the same basic forces operate.

Proposition 5 *Let strategic agents have linear preferences over date 1 consumption and $\mu_i = \mu$. Then the quantity response to a change in price impact satisfies*

$$\frac{\partial a_i(z)}{\partial \mu_i} = \frac{-q'(z)a_i(z)}{(1+\mu)q'(z) + \mu q''(z)a_i(z) - \frac{\partial \Lambda_i(z)}{\partial a_i(z)}}. \tag{7}$$

If preferences are symmetric, then the absolute value of elasticity $\epsilon_i(z)$ is strictly higher for sellers than buyers in every state.

The sorting channel appears in the last term of the denominator in equation (7). By the convexity of marginal utility, the derivative of the state price is smaller in absolute terms for sellers since they have more income.

Of course, the curvature of marginal utility is also directly affected by preferences, hence preference heterogeneity may reverse the sorting channel. This is illustrated by the next example, in which a risk-neutral agent can tilt asset prices in her favor even as a buyer. Hence the *joint* distribution over income and preferences matters for asset prices.

Example 3 (Heterogeneous Risk Aversion) *Let there be two types $i \in \{1, 2\}$, two states $z \in \{l, h\}$ and assume masses satisfy $\mu_1 = \mu_2 = \mu$. There is only aggregate risk, $y_i(h) = y_h$ and $y_i(l) = y_l$ for all i. Type 1 is risk-averse and Type 2 is risk-neutral. The competitive fringe has a constant coefficient of absolute risk aversion γ_f. In equilibrium, expected returns are equal to*

$$\frac{\pi(z)}{q^*(z)} = 1 + \gamma_f \mu a_2^*(z).$$

where $a_2^*(h) > 0$ and $a_2^*(l) < 0$. As a result, prices are distorted upward when the risk-neutral agent is a seller, and distorted downward when the risk-neutral agent is a buyer.

Now consider the benefits of distorting. Price impact measures the extent to which a marginal change in portfolios moves prices. In our model, price impact is equal to the derivative of marginal utility of the competitive fringe. With convex marginal utility, price impact itself depends on the *level* of consumption, which in turn is determined by the aggregate demand of strategic agents. Hence quantity shading by one strategic agent imposes a pricing externality on others. This is most transparently seen using a second order approximation to the pricing function:

$$\Delta q(z) \approx q'(z) \sum_i \Delta a_i(z) + \frac{1}{2} q''(z) \sum_i \sum_j \Delta a_i(z) \cdot \Delta a_j(z) \tag{8}$$

The first term is the change in demands evaluated at the slope of the pricing function. The second term reflects strategic interactions. Recall that buyers reduce demand ($\Delta a(z) < 0$) while sellers decrease supply ($\Delta a(z) > 0$). This leads to the following result.

Proposition 6 *Strategic externalities under non-linear price impact operate as follows:*

(i) For buyers, demand reductions are a strategic substitute. If one buyer reduces demand, price impact falls and other buyers' have weaker incentives to reduce demand.

(ii) For sellers, supply reductions are a strategic complement. If one seller reduces supply price impact increases and other seller's marginal incentive to reduce supply also increases.

(iii) If the preferences of the competitive fringe satisfy constant absolute risk aversion, then states of the world in which prices are high because additional consumption is highly valuable to buyers are most heavily distorted by sellers. Formally, the elasticity of supply is increasing in security prices, $\frac{\partial \epsilon_i}{\partial q(z)} > 0$.

Strategic externalities are reflected in the second term of the denominator in equation (7). The term is negative for sellers (since $a_i < 0$ for sellers) and grows in absolute value as price impact rises. Strategic externalities therefore represent a second channel that favors price increases at the expense of buyers. The following example illustrates.

Example 4 (Strategic Interactions) *There is one buyer and one seller. The seller reduces supply by Δ, and the buyer reduces demand by $\alpha\Delta$ for some $\alpha < 1$. The change in the asset price satisfies*

$$\Delta q(z) \approx \underbrace{q'(z)(1-\alpha)\Delta}_{\text{Net Demand Channel}>0} + \underbrace{\frac{\Delta^2}{2} q''(z)\left(1+\alpha^2-2\alpha\right)}_{\text{Strategic Interaction}>0}$$

The third statement in Proposition 6 shows that sellers most heavily distort the states that are most valuable to buyers, where we measure value by the asset price. Since prices reflect marginal valuations, high prices also indicate a steep marginal utility curve for buyers, and therefore an unwillingness to adjust quantities aggressively. As such, sellers can extract more rents, and so they respond more aggressively. This implies that buyers obtain even fewer scarce securities. We establish this result in the special case of CARA because this implies a constant semi-elasticity of price impact and makes the economic forces particularly transparent. The same mechanisms, however, are present with more general preferences.

Taken together, we have established three channels of strategic interactions that favor price increases. In the case of homogenous preferences, these channels are sufficient to generate an increase in prices state by state. Sufficiently larger heterogeneity in risk aversion, however, can overturn this result if coupled with an appropriate distribution of income across states. We clarify the aggregate implications of these channels in the following subsection.

4.2 Aggregate implications: Risk-free rates and risk premia

We now map price distortions at the security level into predictions for aggregate returns such as the risk-free rate and risk premia. For some results, it will be instructive to consider the special case of a *passive fringe* in state z, by which we mean that the fringe does not sell any assets in the competitive equilibrium in state z where $\mu_i = 0$. This is to ensure that price movements are primarily governed by the strategic interactions between large buyers and sellers. There always exists a fringe endowment process such that this condition is satisfied.

Assumption 1 (Passive Fringe) *$\{w_2(z)\}_{z\in\mathcal{Z}}$ is such that $D(z) \geq 0$ if $\mu_i = 0$ for all i.*

Price impact introduces a divergence between market prices and valuations based on state prices. State prices determine the required returns of an asset if agents could trade it *without price impact*, while market prices determine actually prevailing returns. They coincide in canonical models of perfect competition, but differ here. We characterize both classes of equilibrium returns, and show that their divergence is directly linked to the degree of equilibrium market power. In what follows, we define the state price-implied risk-free rate r_f^* and the market-implied risk-free rate r_f^m as

$$r_f^* = \left[\sum_{z\in\mathcal{Z}} E^*\left[\Lambda_i\left(z\right)\right]\right]^{-1} \qquad \text{and} \qquad r_f^m = \Big[\sum_{z\in\mathcal{Z}} q(z)\Big]^{-1}$$

respectively, where $E^*\left[\cdot\right]$ is the cross-sectional average. Claim z's market-implied excess return is $r\left(z\right) = \frac{\pi(z)}{q(z)} - r_f^m$. The *average market power exerted in state* z is $\overline{mkt}\left(z\right) = q'(z)\frac{A(z)}{N}$. Taking a cross-sectional average of the first-order condition for $a_i(z)$ gives the following relationship between state prices, price impact, and asset prices:

$$q(z) = E^*[\Lambda_i(z)] - \overline{mkt}\left(z\right) \tag{9}$$

The next proposition summarizes price responses to price impact, and states our model's key asset pricing implications. Since price impact is irrelevant for untraded assets that are not traded, statements regarding claim z apply only if the claim is traded.

Proposition 7 *The following results characterize asset prices under price impact:*

(i) Suppose preferences are homogenous and Assumption 1 holds. If not all agents invest in storage, then asset prices are higher state-by-state than in the competitive equilibrium. Moreover, the market-implied risk-free rate r_f^m is lower than in the competitive equilibrium and bounded below by the return on storage R. If all agents employ storage, then some states may have lower prices than in the competitive equilibrium.

(ii) The state-price implied risk-free rate r_f^* is lower than in the competitive equilibrium and is bounded below by the rate of return on storage R.

(iii) The wedge between r_f^* and r_f^m is determined by market power: $\frac{1}{r_f^*} - \frac{1}{r_f^m} = \sum_z \overline{mkt}(z)$.

(iv) The market-implied excess return $ER(z)$ of security z satisfies the following decomposition:

$$ER(z) = \underbrace{-Cov\left(\frac{\bar{u}'(z)}{E[\bar{u}'(z)]}, \frac{1}{q(z)}\right)}_{\text{Risk Premium}} + \underbrace{\left(r_f^m - r_f^* - \overline{mkt}(z)\, r_f^*\right)}_{\text{Relative Market Power}},$$

where $\bar{u}'(z) = E^*\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]$ is the cross-sectional SDF. As such, assets in which sellers exert high market power earn a lower excess return that cannot be explained by risk exposure.

Statement (i) shows that sellers succeed in raising prices for any investment policy with limited use of storage if preferences are homogeneous and the competitive fringe is passive. This a simple application of the sorting channel described in the previous section. The assumption of a passive fringe is useful because it ensures that the total demand of strategic agents equals their total supply in a neighborhood around perfect competition. If the storage constraint binds for all agents, however, then the sum of state prices is fixed at the storage rate, R, and an increase in market concentration instead increases the price of aggregate high marginal utility claims, but lowers them for aggregate low marginal utility claims. This is because of the *selection channel* described in the previous section, whereby high-priced states are most heavily distorted by sellers.

Statement (ii) shows that the state-price implied risk-free rate falls under general conditions. This is because inefficient risk-sharing raises average marginal utility across states and agents. Statement (iii) makes precise the wedge between state price-implied returns and market-implied returns: it is determined by average market power exerted in financial markets. Statement (iv) shows that we can decompose the expected excess return of debt contingent on state z into two components: a risk premium that reflects the covariance between the Arrow-Debreu security and the average marginal utility of purchasing agents, and a bias that reflects the average market power of the agents that trade claims. In this precise sense, market power leads to mispricing.

We now turn to a comparison of excess returns across states. The excess return in state z can be approximated to first-order around the competitive equilibrium as

$$ER(z) - ER^{CE}(z) \approx -\frac{\pi(z)\, q'(z)}{q(z)^2}\,\Delta D(z) + \left(r_f^m\right)^2 \sum_{z'} q'(z')\,\Delta D(z'),$$

where $\Delta D(z)$ is the change in aggregate net demand by strategic agents and $q'(z)$ is price impact evaluated at fringe consumption in the competitive equilibrium.[5]

Market power distorts excess returns along two dimensions: (i) it lowers the return to the Arrow security of state z according to the product of $\frac{\pi(z)q'(z)}{q(z)^2}$ and the net change in strategic demand $\Delta D(z)$; and (ii) it alters the risk-free rate by the sum of price distortions across all states. Hence the excess return earned in a given state increases if its price is more distorted than the average state.

This provides some insights into the market risk premium. The key question is whether high aggregate states, which have lower Arrow prices, are more distorted than average. If market power primarily raises prices of high aggregate states, it contributes to a decline in the market risk premium; if it primarily raises prices of low aggregate states, it contributes to an increase.

The net distortion of sellers relative to buyers is given by $\Delta D(z)$, where $\Delta D(z) > 0$ if sellers distort more aggressively than buyers, and prices are increasing in $D(z)$, all else equal. Holding fixed $\Delta D(z) > 0$, states in which $\frac{\pi(z)q'(z)}{q(z)^2}$ is larger see a sharper decline in excess returns. Under strict convex marginal utility, these are states with high fringe consumption, such as high aggregate states. The following statement is now obvious.

Corollary 1 *All else equal, the excess return in state z is more likely to be higher than in the competitive equilibrium if (i) $\Delta D(z)$ is lower than the average $\Delta D(z)$, and (ii) the consumption of the fringe in the competitive equilibrium is higher than the average.*

It is difficult to characterize the relative magnitudes of $\Delta D(z)$ without further restrictions on production returns. This is because shocks to the income distribution may lead to subtle differences in incentives to distort. In the previous section, we showed that a shock

[5]The derivation of this expression is in the proof of Corollary 1.

which raises incomes lowers the cost of distorting by flattening marginal utility. As such, an aggregate shock that raises incomes for all agents unambiguously leads to a reduction in trading volumes, but the impact on asset prices is more subtle. If buyers and sellers respond symmetrically, then the consumption of the fringe is unchanged and asset prices do not adjust. Change in prices (and the risk premium) therefore require a particular correlation structure between idiosyncratic and aggregate risk. Section 4.4 provides a numerical example where the risk premium is indeed decreasing in price impact.

4.3 Investment: A strategic q-theory

We now study the investment consequences of financial market concentration. Agents make investment decisions taking into account that differences in production technologies make them natural sellers of certain securities and buyers of others. Conversely, distorted state prices affect the marginal value of an additional investment. The franchise value V_i of agent i's production is $V_i = \sum_{z \in \mathcal{Z}} \Lambda_i(z)(y_i(z)k_i + Rs_i)$. Substituting for state prices using first-order conditions for asset positions, we can decompose this value as

$$V_i = \underbrace{\sum_{z \in \mathcal{Z}} q(z)(y_i(z)k_i + Rs_i)}_{\text{Market Value of Operations}} + \underbrace{\sum_{z \in \mathcal{Z}} \mu_i q'(z) a_i(z)(y_i(z)k_i + Rs_i)}_{\text{Trading Profits}} \tag{10}$$

The first term represents the market value of agent i's production. Under perfect competition, this term fully determines agent i's franchise value. Price impact introduces the second term which reflects that agent i can profit from strategically selling claims against its production and buying claims against the production of others in financial markets.

The next proposition shows that our model gives rise to a strategic q-theory of investment in which marginal cost is now equated to not only marginal discounted expected revenue as in standard q-theory, but also the value-added from strategic trading. We define $mv_i \equiv \sum_{z \in \mathcal{Z}} q(z) \frac{(y_i(z)k_i + Rs_i)}{k_i + s_i}$ and $p_i \equiv \sum_{z \in \mathcal{Z}} \mu_i q'(z) a_i(z) \frac{(y_i(z)k_i + Rs_i)}{k_i + s_i}$ to be the market value and trading profit multipliers that generate franchise value V_i given investment k_i in accordance with equation 10. We identify *dominated technologies* by comparing outcomes in the strategic economy to two benchmarks: (i) perfect competition; and (ii) financial autarky (no trading of financial claims). Since the competitive equilibrium is efficient, an agent chooses a larger scale of production under autarky than in the competitive equilibrium only if his technology is dominated by other technologies.

Proposition 8 *At an interior optimum for risky investment under price impact, k_i^* satisfies*

$$mv_i + p_i = 1,$$

while in the competitive equilibrium optimal investment satisfies $mv_i = 1$. As such, an agent who earns negative trading profits ($p_i < 0$) chooses a larger scale of production than in the competitive equilibrium, while an agent who earns positive trading profits ($p_i > 0$) chooses a smaller scale of production. This has the following implications:

(i) an agent with a dominated technology may choose a larger scale of production than in the competitive equilibrium.

