

Speculation and Risk Sharing with New Financial Assets

Alp Simsek*

May 2011

Abstract

While the traditional view of financial innovation emphasizes the risk sharing role of new financial assets, belief disagreements about these assets naturally lead to *speculation*, which represents a powerful economic force in the opposite direction. This paper investigates the effect of financial innovation on risks in an economy when both the risk sharing and the speculation forces are present. I consider this question in a standard CARA-Normal framework. Financial assets provide hedging services but they are also subject to speculation because traders do not necessarily agree about their payoffs. I define the average variance of traders' net worths as a measure of financial stability for this economy, and I decompose it into two components: the *uninsurable variance*, defined as the average variance that would obtain if there were no belief disagreements, and the *speculative variance*, defined as the residual variance that results from speculative trades based on belief disagreements. Financial innovation always decreases the uninsurable variance because new assets increase the possibilities for risk sharing. My main result shows that financial innovation also always increases the speculative variance. This is true even if traders completely agree about the payoffs of new assets. The intuition behind this result is the *hedge-more/bet-more* effect: Traders use new assets to hedge their bets on existing assets, which in turn enables them to place larger bets and take on greater risks. This effect suggests that financial innovation is more likely to be destabilizing in more complete financial markets and when it concerns derivative assets.

In a dynamic setting, financial innovation always reduces the average variance in the long run because traders learn from past asset payoffs. A question emerges as to how new assets should be introduced to minimize their short run impact on the speculative variance. I show that staggering (or delaying) the introduction of new assets is not effective because it reduces traders' learning simultaneously with their speculation. A viable alternative is to set temporary position limits (or taxes) on new assets.

Keywords: financial innovation, heterogeneous beliefs, risk sharing, speculation, hedge-more/bet-more, uninsurable variance, speculative variance, financial instability, Bayesian learning.

JEL Classification: G1, G11, D53.

Preliminary and Incomplete.

*Harvard University (e-mail: asimsek@fas.harvard.edu). I am grateful to Daron Acemoglu, Malcolm Baker, Bengt Holmstrom, Andrei Shleifer, Jeremy Stein, and the seminar participants at Harvard University, MIT, Penn State University, University of Houston, University of Maryland, University of Wisconsin-Madison, Central European University, Yale University for helpful comments. All remaining errors are mine.

1 Introduction

According to the traditional view of financial innovation, new financial assets facilitate the diversification and the sharing of risks.¹ However, this view does not take into account that new assets are often associated with much uncertainty, especially because they do not have a long track record. Belief disagreements come as a natural by-product of this uncertainty and change the implications of risk taking in these markets. In particular, market participants' disagreements about how to value new assets naturally lead to *speculation*, which represents a powerful economic force that tends to increase risks.

An example is offered by the recent crisis. Assets backed by pools of subprime mortgages (e.g., subprime CDOs) became highly popular in the run-up to the crisis. One role of these assets is to allocate the risks to market participants who are best able to bear them. The safer tranches are held by investors that are looking for safety (or liquidity), while the riskier tranches are held by financial institutions who are willing to hold these risks at some price. While these assets (and their CDSs) should have served a stabilizing role in theory, they became a major trigger of the crisis in practice, when a fraction of financial institutions realized losses from their positions. Importantly, the same set of assets also generated considerable profits for a different fraction of market participants,² which suggests that at least some of the trades on these assets were speculative. Moreover, it might be no coincidence that the speculative trades were observed on these relatively new assets. Historical observations, which were scarce for these assets, might be necessary for market participants to converge to consensus (and correct) pricing models.

The recent episode is by no means unique in the history of financial innovation. Gennaioli, Shleifer, and Vishny (2010) describe various cases in which market participants have neglected the risks associated with new financial products. To the extent that participants do not neglect risks at the same level, these episodes also naturally generate speculation. What becomes of the risk sharing role of new assets when market participants disagree about their risks or returns?

I investigate this question in the context of a standard CARA-Normal economy in which traders have both risk sharing and speculative motives for trade. Traders with CARA utilities have endowment risks. They can hedge some of these risks by taking positions in financial assets, which provides the risk sharing motive for trade. On the other hand, traders also have heterogeneous prior beliefs for asset payoffs. They can take positions in assets to bet on their beliefs, which provides the speculation motive for trade. I define the *average variance* of traders'

¹Cochrane (2005) summarizes this view as follows: "Better risk sharing is much of the force behind financial innovation. Many successful new securities can be understood as devices to more widely share risks."

²Lewis (2010) provides a detailed description of investors that took a short position on housing related assets in the run-up to the recent crisis.

net worth as a stability measure for this economy.³ I then decompose the average variance into two components: an *uninsurable variance*, defined as the variance that would obtain if there were no belief disagreements, and a *speculative variance*, defined as the residual amount of variance that results from speculative trades based on belief disagreements. Intuitively, the uninsurable variance captures the stability measure in the ideal scenario in which assets are used only for risk sharing purposes, while the speculative variance captures the deviations from the ideal scenario in view of belief disagreements.

My main result characterizes the effect of financial innovation on each component of average variance. In line with the traditional view, financial innovation always decreases the uninsurable variance because new assets increase the possibilities for risk sharing. Theorem 1 shows that financial innovation also *always* increases the speculative variance. Moreover, there exist economies in which this increase in the speculative variance is sufficiently strong that financial innovation increases the average variance (by an arbitrary amount). This result formalizes the sense in which financial innovation can be destabilizing.

My analysis identifies two main channels by which financial innovation increases the speculative variance. First, new assets lead to new disagreements because they are associated with new uncertainties. Second, and perhaps more importantly, new assets also amplify speculation on existing disagreements. To illustrate the second channel, Theorem 1 shows that new assets increase the speculative variance even if traders completely agree about their payoffs. This result is rather surprising because traders use new assets to hedge some of the speculative risks they have been undertaking from their bets on existing assets. In view of this direct hedging effect, one could expect new assets (on which there is complete agreement) to reduce the speculative variance. This view does not take into account a powerful amplification mechanism, the *hedge-more/bet-more* effect.

To illustrate this effect, consider the following example. Suppose two traders have differing views about the currency, the Swiss Franc, which is highly correlated with the Euro. The optimist believes the Franc will appreciate while the pessimist believes it will depreciate. Traders do not disagree about the Euro, e.g., they disagree about the prospects of the Swiss economy but not about the Euro zone. First suppose traders can only take positions on the Franc and not the Euro. In this case, traders' positions in the Franc will be determined by a standard risk-return trade-off. Traders may not take too large positions on the Franc especially because the Franc is affected by multiple sources of shocks, e.g., those that affect the Swiss economy as well as the shocks to the Euro zone. To bet on their belief differences, traders must bear

³More specifically, I consider the variance of a trader's net worth (calculated according to her own beliefs), and I take the average of this expression over all traders.

all of these risks, which might make them reluctant to take large positions. Suppose instead the Euro is also introduced for trade (which can be interpreted as “financial innovation” in this example). In this case, traders will complement their positions in the Franc by taking the opposite positions in the Euro. This is because the complementary positions enable traders to hedge the risks that also affect the Euro, which they don’t disagree about, and to take purer bets on the Franc. With purer bets, traders bear less risk for each unit position on the Franc, which in turn enables them to take larger positions. Put differently, when traders are able to hedge more, they are induced to bet more. Theorem 1 shows that the hedge-more/bet-more effect is sufficiently strong that the introduction of the Euro in this example (and more generally, a new asset) increases the speculative variance.

While Theorem 1 shows that financial innovation may increase the average variance (if it increases the speculative variance by a sufficient margin), it is silent on the economic environments in which this is more likely. To address this question, I consider the comparative statics of the change in speculative variance with respect to the properties of existing and new assets. I isolate two effects by considering two extreme settings. First, suppose traders completely agree about the payoffs of existing assets (so that financial innovation increases the speculative variance only through the first channel -by generating new disagreements). In this setting, Theorem 2 establishes that financial innovation increases the speculative variance more in economies with more complete financial markets than less complete markets. Intuitively, the hedge-more/bet-more effect not only enables new assets to amplify speculation on old assets, but also enables old assets to amplify speculation on new assets. In more complete financial markets, traders are able to take purer bets on new assets, which leads to a greater increase in the speculative variance. As an alternative, suppose traders completely agree about the payoffs of new assets (so that financial innovation increases the speculative variance only through the second channel -by amplifying existing disagreements). In this setting, Theorem 3 establishes that financial innovation increases the speculative variance more when new assets are correlated with existing assets -as would be the case with new derivative assets. Intuitively, the hedge-more/bet-more effect is operational only if new and old assets are correlated. While obtained in extreme settings, Theorems 2 and 3 create a presumption that financial innovation is more likely to be destabilizing in more complete financial markets and when it concerns derivative assets.

In practice, new financial products are usually designed by financial institutions with profit incentives. It is thus natural to ask whether new assets that will increase the average variance will be *endogenously* innovated in this economy. I address this question by introducing a

profit seeking market maker that innovates new assets for which it subsequently serves as the intermediary. The market maker's expected profits are proportional to traders' perceived surplus from trading new assets. Thus, traders' speculative trading motive, as well as their risk sharing motive, creates innovation incentives for the market maker. To illustrate this point, Theorem 4 identifies an extreme case in which the market maker will innovate assets that *maximize* the average variance among all possible choices -completely disregarding the risk sharing motive for innovation.

While the static model is useful to analyze the speculation generated by a given level of belief disagreements, a natural question is how these disagreements evolve over time. I consider this question in a dynamic extension in which traders update their beliefs by observing past asset payoffs. In this dynamic setting, traders' belief disagreements disappear in the long run. Consequently, new assets always reduce the long run average variance because they provide their risk sharing services without being subject to speculation. Importantly, traders beliefs converge to each other at a faster rate, t (where t is the number of observations for asset payoffs), than the rate, \sqrt{t} , at which they converge to the true distribution of asset payoffs. This is because traders update their beliefs from the same set of observations: Even if those observations do not yet reflect the true distribution, they facilitate convergence of traders' beliefs.

These results suggest that speculation on new assets is likely to be most potent in the short run following their introduction. A question emerges as to how new assets should be introduced to reduce their short run impact on speculation. A plausible idea is to stagger the introduction of assets, so as to introduce them "one at a time." This policy ensures that traders learn about the payoffs of some new assets before they get to trade other new assets. In view of the hedge-more/bet-more effect, this policy reduces traders' speculation on assets that are introduced earlier. It would then appear that an appropriately designed staggering policy could reduce speculation.

This argument is incomplete as it misses the fact that the introduction of new assets also facilitates learning. In particular, traders arguably learn about the payoffs of new assets to the extent that they trade them. Under this assumption, a staggering policy reduces the amount of learning from assets that are introduced later. As traders learn less, they tend to speculate more in subsequent periods. In fact, under mild assumptions, Theorem 6 shows that the staggering policy has no net effect on an appropriate measure of speculative variance. This result applies regardless of assets' correlation structure, suggesting that a staggering policy is unlikely to be an effective tool to reduce speculation. A viable alternative is temporary position limits (or taxes) on new assets, which reduces traders' speculative positions without

precluding learning.

My paper contributes to a recent literature that analyzes the role of financial innovation in generating financial instability. Rajan (2005) and Calomiris (2008) emphasize the effect of financial innovation on agency problems. Gennaioli, Shleifer and Vishny (2010) investigate the neglected risks associated with new financial assets. My paper identifies speculation as additional channel through which new assets may decrease stability. The idea that speculation may create financial instability appears also in Stiglitz (1989), Summers and Summers (1991), and Stout (1995). However, these analyses are mostly informal and they do not derive any results similar to my theorems.

In related work, Brock, Hommes and Wagener (2009) identify another source of instability brought about by financial innovation and the hedge-more/bet-more effect. Their main ingredient is reinforcement learning: That is, they assume traders choose their beliefs according to a fitness measure, such as past profits made by the belief. This ingredient implies that the steady-state corresponding to the fundamental asset price can be dynamically unstable. Brock et al. (2009) show that the introduction of new assets increases the range of parameters for which the steady-state is unstable. In view of the hedge-more/bet-more effect, new assets enable traders to take greater positions on their beliefs. Consequently, a belief that turns out to be correct yields a greater profit, and it is chosen by a greater number of traders in the next period. This in turn makes the steady-state more likely to be dynamically unstable. My paper complements Brock et al. (2009) by showing that financial innovation also leads to a different notion of instability: an increase in the riskiness of traders' net worths. Moreover, this result does not require the reinforcement learning ingredient.

My paper is also related to a sizeable literature on security design and financial innovation, e.g., Van Horne (1985), Miller (1986), Duffie and Jackson (1989), Cuny (1993), Allen and Gale (1994), Duffie and Rahi (1994), Demange and Laroque (1995), Athanasoulis and Shiller (2000, 2001), Tufano (2003). To my knowledge, this literature has not explored the implications of heterogenous beliefs for security design. For example, in their survey of the literature, Duffie and Rahi (1994) note that "one theme of the literature, going back at least to Working (1953) and evident in the Milgrom and Stokey (1982) no-trade theorem, is that an exchange would rarely find it attractive to introduce a security whose sole justification is the opportunity for speculation." From an empirical point of view, this observation is hard to square with speculation on new securities in the run-up to the recent crisis. From a theoretical point of view, Theorem 4 shows that the observation does not apply if traders have heterogeneous priors rather than heterogeneous information. The observation also does not apply if traders have heterogeneous information but security prices do not reveal information fully due to

the presence of noise traders. The analogues of my results can be derived for this alternative setting. The important economic ingredient is that traders continue to have some disagreement after observing asset prices.

Another strand of the literature has analyzed the implications of heterogenous beliefs for asset prices but not for risk sharing.⁴ My model is closest to Lintner (1969), who generalizes CAPM to a model in which beliefs and risk aversion coefficients are heterogeneous. However, Lintner (1969) does not consider the implications of this theory for the riskiness of allocations.

The rest of the paper is organized as follows. Section 2 introduces the basic environment. This section also uses simple examples to illustrate the two channels by which new assets increase the riskiness of traders' net worths. Section 3 completes the description of the environment and characterizes the equilibrium. This section defines the average variance, and decomposes it into the uninsurable and the speculative variances. Section 4 presents the main result which characterizes the effect of financial innovation on the two components of average variance. Section 5 provides the comparative statics for the speculative variance. Section 6 considers endogenous financial innovation. Section 7 presents the dynamic model and Section 8 concludes. Appendix A contains proofs that are omitted from the main text.

2 Basic Environment and Main Channels

Consider an economy with two dates, $\{0, 1\}$, and a single consumption good (dollar). There are a finite number of traders denoted by $i \in I = \{1, 2, \dots, |I|\}$. Traders are endowed with e dollars at date 0, which is constant, and w_i dollars at date 1, which is a Normally distributed random variable. Traders only consume at date 1, and they can transfer resources to date 1 by investing in one of two ways. They can invest in a storage technology that yields one dollar for each dollar invested. Alternatively, they can invest in risky assets denoted by $j \in J = \{1, \dots, |J|\}$. Asset j is in fixed supply, normalized to zero, and it pays a^j dollars at date 1, which is a Normally distributed random variable.⁵ At date 0, the asset is traded in a competitive market at price p^j . The assets' payoff and price vectors are respectively denoted by $\mathbf{a} = (a^1, \dots, a^{|J|})$ and $\mathbf{p} = (p^1, \dots, p^{|J|})$. This vector notation will be used throughout the paper.

