Why trade with Goldman?

James McLoughlin*

March 11, 2012

Abstract

Using a model of repeated trade between a long-lived, informed, price-discriminating market maker and risk averse traders with endogenous hedging demands, I first show that traders are weakly better off trading with an informed dealer, as they may learn something about an asset's value in the process of transacting. Second, while long-term incentives can induce an informed market maker to honestly reveal information and increase risk-sharing, they also enable the market maker to hide her information and extract more rents, reducing price informativeness. This less desirable outcome dominates with respect to both the parameter space and a selection criterion. Finally, measures of market quality, such as the transient component of price volatility (illiquidity), may not accurately reflect welfare.

*jmcloughlin@haas.berkeley.edu, Ph.D. candidate, Haas School of Business, University of California, Berkeley. I thank Christine Parlour, Terry Hendershott, Willie Fuchs, Brett Green, Marcus Opp, Bradyn Breon-Drish, Barney Hartman-Glaser, and Sara Holland for their comments on earlier drafts. All errors are my own.

1 Introduction

The first thing you need to know about Goldman Sachs is that it's everywhere. The world's most powerful investment bank is a great vampire squid wrapped around the face of humanity, relentlessly jamming its blood funnel into anything that smells like money.
–Matt Taibbi, "The Great American Bubble Machine", Rolling Stone, 2009

People lost money in it, but the security itself delivered the specific exposure that the client wanted to have.
–Lloyd Blankfein, C.E.O. Goldman Sachs, 2010

In the financial markets literature, the strategic use of private information is largely viewed through the lens of equity markets. Much of our intuition rests on a paradigm in which competitive, uninformed market makers set prices against which a large number of agents, some of whom have private information, may trade. For example, a dealer quoting prices in a particular firm's stock is unlikely to have inside information about the firm's fundamentals, but may be exposed to adverse selection if the firm's employees or their associates do. However, a large number of asset classes, from traditional assets such as municipal bonds to more exotic asset-backed securities and derivatives, trade in over-the-counter (OTC) markets that do not fit this paradigm: price quotes are private, dealers have market power in setting prices, and it may be more reasonable to assume that they are the smart money.

If a market maker knows more about an asset's value than her clients, how do these clients avoid being taken advantage of when buying or selling the asset? Are they better off transacting with an uninformed dealer? Can longer term incentives, such as those promoted by recent financial regulation, improve client welfare by making the market maker more cooperative? To address these questions I posit a model of repeated trade between an informed market maker and uninformed risk-averse agents interested in hedging their endowment shocks. The market maker is able to use her information strategically because she has pricing power—for simplicity I assume she is a monopolist—and may be more or less focused on short-term profits depending on how she discounts future cash flows.

I find that while long-term incentives can increase risk sharing gains, they are potentially detrimental to price informativeness. In a static setting, as long as consumers of a market maker's risk sharing services are sophisticated—i.e. they have rational expectations—they are not easily exploited and may benefit by learning from an informed dealer's price quotes. Increasing a market maker's concern for future business can improve welfare by promoting more trade and informative prices, but this is not the only possible outcome. In fact, long-term incentives are more likely to increase the dealer's trading gains by allowing her to conceal information.

There are three distinct aspects of my model. First, the price makers (dealers) are more informed about the risky asset's value than the price takers (traders). This information structure may apply for OTC assets because dealers often have a role in designing the securities that satisfy their clients' hedging needs, giving them a better sense of what these assets are worth. For example, following the financial crisis, politicians, regulators, and the media censured investment banks, particularly Goldman Sachs, for betting against their clients using superior information about the value of mortgage-backed securities.[1] Even when dealers don't have superior information about an asset's fundamental value, they may have private information about inventory imbalances that give them an edge: Cao, Evans, and Lyons (2006) show private information about order flow can predict future prices changes in the foreign exchange market.

Second, imperfect competition allows the market maker to strategically use her informational advantage when setting prices. For more exotic OTC assets, there may only be a few players who can accommodate a given hedging need, and, even in more staid asset classes, there is empirical support for imperfect competition—Green, Hollifield, and Schurhoff (2007) find that a significant component of transaction costs in the municipal bond market is due to dealer market power.

Third, the market maker is a long-lived player who trades off shorter versus longer term profits depending on her incentives. I use the market maker's discount factor as a reduced form proxy for the extent to which she cares about current versus future gains. Traders in OTC markets know the counter-party they are dealing with, and should consider a dealer's incentives when forming beliefs about what she is or isn't willing to do.[2]

The beliefs traders entertain about asset value, given an observed price quote, drive the model's results. An informed dealer, who posts these quotes, has long-term incentives that change the trader's viable set of beliefs about asset value, and hence the possible trading outcomes, relative to the myopic case. In my model, a myopic risk-averse trader (the client) wants to hedge an endowment shock of the risky asset each period. For ease of exposition, I assume he wants to sell due to a positive shock. A long-lived, risk-neutral, monopolist

[1]See Taibbi (2009) and Hobson (2010) for media coverage, the U.S. Senate Permanent Subcommittee's analysis of the financial crisis (2010) for political investigations, and Securities and Exchange Commission (2010) for the SEC's criminal complaint against Goldman Sachs.

[2]The role of incentives in financial markets has also received a lot of recent attention: following the crisis regulators in the U.S. and abroad have recommended a shift from short to long-term incentive-based compensation schemes. For details, see proposed implementations of the Dodd-Frank act's section 956 and the U.K. Financial Services Authority's "Principle 5" in The Financial Services Authority (2009).

market maker knows whether the mean asset payoff is high or low, and bids accordingly, potentially signaling its value.

The market maker's preferred bidding strategy depends on the magnitude of hedging demand. First, consider a myopic market maker who only cares about current period profits. When the trader is relatively desperate—for example, if he is highly risk-averse—the market maker can bid low regardless of what she knows about the asset's value and trade occurs; the trader knows there is a chance the asset's value is high, but is still willing to trade at a low bid. In concealing her private information via a pooling bid, the market maker maximizes the trader's uncertainty and the corresponding insurance rents. When hedging demand is less extreme, this pooling strategy is untenable: given his prior beliefs, the trader is unwilling to trade at a lowball bid, and the dealer cannot commit to posting a higher bid because she loses money when the asset value is low. Any high bid therefore reveals the asset value is high. Signaling the asset value via this separating strategy is the only way trade occurs, so the market maker is willing to reveal her information in order to get the client's business. Her profits are lower than those of the pooling strategy: by revealing information, she reduces trader uncertainty and the corresponding insurance rents. She also has to forgo trade in the low-value state because the trader never trusts a low bid.

The myopic case above demonstrates that when traders have rational expectations, they are not easily exploited. If anything, trading with an informed market maker allows them to learn about asset value, and they are therefore weakly better off trading with an informed rather than an uninformed monopolist. In the separating equilibrium, the market maker is engaging in a form of information sale à la Admati and Pfleiderer (1990): she reveals her information, alleviating some of the trader's uncertainty, in exchange for the residual risk sharing gains.

Long-term incentives yield two additional types of equilibria with higher gains from trade. The first is the type of Pareto improvement we usually associate with inter-temporal incentives: low bids are not trustworthy in the static separating equilibrium, but if future risk sharing gains are large enough, a market maker who cares about the long-term can post a credible low bid when the asset value is indeed low. If she deviates by lowballing when asset value is high, traders will no longer trust her, and she will lose out on these additional profits going forward. This separating outcome is Pareto improving because the trader is no worse off—he still learns about asset value and reduces his uncertainty—and the market maker extracts the additional gains associated with trade in the low value state[3].

[3]This outcome is the easiest to support. Equilibria also exist where the trader receives some or all of

Long-term incentives also produce pooling equilibria that are not Pareto improving. When the market maker cares exclusively about current profits, high bids unambiguously reveal her information. However, if hiding her information leads to significant future profits, the market maker can credibly post a high bid that is uninformative. She is willing to incur small losses today in order to keep her information hidden in future periods, some of which will involve large gains on the spread between her pooling bid and the high asset value. Relative to the myopic separating outcome, gains from trade are again increased, but not in a Pareto improving way: because the market maker can now conceal her information, she captures additional risk sharing rents *and* recovers the gains associated with sharing her information in the separating strategy.

When we care about price informativeness in addition to total insurance gains, long-term incentives enable a socially undesirable pooling outcome. How likely is it that this equilibrium occurs relative to the preferable separating improvement? An increase in long-term incentives may be desirable if it enables a shift from the static/constrained separating equilibrium to a full trade separating equilibrium, but not the pooling equilibrium. It turns out this is never the case: generically, any parametrization supporting a separating improvement supports the pooling equilibrium, but the converse isn't true.

As usual, inter-temporal incentives support many possible equilibria in the repeated game. I use the notion of *undefeated* equilibria in Mailath, Okuno-Fujiwara, and Postlewaite (1993) to select the market maker's maximum profit outcome.[4] Intuitively, since the market maker moves first, if she posts a bid in line with her preferred pooling strategy, any separating equilibrium requires that the trader maintain "stubborn" beliefs: under his prior, he would trade, but instead makes the most pessimistic inference possible from the market maker's bid. These stubborn beliefs don't survive the *undefeated* criterion, and this strengthens the above result: long-term incentives only serve to make the market maker better off, at the trader's expense.

The media and the empirical microstructure literature often portray price volatility as a bad thing because it implies higher trading costs for investors.[5] Fixing parameters, I analyze how price volatility varies across the model's equilibria. Decomposing price volatility into its desirable and undesirable components—price discovery and transient illiquidity—I show the latter, a common measure of market quality, does not accurately reflect trader welfare.

the additional risk-sharing gains for a smaller set of parameters—these are also a Pareto improvement. For ease of exposition I focus on repeated game improvements that are most easily supportable.

[4]See Gomes (2000) for an application of this selection criterion in the finance literature.

[5]Recent discussion, for example, has focused on whether algorithmic / high-frequency trading has increased price volatility. See, for example, Chaboud, Chiquoine, Hjalmarsson, and Vega (2009)

The static, partial-trade separating equilibrium endogenously constrains trade at extreme prices, reducing the price-impact of trades and leading to a lower illiquidity measure than the Pareto-dominant full insurance outcome. Huang and Wang (2010) also show that transient volatility doesn't accurately reflect welfare, but they reach this conclusion through an entirely different channel—information is symmetric in their model and participation costs drive the welfare implications. Finding the same result herein, in a different context, makes the point that illiquidity is not an indicator of welfare more robust.

The lack of perfect competition and the endogenous hedging demand in my model are similar in spirit to Glosten (1989), which features a price-discriminating monopolist dealer, and the models of Biais, Martimort, and Rochet (2003), Bernhardt and Hughson (1997), and Dennert (1993), which feature oligopolistic dealers. All of these papers assume that the liquidity demand side is more informed, as is the case in most of the microstructure literature.

There is an earlier literature that considers the strategy of an informed market maker, or, equivalently, a "big" inside trader. Grinblatt and Ross (1985) and Gould and Verrecchia (1985) both analyze the linear pricing strategies of an informed dealer, but they assume she can commit to this strategy ex-ante, meaning she sometimes posts a price that implies expected losses. Laffont and Maskin (1990) use a static model like the myopic case herein to address this commitment problem, and show that prices can be informationally inefficient—pooling is a possible outcome. They restrict the insider to uniform pricing strategies, which constrains trade beyond the adverse selection dimension. I assume full price discrimination, which separates risk-sharing gains from learning effects, and is closer in spirit to how trade occurs in over-the-counter markets, where price schedules are not constant in quantity.[6] More recently, some models of limit-order markets, such as Goettler, Parlour, and Rajan (2009), do allow for informed trade originating from either the liquidity demand or supply side.

I discuss how a *dealer*'s long-term incentives affect prices and welfare in a financial markets context. Several papers show ways in which dealers can use a *trader*'s long-term concerns to their benefit. Seppi (1990) and Benveniste, Marcus, and Wilhelm (1992) look at whether monopolist dealers can screen out informed trade via the threat of reduced-form punishments, while Desgranges and Foucault (2005) explicitly model a dynamic game between a dealer and an informed trader. Bernhardt, Dvoracek, Hughson, and Werner (2005) use a reduced form model—prices are exogenous—along with supporting empirical

[6]For example, Edwards, Harris, and Piwowar (2007) analyze how corporate bond trading costs vary with order size.

evidence to show that dealers offer clients price improvements in exchange for their future order flow.

Finally, there is a significant corporate finance literature that highlights problems associated with managerial myopia (e.g. Stein (1988)), and suggests remedies that keep management focused on long-term shareholder value (e.g. Edmans (2009)). This paper provides an example of how myopia can improve, rather than harm, consumer welfare.

2 Model

Consider the following discrete-time infinite horizon game ($t = 0, 1, 2, \ldots$) in which a risk-neutral informed dealer quotes prices for an uninformed, strategic, risk-averse trader. In each period t, the asset produces an end of period payoff v_t:

$$v_t = v_{t-1} + \eta_t + \epsilon_t, \qquad \mathbb{P}(\eta_t = \Delta) = \mathbb{P}(\eta_t = -\Delta) = \frac{1}{2}$$

where ϵ_t is a zero-mean residual risk term with variance σ_ϵ^2 and η_t is the fundamental component of the asset payoff. The dealer knows η_t before any trade occurs; traders only see its value after payoffs are realized.[7] I refer to the monopolist dealer as M and the trader as T.

Before trade occurs, the trader receives an endowment shock $y_t \in \{-1, 1\}$ of the risky asset, where positive and negative shocks are equally likely. For simplicity, I assume the endowment shock is common knowledge, so that the trader's risk-sharing motive—whether he wants to buy or sell—is known before prices are set, so on a given trade the market maker only quotes one side of the market. The trader's end of period wealth w_t is

$$w_t = (y_t + q_t)v_t - q_t P(q_t)$$

where $q_t \in \{-1, 0, 1\}$ is the amount of the risky asset purchased by the trader from the market maker at price $P(q_t)$. I assume traders choose portfolios myopically, ignoring intertemporal hedging demand, or equivalently that they are short-lived and the history of the game is publicly observable. When traders are indifferent, I assume trade occurs. Trader utility V_t^T is a combination of mean-variance preferences over end-of-period wealth and any benefit from learning about the asset payoff:

[7] This is for simplicity to retain perfect public monitoring of the market maker's actions.