(ii) aggregate risky investment and expected productivity are lower than in the competitive equilibrium, and aggregate investment in storage is higher.

The first statement shows that trading profits and real investment are substitutes. When an agent can earn rents in financial markets by restricting the supply of financial assets, she may choose to also restrict real investment to further boost these rents. Conversely, an agent who must pay rents in financial markets may choose to invest on her own. Since these are agents with poor real investment opportunities, an implication of this result is that agents may invest in dominated technologies in order to escape financial market competition. *Total* risky investment, however, must fall. This is because price impact hampers risk sharing and makes it less attractive to hold risky returns. In the aggregate, agents are therefore less willing to take on risk than in the competitive equilibrium. Moreover, changes in production also feed back into asset markets. If agents make fewer risky investments, aggregate income falls and certain securities become even scarcer. This, in turn, makes it more profitable to distort the supply of financial claims.

We illustrate the misallocation result in a highly stylized example without risk.

Example 5 (Misallocation) *There are two types of agents. Production technologies satisfy $y_1(z) = y^h$ and $y_2(z) = y^l \in (R_f, y_h)$ so that Type 2's production technology is strictly dominated. Let $\mu_1 = \mu_2 = \mu$. There exists a threshold $\bar{\mu}$ such that $k_2^* > 0$ if and only if $\mu > \bar{\mu}$.*

The inefficient technology is never used in the competitive equilibrium because it is socially optimal to share returns from the efficient technology. It is used in the market equilibrium, however, if price impact is sufficiently high because the owner of the efficient technology extracts too many rents. Through this mechanism, our model can give rise to a simultaneous decline in aggregate risky investment as well as misallocation. Interestingly, these outcomes are caused by the same friction that also leads to a decline in the risk-free rate. That is, the joint fall in interest rates and investment are not a puzzling confluence of events, but instead are driven by the same economic forces.

Consistent with our strategic q-theory formulation, Ma (2019) show that firms engage in strategic buybacks of their own debt and equity. Since price impact may inflate asset prices and allow for profits from strategic trading, moreover, firms can command high market values and Price-to-Earnings ratios despite the secular decline in investment since the early 2000s (Gutierrez and Philippon (2017b))).

4.4 A Numerical Example

We now provide a simple calibration exercise to illustrate the impact of market concentration on real and financial outcomes. To transparently describe model mechanisms, we assume the following ex-ante symmetric structure. All strategic agents have additive time-separable CRRA utility. The fringe has CRRA utility over date-2 consumption. There are two types of strategic agents, $i \in \{1,2\}$. The stochastic process for risky investment returns is composed of an aggregate shock $\theta \in \{l,h\}$ and a distributional shock $\delta \in \{1,2\}$. Given a mean aggregate return $\bar{Y}$, the aggregate return process satisfies

$$Y(\theta) = \begin{cases} (1+\Delta)\bar{Y} & \text{if } \theta = h \\ (1-\Delta)\bar{Y} & \text{if } \theta = l \end{cases} \qquad \text{where} \qquad 0 < \Delta < 1.$$

The idiosyncratic return process satisfies

$$y_i(\delta,\theta) = \begin{cases} (1+\alpha)Y(\theta) & \text{if } \delta = i \\ (1-\alpha)Y(\theta) & \text{if } \delta \neq i \end{cases} \qquad \text{where} \qquad 0 < \alpha < 1.$$

All states are equally likely, $\pi(\delta,\theta) = \frac{1}{4}$. This implies that asset prices are invariant to the distributional shock δ, and we can restrict attention to two prices, $q(h)$ and $q(l)$.

Table 1 summarizes the parameters of the model. We calibrate only the risk aversion parameter of the competitive fringe; all others are exogenously fixed. We choose the return process to be broadly in line with U.S. aggregate data. The mean aggregate growth rate is 2% per annum. We pick conservative numbers for aggregate and idiosyncratic volatility in order to illustrate the potential bite of the price impact mechanism. Aggregate volatility is 5%, while idiosyncratic volatility is 10%. This is approximately one quarter of the cross-sectional sales-growth volatility among Compustat firms. Hence we implicitly assume that agents can obtain some diversification without price impact. The return to storage is normalized to unity. Endowments are also normalizations. We exogenously fix the coefficient of relative risk aversion of strategic agents at 5, which again is conservative when compared to standard asset pricing calibrations. We then calibrate the coefficient of relative risk aversion of the competitive fringe to match a cross-sectional average price elasticity of 3 from Koijen and Yogo (2019). We report model results for all $\mu \in (0, 1]$. This allows us to illustrate counterfactual outcomes for various levels of market concentration.

Parameter	**Interpretation**	**Value**
N	Number of strategic types	2
$1/\mu$	Number of strategic agents of each type	5
$\bar{Y}$	Mean gross aggregate return	1.02
Δ	Aggregate volatility	0.05
α	Idiosyncratic Volatility	0.1
R	Gross return to storage	1
e	Endowment of Strategic Agents	2.5
γ_s	Relative risk aversion strategic agents	5
γ_f	Relative risk aversion fringe	2

Table 1: Parameter choices.

Figure 1 shows risk-free rates and the market risk premium. We plot both the market-implied risk-free rate (solid line) and the state-price implied risk-free rate (dashed line). The level of the real risk-free rate is broadly in date with the data, albeit slightly lower on average. The state-price implied risk-free rate declines monotonically until agents begin investing in storage (at around $\mu \approx 0.35$). The decline occurs because price impact hampers risk sharing and thus raises the volatility of the market-wide SDF. For μ sufficiently large, insurance is sufficiently expensive on financial markets and agents begin self-insuring using the inefficient storage condition. Since storage is not subject to price impact, a simple no-arbitrage condition now implies that the state-price implied risk-free rate must be exactly equal to the storage return.

The market-implied risk-free rate is non-monotonic: it declines if agents only invest in risky capital, but begins to rise when agents self-insure with storage. This is because an increase in storage flattens the marginal utility curve in the bad state (by raising incomes in the bad state at the expense of the high state), thereby causing a sharper strategic response. In fact, the gap between the two risk-free rates is accounted for only by the extent of market power exerted in equilibrium (see Proposition 7).

The "risk off" effect is also reflected in the market risk premium, which declines sharply as agents increase their investment in storage. Interestingly, however, we also observe a declining risk premium for small values of μ where no agent invests in storage. This is because sellers are particularly willing to distort states with high income. As a result, the price of the high-state $q(h)$ initially rises faster than that of the low-state claim $q(l)$. While risk compression is qualitatively in line with recent data (Bianchi, Lettau, and Ludvigson (2020)), our model is unable to match the average level of the risk premium. This is not surprising given the simplicity of our framework.

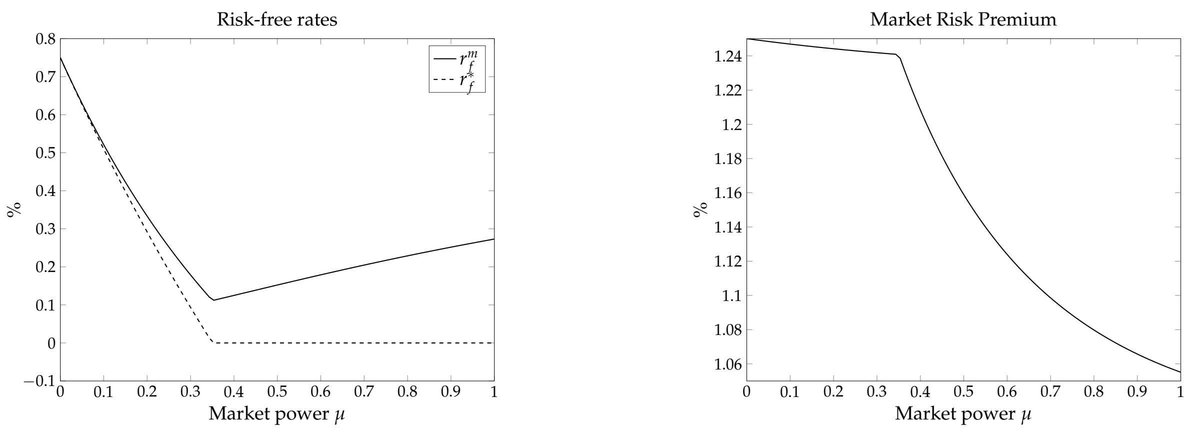

Figure 1: Equilibrium returns.

Figure 2 shows equilibrium investment policies and asset prices. At $\mu = 0$, strategic agents invest about half of their endowment in the risky asset, and consume the rest at date 1. As market power rises, risky investment falls because agents find it more costly to insure against idiosyncratic shocks. The fall is particularly sharp when agents begin investing in storage. Since $\mu = 0.2$ in the benchmark calibration, a doubling of market power to $\mu = 0.4$ leads to a decline in risky investment of about 15%.[6]

[6] Since there are no differences in expected returns across agents because they are ex-ante symmetric, there is no scope for misallocation in this example.

The right panel shows asset prices. Prices depend only on the aggregate state, and the price on the low-state claim is higher because it provides insurance. Both prices rise initially as strategic agents begin to ration insurance, holding the stock of risky investment fixed. The high price begins to fall once storage lowers the aggregate quantity of risk. We have already argued that the low-state price falls as a result of strategic considerations.

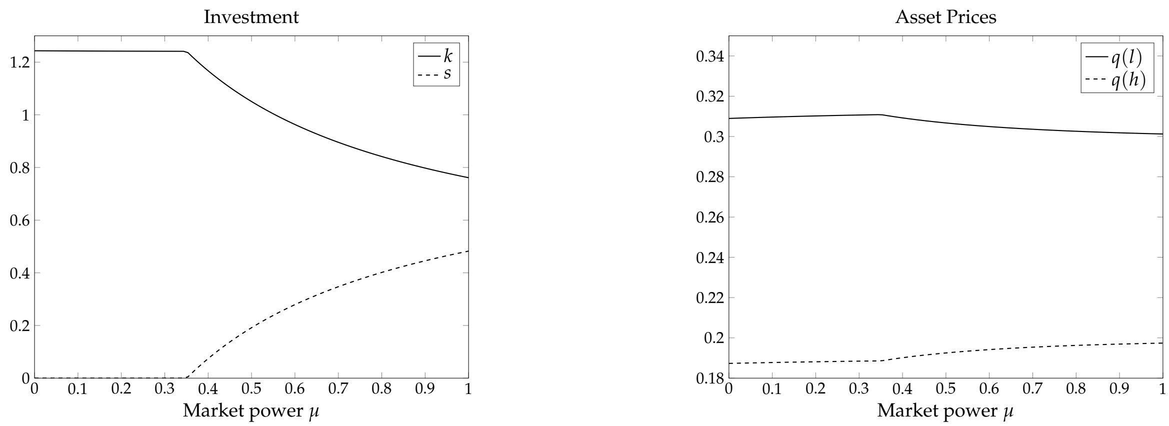

Figure 2: Investment and asset prices.

Figure 3 shows the demand absorbed by the competitive fringe. In the range where risky investment is approximately constant, the fringe is a supplier of insurance in all states. This partially offsets the insurance rationing by strategic agents. The right panel shows that increasing market power unequivocally lowers welfare for strategic agents. (This does not imply a Pareto ranking, as the competitive fringe may obtain higher expected utility when market power grows.)

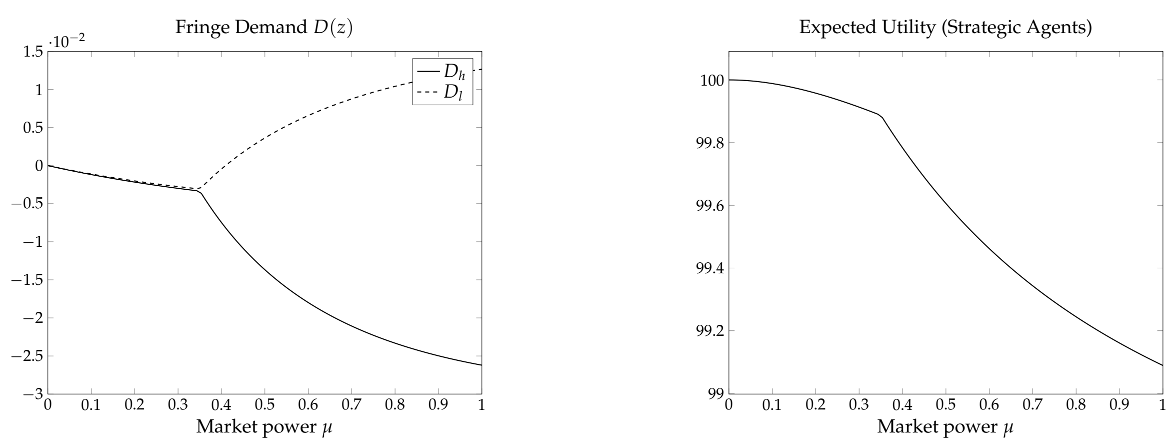

Figure 3: Demand of competitive fringe and strategic agents' expected utility (normalized to 100 at $\mu = 0$).

Taken together, the model gives rise to a sharp non-linearity in the relative response of real quantities and returns. For a small degree of price impact, strategic considerations are primarily reflected in a declining risk-free rate. However, agents are able to obtain enough insurance such that they continue to invest primarily in the risky technology. As a result, the distribution of income across states remains approximately constant, both asset prices are similarly distorted, and the response of the market risk premium is muted. This is no longer the case if price impact is large. Agents now self-insure using storage and there is a sharp decline in risky investment. While this pins down risk-free rates in a narrow band around R, the market risk premium now declines sharply.

More generally, the model predicts that the investment response co-moves strongly with the equity premium response, and particularly so at low interest risk-free rates. This may seem puzzling, since basic intuition would suggest that low risk-free rates and a low required return on risky investment encourage risky investment. Viewed through the lens of the model, however, all three observations are symptoms of the same strategic distortions.

4.5 Efficiency

We now discuss the welfare implications of our macroeconomy with strategic trading. In contrast to economies with missing markets or constraints on trading (i.e. collateral constraints, limited commitment) in which constrained efficiency depends on the pecuniary externalities of the particular market structure, the economy with price impact is generically constrained inefficient because the trading distortions in financial markets are self-inflicted by market participants.