Trader i 's position in asset j is denoted by x_i^j . Given the price vector \mathbf{p} , the trader chooses an asset portfolio, \mathbf{x}_i , and invests the rest of her budget, $e - \mathbf{x}_i' \mathbf{p}$, in the storage technology.

⁴A very incomplete list includes Miller (1977), Harrison and Kreps (1978), Varian (1985,1989), Harris and Raviv (1993), Allen, Morris, Postlewaite (1993), Kandel and Pearson (1995), Zapatero (1998), Chen, Hong and Stein (2003), Scheinkman and Xiong (2003), Geanakoplos (2009), Cao (2010), Simsek (2011).

⁵The normalization of the asset supply to zero is without loss of generality because traders can short sell the asset without frictions. Similarly, asset payoffs can take negative values because the environment is frictionless. In particular, there is no limited liability (negative net worth is allowed) and repayment is enforced by contracts.

With these investment decisions, her net worth at date 1 is a random variable given by:

$$n_i = e - \mathbf{x}'_i \mathbf{p} + w_i + \mathbf{x}'_i \mathbf{a}. \quad (1)$$

Trader i 's preferences are captured by a CARA utility function over net worth at date 1. In particular, she chooses her portfolio to maximize:

$$E_i [-\exp(-\theta_i n_i)],$$

where θ_i denotes her coefficient of absolute risk aversion. Given the CARA-Normal setting, her optimization reduces to the usual mean-variance problem:

$$\max_{\mathbf{x}_i} E_i [n_i] - \frac{\theta_i}{2} \text{var}_i [n_i]. \quad (2)$$

Here, $E_i [\cdot]$ denotes the mean and $\text{var}_i [\cdot]$ denotes the variance according to trader i 's beliefs.

The equilibrium in this economy is a collection of asset prices, \mathbf{p} , and allocations, $(\mathbf{x}_1, \dots, \mathbf{x}_{|I|})$, such that each trader i chooses her portfolio optimally and prices clear asset markets, that is,

$$\sum_i x_i^j = 0 \text{ for each } j \in J. \quad (3)$$

I will capture financial innovation in this economy as an expansion of the set of traded assets, J . The main goal of this paper is to characterize the effects of financial innovation on a measure of financial stability (which will be formally defined in the next section). Before I turn to the general characterization, I use two simple examples that respectively illustrate the two channels by which financial innovation increases the riskiness of traders' net worths.

Example 1: New assets generate new disagreements

This example illustrates that new assets generate speculation when traders disagree about their payoffs. Moreover, this speculation can be sufficiently strong that financial innovation makes traders' net worths riskier (despite the fact that new assets also provide their traditional risk sharing benefits). To illustrate these points, I first consider two benchmarks: A no-asset benchmark and a common-beliefs benchmark. I then consider the setting with belief disagreements and compare it with the benchmarks.

Suppose there are two traders with the same coefficient of risk aversion, i.e., $I = \{1, 2\}$ and $\theta_1 = \theta_2 \equiv \theta$. Traders' date 1 endowments are perfectly negatively correlated. In particular,

let $v \sim N(0, 1)$ denote a standard normal random variable and suppose:

$$w_1 = v \text{ and } w_2 \equiv -v. \quad (4)$$

First consider the case in which there are no assets, i.e., $J = \emptyset$. In this *no-asset benchmark*, traders' date 1 net worths are given by

$$n_1 = e + v \text{ and } n_2 = e - v. \quad (5)$$

Traders' net worths are risky because they are unable to hedge their endowment risks.

Next suppose a new asset is introduced with payoff:

$$a^1 = v. \quad (6)$$

Suppose traders have common (not necessarily correct) beliefs about the asset payoff. In particular, suppose traders' common belief for the random variable v is denoted by the distribution $N(\bar{\mu}, 1 + \frac{1}{\tau})$. Intuitively, the parameter τ capture the precision of traders' prior beliefs. In this *common-beliefs benchmark*, traders' equilibrium portfolios are given by:

$$x_1^1 = -1 \text{ and } x_2^1 = 1 \quad (7)$$

(and the equilibrium price is $p^1 = \bar{\mu}$). Traders' net worths are constant and given by

$$n_1 = n_2 = e. \quad (8)$$

In this benchmark (with common beliefs), financial innovation enables traders to hedge and diversify their idiosyncratic risks.

Next suppose traders have *heterogeneous prior beliefs* for the payoff of the new asset. In particular, trader 1's prior belief for the random variable v is denoted by $N(\bar{\mu} + \varepsilon, 1 + \frac{1}{\tau})$, while trader 2's prior belief is captured by $N(\bar{\mu} - \varepsilon, 1 + \frac{1}{\tau})$, for some $\varepsilon > 0$. The parameter ε captures the level of belief disagreements. Note that trader 1 is optimistic about the asset payoff while trader 2 is pessimistic. In this different-beliefs setting, the equilibrium portfolios are given by:

$$x_1^1 = -1 + \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} \text{ and } x_2^1 = 1 - \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1}. \quad (9)$$

Note that traders' positions are influenced by their hedging demands, as in (7), as well as their belief differences. Note also that traders take greater speculative positions when they are more optimistic, when they are less risk averse, and when they have more precise prior beliefs.

Traders' net worths are given by:

$$n_1 = e + \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} v \text{ and } n_2 = e - \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} v. \quad (10)$$

Comparing the net worth expressions (10) to their counterparts in the earlier benchmarks illustrates two points. First, unlike the common-beliefs benchmark [cf. Eq. (8)], traders' net worths are not constant. That is, a new asset (that is fully correlated with traders' endowments) fails to facilitate full diversification. This is because the speculative motive for trade (driven by belief differences) influences traders' demands and distorts risk sharing.

Second, if $\varepsilon > \theta \frac{\tau+1}{\tau}$, then traders' net worths are even riskier than the no-asset benchmark [cf. Eq. (5)]. In this case, trader 1 is so optimistic about the asset's payoff that she takes a positive position, despite the fact that her endowment is negatively correlated with the asset payoff. Consequently, the new asset increases the riskiness of her net worth. Hence, when traders' disagreements about new assets are sufficiently large, financial innovation makes traders' net worths riskier.

Example 2: New assets amplify speculation on existing disagreements

This example illustrates the hedge-more/bet-more effect, by formalizing the Franc-Euro setting described in the introduction. It also shows that financial innovation increases the riskiness of traders' net worths through a second channel: By amplifying traders' speculations on existing disagreements. To this end, I first consider a benchmark in which traders bet on an existing asset (Franc). I then consider the innovation of a second asset (Euro) which is correlated with the first one and about which traders agree.

As in Example 1, suppose there are two traders with the same coefficient of risk aversion, i.e., $I = \{1, 2\}$ and $\theta_1 = \theta_2 \equiv \theta$. But this time suppose traders' date 1 endowments are constant, $w_1 = w_2 = 0$, which ensures that there is no risk sharing motive for trade. Suppose also that the fundamental risk in this economy is captured by 2 random variables, v_1, v_2 , which are i.i.d. with standard Normal distribution $N(0, 1)$. Traders have common beliefs for v_2 given by $N(\bar{\mu}, 1 + \frac{1}{\tau})$. However, they disagree about v_1 . In particular, trader 1's belief for v_1 is given by $N(\bar{\mu} + \varepsilon, 1 + \frac{1}{\tau})$ while trader 2's belief is given $N(\bar{\mu} - \varepsilon, 1 + \frac{1}{\tau})$.

First suppose there is a single asset for trade, with payoff:

$$a^1 = v_1 + \alpha v_2.$$

In particular, the asset is affected by both sources of fundamental risk, with the weight, α , capturing the relative importance of the second risk. In this case, by symmetry, the equilibrium

price is given by $p^1 = \bar{\mu}$. Substituting this expression, trader 1's mean-variance problem [cf. Eq. (2)] can be written as:

$$\max_{x_1^1} x_1^1 \varepsilon - \frac{\theta}{2} (1 + \alpha^2) \left(1 + \frac{1}{\tau}\right) (x_1^1)^2. \quad (11)$$

The first term in this expression is the trader's expected payoff in equilibrium. The second term is the trader's expected cost from variance of her net worth. Trader 2's mean-variance problem takes a similar form. Traders' equilibrium portfolios can be solved as:

$$x_1^1 = \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} \frac{1}{1 + \alpha^2} \quad \text{and} \quad x_2^1 = -\frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} \frac{1}{1 + \alpha^2}, \quad (12)$$

and their net worths are given by:

$$n_1 = e + \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} \frac{v_1 + \alpha v_2}{1 + \alpha^2} \quad \text{and} \quad n_2 = e - \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} \frac{v_1 + \alpha v_2}{1 + \alpha^2}. \quad (13)$$

Note that traders do not necessarily bet "too much" on their disagreements, because the asset's payoff is influenced by v_2 as well as v_1 . In particular, traders' speculative positions (and the variances of their net worths) are decreasing in α , and they converge to zero as α limits to infinity. Intuitively, the ability to trade asset 1 enables traders to take only *impure bets* because the asset's payoff also responds to risks traders do not disagree about. To bet on their belief disagreements, traders must also hold these additional risks [as formally captured by the α term in problem (11)], which makes them reluctant to take large positions.

Next suppose a new asset is introduced with payoff:

$$a^2 = v_2.$$

In this case, straightforward calculations show that traders' equilibrium portfolios are given by:

$$\text{Trader 1} \quad : \quad x_1^1 = \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1}, \quad x_1^2 = -\alpha \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1}, \quad (14)$$

$$\text{Trader 2} \quad : \quad x_2^1 = -\frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1}, \quad x_2^2 = \alpha \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1},$$

and their net worths are given by:

$$n_1 = e + \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} v_1 \quad \text{and} \quad n_2 = e - \frac{\varepsilon}{\theta} \frac{\tau}{\tau + 1} v_1. \quad (15)$$

Note that the magnitude of traders' positions on asset 1 is greater than the earlier setting in

which asset 2 was not available [cf. Eqs. (14) and (12)]. Importantly, the variances of their net worths are also greater [cf. Eqs. (15) and (13)]. This shows that the innovation of asset 2, about which traders do not disagree, enables traders to take greater speculative positions on asset 1 and makes traders' net worths riskier.

To understand these results, first consider the portfolios in (14). Note that traders complement their speculative positions in asset 1 by taking the opposite positions in asset 2. These complementary positions enable the traders to hedge out the risk, v_2 , which they do not disagree about. This in turn enables traders to take *purer bets* on the risk, v_1 . In fact, traders' net worths (15) are identical to those that would obtain if they could trade an alternative asset that pays:

$$a^{syn} = a^1 - \alpha a^2 = v_1. \quad (16)$$

Traders “create” this synthetic asset by simultaneously investing in multiple assets.

Next consider why traders increase their positions on asset 1 and why their net worths become riskier. To understand these effects, it is useful to consider the analogue of the mean-variance problem (11) in this case. In terms of the synthetic asset in (16), trader 1 solves:

$$\max_{x_1^{syn}} x_1^{syn} \varepsilon - \frac{\theta}{2} \left(1 + \frac{1}{\tau} \right) (x_1^{syn})^2. \quad (17)$$

Note that problems (11) and (17) are very similar, except that the former problem has an additional factor of $(1 + \alpha^2)$ multiplying the cost term. This difference captures the *hedging effect*: The fact that traders use new assets to hedge the risks on their speculative positions tends to reduce the riskiness of their net worths (and the associated costs). In fact, controlling for a trader's speculative position on risk v_1 , i.e., assuming $x_1^{syn} = x_1^1$, the hedging effect always leads to lower variance of net worth. In view of this observation, a naive view could suggest that new assets (on which there is complete agreement) should reduce the riskiness of traders' net worths.

However, the naive view misses an important amplification mechanism: the *hedge-more/bet-more effect*. When traders are able to take purer bets, they also take larger bets. In this example, the “marginal cost” (in terms of additional variance) of taking an additional speculative position on v_1 is lower when the second asset is available. This induces the trader to take a larger speculative position in that case, i.e., $x_1^1 < x_1^{syn}$ for the respective optima of problems (11) and (17).

The amplification of speculative positions (through the hedge-more/bet-more effect) tends to increase the riskiness of traders' net worths. Recall also that the direct hedging effect tends to reduce the riskiness of traders' net worths. A priori, it is not clear that the amplification

effect should be sufficiently strong to overcome the direct hedging effect. However, this is always the case for a CARA-Normal economy. In particular, since the mean-variance problems (11) and (17) are linear-quadratic, their cost terms (calculated at their respective optima) satisfy:

$$(1 + \alpha^2) \frac{\theta}{2} \left(1 + \frac{1}{\tau}\right) (x_1^1)^2 < \frac{\theta}{2} \left(1 + \frac{1}{\tau}\right) (x_1^{syn})^2.$$

That is, the reduction in marginal cost (in terms of variance) generates such a large portfolio reaction that the trader's total costs (and their net worth variances) increase.

3 Environment and Equilibrium

The examples in the previous section illustrated two main channels by which financial innovation tends to increase the riskiness of traders' net worths. As illustrated by Example 1, there is also the traditional channel (risk sharing and diversification with common beliefs) by which financial innovation tends to decrease the riskiness of net worths. The rest of the paper analyzes a general CARA-Normal economy, in which all the channels of the previous section are operational. This section characterizes the equilibrium. It also defines a measure for the riskiness of traders' net worths, and decomposes the measure into two components (which loosely corresponds to speculation and risk sharing motives for trade). The next section uses this decomposition to present the main result, which characterizes the effect of financial innovation on each component of the riskiness measure.

The economy in this section builds upon the environment in Section 2 by providing a specification for agents' beliefs. In particular, the uncertainty in this economy is captured by the m dimensional random variable, $\mathbf{v} = (v_1, \dots, v_m)'$, which has the multivariate standard Normal distribution $N(\mathbf{0}, \mathbf{I})$. Traders' date 1 endowments can be written in terms of \mathbf{v} as:

$$w_i = (\mathbf{W}_i)' \mathbf{v},$$

for some $\mathbf{W}_i \in \mathbb{R}^m$. Asset j 's payoff can also be written in terms of \mathbf{v} as:

$$a^j = (\mathbf{A}^j)' \mathbf{v},$$

for some $\mathbf{A}^j \in \mathbb{R}^m$. The vectors, $\{\mathbf{A}^j\}_j$, are linearly independent, which ensures that the assets are not redundant. Traders have potentially heterogeneous prior beliefs about the random variable, \mathbf{v} . To simplify the analysis, I assume that traders know that \mathbf{v} is Normally distributed and that they agree about the variance of \mathbf{v} (but they may disagree about the mean of \mathbf{v}). In particular:

Assumption (A1). Trader i 's belief for the random variable \mathbf{v} is given by $N(\boldsymbol{\mu}_i^{\mathbf{v}}, \Lambda^{\mathbf{v}})$, where $\boldsymbol{\mu}_i^{\mathbf{v}} \in \mathbb{R}^m$ and $\Lambda^{\mathbf{v}}$ is an $m \times m$ covariance matrix with full row rank.

