DISCUSSION PAPER SERIES

No. 10665

OPACITY AND LIQUIDITY

André Stenzel and Wolf Wagner

FINANCIAL ECONOMICS

Centre for Economic Policy Research

ISSN 0265-8003

OPACITY AND LIQUIDITY

André Stenzel and Wolf Wagner

Discussion Paper No. 10665
June 2015
Submitted 11 June 2015

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801
www.cepr.org

OPACITY AND LIQUIDITY

Abstract

We present a model that links the opacity of an asset to its liquidity. While low opacity assets are liquid, intermediate levels of opacity provide incentives for investors to acquire private information, causing adverse selection and illiquidity. High opacity, however, benefits liquidity by reducing the value of a unit of private information to investors. The cross-section of bid-ask spreads of U.S. firms is shown to be consistent with this hump-shape relationship between opacity and illiquidity. The analysis suggests that uniform disclosure requirements may not be desirable; optimal information provision can be achieved by subsidizing information. The model also delivers predictions about when it is optimal for asset originators to sell intransparent products or pools composed of correlated assets.

JEL Classification: D82, G14 and G18
Keywords: asset liquidity, endogenous information acquisition and opacity

André Stenzel andre.stenzel@gess.uni-mannheim.de
University of Mannheim

Wolf Wagner wagner@uvt.nl
Tilburg University and CEPR

1 Introduction

Opacity and illiquidity are two central concepts in economics. They are, however, rarely distinguished from each other. Both arise from incompleteness of information. An asset is said to be opaque when there is a common lack of knowledge about its pay-offs. By contrast, when some agents know more than others about an asset, the asset tends to be illiquid because of adverse selection problems. The difference between opacity and illiquidity thus boils down to whether the incompleteness of information is of a public or private nature.

How can the two be related? At first, one would expect a positive link between opacity and illiquidity. When there is more opacity, there is more scope for agents having different information sets. Adverse selection should then be more pronounced and liquidity be low. This reasoning is consistent with common thinking among policy makers that transparency is beneficial for the financial system: more public information should deter wasteful private acquisition of information and also reduce the potential for asymmetries among investors.

This argumentation, however, ignores the fact that private information is endogenous. Gathering it is costly; hence it has to be profitable for investors to acquire it. The relationship between opacity and liquidity will thus depend on the scope for private information as well as on the incentives to acquire such information. It is not obvious why the value of information should be higher for opaque assets. Casual observation also throws doubt on an exclusively positive link between opacity and illiquidity. Many opaque assets are frequently traded and have low bid-ask spreads. A case in point is the banking industry. Banking is considered a very opaque business. Nonetheless, the major banks are heavily traded and their stocks display high liquidity.

This paper presents a model that analyzes the link between opacity and liquidity. We consider an investor who holds an asset of given opacity. We define opacity as the fraction of states of the world in which it is publicly not known what the pay-offs are. The investor can decide how much he wants to learn about these states.[1] Doing so incurs a fixed cost

[1]Learning about a state can be thought of as understanding how an asset performs in a certain scenario,

per state. Following this, the state of the world becomes known. The investor may be hit by a liquidity shock that forces him to sell the asset to the public. Illiquidity arises at this stage since market participants anticipate that the investor will sometimes trade opportunistically on his private information.

For a completely transparent asset, there is no scope for private information. Such an asset trades without an adverse selection discount and hence is liquid. At the other extreme, for a very opaque asset the scope for private information is maximal. At the same time, however, the incentives to acquire information are low. The reason is that acquiring knowledge about a certain number of states is then less valuable as these states constitute a smaller share of the overall number of opaque states. For a sufficiently high level of opacity, it can be shown that it is never optimal to acquire any information. Complete symmetry of information is preserved and the asset is liquid. At intermediate values of opacity, however, the investor always acquires information and there is adverse selection.[2]

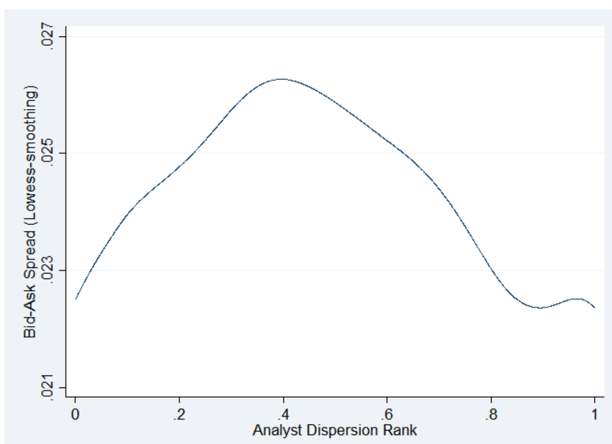

Figure 1: Relationship Between Firm Opacity and Illiquidity

e.g., an oil price shock or an economic downturn.

[2]The seminal paper by Grossman and Stiglitz (1980) considers the incentives of agents to learn about the expected pay-off of an asset. A lower quality of the signal reduces the incentives to become informed, leading to equilibrium prices reflecting fundamentals less well. While in Grossman and Stiglitz (1980) information incompleteness arises with respect to the expected pay-off of the asset (the "fundamentals"), in our analysis of opacity the latter is known. Instead, learning takes place about the mapping between (future) states of the world and pay-offs.

