The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXIII, NO. 2 • APRIL 2018

Liquidity as Social Expertise

PABLO KURLAT*

ABSTRACT

This paper proposes a theory of liquidity dynamics. Illiquidity results from asymmetric information. Observing the historical track record teaches agents how to interpret public information and helps overcome information asymmetry. However, an illiquidity trap can arise: too much asymmetric information leads to the breakdown of trade, which interrupts learning and perpetuates illiquidity. Liquidity falls in response to unexpected events that lead agents to question their valuation models (especially in newer markets) may be slow to recover after a crisis, and is higher in periods of stability.

THIS PAPER PROPOSES A THEORY OF MARKET LIQUIDITY and its evolution over time. This theory is based on the interaction between information asymmetry and social learning.

Liquidity is an elusive concept, with the literature on it plagued by challenges of both definition and measurement.[1] In this paper, I refer to the following notion of liquidity: The liquidity of an asset class is the fraction of the potential gains from trade in that asset class that are realized in equilibrium. If the potential gains from trade are large, it is tautologically true that asset liquidity has important consequences for social welfare.

The theory in this paper is based on a minor modification of an otherwise standard model of trade under asymmetric information in the spirit of Akerlof (1970). The assumption is that rather than being purely uninformed, asset buyers have access to some information, but their ability to make use of that information depends on their collective experience, which I refer to as "social expertise." The model abstracts from the specific features of the assets that are traded. However, to fix ideas, it is useful to think about the model as applying to markets such as the market for initial public offerings (IPOs), the primary market for asset-backed securities, or the primary market for sovereign bonds.

*Pablo Kurlat is with the Department of Economics, Stanford University. An early version of parts of this paper circulated under the title 'Lemons, Market Shutdowns and Learning'. I am grateful to Manuel Amador, Marios Angeletos, Ricardo Caballero, Bengt Holmström, Peter Kondor, Juan Pablo Nicolini, Fabrizio Perri, Victor Rios-Rull, Jean Tirole, Iván Werning, and various seminar participants for helpful comments. I have read the *Journal of Finance's* disclosure policy and have no conflicts of interest to disclose.

[1] See Brunnermeier and Pedersen (2009) for a discussion of the conceptual issues and Goyenko, Holden, and Trzcinka (2009) for a discussion of measurement.

DOI: 10.1111/jofi.12606

The model works as follows. Each period, "sellers" own assets of heterogeneous quality that are independent over time. These assets are more productive if held by "buyers" rather than sellers, so there are gains from trade. Sellers know the asset qualities but buyers cannot observe them, so a classic lemons problem arises.

Buyers can alleviate their information disadvantage by observing public, asset-specific signals. However, for these signals to be useful, buyers need to know how signals correlate with asset qualities. The key assumption in the model is that the joint distribution of signals and asset qualities, summarized by the single parameter μ, is not known exactly. Agents can learn about μ from commonly observed data, based on a sample of past assets. For each of them, the data contain the signal it produced when it was created and an indicator for how it turned out. The idea is that agents turn to the historical track record in making sense of the information available on any given asset, which is standard practice among financial analysts. For instance, the popular textbook by Damodaran (2008) provides guidance on the use of "comparables" to determine which pieces of information about a particular asset one should focus on and shows how to use them in valuation. One of the challenges, in practice, is finding a sufficiently large and sufficiently similar sample of historical precedents. In terms of the model, the question is at what rate are observations added to the historical sample. I assume that this rate depends positively on the volume of trade, a form of learning-by-doing.

The model delivers several predictions about the relationship between information, trading, and liquidity. The first result is that if information asymmetry is sufficiently severe, liquidity is increasing in the precision of agents' estimates of μ. Knowing μ increases traders' ability to extract information from signals, reducing the degree of information asymmetry and increasing liquidity. Hence, liquidity is a function of traders' expertise. Asset prices, the level of investment, and the volume of trade are increasing in liquidity.

Depending on parameters, the model may feature an "illiquidity trap." If at any point in time, estimates of μ are sufficiently imprecise, assets will be completely illiquid and trade will break down. If the learning process is such that data only are generated by trading, markets will generate no data for agents to learn from, which will perpetuate the illiquidity. Whether the economy falls into an illiquidity trap depends on the sample realizations during the first few periods. If the first observations lead to precise and correct estimates of μ, this will increase liquidity and reinforce the learning process, which becomes self-sustaining. If instead the first observations lead to imprecise estimates of μ because they conflict with each other or with agents' prior, signals will become uninformative, which will lead to the illiquidity trap. Even under parameters such that there is no permanent illiquidity trap, market liquidity can be slow to recover after a disruption.

The model also predicts that markets will tend to become more liquid over time, as traders accumulate more observations with which to estimate μ. In the short run, unexpected events disrupt liquidity because they increase buyers' uncertainty as to whether they are using the correct model (i.e., the correct value of μ) to evaluate assets. This increases information asymmetry and lowers liquidity. Furthermore, unexpected events will be more disruptive in newer markets, where traders have not had time to accumulate a long track record of observations and thus revise their beliefs more strongly in light of new information.

I next extend the model to allow the structure of the economy, as captured by μ, to change over time. In this case, unexpected events will be more common in times of structural change and therefore liquidity will be higher when the underlying economy is more stable. Even in the long run, liquidity will remain fragile.

The dynamics of the model follow from positive feedback between learning and trading. The fact that trading generates observations to learn from is simply an assumption, though I provide examples of underlying models that could give rise to it. The main contribution of the paper is to show how, despite the fact that asset payoffs are independent over time, past observations can generate useful expertise that alleviates information asymmetry and increases liquidity.

This paper relates to several strands of literature. First, it builds on the literature following Akerlof (1970) on how asymmetric information can create barriers to trade (Wilson (1980), Kyle (1985), Glosten and Milgrom (1985), Levin (2001), Attar, Mariotti, and Salanié (2011)). The main contribution relative to this literature is to explore the effect of more precise knowledge about the information structure on trade in a simple model that is quite close to Akerlof's (1970) basic setup.

Second, the paper relates to a large literature on social learning. The form of learning of interest here, namely, learning that is dependent on economic activity, has been studied by Veldkamp (2005), van Nieuwerburgh and Veldkamp (2006), Ordoñez (2009), and Fajgelbaum, Schaal, and Taschereau-Dumouchel (2014). These stories focus on agents who need to learn the level of aggregate productivity to make production and investment decisions. This paper focuses instead on agents who learn the right way to evaluate assets, and studies the consequences this has for market liquidity. The mechanics of the illiquidity trap are also related to the informational cascades studied by Banerjee (1992) and Caplin and Leahy (1994).

Finally, the paper contributes to the literature on the sources of liquidity fluctuations. These include theories based on asymmetric information (Daley and Green (2012), Cespa and Foucault (2014), Dang, Gorton, and Holmström (2015), Daley and Green (2016)), theories based on Knightian uncertainty or non-Bayesian learning (Hong, Stein, and Yu (2007), Caballero and Krishnamurthy (2008), Routledge and Zin (2009), Uhlig (2010)), and theories based on balance sheet effects (Shleifer and Vishny (1992), Holmström and Tirole (1997), Kiyotaki and Moore (1997), Brunnermeier and Pedersen (2009)). In my model, the liquidity of an asset class can fluctuate even though all agents are Bayesian, none are financially constrained, and the distribution of payoffs for the asset class remains unchanged.

The rest of the paper is organized as follows. Section I describes the model. Section II characterizes the static outcomes of the model. Section III presents the main results on liquidity dynamics and Section IV offers some final remarks. Appendix A contains proofs, while Appendix B provides microfoundations for the learning mechanism.

I. The Model

A. Agents, Endowments, and Technology

There are two types of agents, sellers and buyers. Each lives for only one period (when I introduce dynamics, I study the continuous time limit where the period is short). Both types of agents have large endowments of the consumption good and are risk neutral.

There is a measure $1+\lambda$ of sellers: a measure one are endowed with one unit of capital and a measure λ are endowed with one "lemon," which looks like capital but is useless. I refer to capital and lemons collectively as "assets." A seller who is endowed with capital may (within the period) use a unit of capital to produce Z consumption goods, where Z is a nonnegative random variable, which is distributed i.i.d. across sellers and over time according to a continuous distribution F. The law of large numbers applies, so F also represents the distribution of Z among sellers. After production, capital depreciates completely. Buyers are not endowed with capital but have access to a technology that produces θ consumption goods with one unit of capital.

ASSUMPTION 1: $F(\theta) = 1$.

Assumption 1 says that sellers are always less productive than buyers. A social planner would want all capital to be operated by buyers.

B. Information

Each seller knows whether he holds capital or a lemon, but this is not observable to either buyers or other sellers. However, each asset emits a publicly observable signal s, which takes one of two possible values, A and B, according to the conditional probabilities

$$\begin{aligned}\Pr[s = A|\text{Capital}] &= \Pr[s = B|\text{Lemon}] = \mu,\\ \Pr[s = B|\text{Capital}] &= \Pr[s = A|\text{Lemon}] = 1 - \mu.\end{aligned} \tag{1}$$

The variable μ is itself random, realized at the beginning of time but initially unknown to the agents, drawn according to

$$\mu = \begin{cases} \bar{\mu} & \text{with probability } g_0 \\ 1-\bar{\mu} & \text{with probability } 1-g_0, \end{cases} \tag{2}$$

where $\bar{\mu} \in (0.5, 1)$.

Occasionally, after its payoff is realized, an asset will emit an additional "ex-post" public signal $z \in \{H, L\}$ with the conditional probabilities

$$\Pr[z = H|\text{Capital}] = \Pr[z = L|\text{Lemon}] = \gamma,$$
$$\Pr[z = L|\text{Capital}] = \Pr[z = H|\text{Lemon}] = 1 - \gamma, \tag{3}$$

where $\gamma \in (0.5, 1)$ is a *known* parameter. When an ex-post signal z is published, agents can contrast it to the original signal s to update their beliefs about μ. Each cohort can observe the entire history of past $\{s, z\}$ pairs. They cannot, however, observe the quantity of A- and B-labeled assets that trade. The publication of ex-post signals takes place randomly over time at an aggregate endogenous Poisson rate ϕ_t.

ASSUMPTION 2: *Let x_t be the fraction of capital that sellers trade at time t. The rate ϕ_t is given by $\phi_t = \phi(x)$ where $\phi(\cdot)$ is an increasing function.*

C. Discussion of Assumptions

The labels "seller" and "buyer" are meant to be interpreted broadly as a pair of agents for whom there are gains from trade. For instance, in the context of IPOs, the gap $\theta - Z$ between their respective productivities can represent the gains from diversification and not just literally higher output. In the context of sovereign debt, it can represent the gains from intertemporal trade.

The assumption that μ is not initially known is meant to capture the idea that market participants might not know which model to apply to a particular asset class. Uncertainty about μ represents uncertainty about how to translate information into valuations. If $\mu = \bar{\mu}$, then equation (1) implies that signal A is more likely to come from a unit of capital and signal B is more likely to come from a lemon, but the reverse holds if $\mu = 1 - \bar{\mu}$. I assume this formulation, where signals reverse their meaning depending on the value of μ, to make things stark. There are types of information for which this type of assumption does not make much sense, surely the signal "this company increased its profits" is a better signal than "this company reduced its profits," even if one is not sure by how much. For more ambiguous types of information, however, this assumption can make more sense. Suppose the signal is "this company has had fast growth but low profit margins." Is this a good or a bad sign? In this particular industry, is it easy to raise profit margins without choking off growth? Would one instead be more optimistic about a slower growing firm that had healthier profit margins? Furthermore, the symmetric formulation has the desirable feature that, for someone who knew the real value of μ, signals would be equally informative no matter what this value was, the only determinant of their actual informativeness is the extent to which agents know the correct way to use the information. This formulation is a parsimonious way of modeling the idea that uncertainty about μ (and not the value of μ itself) lowers the information content of signals.

In using this type of information, it is helpful to have a long track record of initial signals and eventual outcomes. This is the notion of expertise that the model explores: expertise comes from having a sufficiently large historical database that it is possible to estimate μ precisely. Expertise is "social" in the sense that the historical database is assumed to be publicly available.[2]

The two signals s and z play different roles in the model. The original signal s is a composite of hard-to-interpret public information about an asset that is available at the time it trades. In the context of an IPO, s represents the content of the company's prospectus, analyst reports about its industry, the biography of its founders, etc. The ex-post signal z represents the information that is publicly available after the asset trades, for instance, the company's financial statements for the first few years after the IPO. While these latter signals might also not be definitive about whether the company turned out to be a good investment (in terms of the model, γ is less than 1), they have the advantage that there is less ambiguity about how to interpret them. Since it is too late to use this information for trading the asset that generated it, agents only use it to update their valuation models for the next asset that comes along.[3]

Assumption 2 is the key driver of the model's dynamics. It says that there is a form of learning-by-doing that depends on the volume of trade. More than one possible mechanism may give rise to this. Appendix B provides three examples. In the first, there is a small cost of producing the original signals s. For instance, the cost of s could represent the cost of preparing a detailed prospectus for an IPO or a bond issuance. In equilibrium, only sellers that plan to sell their asset incur the cost of producing s. Ex-post signals z could represent subsequent financial statements by the issuing firm, and the example assumes for simplicity that these are free. The number of $\{s, z\}$ pairs that are observed depends on the number of s signals produced, which depends on the volume of trade. In the second example, the signals are assumed to be produced by intermediaries who buy and then resell assets, and in equilibrium disclose their information due to the unraveling logic of Milgrom (1981) and Grossman (1981). Again, the quantity of $\{s, z\}$ pairs that results depends on the volume of trade. In the third example, all assets are assumed to produce both ex-ante and ex-post signals but the volume of investment is endogenous. Better information means investors obtain a higher average value from their assets, which leads to higher investment, trade, and information flow.