Assumption 1. Trader utility V_t^T is given by

$$V_t^T = \mathbb{E}w_t - \frac{\gamma}{2}\mathbf{V}ar(w_t) + L(\mathbf{V}ar(v_t) - \mathbf{V}ar(v_t|\text{price}))$$

where $L(0) = 0$ and $L(\cdot)$ is strictly increasing.

The learning benefit $L(\cdot)$ is a reduced form way of capturing any social gains from price efficiency. For example, more informative prices may allow investors to hedge exposures in correlated assets more efficiently, or may encourage more efficient real investment.[8]

The informed market maker M sets a price schedule $P(q_t; \eta_t)$ that maximizes time-t profits $\pi_t = \mathbb{E}q_t(P(q_t; \eta_t) - v_t)$ and the present value of future profits discounted at rate $\beta \in (0, 1)$. Abstracting away from the time-value of money, the parameter β captures the extent to which the dealer's employees have an incentive to favor long-term firm value (e.g. via restricted stock/option grants or bonus clawbacks) over short term profits (e.g. those encouraged by annual cash bonuses). Denote her normalized utility at time t as

$$V_t^M = (1-\beta)\mathbb{E}\sum_{s=t}^{\infty} \beta^{s-t}\pi_s = (1-\beta)\mathbb{E}\pi_t + \beta\mathbb{E}V_{t+1}^M$$

Given that M has full market power in setting a price schedule, restricting trade to whole units of the asset does not sacrifice much generality in a qualitative sense, and greatly simplifies the specification of equilibrium beliefs; I discuss unrestricted quantities as an extension.[9] We can therefore think in terms of bid/ask prices for one unit of the asset:

$$B(\eta_t) = P(-1, \eta_t) \qquad A(\eta_t) = P(1, \eta_t)$$

After the asset pays off, all information becomes public, and $h^t = \{h^{t-1}, \eta_t, P_t, q_t, v_t\}$ denotes the history of realized play through the end of period t. This ex-post history includes the fundamental innovation η_t, so the game has perfect public monitoring. The trader has rational expectations: his expected utility and optimal portfolio choice are based on beliefs about the market maker's pricing strategy and any information revealed by a given price schedule realization. The history of prior play can affect his beliefs about the information revealed in prices. Let

$$\mu(P_t; h^{t-1})$$

[8] Leland (1992) shows that insider trading and the corresponding information impounded in prices can lead to a social welfare increase via more efficient real investment.

[9] In states of the world where trade is not possible with whole units, trade is possible at a constrained quantity in the more general case. However, the main welfare conclusions of the paper are unchanged.

denote the trader's belief that the fundamental is high ($\eta_t = +\Delta$) given this period's price and the history of prior play. This belief will be specified for all possible price realizations, not just those on the equilibrium path.

In general, the pricing strategy $P_t = P(\eta_t; h^{t-1})$ may involve randomization, in which case prices will partially reveal information about the fundamental. I restrict attention to pure pricing strategies that reveal either no information (pooling) or all information (separating).

Definition 1. A tuple $\{P^*(\cdot,\cdot), \mu^*(\cdot;\cdot), q^*(\cdot)\}$ is a perfect Bayesian equilibrium of the repeated game if, for any time t and any ex-ante history h^{t-1},

- Taking the market maker's pricing rule $P(\eta_t; h^{t-1})$ as given, the beliefs $\mu^*(P_t; h^{t-1})$ about the fundamental realization η_t are formed using Bayes' rule for any price P_t that occurs with non-zero probability under the pricing rule.

- The portfolio choice $q^*(P_t) \in \{-1, 0, 1\}$ maximizes the trader's expected utility over time t wealth given his endowment shock y_t, the observed price P_t, and corresponding belief $\mu^*(P_t; h^{t-1})$.

- The pricing rule $P(\eta_t; h^{t-1})$ maximizes the informed market maker's expected utility given the evolution of trader beliefs $\mu^*(\cdot;\cdot)$, optimal demand $q^*()$, and fundamental realization η_t.

Because information is perfectly revealed at the end of each period, I focus on simple trigger strategies that achieve efficient outcomes. As in most signaling games, beliefs are not restricted off of the equilibrium path without refinements, leading to multiple equilibria for some parameterizations, even in the static game. The *undefeated* selection criterion discussed in Mailath, Okuno-Fujiwara, and Postlewaite (1993) applies naturally to this game, and I use it as a refinement in both the static and dynamic cases.

3 The benefit of trading with an informed dealer

I first consider a myopic market maker and show that traders are weakly better off if she is informed. Because everything in the model is symmetric and endowment shocks y_t are common knowledge, I restrict attention to the bid-side of the market when describing results.[10] This side is relevant when $y_t = 1$: the trader has a positive endowment shock he

[10] With unknown endowments, signaling and beliefs are more complicated because both sides of the quote are relevant strategic choices for the market maker.

wants to hedge by selling the asset, and the market maker need only supply a price on the bid-side of the market. To emphasize that this is the bid side of the quote, I will denote the price P_t offered for one unit as B_t. When discussing empirical implications, I reflect these results to the ask-side of the market. I refer to the states $\eta = +\Delta$ and $\eta = -\Delta$ as the high and low asset value states throughout, and to the market maker as the high or low type in these cases. Proofs are in the appendix.

When $\beta = 0$, the market maker does not care about future consequences and equilibria correspond to static game outcomes. As a benchmark, consider the symmetric information case at time t—when M does not know whether asset value is high or low—given prior asset value realization v_{t-1}:

Lemma 1. *In the static symmetric information game where neither party know η_t, full risk sharing occurs, the market maker captures all gains from trade, and prices are uninformative.*

The trader wants to sell to hedge her endowment shock. Maximum gains from trade are achieved via full risk sharing, and the market maker can capture all of these gains by posting a bid $B^* = v_{t-1} - \frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$ that equals T's no-trade certainty equivalent. Prices are obviously uninformative because neither party knows the fundamental value's realization.

When the market maker is informed about the fundamental component η_t, she is better off, ex-ante, if she posts the uninformative (pooling) bid of Lemma 1. However, she can not always commit to this bid, and a separating strategy that reveals her information is the only viable alternative. Specifically, she is not willing to bid above the asset's expected value $\mathbb{E}v_t = v_{t-1} + \eta_t$, and this constraint may bind in the low-value state. The trader's risk aversion parameter characterizes when these two outcomes—pooling or separating—are possible:

Proposition 1. *Risk aversion thresholds γ_S and γ_P characterize outcomes in the static game:*

1. *A separating equilibrium $\mathcal{S}^0$ with partial trade—trade only occurs when asset value is high—exists for $\gamma < \gamma_S$.*

2. *A pooling equilibrium $\mathcal{P}^0$ featuring full trade exists for $\gamma \geq \gamma_P$.*

3. *$\gamma_P < \gamma_S$, so both equilibria exist for $\gamma \in [\gamma_P, \gamma_S)$*

4. *No mixed-strategy equilibria exist.*

When T is highly risk-averse, and correspondingly more desperate to trade, M can post a low bid regardless of η_t and capture all gains via full insurance. But when T is relatively risk tolerant, he is no longer willing to trade at such low bids unless they credibly reveal the asset value is low. The market maker generally prefers pooling bids that hide her information because they increase uncertainty and hence insurance rents she extracts. When this is not possible, a separating pricing strategy is her next best option.

Any high bid $B_H > v_{t-1} - \Delta$ at which trade is feasible credibly signals that the asset value is high to the trader. Maximum gains are extracted when M offers the trader's outside option:

$$B_H^* = v_{t-1} + \Delta - \frac{\gamma}{2}\sigma_\epsilon^2$$

The intuition is the same as in the symmetric case, except everything is conditional on $\eta_t = \Delta$; the expected payoff is higher, and the trader's uncertainty is reduced, lowering the risk premium M can charge. M will never post this bid when the asset value is low because she loses money on the trade. In the low-value state, any low bid $B_L \leq v_{t-1} - \Delta$ needs to be incentive compatible: M must not be tempted to post a low bid when the asset value is actually high. This is only possible if B_L^*, the equilibrium bid offered when $\eta_t = -\Delta$, results in no trade. Hence risk-sharing is only possible in the high value state.

Figure 1 shows the main structure of the static game on the equilibrium path.[11] To support a separating equilibrium, T correctly believes the asset value is high when B_H^* is posted ($\mu_H = 1$). He also has to believe that any low bid B_L at which trade could occur is not necessarily coming from the low-type—otherwise the high type will deviate to exploit this. Finally, bids that are low enough to prohibit trade, including the equilibrium bid $B_L^* < v_{t-1} - \Delta - \frac{\gamma}{2}\sigma_\epsilon^2$, are credible from the low type ($\mu_L = 0$). Under appropriate beliefs, this partial trade equilibrium is possible as long as γ is below γ_S: when risk aversion is too extreme, the high bid B_H^* is below the low asset value, so the low type will want to mimic this bid, breaking the consistency of T's beliefs.

The pooling equilibrium $\mathcal{P}^0$ is comparable to the symmetric information case in Lemma 1: when the symmetric information bid is below the low asset value $v_{t-1} - \Delta$, the informed market maker can credibly bid this amount in all states of the world and capture the full ex-ante risk sharing gains. The trader knows he may hit this bid in a high value state, but his hedging demand is large enough that he still prefers to trade.

Once the symmetric information bid is above the low asset value because hedging demand is lower, which occurs when $\gamma < \gamma_P$, it is no longer rational for the informed market maker to maintain this bid in the low state. In that case, the market maker can not cred-

[11] A full specification of off-equilibrium beliefs supporting Proposition 1 is given in the appendix.

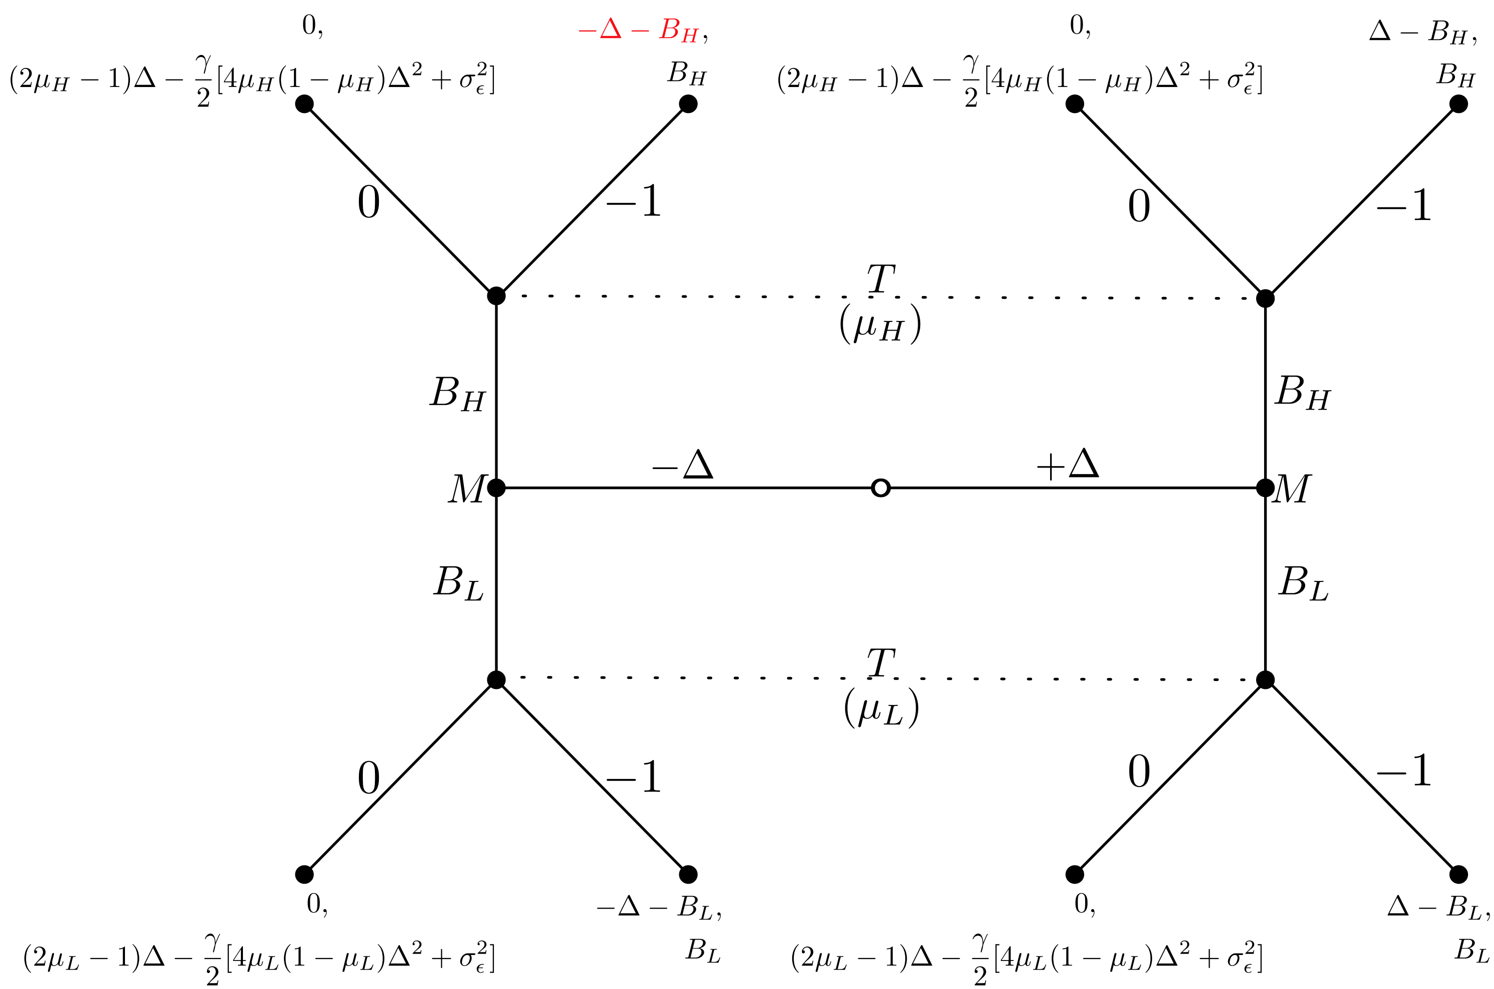

Figure 1: The static game: Nature chooses fundamentals; market maker M chooses either a high or low bid, where $B_H > -\Delta$ and $B_L \leq -\Delta$; trader T forms beliefs μ_H and μ_L about the probability the fundamental is high for each bid and chooses a quantity to trade, 0 or -1. Red font indicates a certain loss.

ibly bid low for the asset, so she foregoes gains in the low value state of the world while retaining them in the high value state via the separating strategy $\mathcal{S}^0$.