Proposition 9 *The market equilibrium is generically constrained Pareto inefficient. To first-order, welfare losses arise from imperfect risk-sharing that can be ameliorated by taxing financial trading. To second-order, welfare losses are driven by inefficient investment which can be ameliorated by investment subsidies.*

The source of inefficiency is that some agents intentionally remain exposed to diversifiable risk for strategic reasons. The primary role for policy is to improve risk-sharing among the strategic agents given the capital and savings decisions, potentially through Pigouvian taxes or transfers. A secondary role, however, is to improve allocative efficiency. This can be accomplished by subsidizing the investment of agents who produce in scarcer states, as measured by buyers' willingness to pay for additional consumption.

4.6 Empirical Implications

Our framework has several clear empirical predictions for the aggregate implications of market power in financial markets. First, investment and investment efficiency both fall because of endogenous incompleteness induced by market power. Second, with regard to asset pricing, market power depresses the risk-free rate and can also compress realized risk premia, which lowers the market-implied Sharpe Ratio. Third, in the cross-section, wealth become more concentrated over time because there is less co-insurance. In the following, we discuss several additional testable predictions.

First, a necessary condition for our mechanism to have aggregate implications is that the strategic decisions of financial actors, such as institutional investors, have real effects. To test this, one could, for instance, examine how the corporate debt purchases by life insurance companies impact bond yields and investment decisions of issuing firms. Ellul, Jotikasthira, and Lundblad (2011) finds that insurance companies have price impact in corporate bond markets; Zhu (2020) and Chakraborty and MacKinlay (2020) provide evidence of the real effects of such supply-side pressures.

Second, market power in financial markets has business cycle implications. Its impact is most pronounced when the dispersion in firm productivity and the returns to intermediary lending are high, as during recessions (Bachmann and Bayer (2014)). As price impact impairs reallocation, our model predicts that downturns exacerbate the ex-post concentration of capital. In addition, as aggregate investment and investment efficiency falls, not only does the quantity of risk fall, but firm security valuations divorce further from their risk-return values because of the wedge in state prices because of price impact.

Our framework also provides predictions as to when market power is most distortive from a social perspective, and allows us to quantify this distortion. To construct one intuitive measure of this social cost in market $z \in \mathcal{Z}$ with price $q(z)$, $\Delta(z)$, let $\bar{A}_B(z) = \frac{1}{|\mathcal{B}|}\sum_{i\in\mathcal{B}} a_i(z) \geq 0$ be the average demand of buyers in market z. Similarly, let $\bar{A}_S(z) \leq 0$ be the average supply of sellers. The average unexploited gains from trade can be measured by state prices differences and are related to price impact by

$$\Delta(z) = \frac{1}{|\mathcal{B}|}\sum_{i\in\mathcal{B}} \Lambda_i(z) - \frac{1}{|\mathcal{S}|}\sum_{j\in\mathcal{S}} \Lambda_j(z) = \mu\gamma_f(z)\, q(z)\left(\bar{A}_B(z) + \left|\bar{A}_S(z)\right|\right) \geq 0,$$

where we recognize that $\frac{\partial q(z)}{\partial A(z)} = \gamma_f(z)\, q(z)$, and where $\gamma_f(z)$ is the fringe's coefficient of absolute risk aversion (CARA) in state z. The average gains from trade are the implied average bid-ask spread in market z, and *Kyle's Lambda* $\times$ *average trading volume* is a measure of the social value of closing this spread. Paradoxically, those that appear to be the deepest markets (highest volumes with lowest price discounts) are those in which the social value to closing the spread is the largest; by revealed preference, agents trade most in the markets in which there are the largest gains from trade.

If CARA is a reasonable approximation to the utility of competitive agents and price impact is similar across states, then $\Delta(z)$ is proportional to the average market value of all trades. Since the Law of one Price holds, equilibrium allocations are invariant to the specific structure of financial markets. As such, one can introduce a redundant security j, and use our measure to quantify the *aggregate* unrealized gains from trade in these markets. Of particular interest is risk-free debt, in which case

$$\Delta = \sum_{z \in \mathcal{Z}} \Delta(z) = \mu \gamma_f \frac{1}{r_f^m} \left(\bar{A}_B + \left| \bar{A}_S \right| \right).$$

The total average distortion from market power in financial markets, Δ, can then be measured as the *price of a risk-free bond* $\times$ *average trading volume* of strategic agents in, for instance, T-Bill markets.

5 Extensions and Discussion

5.1 Market Segmentation and Endogenous Arbitrage

In our baseline model, the Law of One Price holds because assets can be priced using the aggregate consumption process of the competitive fringe. The intuition is analogous to that of Arbitrage Pricing Theory (APT) in which a (small) number of deep-pocketed price-taking investors is sufficient for the Law of One Price to hold.[7]

We now show that this result no longer holds if there is market segmentation among the competitive fringe, by which we mean that the fringe is composed of multiple subgroups, each of which can only trade a strict subset of all assets. The Law of One Price may now fail because price impact leads strategic agents to refrain from closing all

[7] Similar to our work, Carvajal (2018) also studies a setting in which this intuition applies in the presence of large agents that internalize their price impact.

possible arbitrage opportunities because doing so would accrue inframarginal costs from price impact. We summarize these observations in the following proposition.

Proposition 10 *If the fringe is segmented across state-contingent security markets, then the Law of One Price generically fails if the fringe's members do not trade both the redundant asset and the replicating state contingent securities.*

Such a failure of the Law of One Price can potentially rationalize strict arbitrages observed in practice, such as violations of Covered Interest Parity documented over the last decade.[8] Arguably, the increased regulatory burden in the aftermath of the recent financial crisis has made it more difficult for intermediaries to participate in financial markets because of tighter balance sheet constraints and increased reporting requirements. While larger, more strategic banks can hedge their sizable volumes of currency exposures in futures markets, smaller, more competitive agents and more heavily regulated banks have more difficulty offloading their positions. As a consequence, the reduced participation of both strategic and competitive actors brings the currency forward markets closer to our setting with a segmented fringe. This is consistent with Wallen (2019).

5.2 A Market for Capital

To highlight the implications of impaired insurance in financial markets for allocative efficiency, we abstracted from issues of capital illiquidity by assuming it was in elastic supply. If capital were instead in limited supply (i.e. fixed), then impaired risk-sharing would impact each agent's demand for capital, which would feed back into the price of capital. We briefly discuss the implications of this pecuniary externality.

As we characterized in the case of an elastic market for capital, impaired risk-sharing from price impact leads to cross-sectional capital misallocation and to lower aggregate investment. This lowers the social marginal product of capital. With an inelastic capital market, however, lower demand for capital would, in turn, lower the price of capital. This lower price of capital represents an offsetting windfall for the most productive agents, which raises their investment. As such, an inelastic market for capital acts, in part, as a buffer against impaired risk-sharing for the more productive agents. A fluctuating

[8] Covered Interest Parity is the requirement that, since exchanging into a foreign currency, lending in that currency, and then exchanging out of it using a forward contract is a riskless transaction, it should earn the domestic riskfree rate.

market price of capital, however, also adds an additional dimension to the constrained inefficiency of the Nash Equilibrium in Demand Curves, since then markets are endogenously incomplete with two goods (e.g. Hart (1975)).

5.3 Stock Market Equilibrium

Our analysis presumes that control rights are not marketable. Instead, firms act as large entrepreneurs who make production decisions based on their own preferences and share risks with each other. Risk-sharing is distinct from the sharing of control rights, such as in the stock market equilibria of Dréze (1974). In these (incomplete markets) settings, consumers buy shares of equity in firms whose objective is to maximize shareholder value according to some aggregation rule for shareholder preferences.

If all agents in our setting were perfectly competitive, the equilibrium with complete markets would coincide (in allocation) with a stock market equilibria because financial markets span all firms' cash flow risks. With market power, however, strategic agents endogenously leave markets incomplete, which prevents full aligning of their state prices. While partial alignment does bring them closer to agreeing on the socially optimal investment decisions of any agent in the economy, as with incomplete markets, they would still generically disagree with each other's privately optimal investment decisions.

By having agents trade risk rather than control rights, we side-step the unresolved question of how to appropriately define the objective in effectively incomplete markets. Our analysis, however, reveals that complete markets is insufficient for shareholders to agree on the production plan of a firm in the presence of price impact.

Limits to the trading of control rights may obtain for other reasons. For example, there may be frictions in the issuance of equity, as is commonly assumed in the literature on financial intermediation. In this case, the Lagrange multiplier on the equity issuance constraint would act as a wedge that distorts decisions in favor of the inside equity holder, and the objective of the firm may then be closer to that in Grosssman and Hart (1979). This would mean that the de-facto objective function of the firm is similar to the one we have assumed even if control rights were marketable.[9]

[9] We thank Gian Luca Clementi for pointing this out to us.

6 Conclusion

We construct a model of concentrated financial markets in which large, risk-averse agents invest in risky projects and internalize their price impact when trading state contingent claims with each other. As a result of strategic interactions, agents voluntarily under-insure production risk. Such under-insurance adversely impacts capital investment decisions and, despite market completeness, generates dispersion in state prices. In line with recent trends, strategic market incompleteness can lead to a joint decline in the risk-free rate, risk premia, investment, and productivity. Our model is tractable and may be useful for additional applications. We also offer several testable implications.

References

ALEXANDER, L., AND J. EBERLY (2018): "Investment Hollowing Out," *IMF Economic Review*, 6, 5–30.

ALVAREZ, F., AND U. JERMANN (2000): "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68(4), 775–797.

——— (2001): "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints," *Review of Financial Studes*, 14(4), 1117–1151.

BACHMANN, R., AND C. BAYER (2014): "Investment Dispersion and the Business Cycle," *American Economic Review*, 104(4), 1392–1416.

BIANCHI, F., M. LETTAU, AND S. C. LUDVIGSON (2020): "Monetary Policy and Asset Valuation," *mimeo Duke, UC Berkeley, and NYU*.

CAMPBELL, J. Y., M. LETTAU, B. G. MALKIEL, AND Y. XU (2001): "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, 56(1), 1–43.

CARVAJAL, A. (2018): "Arbitrage pricing in non-Walrasian financial markets," *Economic Theory*.

CARVAJAL, A., AND M. WERETKA (2012): "No-arbitrage, state prices and trade in thin financial markets," *Economic Theory*, 50(1), 223–268.

CHAKRABORTY, I., AND A. MACKINLAY (2020): "The Effects of Financing Frictions in Investment-Grade Debt Markets," *mimeo University of Miami*.

CORBAE, D., AND R. LEVINE (2018): "Competition, Stability, and Efficiency in Financial Markets," *Working Paper*.

COVARRUBIAS, MATIAS, G. G. P. T. (2019): "From Good to Bad Concentration? US Industries over the Past 30 Years," *NBER Macroeconomics Annual*, 34.

DE LOECKER, J., J. EECKOUT, AND G. UNGER (2018): "The Rise of Market Power and the Macroeconomic Implications," *mimeo*.

DEL NEGRO, MARCO, G. D. G. M. P., AND A. TAMBALOTTI (2018): "Global Trends in Interest Rates," *FRBNY Staff Report*, 866.

DRECHSLER, I., A. SAVOV, AND P. SCHNABL (2018): "The Deposits Channel of Monetary Policy," *Quarterly Journal of Economics*.

DRÉZE, J. (1974): "Investment under Private Ownership: Optimality, Equilibrium and Stability," in *Allocation Under Uncertainty*. New York:Wiley.

EISENBACH, T., AND G. PHELAN (2020): "Cournot Fire Sales in Real and Financial Markets," *FRBNY Staff Report*, 837.

ELLUL, A., C. JOTIKASTHIRA, AND C. T. LUNDBLAD (2011): "Regulatory Pressure and Fire Sales in the Corporate Bond Market," *Journal of Financial Economics*, 101(3), 596–620.

FERNALD, JOHN G, H. R. E. S. J. H., AND M. W. WATSON (2017): "The Disappointing Recovery of Output After 2009," *Technical report, Brookings Papers on Economic Activity*.

GLEBKIN, S., S. MALAMUD, AND A. TEGUIA (2020): "Asset Prices with Market Power and Non-Gaussian Payoffs," *mimeo INSEAD, EPFL, and UBC Sauder*.

GOLOSOV, MIKHAIL, K. N., AND A. TSYVINSKI (2003): "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 70(3), 569–587.

GROSSSMAN, S. J., AND O. D. HART (1979): "A Theory of Competitive Equilibrium in Stock Market Economies," *Econometrica*, 47(2), 293– 330.

GUTIERREZ, G., AND T. PHILIPPON (2017a): "Declining Competition and Investment in the US," *Working Paper*.

——— (2017b): "Investmentless Growth: An Empirical Investigation," *Brookings Papers on Economic Activity*, pp. 89–169.

HART, O. (1975): "On the Optimality of Equilibrium When the Market Structure is Incomplete," *Journal of Economic Theory*.

HAU, H., P. HOFFMANN, S. LANGFIELD, AND Y. TIMMER (2017): "Discriminatory Pricing of Over-the-Counter Derivatives," *Swiss Finance Institute Research Paper No. 17-70*.

JONES, C., AND T. PHILIPPON (2016): "The Secular Stagnation of Investment?," *Working Paper*.

KEHOE, T. J., AND D. K. LEVINE (1993): "Debt-Constrained Asset Markets," *The Review of Economic Studies*, 60(4), 865–888.

KOCHERLAKOTA, N. R. (1996): "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, 34, 42–71.

KOCHERLAKOTA, N. R., AND L. PISTAFERRI (2009): "Asset Pricing Implications of Pareto Optimality with Private Information," *Journal of Political Economy*, 117(3), 555–590.

KOIJEN, R. S., AND M. YOGO (2019): "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, 127, 1475–1515.

KYLE, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 54(6), 1315–1335.

——— (1989): "Informed Speculation with Imperfect Competition," *The Review of Economic Studies*, 56(3), 317–355.

LAUBACH, T., AND J. WILLIAMS (2016): "Measuring the Natural Rate of Interest Redux," *Business Economics*, 51(2), 57–67.

MA, Y. (2019): "Nonfinancial Firms as Cross-Market Arbitrageurs," *Journal of Finance*, 74(6), 3041–3087.

MALAMUD, S., AND M. ROSTEK (2017): "Decentralized Exchange," *American Economic Review*, 107(11), 3320–3362.

PRITSKER, M. (2005): "Large Investors: Implications for Large Investors: Implications for Equilibrium Asset Returns, Shock Absorption, and Liquidity," *Finance and Economics Discussion Series (FEDS) Working Paper No. 2005-36*.

ROSTEK, M., AND M. WERETKA (2015): "Dynamic Thin Markets," *The Review of Financial Studies*, 28(10), 2946–2992.

WALLEN, J. (2019): "Markups to Financial Intermediation in Foreign Exchange Markets," *mimeo Stanford*.