[2] One example that might be familiar to some readers arises in graduate student admissions. Imagine a recommendation letter for an undergraduate from a university without a track record of students going on to graduate school saying "this is the second best student I have ever had in my class." Without knowing how talented the other students were, how trustworthy the letter-writer is, etc., it is hard to extract much information from this signal. If, instead, there have been many students from the same background, the informational content of the letter is much greater.

[3] Of course, this information will be reflected in secondary markets for the asset, but the focus here is on the primary market. Brown (2015) emphasizes the importance of post-IPO information production, though he focuses on the usefulness of information for the firm's decision-making rather than for providing feedback to IPO investors.

Modeling time as continuous (i.e., taking the limit as each period is short) is for simplicity and does not make much difference for any of the results. The main advantage of this modeling choice is that it allows one to model the ex-post signals as a Poisson process, which ensures that they arrive one at a time even if ϕ is very high. This avoids the slightly more complicated Bayesian updating formula that would arise if more than one ex-post signal is observed in the same period.

II. Static Equilibrium

A. Equilibrium Definition

Since all agents are short-lived, it is possible to define an equilibrium of the static economy that takes place each period. The equilibrium will depend on the beliefs about μ that agents have upon entering the period. As is well known (Wilson (1980), Hellwig (1987)), there is more than one way to formulate a competitive equilibrium in environments with asymmetric information. In particular, the formulation depends on whether the sellers can signal their type, for instance, by committing to retain a fraction of the asset. The definition here implicitly assumes they cannot do so.

Since lemons are useless, sellers who own a lemon will sell it at any positive price. Sellers who own capital will sell it if the price they can get is greater than their own productivity. Letting p_s denote the equilibrium price conditional on signal s, a seller will sell capital if $Z < p_s$. Buyers will make zero profits upon buying A- or B-labeled assets if the following conditions hold:

$$p_A = \frac{\mathbb{E}[\mu \mathbb{I}(Z < p_A)]}{\mathbb{E}[\mu \mathbb{I}(Z < p_A) + \lambda(1-\mu)]}\theta, \tag{4}$$

$$p_B = \frac{\mathbb{E}[(1-\mu) \mathbb{I}(Z < p_B)]}{\mathbb{E}[(1-\mu) \mathbb{I}(Z < p_B) + \lambda\mu]}\theta. \tag{5}$$

In equations (4) and (5), the numerator is the expected number of s-labeled units of capital sold by sellers and the denominator is the expected number of s-labeled assets (capital plus lemons) sold.[4] Equations (4) and (5) could have multiple solutions.

DEFINITION 1: *An equilibrium consists of* $\{p_A, p_B\}$ *that are equal to the highest solution of equations* (4) *and* (5)*, respectively.*

Notice that the quantities of A- and B-labeled assets that are traded depend on the realized value of μ, since that determines how many units of capital and lemons end up carrying each label. The definition of equilibrium implicitly assumes that buyers stand ready to absorb however many units are put up for sale at the equilibrium price. The requirement that prices correspond to

[4] These are *expected* because the true value of μ is uncertain.

the *highest* solution requirement can be derived formally by assuming that there are markets at every possible price and agents select at which prices to trade, as in Kurlat (2016), or by deriving the equilibrium from a game-theoretic formulation with nonexclusive competition, as in Attar, Mariotti, and Salanié (2011).

B. Equilibrium Characterization

Since Z is independent of the realization of μ, (4) and (5) can be rewritten as:

$$p_A = \frac{\hat{\mu} F(p_A)}{\hat{\mu} F(p_A) + \lambda(1 - \hat{\mu})}\theta, \tag{6}$$

$$p_B = \frac{(1 - \hat{\mu}) F(p_B)}{(1 - \hat{\mu}) F(p_B) + \lambda \hat{\mu}}\theta, \tag{7}$$

where $\hat{\mu} \equiv \mathbb{E}(\mu)$.

Notice two features of equations (6) and (7), both of which follow from the binary distribution. First, beliefs about μ enter prices only through their mean $\hat{\mu}$, that is, $\hat{\mu}$ is a sufficient statistic for the problem of inferring whether an asset is a lemon. Agents interpret signals as though they knew that the value of μ was equal to $\hat{\mu}$. Second, equations (6) and (7) are symmetric, meaning that beliefs $\hat{\mu}$ and $1 - \hat{\mu}$ lead to the same prices, except that which price corresponds to which signal is reversed. When $\hat{\mu} > 0.5$, we have $p_A \geq p_B$, and vice versa.

DEFINITION 2: *The informativeness of signals given beliefs about μ is $\tau \equiv 2|\hat{\mu} - 0.5|$.*

The informativeness measure τ is scaled so that $\tau = 0$ means signals convey no information and $\tau = 1$ means they convey perfect information. The first case arises when $\hat{\mu} = 0.5$, so each signal is believed to be equally likely to arise from capital or lemons; the second arises when $\hat{\mu} = 0$ or $\hat{\mu} = 1$, so signals perfectly reveal the type of asset that generates them. Given the symmetry of equations (6) and (7), equilibrium outcomes depend on τ and not on whether $\hat{\mu}$ is above or below 0.5.

Substituting the definition of τ into equations (6) and (7), it immediately follows that the two equilibrium prices (denoted by $p_H(\tau)$ and $p_L(\tau)$) are, respectively, the highest solutions to equations

$$p_H = \frac{(1 + \tau) F(p_H)}{(1 + \tau) F(p_H) + \lambda(1 - \tau)}\theta, \tag{8}$$

$$p_L = \frac{(1 - \tau) F(p_L)}{(1 - \tau) F(p_L) + \lambda(1 + \tau)}\theta. \tag{9}$$

From an ex-ante point of view, without knowing the true value of μ, the total value that sellers obtain from their assets is

$$V(\tau) = \underbrace{\lambda \left[\frac{1-\tau}{2} p_H(\tau) + \frac{1+\tau}{2} p_L(\tau) \right]}_{\text{value from selling lemons}} + \underbrace{\frac{1+\tau}{2} \left[\underbrace{F(p_H(\tau)) p_H(\tau)}_{\text{if sold}} + \underbrace{\int_{Z \geq p_H(\tau)} \mathrm{ZdF}\ (Z)}_{\text{if retained}} \right]}_{\text{value of capital with more favorable signal}}$$

$$+ \underbrace{\frac{1-\tau}{2} \left[\underbrace{F(p_L(\tau)) p_L(\tau)}_{\text{if sold}} + \underbrace{\int_{Z \geq p_L(\tau)} \mathrm{ZdF}\ (Z)}_{\text{if retained}} \right]}_{\text{value of capital with less favorable signal}}. \tag{10}$$

LEMMA 1:

1. *The ex-ante value of assets is in between the autarky value and the first-best value:* $V(\tau) \in [\mathbb{E}(Z), \theta]$.
2. *Under full information, the ex-ante value of assets is the first best value:* $V(1) = \theta$.
3. *If the following condition holds,*

$$\frac{F(p)}{F(p) + \lambda} \theta < p \quad \textit{for all } p > 0, \tag{11}$$

then

(a) $p_L(\tau) = 0$.

(b) *There exists an informativeness cutoff* $\tau^* > 0$ *such that, for* $\tau \leq \tau^*$, *we have* $p_H(\tau) = 0$, *and for* $\tau > \tau^*$, $p_H(\tau)$ *is strictly increasing in* τ.

(c) *If signals are not informative, the ex-ante value of assets is the autarky value:* $V(0) = \mathbb{E}(Z)$.

(d) *The ex-ante value of assets* $V(\tau)$ *is increasing in informativeness* τ.

In autarky, the total value of the assets for sellers is $\mathbb{E}(Z)$, that is, the expected output they can obtain from the assets on their own. In the first-best allocation, the value of the assets is θ, the output buyers can obtain from them. Part 1 of Lemma 1 states that, not surprisingly, the value that sellers obtain in equilibrium is somewhere in between these two extremes (by assumption, buyers always obtain zero surplus). Part 2 states that the upper bound is reached in the limit of perfect informativeness.

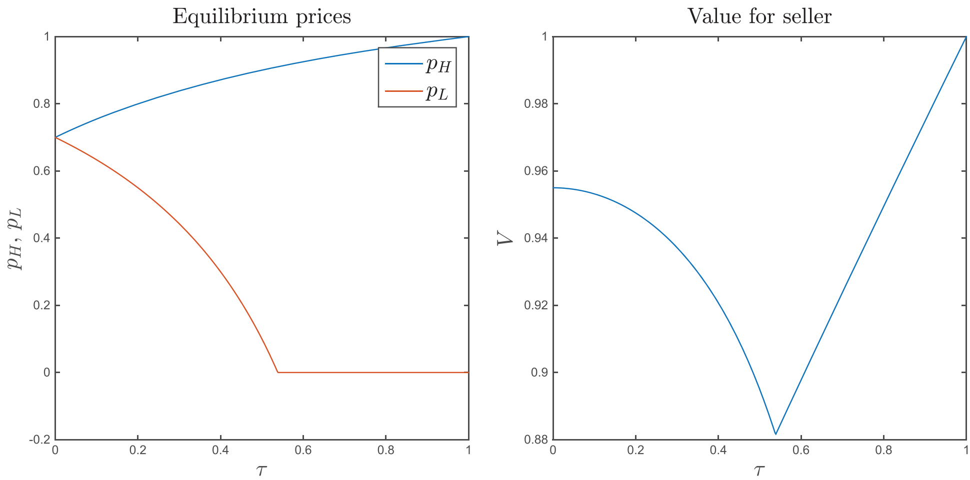

Figure 1. An example in which condition (11) does not hold. At low levels of informativeness, more informativeness lowers the value for sellers. The example uses $\lambda = 0.3$, $\theta = 1$, and $Z \sim U[0, 1]$. (Color figure can be viewed at wileyonlinelibrary.com)

Part 3 of Lemma 1 gives a sufficient condition for more-informative signals to increase sellers' value. As Levin (2001) points out, this is *not* a general feature: it is not always the case that giving the less-informed buyer more information increases efficiency. Figure 1 shows an example where condition (11) does not hold and $V(\tau)$ is decreasing for some range of τ.

What happens in the example is that more-informative signals make the good signal better (leading to a higher price and less inefficient retention by sellers) and the bad signal worse (leading to a lower price and more inefficient retention). In general, the net effect could go either way. Condition (11), which holds if λ is sufficiently high, implies that when signals are uninformative ($\tau = 0$), there is no positive solution to equations (6) and (7), so the price is zero and no capital trades. This, in turn, implies that $p_L = 0$ for any τ. Hence, under condition (11), trade can take place only if signals are sufficiently informative, and only upon observing the good signal. Therefore, there is no harm in making bad signals worse since they would lead to no trade anyway. In this case, it is indeed true that more informative signals lead to higher gains from trade.[5]

In what follows, I assume that condition (11) holds. This means that the model pertains to situations in which the adverse selection problem is so severe that, with purely uninformed buyers, trade would break down.

DEFINITION 3: *The level of liquidity is* $\frac{V(\tau)-\mathbb{E}(Z)}{\theta-\mathbb{E}(Z)}$.

I define the liquidity of assets as the fraction of the potential gains from trade that are realized. Illiquidity, conversely, is the loss of value that results from

[5] The fact that trade breaks down is sensitive to the assumption that assets are of only two kinds and the value of lemons is exactly zero. Hendren (2013) discusses what is the equivalent of condition (11) with a continuous distribution.

asymmetric information. Under this definition, liquidity is not an attribute of the assets themselves but rather of the entire equilibrium. Lemma 1 implies that if condition (11) holds, liquidity is increasing in the informativeness of signals. Furthermore, the fraction of capital that will be traded is

$$x(\tau) = F(p_H(\tau))\frac{1+\tau}{2}, \tag{12}$$

which is also increasing in informativeness and therefore positively associated with liquidity.

The key to the relationship between informativeness and liquidity is the endogeneity of sellers' decision of whether to sell their asset. Imagine that, instead of selling being an endogenous decision, a fixed fraction α of sellers who hold capital were to sell their asset, irrespective of what signal it issued, their own productivity and the equilibrium price. Then equilibrium prices would be:

$$p_A = \frac{\hat{\mu}\alpha}{\hat{\mu}\alpha + \lambda(1-\mu)}\theta,$$
$$p_B = \frac{(1-\hat{\mu})\alpha}{(1-\hat{\mu})\alpha + \lambda\hat{\mu}}\theta,$$

and the total value that sellers obtain would be

$$V(\tau) = \underbrace{\lambda[\hat{\mu}p_B + (1-\hat{\mu})p_A]}_{\textit{value from selling lemons}} \underbrace{\mu\left[\underbrace{\alpha p_A}_{\textit{if sold}} + \underbrace{(1-\alpha)\mathbb{E}(Z)}_{\textit{if retained}}\right]}_{\textit{value of capital with A signal}}$$
$$+ \underbrace{(1-\mu)\left[\underbrace{\alpha p_B}_{\textit{if sold}} + \underbrace{(1-\alpha)\mathbb{E}(Z)}_{\textit{if retained}}\right]}_{\textit{value of capital with B signal}} = \alpha\theta + (1-\alpha)\mathbb{E}(Z),$$

so the assets' liquidity would not depend on the informativeness of signals. The reason that informative signals enhance liquidity is that, by aligning prices with assets' value-in-best-use, they reduce inefficient asset retention by sellers.