Note that any equilibrium has to involve trade in at least one state: under any beliefs the trader holds, a high-type market maker can, at a minimum, post a high bid that signals asset value and captures gains. Similarly, mixed strategy outcomes, which partially reveal information, are not possible. In these cases, M would occasionally signal the high asset value, but in any such equilibrium she always prefers posting a low-bid to exploit the trader's more generous beliefs.

The market maker's ex-ante profits are higher in the pooling equilibrium $\mathcal{P}^0$ because uninformative prices increase uncertainty and allow her to extract more risk-sharing gains. There is not a Pareto dominating outcome, so I use the *undefeated* selection criterion:

Lemma 2. *When $\gamma \in [\gamma_P, \gamma_S)$, so that both equilibria $\mathcal{P}^0$ and $\mathcal{S}^0$ are possible outcomes, $\mathcal{P}^0$ is the only undefeated equilibrium.*

When both equilibria are possible, separation only occurs if T's beliefs are highly pessimistic. He needs to stubbornly believe that any low bid is posted by the high type. This equilibrium is *defeated* because there is another equilibrium (pooling) in which that price is posted by both types, and, further, it is the equilibrium they prefer to play. This selection criterion can also be supported with a single set of "robust" beliefs across all parameterizations:

Corollary 1. *In the static game, the belief specification*

$$\mu^*(B) = \begin{cases} 1, & \forall B_t > v_{t-1} - \Delta \\ \frac{1}{2} & \forall B_t \in [\min_\mu \underline{u}(\mu), v_{t-1} - \Delta] \\ 0, & \textit{otherwise} \end{cases}$$

supports prices and quantities corresponding to the separating equilibrium $\mathcal{S}^0$ for $\gamma < \gamma_P$ and the pooling equilibrium $\mathcal{P}^0$ for $\gamma \geq \gamma_P$.

The separating threshold γ_S is increasing in Δ because higher information asymmetry allows the high type dealer to signal without tempting the low type to mimic this bid. It is decreasing in σ_ϵ^2 because higher residual risk implies a lower bid B_H^*, which is only above $v_{t-1} - \Delta$ if risk-aversion is lower. The static pooling threshold γ_P is monotonic in residual risk, but non-monotonic in fundamental risk:

Remark 1. The comparative statics of γ_P are

$$\frac{\partial \gamma_P}{\partial \sigma_{\epsilon^2}} < 0 \qquad \frac{\partial \gamma_P}{\partial \Delta} \begin{cases} > 0 \text{ for } \Delta < \sigma_\epsilon \\ < 0 \text{ for } \Delta > \sigma_\epsilon \end{cases}$$

As σ_ϵ^2 increases, traders face more risk and are anxious to hedge, so the pooling threshold is lower. For the standard deviation of fundamentals Δ, there are two opposing effects, illustrated in Figure 2. When Δ is very small, the adverse selection problem is minimal and pooling bids are credible for most risk aversion parameters. As Δ increases towards σ_ϵ, the adverse selection effect dominates, pushing up the pooling threshold because only desperate traders with high risk aversion are willing to trade at the uninformative pooling bid. Finally, for high values of Δ, the trader's large fundamental risk exposure and the consequent hedging demand override the adverse selection effect, lowering the pooling threshold.

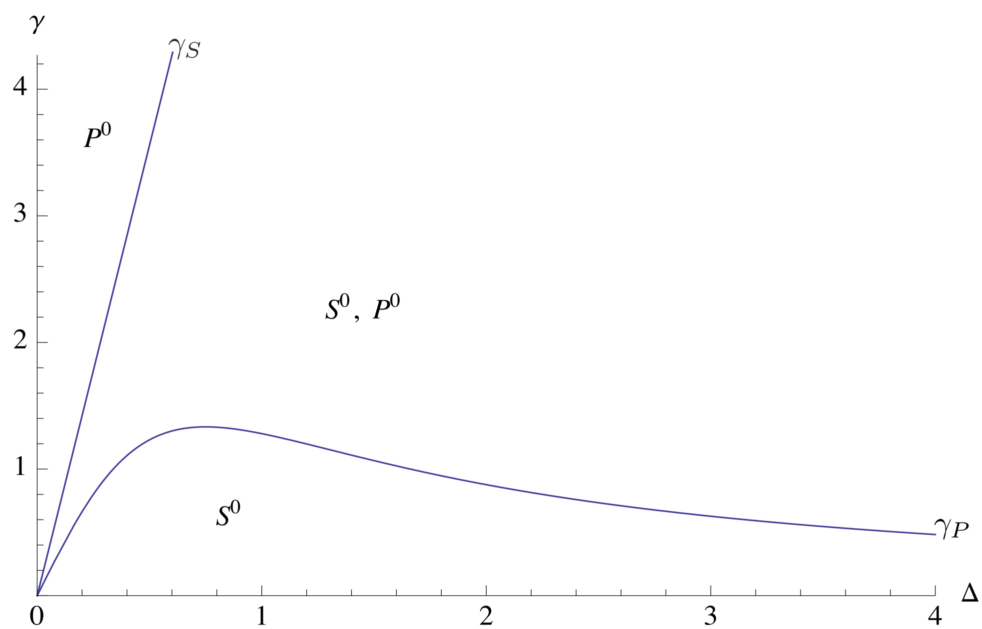

Figure 2: The pooling (γ_P) and separating (γ_S) thresholds as functions of the fundamental standard deviation Δ. Residual risk $\sigma_\epsilon = 0.75$. The pooling equilibrium $\mathcal{P}^0$ occurs above γ_P while the separating equilibrium $\mathcal{S}^0$ occurs below γ_S. In the region where both are feasible, $\mathcal{P}^0$ defeats $\mathcal{S}^0$.

When the pooling equilibrium is a possible outcome, traders are neither better or worse off transacting with an informed monopolist—she extracts the same insurance rents as an uninformed market maker (Lemma 1). For moderate risk aversion $\gamma < \gamma_P$, traders are actually better off facing an informed monopolist, because the adverse selection friction works in their favor. The market maker prefers something to nothing, and is willing to reveal her information via aggressive bids in the high-value state to achieve trade. She still extracts any residual risk-sharing gains, but traders learn something in the process and reduce their uncertainty.

4 The ambiguous effects of long-term incentives

When the market maker is myopic, the preceding results show that consumers (traders) are better off in the separating equilibrium that occurs for moderate risk aversion—even though M extracts all residual insurance gains, the information in prices reduces trader uncertainty. When hedging demand is high, M is able to maintain uninformative prices and extract the maximum possible rents. Do stronger inter-temporal incentives, represented by higher values of the discount rate β, improve trader welfare?

When hedging demand is high ($\gamma \geq \gamma_P$), the pooling outcome $\mathcal{P}^0$ is an equilibrium of the static game and therefore a sub-game perfect equilibrium of the repeated game. It is also the undefeated outcome of the repeated game because there is no credible threat that gives incentives any bite. I impose the following assumption to focus discussion on the additional gains achievable in a repeated context.

Assumption 2. Hedging demand is moderate, i.e. $\gamma < \gamma_P$, where γ_P is defined in Proposition 1.

The only static equilibrium under this restriction is separating and constrains the quantity traded in the low-value state to zero. This is also sub-game perfect in the repeated game and can be used to support more efficient equilibria.

Both equilibria I discuss have the same flavor: the market maker is able to credibly commit to a pricing strategy that achieves additional gains; traders believe this is the strategy being played until evidence arises to the contrary and make trades according to these beliefs; any deviation triggers reversion to the static separating equilibrium where these additional gains are lost. The static separating equilibrium $\mathcal{S}^0$ is the harshest playable trigger and supports efficient equilibria in the repeated game.

There are two ways to increase gains from trade in the repeated context: an additional pooling equilibrium involving small losses for M on bids in the low-value state, and an additional separating equilibrium in which she can credibly signal the asset value in the low state without constraining the quantity traded.[12] Given a choice, the market maker

[12]There are also a set of additional mixed strategy, partially revealing equilibria in between, but they are ignored for simplicity. They are also inferior to the pooling outcome from M's perspective.

prefers the pooling improvement while the trader prefers the separating improvement.

4.1 Incentives may encourage more informative prices

The first dynamic result shows that a Pareto improvement is possible for some parameterizations because M can credibly signal the asset value is low. Traders believe high and low bids correspond to the true value of the asset. If M takes advantage of this by bidding low when the asset value is high, T no longer believes these low bids, and play reverts to the static equilibrium. Under Assumption 2, the static separating equilibrium $\mathcal{S}^0$ is the market maker's outside option should she choose this deviation, implying her ex-ante continuation value $V^M_{\mathcal{S}^0}$ is

$$V^M_{\mathcal{S}^0} = \frac{1}{2} \cdot 0 + \frac{1}{2}\frac{\gamma}{2}\sigma^2_\epsilon = \frac{\gamma}{4}\sigma^2_\epsilon$$

In a full-trade separating equilibrium, she earns the residual spread in both states of the world for a continuation value of $\frac{\gamma}{2}\sigma^2_\epsilon$, so her incentive constraint in the low state is

$$(1-\beta)\frac{\gamma}{2}\sigma^2_\epsilon + \beta\frac{\gamma}{2}\sigma^2_\epsilon \geq (1-\beta)(2\Delta + \frac{\gamma}{2}\sigma^2_\epsilon) + \beta V^M_{\mathcal{S}^0}$$

The first term on the right-hand side is M's most profitable deviation: by posting a low bid when the asset value is high, she makes the spread between high and low asset values (2Δ) plus the residual risk sharing gains. The following proposition summarizes the results of this heuristic argument.

Proposition 2. *If the market maker has significant long-term incentives ($\beta > 2/3$), a full-trade separating equilibrium $\mathcal{S}^\infty$ of the repeated game exists when trader risk aversion γ exceeds the threshold γ^∞_S. γ^∞_S is decreasing in the market maker's long-term incentives β, increasing in information asymmetry Δ, and decreasing in residual risk σ^2_ϵ.*

The market maker is able to credibly signal low-asset value through her bid because she would forgo significant future gains from trade by deceiving the trader. However, this is only supportable when the trader's risk-aversion, and therefore hedging demand, are relatively high. When risk aversion is low, the future gains are not high enough to keep the market maker honest. Similarly, there is a lower bound on how much M cares about future gains, parameterized by β.

The comparative statics of the supporting risk-aversion threshold γ^∞_S are straightforward. It is decreasing in residual risk σ^2_ϵ because it implies higher hedging demand for a given risk aversion level, and hence more gains that are lost by deviating; it is increasing in fundamental risk Δ because this magnifies the benefit from deviating in the present period if the asset value is high; and it is decreasing in β because, as usual, more weight on future gains supports cooperation over a larger range of parameters.

The trader is no better off here, but he still has the same informational benefit—regardless of whether or not trade occurs in each state, both the repeated and static separating equilibria fully reveal the asset's value to him. The market maker achieves additional insurance rents in the low-value state, and for an outside observer, prices are more informative than the static separating case, where prices are unobservable in the low state.

4.2 Incentives also enable information to be concealed

The static pooling equilibrium $\mathcal{P}^0$ breaks down once the symmetric uninformed bid, which captures all ex-ante gains from trade, is above the low asset value and no longer credible for a low-type market maker. In that case, such a bid leads to an expected loss for the market maker when the asset value is low. In a repeated context, the market maker can credibly commit to incurring such losses in exchange for large gains in the high value state, again modulo a restriction on risk aversion so that future gains are large enough to support this strategy.

The market maker is tempted to abstain from bidding in the low value state, forgoing the expected losses in the present period at the expense of future profits. Her outside option is again the static separating continuation value $V^M_{\mathcal{S}^0}$. By staying on the equilibrium path, she always earns the full gains $\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$, so her IC constraint in the low state is

$$(1-\beta)[\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2) - \Delta] + \beta\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2) \geq (1-\beta)\cdot 0 + \beta V^M_{\mathcal{S}^0}$$

where the first term is her loss on the pooling bid—it's negative under Assumption 2. The solution to this constraint characterizes the repeated pooling outcome:

Proposition 3. *A pooling equilibrium $\mathcal{P}^\infty$ of the repeated game exists when trader risk aversion γ exceeds the threshold γ_P^∞, where $\gamma_P^\infty \leq \gamma_P$. γ_P^∞ is decreasing in the market maker's long-term incentives β, and its other comparative statics are analogous to the static threshold γ_P.*

In this case, the market maker's long term incentives allow her to extract full gains for more parameterizations because the pooling threshold is lower than in the static case. M's continuation utility allows her to credibly post a high bid in the low value state. Maximal insurance gains are achieved, but prices are uninformative and the trader is left with no share of the surplus. In this sense, pooling is an undesirable outcome, yet it is more easily supported when the market maker has long term incentives.