WANG, Y., T. M. WHITED, Y. WU, AND K. XIAO (2018): "Bank Market Power and Monetary Policy Transmission: Evidence from a Structural Estimation," *Working Paper*.

ZHU, Q. (2020): "Capital Supply and Corporate Bond Issuances: Evidence from Mutual Fund Flows," *Journal of Financial Economics*, forthcoming.

A Proofs of Propositions

A.1 Proof of Proposition 1:

Given a binding budget constraint and quasi-linear preferences, the first-order condition for $D(z)$ immediately implies the pricing equation

$$q(z) = \pi(z)u_f'\left(c_{f2}(z)\right) = \Lambda_f\left(z\right),$$

where $\Lambda^f\left(z\right)$ is the competitive fringe's state price. Using the market-clearing condition gives $c_{f2}(z) = w_2(z) - A(z)$. Since $u_f\left(z\right)$ is twice continuously differentiable and each agent's position size scales by its mass μ_i, it is then straightforward to verify that price impact is given by:

$$\frac{\partial q(z)}{\partial a_i\left(z\right)} = -\mu_i \pi(z) u_f''\left(c_{f2}(z)\right) = -\mu_i \frac{\partial q(z)}{\partial A\left(z\right)},$$

which also implies that price impact is symmetric across all strategic agents. Alternatively, we can express $\frac{\partial q(z)}{\partial A}\left(z\right)$ as:

$$\frac{\partial q(z)}{\partial A}\left(z\right) = \gamma_f\left(z\right) q\left(z\right),$$

where $\gamma_f\left(z\right)$ is the coefficient of absolute risk aversion of the competitive fringe.

We now prove that the introduction of redundant assets does not change real allocations. To this end, define asset j by its dividend process $d_j(z)$, and let q_j denote its asset prices. By the FOC from the competitive fringe's optimization program:

$$q_j = \sum_z \Lambda^f\left(z\right) d_j\left(z\right)\left(z\right),$$

where $\Lambda^f\left(z\right) = u_f'\left(c_2\left(z\right)\right)$ is the state price of the fringe. For the Arrow-Debreu securities, $d_j\left(z\right) = \mathbf{1}_{\{z'=z\}}$. It then follows that Arrow-Debreu prices in the economy satisfy martingale pricing with $\Lambda^f\left(z\right)$ as the appropriate state price deflator.

Aggregating the FOCs of strategic agents (see Proposition 2), we also have that:

$$q_j = \frac{1}{N}\sum_n \Lambda_n\left(z\right) d_j\left(z\right) + \frac{1}{N}\frac{\partial q_j}{\partial \hat{A}_j}\hat{A}_j.$$

Consider now the introduction of a redundant asset ϕ with redundant payoff $d_{\phi}(z) = \sum_{z \in \mathcal{Z}} \phi(z)\, d(z)$. Let the new asset allocations not have hats. Then:

$$q_{\phi} = \sum_{z \in \mathcal{Z}} d_{\phi}(z)\, \Lambda^{f}(z) = \sum_{z \in \mathcal{Z}} \phi(z)\, 1_{\{z'=z\}} \Lambda^{f}(z) = \sum_{z \in \mathcal{Z}} \phi(z)\, q(z),$$

by Law of One Price, since fringe satisfies martingale pricing. Since $q_{\phi} = \sum_{z \in \mathcal{Z}} \phi(z)\, q(z)$, it must then be the case by equating the aggregated FOCs of the strategic agents for the redundant security and the replicating Arrow-Debreu securities that:

$$\sum_{n} \Lambda_{n}(z) \sum_{z \in \mathcal{Z}} \phi(z)\, \mathbf{1}_{\{z'=z\}} + \frac{\partial q_{\phi}}{\partial A_{\phi}} A_{\phi} = \sum_{z \in \mathcal{Z}} \phi(z) \left(\sum_{n} \Lambda_{n}(z)\, \mathbf{1}_{\{z'=z\}} + \frac{\partial q(z)}{\partial A(z)} A(z) \right),$$

from which follows, by cancellation of the state price terms that:

$$\frac{\partial q_{\phi}}{\partial A_{\phi}} A_{\phi} = \sum_{z \in \mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \phi(z)\, A(z).$$

This immediately implies that there must be trade in the redundant asset if there is trade in the replicating assets. We recognize that for the Arrow-Debreu securities:

$$\frac{\partial q(z)}{\partial A(z)} = \frac{\partial \Lambda^{f}(z)}{\partial A(z)} = -u_{f}''(c_{2}(z)),$$

and, by linearity, for the redundant security:

$$\frac{\partial q_{\phi}}{\partial A_{\phi}} = -\sum_{z \in \mathcal{Z}} \phi(z)\, u_{f}''(c_{2}(z)) = \sum_{z \in \mathcal{Z}} \phi(z) \frac{\partial q(z)}{\partial A(z)}.$$

Notice this argument does not rely on the quasi-linearity of competitive fringe, since if they instead had preferences $u_{f1}(\cdot)$ over initial consumption, we would have that:

$$\frac{\partial q(z)}{\partial A(z)} = -\left(\frac{u_{f}''(c_{2}(z))}{u_{f}'(c_{2}(z))} - \frac{u_{f1}''(c_{1})}{u_{f1}'(c_{1})} \frac{\partial c_{f1}}{\partial A(z)} \right) q(z),$$

where:

$$\frac{\partial c_{f1}}{\partial A(z)} = q(z) + \frac{\partial q(z)}{\partial A(z)} A(z) + \frac{\partial q_{\phi}}{\partial A(z)} A_{\phi}.$$

and again:

$$\frac{\partial q_\phi}{\partial A_\phi} = -\sum_{z\in\mathcal{Z}} \phi\left(z\right) \frac{u_f''\left(c_2\left(z\right)\right)}{u_f'\left(c_2\left(z\right)\right)} q\left(z\right) - \frac{u''\left(c_{f1}\right)}{u'\left(c_{f1}\right)} \frac{\partial c_{f1}}{\partial A_\phi},$$

where:

$$\frac{\partial c_{f1}}{\partial A_\phi} = q_\phi + \frac{\partial q\left(z\right)}{\partial A_\phi} A\left(z\right) + \frac{\partial q_\phi}{\partial A_\phi} A_\phi.$$

Notice that $\frac{\partial c_{f1}}{\partial A_\phi} = \sum_{z\in\mathcal{Z}} \phi\left(z\right) \frac{\partial c_{f1}}{\partial A(z)}$, or else the fringe could make a riskless profit from buying the cheaper of the redundant security or the replicating bundle and selling the other. Such a transaction would leave consumption at date 2 unchanged by matching cash flows, and raise consumption at date 1, which would raise the competitive fringe's utility. This would be a violation of the optimality of the fringe's asset allocation decisions, a contradiction. Therefore, again $\frac{\partial q_\phi}{\partial A_\phi} = \sum_{z\in\mathcal{Z}} \phi\left(z\right) \frac{\partial q(z)}{\partial A(z)}$.

Consequently, what is essential is that the competitive fringe takes prices as given, and consequently ensures no arbitrage across traded assets by the Law of One Price.

Now let us conjecture that the introduction of the redundant asset has no real effects on allocations. If there are no real effects, then it must be the case that $\hat{A}\left(z\right) = A\left(z\right) + \phi\left(z\right) A_\phi$. Since allocations for the fringe are the same both with and without the redundant security, it follows that consumption, c_{f1} and $c_{f2}\left(z\right)$, are the same in both economies. Consequently, state prices $\Lambda^f\left(z\right)$ must be the same in both economies and therefore so are prices since $q_j = \sum_{z\in\mathcal{Z}} \Lambda^f\left(z\right) d_j\left(z\right)$. Substituting the requirement that $\hat{A}\left(z\right) = A\left(z\right) + \phi\left(z\right) A_\phi$ into $\frac{\partial q_\phi}{\partial A_\phi} A_\phi = \sum_{z\in\mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \phi\left(z\right) A\left(z\right)$, and substituting our expression for price impact $\frac{\partial q_\phi}{\partial A_\phi}$, one has that:

$$\sum_{z\in\mathcal{Z}} \phi\left(z\right) \frac{\partial q\left(z\right)}{\partial A\left(z\right)} A_\phi = \sum_{z\in\mathcal{Z}} \frac{\partial q\left(z\right)}{\partial A\left(z\right)} \phi\left(z\right) \left(\hat{A}\left(z\right) - \phi\left(z\right) A_\phi\right),$$

from which follows that the above holds if:

$$A_\phi = \frac{\sum_{z\in\mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \phi\left(z\right)}{\sum_{z\in\mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \left(\phi\left(z\right)^2 + \phi\left(z\right)\right)} \hat{A}\left(z\right),$$

and, substituting this into $\hat{A}(z) = A(z) + \phi(z) A_\phi$, we arrive at:

$$A(z) = \left(1 - \phi(z) \frac{\sum_{z \in \mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \phi(z)}{\sum_{z \in \mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \left(\phi(z)^2 + \phi(z)\right)}\right) \hat{A}(z) = f(z) \hat{A}(z).$$

To complete the argument, we recognize that, since allocations are the same, we can express the price in both economies as:

$$\begin{aligned} q(z) &= \frac{1}{N} \sum_n \Lambda_n(z) \mathbf{1}_{\{z'=z\}} + \frac{1}{N} \frac{\partial q(z)}{\partial \hat{A}(z)} \hat{A}(z), \\ q(z) &= \frac{1}{N} \sum_n \Lambda_n(z) \mathbf{1}_{\{z'=z\}} + \frac{1}{N} \frac{\partial q(z)}{\partial A(z)} A(z). \end{aligned}$$

Equating these two representations, we must last verify the consistency condition that:

$$\frac{\partial q(z)}{\partial A(z)} A(z) = \frac{\partial q(z)}{\partial \hat{A}(z)} \hat{A}(z).$$

Notice that consumption for the fringe with and without the redundant assets is the same, which implies that:

$$\begin{aligned} c_{f1} &= w_1 + \sum_{z \in \mathcal{Z}} q(z) A(z) + q_\phi A_\phi = w_1 + \sum_{z \in \mathcal{Z}} q(z) \left(A(z) + \phi(z) A_\phi\right) = w_1 + \sum_{z \in \mathcal{Z}} q(z) \hat{A}(z), \\ c_{f2}(z) &= w_2(z) - \sum_{z \in \mathcal{Z}} \mathbf{1}_{\{z'=z\}} A(z) - d_\phi(z) A_\phi = w_2(z) - \sum_{z \in \mathcal{Z}} \mathbf{1}_{\{z'=z\}} \left(A(z) + \phi(z) A_\phi\right) \\ &= w_2(z) - \sum_{z \in \mathcal{Z}} \mathbf{1}_{\{z'=z\}} \hat{A}(z). \end{aligned}$$

Matching cash flows, it follows that $\frac{\partial q(z)}{\partial A(z)} = \frac{1}{f(z)} \frac{\partial q(z)}{\partial \hat{A}(z)}$, it follows that:

$$\frac{\partial q(z)}{\partial A(z)} A(z) = \frac{1}{f(z)} \frac{\partial q(z)}{\partial \hat{A}(z)} f(z) \hat{A}(z) = \frac{\partial q(z)}{\partial \hat{A}(z)} \hat{A}(z),$$

as required. Consequently, the consistency condition holds, which confirms our conjecture that redundant assets do not have real effects.

A.2 Proof of Proposition 2

Since no agents default in equilibrium, the shadow value of an increase in debt issuance by agent i today is absorbed by the appropriate Lagrange multiplier, which increases the asset price despite the lack of change in agent $j's$ marginal utility by Kuhn-Tucker. FONCs are then valid. Normalize the Lagrange multiplier on the budget constraint to be φ_i to be $\varphi_i e_i^{\gamma}$. The FONCs for c_{i1}, k_i, s_i, and $a_i(z)$ in the strategic agent's problem are given by:

$$\begin{aligned}
c_{i1} &: u'(c_{i1}) - \varphi_i e_i^{\gamma} \leq 0 \; (= \; if \; c_{i1} > 0), \\
k_i &: \sum_{z \in \mathcal{Z}} \pi(z) u_2'(c_{i2}(z)) y_i(z) - \varphi_i e_i^{\gamma} \leq 0 \; (= \; if \; k_i > 0), \\
s_i &: \sum_{z \in \mathcal{Z}} \pi(z) u_2'(c_{i2}(z)) R - \varphi_i e_i^{\gamma} \leq 0 \; (= \; if \; s_i > 0), \\
a_i(z) &: -\pi(z) u_2'(c_{i2}(z)) + \varphi_i e_i^{\gamma} \left(q(z) + \frac{\partial q(z)}{\partial a_i(z)} a_i(z) \right) = 0 \; (a_i(z) > \bar{a}_i(z)).
\end{aligned}$$

Therefore, the above represents the FONCs for agent $i's$ problem. Since $u_1(\cdot)$ satisfies the Inada condition, $c_{i1} > 0$ and the first FOC binds with equality.

Define $h_i = \frac{e_i}{\sum_{i=1}^{N} e_i}$ to be the effective share of scaled equity of agent i. Let us conjecture that each agent's optimal policies are such that, for policy s_i, we can decompose s_i into $s_i = \hat{s}_i e_i$, where $\hat{s}_i$ is independent of e_i and the level of equity of any other agent. Imposing market-clearing implies

$$q(z) = \pi(z) u_f' \left(w_2(z) - \left(\sum_{i=1}^{N} e_i \right) \sum_{i=1}^{N} \hat{a}_i(z) h_i \right).$$

Since utility is homogeneous in degree e_i^{γ}, we can then rewrite the FONCs as:

$$\begin{aligned}
\hat{c}_{i1} &: u_1'(\hat{c}_{i1}) - \varphi_i = 0, \\
\hat{k}_i &: \sum_{z \in \mathcal{Z}} \pi(z) \frac{u_2'(\hat{c}_{i2}(z))}{u_1'(\hat{c}_{i1})} y_i(z) \leq 1 \; \left(= \; if \; \hat{k}_i > 0\right), \\
\hat{s}_i &: \sum_{z \in \mathcal{Z}} \pi(z) \frac{u_2'(\hat{c}_{i2}(z))}{u_1'(\hat{c}_{i1})} R \leq 1 \; (= \; if \; \hat{s}_i > 0), \\
\hat{a}_i(z) &: -\pi(z) \frac{u_2'(\hat{c}_{i2}(z))}{u_1'(\hat{c}_{i1})} + q(z) - \frac{\partial q(z)}{\partial \hat{a}_i(z)} \hat{a}_i(z) = 0 \; \left(\hat{a}_i(z) > \widehat{\bar{a}}_i e_i\right),
\end{aligned}$$

and by definition of the budget constraint at $t = 2$:

$$\hat{c}_i(z) = y_i(z) \hat{k}_i + \hat{a}_i(z) + R\hat{s}_i.$$

The budget constraint at $t = 1$ can be expressed as

$$\hat{c}_{i1} + \hat{k}_i + \sum_{z \in \mathcal{Z}} q(z)\,\hat{a}_i(z) + \hat{s}_i = 1.$$

The rewritten FONCs are independent of $e_i \ \forall\ i \in \{1, ..., N\}$, and, defining $\Lambda_i(z) = \pi(z) \frac{u_2'(\hat{c}_{i2}(z))}{u_1'(\hat{c}_{i1})}$ to be the effective state price of agent i in state z, confirm our conjecture.