III. Dynamics

The only link between different periods in the economy comes from the evolution of beliefs about μ, which depend on the history of $\{s, z\}$ pairs that are publicly observed. At any given point in time, beliefs about μ can be summarized by $g_t = \Pr_t(\mu = \bar{\mu})$. Using Definition 2 and equation (2), the resulting informativeness of signals is given by

$$\tau_t = 2|g_t\bar{\mu} + (1-g_t)(1-\bar{\mu})|$$
$$= 4|g_t - 0.5|(\bar{\mu} - 0.5).$$

Maximal informativeness is reached when the value of μ is known ($g_t = 1$ or $g_t = 0$). In this case, informativeness is equal to

$$\bar{\tau} \equiv 2(\bar{\mu} - 0.5).$$

As g_t approaches 0.5, informativeness decreases, and becomes zero when $g_t = 0.5$.

ASSUMPTION 3: $p_H(\bar{\tau}) > 0$.

Assumption 3 says that the maximal informativeness $\bar{\tau}$ is sufficiently high that if agents knew the value of μ, this would result in trade at positive prices. In the absence of this assumption, assets would always be completely illiquid.[6]

The evolution of beliefs g_t is easy to characterize. Beliefs are revised whenever an ex-post signal z is observed from some asset and remain constant in the intervals in between. Whenever an ex-post signal is observed, agents will compare it with the original signal s. Letting m be the indicator of the event $\{s, z\} \in \{\{A, H\} \cup \{B, L\}\}$, its conditional probability is $\Pr[m|\mu] = \omega(\mu) \equiv \mu\gamma + (1-\mu)(1-\gamma)$. Letting $\bar{\omega} \equiv \omega(\bar{\mu})$, $\bar{\omega} - 0.5$ measures the information content of ex-post signals for inferring μ. It is increasing in γ because more accurate ex-post signals provide more feedback about the correct interpretation of the original signals. By Bayes's rule, the posterior upon observing an ex-post signal is

$$g_t' = \frac{[\bar{\omega}^m(1-\bar{\omega})^{1-m}]g_t}{[\bar{\omega}^m(1-\bar{\omega})^{1-m}]g_t + [\bar{\omega}^{1-m}(1-\bar{\omega})^m](1-g_t)}. \tag{13}$$

The model's dynamics follow from Assumption 2, which says that the Poisson arrival rate of ex-post signals depends on the volume of trade x. Since by equation (12), trading volume depends on informativeness τ, it is possible to express the arrival rate of signals simply as a function of τ:

$$\phi(\tau) \equiv \phi(x(\tau)).$$

Since $\phi(x)$ and $x(\tau)$ are increasing, $\phi(\tau)$ is increasing as well.

Depending on parameters, the model may have the feature that when τ is sufficiently low, no learning will take place, which implies that τ will remain low. I refer to this outcome as an illiquidity trap.

DEFINITION 4: *There is an illiquidity trap if $\phi(\tau) = 0$ for sufficiently low τ.*

[6] Assumption 3 can be stated directly in terms of the primitives of the model as

$$\max_p (1-\bar{\tau})F(p)(\theta - p) - p\lambda(1-\bar{\tau}) > 0.$$

If condition (11) holds, trade breaks down for sufficiently low τ. If, in addition, $\phi(x) = 0$ for $x = 0$ (i.e., only traded assets generate ex-post signals), the model will feature an illiquidity trap. If, instead, $\phi(x) > 0$ for all x (because retained assets also generate ex-post signals, even if infrequently), an illiquidity trap will not obtain. The examples in Appendix B include two cases in which $\phi(\tau)$ gives rise to an illiquidity trap and one case in which it does not.

The mechanics of the illiquidity trap are similar to the information cascades described by Banerjee (1992), though the logic is somewhat different. Here, the learning process may be interrupted not because agents stop paying attention to information, but rather because trade collapses and no further information is generated. This is closer to the mechanism studied by Caplin and Leahy (1994), who show that inertia in decision-making can also interrupt the flow of information.

A. Short-Run Dynamics

Turn now to the short-run dynamics of the model, assuming there is no illiquidity trap. The first result says that, on average, informativeness increases over time.

PROPOSITION 1: *Suppose there is no illiquidity trap. Then, for any $t' > t$, $\mathbb{E}(\tau_{t'}|\tau_t) > \tau_t$.*

Proposition 1 provides a characterization of how the learning dynamics evolve in the short run. The law of iterated expectations implies that beliefs g_t must be a martingale; since informativeness τ is a convex function of g_t, in expectation, it increases. On average, new observations push beliefs toward either $g_t = 0$ or $g_t = 1$, increasing agents' confidence in their estimates of μ and increasing liquidity. This explains how markets become more mature. As market participants' experience increases, they are better able to use the available information, which increases valuation accuracy, alleviates information asymmetry, and increases liquidity. This is the sense in which liquidity is a function of social expertise, which accumulates as observations are added to the publicly available historical track record.

Whenever $g \neq 0.5$, one of the two possible values of μ is considered more likely, and thus some realizations of $\{s, z\}$ are more likely than others.

DEFINITION 5: *An observation $\{s, z\}$ is unlikely if either (i) $g > 0.5$ and $\{s, z\} \in \{\{A, L\} \cup \{B, H\}\}$ or (ii) $g < 0.5$ and $\{s, z\} \in \{\{A, H\} \cup \{B, L\}\}$.*

An unlikely event contradicts the model of the world that the agents think is more likely to be correct. Liquidity will react differently to likely and unlikely events, especially if the level of informativeness required for trade is large.

DEFINITION 6: *τ^* is large if*

$$\tau^* > \bar{\tau}\frac{\left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} - 1}{\left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} + 1}. \tag{14}$$

Condition (14) relates the cutoff level of informativeness required for positive liquidity τ^* to the information content of any given observation $\bar{\omega}$. If the condition holds, then no single ex-post observation carries sufficient information to overturn agents' beliefs to the point that markets remain liquid while reversing which of the two signals is considered good. The model is very stark in that there are only two possible models of the world, one in which $\mu = \bar{\mu}$ and one in which $\mu = 1 - \bar{\mu}$. Condition (14) guarantees that questioning one model cannot lead to immediately adopting the other with a sufficiently high degree of confidence that markets remain liquid. In a limit where each single piece of information is small (with $\bar{\omega}$ close to 0.5), condition (14) always holds: evidence against the preferred model leads to uncertainty before it leads to adoption of the alternative model, which implies that τ^* is large. The cutoff τ^* can fail to be large only if information comes in sufficiently large discrete pieces.

PROPOSITION 2: *If τ^* is large, then:*

1. *If liquidity is positive, an unlikely observation leads to lower liquidity.*
2. *Starting from any beliefs, there is a finite n such that a sequence of n unlikely observations will make assets completely illiquid.*
3. *The number of unlikely observations needed to make assets completely illiquid is increasing in τ_t.*

Assets are more liquid the more agents understand (or think they understand) the information structure of the economy. When unlikely events take place, they cast doubt on whether agents are using the right valuation model to guide their actions. The increased uncertainty reduces liquidity by making the information asymmetry between sellers and buyers more pronounced. Proposition 2 shows that this effect makes asset liquidity fragile: a sufficiently unlikely sequence of observations will push beliefs from levels that support positive liquidity to levels at which assets are completely illiquid. The more confidence that agents have to begin with, the less their beliefs will be swayed by contradictory information and therefore a greater number of unlikely events are necessary to render markets illiquid. Given that, by Proposition 1, τ_t tends to increase over time, this predicts that market liquidity tends to become less fragile when markets are more mature. Several studies document this feature of market maturation. For instance, Buckley (1997a, 1997b) describes how this took place in emerging market debt markets and Anderson and Gascon (2009) do so for the commercial paper market.

The result relies on τ^* being large, but only in a technical sense. If τ^* is not large, then there are values of τ such that any observation, even an unlikely one, increases liquidity, but only if liquidity was very low to begin with.

Proposition 2 establishes that, holding parameter values fixed, if assets are more liquid, this liquidity will be less fragile in response to unlikely events. Comparing across different parameter values, the same level of liquidity can be more or less fragile depending on the exact circumstances that give rise to it.

PROPOSITION 3: *Suppose there are two otherwise identical economies that differ in $\bar{\mu}$ and g_t so that informativeness $\tau_t = 4|g_t - 0.5|(\bar{\mu} - 0.5)$ is the same in both. An unlikely observation will make liquidity fall more in the economy with higher $\bar{\mu}$ and lower $|g_t - 0.5|$.*

Informativeness depends on (i) how informative signals would be if agents knew the information structure (as measured by the maximal informativeness $\bar{\tau} = 2(\bar{\mu} - 0.5)$) and (ii) how confident agents are that they know the true value of μ (as measured by $|g_t - 0.5|$). Proposition 3 shows that liquidity is more resilient to unlikely events when it results from a high degree of certainty about the information structure, even if the signals themselves are less informative. In contrast, a market in which the signals, properly interpreted, carry a lot of information but the right value of μ is uncertain faces more illiquidity-inducing surprises. This implies that, even controlling for how liquid they are at a given point in time, liquidity in more mature markets is less fragile.[7]

The link between surprising events in relatively new markets and the onset of financial crises has been emphasized by Caballero and Krishnamurthy (2008) and Caballero and Kurlat (2009). One historical example of this dynamic at play is the bankruptcy of Penn Central in 1970. Penn Central was a large issuer of commercial paper, which was a relatively new asset class, total outstanding issues having risen from about \$10 billion to about \$40 billion in four years. According to Schadrack and Breimyer (1970), sellers mistakenly believed that paper issued by large corporations was safe, so when Penn Central filed for bankruptcy in June 1970, this came as a large surprise. According to the account by Calomiris (1993, p. 13), it then became "necessary for the market to reevaluate its methods for pricing paper generally in light of this surprising event." The resulting uncertainty led to a freeze in new issues of commercial paper, with the outstanding stock falling by about 10% in the first month after Penn Central's bankruptcy.

A more recent example is the reaction of the asset-backed securities market to the downturn in the housing market. Foote, Gerardi, and Willen (2012) cite financial industry reports around 2005 that discuss different scenarios for the performance of the various tranches of mortgage-backed securities. The most pessimistic scenarios in these reports (labeled "meltdown" in one of them) far underestimate the losses that materialized in the following years. In other words, the housing crisis was considered an unlikely event. As documented by Adrian and Shin (2009) and Brunnermeier (2008), issuance of all kinds of asset-backed securities fell almost to zero when this unlikely event materialized in 2008.[8]

[7] The converse of Proposition 3 is also true. Holding τ_t constant, a *likely* observation will raise liquidity more in the less mature economy with higher $\bar{\mu}$ and lower $|g_t - 0.5|$. If a mature market in which μ is known with high precision is not very liquid, adding more observations to the sample will not improve liquidity much, even if they are consistent with what agents believe. Instead, in a less mature market, an observation that confirms agents' beliefs will significantly increase the precision of their estimates of μ, leading to a larger increase in liquidity.

[8] Caballero and Krishnamurthy (2008) describe a similar pattern following the stock market crash of 1987, the LTCM crisis in 1998, and the terrorist attacks of September 11, 2001.

Through the lens of the model, these episodes can be interpreted as follows. Before the onset of each crisis, sellers believed that they had reliable (high $\bar{\mu}$) indicators of asset quality (the size of the corporation issuing the commercial paper, credit ratings for mortgage-backed securities). Asset buyers were therefore confident that they were not at a large information disadvantage and markets were liquid, that is, most of the potential gains from trade were realized. Then, the unexpected events took place. Since the markets were relatively new (in model terms, g_t had not had time to converge all the way to zero or one yet), the unexpected events could not be dismissed as outliers, which led sellers to question their valuation models. Without confidence in a valuation model, the logic of asymmetric information took over and markets became illiquid, at least until confidence in a revised valuation model was built up.

In the above examples, the unlikely events are also bad events: assets that were expected to be good were revealed to be bad. However, what matters in light of the model is not that the news was bad, but rather that the news contradicted traders' prior, leaving them confused. Still, mostly negative surprises are consistent with what one would expect from the model. Traded assets are more likely to produce ex-post signals and, under condition (11), only assets with signals believed to be good trade. Therefore, most ex-post signals will come from assets that were originally believed to be good and most surprises will be unpleasant ones.