In sum, long-term incentives can have two effects: an informed market maker may be willing to honestly reveal asset value by signaling through prices, but she is also willing to incur losses in order to maintain uninformative prices. Is the uninformative pooling equilibrium $\mathcal{P}^{\infty}$ a legitimate concern? The following proposition shows that, in terms of the parameter space, it is the more likely result of long-term incentives:

Lemma 3. *The pooling threshold supported by repeated play, γ_P^{∞}, is always lower than the repeated separating threshold γ_S^{∞}. γ_S^{∞} is only below γ_P when the market maker has sufficient long-term incentives:*

$$\gamma_S^{\infty} < \gamma_P \iff \beta > \underline{\beta}$$

where $\underline{\beta} \in (\frac{4}{5}, 1)$. $\underline{\beta}$ is increasing in information asymmetry Δ and decreasing in residual risk σ_ϵ^2

The full-trade separating equilibrium $\mathcal{S}^{\infty}$ is harder to support because M's short-term gains from deviating in the low value state are significantly larger than the losses she incurs in the pooling equilibrium $\mathcal{P}^{\infty}$. The separating improvement is only possible for moderate risk aversion ($\gamma < \gamma_P$), so inter-temporal incentives (β) must be sufficiently high to support $\mathcal{S}^{\infty}$. When information asymmetry is high, M's profitable deviation is more attractive, so she needs to care more about long-term gains in order to resist it. When residual risk is high, future gains from staying on the equilibrium path are higher, and cooperation can be supported for lower long-term incentives.

Figure 3 shows the regions of the parameter space in which each equilibrium outcome is possible. It is the repeated-game analog of Figure 2. When the market maker has long term incentives ($\beta = 0.9$ in the figure), the repeated separating equilibrium $\mathcal{S}^{\infty}$ featuring maximal gains, a fair surplus split, and fully informative prices can be supported for some moderate risk aversion levels $\gamma \in [\gamma_S^{\infty}, \gamma_P)$. However, whenever $\mathcal{S}^{\infty}$ is possible, so is $\mathcal{P}^{\infty}$. Whether long-term incentives are desirable depends on how the equilibria discussed thus far compare to each other in a welfare sense.

4.3 Welfare Summary

The pooling equilibrium occurs for high risk aversion regardless of long-term incentives using the *undefeated* criterion. For more moderate risk aversion (Assumption 2), the partial-

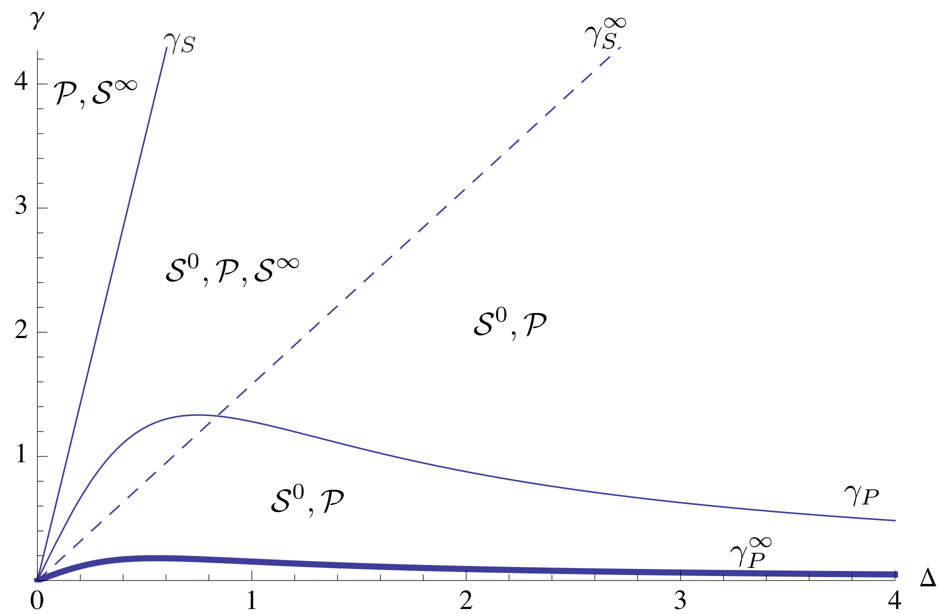

Figure 3: Static thresholds γ_P and γ_S (thin lines), γ_S^∞ (dashed line), and γ_P^∞ (thick line) as a function of the fundamental standard deviation Δ, for residual risk $\sigma_\epsilon = 0.75$ and discount factor $\beta = 0.9$. $\mathcal{P}$ denotes a region where either of the pooling outcomes, $\mathcal{P}^0$ or $\mathcal{P}^\infty$, are possible.

trade separating equilibrium $\mathcal{S}^0$ occurs absent long-term incentives. This equilibrium does not achieve full insurance, but does provide the trader with some gains via informative prices.

Do long-term incentives improve things in this parameter region? If risk aversion and incentives are high enough to support $\mathcal{S}^\infty$, and beliefs can be influenced such that this is the equilibrium played, then incentives can achieve a Pareto improvement—traders still learn from prices, and more insurance occurs.[13] However, the feasibility of $\mathcal{S}^\infty$ implies the pooling equilibrium $\mathcal{P}^\infty$ is also possible (Lemma 3). The converse is not true—in this sense, pooling improvements dominate separating improvements with respect to the underlying parameter space.

Applying the *undefeated* selection criterion in the repeated game—that traders interpret dealer actions in the context of all possible equilibria they might be playing—long-term incentives hurt the trader because they only enable $\mathcal{P}^\infty$. A move from $\mathcal{S}^0$ to $\mathcal{P}^\infty$ does generate additional gains from trade, but it allocates these gains to the market maker while reducing price informativeness.

[13] Depending on the beliefs, traders can also receive some of the additional insurance gains, but these outcomes are harder to support.

Proposition 4. *Total risk sharing gains are highest in the pooling equilibria* $(\mathcal{P}^0, \mathcal{P}^\infty)$, *while the trader surplus is highest in the separating equilibria* $\mathcal{S}^0$ *and* $\mathcal{S}^\infty$.

Equilibrium	*Dealer Surplus*	*Trader Surplus*
$\mathcal{P}^0, \mathcal{P}^\infty$	$\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$	0
$\mathcal{S}^0$	$\frac{\gamma}{4}\sigma_\epsilon^2$	$L(\Delta^2)$
$\mathcal{S}^\infty$	$\frac{\gamma}{2}\sigma_\epsilon^2$	$L(\Delta^2)$

There is a fundamental tradeoff between price efficiency and the amount of risk sharing that takes place. Because they both have the same welfare implications, let $\mathcal{P}$ denote either of the pooling equilibria. Thus far, I've argued that these pooling equilibria are socially undesirable because they are bad for the consumer and result in uninformative prices. When we care about the trader's take or price informativeness, these equilibria are inferior to the full-trade separating equilibrium ($\mathcal{S}^\infty \succ \mathcal{P}$). The repeated separating equilibrium is a Pareto improvement over the static separating outcome, so $\mathcal{S}^\infty \succ \mathcal{S}^0$.

I've used a monopolist market maker as a proxy for imperfect competition, but it is reasonable to argue that some of the additional gains generated in the pooling equilibrium would accrue to consumers under a more competitive market structure. To that end, consider a weighted social objective

$$V = V^T + \alpha V^M$$

where V^T and V^M are the ex-ante utilities of the trader and dealer in a given equilibrium, and $\alpha \in (0,1)$ denotes the level of competition in the dealer market. When $\alpha = 0$, $\mathcal{S}^0 \succ \mathcal{P}$. But even when $\alpha \in (0,1)$, there are parameterizations which result in the same social preference—as long as the market is not perfectly competitive, additional risk sharing gains enabled by the pooling equilibrium are potentially offset by the loss in price informativeness $L(\Delta^2)$. This occurs when information asymmetry (Δ), and hence the gain from informative prices, is large relative to residual risk σ_ϵ. As an example, consider the simple Pareto frontier in Figure 4 (where $L(x) = \frac{\gamma}{2}x$): in this case, if the introduction of long-term incentives induces a shift from $\mathcal{S}^0$ to $\mathcal{P}^\infty$, there is a small increase in total social gains at the expense of a large loss in trader welfare.

For a social planner instituting longer-term incentives, even if the selection criterion is relaxed, it's ambiguous which outcome will occur when both full-trade equilibria are possible. If the planner is not sure of the underlying parameters, it's more likely that incentives will support $\mathcal{P}^\infty$ rather than $\mathcal{S}^\infty$ because the former occurs for a strict superset of the parameter space. Hence it's possible that the introduction of these incentives only

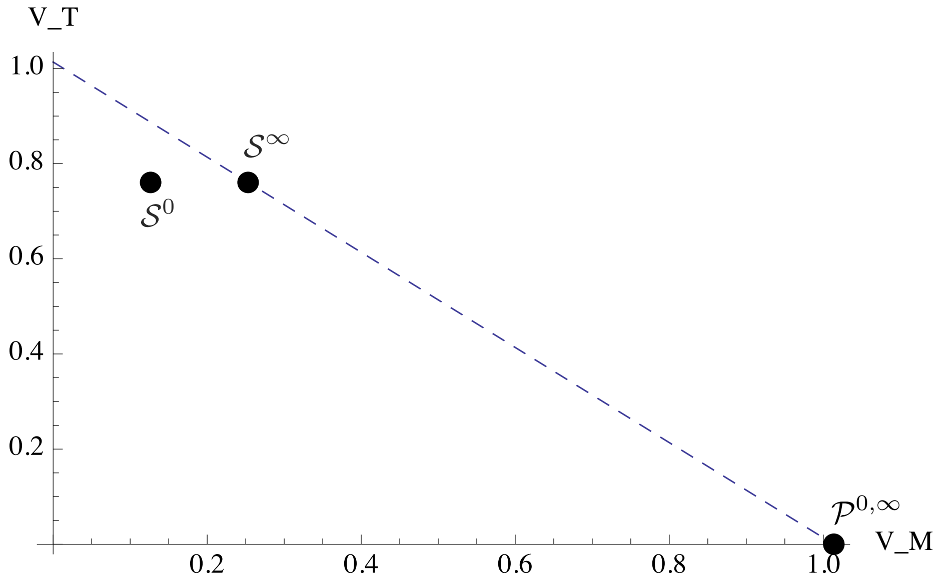

Figure 4: The Pareto frontier: ex-ante trader utility V^T (y-axis) vs. market maker utility V^M for model parameters $\sigma_\epsilon = 0.75, \gamma = 0.9, \Delta = 1.3$ and learning benefit $L(x) = \frac{\gamma}{2}x$

serves to make a less desirable pooling outcome possible—this is the case, for example, in the region of Figure 3 where $\mathcal{S}^0$ and $\mathcal{P}^\infty$ are the only viable outcomes.

The upshot, then, is that whether longer-term incentives make sense depends on a social tradeoff between risk sharing and price informativeness. When we care significantly about the latter, longer-term incentives can have negative welfare consequences, at least if competition is imperfect. In this case, it only makes sense to introduce them if some other policy will also influence beliefs and actions towards the Pareto improving outcome $\mathcal{S}^\infty$, or if some other transfer will restore the consumer's lost gains.

5 Welfare and price volatility

Given the welfare conclusions of the previous section, what can we expect to see in the data? Specifically, are standard measures of market quality aligned with welfare? For the sake of argument, fix the model parameters and suppose social preferences over equilibria are given by the example in Figure 4 ($\mathcal{S}^\infty \succ \mathcal{S}^0 \succ \mathcal{P}$)—the conclusions of this section are unchanged if we reverse the last two.

Now consider price volatility as a measure of market quality: recent media attention has focused on the presumably negative effects of commodity speculation and high-frequency trading on volatility—Chaboud, Chiquoine, Hjalmarsson, and Vega (2009) test the latter claim. Volatility is not necessarily a bad thing if it is associated with price discovery, and I will therefore focus on the transient "illiquidity" component of price volatility. In Roll's (1984) classic model of the bid-ask bounce, this measure corresponds to fixed trading costs or markups.