It then follows that the indirect utility of agent i is

$$U_1^i(e_i, \vec{e}_{-i}) = e_i^{1-\gamma} u_1(\hat{c}_{i1}) + e_i^{1-\gamma} \sum_{z \in \mathcal{Z}} \pi(z)\, u_2(\hat{c}_{i2}(z)) = \tilde{U}_1^i\left(h_i, \vec{h}_{-i}\right) e_i^{1-\gamma}.$$

To see that all the controls, normalized by the equity of agent i e_i are bounded, we consider the cases in which agent i does and does not issue state-contingent claims. When agent i does not issue any state-contingent claims, then by the balance-sheet constraint, it can only self-finance its uses of funds, and therefore $\hat{c}_{i1}, -\hat{a}_i(z), \hat{k}_i, \hat{a}_j(z) \ \le 1$. When agent i issues state-contingent claims, then the total normalized resources of all agents in the network is $\sum_{i=1}^N e_i$ plus borrowing from the competitive fringe.

By assumption, the credit that can be extracted from the competitive fringe in any state, $q(z)\,D(z)$, is bounded from above since $q(z) = u_f'\left(c_{2f}(z)\right)$ and $u_f'(x)\,x$ is bounded for $x > 0$, and $c_{2f}(z) > 0$.[10] It then follows that:

$$-q(z)\,\hat{a}_i(z) \le \hat{e} = \frac{1}{e_i}\left(\sum_{j \neq i} e_j + \max_{D(z)} q(z)\,D(z)\right).$$

Since $q(z)$ is bounded from below by assumption, it follows that $\hat{a}_i(z)$ is bounded. Consequently, normalized savings, capital, and consumption $\hat{s}_i, \hat{k}_i$, and $\hat{c}_{i1}$, are all bounded from above. They are also all trivially bounded from below because of the nonnegativity constraints on investment and consumption, and because total asset purchases by agent i cannot exceed e_i. Since all controls lie in a closed and bounded set, they lie in a compact set by the Heine-Borel Theorem.

Define $\mathcal{X}_i$ to be the state space for agent i $\mathcal{X}_i = \mathbb{R}_+^N \times \Sigma(-i)$, where $\Sigma(-i)$ is the space of strategies of the other strategic agents. Notice that, given the homotheticity of agent utility, it follows that $U_0^i(e_i, \vec{e}_{-i}) = U_0^i\left(h_i, \vec{h}_{-i}\right) e_i^{1-\gamma}$, where $h_j \in [0, 1]\ \forall\ j \in$

[10] In the case of power utility with CRRA index > 1, the Inada condition guarantees that $c_{f1} > 0$, in equilibrium, and therefore that the condition holds because of declining marginal utility.

$\{1, ..., N\}$, and the range of the controls $\hat{k}_i$, $\hat{s}_i$, and $\hat{a}_i(z)\, \forall z \in \mathcal{Z}$ lies in a compact, convex set. As such, the set of admissible controls normalized by equity $\mathcal{M} \subseteq \mathbb{R}^{2+|\mathcal{Z}|}$ is compact and convex, and let $g_i \in \mathcal{M}$ is a $|\mathcal{Z}|+2$ tuple $g_i = \begin{bmatrix} \hat{s}_{i1} & \hat{k}_i & \hat{a}_i(z_1) & \dots & \hat{a}_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix}'$, which maps to the controls in the original problem as $g_i e_i$. Then $G_i : \mathcal{X}_i \to \mathcal{M}$ is a compact-valued correspondence.

Since the objective $e_i^{1-\gamma} u_1^i(\hat{c}_{i1}) + e_i^{1-\gamma} \sum_{z \in \mathcal{Z}} \pi(z)\, u_2^i(\hat{c}_{i2}(z))$ is continuous in the states, and consequently in the controls, and the correspondence G_i is compact-valued, then by Berges' Theory of the Maximum there exists a solution to the agent's problem and the optimal G_i^* is an upper-hemicontinuous correspondence. It is the homogeneity of the optimal controls with respect to the equity of agent i e_i that allows us to express the space of controls as a compact-valued correspondence to satisfy the theorem if initial equity is allowed to be potentially unbounded.

Given that G_i^* is an upper-hemicontinuous correspondence $\forall\, i \in \{1, ..., N\}$, applying Kakutani's Fixed Point Theorem on the market-clearing conditions for agent debt establishes the existence of a fixed point that is an equilibrium.

A.3 Proof of Proposition 3

For the only if part, if the competitive and strategic equilibria coincide, then $q(z)$ and $\Lambda_i(z)$ must coincide. From the FONCs for asset positions, $a_i^{CE}(z)$, this can only be the case if either price impact is zero ($\partial q(z) / \partial a_i(z)$ or, more generically, when $a_i^{CE}(z) = 0$ in the competitive equilibrium. When $a_i^{CE}(z) = 0$, however, then there are no gains from trade, and the only equilibrium is autarky. For the if part, if there are no gains from trade in the competitive equilibrium, then $a_i^{CE}(z) = 0$ and the competitive and strategic equilibria trivially coincide.

It is trivial from the first-order conditions for $a_i(z)$ that sellers have an incentive to raise $a_i(z)$ while buyers have an incentive to lower. By definition of state prices, this immediately leads to state price dispersion starting from the competitive equilibrium baseline without any such dispersion.

Now consider a competitive equilibrium and fix capital allocation decisions. Suppose now that we alter agents' production technologies to reduce productivity in states in which agents overlap in production and increase it in states in which they do not, such that output in each state remains unchanged. Let the new productivities be indexed by

$\{\tilde{y}_i(z)\}_{i=1}^N$. For instance, with two agents and two states of production, we can shift productivity so that agent i now produces all output in state 1, $y_i(1) + \frac{k_j}{k_i} y_j(1)$, and agent j produces all output in states 2, $y_j(2) + \frac{k_i}{k_j} y_i(2)$.

Since agents can insure each other against states in which they differ in production compared to states in which they jointly produce, gains from trade increase in the economy. This is reflected by the lower correlation in state prices for the same asset positions since agents could, in principle, trade until they have the same relative marginal valuations of consumption across states (same state prices). With perfect competition, agents would trade until marginal utilities or state prices are equalized, and consequently the redistribution of productivity is irrelevant to the consumption allocation. The trading volume in their asset positions for claims in state z is then $\sum_{i=1}^N |\tilde{a}_i(z) - a_i(z)|$, where $\tilde{a}_i(z) = c_i(z) - \tilde{y}_i(z) k_i - Rs_i$ and $c_i(z)$ is agent $i's$ consumption in the competitive equilibrium.

Suppose instead we consider the same perturbation from the same capital and asset allocation decisions, but instead agents are strategic. Since agents now internalize their price impact, however, they ration their asset demands and supplies to extract rents ($\frac{\partial q(z)}{\partial A(z)} \Delta a_i(z)$ for a change in position of $\Delta a_i(z)$) from manipulating prices. As such, total trading volume is bounded from above by $\sum_{z \in \mathcal{Z}} \sum_{i=1}^N |\tilde{a}_i(z) - a_i(z)|$, which are positive related to total gains from trade. Since agents can always choose to trade fully their differences in risk exposures, the distortions must be (weakly) larger with this increase in pure firm-specific risk.

A similar argument applies if, instead of changing agent production technologies, we change their risk aversions. Increasing the risk aversion of a subset of agents and increasing it for others introduces gains from trade to shift aggregate risk from the more risk averse to the more risk tolerant. Similar arguments then establish that distortions from market power are (weakly) larger from this shift.

Lastly consider distortions in state prices. Since there are no cross-holdings, suppose that agent i buys claims from agent j in state z. Then FONCs imply

$$\Lambda_i(z) - \frac{\partial q(z)}{\partial a_i(z)} a_i(z) = q(z) = \Lambda_j(z) - \frac{\partial q(z)}{\partial a_j(z)} a_j(z),$$

from which follows that:

$$\Lambda_i(z) - \Lambda_j(z) = -\frac{\partial q(z)}{\partial a_j(z)} a_j(z) + \frac{\partial q(z)}{\partial a_i(z)} a_i(z) > 0,$$

and $\Lambda_i(z) > \Lambda_j(z)$. Consequently risk-sharing is imperfect.

A.4 Proof of Proposition 4

Consider the Nash Equilibrium in demand schedules allocation in the economy with imperfect competition, $\left\{ \left(c_{i1}, k_i, s_i, \{c_i(z)\}_{z \in \mathcal{Z}} \right) \right\}_{i=1}^{N}$. Define

$$\Lambda_i(z) = q(z) + \frac{\partial q(z)}{\partial a_i(z)} a_i(z) \ \forall \ (i, z),$$

to be the implied state price deflator of agent i in state z. Notice the implied state price deflator can include a collateral premium.

Now consider a fictitious incomplete markets economy in which all agents are instead competitive. We will try to derive the implied market structure, indexed by a set of M securities with arbitrary $M \times Z$ payoff matrix D. Since there are Z possible states of the world, it follows that we need at most Z linearly independent securities for markets to be complete, and consequently $M \leq Z$. If markets are incomplete, then these M securities are rank deficient, and consequently specifying at most Z potentially tradeable securities is without loss. Since assets cannot distinguish between some (subset of) states, the market-implied asset span spans M (linear combinations of the Z) states.

Let the price vector for these M assets be $\vec{p}$ and define d to be the dividend yield matrix with elements d_{ij}, where d_{ij} is the payoff of the i^{th} security in state j, and $q()$ is the price of security i. We follow the convention that, since by no arbitrage $p_i = 0$ iff $d_{ij} = 0$ $\forall\, j \in \{1, ..., M\}$, that $d_{ij} = 0$ in this contingency. No arbitrage then requires that

$$\mathcal{M} d = \iota_N \iota_M^T,$$

where $\mathcal{M}$ is the $N \times Z$ matrix with elements $\mathcal{M}_{ij} = \Lambda_i(j)$ for intermediary i and state j, and ι_M is the $M \times 1$ vector of one. For fixed M, this equation can be rewritten as a matrix equation:

$$(I_M \otimes \mathcal{M})\, vec(d) = vec\left(\iota_N \iota_M^T\right).$$

Since $rank\left(I_M \otimes \mathcal{M}\right) = M + rank\left(\mathcal{M}\right)$, it follows that there is a unique solution for d if $N \geq Z$, while the system is under-identified if $N < Z$.

As the synthetic assets are potentially derivatives, d can have negative entries. We then choose the largest M such that the recovered d has rank M. We choose the largest M since, if there were an additional asset that replicated the asset span that we ignored, then introducing this additional asset would not initiate trade, since it would already be priced at its no arbitrage value by all intermediaries. Consequently, $M \in \max_{m\in\{0,\ldots,K\}}\left\{rank\left(d\right) = m\right\}$.

With this payoff matrix, we have constructed an incomplete market structure that measures the degree of market incompleteness by replicating the effective asset span of the Cournot-Nash Equilibrium with complete markets. Since at least one state price is misaligned among intermediaries, the rank of d must be less than Z, and markets must be incomplete.

Finally, we recognize that the state prices of agent i and j are related by:

$$\Lambda_i\left(z\right) = \Lambda_j\left(z\right) + \frac{\partial q\left(z\right)}{\partial a_i\left(z\right)}\left(a_j\left(z\right) - a_i\left(z\right)\right).$$

Substituting for $\Lambda_j\left(z\right)$ for each $j \neq i$, we can rewrite the condition:

$$\mathcal{M}d = \iota_N \iota_M^T,$$

as $d^T\vec{\Lambda}_i = \iota_M$ and:

$$d^T\left(\vec{q'}_i \odot \left(\vec{a}_j - \vec{a}_j\right)\right) = 0_{M\times 1}\forall j,$$

where $\vec{q'}_i = \left[\mu_i q'(1), \mu_i q'(2), \ldots, \mu_i q'(Z)\right]^T$ is the $Z \times 1$ vector with entries $\frac{\partial q(z)}{\partial a_j(z)}$ and $\vec{a}_j$ is the the $Z \times 1$ vector with entries $a_j\left(z\right)$.

For the second part of the claim, consider the Hansen and Jagannathan (1991) decomposition of an admissible state price deflator, $\Lambda_i\left(z\right)$, into:

$$\Lambda_i\left(z\right) = \Lambda\left(z\right) + \left(\Lambda_i\left(z\right) - \Lambda\left(z\right)\right),$$

where a generic $\Lambda\left(z\right)$ is a state price deflator implied by market prices if:

$$\Lambda\left(z\right) = \bar{\Lambda} + \left(d\left(z\right) - d\Pi\right)\beta\left(\bar{\Lambda}\right),$$

$\bar{\Lambda}$ is the market-implied mean of the state price deflator, $\bar{\Lambda} = \left[\sum_{z\in\mathcal{Z}} q(z)\right]$, or the inverse of the market-implied riskfree rate, and Π is the $|\mathcal{Z}| \times 1$ vector with entries $\pi(z)$. As any admissible stochastic discount factor (SDF), $\Lambda(z)/p(z)$, has a Reisz Representation when projected onto the marketed space, or the space of dividend processes spanned by asset payoffs, this market-implied SDF necessarily takes this functional form, is the minimum variance SDF for a given SDF mean (or inverse of an implied fixed risk-free rate), and is unique.