The assumption that $\Lambda^{\mathbf{v}}$ has full row rank ensures that traders have finite demands. This completes the specification of traders' beliefs. An economy is formally denoted by $\mathcal{E}(J) = \left(J; \{\mathbf{A}^j\}_{j \in J}; \{\mathbf{W}_i, \boldsymbol{\mu}_i^{\mathbf{v}}, \Lambda^{\mathbf{v}}\}_i \right)$.

Note also that trader i believes the asset payoffs are Normally distributed, $N(\boldsymbol{\mu}_i, \Lambda)$, with:

$$\boldsymbol{\mu}_i \equiv \mathbf{A}'\boldsymbol{\mu}_i^{\mathbf{v}} \text{ and } \Lambda \equiv \mathbf{A}'\Lambda^{\mathbf{v}}\mathbf{A}.$$

As with the underlying uncertainty, traders agree about the variance of asset payoffs while they may disagree about their mean. In addition, trader i believes that her endowment is Normally distributed, and that the covariance of her endowment with the asset payoffs is given by $\boldsymbol{\lambda}_i = \mathbf{A}'\Lambda^{\mathbf{v}}\mathbf{W}_i$. In this economy, traders' beliefs for asset payoffs, $\{N(\boldsymbol{\mu}_i, \Lambda)\}_i$, and the covariance terms, $\{\boldsymbol{\lambda}_i\}_i$, provide a sufficient statistic for equilibrium.

3.1 Characterization of Equilibrium

From the mean variance problem, (2), trader i 's portfolio demand can be solved as:

$$\mathbf{x}_i = \Lambda^{-1} \left[\frac{\boldsymbol{\mu}_i - \mathbf{p}}{\theta_i} - \boldsymbol{\lambda}_i \right]. \quad (18)$$

This expression illustrates the two sources of demand. First, the trader tends to hold a long position on an asset to the extent that its payoff is negatively correlated with her endowment, as captured by the term, $\boldsymbol{\lambda}_i$. Second, she also tends to hold a long position to the extent that she thinks the asset is underpriced, as captured by the term, $\frac{\boldsymbol{\mu}_i - \mathbf{p}}{\theta_i}$.

Next consider the determination of the equilibrium price vector, \mathbf{p} . Aggregating Eq. (18) and using market clearing, prices can be solved in closed form as:

$$\mathbf{p} = \frac{1}{|I|} \sum_{i \in I} \left(\frac{\bar{\theta}}{\theta_i} \boldsymbol{\mu}_i - \bar{\theta} \boldsymbol{\lambda}_i \right), \quad (19)$$

where $\bar{\theta} \equiv \left(\sum_{i \in I} \theta^{-1} / |I| \right)^{-1}$ is the Harmonic mean of traders' absolute risk aversion coefficients. Intuitively, the price of an asset is high either if the asset is negatively correlated with traders' endowments (captured by the $\boldsymbol{\lambda}_i$ term) or if traders on average believe that the asset will yield a high payoff (captured by the $\boldsymbol{\mu}_i$ term). The beliefs of more risk averse traders have a smaller effect on the price since they bet relatively less on their opinions.

Using the price expression (19), trader i 's portfolio in (18) can also be solved in closed form. In addition, trader i 's portfolio can be decomposed into two components that capture

the two motives for trade in this economy. In particular, $\mathbf{x}_i = \mathbf{x}_i^R + \mathbf{x}_i^S$ where

$$\mathbf{x}_i^R = -\Lambda^{-1}\tilde{\boldsymbol{\lambda}}_i \text{ and } \mathbf{x}_i^S = \Lambda^{-1}\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i}. \quad (20)$$

Here,

$$\tilde{\boldsymbol{\lambda}}_i = \boldsymbol{\lambda}_i - \frac{\bar{\theta}}{\theta_i} \frac{1}{|I|} \sum_{i \in I} \boldsymbol{\lambda}_i$$

denotes the relative covariance of trader i 's endowment, and

$$\tilde{\boldsymbol{\mu}}_i = \boldsymbol{\mu}_i - \frac{1}{|I|} \sum_{i \in I} \frac{\bar{\theta}}{\theta_i} \boldsymbol{\mu}_i \quad (21)$$

denotes the relative optimism of trader i . Note that $\{\mathbf{x}_i^R\}_i$ would be the equilibrium allocation if there were no belief differences (i.e., if $\tilde{\boldsymbol{\mu}}_i = 0$ for each i). Hence, I refer to \mathbf{x}_i^R as the *risk sharing portfolio* of trader i . On the other hand, the residual allocations, $\{\mathbf{x}_i^S\}_i$, are purely driven by belief differences. Thus, I refer to \mathbf{x}_i^S as the *speculative portfolio* of trader i .

The main goal of this paper is to analyze the effect of financial innovation on the riskiness of traders' net worths. To this end, I first formalize the notion of riskiness. In particular, I define the *average variance* of traders' net worths as:⁶

$$\begin{aligned} \Omega &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \text{var}_i(n_i), \\ &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^{\mathbf{v}} \mathbf{W}_i + \mathbf{x}'_i \Lambda \mathbf{x}_i + 2\mathbf{x}'_i \boldsymbol{\lambda}_i). \end{aligned} \quad (22)$$

Note that the average variance is a summary measure that assigns a greater weight to traders that are more risk averse. While this is reasonable, the exact form (linearity) of the weights might seem arbitrary. The following lemma justifies the choice of weights. It shows that traders' risk sharing portfolios, $\{\mathbf{x}_i^R\}_i$, minimize the average variance expression in (22) among all feasible portfolios.

Lemma 1. *The risk sharing portfolios, $\{\mathbf{x}_i^R\}_i$, solve the following problem:*

$$\min_{\{\mathbf{x}_i \in \mathbb{R}^{|J|}\}_i} \Omega \quad \text{s.t.} \quad \sum_i \mathbf{x}_i = 0. \quad (23)$$

This result shows that the average variance, Ω , provides a natural measure for the riskiness of net worths. In particular, this measure would be minimized (given the available set of

⁶The fact that variances are calculated according to traders' own beliefs is without loss of generality in view of Assumption 1. In particular, since all traders agree about the covariance matrix of \mathbf{v} , the average variance calculated according to an arbitrary trader's belief would be the same as the expression in (22).

assets) had there been only the risk sharing motive for trade. Consequently, I let Ω^R denote the minimum for problem (23) and refer to it as the *uninsurable variance*. The extent to which Ω deviates from Ω^R could be viewed as the effect of speculation. I thus define $\Omega^S = \Omega - \Omega^R$ and refer to it as the *speculative variance*. The next lemma characterizes the two components of average variance, $\Omega = \Omega^R + \Omega^S$, in terms of the exogenous parameters of the model.

Lemma 2. *The uninsurable variance is given by:*

$$\Omega^R = \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\theta} \left(\mathbf{w}_i' \Lambda^{\mathbf{v}} \mathbf{w}_i - \tilde{\boldsymbol{\lambda}}_i' \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i \right), \quad (24)$$

and the speculative variance is given by:

$$\Omega^S \equiv \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\theta} \left(\frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \Lambda^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} \right). \quad (25)$$

The forms of Ω^R and Ω^S are intuitive. Eq. (24) illustrates that, loosely speaking, the uninsurable variance will be lower to the extent that the assets provide better risk sharing opportunities (captured by $\tilde{\boldsymbol{\lambda}}_i$). Similarly, Eq. (25) illustrates that the speculative variance will be greater to the extent that the assets feature greater belief disagreements (captured by $\tilde{\boldsymbol{\mu}}_i$). The next sections characterize the effect of financial innovation on Ω^R and Ω^S .

4 Effect of Financial Innovation on Average Variance

I model financial innovation as an expansion of the set of traded assets. To this end, it will be useful to define economies in which only a subset of the assets in J are traded. In particular, given $\hat{J} \subset J$, let $\mathcal{E}(\hat{J}) = \left(\hat{J}; \{\mathbf{A}^j\}_{j \in \hat{J}}; \{\mathbf{W}_i, \boldsymbol{\mu}_i^{\mathbf{v}}, \Lambda^{\mathbf{v}}\}_i \right)$ denote the economy in which the asset set is given by \hat{J} . Where it does not create confusion, I also use the notation, $z(\hat{J})$, to refer to the equilibrium variable z for the economy $\mathcal{E}(\hat{J})$. To capture financial innovation, suppose J can be broken down into a set of old assets, J_O , and a set of new assets, J_N (formally, $J = J_O \cup J_N$ where J_O and J_N are disjoint sets). The differences between the economies $\mathcal{E}(J_O)$ and $\mathcal{E}(J_O \cup J_N)$ can be attributed to financial innovation. I next present the main result.

Theorem 1. *Consider the average variance and its components respectively for the economies $\mathcal{E}(J_O)$ and $\mathcal{E}(J_O \cup J_N)$.*

(i) *Financial innovation always reduces the uninsurable variance, that is:*

$$\Omega^R(J_O \cup J_N) \leq \Omega^R(J_O).$$

(ii) *Financial innovation always increases the speculative variance, that is:*

$$\Omega^S(J_O \cup J_N) \geq \Omega^S(J_O).$$

(ii) *For any $\eta > 0$, there exist economies in which the increase in the speculative variance is sufficiently large that financial innovation increases the average variance by at least η , that is, $\Omega(J_O \cup J_N) > \Omega(J_O) + \eta$.*

The first part of this theorem is a corollary of Lemma 1, and it shows that financial innovation always provides some risk sharing benefits. This part formalizes the traditional view of financial innovation in the context of this model. On the other hand, the second part of the theorem identifies a second force which always operates in the opposite direction. In particular, when there is belief heterogeneity, financial innovation also always increases the speculative variance. Hence, the net effect of financial innovation on average variance is ambiguous, and it depends on the relative strength of the two forces.

Most of the literature on financial innovation considers the special case without belief heterogeneity. Theorem 1 shows that the common-beliefs assumption is restrictive, as it shuts down an important economic force by which financial innovation always has a positive effect on the average variance. Furthermore, the third part of the theorem shows that it is indeed possible for this force from belief heterogeneity to dominate the traditional force.

It is also worth emphasizing the generality of the second part of Theorem 1. The result applies for all sets of existing and new assets, J_O and J_N , with no restrictions on the joint distribution of asset payoffs or traders' beliefs for \mathbf{v} [except for the relatively mild Assumption (A1)]. For example, Theorem 1 shows that financial innovation increases the speculative variance even if there is no belief disagreement about new assets (i.e., for each $j \in J_N$, μ_i^j is a constant independent of i). This is surprising because, as illustrated by Example 2, new assets are used to hedge (to some extent) the risks from traders' speculation on their existing disagreements. Put differently, the direct *hedging effect* of new assets tends to reduce Ω^S . However, as illustrated by Example 2, there is also the *hedge-more/bet-more effect* which tends to increase Ω^S . Theorem 1 shows that, in the standard CARA-Normal setting, the hedge-more/bet-more effect is sufficiently strong that new assets always increase Ω^S .

The rest of this section provides a sketch proof for the second part of Theorem 1 (the proofs for the first and the third parts relegated to the Appendix). To facilitate intuition, it is useful to think of the speculative portfolios as corresponding to equilibrium portfolios in an economy with no endowment risks, as formalized by the following lemma.

Lemma 3. Given an economy $\mathcal{E}(\hat{J})$, consider the hypothetical economy that is identical except that traders have no endowment risk, i.e., $\mathbf{W}_i = 0$ for all $i \in I$.

(i) The speculative portfolios in economy $\mathcal{E}(\hat{J})$ are identical to the portfolios in the hypothetical economy, i.e., $\mathbf{x}_i^S = \mathbf{x}_{i,H}$ for each i . Moreover, $\mathbf{x}_{i,H}$ solves:

$$\max_{\mathbf{x}_i \in \mathbb{R}^{|\hat{J}|}} (\tilde{\boldsymbol{\mu}}_i)' \mathbf{x}_i - \frac{\theta_i}{2} \mathbf{x}_i' \Lambda \mathbf{x}_i. \quad (26)$$

(ii) The speculative variance in economy $\mathcal{E}(\hat{J})$ is identical to the average variance in the hypothetical economy, i.e., $\Omega^S = \Omega_H$, which is given by:

$$\Omega_H = \frac{1}{|I|} \sum_i \frac{\theta_i}{\theta} \mathbf{x}_{i,H}' \Lambda \mathbf{x}_{i,H}. \quad (27)$$

In view of this lemma, it suffices to show that financial innovation always increases the average variance in the hypothetical economy. I next show this in four steps.

First, note that financial innovation “expands” the constraint set of problem (26). That is, when the asset set is $\hat{J} = J^O \cup J^N$, traders are able to make all the speculative trades they could make when the asset set is $\hat{J} = J^O$ (and more). Put differently, new assets expand the “betting possibilities frontier” for traders. This observation, which is central for the result, further reinforces the messages of Examples 1 and 2. In particular, financial innovation increases the betting possibilities frontier through two distinct channels. As emphasized by Example 1, new assets are likely to generate new disagreements. In addition, as emphasized by Example 2, even if new assets do not generate disagreements themselves, they enable the traders to take purer bets on existing disagreements. Both of these channels manifest themselves as an expansion of the constraint set of problem (26).

Second, as the constraint set of problem (26) expands, its optimum value also increases. Put differently, financial innovation increases traders’ expected (certainty-equivalent) payoffs, because it increases traders’ betting options.

Third, since problem (26) is a quadratic optimization problem, at the optimum allocation traders’ expected returns, $(\tilde{\boldsymbol{\mu}}_i)' \mathbf{x}_i$, are proportional to their variance costs, $\frac{\theta_i}{2} \mathbf{x}_i' \Lambda \mathbf{x}_i$. In particular:

$$(\tilde{\boldsymbol{\mu}}_i)' \mathbf{x}_{i,H} = 2 \frac{\theta_i}{2} \mathbf{x}_{i,H}' \Lambda \mathbf{x}_{i,H}. \quad (28)$$

For a basic intuition, consider a trader that is relatively less risk averse, i.e., with lower θ_i than average. Naturally, this trader holds an equilibrium portfolio with a higher expected return than the average trader. According to (28), the trader also takes greater risks (captured by a greater variance for net worth). This intuition is one of the central insights of CAPM: Traders

obtain greater expected returns by holding riskier portfolios. Eq. (28) shows that a similar intuition applies to other parameters of the model. In particular, consider a trader that is more optimistic (or pessimistic) than the average trader. This trader naturally holds an equilibrium portfolio with a higher (perceived) expected return than the average trader. According to (28), she also takes greater risks. At the optimal allocation, higher expected returns go hand-in-hand with higher risks.