I have described the results so far from the bid side of the market, which occurs when the endowment shock is $y_t = 1$ and the trader wants to sell. This is without loss of generality because I've assumed the market maker knows the desired trade direction—she only has to quote one side of the market. This is often the case in OTC markets, as dealers are not necessarily required to quote both sides. The results on the ask side of the market are therefore symmetric, and that implies a data generating process for observed prices in each of the equilibria discussed so far:

$$\begin{aligned}
\mathcal{P} : P_t &= v_{t-1} - y_t \frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2) \\
\mathcal{S}^\infty : P_t &= v_{t-1} + \eta_t - y_t \frac{\gamma}{2}\sigma_\epsilon^2 \\
\mathcal{S}^0 : P_{\tau_n} &= v_{\tau_{n-1}} + \eta_{\tau_n} - y_{\tau_n} \frac{\gamma}{2}\sigma_\epsilon^2
\end{aligned}$$

where $\{\tau_n\}_{n=1}^\infty$ is the sequence of realized trade times for $\mathcal{S}^0$ because trade doesn't occur in every state. These price processes can be use to decompose the variance of price changes in each equilibrium into permanent and transient components:

Proposition 5. *For each equilibrium, let* Σ *denote the total variance of price changes and* λ *the transient component. Then,*

$$\begin{aligned}
\Sigma_{\mathcal{P}} &= (\Delta^2 + \sigma_\epsilon^2) + \underbrace{\frac{\gamma^2}{2}(\Delta^2 + \sigma_\epsilon^2)^2}_{\lambda_{\mathcal{P}}} \\
\Sigma_{\mathcal{S}^\infty} &= (\Delta^2 + \sigma_\epsilon^2) + \underbrace{\frac{\gamma^2}{2}\sigma_\epsilon^4}_{\lambda_{\mathcal{S}^\infty}} \\
\Sigma_{\mathcal{S}^0} &= 2(\Delta^2 + \sigma_\epsilon^2) + \underbrace{\frac{\gamma^2}{2}\sigma_\epsilon^4 - \Delta\gamma\sigma_\epsilon^2}_{\lambda_{\mathcal{S}^0}}
\end{aligned}$$

and $\lambda_{\mathcal{P}} > \lambda_{\mathcal{S}^\infty} > \lambda_{\mathcal{S}^0}$

The details are in the appendix, but, as in Roll (1984), the transient component λ corresponds to the spread between price and fundamental value, which reverts towards fundamentals. The ranking of λ across equilibria is not consistent with the posited social preferences. Note that comparing illiquidity across the two full-trade equilibria, $\mathcal{P}$ and $\mathcal{S}^\infty$ matches welfare: the pooling equilibria $\mathcal{P}$ is more illiquid because uninformative prices induce higher trader uncertainty, allowing the market maker to charge greater spreads. The confounding effect comes from the fact that $\lambda_{\mathcal{S}^\infty} > \lambda_{\mathcal{S}^0}$: the full-trade separating equilibrium is Pareto dominant, but displays higher illiquidity. The reason is that trade in $\mathcal{S}^0$ is constrained to preclude extreme price impact: sells at the bid only happen when the asset value is high, while buys at the ask only happen when it is low. Adverse selection endogenously constrains trade and dampens the transient component of price volatility.

Huang and Wang (2010) find an analogous result—that lower transient volatility does not necessarily represent higher welfare—via an entirely different channel. In their model, prices are competitive, information is symmetric, and the friction is participation costs that restrict the supply and demand of liquidity. Here, prices are not competitive and adverse selection is the cause of inefficiency. The common theme in both cases is that an endogenous constraint on trade simultaneously reduces both welfare and the transient component of price volatility. The illiquidity measure λ defined above also corresponds directly to the measure of corporate bond illiquidity proposed by Bao, Pan, and Wang (2009).[14].

In sum, decomposition of volatility into good and bad components does not accurately reflect underlying welfare[15]. This suggests that such measures of market quality are not necessarily a good intermediate policy objective, at least in markets characterized by the features of this model, and they may not be good indicators of whether a change in incentives has had a positive or negative effect on welfare.

6 Robustness and Extensions

6.1 Divisible shares

Thus far, I've assumed that trading the asset is an all or nothing proposition: the trader is unable to sell a fractional share of the asset.[16] This greatly simplifies the specification of equilibrium strategies and beliefs. In this section, I show that the main welfare implication of the paper—that inter-temporal incentives reduce trader welfare—is robust to this

[14]Their measure γ is the negative of the first-order auto-covariance of price changes, which equals $\lambda/2$

[15]The measure used here is strictly price based. A more nuanced decomposition of these components using both prices and order flow as in **?** does not substantially change the results: for a large portion of the parameter space, the ranking is unchanged.

[16]He will never be able to sell more than his share in equilibrium, as there are no gains beyond full insurance.

assumption.

Pooling outcomes $\mathcal{P}^0$ and $\mathcal{P}^\infty$ are the same when they are possible: M posts uninformative prices at which full insurance occurs. Similarly, the full-trade separating improvement $\mathcal{S}^\infty$, possible for high β, has full-trade in both states at the same prices. What changes is the static separating outcome $\mathcal{S}^0$: when the asset value is low, M can make an incentive compatible bid that results in trade. This is achieved by constraining the quantity traded in the low value state such that a high type market maker prefers to trade the full quantity at her high bid rather than deviate by bidding low—she makes a larger spread in doing so, but the reduced quantity negates this gain.

Proposition 6. *The risk aversion threshold $\hat{\gamma}_P$ characterizes outcomes in the static game with arbitrary quantities:*

1. *A full insurance pooling equilibrium $\mathcal{P}^0$ exists for $\gamma \geq \hat{\gamma}_P$.*
2. *A separating equilibrium $\mathcal{S}^0$ with partial trade always exists. Full residual insurance occurs in the high value ($\eta_t = \Delta$) state, while trade in the low value state is constrained to a quantity $Q_L^* \in (0,1)$.*

When quantities are unconstrained, M can fully price discriminate, so it is without loss of generality to consider her action space as offering a price-quantity pair (B, Q), where Q is the quantity she is willing to buy. Trader T then chooses an optimal sale quantity $q \in \{0, Q\}$—let q be positive when he is selling. Beliefs $\mu(B, Q; h_{t-1})$ that the asset value is high are formed over both price and quantity dimensions, and the equilibrium concept is identical to Definition 1, other than the change from an optimal pricing strategy $B^*(\eta_t; h_{t-1})$ to a price-quantity strategy $(B^*, Q^*)(\eta_t; h_{t-1})$.

Any high bid $B_H > v_{t-1} - \Delta$ credibly signals that the asset value is high (M loses money if she posts this in the low state), and maximum gains are extracted when M offers the trader's outside option:

$$(B_H^*, Q_H^*) = (v_{t-1} + \Delta - \frac{\gamma}{2}\sigma_\epsilon^2, 1)$$

In the low-value state, any bid $B_L \leq v_{t-1} - \Delta$ needs to be incentive compatible: M must not be tempted to use this bid when the asset value is actually high. This is possible if Q_L^*, the equilibrium quantity offered when $\eta_t = -\Delta$, is constrained.

First ignore incentive compatibility: given any quantity $q \in [0, 1]$, the low bid price $B_L(q)$ needs to meet the trader's participation constraint (IR_q) by offering his autarky

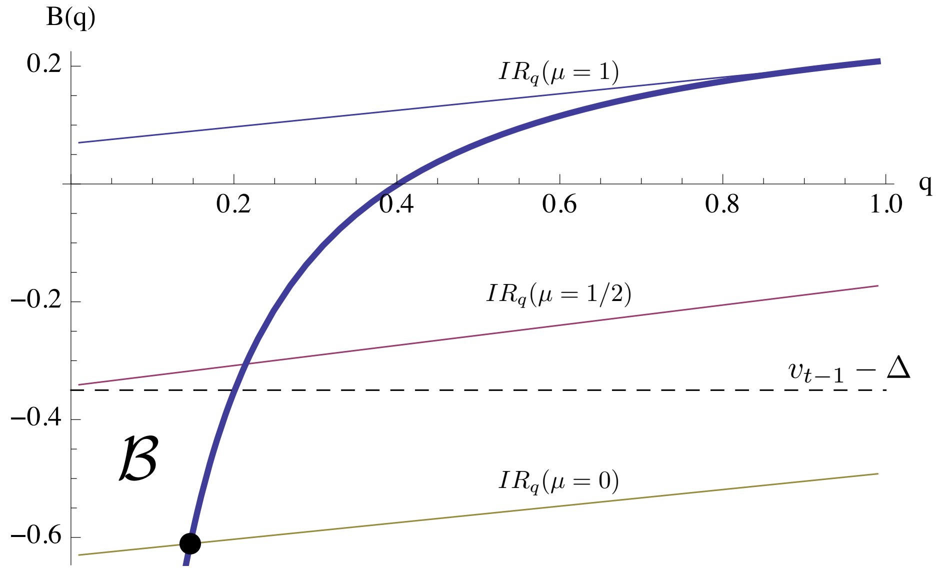

Figure 5: The static separating bid with partial shares. The thin slanted lines are the participation constraints $IR_q(\mu)$ for different trader beliefs μ. The horizontal dashed line denotes the low asset value $v_{t-1} - \Delta$—bids below this are profitable in the low state. The dark, curved line is the IC_q constraint. $\mathcal{B}$ denotes the region of incentive compatible, rational low-state bids. The intersection of $IR_q(\mu = 0)$ and IC_q is the optimal bid-quantity pair in the low state. Parameters: $\gamma = 0.5, \sigma_\epsilon = 0.75, \Delta = 0.35$.

value, implying:

$$B_L(q) \geq v_{t-1} - \Delta - (2 - q)\frac{\gamma}{2}\sigma_\epsilon^2 \qquad (IR_q)$$

This price will be incentive compatible for small q, as low bid profits $\pi_L(q)$ are lower than the full insurance profits $\pi_H^* = \frac{\gamma}{2}\sigma_\epsilon^2$ at (B_H^*, Q_H^*). However, as q increases, the market maker is tempted to deviate when the asset value is high and capture the much larger spread provided by B_L, requiring the incentive constraint $\pi_H^* \geq \pi_L(q)$ when asset value is high. This constraint implies a lower bound on B_L for a given quantity

$$B_L(q) \geq v_{t-1} + \Delta - \frac{\gamma}{2q}\sigma_\epsilon^2 \qquad (IC_q)$$

Figure 5 shows the relationship between these constraints: there is a $q^* \in (0, 1)$ such that for quantities higher than q^*, the IC constraint binds. Trader T believes that any bid-quantity pair above the IC constraint credibly signals the asset value is low (region $\mathcal{B}$). As q increases from 0, M only has to ensure participation, and profits are increasing. Once the incentive constraint binds at q^*, profits decrease. Hence she will select this quantity in the separating equilibrium, and bid such that the trader participates.

The modification of $\mathcal{S}^0$ when partial shares can be bought/sold does not qualitatively change the results regarding long-term incentives. The additional constrained trade in the low value state improves M's outside option, but this options supports both the separating (γ_S^∞) and pooling (γ_P^∞) equilibria. It is still the case that the pooling equilibrium $\mathcal{P}^\infty$ dominates $\mathcal{S}^\infty$ because $\gamma_S^\infty > \gamma_P^\infty$. The welfare conclusions that follow are unchanged.

6.2 Asymmetric fundamental shocks

To derive the main results of the paper, I use a model with symmetric, equally likely shocks to the privately known component of asset value η_t. I now relax this assumption and show the results of the paper are largely unchanged. Suppose asset value follows a process with asymmetric shocks

$$v_t = v_{t-1} + \eta_t + \epsilon_t, \qquad \mathbb{P}(\eta_t = \Delta_H) = p, \mathbb{P}(\eta_t = \Delta_L) = 1 - p$$

where $\Delta_H > \Delta_L$ and $p \in (0, 1)$. Under these dynamics, it is no longer without loss of generality to consider one side of the market exclusively. There is still symmetry in the sense that ask-side results when $\mathbb{P}(\eta_t = \Delta_H) = p$ correspond to bid-side results where $\mathbb{P}(\eta_t = \Delta_H) = 1 - p$. This implies different pooling thresholds for each side of the market:

Proposition 7. *Risk aversion thresholds γ_S, $\gamma_{P,A}$, $\gamma_{P,B}$ characterize outcomes in the static game with asymmetric fundamental shocks:*

1. *A separating equilibrium $\mathcal{S}^0$ with partial insurance—sells only occur when asset value is high, buys when asset value is low—exists for $\gamma < \gamma_S$. This threshold is independent of whether the trader is buying or selling.*

2. *A pooling equilibrium $\mathcal{P}_B^0$ featuring full insurance on the bid-side of the market ($y_t = 1$) exists for $\gamma \geq \gamma_{P,B}$.*

3. *A pooling equilibrium $\mathcal{P}_A^0$ featuring full insurance on the ask-side of the market ($y_t = -1$) exists for $\gamma \geq \gamma_{P,A}$.*

4. *A pooling equilibrium $\mathcal{P}^0$ featuring full insurance on the both sides of the market exists for $\gamma \geq \max\{\gamma_{P,A}, \gamma_{P,B}\}$.*

5. *γ_S is greater than both pooling thresholds, so pooling and separating are possible for $\gamma \in [\gamma_{P,A}, \gamma_S)$ on the ask-side and $\gamma \in [\gamma_{P,B}, \gamma_S)$ on the bid-side.*

6. *$\gamma_{P,A} < \gamma_{P,B} \iff p > \frac{1}{2}$.*

7. No mixed-strategy equilibria exist.

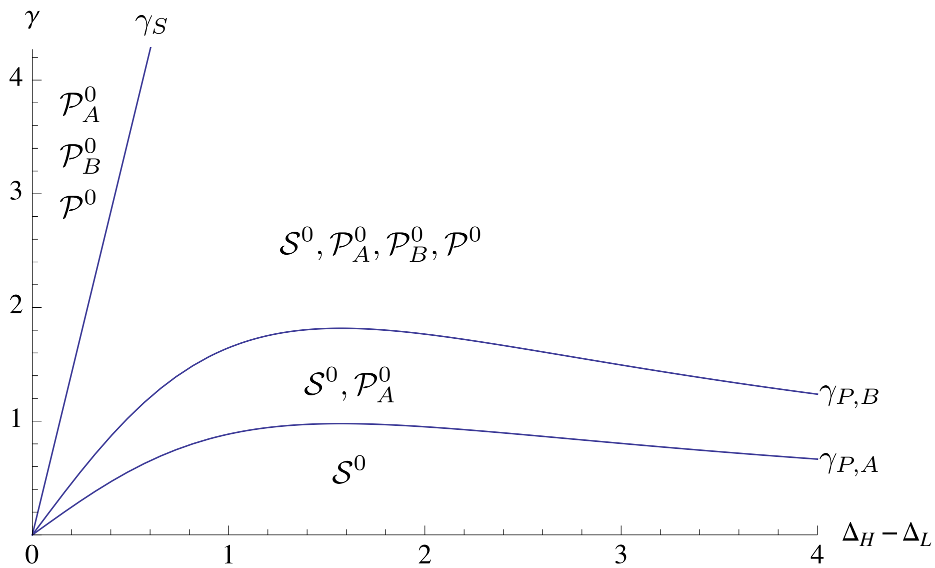

Figure 6: The pooling ($\gamma_{P,A}, \gamma_{P,B}$) and separating ($\gamma_S$) thresholds as functions of the private information spread $\Delta_H - \Delta_L$. Residual risk $\sigma_\epsilon = 0.75$ and high-state probability $p = 0.65$. Pooling on both sides of the market is possible for $\gamma > \gamma_{P,B}$, but only on the ask-side for $\gamma \in [\gamma_{P,A}, \gamma_{P,B})$. $\mathcal{P}^0$ is undefeated among all equilibria, while $\mathcal{P}_A^0$ and $\mathcal{P}_B^0$ defeat $\mathcal{S}^0$.