From Hansen and Jagannathan (1991), the minimum variance choice of (SDF) $\Lambda(z)/\pi(z)$, $\Lambda^*(z)/\pi(z)$, satisfies:

$$\beta(\bar{\Lambda}) = \Sigma^{-1}(\iota_M - d\Pi\bar{\Lambda}) = \vec{q}E\left(DD'\right)^{-1}D,$$

where Σ is the covariance matrix of returns. By construction, one has that:

$$Cov(\Lambda^*(z), \Lambda_i(z) - \Lambda^*(z)) = 0. \tag{11}$$

For any $\Lambda_i(z)$ and $\Lambda_j(z)$, we then have

$$Cov\left(\Lambda^*(z), \Lambda_i(z) - \Lambda_j(z)\right) = 0. \tag{12}$$

Since:

$$\Lambda^* = \bar{\Lambda} + (d - \Pi d)\Sigma^{-1}(\iota_M - \Pi d\bar{\Lambda}),$$

it follows that:

$$Cov\left(d, \Lambda_i(z) - \Lambda_j(z)\right) = 0, \tag{13}$$

which is our orthogonality condition identifying d along with $\Lambda^*\iota_M = \left[\sum_{z\in\mathcal{Z}} q(z)\right]$. Consequently, we can express the implied market incompleteness by equation (12).

A.5 Proof of Proposition 5:

Writing the FOC for agent i as

$$q(z) + \mu\frac{\partial q(z)}{\partial a_i(z)}a_i(z) - \Lambda_i(z) = 0.$$

Let $\partial_A \vec{q}$ be the $|\mathcal{Z}| \times 1$ vector with entries $[\partial_A \vec{q}]_z = \frac{\partial q(z)}{\partial A(z)}$ and $\odot$ be the Hadamard product. Then by the Implicit Function Theorem across the $|\mathcal{Z}|$ markets

$$\partial_A \vec{q} = -\left(\begin{bmatrix} x_i(z_1)^{-1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & x_i\left(z_{|\mathcal{Z}|}\right)^{-1} \end{bmatrix} + \frac{\partial \vec{\Lambda}}{\partial c_{i,1}} \left(\vec{q} + \partial_A \vec{q} \odot \vec{a}_i\right)^T \right)^{-1} \vec{a}_i,$$

where $\partial_{c_{i,1}} \vec{\Lambda}$ is the $|\mathcal{Z}| \times 1$ vector with entries $\left[\frac{\partial \vec{\Lambda}}{\partial c_{i,1}}\right]_k = \frac{\partial \Lambda_i(z_k)}{\partial c_{i,1}}$, and

$$x_i(z) = \frac{1}{(1+\mu)\frac{\partial q(z)}{\partial A(z)} + \mu \frac{\partial^2 q(z)}{\partial A(z)^2} a_i(z) - \frac{\partial \Lambda_i(z)}{\partial c_i(z)}},$$

This represents the partial equilibrium response to a change in market power in all markets.

By the Sherman-Morrison formula, since u^T is an outer product of column matrices, we can express

$$\left(C + \vec{u}\vec{v}^T\right)^{-1} = C^{-1} - \frac{C^{-1}\vec{u}\vec{v}^T C^{-1}}{1 + \vec{v}^T C^{-1} \vec{u}},$$

from which follows that

$$\left(\begin{bmatrix} x_i(z_1)^{-1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & x_i\left(z_{|\mathcal{Z}|}\right)^{-1} \end{bmatrix} + \frac{\partial \vec{\Lambda}}{\partial c_{i,1}} \left(\vec{q} + \partial_A \vec{q} \odot \vec{a}_i\right)^T \right)^{-1}$$

$$= \begin{bmatrix} x_i(z_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & x_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix} - \frac{\begin{bmatrix} \frac{\partial \Lambda_i(z_1)}{\partial c_{i,1}} x_i(z_1) \\ \vdots \\ \frac{\partial \Lambda_i\left(z_{|\mathcal{Z}|}\right)}{\partial c_{i,1}} x_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix} \begin{bmatrix} \left(q(z_1) + \frac{\partial q(z_1)}{\partial A(z_1)} a_i(z_1)\right) x_i(z_1) \\ \vdots \\ \left(q\left(z_{|\mathcal{Z}|}\right) + \frac{\partial q\left(z_{|\mathcal{Z}|}\right)}{\partial A\left(z_{|\mathcal{Z}|}\right)} a_i\left(z_{|\mathcal{Z}|}\right)\right) x_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix}^T}{1 + \sum_{z \in \mathcal{Z}} \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z)\right) \frac{\partial \Lambda_i(z)}{\partial c_{i,1}} x_i(z)},$$

and therefore

$$\partial_A \vec{q} = -\mu \left(\begin{bmatrix} \frac{\partial q(z_1)}{\partial A(z_1)} x_i(z_1) \\ \vdots \\ \frac{\partial q\left(z_{|\mathcal{Z}|}\right)}{\partial A\left(z_{|\mathcal{Z}|}\right)} x_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix} - \frac{\sum_{z \in \mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z)\right) x_i(z)}{1 + \sum_{z \in \mathcal{Z}} \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z)\right) \frac{\partial \Lambda_i(z)}{\partial c_{i,1}} x_i(z)} \begin{bmatrix} \frac{\partial \Lambda_i(z_1)}{\partial c_{i,1}} x_i(z_1) \\ \vdots \\ \frac{\partial \Lambda_i\left(z_{|\mathcal{Z}|}\right)}{\partial c_{i,1}} x_i\left(z_{|\mathcal{Z}|}\right) \end{bmatrix} \right).$$

Since $\frac{\partial \Lambda_i(z)}{\partial c_{1i,}} = -\frac{u''\left(c_{i,1}\right)}{u'\left(c_{i,1}\right)}\Lambda_i\left(z\right)$, this further reduces to

$$\partial_A \vec{q} = -\begin{bmatrix} \frac{\frac{\partial q(z_1)}{\partial A(z_1)} - \frac{u''\left(c_{i,1}\right)}{u'\left(c_{i,1}\right)} \sum_{z\in\mathcal{Z}} \left(\frac{\partial q(z_1)}{\partial A(z_1)} \Lambda_i(z) - \frac{\partial q(z)}{\partial A(z)} \Lambda_i(z_1) \right) \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z) \right) x_i(z)}{1 - \frac{u''\left(c_{i,1}\right)}{u'\left(c_{i,1}\right)} \sum_{z\in\mathcal{Z}} \Lambda_i(z) \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z) \right) x_i(z)} x_i\left(z_1\right) \odot \vec{a}_i \\ \vdots \\ \frac{\frac{\partial q\left(z_{|\mathcal{Z}|}\right)}{\partial A\left(z_{|\mathcal{Z}|}\right)} - \frac{u''\left(c_{i,1}\right)}{u'\left(c_{i,1}\right)} \sum_{z\in\mathcal{Z}} \left(\frac{\partial q\left(z_{|\mathcal{Z}|}\right)}{\partial A\left(z_{|\mathcal{Z}|}\right)} \Lambda_i(z) - \frac{\partial q(z)}{\partial A(z)} \Lambda_i\left(z_{|\mathcal{Z}|}\right) \right) \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z) \right) x_i(z)}{1 - \frac{u''\left(c_{i,1}\right)}{u'\left(c_{i,1}\right)} \sum_{z\in\mathcal{Z}} \Lambda_i(z) \left(q(z) + \frac{\partial q(z)}{\partial A(z)} a_i(z) \right) x_i(z)} x_i\left(z_{|\mathcal{Z}|}\right) \odot \vec{a}_i \end{bmatrix},$$

where $\odot$ is the Hadamard product.

In the special case of quasi-linear preferences, the above reduces to:

$$\frac{\partial a_i(z)}{\partial \mu_i} = \frac{-q'(z)a_i(z)}{(1+\mu)q'(z) + \mu q''(z)a_i(z) - \frac{\partial \Lambda_i(z)}{\partial a_i(z)}},$$

We next recognize that sellers distort upwards and buyers distort downward. We therefore define the relevant elasticity to be

$$\epsilon_i = \frac{q'(z)\mu}{(1+\mu)q'(z) + \mu q''(z)a_i(z) - \frac{\partial \Lambda_i(z)}{\partial a_i(z)}},$$

where $\frac{\partial \Lambda_i(z)}{\partial a_i(z)} < 0$. This is essentially the absolute value of the elasticity for the particular case where quantities can be positive and negative.

Let agent i be a buyer and j be a seller. When agents have the same preferences, which determines $\Lambda_i\left(z\right)$, then for buyers $a_i\left(z\right) > 0 > a_j\left(z\right)$ and $\frac{\partial \Lambda_i(z)}{\partial a_i(z)} < \frac{\partial \Lambda_j(z)}{\partial a_j(z)} < 0$. Since $q''\left(z\right) > 0$, it follows that:

$$q''\left(z\right)a_i\left(z\right) - \frac{\partial \Lambda_i(z)}{\partial a_i(z)} > q''\left(z\right)a_j\left(z\right) - \frac{\partial \Lambda_j(z)}{\partial a_j(z)}.$$

Consequently, it follows that $\epsilon_j > \epsilon_i$.

A.6 Proof of Proposition 6

The first two statements are straightforward. For the final part of the claim, assume that fringe has CARA preferences. Then, $q''(z)/q'(z) = \gamma_f$ and $q'\left(z\right) = \gamma_f q\left(z\right)$, and we can write the elasticity as

$$\epsilon_i = \frac{\mu}{1 + \mu + \mu\gamma_f a_i(z) - \gamma_f^{-1}\frac{\Lambda_i'(z)}{q(z)}}.$$

We now consider the following thought experiment: if other agent's valuation rise so that $q(z)$ increases, what is the effect on the elasticity of substitution? We have

$$\frac{\partial\epsilon_i}{\partial q(z)} = \frac{-\mu\left(\mu\gamma_f\frac{\partial a_i(z)}{\partial q(z)} - \frac{1}{\gamma_f q(z)}\frac{\partial\Lambda_i'(z)}{\partial q(z)} + \frac{\Lambda_i'(z)}{\gamma_f q(z)^2}\right)}{\left(1 + \mu - \mu\gamma_f a_i(z) + \frac{\Lambda_i'(z)}{\gamma_f q(z)}\right)^2} = \frac{\epsilon_i^2}{\mu}\left(-\mu\gamma_f\frac{\partial a_i\left(z\right)}{\partial q\left(z\right)} + \frac{1}{\gamma_f q\left(z\right)}\frac{\partial\Lambda_i'\left(z\right)}{\partial q\left(z\right)} - \frac{\Lambda_i'(z)}{\gamma_f q(z)^2}\right).$$

If $a\left(z\right) > 0$, then an increase in $a\left(z\right)$ reduces marginal utility, but a higher price $q\left(z\right)$ counteracts this (all else equal, buyer buys on the margin less), so $\frac{\partial\Lambda_i'(z)}{\partial q(z)} > 0$. Similarly, $\frac{\partial a_i(z)}{\partial q(z)} < 0$. It then follows that

$$\frac{\partial\epsilon_i}{\partial q(z)} = \frac{\epsilon_i^2}{\mu}\left(-\mu\gamma_f\frac{\partial a_i\left(z\right)}{\partial q\left(z\right)} + \frac{1}{\gamma_f q\left(z\right)}\frac{\partial\Lambda_i'\left(z\right)}{\partial q\left(z\right)} - \frac{\Lambda_i'(z)}{\gamma_f q(z)^2}\right) > 0.$$

This implies that agents distort more aggressively when other agents' valuations are high, reinforcing the strategic complementarity in their asset demands.

A.7 Proof of Proposition 7:

(i) We start with the aggregated FOCs:

$$q\left(z\right) + \frac{\mu}{N+1}\frac{\partial q\left(z\right)}{\partial A\left(z\right)}A\left(z\right) = E^*\left[\Lambda_i\left(z\right)\right].$$

Compare now an economy along the sequence of $\mu's$ and the competitive equilibrium (superscript CE):

$$q\left(z\right) + \frac{\mu}{N+1}\frac{\partial q\left(z\right)}{\partial A\left(z\right)}A\left(z\right) = E^*\left[\Lambda_i\left(z\right)\right] - E^*\left[\Lambda_i^{CE}\left(z\right)\right].$$

In competitive equilibrium, $E^*\left[\Lambda_i^{CE}\left(z\right)\right] = \Lambda^{CE}\left(z\right)$, i.e. state prices are all aligned. We invoke the following lemma.

Lemma 1: When capital decisions are fixed or there is no storage, the average state price across all $N+1$ agents is (weakly) higher with market power than in the competitive equi-

librium, or $E^{*}\left[\Lambda_{i}\left(z\right)\right]>\Lambda^{CE}\left(z\right)$. With storage, state prices of high marginal utility states are higher while those of low marginal utility states are lower, so that their sum is fixed at $\frac{1}{R}$.

It therefore follows that:

$$q\left(z\right)+\frac{\mu}{N+1}\frac{\partial q\left(z\right)}{\partial A\left(z\right)}A\left(z\right)\geq 0.$$

Since the state price $q\left(z\right)$ is related to the marginal utility of the competitive fringe at date 2, and the competitive fringe has strictly convex marginal utility, $\frac{\partial q(z)}{\partial A(z)}>0$. If $A\left(z\right)\leq 0$, it is then immediate that $q\left(z\right)$ is higher than in the competitive equilibrium.

Since $A\left(z\right)=0\forall\ z$ in the competitive equilibrium by Assumption 1, it follows that $\sum_{z\in\mathcal{Z}}q\left(z\right)$ is larger than in the competitive equilibrium. Since r_{f}^{m} is the inverse of the sum of the state prices, the claim then follows. Moreover, $r_{f}^{m}\geq R$ because storage can be traded without market impact.

(ii) The second part of the claim follows directly from Lemma 1. The floor of r_{f}^{*} is also R since, from the FOCs in the trading game that, if there is storage, then $E\left[u_{2}'\left(\frac{c_{i2}(z)}{c_{i1}}\right)\right]=\frac{1}{R}$. To see that this is a lower bound, a necessary condition that agent i to invest in capital in its production technology is $E\left[y_{i}\left(z\right)\right]>R$, with the strict inequality necessary to embed a risk premium for the agent. Consequently, as long as agent i holds storage, then $E\left[u_{2}'\left(\frac{c_{i2}(z)}{c_{i1}}\right)\right]=\frac{1}{R}$, and only once it exhausts all its resources in state contingent claims and capital, then $E\left[u_{2}'\left(\frac{c_{i2}(z)}{c_{i1}}\right)\right]<\frac{1}{R}$. Since this holds for all agents, it follows that $r^{*}\geq R$.[11]

Furthermore, notice that the expected gross return to capital for agent i is $E\left[y_{i}\left(z\right)-1\right]$, since capital is supplied elastically to all agents at unit cost. It then follows that, if the riskless rate is depressed, that the expected excess return to capital, $E\left[A\left(z\right)-r^{*}\right]$, increases in the economy. Consequently, the expected excess return on investment increases in the economy, raising Tobin's q.