The final step is obtained by combining the second and the third steps. Since financial innovation always increases certainty equivalent payoffs, it increases both the expected return and the variance costs in Eq. (28) for each trader. It follows that financial innovation increases the average variance in the hypothetical economy [cf. Eq. (27)]. Since $\Omega_H = \Omega^S$, financial innovation also increases the speculative variance in the original economy (see Appendix A for an alternative proof based on matrix algebra).

5 Comparative Statics of the Change in Speculative Variance

Theorem 1 characterizes the effect of financial innovation on the components of average variance, Ω^R and Ω^S . The result also shows that the net effect on average variance, Ω , (in particular, whether it decreases or increases) depends on the strengths of the effects on Ω^R and Ω^S . This section analyzes the comparative statics of the two effects, which are crucial to identify the economic environments in which financial innovation is likely to increase the average variance.

A number of straightforward conclusions can be drawn by inspecting the functional forms for Ω^S and Ω^R in (25) and (24). To give two examples, increasing traders' belief disagreements (by scaling up $\{\mu_i\}_i$ proportionally) or decreasing their risk aversion coefficients (by scaling down $\{\theta_i\}$) increase the effect, $\Omega^S(J_O \cup J_N) - \Omega^S(J_O)$, while they do not change the effect, $\Omega^R(J_O) - \Omega^R(J_O \cup J_N)$. Thus, these changes make financial innovation more likely to increase the average variance.

In addition to these results, the hedge-more/bet-more effect suggests that the change in speculative variance, $\Omega^S(J_O \cup J_N) - \Omega^S(J_O)$, should also depend on the properties of old and new assets. I illustrate two effects by considering two extreme settings. First, suppose traders completely agree about the payoffs of old assets (so that financial innovation increases the speculative variance only by generating new disagreements). In this setting, the following result shows that financial innovation increases the speculative variance more in economies with more complete financial markets than those with less complete markets.

Theorem 2. *Consider two economies $\mathcal{E}(J = J_O \cup J_N)$ and $\mathcal{E}(\hat{J} = \hat{J}_O \cup J_N)$ that share the*

same set of new assets but potentially different set of old assets. Suppose in each economy traders agree about the payoffs of old assets, that is, for each $j \in J_O$ (resp. $\hat{j} \in \hat{J}_O$), μ_i^j (resp. $\mu_i^{\hat{j}}$) is a constant independent of i . Suppose also that $\hat{J}_O \subset J_O$, so the economy $\mathcal{E}(J)$ has a more complete set of existing assets than the economy $\mathcal{E}(\hat{J})$. Then, financial innovation increases the speculative variance more in economy $\mathcal{E}(J)$ than economy $\mathcal{E}(\hat{J})$, that is:

$$\Omega^S(J_O \cup J_N) - \Omega^S(J_O) \geq \Omega^S(\hat{J}_O \cup J_N) - \Omega^S(\hat{J}_O).$$

The intuition for this result is the hedge-more/bet-more effect. This effect not only enables new assets to amplify speculation on old assets (as illustrated in Example 2), but it also enables old assets to amplify speculation on new assets. In more complete financial markets, traders are able to take purer bets on new asset, which leads to a greater increase in the speculative variance.

I next consider the other extreme, by assuming that traders completely agree about the payoffs of new assets (so that financial innovation increases the speculative variance only by amplifying speculation on existing disagreements). In this setting, the following result shows that financial innovation increases the speculative variance more when new assets are correlated with existing assets -as would be the case with new derivative assets.

Theorem 3. Consider two economies $\mathcal{E}(J = J_O \cup \{j_N\})$ and $\mathcal{E}(\hat{J} = J_O \cup \{\hat{j}_N\})$ that share the same set of old assets but that have different new assets: j_N and \hat{j}_N . Suppose there is no disagreement about either j_N or \hat{j}_N , that is, $\mu_i^{j_N}$ and $\mu_i^{\hat{j}_N}$ are constants independent of i . Suppose also that j_N is correlated with at least one existing asset $j \in J_O$ while \hat{j}_N is not correlated with any of the assets in J_O (that is, $(\mathbf{A}^j)' \Lambda^{\mathbf{v}} \mathbf{A}^{j_N} \neq 0$ for some $j \in J_O$, while $(\mathbf{A}^j)' \Lambda^{\mathbf{v}} \mathbf{A}^{\hat{j}_N} = 0$ for all $j \in J_O$). Then, financial innovation increases the speculative variance more in economy $\mathcal{E}(J)$ than economy $\mathcal{E}(\hat{J})$, that is:

$$\Omega^S(J_O \cup \{j_N\}) - \Omega^S(J_O) \geq \Omega^S(J_O \cup \{\hat{j}_N\}) - \Omega^S(J_O).$$

Intuitively, the hedge-more/bet-more effect is operational in economy $\mathcal{E}(J)$, but not in economy $\mathcal{E}(\hat{J})$ because the new and old assets are uncorrelated. When there is agreement about new assets, financial innovation increases the speculative variance only because of this effect. Consequently, financial innovation increases the speculative variance in economy $\mathcal{E}(J)$, but not in economy $\mathcal{E}(\hat{J})$.

While Theorems 2 and 3 apply in extreme settings, they are suggestive of more general forces that would also be operational under less extreme assumptions. Taken together, these

results create a presumption that financial innovation is more likely to be destabilizing when it concerns more developed financial markets and derivative assets.

6 Endogenous Financial Innovation

In practice, many financial products are introduced by financial institutions with profit incentives. This section endogenizes financial innovation by introducing a profit seeking market maker. It shows that new assets that increase the average variance can endogenously emerge in view of the profit incentives of the market maker.

Consider the earlier model with the only difference that the available assets, J , are introduced by a market maker. At the beginning of date 0, the market maker can choose the set of available assets, J , from a larger set, \mathcal{J} , which can be roughly interpreted as the technology frontier for financial innovation. The market maker faces a capacity constraint: She can introduce M assets for free, but it is infinitely costly to introduce $M + 1$ assets. This assumption is made only for simplicity and does not change any of the qualitative conclusions. Once the market maker chooses J , these assets are traded in a competitive market similar to the previous sections. The main difference is that trade in these asset is also intermediated by the market maker, which enables the market maker to earn profits. To capture this idea in a simple way, suppose the market maker can determine whether a trader can participate in the competitive market for assets J . In particular, the market maker sets a fixed membership fee π_i for each trader i and makes a take it or leave it offer. If trader i accepts the offer, then she can trade the available assets in the competitive market. Otherwise, trader i is out of the market, and her net worth is given by her endowment, $e + \mathbf{W}_i \mathbf{v}$.⁷

The equilibrium of this economy can be characterized backwards. First consider the competitive equilibrium after the market maker has chosen J and traders decided whether or not to participate in the market. Without loss of generality, assume that all traders have accepted the offer (this will be the case in equilibrium). In view of the CARA utility (that features no wealth effects), traders' portfolio choices are independent of the fixed fees they have paid. In particular, the equilibrium allocation, $\{\mathbf{x}_i(J)\}_i$, is characterized as in earlier sections.

Next consider the fixed fees the market maker charges for a given choice of assets, J . If trader i rejects the offer, she receives the certainty equivalent payoff from her endowment. Otherwise, she receives the certainty equivalent payoff from her equilibrium allocation net of the fixed fee, $\pi_i(J)$. The market maker sets $\pi_i(J)$ so that the trader is just indifferent to accept the offer (and in the unique equilibrium, she accepts). Straightforward calculations (relegated

⁷ Put differently, the market maker can extract the full surplus from traders. I adopt this assumption because it simplifies the analysis considerably while capturing the essential trade-offs regarding innovation incentives.

to Appendix A.4) show that the market maker's expected profits are given by:

$$\sum_i \pi_i(J) = \sum_i \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i(\mathbf{J})}{\theta_i} - \tilde{\lambda}_i(\mathbf{J}) \right) \mathbf{\Lambda}(J)^{-1} \left(\frac{\tilde{\mu}_i(\mathbf{J})}{\theta_i} - \tilde{\lambda}_i(\mathbf{J}) \right). \quad (29)$$

In this model, the market maker's profits correspond to traders' perceived surplus from trading assets. Consequently, the expression (29) reflects the two motives for trade in this economy. The market maker's profits are larger for assets that facilitate better risk sharing [i.e., larger $\tilde{\lambda}_i(J)$], or for assets that generate greater belief disagreements [i.e., larger $\tilde{\mu}_i(J)$].

Finally consider the market maker's innovation decision at the beginning of date 0. The market maker chooses the set of assets, J , which maximizes profits from all traders. That is, J solves:

$$\max_{\hat{J} \subset \mathcal{J}, |\hat{J}|=M} \sum_{i \in I} \pi_i(\hat{J}). \quad (30)$$

The next result characterizes the optimal innovation decision in two extreme cases.

Theorem 4. *Consider the above described economy with endogenous financial innovation.*

(i) *Suppose traders do not disagree about the assets in \mathcal{J} , that is, for each $j \in \mathcal{J}$, μ_i^j is a constant independent of i . Then, the market maker innovates assets that minimize the average variance, i.e., J solves:*

$$\min_{\hat{J} \subset \mathcal{J}, |\hat{J}|=M} \Omega(J).$$

(ii) *Suppose traders disagree on some of assets in \mathcal{J} , that is, there exists $i, \hat{i} \in I$ and $j \in \mathcal{J}$ such that $\mu_i^j \neq \mu_{\hat{i}}^j$. Suppose also that traders' beliefs are parameterized by $K \in \mathbb{R}_+$, with $\mu_{i,K}^{\mathcal{J}} = K \mu_i^{\mathcal{J}}$ for all i . Then, there exists \bar{K} such that if $K > \bar{K}$, then the market maker innovates assets that maximize the average variance, i.e., J solves:*

$$\max_{\hat{J} \subset \mathcal{J}, |\hat{J}|=M} \Omega(J).$$

The first part of the theorem formalizes the traditional view of financial innovation: In a common-beliefs benchmark, the market maker introduces assets that minimize the riskiness of traders' net worths. On the other hand, the second part illustrates that, when traders' disagreements are sufficiently strong, the market makers' incentives are tilted towards innovating assets that will increase the riskiness of traders' net worths. This result complements the earlier findings (especially, Theorem 1) by identifying speculation as a potentially important factor in financial innovation.

7 Learning and Dynamic Equilibrium

While the static model is useful to analyze the speculation generated by a given level of belief disagreements, a natural question is how these disagreements evolve over time. The first part of this section addresses this question by considering a simple dynamic extension in which traders learn from past asset payoffs. It shows that financial innovation always decreases the average variance in the long run. A question thus emerges as to how new assets should be introduced to reduce their short run impact on the speculative variance. A plausible idea to stagger the introduction of new assets, so as to introduce them “one at a time.” The second part of this section analyzes the implications of this policy.

7.1 Dynamic Evolution of Speculative Variance

Consider a simple dynamic extension in which the static economy of earlier sections is repeated in each period $t \in \{0, 1, \dots, \infty\}$. In particular, traders receive endowments at the beginning of period t . They trade assets in period t , and they consume their net worth at the end of period t (so they do not solve any intertemporal optimization problem).

The central feature of the dynamic model is traders’ learning about past asset payoffs. To introduce learning in a tractable way, it is useful to put more structure on the uncertainty model described in Section 2. In particular, suppose the underlying uncertainty consists of two parts, $\mathbf{v} = (\mathbf{v}_{endow}, \mathbf{v}_{other})$, where \mathbf{v}_{endow} captures traders’ endowment risks and \mathbf{v}_{other} captures other risks. In particular, suppose

$$w_i = (\mathbf{W}_i)' \mathbf{v}_{endow}, \text{ and } a^j = (\mathbf{A}^j)' \mathbf{v},$$

so that endowments only respond to \mathbf{v}_{endow} , while asset payoffs may respond to both types of risks. Suppose also that traders know (and agree about) the distribution of \mathbf{v}_{endow} , while they face uncertainty about the distribution of \mathbf{v}_{other} (which will be described shortly). This assumption is made for tractability, but it can also be motivated by noting that traders are more familiar with their endowment risks compared to other risks that might affect asset payoffs. In this setting, traders know the covariance terms, $\{\boldsymbol{\lambda}_i\}_i$. Moreover, traders do not receive any information from observing endowment realizations. In contrast, traders update their beliefs for \mathbf{v}_{other} by observing asset payoffs.

It is useful to specify traders’ beliefs directly in terms of the asset payoffs (instead of the underlying uncertainty, \mathbf{v}_{other}) which is without loss of generality. In particular, let $N(\boldsymbol{\mu}_{true}, \Lambda_{true})$ denote the true distribution of asset payoffs. Suppose that traders know Λ_{true} (for simplicity), but they do not know the mean of asset payoffs. In particular, trader i believes

that $\boldsymbol{\mu}_{true}$ is distributed according to $N(\boldsymbol{\mu}_{i,prior}, \Lambda_{prior})$. As trader i observes asset payoffs, she updates her prior about $\boldsymbol{\mu}_{true}$. In particular, let $\{\mathbf{a}_i(t)\}_{t=1}^{t-1}$ denote the asset payoff realizations up to date t , and $\bar{\mathbf{a}}(t) = \frac{\sum_{\tilde{i}=0}^{t-1} \mathbf{a}(\tilde{i})}{t}$ denote the average realization. Using Bayes' rule, trader i 's posterior for $\boldsymbol{\mu}_{true}$ is also Normally distributed, $N(\boldsymbol{\mu}_{i,post}(t), \Lambda_{post}(t))$, with mean and variance given by:

$$\begin{aligned}\Lambda_{post}(t) &= \left(\Lambda_{prior}^{-1} + t\Lambda_{true}^{-1}\right)^{-1} \text{ and} \\ \boldsymbol{\mu}_{i,post}(t) &= \Lambda_{post}(t) \left(\Lambda_{prior}^{-1}\boldsymbol{\mu}_{i,prior} + t\Lambda_{true}^{-1}\bar{\mathbf{a}}(t)\right).\end{aligned}\tag{31}$$

Note that traders update their beliefs by combining their prior information with the information in observations. The relative weight on the prior is given by the prior precision matrix, Λ_{prior}^{-1} , while the relative weight on the observations are given by the precision matrix for the observation, Λ_{true}^{-1} .

Given traders' posterior beliefs for $\boldsymbol{\mu}_{true}$, their beliefs for the asset payoffs can also be calculated. In particular, trader i at date t believes that the asset is Normally distributed, $N(\boldsymbol{\mu}_i(t), \Lambda(t))$, with mean and variance given by:

$$\boldsymbol{\mu}_i(t) = \boldsymbol{\mu}_{i,post}(t) \text{ and } \Lambda(t) = \Lambda_{true} + \Lambda_{post}(t).\tag{32}$$

Note that traders' beliefs for period t , $\{N(\boldsymbol{\mu}_i(t), \Lambda_i(t))\}_{i \in I}$ (which are stochastic processes), always satisfy Assumption (A1) of Section 3. Consequently, conditional on beliefs, the equilibrium in each period is characterized as in the earlier sections.