Negative surprises would have bad consequences even in a frictionless market. However, they would not necessarily lead to sharp drops in trading volume, just to lower prices. The breakdown of trading is consistent with a model with underlying information asymmetry. Dang, Gorton, and Holmström (2015) study a slightly different mechanism, also based on asymmetric information, through which a negative surprise can lead to illiquidity. Debt-like securities become more information-sensitive when they are closer to default, so commonly observed bad news increases the importance of noncommonly observed pieces of information, possibly leading to the breakdown of trade. In any given instance, it is hard to distinguish empirically which of the two mechanisms is at play since they both lead to the same basic prediction, and both forces could be operating at once. However, they do have some different implications that could be used to tell the mechanisms apart. First, the two mechanisms have different predictions about how the volume of trade reacts to negative news about aggregate fundamentals. Suppose buyers' productivity θ and all sellers' productivity Z fell in the same proportion. It is immediate from equation (8) that p_H would fall by the same proportion and from equation (12) that the fraction of capital traded would be unchanged. Unlike bad news about realizations of assets thought to be good, this type of bad news would have no effect on volume. Instead, if the underlying issue was the information sensitivity of debt-like securities, volume of trade would react to anything that moved beliefs toward the information-sensitive region, even if the relative valuations of buyers and sellers were unaffected. Second, in principle, one could look for instances of *positive* surprises. These should not lead to illiquidity if the underlying mechanism is the information sensitivity of debt-like securities but would lead to illiquidity if they led traders to distrust their valuation models. The difficulty with this is that if the model is correct, surprises *should* tend to be negative ones. Finally, one could look at markets that have the opposite payoff profile as debt securities, such as in-the-money convertible preferred shares. These become *less* information sensitive following bad news. If negative surprises lead to lower liquidity in these types of markets, that would be consistent with the view that traders have lost confidence in their valuation models.

Daley and Green (2012, 2016) study a related mechanism through which bad news leads to illiquidity. They consider a dynamic trading environment in which sellers have the option to retain an asset, while they wait for the market to receive more news about the asset's quality. In this environment, it is also true that prices fall and sellers of good assets choose not to sell in response to bad news, so liquidity (in the sense of the fraction of realized gains from trade) falls. Indeed, they also find a region of beliefs such that trade stops completely, although in their setting, it is because the option value of waiting for good news raises sellers' reservation prices, while in my setting, it is because under condition (11), the static model has a breakdown of trade. Daley and Green focus on a setting with only one asset, or with many identical assets, so that any news pertains to the payoff of that asset. In my setting, given that payoffs are i.i.d. over time, simply observing assets' payoffs does not provide buyers with useful information. Rather, it is the track record of ex-ante and ex-post signal *pairs* that allows them to learn about μ for the entire asset class and it is this knowledge, combined with the ex-ante signals of the current generation of assets, that provides useful information.

Another interpretation that has been proposed is that markets' reaction to negative surprises is related to some form of non-Bayesian assessment. Routledge and Zin (2009) model an uncertainty-averse intermediary setting bid-ask spreads optimally, and show how extreme events can lead to large declines in market liquidity. Hong, Stein, and Yu (2007) explore the asset pricing implications of "paradigm shifts" in a not fully Bayesian learning model in which investors base forecasts on a misspecified univariate model and reject it in favor of another one when the forecast errors are large. Uhlig (2010) studies how bank runs take place in a model in which the potential buyers of bank assets are uncertainty-averse. Caballero and Krishnamurthy (2008) study a model in which traders' Knightian uncertainty triggers a flight to quality. In these Knightian environments, public policy can play a valuable role. If the government can design a way to insure traders against the worst-case scenario they are afraid of, then it can restore market confidence, possibly without making losses. Instead, if the market breakdown is caused by increased information asymmetry, merely insuring against tail events will not do much. The government can still intervene to restore trade (for instance, by outright purchases at above-market prices), but this will require absorbing losses.

Another mechanism through which negative surprises can lead to declines in liquidity works through balance sheets. Shleifer and Vishny (1992), Kiyotaki and Moore (1997), Holmström and Tirole (1997), and Brunnermeier and Pedersen (2009), among others, study environments in which specialized traders in a particular asset are financially constrained. Since they hold the asset on their balance sheet, their constraints tighten if there is a negative shock to the asset. This diminishes their ability to carry out or intermediate trades, and hence reduces the asset's liquidity. In these environments, transfers to the constrained traders can restore liquidity. For this channel to matter, the asset in question must have a large weight in traders' portfolio. Instead, if the effect works through information asymmetry and a loss of confidence in valuation models, the weight of the asset in traders' portfolio would not be relevant and transfers to traders would not help.

B. Long-Run Outcomes

The long-run properties of the model depend on whether there is an illiquidity trap and, if so, on initial conditions and the size of the trap.

PROPOSITION 4:

1. *If there is no illiquidity trap, then informativeness will reach its maximum, that is,* $\tau_t \rightarrow_{a.s.} \bar{\tau}$ *for any initial beliefs* g_0.
2. *If there is an illiquidity trap and* $\tau_0 \leq \tau^*$, *then the economy will remain in the trap, that is,* $\tau_t = \tau_0$ *for all* t.
3. *If there is an illiquidity trap,* τ^* *satisfies condition (14), and* $\tau_0 > \tau^*$, *then the economy will either reach maximum informativeness or fall into the trap. The probability that it reaches maximum informativeness can be bounded by:*

$$\lim_{t\to\infty} \Pr[\tau_t \geq \bar{\tau} - \epsilon] \in \left[\frac{\tau_0 - \tau^*}{\bar{\tau} - \tau^*}, \frac{\tau_0 + \tau^*}{\bar{\tau} - \epsilon + \tau^*}\right] \tag{15}$$

   for any $\epsilon > 0$.
4. *If there is an illiquidity trap but* τ^* *does not satisfy condition (14), then there are values of* $\tau_0 > \tau^*$ *such that* $\lim_{t\to\infty} \Pr[\tau_t \leq \tau^*] > 0$ *and values of* $\tau_0 > \tau^*$ *such that* $\Pr[\tau_t \leq \tau^*] = 0$ *for all* t.

Proposition 4 establishes the conditions under which assets eventually become as liquid as possible or fall into the illiquidity trap. Part 1 says that if parameters are such that there is no illiquidity trap, then learning always takes place. Eventually, large samples will accumulate and agents will learn the value of μ perfectly. Informativeness will converge to its maximal level $\bar{\tau}$ and so will liquidity. Part 2 says that if there is an illiquidity trap and the economy begins in that state, then learning never takes place and assets will remain illiquid, that is, the economy will repeat the initial period static equilibrium forever.

If there is an illiquidity trap but the economy begins outside it, some learning will take place at first. Since the sequence of realizations that agents learn from is random, there is a positive probability that it leads beliefs g_t toward 0.5. Parts

3 and 4 distinguish between two possibilities. Part 3 says that if the illiquidity trap is large, then as g_t approaches 0.5, the economy will fall into the trap, so the probability of falling into it is positive no matter what the initial beliefs were. Part 4 says that if the trap is not large, a single observation may lead beliefs to jump across the illiquidity trap; in this case, there are initial beliefs such that no sequence of observations leads into the trap. Conversely, if initial observations lead g_t away from 0.5, liquidity increases, and there is a positive probability that it converges to $\bar{\tau}$. The probabilities that τ_t falls into the trap or converges to $\bar{\tau}$ can be bounded. Condition (15) says that the closer the economy begins to the edge of the illiquidity trap, the higher the probability that it will fall into it, while the further away from the illiquidity trap it begins, the more likely it is to avoid it.[9]

The reason for the possibility of two long-run outcomes is that the model features a form of dynamic strategic complementarities in trading decisions. Higher trading at date t results in a higher rate of learning, which leads over time to more precise estimates of μ and therefore more informativeness. This informativeness sustains higher levels of trading and learning. Due to this complementarity, there is path dependence: randomness in the early stages of the learning process may have long-term consequences for how financial markets develop, a feature of many models with complementarities. This suggests a possible role for government intervention to get the learning cycle started (or restarted if it starts but then enters the illiquidity trap), for instance, by subsidizing the early stages of the financial industry. This is not unlike other "big-push" policies based on learning-by-doing externalities (see Easterly (2006) for a skeptical look at this type of argument).

The relationship between measures of market volume and liquidity has been the subject of a large empirical literature (Easley, Kiefer, and Paperman (1996), Chordia, Roll, and Subrahmanyam (2001), Amihud (2002)). The positive association is often attributed to some form of thick market effect through easier search or increased competition. Here, the effects are through information asymmetry and are both static and dynamic. The static effect is standard: information asymmetry interferes with trade, lowering volume and liquidity. Glosten and Milgrom (1985), Kyle (1985), and indeed Akerlof (1970) examine aspects of this effect.

The dynamic effect comes from the learning-by-doing assumption: higher volume leads to more learning, which leads to lower information asymmetry and higher liquidity in future periods. A related dynamic effect is at play in the model analyzed by Glosten and Milgrom (1985). In that model, traders learn from observing whether the previous trader chose to buy at the ask price, sell at the bid price, or not trade, which partly reveals that trader's information. In that world, it is also true that trading volume determines the speed of learning and future information asymmetry. The key difference between the mechanisms is that Glosten and Milgrom (1985) assume that

[9] This probability can also be computed exactly, but the value depends on $\bar{\omega}$. The bounds in (15) hold for any $\bar{\omega}$.

trades are observable, whereas I assume that they are not. Which assumption is more appropriate depends on the application. Observability is probably a good assumption for relatively centralized markets with transparent record-keeping, such as the stock market, but less so for other contexts. By assuming that trades are not observable, I focus on information that is produced as a *by-product* of trade, for instance, through the mechanisms described in Appendix B. Note that, in my setting, if traders could perfectly observe how many A-labeled assets are actually sold, they could then infer how many of them were lemons, and from this, they would infer μ exactly since there is no other source of noise.

The learning-by-doing mechanism at the heart of the model is similar to that assumed by Veldkamp (2005), van Nieuwerburgh and Veldkamp (2006), Ordoñez (2013), and Fajgelbaum, Schaal, and Taschereau-Dumouchel (2014), all of whom assume that there is an association between the level of economic activity and the rate at which agents learn. This literature focuses on the problem of agents who need to learn the level of aggregate productivity to make production and investment decisions. Fajgelbaum, Schaal, and Taschereau-Dumouchel (2014) combine this idea with the assumption of irreversible investment and show that such a model can generate uncertainty traps, where investment is persistently low. Just like the illiquidity traps I describe, there are two parts to this logic: a link from uncertainty to a lack of investment and a link from a lack of investment to a lack of learning. The latter works the same way in both models. The former works differently. The uncertainty trap model relies on the option value of waiting, so it is well suited to thinking about physical investment, where irreversibility is a natural assumption. Instead, the illiquidity trap relies on the logic of Akerlof (1970), so it is well suited to thinking about environments in which information asymmetries are important, even if there are no irreversibilities, as is arguably the case for financial assets. Note also that uncertainty about productivity (i.e., about θ) would have no effect on the model. Buyers are risk neutral, so formulas (6) and (7) would still apply, with $\mathbb{E}(\theta)$ replacing θ. Good news about θ would increase the gains from trade and equilibrium prices but the resolution of uncertainty per se would not be important.[10] For instance, one would not expect systematic increases in activity after the release of aggregate economic data the way one would under the uncertainty logic. Instead, one should expect trade to increase following the release of ex-post performance indicators of recently traded individual assets.

More broadly, models in which agents learn about aggregate fundamentals have different implications for the relationship between learning and the volume of trade depending on how the reasons for trading are modeled. For instance, if trade results from (price-sensitive) noise traders, then when rational traders learn precise information about fundamentals, they become more willing to accommodate noise trader demand at a smaller price impact, which will increase the total volume of trade. Conversely, in models in which trade results

[10] And, as discussed above, equally good news about θ and the seller's productivity Z would just scale everything up without affecting the volume of trade.

from differences in traders' priors, learning about fundamentals will reduce the scope of disagreement and therefore the volume of trade.

Learning in the model has a time dimension: traders observe signals and update their beliefs, and these beliefs affect trading behavior in future periods. A related literature (Admati (1985), Cespa and Foucault (2014), Chabakauri, Yuan, and Zachariadis (2016)) has studied multiasset noisy rational expectations models, in which learning has a cross-sectional dimension: observing the prices of asset A is informative about asset B. Cespa and Foucault (2014) show that, in such a model, a decrease in liquidity in the market for asset A makes its price less informative about asset B and, in turn, lowers asset B's liquidity. This is a similar link between trading in one market and liquidity in another, except on a cross-sectional rather than temporal dimension. There are three main differences between this mechanism and the dynamic feedback that I study. The first is more technical than substantive: the dynamic learning channel is fundamentally asymmetric, with learning affecting future markets but not past markets. This means that, instead of involving a fixed point with all the markets at once, the model can be solved by carrying a state variable that represents the current state of knowledge. Given the simplicity of formulation (2), this is just a scalar with an easily derived law of motion. Second, in noisy rational expectations models, the relationship between trading volume and information generation could go in either direction, depending on what drives the variation. If the variation is driven by traders' risk aversion, there is a positive association: less risk-averse traders trade more aggressively on their information, which increases volume and makes prices more informative. Instead, if the variation is driven by the volume of noise traders, there is a negative association: more noise traders lead to higher volume but less informative prices. In contrast, I assume that trading activity itself is what generates information, so the association is positive by assumption. Finally, the cross-market learning mechanism in the noisy rational expectations models is driven by the fact that asset payoffs are correlated. In my setting, since the object that agents are learning about is a parameter of the information structure instead of the payoff of a particular asset, learning it is useful in future periods even though the assets themselves are i.i.d.