The effects of asymmetric fundamental shocks in the myopic case are illustrated in Figure 6. The results largely correspond to the symmetric case (Figure 2), but the spread $\Delta_H - \Delta_L$ determines the volatility of fundamentals and hence the degree of information asymmetry. The previous comparative statics, where this spread replaces the parameter Δ, are unchanged. The bid threshold $\gamma_{P,B}$ is increasing in p, while the ask threshold $\gamma_{P,A}$ is decreasing in p—as the trader is more sure the asset value is high, it is harder to support a low pooling bid.

The major difference is the asymmetry between when pooling is supported on each side of the market. In the figure, the region between the two pooling thresholds supports pooling at the ask, but not at the bid (for $p < 1/2$, the results are reversed). This region is largest when p takes on extreme values. The ask-side pooling outcome $\mathcal{P}_A^0$ defeats $\mathcal{S}^0$ because M prefers pooling where possible, so it will be the equilibrium in that region under my selection criteria, and the market maker has in improved outside option in the repeated game. I have not yet analyzed how this affects the feasibility of improvements in the repeated game. However, when $\gamma < \min\{\gamma_{P,B}, \gamma_{P,A}\}$, separating prices are still the only possible outcome on both sides of the market, so M's outside option is the same as in the symmetric case. In this region, the results are qualitatively the same as the symmetric information case:

Proposition 8. *In the repeated game, for* $\gamma < \min\{\gamma_{P,B}, \gamma_{P,A}\}$*, there exists a risk aversion threshold* γ_S^∞ *supporting a separating improvements on both the bid and ask sides of the market. Thresholds* $\gamma_{P,B}^\infty$ *and* $\gamma_{P,A}^\infty$ *support pooling improvements on the bid and ask sides, respectively, where*

$$\gamma_{P,B}^\infty < \gamma_S^\infty \qquad \gamma_{P,A}^\infty < \gamma_S^\infty$$

This confirms that, if the static separating equilibrium $\mathcal{S}^0$ is the market maker's outside option, whenever long-term incentives enable a separating improvement, they also enable a pooling improvement. While the separating outcome is a Pareto improvement, the pooling outcome reduces trader welfare. Hence the qualitative conclusions about the effects of long-term incentives on trader welfare remain unchanged.

7 Conclusion

Using a stylized model of repeated trade, I've shown that traders interested in hedging a risk are weakly better off trading with an informed dealer. If they've got rational expectations, the dealer can't use the information to her advantage, and may actually share what she knows to facilitate trade. Surprisingly, when the dealer has long-term incentives—she values future profits highly—traders may be worse off because the dealer is willing to incur short-term losses to conceal her private information.

Using simple assumptions on the dynamics of asset value, I decompose price volatility to show that liquidity measures do not necessarily reflect trader welfare (or price efficiency). Enriching the fundamental dynamics will allow for more precise statements about which types of assets are more likely to benefit or suffer from long term incentives in this context (for example those with bond-like versus stock like payoffs), as well as sharper empirical implications. For example, fundamentals could more closely track the dividends of a consumption tree, where the market maker knows whether consumption growth will be high or low. I also made a strong perfect monitoring assumption; relaxing it will not necessarily change the welfare results, but it will imply other interesting price dynamics: when traders experience significant losses, temporary periods of lower market liquidity may follow if they infer the market maker has deviated from a given pricing strategy.

A Appendix - proofs

Proof of Proposition 1. For a given bid B, let μ be trader T's belief about whether the asset value is high. His certainty equivalent $\underline{u}(\mu)$ from abstaining ($q = 0$) is

$$\begin{aligned} \mathbb{E}[v_t|\mu] &= v_{t-1} + (2\mu - 1)\Delta \\ \mathbf{V}ar[v_t|\mu] &= 4\mu(1-\mu)\Delta^2 + \sigma_\epsilon^2 \\ \Longrightarrow \underline{u}(\mu) &= v_{t-1} + (2\mu - 1)\Delta - \frac{\gamma}{2}[4\mu(1-\mu)\Delta^2 + \sigma_\epsilon^2] \end{aligned}$$

First note that any equilibrium has to involve trade in the high value state $\eta_t = \Delta$: T is willing to trade regardless of his beliefs if M's bid exceeds $\underline{u}(1)$, so M can profitably deviate from any no-trade outcome by posting this bid. Any outcome involving trade is either separating, pooling, or partially revealing (mixed).

1. **Separating Equilibrium** $\mathcal{S}^0$: Separation requires that M credibly reveal her information, which is only possible for bids $B_H > v_{t-1} - \Delta$ at which trade occurs because the low type always loses money on such a bid. Hence beliefs for high bids are $\mu([\underline{u}(1), \Delta]) = 1$ and M will extract maximum gains for these beliefs: $B_H^* = \underline{u}(1)$. These beliefs also need to hold for out-of-equilibrium bids in $(-\Delta, \underline{u}(1))$—otherwise high types have a profitable deviation by bidding on this interval, while low types do not, implying any such beliefs are inconsistent. Trade can not occur for any bid $B_L < -\Delta$, as high types will then have a profitable deviation posting B_L instead of B_H^*, so we can arbitrarily specify a no-trade bid $B_L^* < \underline{u}(0)$ for which beliefs are separating. Summarizing equilibrium bids, beliefs and trading strategies:

$$B^*(\eta_t) = \begin{cases} \underline{u}(1), \eta_t = \Delta \\ < \underline{u}(0), \eta_t = -\Delta \end{cases} \quad \mu^*(B) = \begin{cases} 1, \forall B \geq \underline{u}(0) \\ 0, otherwise \end{cases} \quad q^*(B) = \begin{cases} -1, B \geq \underline{u}(1) \\ 0, otherwise \end{cases}$$

   For some parameterizations, there are alternative out-of-equilibrium beliefs that support the same outcome, but the above hold generically and satisfy the "intuitive criterion" of Cho and Kreps (1987). Finally, this separating equilibrium is only feasible if $\underline{u}(1) > v_{t-1} - \Delta$, as otherwise the low type can also profitably post this bid, implying these beliefs are inconsistent with separating beliefs. This implies the parameter restriction $\gamma < \gamma_S$, where

$$\gamma_S \equiv \frac{4\Delta}{\sigma_\epsilon^2}$$

2. **Pooling Equilibrium** $\mathcal{P}^0$: In a pooling equilibrium, bids are uninformative, so T's beliefs are $\mu = \frac{1}{2}$ for equilibrium bids. M will extract the maximum possible gains

by bidding $\underline{u}(\frac{1}{2})$ in either state:

$$B^*(\eta_t) = \underline{u}(\frac{1}{2}) \quad \mu^*(B) = \begin{cases} 1, \forall B > v_{t-1} - \Delta \\ \frac{1}{2}, \textit{otherwise} \end{cases} \quad q^*(B) = \begin{cases} -1, B \geq \underline{u}(1) \\ -1, B \in [\underline{u}(\frac{1}{2}), v_{t-1} - \Delta] \\ 0, \textit{otherwise} \end{cases}$$

Again, off-equilibrium beliefs are consistent with the intuitive criterion, and their exist inconsequential modifications to beliefs supporting the same outcome. For this equilibrium to be viable, the high type needs to be willing to post the pooling bid in both states, which is only possible if $\underline{u}(\frac{1}{2}) \leq v_{t-1} - \Delta$ i.e.

$$\gamma \geq \gamma_P \equiv \frac{2\Delta}{\Delta^2 + \sigma_\epsilon^2}$$

3. By inspection, $\gamma_S > \gamma_P$.

4. Any mixed strategy equilibrium entails M signaling—posting a high bid when asset value is high—some, but not all, of the time. That implies that beliefs for any low bid $B_L \leq v_{t-1} - \Delta$ will assign positive probability to the asset being the high type. Trade can not occur at any such bid, because the high type would always prefer to trade at B_L. But if low bids do not involve trade, the high type will strictly prefer to signal via a high bid $B_H \geq \underline{u}(1)$ to capture some gains, implying a separating equilibrium.

□

Proof of Corollary 1. Note that the lower bound $\min_\mu \underline{u}(\mu)$ provides the lowest possible bid price at which trade could occur under any belief. For example, when $\gamma > 2/\Delta$, $\underline{u}(0) > \underline{u}(\frac{1}{2})$, and vice versa. When $\gamma < \gamma_P$, the pooling bid $\underline{u}(\frac{1}{2}) > v_{t-1} - \Delta$, so low bids $B_L \leq v_{t-1} - \Delta$ will not result in trade, and the only feasible outcome with trade is separating, which high types will choose. When $\gamma \geq \gamma_P$, pooling bids are viable and will be chosen by both types. This is consistent with the beliefs because the pooling bid $\underline{u}(\frac{1}{2}) \in [\min_\mu \underline{u}(\mu), v_{t-1} - \Delta]$. □

Proof of Proposition 2. The market maker's outside option is the static separating equilibrium $\mathcal{S}^0$ for $\gamma < \gamma_S$—for $\gamma \geq \gamma_S$, pooling is the only outcome and cannot be improved upon. Trade in a full-trade separating equilibrium occurs at both low and high bids, $B_L \leq v_{t-1} - \Delta$ and $B_H > v_{t-1} - \Delta$, and can only be maintained if M will not deviate by bidding B_L when asset value is high. Let s_L and s_H be the spreads she earns on each bid; so, $B_L = v_{t-1} - \Delta - s_L$. Her incentive compatibility constraint is then

$$(1-\beta)s_H + \frac{\beta}{2}(s_L + s_H) \geq (1-\beta)(2\Delta + s_L) + \beta\frac{\gamma}{4}\sigma_\epsilon^2$$

where the last term is the discounted ex-ante gains extracted for her outside option, providing a lower bound on the high spread

$$s_H \geq \frac{s_L(2-3\beta) + 4(1-\beta)\Delta + \beta\frac{\gamma}{2}\sigma_\epsilon^2}{(2-\beta)}$$

For trade to occur, bids need to satisfy the trader's participation constraints $B_L \geq \underline{u}(0)$ and $B_H \geq \underline{u}(1)$, hence both spreads s_L and s_H must be less than $\frac{\gamma}{2}\sigma_\epsilon^2$ implying

$$s_L(2-3\beta) \leq (1-\beta)(\gamma\sigma_\epsilon^2 - 4\Delta)$$

There are two cases to consider:

1. $\beta \leq \frac{2}{3}$: In this case,

$$s_L \leq \frac{(1-\beta)(\gamma\sigma_\epsilon^2 - 4\Delta)}{(2-3\beta)}$$

   but this spread has to be weakly positive, otherwise M is no better off than equilibrium $\mathcal{S}^0$, implying

$$\gamma \geq \frac{4\Delta}{\sigma_\epsilon^2} = \gamma_S$$

   for which $\mathcal{S}^0$ is not a credible outside option, and hence no separating improvement is possible.

2. $\beta > \frac{2}{3}$: In this case

$$s_L \geq \frac{(1-\beta)(\gamma\sigma_\epsilon^2 - 4\Delta)}{(2-3\beta)}$$

   which needs to be lower than the maximum spread $\frac{\gamma}{2}\sigma_\epsilon^2$. This is only possible if $\gamma \geq \gamma_S^\infty$, where

$$\gamma_S^\infty = \frac{8(1-\beta)\Delta}{\beta\sigma_\epsilon^2}$$

   and for any beliefs supporting separation, M maximizes profit via $s_L^* = \frac{\gamma}{2}\sigma_\epsilon^2$. The comparative statics of this threshold follow immediately.