Consider next the evolution of the speculative variance, $\Omega^S(t)$. Using Eqs. (31) and (32), traders' belief differences can be written as:

$$\tilde{\boldsymbol{\mu}}_i(t) = \left(\Lambda_{prior}^{-1} + t\Lambda_{true}^{-1}\right)^{-1} \Lambda_{prior}^{-1} \tilde{\boldsymbol{\mu}}_{i,prior}.\tag{33}$$

In particular, even though traders' beliefs are stochastic, their belief differences evolve deterministically. This is because all traders update their priors by observing the same payoffs, $\{\mathbf{a}_i(t)\}_{t=1}^{t-1}$ (and their priors have the same precision by assumption). Note also that $\tilde{\boldsymbol{\mu}}_i(t)$ limits to 0 as $t \rightarrow \infty$. Intuitively, as traders learn by observing asset returns, their prior beliefs converge and their belief disagreements disappear in the long run. By Eq. (25), it follows that the speculative variance also limits to zero:

$$\lim_{t \rightarrow \infty} \Omega^S(t) = 0.\tag{34}$$

Moreover, Eq. (33) shows that traders belief differences converge to 0 at rate t [and thus,

by Eq. (25), the speculative variance converges to 0 at rate t^2 . In contrast, the central limit theorem implies that traders' prior beliefs converge to the true mean, $\boldsymbol{\mu}_{true}$, only at rate \sqrt{t} . It follows that traders' belief differences in this model (and the speculative variance) are likely to disappear long before traders learn the true distribution of asset payoffs.

It is useful to illustrate this result (as well as the updating of beliefs) with an example. Consider Example 2 from Section 2 in which asset 1 has the payoff, $a^1 = v^1 + \alpha v^2$, and asset 2 has the payoff $a^2 = v^2$. Recall that (v^1, v^2) is a standard Normal random variable, and that traders disagree on v^1 but not on v^2 . In the notation of this section, asset payoffs are given by:

$$\boldsymbol{\mu}_{true} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \text{ and } \Lambda_{true} = \begin{bmatrix} 1 + \alpha^2 & \alpha \\ \alpha & 1 \end{bmatrix},$$

and traders' prior beliefs are given by:

$$\boldsymbol{\mu}_{1,prior} = \begin{bmatrix} \bar{\mu} + \varepsilon \\ \bar{\mu} \end{bmatrix} \text{ and } \boldsymbol{\mu}_{2,prior} = \begin{bmatrix} \bar{\mu} - \varepsilon \\ \bar{\mu} \end{bmatrix}, \text{ with } \Lambda_{prior} = \frac{1}{\tau} \Lambda_{true}.$$

Trader i 's initial belief for the asset payoffs are given by $N(\boldsymbol{\mu}_{i,prior}, (1 + \frac{1}{\tau}) \Lambda_{true})$, which is the same as in Section 2. Eqs. (31) and (32) of this section show that trader i 's period t belief for the asset payoff has mean and variance given by:

$$\boldsymbol{\mu}_i(t) = \frac{\tau \boldsymbol{\mu}_{i,prior} + t \bar{\mathbf{a}}(\tilde{t})}{\tau + t} \text{ and } \Lambda(t) = \left(1 + \frac{1}{\tau + t}\right) \Lambda_{true}.$$

In particular, note that the mean of traders' period t belief is a weighted average of the mean of their prior belief and the average of observations, with relative weights given by respectively the precision parameter, τ , and the number of periods, t . Note also that the variance of traders' beliefs for asset payoffs gradually declines and converges to the true asset variance, Λ_{true} . Traders' belief differences can also be calculated as $\tilde{\boldsymbol{\mu}}_1(t) = -\tilde{\boldsymbol{\mu}}_2(t) = \begin{bmatrix} \frac{\tau}{\tau+t} \varepsilon \\ 0 \end{bmatrix}$. Figure 1 illustrates the evolution of traders' beliefs and the speculative variance in this example. Note that traders' beliefs converge to each other at a faster rate than the rate at which they converge to the true mean.

These results imply that financial innovation has a negligible effect on speculative variance in the long run. On the other hand, by Theorem 1, financial innovation reduces the uninsurable variance, $\Omega^R(t)$, in the short run as well as the long run. The following result uses these observations to provide an analogue of Theorem 1 regarding the long run effects of new assets. As before, let $J = J_O \cup J_N$ where J_O denotes the set of old assets and J_N denotes a set of new assets introduced in period 0.




Figure 1: Dynamic evolution of traders’ beliefs and the speculative variance in Example 2.

Theorem 5. *Consider the dynamic equilibrium described in this section respectively for the initial set of assets J_O and $J_O \cup J_N$. Financial innovation reduces the average variance in the long run, that is:*

$$\lim_{t \rightarrow \infty} \Omega(t; J_O \cup J_N) \leq \lim_{t \rightarrow \infty} \Omega(t; J_O).$$

Intuitively, in the long run, new assets provide their risk sharing services without being subject to speculation. Consequently, while financial innovation may temporarily increase the average variance, it always decreases it in the long run. Hence, a sufficiently patient social planner that cares about the reduction of average variance would always prefer to introduce new assets. A natural question is whether the planner could also mitigate the short run impact of financial innovation, which I turn to next.

7.2 Staggering the Introduction of New Assets

A plausible policy is to stagger (or delay) the introduction of assets, so as to introduce them “one at a time.” This policy ensures that traders learn about the payoffs of some new assets before they get to trade other new assets. In view of the hedge-more/bet-more effect, this policy would reduce speculation on the assets that are introduced earlier. On the other hand, this policy could also adversely affect traders’ learning. Arguably, traders learn about the payoffs of new assets to the extent that they trade them. If this is the case, then a staggering policy reduces the amount of learning on the assets that are introduced later. Hence, a staggering policy represents a trade-off between speculation and learning.

I investigate the implications of this trade-off in a slightly modified version of the earlier dynamic model. Suppose the asset j is introduced for trade in period t^j . More specifically, asset j can be traded in all periods $t \geq t^j$ but not before. Let $\mathbf{t} = \{t^j\}_{j \in J}$ denote the asset introduction policy, and $J(t)$ denote the set of assets that are available for trade in period t . A staggering policy corresponds to $t^1 < \dots < t^{|J|}$. An immediate introduction policy is denoted by $t^1 = t^2 = \dots = t^{|J|} = 0$.

The key assumption of this subsection is that traders observe the payoff of a new asset if and only if it is introduced for trade. More specifically, at the end of a period t , traders only observe the payoffs of assets that were traded in that period, $\mathbf{a}^{J(t)}(t)$. This assumption implies that traders posterior beliefs for $\boldsymbol{\mu}_{true}$ is different than (31), and given by:

$$\begin{aligned} \Lambda_{post}(t) &= \left(\Lambda_{prior}^{-1} + \sum_{\tilde{t}=0}^{t-1} \left(\Lambda_{true|J(\tilde{t})} \right)^{-1} \right)^{-1}, \text{ and} \\ \boldsymbol{\mu}_{i,post}(t) &= \Lambda_{post}(t) \left(\Lambda_{prior}^{-1} \boldsymbol{\mu}_{i,prior} + \sum_{\tilde{t}=0}^{t-1} \left(\Lambda_{true|J(\tilde{t})} \right)^{-1} \mathbf{a}^{J(\tilde{t})}(\tilde{t}) \right). \end{aligned} \quad (35)$$

Here, the matrix $\Lambda_{true|J(t)}$ corresponds to the sub-matrix of Λ_{true} corresponding to the set of available assets, $J(t)$. Note that traders incorporate only the information provided by available assets. Given these posteriors, traders' beliefs for asset payoffs are given by (32). The rest of the equilibrium is characterized as in the previous subsection.

The goal of this subsection is to investigate the effect of the introduction policy, \mathbf{t} , on an appropriately defined measure of short run speculative variance. I consider the following measure:

$$\sum_{t=0}^{\infty} \Omega^S(t),$$

which converges to a finite constant in view of the analysis in the previous subsection. As this measure does not have discounting, it provides a fair comparison between a staggering policy and an immediate introduction policy. The next result shows that, under a mild assumption on the variance of traders' prior beliefs, the introduction policy has no effect on this measure.

Assumption (A1^D). The variance of traders' prior beliefs for $\boldsymbol{\mu}_{true}$ can be written as $\Lambda_{prior} = \frac{1}{\tau} \Lambda_{true}$ for some $\tau > 0$.

This assumption simplifies the learning dynamics in (35) by ensuring that traders' prior beliefs have the same variance structure as their observations. The parameter, τ , provides a measure of the precision of traders' prior beliefs relative to the precision of payoff realizations.

Theorem 6. Consider the dynamic equilibrium described in this subsection and suppose As-

assumption (A1^D) holds. The sum of the speculative variance, $\sum_{t=0}^{\infty} \Omega^S(t)$, is a constant independent of the asset introduction policy, \mathbf{t} .

This result is relatively easy to see when assets are not correlated with each other, i.e., when Λ_{true} is a diagonal matrix. In this case, Assumption (A1^D) and Eq. (35) imply that learning is separable across assets. In particular, a trader i 's posterior belief for the mean payoff of an asset j only depends on the past observations of this asset (in particular, it is independent of the availability of other assets). Since $\Lambda(t) = \Lambda_{true} + \Lambda_{post}(t)$ is diagonal, Eq. (25) implies that speculation is also separable across assets. In particular, the contribution of asset j to $\Omega^S(t)$ is independent of the availability of other assets. Combining these observations, it follows that the contribution of each asset j to the sum, $\sum_{t=0}^{\infty} \Omega^S(t)$, can be calculated separately. Since the environment is stationary, the introduction policy, t^j , simply shifts the beginning of traders' learning and speculation with no effect on the contribution of j to $\sum_{t=0}^{\infty} \Omega^S(t)$. Repeating this argument for all assets $j \in J$, it follows that the sum, $\sum_{t=0}^{\infty} \Omega^S(t)$, is a constant that is independent of the introduction policy [see Eq. (A.11) in the appendix].

While this argument provides a sketch proof for a special case, it does not apply when assets are correlated. In this case, speculation on an asset j is no longer independent of the availability of other assets. More specifically, in view of the hedge-more/bet-more effect, a staggering policy should lower the amount of speculation on assets that are introduced earlier. To illustrate this idea, consider once again Example 2 in which asset payoffs are given by $a^1 = v^1 + \alpha v^2$ and $a^2 = v^2$, and disagreements are on v^1 . In this setting, consider a staggering policy of introducing asset a^1 in period $t^1 = 0$ and asset a^2 in a later period $t^2 > 0$. With this policy, in periods $t < t^2$, the speculative variance is lower than the case in which both assets are available. Moreover, in these periods traders learn about the distribution of a^1 on which they have initial disagreements. It could then appear that a staggering policy with a sufficiently large $t^2 > 0$ should lead to a lower speculative variance compared to an immediate introduction policy.

Theorem 6 implies that this argument is incorrect. To see why this is so, consider Figure 2 which provides a further diagnosis of the example. The first two panels of this figure show that, while the delayed introduction policy decreases the belief disagreements on a^1 , it simultaneously increases belief disagreements on a^2 (on which there is initially no disagreement). Intuitively, by observing only $a^1 = v^1 + \alpha v^2$ (which is a noisy measure of v^1), traders do not fully learn v^1 which is the initial source of their disagreements. Consequently, as they learn a^1 better, their remaining disagreements about v^1 induce them to start to disagree about the residual, $v^2 = \frac{a^1 - v^1}{\alpha}$. It follows that, when the asset $a^2 = v^2$ is introduced at date t^2 , there will be an

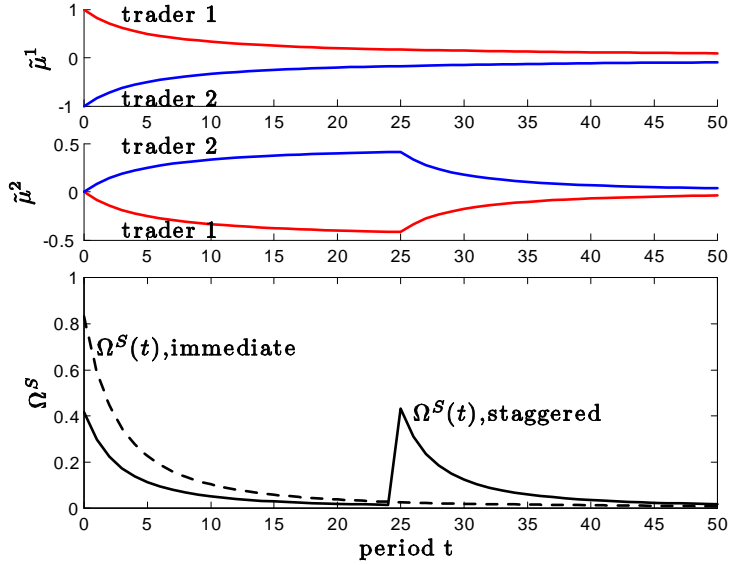


Figure 2: The first two panels illustrate the evolution of traders' belief differences for assets a^1 and a^2 in Example 2 with the staggered introduction policy ($t^1 = 0$ and $t^2 > 0$). The last panel illustrates the evolution of the speculative variance corresponding to the staggered introduction policy as well as the immediate introduction policy ($t^1 = t^2 = 0$).

additional spike in speculative variance. Put differently, while the staggering policy reduces the speculative variance initially, it leads to a greater speculative variance in subsequent periods, as illustrated by the last panel of Figure 2. It is no longer obvious whether the staggering policy or the immediate introduction policy should lead to a lower sum of speculative variance.

I next provide a sketch proof that the two policies lead to the same sum of speculative variance in this example. Define a set of alternative assets $(\tilde{a}^1, \tilde{a}^2)$ as follows:

$$\begin{aligned}\tilde{a}^1 &= a^1 = v^1 + \alpha v^2, \\ \tilde{a}^2 &= a^1 - \left(\alpha + \frac{1}{\alpha}\right) a^2 = v^1 - \frac{1}{\alpha} v^2.\end{aligned}$$

It can be checked that, for any introduction policy \mathbf{t} that satisfies $t^1 \leq t^2$, the alternative assets, $(\tilde{a}^1, \tilde{a}^2)$, represent the same set of trading possibilities and provide the same information as the original assets, (a^1, a^2) . Moreover, the alternative assets, \tilde{a}^1 and \tilde{a}^2 , are uncorrelated by construction. Invoking the earlier argument for uncorrelated assets, it follows that the sum of speculative variance, $\sum_t \Omega^S(t)$, does not depend on the introduction policy also for the original assets, (a^1, a^2) .

The proof of Theorem 6 in the appendix formalizes this argument (and extends it to a

general set of assets and policies). Intuitively, each asset that is introduced for trade opens a new dimension for speculation and learning (captured \tilde{a}^i in the above proof) which is orthogonal to the dimensions that are already opened. Staggering (or delaying) the introduction of new assets is ineffective because it delays simultaneously the speculation as well as the learning in the corresponding dimensions.