Figure 2 provides an example where the economy may or may not fall into an illiquidity trap depending on the initial observations. In this example, the frequency of observations ϕ is proportional to the total volume of trade x, so that if trade breaks down, learning breaks down too. Initial beliefs imply $\tau_0 > \tau^*$, so the economy begins outside the illiquidity trap. The figure shows two possible realizations. In both cases, beliefs oscillate in response to the first few observations, which include instances of both $\{s, z\} \in \{\{A, H\} \cup \{B, L\}\}$ (which raises g) and $\{s, z\} \in \{\{A, L\} \cup \{B, H\}\}$ (which lowers g). Since the true value of μ is $\bar{\mu}$, the first type of observation is more common, so over time g drifts upward toward $g = 1$, which leads to the maximum possible level of informativeness, asset prices, and trading. However, in one of the realizations, the early observations include many instances of $\{s, z\} \in \{\{A, L\} \cup \{B, H\}\}$, upon which g decreases and τ falls below τ^*. The economy thus falls into the illiquidity trap: there is no

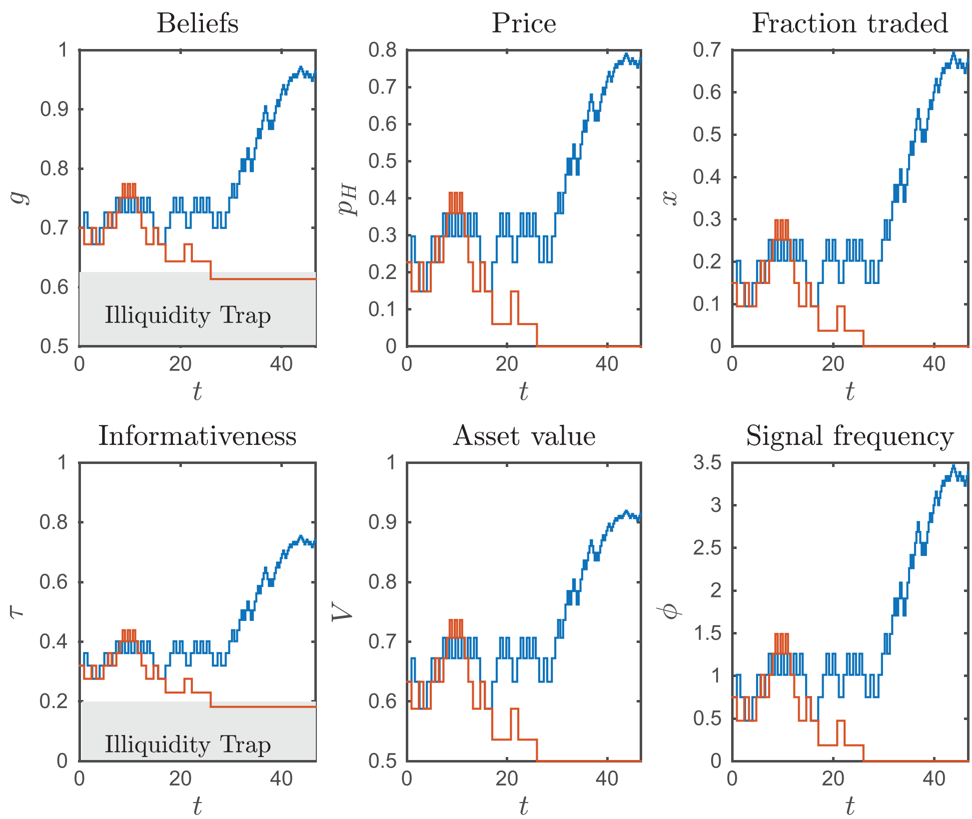

Figure 2. Two possible long-run outcomes. The example uses $\mu = \bar{\mu} = 0.9$, $g_0 = 0.7$, $\lambda = 1.5$, $Z \sim U[0, 1]$, $\theta = 1$, $\gamma = 0.54$, and $\phi = 5x$. (Color figure can be viewed at wileyonlinelibrary.com)

trade and therefore no learning. Note that, even before actually entering the illiquidity trap, whenever τ is low, the frequency of observations ϕ is low, so the learning process slows down before stopping entirely.

C. Changing μ

Suppose that, instead of being fixed forever, μ switches randomly between its two possible values at a Poisson rate δ. Since traders do not observe the changes in μ directly, they face a filtering problem: keeping track of the changing value of μ on the basis of their observations. The solution of this problem is as follows. During any interval in which no ex-post signals are observed, g_t evolves according to

$$\frac{dg_t}{dt} = \delta(1 - 2g_t), \tag{16}$$

and any time there is a new observation, g_t jumps according to (13). Equation (16) says that, during periods with no new observations, g_t mean-reverts toward 0.5.

The assumption that μ changes value is meant to capture the idea that the structure of the economy, which determines the correct way to use information, may change over time. This means, in turn, that, just like the expertise of doctors, lawyers, or car mechanics, the expertise of traders depreciates: unless new observations are added, agents are rationally aware that the true model may have changed so that their expertise is based on possibly outdated data, and hence they become less confident. As a result, even if the historical track record includes a large number of observations, g_t does not converge to zero or one over time. Even experienced traders will react to unlikely events by wondering whether these events are outliers or whether the model of the world that they are used to relying on has stopped working, which lowers liquidity. It is still the case that, on average, liquidity increases over time, but it will remain fragile even in mature markets.

LEMMA 2: *If $\delta > 0$ and there is an illiquidity trap, then $\tau_t \rightarrow_{a.s.} 0$.*

Lemma 2 says that if there is an illiquidity trap, the economy will eventually fall into it and never emerge. This is in contrast to Proposition 4, which says that with a constant μ, the economy can avoid the illiquidity trap in the long run with positive probability. The difference is that with constant μ, experience does not depreciate, so as beliefs converge to either $g = 0$ or $g = 1$ the sequence of unlikely events that is needed for the economy to fall back into the illiquidity trap gets longer and less probable. Instead, when μ can change, unlikely events are not the only way to drive the economy into the illiquidity trap: periods with no observations prevent beliefs from converging to $g = 0$ or $g = 1$, so the sequence of unlikely observations that leads to the illiquidity trap remains bounded.

If there is no illiquidity trap, standard arguments imply that informativeness will converge to an invariant distribution. In the long run, informativeness will fluctuate. Sometimes, there will be many recent observations that all point in the same direction, giving agents high confidence in their estimates of μ, and leading, in turn, to high informativeness and liquidity; at other times, there will be few or contradictory observations, pushing beliefs toward $g = 0.5$ and lowering liquidity. Prices, the volume of trade, and learning rates will therefore also fluctuate.

One of the sources of such fluctuations is actual changes in the true value of μ. In expectation, these changes lower informativeness.

PROPOSITION 5: *Starting from any beliefs $g_t \neq 0.5$, let $t + T$ be the first time an ex-post signal is observed after t. Expected informativeness at time $t + T$ is higher if μ has not changed between t and $t + T$, that is, $\mathbb{E}(\tau'_{t+T} | \mu_{t+T} = \mu_t) > \mathbb{E}(\tau'_{t+T} | \mu_{t+T} \neq \mu_t)$.*

Proposition 5 establishes that informativeness, and therefore liquidity, is higher when the economy remains stable than when it suffers shocks. If the

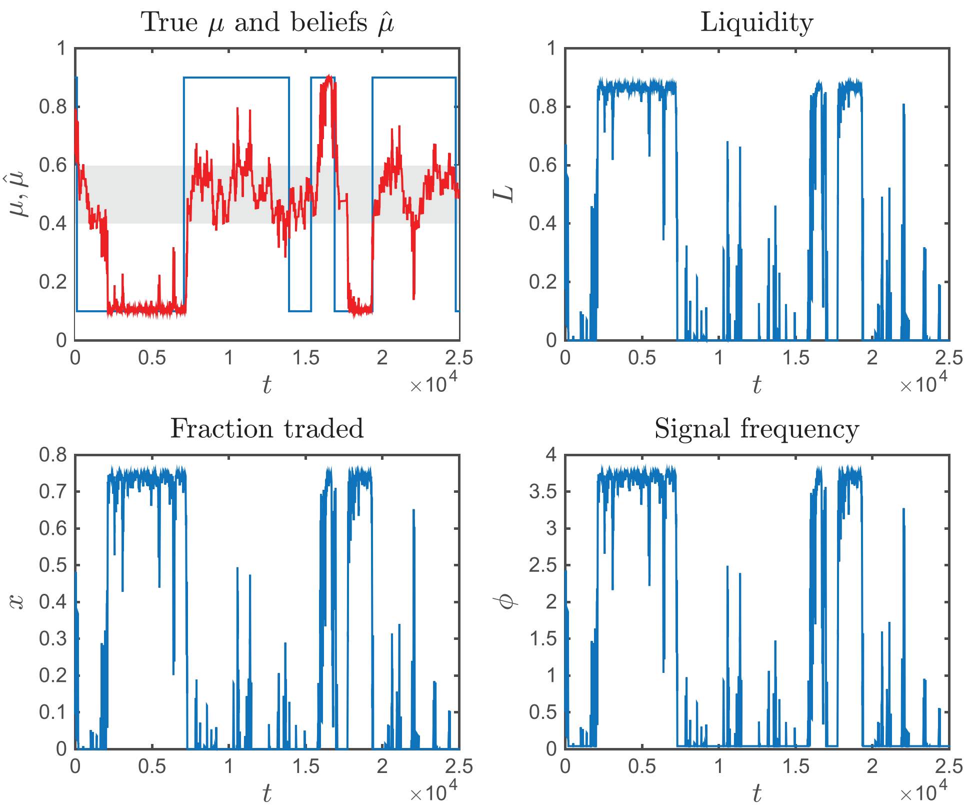

Figure 3. Simulated path. The example uses $\mu = \bar{\mu} = 0.9$, $g_0 = 0.7$, $\lambda = 1.5$, $Z \sim U[0, 1]$, $\theta = 1$, $\gamma = 0.54$, $\phi = 5x + 0.04(1 - x)$, and $\delta = 0.0001$. The gray area in the first panel shows the region of $\hat{\mu}$ in which trade breaks down. (Color figure can be viewed at wileyonlinelibrary.com)

true value of μ has changed since the last observation, it is more likely that the next observation will conflict with agents' beliefs, which will make agents uncertain about the true value of μ, and thus decrease liquidity. In periods of stability, it is more likely that new observations will reaffirm agents' priors, which will shift beliefs away from $g_t = 0.5$, increasing the informativeness of signals and, in turn, liquidity.

Figure 3 plots a simulated path for the same economy as in Figure 2 with two minor differences. First, the arrival rate of ex-post signals is $\phi = 5x + 0.04(1 - x)$, which means that it is positive even when $x = 0$ (i.e., when trade breaks down). This implies that there is no illiquidity trap. Second, the value of μ changes with a Poisson intensity $\delta = 0.0001$, so expertise depreciates at a small positive rate.

In the example, at first the signals move the economy toward the region of beliefs where assets are completely illiquid; this decreases the rate of learning process, so the economy remains there for a long time. Eventually, a string of signals that all indicate that $\mu = 1 - \bar{\mu}$ arrive close to each other and the

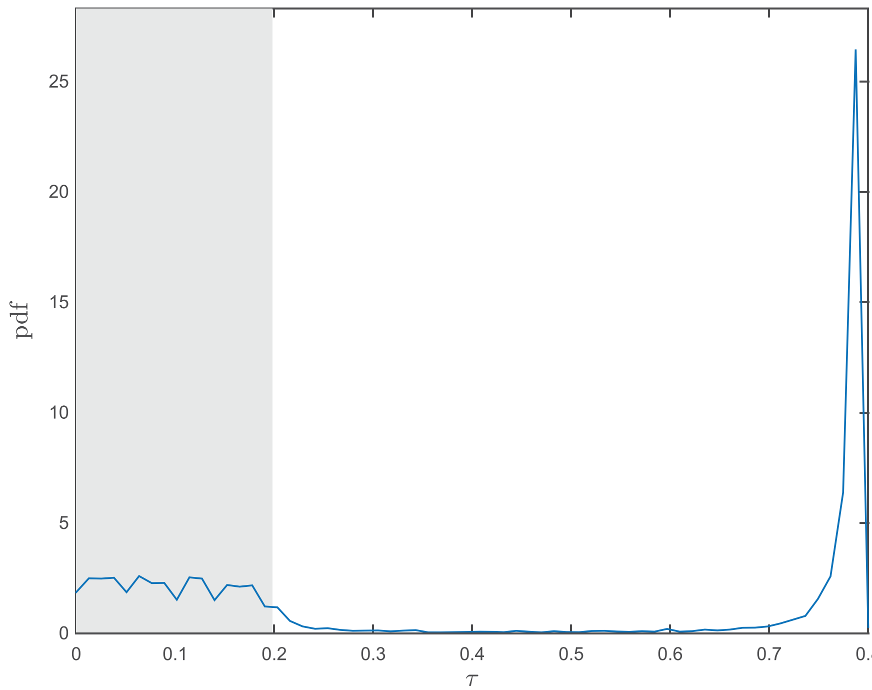

Figure 4. Long-term distribution of informativeness. This distribution is computed by sampling 10,000 equally spaced observations over a simulation of 100,000 periods in calendar time and approximating the density with a Gaussian kernel. The gray area shows the region of τ in which trade breaks down. (Color figure can be viewed at wileyonlinelibrary.com)

economy escapes the illiquid region. The learning rate accelerates, so agents receive a lot of signals that, on average, confirm μ has not changed and liquidity remains high for many periods. Eventually, a string of signals arrives that points (correctly) to μ having changed sign, so the economy becomes illiquid again and the cycle repeats itself, with relatively long periods of sustained liquidity and sustained illiquidity. Liquidity and the arrival rate of signals all comove with beliefs.