Hence for $\beta > \frac{2}{3}$, the prices, beliefs and demands supporting the full trade separating equilibrium $\mathcal{S}^\infty$ are

$$B^*(\eta_t; h^{t-1}) = \begin{cases} \underline{u}(1), \eta_t = \Delta \\ \underline{u}(0), \eta_t = -\Delta \end{cases} \qquad \mu^*(B, h^{t-1}) = \begin{cases} 1, \forall B > v_{t-1} - \Delta \\ 0, otherwise \end{cases}$$

$$q^*(B) = \begin{cases} -1, B \geq \underline{u}(1) \\ -1, B \in [\underline{u}(0), v_{t-1} - \Delta] \\ 0, otherwise \end{cases}$$

for any history h^{t-1} which entails exclusive play of the above bidding strategy by M, and reversion to the static equilibrium $\mathcal{S}^0$ following any period in which M deviates by bidding low when asset value is high. □

Proof of Proposition 3. A pooling equilibrium of the repeated game that increases trade is only possible for $\gamma < \gamma_P$. Above that value, the static pooling equilibrium achieves the same gains. For such γ, $\mathcal{S}^0$ is M's outside option following any deviation. M will maintain the pooling bid, even at a loss in the low state, if future gains are high enough. For a pooling bid $B = v_{t-1} - s$, where $s < \Delta$ is the spread charged, M is tempted to abstain from posting this bid in the low state, requiring that

$$(1-\beta)(s-\Delta) + \beta s \geq (1-\beta)\cdot 0 + \beta\frac{\gamma}{4}\sigma_\epsilon^2$$

so that the spread must satisfy

$$s \geq (1-\beta)\Delta + \beta\frac{\gamma}{4}\sigma_\epsilon^2$$

Trade in any pooling equilibrium requires a bid $B \geq \underline{u}(\frac{1}{2})$, placing an upper bound on the spread:

$$(1-\beta)\Delta + \beta\frac{\gamma}{4}\sigma_\epsilon^2 \leq \frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$$

This condition is only satisfied if $\gamma \geq \gamma_P^\infty$, where

$$\gamma_P^\infty = \frac{4(1-\beta)\Delta}{2\Delta^2 + (2-\beta)\sigma_\epsilon^2}$$

It is easy to show $\gamma_P^\infty \leq \gamma_P$. A full specification of the actions and beliefs supporting this equilibrium ($\mathcal{P}^\infty$) is given by

$$B^*(\eta_t; h^{t-1}) = \underline{u}(\frac{1}{2}) \qquad \mu^*(B, h^{t-1}) = \frac{1}{2} \qquad q^*(B) = \begin{cases} -1, B \geq \underline{u}(\frac{1}{2}) \\ 0, \textit{otherwise} \end{cases}$$

for any history h^{t-1} which entails exclusive play of the above bidding strategy by M, and reversion to the static equilibrium $\mathcal{S}^0$ following any period in which M deviates by bidding low when asset value is high. The specified beliefs do not necessarily survive refinements, but any more nuanced specification results in the same outcome. □

Proof of Lemma 3. The fact that $\gamma_S^\infty > \gamma_P^\infty$ for all parameterizations follows from inspection of their definitions in Propositions 2 and 3. Comparing γ_S^∞ to the static pooling threshold γ_P defined in Proposition 1 produces the lower bound $\underline{\beta}$:

$$\underline{\beta} = \frac{4(\Delta^2 + \sigma_\epsilon^2)}{4\Delta^2 + 5\sigma_\epsilon^2}$$

□

Proof of Proposition 4. First consider the consumer surplus: in any pooling equilibrium, T receives $\underline{u}(\frac{1}{2})$ for the asset, which is equal to his autarky certainty equivalent, so the consumer surplus is zero. For either of the separating equilibria, T's ex-ante gain over autarky ($\underline{u}(1/2)$) is

$$\mathbb{E}[\underline{u}(\mu(B))] - \frac{\gamma}{2}\mathbf{V}ar(\underline{u}(\mu(B))) + L(\Delta^2) - \underline{u}(1/2) = L(\Delta^2)$$

In the pooling equilibrium, full gains are extracted by M for a surplus of $\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$. In the full-trade separating equilibrium $\mathcal{S}^\infty$, M extracts the residual risk sharing gains $\frac{\gamma}{2}\sigma_\epsilon^2$ in both states of the world, while in the partial trade separating equilibrium $\mathcal{S}^0$, residual risk sharing gains of $\frac{\gamma}{2}\sigma_\epsilon^2$ are lost in the low-state, which occurs half of the time, reducing ex-ante gains accordingly. □

Proof of Proposition 5. Negative endowment shocks $y_t = -1$, which are equally likely, lead to purchases at ask prices symmetric to the bids in the previous propositions, implying price processes for the various equilibria. Letting Σ denote the variance of price changes for each of them:

1. *Pooling equilibria* $\mathcal{P}^0, \mathcal{P}^\infty$: Prices are given by

$$P_t = v_{t-1} - y_t \frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$$

   implying price differences

$$P_t - P_{t-1} = \eta_{t-1} + \epsilon_{t-1} - (y_t - y_{t-1})\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2)$$

   All of the innovations are independent of each other, implying

$$\Sigma_{\mathcal{P}} = (\Delta^2 + \sigma_\epsilon^2)\left[1 + \frac{\gamma^2}{2}(\Delta^2 + \sigma_\epsilon^2)\right]$$

2. *Full-trade separating equilibrium* $\mathcal{S}^\infty$: Similarly, again using the independence of innovations, prices and changes are given by

$$\begin{aligned} P_t &= v_{t-1} + \eta_t - y_t \frac{\gamma}{2}\sigma_\epsilon^2 \\ P_t - P_{t-1} &= \eta_{t-1} + \epsilon_{t-1} + (\eta_t - \eta_{t-1}) - (y_t - y_{t-1})\frac{\gamma}{2}\sigma_\epsilon^2 \\ \Sigma_{\mathcal{S}^\infty} &= \Delta^2 + \sigma_\epsilon^2 + \frac{\gamma^2}{2}\sigma_\epsilon^4 \end{aligned}$$

3. *Partial-trade separating equilibrium* $\mathcal{S}^0$: For this equilibrium, trade does not occur in some states of the world, so there are periods where prices are not observed. Let $\{\tau_n\}_{n=1}^\infty$ denote the sequence of realized trade times. Prices are then given by:

$$P_{\tau_n} = v_{\tau_n - 1} + \eta_{\tau_n} - y_{\tau_n}\frac{\gamma}{2}\sigma_\epsilon^2$$

with differences

$$P_{\tau_n} - P_{\tau_{n-1}} = v_{\tau_n - 1} - v_{\tau_{n-1}-1} + (\eta_{\tau_n} - \eta_{\tau_{n-1}}) - (y_{\tau_n} - y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2$$
$$= \sum_{s=\tau_{n-1}}^{\tau_n - 1} (\eta_s + \epsilon_s) + (\eta_{\tau_n} - \eta_{\tau_{n-1}}) - (y_{\tau_n} - y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2$$

These differences still have mean zero, so the variance is given by

$$\Sigma_{\mathcal{S}^0} = \mathbb{E}\left\{\sum_{s=\tau_{n-1}}^{\tau_n - 1}(\eta_s + \epsilon_s)\right\}^2 + 2\mathbb{E}\left\{(\eta_{\tau_n} - \eta_{\tau_{n-1}}) - (y_{\tau_n} - y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2\right\}\sum_{s=\tau_{n-1}}^{\tau_{n-1}}(\eta_s + \epsilon_s)$$
$$+ \mathbb{E}\left\{(\eta_{\tau_n} - \eta_{\tau_{n-1}}) - (y_{\tau_n} - y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2\right\}^2$$

It's easiest to calculate these terms separately. Conditional on being in a trade state at time t, which occurs when $(\eta_t, y_t) \in \{(\Delta, 1), (-\Delta, -1)\}$, let $\mathcal{T}$ be the stopping time at which the next trade occurs. Trade occurs with probability $1/2$ independent of the current state, so $\mathbb{E}\mathcal{T} = 2$ (this is a trivially stationary, irreducible, finite-state Markov chain). Then

$$\mathbb{E}\left\{\sum_{s=\tau_{n-1}}^{\tau_n - 1}(\eta_s + \epsilon_s)\right\}^2 = \mathbb{E}\sum_{j=1}^{\infty}\mathbf{1}_{\mathcal{T}=j}\mathbb{E}\left\{\sum_{s=\tau_{n-1}}^{\tau_{n-1}+j}(\eta_s + \epsilon_s)\right\}^2$$
$$= \mathbb{E}\sum_{j=1}^{\infty}\mathbf{1}_{\mathcal{T}=j}\sum_{s=\tau_{n-1}}^{\tau_{n-1}+j}\mathbb{E}(\eta_s + \epsilon_s)^2$$
$$= \mathbb{E}\sum_{j=1}^{\infty}\mathbf{1}_{\mathcal{T}=j}\sum_{s=\tau_{n-1}}^{\tau_{n-1}+j}(\Delta^2 + \sigma_\epsilon^2)$$
$$= (\Delta^2 + \sigma_\epsilon^2)\mathbb{E}\sum_{j=1}^{\infty} j \cdot \mathbf{1}_{\mathcal{T}=j}$$
$$= (\Delta^2 + \sigma_\epsilon^2)\mathbb{E}\mathcal{T}$$
$$= 2(\Delta^2 + \sigma_\epsilon^2)$$

The second term is

$$2\mathbb{E}\left\{(\eta_{\tau_n} - \eta_{\tau_{n-1}}) - (y_{\tau_n} - y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2\right\}\sum_{s=\tau_{n-1}}^{\tau_{n-1}}(\eta_s + \epsilon_s) = 2\mathbb{E}\left(-\eta_{\tau_{n-1}}^2 + \eta_{\tau_{n-1}}y_{\tau_{n-1}}\frac{\gamma}{2}\sigma_\epsilon^2\right)$$
$$= 2(-\Delta^2 + \Delta\frac{\gamma}{2}\sigma_\epsilon^2)$$
$$= \Delta\gamma\sigma_\epsilon^2 - 2\Delta^2$$

The third term is

$$\mathbb{E}\left\{(\eta_{\tau_n}-\eta_{\tau_{n-1}})-(y_{\tau_n}-y_{\tau_{n-1}})\frac{\gamma}{2}\sigma_\epsilon^2\right\}^2 = \mathbb{E}(\eta_{\tau_n}-\eta_{\tau_{n-1}})^2-\gamma\sigma_\epsilon^2\mathbb{E}(\eta_{\tau_n}-\eta_{\tau_{n-1}})(y_{\tau_n}-y_{\tau_{n-1}}) + \frac{\gamma^2}{4}\sigma_\epsilon^4\mathbb{E}(y_{\tau_n}-y_{\tau_{n-1}})^2$$
$$= 2\Delta^2-2\Delta\gamma\sigma_\epsilon^2+\frac{\gamma^2}{2}\sigma_\epsilon^4$$
$$= \frac{1}{2}(2\Delta-\gamma\sigma_\epsilon^2)^2$$

where we've used the fact that $\mathbb{E}\eta_{\tau_n}y_{\tau_n}=\Delta$ conditional on trade occurring. Combining these three terms:

$$\Sigma_{\mathcal{S}^0} = 2(\Delta^2+\sigma_\epsilon^2)+\Delta\gamma\sigma_\epsilon^2-2\Delta^2+\frac{1}{2}(2\Delta-\gamma\sigma_\epsilon^2)^2$$
$$= 2(\Delta^2+\sigma_\epsilon^2)+\frac{\gamma^2}{2}\sigma_\epsilon^4-\Delta\gamma\sigma_\epsilon^2$$

A similar calculation of the first-order auto-covariance ϕ_1 produces:

$$\mathcal{P}:\phi_1 = -\frac{\gamma^2}{4}(\Delta^2+\sigma_\epsilon^2)^2$$
$$\mathcal{S}^\infty:\phi_1 = -\frac{\gamma^2}{4}\sigma_\epsilon^4$$
$$\mathcal{S}^0:\phi_1 = \frac{\gamma}{2}\Delta\sigma_\epsilon^2-\frac{\gamma^2}{4}\sigma_\epsilon^4$$

of each price series identifies the permanent component as $\Sigma+2\phi_1$ and the transitory component as $\lambda=-2\phi_1$. □

Proof of Proposition 6. Consider the bid-side—this is without loss of generality—and a price-quantity offer (B,Q), where Q is the quantity M is willing to buy at price B per unit (the total transfer). This is equivalent to a bid-schedule $B(q)$, where $B(q)$ is an unacceptably low bid for $q\neq Q$.[17] For a given price-schedule $B(q)$, let μ be T's belief that the asset value is high. His certainty equivalent by selling q units $\underline{u}(\mu,q)$ is

$$\underline{u}(\mu,q)=(1-q)[v_{t-1}+(2\mu-1)\Delta]+qB(q)-\frac{\gamma}{2}(1-q)^2[4\mu(1-\mu)\Delta^2+\sigma_\epsilon^2]$$

and his certainty equivalent from abstaining is then $\underline{u}(\mu,0)$. Note, as in the main case, autarky is not an equilibrium: M can always offer the bid-quantity pair $(B,Q)=(\underline{u}(1,1),1)$ and T will accept regardless of his beliefs.

1. **Pooling Equilibrium** $\mathcal{P}^0$: As in the binary ($q\in\{0,1\}$) case, risk sharing gains are maximized at $q=1$, and M can extract all of these gains by offering $\underline{u}(\frac{1}{2},1)$ under

[17] I'm abusing notation here a bit by using $q>0$ to indicate the quantity sold; elsewhere, I've used $q<0$.

the pooling belief $\mu = \frac{1}{2}$, and a full specification of actions and beliefs supporting this outcome is

$$(B^*, Q^*) = \left(\underline{u}(\frac{1}{2}, 1), 1\right) \qquad \mu^*(B,Q) = \begin{cases} 1, \forall (B,Q) : B > v_{t-1} - \Delta \\ \frac{1}{2}, \textit{otherwise} \end{cases}$$

$$q^*(B,Q) = \begin{cases} Q, \forall (B,Q) : B \geq v_{t-1} + \Delta - (2-Q)\frac{\gamma}{2}\sigma_\epsilon^2 \\ Q, \forall (B,Q) : B \in [v_{t-1} - \Delta - (2-Q)\frac{\gamma}{2}(\Delta^2 + \sigma_\epsilon^2), v_{t-1} - \Delta] \\ 0, \textit{otherwise} \end{cases}$$

where the off-equilibrium beliefs and actions satisfy the intuitive criterion. For the full insurance pooling equilibrium to be viable, the pooling bid needs to be profitable in both states ($\underline{u}(\frac{1}{2}, 1) \leq v_{t-1} - \Delta$), requiring that $\gamma \geq \hat{\gamma}_P$

$$\hat{\gamma}_P \equiv \frac{2\Delta}{\Delta^2 + \sigma_\epsilon^2}$$

2. **Separating Equilibrium $\mathcal{S}^0$**: The high state mirrors that of Proposition 1. Let (B_H, Q_H) be the price-quantity pair offered in the high value state. $B_H > v_{t-1} - \Delta$ signals asset value is high, and the market maker extracts all gains in this state when $(B_H^*, Q_H^*) = (\underline{u}(1,0), 1)$. In the low value state, any bid-quantity pair needs to satisfy two constraints: First, an IR constraint that ensures trader T gets at least his outside option, $\underline{u}(0,q) \geq \underline{u}(0,0)$, implying

$$B_L(q) \geq v_{t-1} - \Delta - (2-q)\frac{\gamma}{2}\sigma_\epsilon^2 \qquad (IR_q)$$