This analysis also points to alternative policy tools that might be more effective in reducing the short run speculative variance. A desirable tool should reduce traders' speculation without precluding learning. One alternative is to set temporary position limits (or taxes) on new assets. These position limits naturally restrict traders' speculative positions. However, unlike staggering, position limits do not reduce learning (assuming that traders' learning does not depend on the size of their positions). Moreover, the position limits are also practical as they can be implemented by imposing tighter capital adequacy requirements for new assets than for older assets. These requirements could be phased away over time to ensure that new assets provide their risk sharing benefits in the long run.

8 Conclusion

This paper theoretically analyzed the effect of financial innovation on the allocation of risks in a CARA-Normal economy in which both the speculation and risk sharing forces are present. Financial innovation increases the opportunities for risk sharing and diversification because new assets are correlated with traders' endowment risks. However, financial innovation also increases speculation through two channels: By generating new disagreements, and by amplifying traders' speculation on existing disagreements. The second channel stems from an important economic force, the *hedge-more/bet-more* effect: Traders use new assets to hedge their bets on existing assets (i.e., to take purer bets), which enables them to take larger bets. The hedge-more/bet-more effect suggests that financial innovation is more likely to increase the riskiness of traders' net worths in more complete financial markets and when it concerns derivative assets. A fruitful future research direction may be to investigate in more detail these and other implications of the hedge-more/bet-more effect.

In a dynamic setting, speculation decreases in the long run as traders learn about past asset payoffs. Thus, speculation on new assets is likely to be more potent in the short run following their introduction. This naturally raises the question of how new assets should be introduced to mitigate their short run impact on speculation. Staggering the introduction of new assets is likely to be ineffective because it reduces traders' speculation simultaneously with their learning. A more viable alternative appears to be temporary position limits on new

assets, which can be implemented with temporary capital requirements.

This paper has touched upon policy issues only briefly. This is because the equilibrium in this model is Pareto efficient despite the fact that trade in new securities may increase the average variance. In view of belief disagreements, each trader perceives a large expected return from her trades that justifies the additional risks that she is taking. Despite the Pareto efficiency of equilibrium, it is important to analyze the potential policy interventions for at least two reasons. First, while this paper illustrates the results in a standard CARA-Normal framework without externalities, the main mechanisms apply also in richer environments that may feature externalities. For example, if the traders are financial intermediaries that do not fully internalize the social costs of their losses (or bankruptcies), then an increase in speculation may lead to a Pareto inefficiency. Appendix B presents a stylized model that formalizes this argument. Second, the notion of Pareto efficiency with heterogeneous priors is somewhat unsatisfactory. This is because while all traders perceive a large expected return, at most one of these expectations can be correct. As noted by Stiglitz (1989), “there are real difficulties in interpreting the welfare losses associated with impeding trades based on incorrect expectations.” I leave policy analysis for future work.

A Appendix: Omitted Proofs

A.1 Omitted proofs for Section 3

Proof of Lemma 1. Recall that the objective function for Problem (23) is given by

$$\Omega = \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^\nu \mathbf{W}_i + \mathbf{x}'_i \Lambda \mathbf{x}_i + 2\mathbf{x}'_i \boldsymbol{\lambda}_i). \quad (\text{A.1})$$

The first order conditions are given by:

$$\Lambda \mathbf{x}_i + \boldsymbol{\lambda}_i = \gamma \frac{\bar{\theta}}{\theta_i} \text{ for each } i \in I,$$

where $\gamma \in \mathbb{R}^{|J_E|}$ is a vector of Lagrange multipliers. Note that $\mathbf{x}_i^R = -\Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i$ satisfies these first order conditions for the Lagrange multiplier $\gamma = (\sum_{i \in I} \boldsymbol{\lambda}_i) / |I|$. This shows that $\{\mathbf{x}_i^R\}_i$ is the unique solution to Problem (23).

Proof of Lemma 2. Plugging in $\mathbf{x}_i^R = -\Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i$ into the objective function (A.1), the optimal value, Ω^R , is given by:

$$\begin{aligned} \Omega^R &= \frac{1}{|I|} \sum_i \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^\nu \mathbf{W}_i + \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i) - \frac{2}{|I|} \sum_i \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \frac{\theta_i}{\bar{\theta}} \boldsymbol{\lambda}_i, \\ &= \frac{1}{|I|} \sum_i \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^\nu \mathbf{W}_i + \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i) - \frac{2}{|I|} \sum_i \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \left(\frac{\theta_i}{\bar{\theta}} \tilde{\boldsymbol{\lambda}}_i + \frac{1}{|I|} \sum_{i \in I} \boldsymbol{\lambda}_i \right) \\ &= \frac{1}{|I|} \sum_i \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^\nu \mathbf{W}_i + \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i) - \frac{2}{|I|} \sum_i \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \frac{\theta_i}{\bar{\theta}} \tilde{\boldsymbol{\lambda}}_i \\ &= \frac{1}{|I|} \sum_i \frac{\theta_i}{\bar{\theta}} (\mathbf{W}'_i \Lambda^\nu \mathbf{W}_i - \tilde{\boldsymbol{\lambda}}'_i \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i). \end{aligned}$$

Here, the second line uses the definition of $\tilde{\boldsymbol{\lambda}}_i$ to replace $\boldsymbol{\lambda}_i$; the third line uses the facts that $\sum_i \tilde{\boldsymbol{\lambda}}_i = 0$ and that $\frac{1}{|I|} \sum_{i \in I} \boldsymbol{\lambda}_i$ is a constant independent of i ; and the last line uses simple algebra. This completes the derivation of Eq. (24).

To derive Eq. (25), first consider the expression $|I| \left(\Omega - \frac{1}{|I|} \sum_i \frac{\theta_i}{\bar{\theta}} \mathbf{W}'_i \Lambda^\nu \mathbf{W}_i \right)$. Using the definition of the average variance in (22), this expression can be written as:

$$\begin{aligned}
\sum_i \frac{\theta_i}{\bar{\theta}} \mathbf{x}_i' \Lambda^{-1} \mathbf{x}_i + 2 \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \mathbf{x}_i' \boldsymbol{\lambda}_i &= \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left[\left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right)' \Lambda^{-1} \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right) + 2 \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right)' \Lambda^{-1} \boldsymbol{\lambda}_i \right] \\
&= \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left[\left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right)' \Lambda^{-1} \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right) + 2 \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right)' \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i \right] \\
&= \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left[\left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \tilde{\boldsymbol{\lambda}}_i \right)' \Lambda^{-1} \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} + \tilde{\boldsymbol{\lambda}}_i \right) \right] \\
&= \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \Lambda^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \tilde{\boldsymbol{\lambda}}_i' \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i.
\end{aligned}$$

Here, the first line substitutes for the portfolio demands from (20); the second line uses the definition of $\tilde{\boldsymbol{\lambda}}_i$ to replace $\boldsymbol{\lambda}_i$ and simplifies the resulting expression (as in the first part of the proof); and the next two lines follow by simple algebra. Next, using the fact that the last line equals $|I| \left(\Omega - \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \mathbf{W}_i' \Lambda^\nu \mathbf{W}_i \right)$, the average variance can be written as:

$$\Omega = \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left(\mathbf{W}_i' \Lambda^\nu \mathbf{W}_i - \tilde{\boldsymbol{\lambda}}_i' \Lambda^{-1} \tilde{\boldsymbol{\lambda}}_i \right) + \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \Lambda^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i}.$$

Using the definition of Ω^R in (24), it follows that the speculative variance, $\Omega^S = \Omega - \Omega^R$, is given by the expression in (25).

A.2 Omitted proofs for Section 4

Proof of Lemma 3. To prove the first part, first note the hypothetical economy features $\boldsymbol{\lambda}_i = 0$ for all $i \in I$. Next note that $\mathbf{x}_{i,H}$ solves the mean-variance problem, (2), for the hypothetical economy. Using $\boldsymbol{\lambda}_i = 0$, $\mathbf{W}_i = 0$, and $\mathbf{p} = \frac{1}{|I|} \sum_{i \in I} \frac{\bar{\theta}}{\theta_i} \boldsymbol{\mu}_i$ [by Eq. (19)], this mean-variance problem is equivalent to (26). Finally, consider the first order condition for problem (26), which implies that

$$\mathbf{x}_{i,H} = \Lambda^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} = \mathbf{x}_i^S.$$

Here, the last equality follows from Eq. (20), completing the proof for the first part. For the second part, note that Eq. (27) follows from (22) after substituting $\boldsymbol{\lambda}_i = 0$, $\mathbf{W}_i = 0$.

Proof of Theorem 1. Part (i). By definition, Ω^R , is the optimal value of the minimization problem (23). Financial innovation expands the constraint set of this problem. Thus, it also decreases the optimal value, proving $\Omega^R(J_O \cup J_N) \leq \Omega^R(J_O)$.

Part (ii). The proof is provided in the text. Here, to demonstrate an alternative approach,

I provide a second proof using matrix algebra. First note that the definition of Ω^S in (25) implies that

$$\begin{aligned}\Omega^S(J) - \Omega^S(J_O) &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left(\frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \Lambda^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i} - \frac{\tilde{\boldsymbol{\mu}}_i(J_O)'}{\theta_i} \Lambda(J_O)^{-1} \frac{\tilde{\boldsymbol{\mu}}_i(J_O)}{\theta_i} \right) \\ &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \left(\Lambda^{-1} - \begin{bmatrix} \Lambda(J_O)^{-1} & 0 \\ 0 & 0 \end{bmatrix} \right) \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i}.\end{aligned}\quad (\text{A.2})$$

I next claim that the matrix in the parenthesis,

$$\Lambda^{-1} - \begin{bmatrix} \Lambda(J_O)^{-1} & 0 \\ 0 & 0 \end{bmatrix}, \quad (\text{A.3})$$

is positive semidefinite. In view of this claim, Eq. (A.2) implies $\Omega^S(J) \geq \Omega^S(J_O)$, providing an alternative proof for this part.

This claim follows from Lemma 5.16 in Horn and Johnson (2007). This lemma considers a positive definite matrix partitioned into submatrices of arbitrary dimension, $A = \begin{bmatrix} A_{11} & A_{12} \\ A_{12}^T & A_{22} \end{bmatrix}$. It shows that the matrix A^{-1} is weakly greater than the matrix $\begin{bmatrix} (A_{11})^{-1} & 0 \\ 0 & 0 \end{bmatrix}$ in positive semidefinite order. This in turn implies that the matrix, $A^{-1} - \begin{bmatrix} (A_{11})^{-1} & 0 \\ 0 & 0 \end{bmatrix}$, is positive semidefinite. Invoking this lemma for $A = \Lambda$ and $A_{11} = \Lambda(J_O)$ shows that the matrix in Eq. (A.3) is positive semidefinite, completing the alternative proof.

Part (iii). Consider an economy, $\mathcal{E}(J)$, with two properties: (i) There is no disagreement on existing assets, i.e., $\mu_i^j = \mu_{\hat{i}}^j$ for each $j \in J_O$ and $i, \hat{i} \in I$, and (ii) There is some disagreement about new assets, i.e., $\mu_i^j \neq \mu_{\hat{i}}^j$ for some $j \in J_N$ and $i, \hat{i} \in I$. Let $\mathcal{E}_K(J)$ denote the economy which is identical except that traders' beliefs for the underlying uncertainty, \mathbf{v} , are scaled by the factor K , that is: $\boldsymbol{\mu}_{i,K}^{\mathbf{Y}} = K \boldsymbol{\mu}_i^{\mathbf{Y}}$ for each i . I claim that there exists $K > 0$ such that the result holds for the economy $\mathcal{E}_K(J)$, that is:

$$\Omega_K(J_O \cup J_N) > \Omega_K(J_O) + \eta. \quad (\text{A.4})$$

To show this claim, first note that by assumption:

$$\tilde{\boldsymbol{\mu}}_i^{J_O} = \mathbf{0} \text{ and } \tilde{\boldsymbol{\mu}}_i^{J_N} \neq \mathbf{0}. \quad (\text{A.5})$$

Here \mathbf{z}^J denotes the vector (z^1, \dots, z^J) and $\mathbf{0}$ denote the zero vector. Next note that traders' beliefs for asset payoffs in economy $\mathcal{E}_K(J)$ are scaled by a factor of K , i.e., $\boldsymbol{\mu}_{i,K} = K \boldsymbol{\mu}_i$. Using

(A.5) and (25), this implies:

$$\Omega_K^S(J_O) = K^2 \Omega^S(J_O) = 0 \text{ and } \Omega_K^S(J) = K^2 \Omega^S(J) > 0.$$

These expressions further imply:

$$\lim_{K \rightarrow \infty} \Omega_K^S(J) - \Omega_K^S(J_O) = \infty.$$

Finally, note from Eq. (24) that K does not affect the uninsurable variance Ω^R . In particular:

$$\Omega_K^R(J_O) - \Omega_K^R(J) = \Omega^R(J_O) - \Omega^R(J).$$

Using the last two displayed equations, it follows that there exists a sufficiently large $K > 0$ such that the inequality in (A.4) holds, proving the claim.

A.3 Omitted proofs for Section 5

Proof of Theorem 2. Consider an alternative economy in which $\hat{J}_O \cup J_N$ is the set of old assets and $(J_O \cup J_N) \setminus (\hat{J}_O \cup J_N)$ is the set of new assets. Applying the second part of Theorem 1 to this alternative economy implies $\Omega^S(\hat{J}_O \cup J_N) \leq \Omega^S(J_O \cup J_N)$. In addition, the assumption that traders agree about old assets implies $\Omega^S(\hat{J}_O) = \Omega^S(J_O) = 0$ [cf. Eq. (25)]. Combining these observations proves the theorem.

Proof of Theorem 3. I claim that $\Omega^S(J_O \cup \{\hat{j}_N\}) = \Omega^S(J_O)$, that is, financial innovation does not increase the speculative variance in economy $\mathcal{E}(\hat{J})$. The result follows from this claim since financial innovation always (weakly) increases the speculative variance in economy $\mathcal{E}(J)$.

To prove the claim, note that Eq. (25) implies:

$$\Omega^S(J_O \cup \{\hat{j}_N\}) = \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{\tilde{\boldsymbol{\mu}}_i'}{\theta_i} \begin{bmatrix} \Lambda(J_O) & 0 \\ 0 & \Lambda^{\hat{j}_N \hat{j}_N} \end{bmatrix}^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\theta_i}.$$

The off-diagonal entries are zero because the asset \hat{j}_N is not correlated with the assets in J_O . Since the matrix in the middle is block-diagonal, this equation further implies:

$$\begin{aligned} \Omega^S(J_O \cup \{\hat{j}_N\}) &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left(\frac{\tilde{\boldsymbol{\mu}}_i(J_O)'}{\theta_i} \Lambda(J_O)^{-1} \frac{\tilde{\boldsymbol{\mu}}_i(J_O)'}{\theta_i} + \tilde{\boldsymbol{\mu}}_i^{\hat{j}_N'} \left(\Lambda^{\hat{j}_N \hat{j}_N} \right)^{-1} \tilde{\boldsymbol{\mu}}_i^{\hat{j}_N} \right) \\ &= \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \left(\frac{\tilde{\boldsymbol{\mu}}_i(J_O)'}{\theta_i} \Lambda(J_O)^{-1} \frac{\tilde{\boldsymbol{\mu}}_i(J_O)'}{\theta_i} \right) = \Omega^S(J_O). \end{aligned}$$

Here, the second line follows since $\tilde{\mu}_i^{\hat{N}} = 0$ in view of the assumption that traders agree about the new asset. This completes the proof of the claim and the theorem.