Figure 4 plots the long-term distribution of informativeness. Since both the rate of learning when there is no trade and the depreciation rate of expertise are low, the economy spends most of the time in states in which informativeness is either close to its maximum $\bar{\tau}$ or at levels that are insufficient to sustain trade. This shows that the model's predictions are continuous with respect to whether there is an illiquidity trap. If a lack of trade leads to a large slowdown in the rate of learning but not quite a complete shutdown, then liquidity will eventually recover, but will spend a large amount of time in an illiquid state.

Figure 5 plots an impulse response of the economy if, starting from the median level of liquidity, it receives a series of signals that push beliefs into the region in which assets are illiquid. In the example, illiquidity makes the

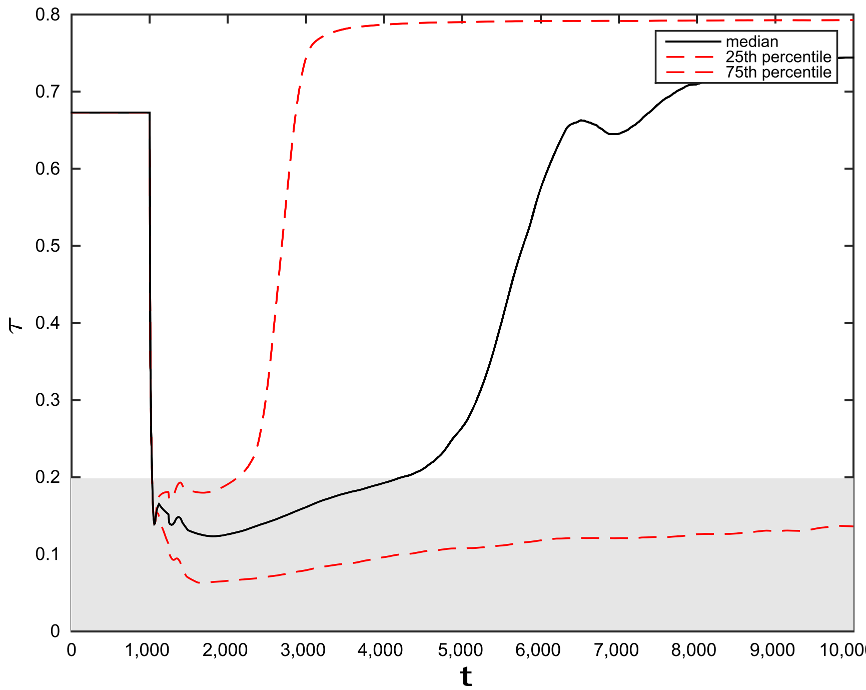

Figure 5. Impulse response. The figure plots the impulse response after, in period 1,000, a series of consecutive unlikely observations shift the economy into the illiquid region. The quantiles are computed over 5,000 simulations, each 10,000 periods long in calendar time. (Color figure can be viewed at wileyonlinelibrary.com)

learning process very slow. It therefore takes (on average) a long time for liquidity to recover after the shock. Note that the example is meant as an illustration and not as a quantitative statement. The point is that if signal frequency $\phi(\tau)$ is very low when $\tau < \tau^*$ (i.e., when trade breaks down), then recovery after assets become illiquid may take a long time. This mechanism could help explain why recoveries after financial crises tend to be slow (Cerra and Saxena (2008), Reinhart and Rogoff (2009)): unexpected shocks make financial expertise outdated, liquidity falls, and recovery requires rebuilding the stock of expertise.

In this example, shocks affect information directly, but the learning dynamics could also transmit shocks that originate elsewhere. For instance, negative shocks to buyers' productivity can reduce trade, which would decrease the rate of learning and lead to lower liquidity in the future. Kurlat (2010) shows how this effect can create persistence in a business cycle model.

IV. Conclusion

Introducing social learning dynamics into a standard model of trade under asymmetric information yields a rich set of implications. It provides a dynamic link between liquidity and trading volume, a theory of why liquidity increases over time, why it can be volatile, why it falls after unexpected events, and why illiquidity can become long-lasting.

Initial submission: September 25, 2016; Accepted: April 3, 2017
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

Appendix A: Proofs

A. *Proof of Lemma 1*

1. Using (8) and (9), equation (10) can be rewritten as

$$V(\tau) = \frac{1+\tau}{2}\left[F(p_H)\theta + \int_{Z\geq p_H(\tau)} Z dF(Z)\right] + \frac{1-\tau}{2}\left[F(p_L)\theta + \int_{Z\geq p_L(\tau)} Z dF(Z)\right],$$

   so the inequalities follow immediately from Assumption 1.
2. Plugging $\tau = 1$ into (8) and (9) gives $p_H = \theta$ and $p_L = 0$. Replacing this in (10) gives the result.
3. (a) For $\tau = 0$, both (8) and (9) reduce to

$$p = \frac{F(p)}{F(p)+\lambda}\theta.$$

   If condition (11) holds, then there is no positive solution, and thus the only solution is $p = 0$. The right-hand side of equation (9) is decreasing in τ and lies below the 45-degree line for $\tau = 0$. It must therefore lie below the 45-degree line for any τ, so the only solution of equation (9) is $p = 0$.

   (b) Since condition (11) is a strict inequality, by continuity, there is a neighborhood around $\tau = 0$ where (8) has no positive solution. The fact that $p_H(\tau)$ is strictly increasing whenever it is positive results from the fact that the right-hand side of equation (8) is increasing in τ. This, in turn, implies that the cutoff τ^* is unique.

   (c) This follows by replacing $p_L = p_H = 0$ in equation (10).

   (d) Since under condition (11), we have $p_L = 0$, equation (10) reduces to:

$$V(\tau) = \frac{1+\tau}{2}\left[F(p_H)\theta + \int_{Z\geq p_H(\tau)} Z dF(Z)\right] + \frac{1-\tau}{2}\mathbb{E}(Z),$$

so

$$\frac{dV}{d\tau} = \frac{1}{2}\left[F(p_H)\theta + \int_{Z \geq p_H(\tau)} ZdF(Z) - \mathbb{E}(Z)\right] + \frac{1+\tau}{2} f(p_H)(\theta - p_H)\frac{dp_H}{d\tau}$$

$$= \frac{1}{2}\left[\int_{Z \leq p_H(\tau)} (\theta - Z)dF(Z) + (1+\tau) f(p_H)(\theta - p_H)\frac{dp_H}{d\tau}\right] \geq 0,$$

where the last inequality follows from the fact that $p_H(\tau) < \theta$ and $\frac{dp_H(\tau)}{d\tau} \geq 0$ by part 3b.

B. Proof of Proposition 1

Assume w.l.o.g. that $g_t > 0.5$, so $\tau_t = 4(g_t - 0.5)(\bar{\mu} - 0.5)$. Let n be the number of ex-post signals that agents observe between t and t'; n is a random variable with a Poisson distribution with parameter $\bar{\phi} = \int_t^{t'} \phi_s ds$. Let r be the fraction of these observations that are either $\{s = A, z = H\}$ or $\{s = B, z = L\}$. Then, by (13), we have

$$g_{t'} = \frac{(\bar{\omega}^r(1-\bar{\omega})^{1-r})^n g_t}{(\bar{\omega}^r(1-\bar{\omega})^{1-r})^n g_t + ((1-\bar{\omega})^r \bar{\omega}^{1-r})^n(1-g_t)}. \tag{A1}$$

Now compute $\mathbb{E}(\tau_{t'})$

$$\begin{aligned}
\mathbb{E}(\tau_{t'}) &= 4(\bar{\mu} - 0.5)\mathbb{E}(|g_{t'} - 0.5|) \\
&> 4(\bar{\mu} - 0.5)\mathbb{E}(g_{t'} - 0.5) \\
&= 4(\bar{\mu} - 0.5)(g_t - 0.5) \\
&= \tau_t.
\end{aligned}$$

The first line is true by definition. The second line follows from the fact that if there is no illiquidity trap, then $\bar{\phi} > 0$. Since the Poisson distribution is unbounded, there is a strictly positive probability of n large enough and r low enough that, using (A1), $g_{t'} < 0.5$. The third line follows from the law of iterated expectations. The last line is again by definition.

C. Proof of Proposition 2

1. Assume w.l.o.g. that $g_t > 0.5$. Equation (13) implies that after an unlikely observation, beliefs are given by

$$g_t' = \frac{(1-\bar{\omega})g_t}{(1-\bar{\omega})g_t + \bar{\omega}(1-g_t)} < g_t. \tag{A2}$$

If $g_{t'} \geq 0.5$, this immediately implies that $\tau_{t'} < \tau_t$. If $g_{t'} < 0.5$, $\tau_{t'} < \tau_t$ holds unless

$$0.5 - g_{t'} \geq g_t - 0.5. \tag{A3}$$

Assume (A3) holds. Then, we have

$$\begin{aligned} 1 &\geq g_t + \frac{g_t}{g_t + \frac{\bar{\omega}}{1-\bar{\omega}}(1-g_t)} \\ &\geq \frac{\tau^*}{4(\bar{\mu}-0.5)} + 0.5 + \frac{0.5 + \frac{\tau^*}{4(\bar{\mu}-0.5)}}{0.5 + \frac{\tau^*}{4(\bar{\mu}-0.5)} + \frac{\bar{\omega}}{1-\bar{\omega}}\left(0.5 - \frac{\tau^*}{4(\bar{\mu}-0.5)}\right)} \\ \Rightarrow \tau^* &\leq \bar{\tau}\,\frac{\left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} - 1}{\left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} + 1}, \end{aligned}$$

which contradicts (14). The second inequality follows from the assumption that $\tau_t \geq \tau^*$, and the last results from rearranging and using $\bar{\tau} \equiv 2(\bar{\mu} - 0.5)$. Hence, (A3) cannot hold and $\tau_t' < \tau_t$ after an unlikely observation. Since, by Lemma 1, liquidity is increasing in τ, the result follows.

2. Assume w.l.o.g. that $g_t > 0.5$. After n unlikely observations, beliefs are given by

$$g_t' = \frac{g_t}{g_t + \left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^n (1-g_t)},$$

so for

$$n > \frac{\log\left(\frac{2(\bar{\mu}-0.5)-\tau^*}{2(\bar{\mu}-0.5)+\tau^*}\right) + \log\left(\frac{2(\bar{\mu}-0.5)+\tau_t}{2(\bar{\mu}-0.5)-\tau_t}\right)}{\log\left[\frac{\bar{\omega}}{1-\bar{\omega}}\right]}, \tag{A4}$$

$g_t' < \frac{\tau^*}{4(\bar{\mu}-0.5)} + 0.5$. If $g_t' > 0.5$, this implies that $\tau_t' < \tau^*$. Furthermore, the argument from part 1 shows that if $g_t > 0.5$ and, after the next observation, $g_t' < 0.5$, then $\tau_t' < \tau^*$, so the inequality holds in this case as well.

3. This follows directly from the fact that the right-hand side of (A4) is increasing in τ_t.

D. *Proof of Proposition 3*

Using (A2), informativeness after an unlikely observation is given by

$$\begin{aligned} \tau_t' &= 4\left(\frac{\frac{\tau_t}{4(\bar{\mu}-0.5)} + 0.5}{\frac{\tau_t}{4(\bar{\mu}-0.5)} + 0.5 + \frac{\bar{\omega}}{1-\bar{\omega}}\left(0.5 - \frac{\tau_t}{4(\bar{\mu}-0.5)}\right)} - 0.5\right)(\bar{\mu} - 0.5) \\ &= \left(\frac{\tau_t - (2\bar{\omega}-1)\bar{\tau}}{1 - (2\bar{\omega}-1)\frac{\tau_t}{\bar{\tau}}}\right). \end{aligned}$$

This expression is decreasing in $\bar{\tau} = 2(\bar{\mu} - 0.5)$ and in $\bar{\omega} = \bar{\mu}\gamma + (1-\bar{\mu})(1-\gamma)$, and therefore is decreasing in $\bar{\mu}$.