(iii) Follows from a summation over states of Equation (9). Notice that $R\left(z\right)$ is the ex-

[11] Though agents will not trade risk-free claims with each other because of market power, they may still trade with outside investors as a result of their difference in preferences.

pected excess payoff in state z. Summing across all states, it follows that:

$$\sum_{z\in\mathcal{Z}} R(z) = \sum_{z\in\mathcal{Z}} E^{*}\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right] \geq 0,$$

since $\sum_{z\in\mathcal{Z}} Cov\left(\bar{u}_2'(z) + \overline{Lev}(z), \delta_j(z)\right) = 0$. Consequently, $\sum_{z\in\mathcal{Z}} R(z)$ is nonzero despite being a riskless expected excess return derived from assembling all state contingent securities. It must therefore be the case that the market-implied riskless rate $\left[\sum_{z\in\mathcal{Z}} q(z)\right]^{-1}$ is biased based on the average market power, $\sum_{z\in\mathcal{Z}} E^{*}\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right]$, in financial markets. Since there is anonymity in financial markets, in that price impact is symmetric across agents ($\frac{\partial q(z)}{\partial a_i(z)} = \frac{\partial q(z)}{\partial a_j(z)}$), the bias is determined by the probability-weighted average of the net demand of agents, $\sum_{z\in\mathcal{Z}} \frac{\partial q(z)}{\partial a_i(z)} E^{*}[a_i(z)] = \sum_{z\in\mathcal{Z}} \frac{\partial q(z)}{\partial A(z)} A(z)$. The bias is positive (negative) when agent net demand is positive (negative), which is the supply of the outside agent. [12]

(iv) We start by defining the expected excess payoff to state z be $R^e(z) = \pi(z) - r_f^{*} q(z)$, where $\frac{1}{r^{*}} = E\left[E^{*}\left[\Lambda_i(z)\right]\right] = E\left[E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]\right]$ and $E^{*}[\cdot]$ is the cross-sectional average $\frac{1}{|N|}\sum_{i=1}^{|N|}[\cdot]$. r^{*} is the state-price implied risk-free rate constructed from the cross-sectional average of agent state prices. Then, it follows that we can express this return from the FONCs for $a_i(z)$ as:

$$\begin{aligned} R^e(z) &= \pi(z) - r_f^{*}\pi(z) E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right] + r_f^{*} E^{*}\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right] \\ &= -Cov\left(\frac{E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]}{E\left[E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]\right]}, \delta(z)\right) + r_f^{*} E^{*}\left[h_i \frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right], \end{aligned}$$

since $E\left[\frac{E^{*}[\Lambda_i(z)]}{\sum_{z\in\mathcal{Z}} E^{*}[\Lambda_i(z)]}\right] E[\delta(z)] = \pi(z)$ and $E\left[\frac{E^{*}[\Lambda_i(z)]}{\sum_{z\in\mathcal{Z}} E^{*}[\Lambda_i(z)]}\delta(z)\right] = \frac{E^{*}[\Lambda_i(z)]}{\sum_{z\in\mathcal{Z}} E^{*}[\Lambda_i(z)]}\pi(z)$. We can further decompose this average excess payoff as:

$$R^e(z) = -Cov\left(\frac{E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]}{E\left[E^{*}\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]\right]}, \delta(z)\right) + r_f^{*} E^{*}\left[\frac{\partial q(z)}{\partial a_i(z)}\right].$$

Dividing by the price of the claim, $q(z)$, and substituting with the market-implied

[12] If instead there were internal market-clearing among the large agents, then asymmetry in price impact would again give rise to a nontrivial bias.

interest rate, we arrive at the expected excess market-implied return, $r(z)$:

$$r(z) = -Cov\left(\frac{\bar{u}'(z)}{E[\bar{u}'(z)]}, \frac{\delta(z)}{q(z)}\right) + \overline{mkt}(z),$$

where:

$$\begin{aligned}
\overline{mkt}(z) &= r^* E^*\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right] + r_f^* - \left[\sum_{z\in\mathcal{Z}} q(z)\right]^{-1} \\
&= \left(E^*\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right] + 1 - \frac{\sum_{z\in\mathcal{Z}} E^*[\Lambda_i(z)]}{\sum_{z\in\mathcal{Z}} q(z)}\right) r_f^* \\
&= \left(E^*\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right] - \frac{\sum_{z\in\mathcal{Z}} E^*\left[\frac{\partial q(z)}{\partial a_i(z)} a_i(z)\right]}{\sum_{z\in\mathcal{Z}} q(z)}\right) r_f^* \\
&= \left(r_f^m \sum_{z\in\mathcal{Z}} \frac{q(z)}{N\eta(z)} - -\frac{q(z)}{N\eta(z)}\right) r_f^*
\end{aligned}$$

with $\bar{u}'(z) = E^*\left[\frac{u_2'(c_{i2}(z))}{u_1'(c_{i1})}\right]$.

A.8 Proof of Proposition 8:

From the main text, we recognize that we can decompose the Euler Equation of strategic agent i as

$$1 = \sum_{z\in\mathcal{Z}} q(z) \frac{(y_i(z) k_i + Rs_i)}{k_i + s_i} + \sum_{z\in\mathcal{Z}} \mu_i q'(z) a_i(z) \frac{(y_i(z) k_i + Rs_i)}{k_i + s_i},$$

where the first term on the RHS is the envelope condition arising from the conventional Euler Equation. Defining

$$mv_i = \sum_{z\in\mathcal{Z}} q(z) \frac{(y_i(z) k_i + Rs_i)}{k_i + s_i}, p_i = \sum_{z\in\mathcal{Z}} \mu_i q'(z) a_i(z) \frac{(y_i(z) k_i + Rs_i)}{k_i + s_i},$$

we arrive at the statement in the proposition.

It then follows that, if the second term in (A.8) is positive, then the first term is less than unity, from which follows that the firm chooses a smaller scale of production than if it were perfectly competitive. Similarly, if the second term is negative, then the first term is greater than unity, and the firm chooses a larger scale of production than if it were perfectly competitive.

Now, let k_i^{CE} be the scale of capital an agent would choose in the competitive equilibrium without price impact. Further, let k_i^{Aut} be the scale of capital an agent would choose in autarky, or if it could not trade state contingent claims with other agents.

There are two forces impacting the choice of capital with price impact, the first indirect and the second direct. First, since $\Lambda_i(z) > \Lambda_j(z)$ in the presence of strategic interaction, it follows that j cannot sell as many claims to i as it would, all else equal, in a competitive environment. By a symmetric argument, j is not able to buy as claims from i as it would in any state z' in which it would like to buy, such that $\Lambda_j(z) > \Lambda_i(z)$, as compared to a competitive environment. Consequently, j is more exposed to the output of its own production than, all else equal, in a competitive environment. Since its state price $\Lambda_j(z)$ is more exposed to its own output, $y_j(z)k_j$, because of diminished risk-sharing opportunities ($\Lambda_j(z)$ lower when $y_j(z)$ is higher, and vice versa), it follows from the FONC for capital that agent j chooses a level of capital k_j that approaches its level under autarky, k_j^{AUT}. Second, as a result of price impact, agent j has an incentive to restrict its supply of state contingent securities to earn oligopoly rents as a seller. This force also lowers the optimal choice of capital. Consequently, given these two forces, if $k_j^{CE} > k_j^{Aut}$, then k_j decreases because of price impact.

If instead $k_j^{CE} \leq k_j^{Aut}$, then agent j's capital choice is more ambiguous. The first force leads k_j to increase toward the autarky capital choice, while the second reduces k_j because of strategic rationing. As a result, k_j may be increasing (first force), decreasing (second force), or U-shaped (both forces) in μ.

Given these cross-sectional results, we now characterize the aggregate implications. For agents for whom $k_j^{CE} > k_j^{Aut}$, the set $\mathcal{J}$, their total capital investment, $\sum_{j\in\mathcal{J}} k_j$, is decreasing in price impact. For agents for which $k_j^{CE} \leq k_j^{Aut}$, the set $1...,N/\mathcal{J}$, their total capital investment, $\sum_{j\in 1...,N/\mathcal{J}} k_j$, may be increasing in price impact. Since price impact acts as an effective tax on the joint production among agents, capital misallocation rises, which lowers total output in each state and, consequently, productivity. That total capital investment also falls follows from the asymmetry that more productive agents, $k_j^{CE} > k_j^{Aut}$, rely on risk-sharing for financing their capital investment, while less productive agents, $k_j^{CE} \leq k_j^{Aut}$, rely on risk-sharing to reduce their scale.

When capital investment is (weakly) lower, savings is (weakly) higher.

A.9 Proof of Corollary 1

Consider the expected excess return to each state z :

$$ER(z) = \frac{\pi(z)}{q(z)} - \frac{1}{\sum_z q(z)}.$$

Compared to their competitive equilibrium counterparts, $ER^{CE}(z)$, it follows that:

$$ER(z) - ER^{CE}(z) = \frac{\pi(z)}{q(z)} - \frac{\pi(z)}{q^{CE}(z)} + \frac{1}{\sum_z q^{CE}(z)} - \frac{1}{\sum_z q(z)},$$

where $\pi(z)/q(z) - \pi(z)/q^{CE}(z) < 0$ and $1/\sum_z q^{CE}(z) - 1/\sum_z q(z) > 0$ from the previous proposition. Notice that $q(z) = \pi(z) u'(c_{f,2}(z))$ from market-clearing, it follows that $\frac{\pi(z)}{q(z)} = u'(c_{f,2}(z))^{-1}$, and consequently $\pi(z)/q(z)$ is driven by the marginal utility of the fringe and not the probability of each state.

Recognizing that $c_{f,2}(z) = c_{f,2}^{CE}(z) - \Delta D(z)$ is the change in the competitive fringe's position, we can rewrite the difference in returns as:

$$\begin{aligned} ER(z) - ER^{CE}(z) &= \frac{1}{u'\left(c_{f,2}^{CE}(z) - \Delta D(z)\right)} - \frac{1}{u'\left(c_{f,2}^{CE}(z)\right)} \\ &\quad + \frac{1}{\sum_z \pi(z) u'\left(c_{f,2}^{CE}(z)\right)} - \frac{1}{\sum_z \pi(z) u'\left(c_{f,2}^{CE}(z) - \Delta D(z)\right)}. \end{aligned}$$

Since

$$\frac{d}{d\varepsilon}\left(\frac{1}{x+\varepsilon} - \frac{1}{x}\right) = \frac{1}{x^2} - \frac{1}{(x+\varepsilon)^2} > 0,$$

it follows that $\frac{\pi(z)}{q(z)} - \frac{\pi(z)}{q^{CE}(z)}$ is more negative for states with lower $q^{CE}(z)$, all else equal.

Taking a first-order approximation around $c_{f,2}^{CE}$, we have that:

$$ER(z) - ER^{CE}(z) \approx -\frac{\gamma(z)}{u'\left(c_{f,2}^{CE}(z)\right)} \Delta D(z) + \frac{\sum_{z'} \pi(z') u'\left(c_{f,2}^{CE}(z')\right) \gamma(z') \Delta D(z')}{\left(\sum_{z'} \pi(z') u'\left(c_{f,2}^{CE}(z')\right)\right)^2},$$

where $\gamma_f(z)$ is the fringe's CARA coefficient evaluated at $c_{f,2}^{CE}(z)$. Recognizing $\frac{\gamma_f(z)}{u'\left(c_{f,2}^{CE}(z)\right)} = \frac{\pi(z) q'(z)}{q(z)^2}$, $\gamma_f(z) u'\left(c_{f,2}^{CE}(z)\right) = q'(z)$ and $r_f^m = \left[\sum_{z'} \pi(z') u'\left(c_{f,2}^{CE}(z')\right)\right]^{-1}$, we arrive at the expression in the main text.

While $\gamma(z)$ is non-increasing in wealth, and $\frac{1}{u'\left(c_{f,2}^{CE}\right)}$ is increasing by the strict convexity of $u'(\cdot)$, it follows the overall relation of $\frac{\gamma(z)}{u'\left(c_{f,2}^{CE}\right)}$ with $c_{f,2}^{CE}$ depends on the utility function (since the expression is equivalent to $du'\left(c_{f,2}^{CE}(z)\right)^{-1}/dc_{f,2}^{CE}(z)$).

If $du''\left(c_{f,2}^{CE}(z)\right)^{-1}/dc_{f,2}^{CE}(z)^2 > 0$, then $\frac{\pi(z)q'(z)}{q(z)^2}$ is increasing in $c_{f,2}^{CE}(z)$, and consequently decreasing in the marginal utility of the fringe. Otherwise, it is decreasing in the marginal utility of the fringe. For strictly convex marginal utility, CARA and CRRA utility functions with CRRA coefficient ≥ 1, for instance, $u'\left(c_{f,2}^{CE}(z)\right)^{-1}$ is also convex, and consequently $du''\left(c_{f,2}^{CE}(z)\right)^{-1}/dc_{f,2}^{CE}(z)^2 > 0$.

A.10 Proof of Proposition 9:

We consider a perturbation from the perspective of a social planner who can design a mechanism to implement any allocation that respects the FOCs of all agents and the market structure that features imperfect competition. Expanding the differential of the indirect utility function for agent i, U_1^i, we see that:

$$\Delta U_i = u_1^{i\prime}(c_{i1})\,\Delta c_{i1} + \sum_{z\in\mathcal{Z}} \pi(z)\,u_2^{i\prime}(c_{i2})\,\Delta c_{i2}(z) + \Delta\left(\phi_i^k k_i\right) + \Delta\left(\phi_i^s s_i\right),$$

where ϕ_i^k is the nonnegativity Lagrange multiplier on k_i, and similarly for ϕ_i^s. Since the prices of storage and capital are both fixed at unity, the Planner cannot manipulate them to improve welfare. Substituting with the FOCs for capital and storage in Proposition 2, and invoking complementarity slackness, the above reduces to:

$$\Delta U_i = u_1^{i\prime}(c_{i1}) \sum_{z\in\mathcal{Z}} \left(\Lambda_i(z)\,\Delta a_i(z) - \Delta\left(q(z)\,a_i(z)\right)\right),$$

where $\Lambda_i(z)$ is the state price of agent i.

By the First Welfare Theorem, when markets are competitive, the competitive equilibrium solves the Planner's Problem and achieves the first-best outcome. It follows that the equilibrium with competitive agents is Pareto efficient.

Important for this result is that, when markets are competitive, $q(z) = \Lambda_i(z) = \Lambda(z)\ \forall i, j$, since state contingent claims align state prices across all agents. Since the state contingent claim price $q(z)$ equals the marginal benefit for buyers $\Lambda_i(z)$ and the marginal cost for sellers $\Lambda_j(z)$, the marginal benefit for buyers equals the marginal cost for sellers. Consequently, all gains from trade are exhausted.