A.4 Omitted proofs for Section 6

Derivation of market maker's expected profit. First note that trader i 's payoff from rejecting the market maker's offer is the certainty equivalent payoff from her endowment:

$$e + \mathbf{W}_i \boldsymbol{\mu}_i^{\mathbf{v}} - \frac{\theta_i}{2} \mathbf{W}_i' \boldsymbol{\Lambda}^{\mathbf{v}} \mathbf{W}_i. \quad (\text{A.6})$$

Next consider trader i 's payoff from accepting the offer, which is given by the certainty equivalent of her equilibrium net worth, n_i , net of the fixed fee, π_i . Using the equilibrium demand (20), trader i 's net worth can be written as:

$$n_i = e - \mathbf{x}_i' \mathbf{p} + \left[\mathbf{W}_i + \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda} \right] \mathbf{v}.$$

The certainty equivalent of this expression is given by:

$$\begin{aligned} e - \mathbf{x}_i' \mathbf{p} + \mathbf{W}_i \boldsymbol{\mu}_i^{\mathbf{v}} + \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \boldsymbol{\mu}_i \\ - \frac{\theta_i}{2} \mathbf{W}_i' \boldsymbol{\Lambda}^{\mathbf{v}} \mathbf{W}_i - \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) - \theta_i \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \boldsymbol{\lambda}_i. \end{aligned} \quad (\text{A.7})$$

Since the fixed fee makes the trader indifferent, it is equal to the difference of the expression in (A.7) from the expression in (A.6). That is:

$$\pi_i = -\mathbf{x}_i' \mathbf{p} - \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) + \theta_i \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\mu_i}{\theta_i} - \lambda_i \right).$$

Next consider the sum of the fixed fees over all i :

$$\begin{aligned} \sum_i \pi_i &= -\sum_i \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) + \sum_i \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} (\mu_i - \theta_i \lambda_i) \\ &= -\sum_i \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) + \sum_i \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\tilde{\mu}_i - \theta_i \tilde{\lambda}_i - \frac{1}{|I|} \bar{\theta} \sum_{i \in I} \frac{\mu_i}{\theta_i} \right) \\ &= -\sum_i \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) + \sum_i \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} (\tilde{\mu}_i - \theta_i \tilde{\lambda}_i) \\ &= \sum_i \frac{\theta_i}{2} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\mu}_i}{\theta_i} - \tilde{\lambda}_i \right). \end{aligned}$$

Here, the first line uses the market clearing condition (3); the second line uses the defini-

tion of $\tilde{\boldsymbol{\mu}}_i$ and $\tilde{\boldsymbol{\lambda}}_i$ in the last term; the third line uses the facts that $\sum_i \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\boldsymbol{\theta}_i} - \tilde{\boldsymbol{\lambda}}_i \right) = \mathbf{0}$ and that $\frac{1}{|\mathcal{I}|} \bar{\boldsymbol{\theta}} \sum_{i \in \mathcal{I}} \frac{\boldsymbol{\mu}_i}{\boldsymbol{\theta}_i}$ is a constant independent of i ; and the last line follows by simple algebra. This shows that the expected profits of the market maker are given by (29), completing the derivation.

Proof of Theorem 4. Part (i). Note that $\tilde{\boldsymbol{\mu}}_i(J) = 0$ for any $J \subset \mathcal{J}$, since traders agree about the asset payoffs. Using this observation, traders' expected profits (29) can be written as:

$$\sum_i \pi_i = \sum_i \frac{\boldsymbol{\theta}_i}{2} \tilde{\boldsymbol{\lambda}}_i \boldsymbol{\Lambda}^{-1} \tilde{\boldsymbol{\lambda}}_i.$$

From Eq. (25), this expression is equal to $c_1 - c_2 \Omega^R(J)$ for some constant c_1 and positive constant c_2 . Thus, maximizing $\sum_i \pi_i$ is equivalent to minimizing $\Omega^R(J)$. Finally, note from Eq. (25), that $\Omega^S(J) = 0$. This further implies $\Omega(J) = \Omega^R(J)$. It follows that the market maker innovates assets that minimize $\Omega(J)$.

Part (ii). Note that

$$\tilde{\boldsymbol{\mu}}_{i,K}(J) = K \tilde{\boldsymbol{\mu}}_i(J) \tag{A.8}$$

for each K and J . I first claim that there exists sufficiently large \bar{K}' such that for each $K > \bar{K}'$, the optimum set, J , satisfies the property:

$$\tilde{\boldsymbol{\mu}}_i(J) \neq \mathbf{0}. \tag{A.9}$$

First note that there exist such a subset of \mathcal{J} (since traders disagree about some assets). Next, note that the expected profits for subsets with this satisfy:

$$\begin{aligned} \lim_{K \rightarrow \infty} \frac{1}{K^2} \sum_i \pi_{i,K} &= \lim_{K \rightarrow \infty} \sum_i \frac{\boldsymbol{\theta}_i}{2} \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\boldsymbol{\theta}_i} - \frac{\tilde{\boldsymbol{\lambda}}_i}{\mathbf{K}} \right) \boldsymbol{\Lambda}^{-1} \left(\frac{\tilde{\boldsymbol{\mu}}_i}{\boldsymbol{\theta}_i} - \frac{\tilde{\boldsymbol{\lambda}}_i}{\mathbf{K}} \right) \\ &= \lim_{K \rightarrow \infty} \sum_i \frac{\boldsymbol{\theta}_i}{2} \frac{\tilde{\boldsymbol{\mu}}_i}{\boldsymbol{\theta}_i} \boldsymbol{\Lambda}^{-1} \frac{\tilde{\boldsymbol{\mu}}_i}{\boldsymbol{\theta}_i} > \mathbf{0}. \end{aligned} \tag{A.10}$$

Here, the first line uses Eqs. (29) and (A.8) and the second line uses the fact that $\tilde{\boldsymbol{\mu}}_i \neq \mathbf{0}$ and $\lim_{K \rightarrow 0} \frac{\tilde{\boldsymbol{\lambda}}_i}{K} = 0$. From the last inequality, it follows that the profits from subsets with property (A.9) limit to ∞ as K limits to ∞ . As an alternative, consider a subset that does not satisfy property (A.9), i.e., suppose $\tilde{\boldsymbol{\mu}}_i(J) = \mathbf{0}$. In this case, Eq. (29) shows that the market maker's profits are given by $\sum_i \frac{\boldsymbol{\theta}_i}{2} \tilde{\boldsymbol{\lambda}}_i \boldsymbol{\Lambda}^{-1} \tilde{\boldsymbol{\lambda}}_i$, which is a finite number. It follows that the optimum set satisfies property (A.9) for sufficiently large K , proving the claim.

Next suppose $K > \bar{K}'$. Without loss of generality, restrict attention to sets, J , that satisfy

property (A.9) holds at the optimum. Using Eqs. (24) and (25), note also that:

$$\Omega_K^S(J) = K^2 \Omega^S(J) \text{ and } \Omega_K^R(J) = \Omega^R(J).$$

Using these equations and $\Omega = \Omega^R + \Omega^S$, note:

$$\begin{aligned} \lim_{K \rightarrow \infty} \frac{1}{K^2} \Omega(J) &= \lim_{K \rightarrow \infty} \Omega^S(J) + \frac{\Omega_K^R(J)}{K^2} \\ &= \Omega^S(J) = \sum_i \frac{\theta_i}{2} \frac{\tilde{\mu}_i}{\theta_i} \Lambda^{-1} \frac{\tilde{\mu}_i}{\theta_i}, \end{aligned}$$

where the last equality uses Eq. (25). Substituting this into (A.10) implies:

$$\lim_{K \rightarrow \infty} \frac{1}{K^2} \sum_i \pi_{i,K} = \lim_{K \rightarrow \infty} \frac{1}{K^2} \Omega(J),$$

where the limits are finite. It follows that, for sufficiently large K , maximizing $\sum_i \pi_{i,K}$ is equivalent to maximizing $\Omega(J)$. More specifically, there exists $\bar{K} > \bar{K}'$ such that for each $K > \bar{K}$, the optimum set, J , maximizes the average variance, $\Omega(J)$. This completes the proof of the theorem.

A.5 Omitted proofs for Section 7

Proof of Theorem 5. Note that the speculative variance limits to zero in either case. Thus, it suffices to show that:

$$\lim_{t \rightarrow \infty} \Omega^R(t; J_O \cup J_N) \leq \lim_{t \rightarrow \infty} \Omega^R(t; J_O).$$

To see this, first note by Eqs. (31) and (32) that $\lim_{t \rightarrow \infty} \Lambda(t) = \Lambda_{true}$. Using Eq. (24) and Theorem 1, it follows that new assets reduce the uninsurable variance in the long run, completing the proof.

Proof of Theorem 6. I first complete the proof provided in the text for the case in which $*_{true}$ is diagonal. In view of assumption (A1^D), traders belief updating in each asset is independent of the availability of other assets. In particular, using Eqs. (25) and (35), trader i 's belief difference for asset j at some period $t \geq t^j$ is given by:

$$\tilde{\mu}^j = \frac{\tau}{\tau + t - t^j} \tilde{\mu}_{i,prior}^j.$$

Note also that the variance of traders' beliefs in period t is given by

$$\Lambda(t) = \Lambda_{true} \left(1 + \frac{1}{\tau + t - t^j} \right),$$

which is a diagonal matrix. Using the last two displayed equations, the speculative variance (25) can be simplified. In particular, the contribution of each asset j to the sum, $\sum_t \Omega^S(t)$, can be calculated as:

$$\sum_{t=t^j}^{\infty} \frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{1}{\theta_i^2} \left(\frac{\tau}{\tau + t - t^j} \right)^2 \left(\frac{1}{\tau + t - t^j} + 1 \right) \frac{\left(\tilde{\boldsymbol{\mu}}_{i,prior}^j \right)^2}{\Lambda_{true}^{jj}}.$$

A change of variables, $\tilde{t} = t - t^j$, shows that this sum is independent of t^j . It follows that the sum, $\sum_t \Omega^S(t)$, is also independent of t^j , and it is given by:

$$\left[\sum_{t=0}^{\infty} \left(\frac{\tau}{\tau + t} \right)^2 \left(\frac{1}{\tau} + 1 \right) \right] \left[\frac{1}{|I|} \sum_{i \in I} \frac{\theta_i}{\bar{\theta}} \frac{1}{\theta_i^2} \sum_{j \in J} \frac{\left(\tilde{\boldsymbol{\mu}}_{i,prior}^j \right)^2}{\Lambda_{true}^{jj}} \right]. \quad (\text{A.11})$$

This term converges to a finite constant because the second bracketed term is constant and the first bracketed term has the same convergence properties as the sum $\sum_{t=0}^{\infty} \frac{1}{t^2}$ (which converges to $\frac{\pi^2}{6}$).

I next provide a proof for the more general case in which Λ_{true} is not necessarily diagonal. Consider an arbitrary policy \mathbf{t} that satisfies the ordering $t^1 \leq t^2 \dots \leq t^{|J|}$ (which is without loss of generality, since the assets can always be relabeled). Recall that the uncertainty is captured by the m dimensional standard random variable \mathbf{v} . The idea is to construct an alternative set of assets, $\left\{ \tilde{a}^j = \left(\tilde{\mathbf{A}}^j \right)' \mathbf{v} \right\}_j$, which are (i) uncorrelated, and (ii) economy $\tilde{\mathcal{E}}^D$ with the alternative assets has the same equilibrium as the economy \mathcal{E} with the original assets for any policy $\tilde{\mathbf{t}}$ that also satisfies the ordering $\tilde{t}^1 \leq \tilde{t}^2 \dots \leq \tilde{t}^{|J|}$. Since the alternative assets are uncorrelated, from the first part it follows that the economy $\tilde{\mathcal{E}}^D$ has the same sum of speculative variance for any introduction policy. Since the two economies have the same equilibria, it follows that the original economy \mathcal{E}^D also has the same sum of speculative variance within the lass of policies that satisfy $\tilde{t}^1 \leq \tilde{t}^2 \dots \leq \tilde{t}^{|J|}$. In particular, the sum of speculative variance with the initial policy \mathbf{t} is the same as the sum of speculative variance with the immediate introduction policy $\tilde{\mathbf{t}} = \mathbf{0}$. Since the initial policy was arbitrary, this provides a proof of the theorem.

The remaining step is to construct the alternative set of assets, $\{\tilde{a}^j\}_j$, with properties (i) and (ii). This set is constructed recursively. Let $\tilde{\mathbf{A}}^1 = \mathbf{A}^1$. For each $k > 1$ and $k \leq |J|$, let $\tilde{\mathbf{A}}^k \in \mathbb{R}^m$ be given by the component of \mathbf{A}^k that is orthogonal to the subspace spanned by the

vectors $\{\tilde{\mathbf{A}}^1, \dots, \tilde{\mathbf{A}}^{k-1}\}$. Formally, let:

$$\tilde{\mathbf{A}}^k = \mathbf{I} - \tilde{\mathbf{A}}_{k-1} \left(\tilde{\mathbf{A}}'_{k-1} \tilde{\mathbf{A}}_{k-1} \right)^{-1} \tilde{\mathbf{A}}'_{k-1} \mathbf{A}^k,$$

where $\tilde{\mathbf{A}}_{k-1} = [\tilde{\mathbf{A}}^1, \dots, \tilde{\mathbf{A}}^{k-1}]$. Note that the assets $\{\tilde{a}^j\}_j$ satisfy property (i) because they are uncorrelated by construction. To show property (ii), consider an arbitrary period t with a set of available assets $k = J(t)$. The available assets in economies $\tilde{\mathcal{E}}^D$ and \mathcal{E}^D are respectively given by $\{\tilde{a}^j\}_{j=1}^k$ and $\{a^j\}_{j=1}^k$. By construction, for any $k \in \{1, \dots, |J|\}$, these two sets of assets can be written as linear combinations of one another. Consequently, these assets provide the same trading opportunities within period t . Their payoff realizations also provide the same information at the end of period t . It follows that the economies $\tilde{\mathcal{E}}^D$ and \mathcal{E}^D have the same equilibrium, completing the proof.