E. Proof of Proposition 4

1. Let N_t be the cumulative number of observations up to and including time t and M_t be the number of those observations that are $\{s, z\} \in \{\{A, H\} \cup \{B, L\}\}$. Then, equation (13) implies that beliefs at t are given by

$$g_t = \frac{\bar{\omega}^{M_t}(1-\bar{\omega})^{N_t-M_t} g_0}{\bar{\omega}^{M_t}(1-\bar{\omega})^{N_t-M_t} g_0 + (1-\bar{\omega})^{Mt}\bar{\omega}^{N_t-M_t}(1-g_0)}. \tag{A5}$$

   If there is no illiquidity trap, then $\lim_{t\to\infty} N_t = \infty$ and the law of large numbers implies that $\frac{M_t}{N_t}$ converges almost surely to either $\bar{\omega}$ or $1-\bar{\omega}$. Equation (A5) implies that g_{t+1} converges almost surely to either one or zero, which implies that $\tau_t \to \bar{\tau}$.
2. This part is immediate from the definition of an illiquidity trap.
3. Suppose $\tau > \tau^*$ and w.l.o.g. assume that $g > 0.5$. Suppose it were the case that an observation of $\{s, z\} \in \{\{A, L\} \cup \{B, H\}\}$ led to $g' < 0.5$. Using (13), we have

$$g' = \frac{(1-\bar{\omega})g}{(1-\bar{\omega})g + \bar{\omega}(1-g)}.$$

   Using $\tau = 4(g-0.5)(\bar{\mu}-0.5)$ and $\tau' = 4(0.5-g')(\bar{\mu}-0.5)$, we have

$$\tau' = \frac{2\bar{\tau}\bar{\omega} - (\tau + \bar{\tau})}{2\bar{\omega} - (2\bar{\omega}-1)\left(\frac{\tau}{\bar{\tau}}+1\right)}. \tag{A6}$$

   If $\tau \geq \tau^*$ and $\tau' \geq \tau^*$, this implies that

$$\tau^* \leq \bar{\tau}\,\frac{\left(\dfrac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} - 1}{\left(\frac{\bar{\omega}}{1-\bar{\omega}}\right)^{0.5} + 1},$$

   which contradicts (14). Therefore, if the illiquidity trap is large, it cannot be the case that beliefs g cross 0.5 while remaining outside the illiquidity trap. Now assume w.l.o.g. that $g_0 > 0.5$ and let $g^* = \frac{\tau^*}{4(\bar{\mu}-0.5)} + 0.5$ be the beliefs at the edge of the illiquidity trap. Since g_t cannot cross the illiquidity trap, we have

$$\Pr(g_t < 1 - g^*) = 0 \quad \forall t. \tag{A7}$$

   Furthermore, as long as g_t remains outside the illiquidity trap, the number of observations converges to infinity, so by the law of large numbers g_t must converge to either zero or 1. Since it cannot converge to zero because that would involve crossing the trap, it follows that for any ϵ,

$$\lim_{t\to\infty} \Pr(g_t > 1 - \epsilon | g_t > g^*) = 1. \tag{A8}$$

By the law of iterated expectations, $\mathbb{E}(g_t) = g_0$, which implies that

$$\begin{aligned} g_0 = {} & \mathbb{E}(g_t | g_t \geq 1-\epsilon)\Pr(g_t \geq 1-\epsilon) \\ & + \mathbb{E}(g_t | g_t \in (g^*, 1-\epsilon))\Pr(g_t \geq \in (g^*, 1-\epsilon)) \\ & + \mathbb{E}(g_t | g_t \in [1-g^*, g^*])\Pr(g_t \in [1-g^*, g^*]) \\ & + \mathbb{E}(g_t | g_t < 1-g^*)\Pr(g_t < 1-g^*). \end{aligned}$$

Taking limits and using (A7) and (A8), it follows that

$$\begin{aligned} g_0 = {} & \lim_{t\to\infty} [\mathbb{E}(g_t | g_t \geq 1-\epsilon)\Pr(g_t \geq 1-\epsilon) \\ & + \mathbb{E}(g_t | g_t \in [1-g^*, g^*])\Pr(g_t \in [1-g^*, g^*])], \end{aligned}$$

and therefore

$$g_0 \leq \lim_{t\to\infty} \Pr(g_t \geq 1-\epsilon) + g^* \left(1 - \lim_{t\to\infty} \Pr(g_t \geq 1-\epsilon)\right),$$

$$g_0 \geq (1-\epsilon) \lim_{t\to\infty} \Pr(g_t \geq 1-\epsilon) + (1-g^*) \left(1 - \lim_{t\to\infty} \Pr(g_t \geq 1-\epsilon)\right).$$

Rearranging gives (15).

4. Formula (A5) can be rewritten as

$$g_t = \frac{\bar{\omega}^{D_t} g_0}{\bar{\omega}^{D_t} g_0 + (1-\bar{\omega})^{D_t}(1-g_0)},$$

where $D_t = 2M_t - N_t$ takes integer values. Suppose that $g_0 = g^* + \epsilon$ for $\epsilon > 0$. Then, for any positive D_t, $g_t > g^*$, so beliefs are outside the illiquidity trap. For $D_t = -1$, equation (A6) applies. If the illiquidity trap is not large and ϵ is small enough, this implies that $g_t < 0.5$ and beliefs are also outside the illiquidity trap. For $D_t < -1$, g_t is even lower, so it is also outside the illiquidity trap. Therefore, no sequence of observations can lead the economy into the trap. Conversely, if $g_0 = \bar{\omega}$, then for $D_t = 1$, we have $g_t = 0.5$, which must be inside the trap and hence the trap is reached with positive probability.

F. Proof of Lemma 2

By (16) after an interval of length T with no observations starting at time t,

$$g_{t+T} = 0.5 + \frac{g_t - 0.5}{e^{2\delta T}}.$$

This implies that, for any g_t, if $T > \frac{1}{2\delta}\log(\frac{\bar{\tau}}{\tau^*})$, then $\tau_{t+T} < \tau^*$. Since in a Poisson process, an interval of length T with no observations happens with positive probability, then by the law of large numbers, it happens almost surely. After it happens, since there is an illiquidity trap, no more ex-post signals are observed and $\tau_t \to 0$.

G. Proof of Proposition 5

Assume w.l.o.g. that $g_t > 0.5$. By (16), we have

$$g_{t+T} = 0.5 + \frac{g_t - 0.5}{e^{2\delta T}} > 0.5.$$

Equation (13) implies that beliefs are given by

$$g'_{t+T}(m=1) = \frac{\bar{\omega} g_{t+T}}{\bar{\omega}^m g_{t+T} + (1-\bar{\omega})^m (1-g_{t+T})}$$

if $\{s = A, z = H \text{ or } s = B, z = L\}$ is observed and

$$g'_{t+T}(m=0) = \frac{(1-\bar{\omega}) g_{t+T}}{(1-\bar{\omega}) g_{t+T} + \bar{\omega}^{1-m}(1-g_{t+T})}$$

otherwise. The fact that $g_{t+T} > 0.5$ therefore implies that

$$\tau'_{t+T}(m=1) > \tau'_{t+T}(m=0).$$

Taking expectations, we have

$$\begin{aligned}
\mathbb{E}(\tau'_{t+T} | \mu_{t+T} = \mu_t) &= \tau'_{t+T}(m=1)\Pr(m=1|\mu_{t+T} = \bar{\mu})\Pr(\mu_t = \bar{\mu}) \\
&\quad + \tau'_{t+T}(m=1)\Pr(m=1|\mu_{t+T} = 1-\bar{\mu})\Pr(\mu_t = 1-\bar{\mu}) \\
&= \tau'_{t+T}(m=0)\Pr(m=0|\mu_{t+T} = \bar{\mu})\Pr(\mu_t = \bar{\mu}) \\
&\quad + \tau'_{t+T}(m=0)\Pr(m=0|\mu_{t+T} = 1-\bar{\mu})\Pr(\mu_t = 1-\bar{\mu}) \\
&= \tau'_{t+T}(m=1)[\bar{\omega} g_t + (1-\bar{\omega})(1-g_t)] \\
&\quad + \tau'_{t+T}(m=0)[(1-\bar{\omega}) g_t + \bar{\omega}(1-g_t)]. \qquad \text{(A9)}
\end{aligned}$$

Similarly,

$$\begin{aligned}
\mathbb{E}(\tau'_{t+T} | \mu_{t+T} \neq \mu_t) &= \tau'_{t+T}(m=1)[\bar{\omega}(1-g_t) + (1-\bar{\omega}) g_t] \\
&\quad + \tau'_{t+T}(m=0)[(1-\bar{\omega})(1-g_t) + \bar{\omega} g_t]. \qquad \text{(A10)}
\end{aligned}$$

Subtracting (A10) from (A9), we have

$$\begin{aligned}
\mathbb{E}(\tau'_{t+T} | \mu_{t+T} = \mu_t) - \mathbb{E}(\tau'_{t+T} | \mu_{t+T} \neq \mu_t) &= (\tau'_{t+T}(m=1) \\
&\quad - \tau'_{t+T}(m=0))(2\bar{\omega} - 1)(2g_t - 1) > 0.
\end{aligned}$$

Appendix B: Microfoundations for Assumption 2

In this appendix, I present three examples of models that would produce a positive relationship between the rate of information flow ϕ and informativeness τ. Given that, by equation (12), the fraction of capital traded x is an increasing function of τ, each example can also be restated in terms of a positive association between ϕ and x, justifying Assumption 2. The first two examples lead to functions $\phi(x)$ that give rise to an illiquidity trap; the third one does not.

A. Example 1: A Small Cost of Producing Ex-Ante Signals

Suppose that, instead of assets simply emitting the public signal s, the seller who owns the asset must decide whether to produce this signal and make it available to the public, and producing this signal entails a cost ϵ. The timing is as follows. First, the seller learns the type of his asset and his productivity Z. Next, he must decide whether to incur the cost of issuing signal s. If he decides to issue the signal, he learns its value (A or B) at the same time as the public, and then decides whether to sell the asset at the signal-conditional price. Assume that ex-post signals z are issued for free by all assets regardless of whether or not they trade.

Let

$$S_C(p_H, p_L, p_\emptyset; Z) = \mathbb{I}\left[\frac{1+\tau}{2}\max\{p_H, Z\} + \frac{1-\tau}{2}\max\{p_L, Z\} - \epsilon > \max\{p_\emptyset, Z\}\right], \tag{B1}$$

$$S_L(p_H, p_L, p_\emptyset) = \mathbb{I}\left[\frac{1-\tau}{2}p_H + \frac{1-\tau}{2}p_L - \epsilon > p_\emptyset\right]. \tag{B2}$$

The terms p_H and p_L denote the price conditional on good and bad signals, respectively, and $p_\emptyset$ denotes the price conditional on not issuing a signal. The indicator functions S_C and S_L determine whether sellers who own capital and lemons, respectively, will find it optimal to produce the signal s, depending on prices.

Definition B1: *Given τ, an equilibrium is given by $\{p_H, p_L, p_\emptyset\}$ such that*

1. *p_H is the highest solution to:*

$$p_H = \frac{(1+\tau)\int_{Z\le p_H} S_C(p_H, p_L, p_\emptyset; Z)dF(Z)}{(1+\tau)\int_{Z\le p_H} S_C(p_H, p_L, p_\emptyset; Z)dF(Z) + \lambda(1-\tau)S_L(p_H, p_L, p_\emptyset)}\theta, \tag{B3}$$

   taking p_L and $p_\emptyset$ as given, if such a solution exists, and $p_H = 0$ otherwise.

2. *p_L is the highest solution to*

$$p_L = \frac{(1-\tau)\int_{Z\le p_L} S_C(p_H, p_L, p_\emptyset; Z)dF(Z)}{(1-\tau)\int_{Z\le p_L} S_C(p_H, p_L, p_\emptyset; Z)dF(Z) + \lambda(1+\tau)S_L(p_H, p_L, p_\emptyset)}\theta, \tag{B4}$$

   taking p_H and $p_\emptyset$ as given, if such a solution exists, and $p_L = 0$ otherwise.

3. *$p_\emptyset$ is the highest solution to*

$$p_\emptyset = \frac{\int_{Z<p_\emptyset}[1 - S_C(p_H, p_L, p_\emptyset; Z)]dF(Z)}{\int_{Z<p_\emptyset}[1 - S_C(p_H, p_L, p_\emptyset; Z)]dF(Z) + \lambda(1 - S_L(p_H, p_L, p_\emptyset))}\theta, \quad \text{(B5)}$$

taking p_L and p_H as given, if such a solution exists, and $p_L = 0$ otherwise.

Implicitly, the equilibrium definition assumes that buyers move first and offer the highest prices consistent with zero profits, taking into account the best response of sellers in terms of both whether to produce the signal and whether to sell the asset.

PROPOSITION B1: *Let p_H^* be the highest solution to*

$$p_H = \frac{(1+\tau)F\left(p_H - \frac{2}{1+\tau}\epsilon\right)}{(1+\tau)F\left(p_H - \frac{2}{1+\tau}\epsilon\right) + \lambda(1-\tau)}. \quad \text{(B6)}$$

If condition (11) holds and $\epsilon < \frac{1-\tau}{2}p_H^*$, then there is a unique equilibrium, with $p_H = p_H^*$ and $p_L = p_\emptyset = 0$. Holders of capital with $Z < p_H - \frac{2}{1+\tau}\epsilon$ and all holders of lemons produce signals, while holders of capital with $Z > p_H - \frac{2}{1+\tau}\epsilon$ do not.