Second, an IC constraint ensuring that the high-type market maker is not tempted to deviate by quoting the low bid-quantity pair. His profit in the high state is $\pi_H^* = \frac{\gamma}{2}\sigma_\epsilon^2$, the residual risk sharing gains, so bid-quantity pairs have to satisfy

$$\pi_H \geq q(v_{t-1} - \Delta - B_L(q))$$
$$\implies B_L(q) \geq v_{t-1} + \Delta - \frac{\gamma}{2q}\sigma_\epsilon^2 \qquad (IC_q)$$

For a given quantity q, M maximizes profits when she bids the minimum amount satisfying both constraints. In the low state, her profits $\pi_L(q)$ for a given q are therefore

$$\pi_L(q) = q\left(v_{t-1} - \Delta - \max\left\{v_{t-1} - \Delta - (2-q)\frac{\gamma}{2}\sigma_\epsilon^2, v_{t-1} + \Delta - \frac{\gamma}{2q}\sigma_\epsilon^2\right\}\right)$$

For small q, the IR_q constraint binds. It's straightforward to show that $\pi_L(q)$ is increasing on q when this is the case, so M increases her offer. Solving for the q that equalizes the constraints shows that there is a single crossing point at

$$q^* = 1 - \frac{\sqrt{\Delta^2 + \Delta\gamma\sigma_\epsilon^2} - \Delta}{\frac{\gamma}{2}\sigma_\epsilon^2} \in (0,1)$$

after which IC_q determines the bid. $\pi_L(q)$ is decreasing as q increases for this case, so M will optimally bid $B_L(q^*)$. Define the set of incentive compatible bids $\mathcal{B}$ as

$$\mathcal{B} = \{(B,Q) : B \geq v_{t-1} + \Delta - \frac{\gamma}{2Q}\sigma_\epsilon^2\}$$

A full equilibrium specification is given by

$$(B^*, Q^*)(\eta_t) = \begin{cases} (\underline{u}(1,0), 1), \eta_t = \Delta \\ (v_{t-1} - \Delta - (2-q^*)\frac{\gamma}{2}\sigma_\epsilon^2, q^*), \eta_t = -\Delta \end{cases}$$

$$\mu^*(B,Q) = \begin{cases} 1, & \forall (B,Q) : B > v_{t-1} - \Delta \cup (B,Q) \notin \mathcal{B} \\ 0, & \forall (B,Q) : B \leq v_t - 1 - \Delta \cap (B,Q) \in \mathcal{B} \end{cases}$$

$$q^*(B,Q) = \begin{cases} Q, \forall (B,Q) : B \geq v_{t-1} + \Delta - (2-Q)\frac{\gamma}{2}\sigma_\epsilon^2 \\ Q, \forall (B,Q) : B \leq v_t - 1 - \Delta \cap (B,Q) \in \mathcal{B} \\ 0, otherwise \end{cases}$$

For some parameterizations, there are more complex out-of-equilibrium belief specifications that support the same outcome—they do not change prices and quantities. Unlike in the binary case, this separating equilibrium always exists: the low type M is never tempted to post the high-type bid, even if the bid $B_H < v_{t-1} - \Delta$ so that is profitable.

□

Proof of Proposition 7. Most of the logic here mirrors Proposition 1, so I omit details on belief specifications, etc. Let μ be trader T's belief about whether the asset value is high. His certainty equivalents $\underline{u}(\mu)$ and $\overline{u}(\mu)$ from abstaining ($q = 0$) following positive and negative endowment shocks, respectively, are

$$\begin{aligned} \mathbb{E}[v_t|\mu] &= v_{t-1} + \Delta_L + p(\Delta_H - \Delta_L) \\ \mathbf{V}ar[v_t|\mu] &= \mu(1-\mu)(\Delta_H - \Delta_L)^2 + \sigma_\epsilon^2 \\ \implies \underline{u}(\mu) &= v_{t-1} + \Delta_L + p(\Delta_H - \Delta_L) - \frac{\gamma}{2}[\mu(1-\mu)(\Delta_H - \Delta_L)^2 + \sigma_\epsilon^2] \\ \overline{u}(\mu) &= -[v_{t-1} + \Delta_L + p(\Delta_H - \Delta_L)] - \frac{\gamma}{2}[\mu(1-\mu)(\Delta_H - \Delta_L)^2 + \sigma_\epsilon^2] \end{aligned}$$

1. **Separating Equilibrium $\mathcal{S}^0$**: Separation requires that M credibly reveal her information, which is only possible for bids $B_H > v_{t-1} + \Delta_L$ and asks $A_L < v_{t-1} + \Delta_H$. Because the other types lose money at these quotes, beliefs are $\mu(\underline{u}(1), \Delta_H) = 1$ for high bids and $\mu(\overline{u}(0), \Delta_H) = 0$ for low asks. M will extract maximum gains for these

beliefs: $B_H^* = \underline{u}(1), A_L^* = -\overline{u}(0)$. Trade can not occur in any separating equilibrium for low bids $B_L < v_{t-1} + \Delta_L$ or high asks $A_H > v_{t-1} + \Delta_H$.

As before, the separating bid quotes are only feasible if $\underline{u}(1) > v_{t-1} + \Delta_L$. Similarly, separating ask quotes are only feasible $A_L^* < v_{t-1} + \Delta_H$. Each of these constraints implies the same parameter restriction $\gamma < \gamma_S$, where

$$\gamma_S \equiv \frac{2(\Delta_H - \Delta_L)}{\sigma_\epsilon^2}$$

2. **Bid-side Pooling Equilibrium** $\mathcal{P}_B^0$: For M to commit to a pooling bid in either state, it is necessary that $\underline{u}(p) \leq v_{t-1} + \Delta_L$, requiring that $\gamma \geq \gamma_{P,B}$, where

$$\gamma_{P,B} \equiv \frac{2p(\Delta_H - \Delta_L)}{\sigma_\epsilon^2 + p(1-p)(\Delta_H - \Delta_L)^2}$$

3. **Ask-side Pooling Equilibrium** $\mathcal{P}_A^0$: Similarly, on the ask-side, pooling requires $-\overline{u}(p) \geq v_{t-1} + \Delta_H$, implying that $\gamma \geq \gamma_{P,A}$, where

$$\gamma_{P,A} \equiv \frac{2(1-p)(\Delta_H - \Delta_L)}{\sigma_\epsilon^2 + p(1-p)(\Delta_H - \Delta_L)^2}$$

4. When γ exceeds both pooling thresholds, then obviously pooling on both sides of the market is feasible.

5. Follows from inspection of the thresholds.

6. Follows from inspection of the thresholds.

7. The logic is the same as Proposition 1: any mixed equilibrium implies partially revealing prices in some states of the world, but for induced under any of those prices, M will deviate from his mixed strategy.

□

Proof of Proposition 8. The logic follows Propositions 2-3 and Lemma 2, so many details are omitted. First consider the separating improvements on the bid-side: the easiest improvement to support gives M all of the residual risk sharing gains on every trade. His tempted deviation is to bid low when the asset has a high value, which triggers the static separating outcome, requiring

$$(1-\beta)\frac{\gamma}{2}\sigma_\epsilon^2 + \beta\frac{\gamma}{2}\sigma_\epsilon^2 \geq (1-\beta)(\Delta_H - \Delta_L + \frac{\gamma}{2}\sigma_\epsilon^2) + \beta\frac{\gamma}{4}\sigma_\epsilon^2$$

implying that $\gamma \geq \gamma_S^\infty$, where

$$\gamma_S^\infty \equiv \frac{4(1-\beta)(\Delta_H - \Delta_L)}{\beta\sigma_\epsilon^2}$$

Note the constraint is identical on the ask-side, as the temptation to deviate is independent of the high-state probability p, so the separating threshold is the same.

For a pooling bid, let s denote the spread between the trader's prior expectation and the bid: $B = \mathbb{E}[v_t|p] - s$. In a pooling equilibrium, s is the market maker's ex-ante expected per-period profit, but she loses money on this spread in the low-state when $\gamma < \gamma_{P,B}$. She is tempted to deviate by abstaining from trade in this state because she knows the asset value is low:

$$\begin{aligned}(1-\beta)(v_{t-1} + \Delta_L - B) + \beta s &\geq (1-\beta)\cdot 0 + \beta\frac{\gamma}{4}\sigma_\epsilon^2 \\ \Rightarrow s &\geq (1-\beta)p(\Delta_H - \Delta_L) + \beta\frac{\gamma}{4}\sigma_\epsilon^2\end{aligned}$$

In order for trade to occur, the bid has to exceed the trader's certainty equivalent given his prior: $B \geq \underline{u}(p)$, implying that

$$s \leq \frac{\gamma}{2}[\mu(1-\mu)(\Delta_H - \Delta_L)^2 + \sigma_\epsilon^2]$$

which combined with the previous IC constraint requires that $\gamma \geq \gamma_{P,B}^\infty$, where

$$\gamma_{P,B}^\infty \equiv \frac{4(1-\beta)p(\Delta_H - \Delta_L)}{(2-\beta)\sigma_\epsilon^2 + 2p(1-p)(\Delta_H - \Delta_L)^2}$$

The logic on the ask side is symmetric with $p \to (1-p)$, so

$$\gamma_{P,A}^\infty \equiv \frac{4(1-\beta)(1-p)(\Delta_H - \Delta_L)}{(2-\beta)\sigma_\epsilon^2 + 2p(1-p)(\Delta_H - \Delta_L)^2}$$

That both of these pooling thresholds are lower than γ_S^∞ follows by inspection. □

References

Admati, Anat R., and Paul Pfleiderer, 1990, Direct and indirect sale of information, *Econometrica* 58, 901–928.

Bao, Jack, Jun Pan, and Jiang Wang, 2009, Liquidity of corporate bonds, *Working Paper*.

Benveniste, Lawrence M., Alan J. Marcus, and William J. Wilhelm, 1992, What's special about the specialist?, *Journal of Financial Economics* 32, 61–86.

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid M. Werner, 2005, Why do larger orders receive discounts on the london stock exchange?, *The Review of Financial Studies* 18, 1343–1368.

Bernhardt, Dan, and Eric Hughson, 1997, Splitting orders, *The Review of Financial Studies* 10, 69–101.

Biais, Bruno, David Martimort, and Jean-Charles Rochet, 2003, Competing mechanisms in a common value environment, *Econometrica* 68, 799–837.

Cao, H. Henry, Martin D.D. Evans, and Richard K. Lyons, 2006, Inventory information, *Journal of Business* 79, 325–364.

Chaboud, Alain, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega, 2009, Rise of the machines: Algorithmic trading in the foreign exchange market, *Working Paper, Federal Reserve Board*.

Cho, In-Koo, and David M. Kreps, 1987, Signaling games and stable equilibria, *The Quarterly Journal of Economics* 102, 179–222.

Dennert, Jürgen, 1993, Price competition between market makers, *The Review of Economic Studies* 60, 735–751.

Desgranges, Gabriel, and Thierry Foucault, 2005, Reputation-based pricing and price improvements, *Journal of Economics and Business* 57, 493–527.

Edmans, Alex, 2009, Blockholder trading, market efficiency, and managerial myopia, *The Journal of Finance* LXIV.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *The Journal of Finance* 62, 1421–1451.

Glosten, Lawrence R., 1989, Insider trading, liquidity, and the role of the monopolist specialist, *The Journal of Business* 62, 211–235.

Goettler, Ronald, Christine Parlour, and Uday Rajan, 2009, Informed traders and limit order markets, *Journal of Financial Economics* 93, 67–87.

Gomes, Armando, 2000, Going public without governance: Managerial reputation effects, *The Journal of Finance* LV, 615–646.

Gould, John P., and Robert E. Verrecchia, 1985, The information content of specialist pricing, *Journal of Political Economy* 93, 66–83.

Green, Richard C., Burton Hollifield, and Norman Schurhoff, 2007, Financial intermediation and the costs of trading in an opaque market, *The Review of Financial Studies* 20, 275–314.

Grinblatt, Mark S., and Stephen A. Ross, 1985, Market power in a securities market with endogenous information, *The Quarterly Journal of Economics* 100, 1143–1167.

Hasbrouck, Joel, 1993, Assessing the quality of a security market: a new approach to transaction-cost measurement, *Review of Financial Studies* 6, 191–212.

Hobson, Jeremey, 2010, Why investors trade with goldman, *NPR's Marketplace*.

Huang, Jennifer, and Jiang Wang, 2010, Market liquidity, asset prices, and welfare, *Journal of Financial Economics* 95, 107 – 127.

Laffont, Jean-Jacques, and Eric S. Maskin, 1990, The efficient market hypothesis and insider trading on the stock market, *Journal of Political Economy* 98, 70–93.

Leland, Hayne E., 1992, Insider trading: Should it be prohibited?, *Journal of Political Economy* 100, 859–887.

Mailath, George, Masahiro Okuno-Fujiwara, and Andrew Postlewaite, 1993, Belief-based refinements in signalling games, *Journal of Economic Theory* 60, 241.

Permanent Subcommittee On Investigations, 2010, Wall street and the financial crisis: The role of investment banks, Discussion paper U.S. Senate Washington, DC.

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *The Journal of Finance* 39, 1127–1139.

Securities and Exchange Commission, 2010, Securities and exchange commission v. goldman, sachs co. and fabrice tourre, 10 civ. 3229, Discussion paper Litigation Release No. 21489.

Seppi, Duane J., 1990, Equilibrium block trading and asymmetric information, *The Journal of Finance* 45, 73–94.

Stein, Jeremy C., 1988, Takeover threats and managerial myopia, *The Journal of Political Economy* 96, 61–80.

Taibbi, Matt, 2009, The great american bubble machine, *Rolling Stone* 1082-1083.

The Financial Services Authority, 2009, Reforming remuneration practices in financial services, Discussion paper.