B Appendix: Speculation and Pareto Inefficiency

The CARA-Normal framework of the main text is useful to analyze the positive effects of new assets on the allocation of risks. However, this framework features no externalities, and the equilibrium is Pareto efficient. Nonetheless, the main mechanisms continue to apply in richer environments that may feature externalities. For example, if the traders are financial intermediaries (banks, for short) that do not fully internalize the social costs of their losses (or bankruptcies), then an increase in speculative variance may be socially undesirable. This appendix develops a stylized model along these lines in which financial innovation is Pareto inefficient. The key idea is that speculation between banks leads to a misallocation of capital because of banks' borrowing constraints. This leads to a reduction in welfare for banks' trading partners (captured by labor in this simple model), which does not enter banks' risk-return trade-off. New assets increase banks' speculation, which leads to a Pareto inefficiency (under appropriate assumptions).

B.1 Benchmark environment and equilibrium

Consider an economy with periods $t \in \{0, 1, \dots\}$, a single consumption good (dollar), and two factors of production (capital and labor). There are two infinitely lived banks, denoted by subscript $i \in \{1, 2\}$. There are also hand-to-mouth workers (one for each period) that live only in period t and that supply labor inelastically.

In period 0, production of the final good takes place in two sectors also denoted by $i \in \{1, 2\}$.

The production function for the final good is given by a standard Cobb-Douglas technology:

$$y_{t,i} = \frac{1}{2} k_{t,i}^\alpha l_{t,i}^{1-\alpha} \text{ for } i \in \{1, 2\}.$$

Here, $k_{i,t}$ denotes capital and $l_{i,t}$ denotes labor allocated to sector i . Labor in each sector is supplied inelastically by workers. To keep expressions simple, suppose $l_1 = l_2 = 1$. The factor prices are competitive:

$$R_{t,i} = \frac{\alpha}{2} k_{t,i}^{\alpha-1} \text{ and } w_{t,i} = \frac{1-\alpha}{2} k_{t,i}^\alpha. \quad (B.1)$$

Banks are initially allocated $n_{0,i}$ units of capital. Banks specialize in production in their own sector. That is, bank i can invest its capital only in sector i . Importantly, there is also a borrowing constraint: Banks cannot borrow capital from each other. Formally, banks production of the final good is not pledgeable. This friction generates a potential misallocation of capital across the two sectors. In particular, it is possible that $R_{t,1} \neq R_{t,2}$, i.e., marginal products of capital are not equated across the two sectors.

Suppose capital fully depreciates after production (for simplicity). At the end of period t , banks have $R_i n_{t,i}$ dollars. Banks choose to allocate their dollars between consumption, $c_{t,i}$, and saving, $s_{t,i}$. Banks have access to a storage technology that generates 1 unit of capital in period $t+1$ for each dollar invested in period t . Bank's capital in period $t+1$ is given by:

$$n_{i,t+1} = s_{t,i}. \quad (B.2)$$

Banks make their consumption-savings decisions to maximize their lifetime utility, given by:

$$\sum_{t=0}^{\infty} \beta^t \log(c_{t,i}).$$

The equilibrium in this model is a collection of allocations, $\{n_{t,i}, k_{t,i}, c_{t,i}\}_{t,i}$, and factor prices, $\{R_{t,i}, w_{t,i}\}_{t,i}$, such that banks choose their allocations optimally and factor markets clear [cf. Eq. (B.1)].

Consider the characterization of equilibrium. In view of log utility, banks choose to save a fraction of their net worth at date 0, that is: $s_{t,i} = \beta R_{t,i} n_{t,i}$. Plugging this into Eq. (B.2) and using the factor prices from (B.1), the evolution of bank i 's capital is given by:

$$n_{i,t+1} = \frac{\beta\alpha}{2} n_{t,i}^\alpha.$$

Consequently, starting from any initial condition, $(n_{0,1}, n_{0,2})$, each bank's capital allocation converges to $n^* = \left(\frac{\beta\alpha}{2}\right)^{\frac{1}{1-\alpha}}$.

Note that the steady-state features $R_1^* = R_2^* = \frac{1}{\beta}$ for each i . Put differently, cross-sectional allocation of capital is optimal in the steady-state. Banks' intertemporal optimization provides one mechanism to counter the cross-sectional misallocation of capital. Note, however, that this mechanism might be slow. In particular, starting from an initial condition with $R_1 \neq R_2$, it might take some time for this mechanism to equate the returns across the two sectors.

B.2 Idiosyncratic Shocks and Assets' Hedging Role

Suppose the economy starts on the steady-state. More specifically, consider the end of period -1 at which banks have R^*n^* dollars and they are about to make their consumption/saving decision for period 0. Suppose banks learn that in period 0 (and only in period 0), each sector will experience an idiosyncratic shock. In particular, banks' date 0 capital is given by:

$$n_{0,i} = n^* + s_{-1} + \varepsilon_i,$$

where $\varepsilon_1 = \bar{\varepsilon}v$ and $\varepsilon_2 = -\bar{\varepsilon}v$, where

$$v = \begin{cases} 1, & \text{with prob. } 1/2 \\ -1, & \text{with prob. } 1/2 \end{cases}. \quad (\text{B.3})$$

In view of log utility, banks continue to save the same amount as before, which leads to date 0 capital, $n_{0,i} = n^* + \varepsilon_i$. That is, banks' initial capital allocation becomes risky. The rest of the equilibrium is calculated as in the previous subsection.

In this setting, banks' utility is decreasing in the size of the shock, $\bar{\varepsilon}$. Consider also the aggregate output at date 0, which is constant and given by:

$$Y_0 = \frac{1}{2} (n_0 + \bar{\varepsilon})^\alpha + \frac{1}{2} (n_0 - \bar{\varepsilon})^\alpha.$$

It can be seen that the output is decreasing in $\bar{\varepsilon}$. In particular, from Jensen's inequality, the output is smaller than the case without the shock, n_0^α . Consequently, consumption of workers (which receive a constant share of output) is also lower. Importantly, the idiosyncratic shock adversely affects not just the banks but also the workers.

In this setting, suppose an asset that is perfectly correlated with ε_1 is available for trade at the end of date -1 . In particular, the asset pays $a = v$ [cf. Eq. (B.3)]. Let x_i denote bank 1's investment in the asset. Bank 1's date 0 capital is given by:

$$n_{0,i} = n^* + s_{-1,i} + \varepsilon_i + x_i a.$$

It can be checked that the asset's equilibrium price is given by $p = 0$ (by symmetry). Moreover,

banks' investment decisions are given by:

$$s_{-1,i} = s^* \text{ and } x_1 = -x_2 = -\bar{\varepsilon}.$$

That is, banks trade the asset to hedge their endowment risks fully, and they continue to save the same amount as before. Consequently, banks' date 0 capital is riskless and given by $n_{0,i} = n^*$. The rest of the equilibrium is identical to the case without the idiosyncratic shock. It follows that the introduction of the new asset increases aggregate output, and makes both the banks and the workers better off. This result illustrates the risk sharing benefits provided by new assets.

B.3 Speculation and Assets' Betting Role

Consider the same example with two differences. First, $\bar{\varepsilon} = 0$, i.e., banks do not face an idiosyncratic shock. Nonetheless, the asset a is still available for trade. Second, suppose banks have heterogeneous prior beliefs about the asset a . In particular, bank 1 believes that $a = 1$ with probability $\pi > 1/2$, while bank 2 believes that $a = -1$ with the same probability π .

In this case, the asset price is still given by $p = 0$ (by symmetry). The appendix derives banks' optimal portfolio allocation, which is given by:

$$s_{-1,i} = s^* \text{ and } x_1 = -x_2 = n^* (2\pi - 1).$$

That is, banks save the same amount as before but they also invest in the risky asset in line with their beliefs. Their level of investment is determined by a standard risk-return trade-off. In particular, the greater banks' belief disagreements (captured by a higher π), the more they invest in the risky asset. Consequently, banks' date 0 capital is given by:

$$\begin{cases} n_{0,1} = 2\pi n^* \text{ and } n_{0,2} = 2(1 - \pi)n^*, & \text{if } v = 1, \\ n_{0,1} = 2(1 - \pi)n^* \text{ and } n_{0,2} = 2\pi n^*, & \text{if } v = -1. \end{cases}$$

Note that, by symmetry, the aggregate output is independent of the realization of the shock and is given by:

$$Y_0 = \frac{1}{2} (2\pi n^*)^\alpha + \frac{1}{2} (2(1 - \pi)n^*)^\alpha.$$

It can also be seen that $Y_t < Y^*$ for each t [and that $\lim_{t \rightarrow \infty} Y_t = Y^*$]. Figure 3 illustrates the evolution of output. Intuitively, there will be misallocation of capital regardless of the realization of the shock (and regardless of which bank has "correct" beliefs).

The introduction of the new asset increases the ex-ante welfare of banks. However, it reduces the ex-ante welfare of workers in all periods. The question arises whether the introduction of




Figure 3: The top panel illustrates the effect of financial innovation on aggregate output. The bottom panel illustrates the effect on banks' capital levels.

the new asset could be Pareto inefficient. In particular, could a social planner prevent trade in the new asset (while also using lump sum transfers) to make all agents better off? The following proposition answers this question in the affirmative.

Proposition 1. *Consider the economy described in this section with belief disagreements (i.e., $\bar{\varepsilon} = 0$ and $\pi > 1/2$). Suppose a planner prevents trade in the new asset, a , and compensates banks by transferring T dollars to each bank at the end of date 0. Suppose the transfers are financed by $2T$ dollars of lump-sum taxes on date-0 workers. There exists $\bar{\alpha}$, such that if $\alpha < \bar{\alpha}$, then there is a transfer level $T(\alpha) > 0$ that makes all agents (including the workers in future periods) better off. In particular, when the share of labor in output is sufficiently high, the competitive equilibrium is Pareto inefficient.*

The key intuition behind this result is that banks' speculation combined with their future borrowing constraints represents a pecuniary externality on future workers. Regardless of which bank is right, there will be misallocation of capital and the wages of future workers will decrease. Banks do not take into account the effect of their betting on workers' future wages. When there are borrowing constraints, this pecuniary externality leads to a Pareto inefficiency.

References

- Allen, F. and D. Gale (1994), *Financial Innovation and Risk Sharing*, MIT Press, Cambridge, MA.
- Allen, F., S. Morris and A. Postlewaite (1993), "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, 61, p.206-29.
- Athanasoulis, G. and R. J. Shiller (2000), "The Significance of the Market Portfolio," *Review of Financial Studies*, 13(2), p.301-329.
- (2001), "World Income Components: Measuring and Exploiting Risk-Sharing Opportunities," *American Economic Review*, 91(4), p.1031-1054.
- Brock, W. H. and F. Wagener (2009), "More Hedging Instruments May Destabilize Markets," *Journal of Economic Dynamics and Control*, Volume 33, Issue 11, p.1912-1928.
- Calomiris, C. W. (1998), "Review of *Stock Market Crashes and Speculative Manias*," edited by Eugene N. White, *Journal of Economic History*, 58, p.614-617.
- Chen, J., H. Hong and J. C. Stein (2002), "Breadth of Ownership and Stock Returns," *Journal of Financial Economics*, Vol. 66, No. 2-3, p.171-205.
- Cochrane, J. H. (2001), "Asset Pricing," *Princeton University Press*, Princeton, NJ.
- Cuny, C. J. (1993), "The Role of Liquidity in Futures Market Innovations," *Review of Financial Studies*, 6, p.57-78.
- Cao, D. V. (2010), "Collateral Shortages and Asset Price Volatility with Heterogeneous Expectations," Working Paper, MIT.
- Demange, G. and G. Laroque (1995), "Optimality of Incomplete Markets," *Journal of Economic Theory*, 65, p.218-232.
- Duffie, D. and M. Jackson (1989), "Optimal Innovation of Futures Contracts," *Review of Financial Studies*, 2, p.275-296.
- Duffie, D. and R. Rahi (1995), "Financial Market Innovation and Security Design: an Introduction," *Journal of Economic Theory*, 65, p.1-42.
- Gennaioli, N., A. Shleifer and R. Vishny (2010), "Neglected Risks, Financial Innovation, and Financial Fragility," Harvard University Working Paper.
- Geanakoplos, J. (2009), "The Leverage Cycle," in Acemoglu D., K. Rogoff, and M. Woodford, eds., *NBER Macroeconomics Annual*.
- Harris, M. and A. Raviv (1993), "Differences of Opinion Make a Horserace," *Review of Financial Studies*, 6, p.473-506.
- Harrison, J. M., and D. M. Kreps (1978), "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, p.323-36.

- Horn R. A. and C. R. Johnson (2007), *Matrix Analysis*, Cambridge University Press, Cambridge, UK.
- Kandel, E. and N. D. Pearson (1995), “Differential Interpretation of Public Signals and Trade in Speculative Markets,” *Journal of Political Economy*, 103, p.831–72.
- Lintner, J. (1969) “The Aggregation of Investor’s Diverse Judgments and Preferences in Purely Competitive Security Markets,” *Journal of Financial and Quantitative Analysis*, 4, p.347-400.
- Milgrom, P. and N. Stokey (1982), “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, 26, p.17-27.
- Miller, E. (1977), “Risk, Uncertainty, and Divergence of Opinion,” *Journal of Finance*, Vol. 32, No. 4, p.1151-1168.
- Miller, M. H. (1986), “Financial Innovation: The Last Twenty Years and the Next,” *Journal of Financial and Quantitative Analysis*, 21, p.459-471.
- Rajan, R. (2005), “Has Financial Development Made the World Riskier?” Proceedings of the Federal Reserve Bank of Kansas City Symposium at Jackson Hole.
- Scheinkman, J. A. and W. Xiong (2003), “Overconfidence and Speculative Bubbles,” *Journal of Political Economy*, 111, p.1183-1219.
- Stiglitz, J. E. (1989), “Using Tax Policy to Curb Speculative Trading,” *Journal of Financial Services Research*, 3 , p.101-115.
- Stout, L. (1995) “Are Stock Markets Costly Casinos? Disagreement, Market Failure, and Securities Regulation,” 81, p.611-712.
- Simsek, A. (2011), “Belief Disagreements and Collateral Constraints,” Working Paper, Harvard University.
- Summers, L. H. and V. P. Summers (1989), “When Financial Markets Work Too Well: A Cautious Case for a Securities Transaction Tax,” *Journal of Financial Services Research*, 3(2-3), p.261-86.
- Van Horne, J. C. (1985), “Of Financial Innovations and Excesses,” *The Journal of Finance*, 40, p.620-631.
- Varian H. (1985), “Divergence of Opinion in Complete Markets: A Note,” *Journal of Finance*, 40, pp. 309-317.
- (1989), “Differences of Opinion in Financial Markets,” In *Financial Risk: Theory, Evidence and Implications*. C. Stone (ed.). Kluwer Academic Publications.
- White, E. N. (1996), *Stock Market Crashes and Speculative Manias*, in: The International Library of Macroeconomic and Financial History, Vol. 13, An Elgar Reference Collection, Cheltenham, UK; Brookfield, US.
- Working, H. (1953), “Futures Trading and Hedging,” *American Economic Review*, 43, p.314-

43.

Zapatero F. (1998), "Effects of Financial Innovations on Market Volatility When Beliefs are Heterogeneous," *Journal of Economic Dynamics and Control*, 22, p.597-626.