With imperfect competition, however, there is now a deadweight loss (Harbinger Triangle) since buyers in a security act as "local" oligopsonists, and sellers act as "local" oligopolists. For sellers, $q(z) > \Lambda_j(z)$, and they sell securities against state z above their marginal cost ($\Lambda_j(z)$), while for buyers, $q(z) < \Lambda_i(z)$, and they buy securities against state z below their marginal benefit ($\Lambda_i(z)$). Consequently, marginal benefit $>$ price $>$ marginal cost since, intuitively, both sides of the market set marginal revenue instead of the price equal to marginal cost. This reduces the quantity traded and leaves potential gains from trade unrecognized.

Consequently, the Planner can improve on the noncompetitive equilibrium state-by-state by having either buyers purchase the next extra marginal unit at their marginal benefit, $\Lambda_i(z)$, or sellers sell their next extra marginal unit at marginal cost, $\Lambda_j(z)$. This can be implemented by having sellers and buyers trade an additional $\epsilon(z)$ of securities against state z at price $q(z)$ and taxing buyers $\frac{\partial q(z)}{\partial A(z)}\epsilon(z)$ when $q(z)$ is below its competitive equilibrium counterpart, and instead sellers when it is above. By subsidizing the trades of the other side of the market with the taxed revenue, this leaves the taxed agents' welfare unchanged and those who receive the trading subsidy strictly better off. Notice that, since price impact is anonymous, $\frac{\partial q(z)}{\partial a_i(z)} = \frac{\partial q(z)}{\partial a_j(z)} = \frac{\partial q(z)}{\partial A(z)}$, the Planner only needs to observe whether agents are long or short in the market to implement the policy. Since these units are still traded at price $q(z)$, this also leaves the utility of the fringe unchanged.

Generically, then, the equilibrium in the trading game with imperfect competition is constrained Pareto inefficient.

We now turn to policy. The first-order gains follow from the previous argument. By implementing two-sided trading costs / subsidies, the Planner can improve risk-sharing among the strategic agents. Alternatively, as is more standard in the constrained efficiency literature (Dréze (1974)), these income transfers could be implemented through side payments equal to the proposed subsidies.

To examine second-order gains to policy, we proceed by a perturbation argument and consider a planner who influences agents' production and storage decisions. Suppose the Planner subsidizes agent $i's$ production k_i by taxing its savings s_i such the overall impact is revenue neutral. Suppose this is more efficient in that risky production has a higher expected return than storage. Then, to first-order, the change in production is Δ_i and in savings is $-\Delta_i$. Through the portfolio choice of the agent, this raises its supply of securities in states in which $y_i(z) > R$ and raises its demand in states in which $y_i(z) < R$. For small enough Δ_i, this leaves consumption $c_{i,2}(z)$ in each state, and the indirect utility of agent i, unchanged.

The impact of this change in security demand, however, changes security prices. Specifically, it lowers prices in states in which agent i supplies more, and raises prices in states in which agent i demands more. This impacts the prices that the other agents face, increasing their demand in the states with now lower prices and lowering it (increasing supply) in states with now higher prices. This change in asset demand feeds back into the production decisions of the other agents. Those who have production that complements agent i (high production when agent $i's$ is low and vice versa) now have incentive to increase their own production, since they can buy insurance from agent i cheaper and sell insurance to agent i more dearly. As such, risky production, which is more efficient than storage, increases in the economy, shifting outward the Pareto frontier. Since there are now more resources in the economy, which can always be disposed of freely, at least one agent must be (weakly) better off. This perturbation analysis is of second-order, since it involves the interaction of production and financial markets.

A.11 Proof of Proposition 10:

Suppose instead of a competitive fringe that trades in all assets, there are segmented fringes that trade in different subsets of the state contingent securities.

Since the competitive fringe is disjoint, notice that the Law of One Price need not hold. All that is imposed from the FONCs of the strategic agents with the introduction of a redundant security j such that $d_j(z) = \sum_{z \in \mathcal{Z}} \alpha(z) \mathbf{1}_{\{z'=z\}}$ is that

$$q_j - \frac{\partial q_j}{\partial a_{i,j}} a_{i,j} = \sum_k \alpha_k \left(q_k - \frac{\partial q_k}{\partial a_{i,k}} a_{i,k} \right),$$

since $\sum_{z \in \mathcal{Z}} \Lambda_i(z) d_j(z) = \sum_{z \in \mathcal{Z}} \Lambda_i(z) \sum_{z \in \mathcal{Z}} \alpha(z) \mathbf{1}_{\{z'=z\}}$. This condition, in itself, admits many possible solutions, and does not impose real neutrality.

When the fringe that trades in the redundant security also trades in all replicating state contingent securities, then by Proposition ___ $q_j = \sum_{z \in \mathcal{Z}} \alpha(z) q(z)$, and real neutrality holds. When the fringe that trades in the redundant security does not trade in all replicating state contingent securities, however, then the Law of One Price holds only when

$$\sum_{z\in\mathcal{Z}}\Lambda_{fj}(z)\,d_j(z)=\sum_{z\in\mathcal{Z}}\Lambda_{fz}(z)\,\alpha(z)\,,$$

where $\Lambda_{fj}(z)$ indexes the member of the fringe that trades the redundant security and $\Lambda_{fz}(z)$ references the member of the fringe that trades in the state contingent security in state z. Generically, however, this condition is not satisfied. As such, we can define $q_j=\sum_k \alpha_k q_k+\delta_j$, for some nonzero wedge δ_j that is unrelated to any exposure to risk.

A.12 Proof of Lemma 1:

We first fix the capital and savings decisions of the N strategic agents, and assume that no storage is employed. We then have an endowment economy. We consider the Planner's problem. Wilson's optimal sharing rule for risk-averse agents reveals that:

$$\begin{aligned} u'\left(\eta_i^{-1}\right)u'\left(c_i(z)\right) &= \lambda(z)\,, \\ u'\left(\eta_i^{-1}\right)u'\left(c_{i,1}\right) &= \lambda_1, \end{aligned}$$

for Pareto weights $u'\left(\eta_{i-1}^{-1}\right)$, and for the risk-neutral at date 1 :

$$u'\left(\eta_f^{-1}\right)=\lambda_1.$$

It follows for date 2 :

$$e(z)=\sum_{i=1}^{f}c_i(z)=u'^{-1}\left(\lambda(z)\right)\sum_{i=1}^{f}\eta_i$$

from which follows that:

$$c_i(z)=\frac{\eta_i}{\sum_{i=1}^{N}\eta_i}e(z)\,.$$

Consequently, all agents consume a fraction of the aggregate endowment. From First Welfare Theorem, this is equivalent to competitive equilibrium allocation, and there exist Pareto weights $\{\eta_i\}_{i=1}^{f}$ that centralize the competitive equilibrium.

At date 1, we modify the argument:

$$e_1 - c_{f,1} = \sum_{i=1}^{N} c_{i,1} = u'^{-1}(\lambda_1) \sum_{i=1}^{N} \eta_i,$$

Since $u'^{-1}(\lambda_1)\,\eta_f = 1$, it follows that:

$$e_1 + 1 - c_{f,1} = u'^{-1}(\lambda_1) \sum_{i=1}^{N+1} \eta_i,$$

and therefore:

$$c_{i,1} = \frac{\eta_i}{\eta_{N+1}} = \frac{\eta_i}{\sum_{i=1}^{N+1} \eta_i} \left(e_1 + 1 - c_{f,1}\right).$$

All risk-averse agents consume a fixed fraction of residual endowment.

From the resource constraint, it follows that:

$$c_{f,1} = e_1 - \frac{\sum_{i=1}^{N} \eta_i}{\eta_f},$$

Consequently, in competitive equilibrium $\Lambda_i(z) = \pi(z)\frac{u'(c_i(z))}{u'(c_{i,1})} = \pi(z)\, u'\left(\frac{e(z)}{e_1+1-c_{f,1}}\right)$ is the same across all risk averse agents, while for quasilinear $\Lambda_f(z) = \pi(z)\, u'\left(c_f(z)\right)$. Notice:

$$\begin{aligned}
u'\left(c_f(z)\right) &= u'\left(\frac{\eta_f}{\sum_{i=1}^{N+1} \eta_i}\right) u'(e(z)) \\
&= u'\left(\frac{u'^{-1}(\lambda_1)\,\eta_{N+1}}{u'^{-1}(\lambda_1)\sum_{i=1}^{N} \eta_i + u'^{-1}(\lambda_1)\,\eta_{+1}}\right) u'(e(z)) \\
&= u'\left(\frac{u'^{-1}(\lambda_1)\,\eta_f}{e_1 - c_{f,1} + u'^{-1}(\lambda_1)\,\eta_{+1}}\right) u'(e(z)) \\
&= u'\left(\frac{e(z)}{e_1 + 1 - c_{f,1}}\right)
\end{aligned}$$

since $\frac{1}{\eta_f} = u'^{-1}(\lambda_1)$, as required. Consequently, state prices are aligned across all agents.

Consider now consider an economy with market power, homothetic preferences, and strictly convex marginal utility. Applying Jensen's Inequality:

$$E^*\left[\Lambda_i(z)\right] = E^*\left[u'\left(\frac{c_{i,1}}{c_i(z)}\right)^{-1}\right] \geq \pi(z)\, u'^{-1}\left(E^*\left[\frac{c_{i,1}}{c_i(z)}\right]\right)^{-1},$$

since $c_{f,1} \geq 1$ and strict convexity and concavity are preserved for the inverse of marginal utility $u(x)^{-1}$. In the above, $c_{f,1}$ is understood as 1 in the above formulation. Notice that the aggregate consumption of the N strategic agents is $e_1 - c_{f,1}$, where e_1 is the aggregate endowment. Applying Jensen's Inequality again, it follows that:

$$E^*\left[\Lambda_i(z)\right] \geq \pi(z)\, u'\left(\frac{e_1 - c_{f,1} + 1}{e(z)}\right)^{-1} = \pi(z)\, u'\left(\frac{e(z)}{e_1 + 1 - c_{f,1}}\right) = \Lambda^{CE\prime}(z),$$

where the 1 arises from the effective consumption of the fringe since it has quasilinear preferences. Therefore,

$$E^*\left[\Lambda_i(z)\right] > \Lambda^{CE\prime}(z),$$

where $\Lambda^{CE\prime}(z)$ references the state price of the representative agent in a pseudo-economy that holds fixed the investments of the N agents as endowments. The above inequality is strict by the strict concavity of agent utility. This holds for all z since state prices are more misaligned state-by-state.

Suppose now that storage is employed. For those agents that use storage the sum of their state prices is fixed at $\frac{1}{R}$, or

$$E\left[\Lambda_i(z)\right] = \frac{1}{R} \forall i \in 1, ..., N.$$

Since only the sum of their state prices are constrained, state prices still are higher by Jensen's Inequality. If all agents employ storage, then the sum of all state prices is constrained, and so is the cross-sectional average of the sum. Then, since average state prices cannot all be higher than in the psuedo-competitive equilibrium, the cross-sectional average must cease to change. Otherwise, the sum would exceed $\frac{1}{R}$, a contradiction.

Since we have established that state prices are (weakly) higher with market power than in an equivalent endowment economy with the same capital and savings decisions, what remains is to compare the competitive equilibrium to the pseudo-competitive equilibrium with perfect risk-sharing but the capital allocation decisions from the strategic equilibrium. Then, our claim will follow by transitivity.

First, consider the case in which storage is not employed in the competitive or in the strategic equilibrium. Then, only capital adjusts. As a result of the aggregation of "mongrel" preferences, based on the state price $\pi(z)\, u'\left(\frac{e(z)}{e_1+1-c_{f,1}}\right)$ in the competitive equilibrium, to arrive at the representative agent, the effective consumption of all agents at date 1 is $e_1 + 1 - c_{f,1} = \frac{\sum_{i=1}^{N+1}\eta_i}{\eta_f}$. The representative agent's preferences, however, do preserve the same cardinal utility functions as the risk-averse agents. Intuitively, since the fringe has quasi-linear preferences, it will fully insure the N risk-averse strategic agents to the extent that the Planner values its consumption vs that of the N agents, as reflected in the Pareto weights.

With a representative agent and competitive markets, we can solve the Planner's Problem again by the First Welfare Theorem:

$$
\begin{aligned}
U_0 &= \sup_{\{k_i\}_{i=1}^{N}, s} E\left[u_1\left(\frac{\sum_{i=1}^{N+1}\eta_i}{\eta_f}\right) + u_2\left(c_2\right)\right] \\
s.t. &: \sum_{i=1}^{N} k_i + s = (N+1)\, e - \frac{\sum_{i=1}^{N+1}\eta_i}{\eta_f}, \\
&: c_2(z) = \sum_{i=1}^{N} y_i(z)\, k_i + Rs,
\end{aligned}
$$

with indirect utility of the representative agent U_0.

By construction, the distribution of investment with market power can achieve no higher utility than under the optimal centralized policy with the same resource constraint. As such, there exist improvements that (weakly) raise $c_2(z)$ in all states by shifting investment away from less toward more productive technologies.

From Proposition 8 and Corollary 5, agents that would invest in capital in the first-best under-invest in the noncompetitive economy because of strategic frictions, while those that do not invest may start to invest because of diminished opportunities to finance the production of other agents. For the same resources transferred intertemporally, $\sum_{i=0}^{N} e_i - c_i$, the first-best employs more efficient technologies without cross-sectional misallocation. As such, while some agents may under-invest and others over-invest with price impact, aggregate investment and investment efficiency falls. As such, state prices are (weakly) higher in all states in the representative agent competitive equilibrium.

Finally, we consider the role of storage. If storage is used in the competitive equilibrium, then it is also employed in the strategic equilibrium, although the converse need not be true. When storage is employed by all agents, then the sum of their state prices is constrained to be $\frac{1}{R}$. As such, since capital allocation is less efficient in the strategic equilibrium, and the sum of state prices is constrained, it follows that prices increase for high marginal utility and decrease for lower marginal utility states to leave the average unchanged. If there is storage in the competitive equilibrium, then this effect is there for all μ, otherwise it becomes operative once all agents employ storage. It then follows that:

$$E^{*}\left[\Lambda_{i}\left(z\right)\right] > \Lambda^{CE\prime}\left(z\right) > \Lambda^{CE}\left(z\right),$$

where $\Lambda^{CE}\left(z\right)$ is the state price of the representative agent in the competitive equilibrium.