PROOF:

1. Equations (B3) and (B4) imply that, in any equilibrium, $p_H \geq p_L$. Using (B1), for a seller of capital with productivity Z to choose not to produce a signal and sell at $p_\emptyset$ requires that

$$\begin{aligned} p_\emptyset &> \frac{1+\tau}{2}\max\{p_H, Z\} + \frac{1-\tau}{2}\max\{p_L, Z\} - \epsilon \\ &\geq \frac{1+\tau}{2}p_H + \frac{1-\tau}{2}p_L - \epsilon, \end{aligned}$$

whereas using (B2), a seller of a lemon will choose not to produce a signal if

$$p_\emptyset > \frac{1-\tau}{2}p_H + \frac{1-\tau}{2}p_L - \epsilon.$$

Since $p_L \leq p_H$, this threshold is lower, so any $p_\emptyset$ such that $S_C(p_H, p_L, p_\emptyset; Z) = 0$ for some $Z < p_\emptyset$ will have $S_L(p_H, p_L, p_\emptyset) = 0$. Using (B5), this implies that

$$p_\emptyset \leq \frac{F(p_\emptyset)F(p_\emptyset) + \lambda}{\theta},$$

so, by Lemma (1), there is no solution to (B4) with $p_\emptyset > 0$. Hence, any equilibrium must have $p_\emptyset = 0$.

2. Replacing $p_\emptyset = 0$ in (B1), we have

$$\begin{aligned} S_C(p_H, p_L, p_\emptyset; Z) &= \mathbb{I}\left[\frac{1+\tau}{2}\max\{p_H, Z\} + \frac{1-\tau}{2}\max\{p_L, Z\} - \epsilon > Z\right] \\ &\geq \mathbb{I}\left[\frac{1+\tau}{2}\max\{p_H, Z\} + \frac{1-\tau}{2}Z - \epsilon > Z\right] \\ &= \mathbb{I}\left[Z < p_H - \frac{2}{1+\tau}\epsilon\right], \end{aligned} \tag{B7}$$

which, using (B3), implies that

$$p_H \geq \frac{(1+\tau)F\left(p_H - \frac{2}{1+\tau}\epsilon\right)}{(1+\tau)F\left(p_H - \frac{2}{1+\tau}\epsilon\right) + \lambda(1-\tau)} \tag{B8}$$

and therefore in any equilibrium $p_H \geq p_H^*$.
3. Since, by assumption $\epsilon < \frac{1-\tau}{2}p_H^*$, using (B2) implies that in any equilibrium, $S_L(p_H, p_L, p_\emptyset) = 1$ for any p_L.
4. Replacing $S_L = 1$, noting that S_C is bounded above by one and using (B4) implies that

$$p_L \leq \frac{(1-\tau)F(p_L)}{(1-\tau)F(p_L) + \lambda(1+\tau)}\theta,$$

so by (1) in any equilibrium $p_L = 0$.
5. Using $p_L = 0$ and $p_\emptyset = 0$ implies that (B7) and therefore (B8) hold as equalities. ■

In equilibrium, therefore, the owners of capital who plan to sell if the signal is positive and the owners of lemons produce ex-ante signals s, while the owners of capital who plan not to sell do not. Since all assets produce ex-post signals z, the number of $\{s, z\}$ pairs that are available for learning is equal to the number of ex-ante signals s that are produced. Letting $\epsilon \to 0$, this is equal to

$$\phi(\tau) = \begin{cases} F(p_H(\tau)) + \lambda & \text{if } p_H(\tau) > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{B9}$$

which is increasing in τ.

B. Example 2: Ex-Post Signals Produced by Intermediaries

Suppose that, instead of trading directly with buyers, sellers instead can trade only with intermediaries. These intermediaries are deep-pocketed, risk-neutral, and competitive, and have no valuation for any asset. They are the only ones who can observe (at no cost) the ex-ante signals s, which are produced automatically by all assets. After buying an asset, an intermediary can privately observe (also at no cost) the ex-post signal z for the asset, and can then sell it to buyers. If they want to, intermediaries can credibly disclose both the ex-ante and the ex-post signals that they observed for each asset. The standard unraveling logic of Milgrom (1981) and Grossman (1981) applies, so intermediaries will choose to fully disclose both s and z. Furthermore, since they have no valuation for the asset, they will sell all the assets they buy. Since no further selection takes place, the law of iterated expectations implies that the expected price, conditional on s, that buyers will be willing to pay after observing $\{s, z\}$ is equal to what they would have been willing to pay having observed s only. Therefore, intermediaries will make zero profits by buying from sellers at the equilibrium price p_S from the baseline model, disclosing the $\{s, z\}$ signal pair, and reselling. The volume of signal pairs produced will therefore be equal to the total volume of trade. Assume w.l.o.g. that $g_t > 0.5$, so only A-labeled assets trade. The total volume of trade (and therefore the rate of flow of signals) will be

$$\phi = \underbrace{\mu F(p_H(\tau))}_{A\text{-labeled capital}} + \underbrace{\lambda(1-\mu)\mathbb{I}(p_H(\tau) > 0)}_{A\text{-labeled lemons}},$$

which is increasing in τ.

C. Example 3: Variable Investment

Suppose that, instead of being endowed with assets, sellers had to incur a cost to produce them. Let $c(K)$ be the cost of producing K units of capital, where $c(\cdot)$ is increasing, concave, and differentiable. Assume that λK lemons are produced as a side product of producing K units of capital, and each seller produces a representative portfolio that includes the aggregate proportion of each type of asset with each type of signal.[11] Furthermore, suppose that all assets always produce ex-ante and ex-post signals that are publicly observable at no cost. Since the relative supply of capital and lemons is the same as in the baseline model, the prices given by equations (8) and (9) and the ex-ante value given by (10) still apply. A seller's investment problem is then simply

$$\max_K V(\tau)K - c(K),$$

with first-order condition

$$c'(K) = V(\tau),$$

which immediately implies that $K(\tau) = (c')^{-1}(V(\tau))$ is increasing in τ. Since all assets produce $\{s, z\}$ pairs, the flow rate of signals is

$$\phi = (1+\lambda)K(\tau),$$

which is increasing in τ.

[11] Assuming instead that the supply of lemons is fixed independent of K would not change the conclusion. Caramp (2016) studies how market conditions differentially affect incentives to produce high- and low-quality assets.

REFERENCES

Admati, Anat R., 1985, A noisy rational expectations equilibrium for multi-asset securities markets, *Econometrica* 53, 629–657.

Adrian, Tobias, and Hyun Song Shin, 2009, Money, liquidity, and monetary policy, *American Economic Review* 99, 600–605.

Akerlof, George A., 1970, The market for "lemons": Quality uncertainty and the market mechanism, *Quarterly Journal of Economics* 84, 488–500.

Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Anderson, Richard G., and Charles S. Gascon, 2009, The commercial paper market, the Fed, and the 2007-2009 financial crisis, *Federal Reserve Bank of St. Louis Review* 91, 589–612.

Attar, Andrea, Thomas Mariotti, and Francois Salanié, 2011, Nonexclusive competition in the market for lemons, *Econometrica* 79, 1869–1918.

Banerjee, Abhijit V., 1992, A simple model of herd behavior, *Quarterly Journal of Economics* 107, 797–817.

Brown, David C., 2015, Investing in security price informativeness: The role of IPO underpricings, Working paper, University of Arizona.

Brunnermeier, Markus K., 2008, Deciphering the liquidity and credit crunch 2007-08, NBER Working Paper No. 14612.

Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Buckley, Ross P., 1997a, The facilitation of the Brady Plan: Emerging markets debt trading from 1989 to 1993, *Fordham International Law Journal* 21, 1802–1889.

Buckley, Ross P., 1997b, The transformative potential of a secondary market: Emerging markets debt trading from 1983 to 1989, *Fordham International Law Journal* 21, 1152–1238.

Caballero, Ricardo J., and Arvind Krishnamurthy, 2008, Collective risk management in a flight to quality episode, *Journal of Finance* 63, 2195–2230.

Caballero, Ricardo J., and Pablo D. Kurlat, 2009, The "surprising" origin and nature of financial crises: A macroeconomic policy proposal, in *Financial Stability and Macroeconomic Policy* (Federal Reserve Bank of Kansas City, Kansas City, MO).

Calomiris, Charles W., 1993, Is the discount window necessary? A Penn-Central perspective, NBER Working Paper No. 4573.

Caplin, Andrew, and John Leahy, 1994, Business as usual, market crashes, and wisdom after the fact, *American Economic Review* 84, 548–565.

Caramp, Nicolas, 2016, Sowing the seeds of financial crises: Endogenous asset creation and adverse selection, Working paper, MIT.

Cerra, Valerie, and Sweta Chaman Saxena, 2008, Growth dynamics: The myth of economic recovery, *American Economic Review* 98, 439–457.

Cespa, Giovanni, and Thierry Foucault, 2014, Illiquidity contagion and liquidity crashes, *Review of Financial Studies* 27, 1615–1660.

Chabakauri, Georgy, Kathy Yuan, and Konstantinos E. Zachariadis, 2016, Multi-asset noisy rational expectations equilibrium with contingent claims, Working paper, LSE.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2001, Market liquidity and trading activity, *The Journal of Finance* 56, 501–530.

Daley, Brendan, and Brett Green, 2012, Waiting for news in the market for lemons, *Econometrica* 80, 1433–1504.

Daley, Brendan, and Brett Green, 2016, An information-based theory of time-varying liquidity, *Journal of Finance* 71, 809–870.

Damodaran, Aswath, 2008, *Damodaran on Valuation* (John Wiley & Sons, Hoboken, NJ).

Dang, Tri Vi, Gary Gorton, and Bengt Holmström, 2015, Ignorance, debt and financial crises, Working paper, Columbia University.

Easley, David, Maureen Kiefer, Nicholas M. O'Hara, and Joseph B. Paperman, 1996, Liquidity, information, and infrequently traded stocks, *Journal of Finance* 51, 1405–1436.

Easterly, William, 2006, Reliving the 1950s: The big push, poverty traps, and takeoffs in economic development, *Journal of Economic Growth* 11, 289–318.

Fajgelbaum, Pablo, Edouard Schaal, and Mathieu Taschereau-Dumouchel, 2014, Uncertainty traps, NBER Working Paper No. 19973.

Foote, Christopher L., Kristopher S. Gerardi, and Paul S. Willen, 2012, Why did so many people make so many ex post bad decisions? The causes of the foreclosure crisis, NBER Working Paper No. 18082.

Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka, 2009, Do liquidity measures measure liquidity?, *Journal of Financial Economics* 92, 153–181.

Grossman, Sanford J., 1981, The informational role of warranties and private disclosure about product quality, *Journal of Law and Economics* 24, 461–483.

Hellwig, Martin, 1987, Some recent developments in the theory of competition in markets with adverse selection, *European Economic Review* 31, 319–325.

Hendren, Nathaniel, 2013, Private information and insurance rejections, *Econometrica* 81, 1713–1762.

Holmström, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *Quarterly Journal of Economics* 112, 663–691.

Hong, Harrison, Jeremy C. Stein, and Jialin Yu, 2007, Simple forecasts and paradigm shifts, *Journal of Finance* 62, 1207–1242.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Kurlat, Pablo, 2010, Lemons, market shutdowns and learning, Working paper, Stanford University.

Kurlat, Pablo, 2016, Asset markets with heterogeneous information, *Econometrica* 84, 33–85.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Levin, Jonathan, 2001, Information and the market for lemons, *RAND Journal of Economics* 32, 657–66.

Milgrom, Paul R., 1981, Good news and bad news: Representation theorems and applications, *Bell Journal of Economics* 12, 380–391.

Ordoñez, Guillermo, 2013, The asymmetric effects of financial frictions, *Journal of Political Economy* 121, 844–895.

Ordoñez, Guillermo L., 2009, Larger crises, slower recoveries: The asymmetric effects of financial frictions, Staff Report 429, Federal Reserve Bank of Minneapolis. Available at https://www.minneapolisfed.org/research/sr/sr429.pdf.

Reinhart, Carmen M., and Kenneth Rogoff, 2009, *This Time Is Different: Eight Centuries of Financial Folly* (Princeton University Press, Princeton, NJ).

Routledge, Bryan R., and Stanley E. Zin, 2009, Model uncertainty and liquidity, *Review of Economic Dynamics* 12, 543–566.

Schadrack, Frederick, and Frederick Breimyer, 1970, Recent developments in the commercial paper market, *Federal Reserve Bank of New York Monthly Review*, October, 280–291.

Shleifer, Andrei, and Robert W. Vishny, 1992, Liquidation values and debt capacity: A market equilibrium approach, *Journal of Finance* 47, 1343–1366.

Uhlig, Harald, 2010, A model of a systemic bank run, *Journal of Monetary Economics* 57, 78–96.

van Nieuwerburgh, Stijn, and Laura Veldkamp, 2006, Learning asymmetries in real business cycles, *Journal of Monetary Economics* 53, 753–772.

Veldkamp, Laura L., 2005, Slow boom, sudden crash, *Journal of Economic Theory* 124, 230–257.

Wilson, Charles, 1980, The nature of equilibrium in markets with adverse selection, *Bell Journal of Economics* 11, 108–130.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Replication Code