A key prediction of the model is a hump-shape relationship between opacity and illiquidity. We examine the cross-section of listed firms in the U.S. and find this prediction confirmed. Figure 1 summarizes the relationship using analyst dispersion as a proxy for opacity and the bid-ask spread as a proxy for illiquidity.

The main analysis considers a fixed degree of opacity of the asset held by the investor. However, since opacity affects information acquisition and liquidity, an investor's valuation of an asset will depend on its opacity. This in turn affects the incentives of originators of assets. We consider the question of how much information an original owner of an asset wants to publicly release, prior to selling to the investor. The issuer's decision is guided by two motives. First, he wants to sell an asset that maximizes value to the investor, as this will benefit him through a higher sale price. Second, he wants to minimize costs associated with releasing information to the public (arising, for example, because third parties have to be hired to certify information).

Two conclusions can be drawn from the analysis of endogenous opacity. First, it can be (privately and socially) optimal to issue opaque assets such as to deter information acquisition. This may explain why opacity in the financial system has remained high, despite the enormous improvements in information dissemination technologies in recent decades (which should have, by themselves, led to much better public information and lower opacity). It can even be desirable to artificially increase an asset's opacity beyond its natural level (for example, by resorting to complex securitization structures). Second, issuers may privately choose opacity levels that are higher than the ones that are desirable for the financial system. This occurs because issuers have to fully bear the cost of reducing opacity, but only partially internalize any benefits for other agents.

Our framework can be applied to understand other decisions of issuers. Consider for instance an originator who wants to sell a pool of assets. From a diversification perspective, such a pool should contain assets of different risk profiles. This is however in sharp contrast to the observed practice of pooling mostly similar assets.[3] Our analysis suggests that issuing correlated assets has a benefit because it lowers information costs: correlated pools avoid duplication of information because learning about one asset is then informative about the rest of the pool. However, the incentives to acquire information are higher in correlated pools, so there is a trade-off. For certain parameter values, the model predicts that informational costs are minimized by selling a correlated pool. The model also delivers testable predictions for other characteristics of asset sales, such as the decision whether to sell assets in a pool or separately.

There are several implications for policy. Uniformly mandated increases in transparency are not desirable because of the non-monotonic nature of the relationship between opacity and liquidity (which coincides with welfare in our setting). In principle, a two-class policy where regulators distinguish between assets according to their opacity can achieve efficiency. For assets that are fairly transparent, the standard policy prescription applies that more transparency increases efficiency. However, assets that are relatively intransparent to start with should not be forced to higher levels of transparency. Such a conditional transparency regime seems, however, informationally demanding.[4] We argue that a better approach is to provide subsidies for issuers to voluntarily increase transparency. Subsidies are efficiency-enhancing regardless of a firm's opacity level since they directly address the source of inefficient information choices of issuers (the positive externality of information for other agents in the financial system). They may, for example, take the form of governments sponsoring infrastructure for services that promote transparency, such as public information repositories.

[3] Gorton and Metrick (2012) refer to this as one of the major puzzles of securitization.

[4] Although differentiated disclosure policies exist in practise (for example, different standards for listed firms).

1.1 Related Literature

Our setting is closely related to recent literature which has analyzed how security design affects information acquisition by investors. While the focus in the present paper is on the question of how much information should be released about an asset, the security design literature studies how an asset's pay-off streams can be separated into different parts to make information acquisition less attractive. A central theme in this literature is the optimality of debt contracts: because debt has a flat payoff for most of the domain (and otherwise its payoff is determined by limited liability), it minimizes the benefits of acquiring private information.[5] Dang et al. (2013a) formally introduce the concept of the *information sensitivity of a security* and show in a model of strategic security design and multiple trading rounds that debt contracts minimize market participants' incentives to acquire information. Using a generalized information structure, Yang (2012) finds standard debt to be least sensitive to private information, irrespective of the composition of the underlying asset pool. Farhi and Tirole (2014) highlight the importance of commonality of information. They show that for an asset to be liquid it is important that information is symmetric. This can be achieved either by common knowledge or by common ignorance. In our paper, informational symmetry arises either for very transparent assets (common knowledge) or for very opaque assets (because of common ignorance). Intermediate levels of opacity, in contrast, lead to one-sided information and cause adverse selection.

There is a small but growing literature that analyzes asset opacity. Kaplan (2006) examines a bank's choice of whether to release information about assets at an interim stage. The paper shows that it can be efficient for the bank to commit to keep information secret, even though this forces the bank to offer non-contingent deposit contracts ex-ante. The reason is that the cost of revealing negative information at an interim stage can outweigh the benefits of positive information. Sato (2014) considers a setup with opacity at the fund

[5]The literature mostly considers situations where only one party in a potential trade can become informed, in which case information acquisition is welfare-reducing. Farhi and Tirole (2014) study information acquisition on both sides. In this case, information acquisition can improve liquidity as it can increase symmetry across agents.

and the asset level. He finds that opaque funds invest in opaque assets and that such funds can trade at a premium. The reason is that managers of opaque funds inflate investors' beliefs about future returns by (secretly) overinvesting in opaque assets and levering up.

Pagano and Volpin (2012) analyze a model where investors differ in their ability to process information. Releasing information about assets is subject to a trade-off. On the one hand, information decreases primary market liquidity because it induces a "winner's curse" problem for unsophisticated investors who cannot parse information. On the other hand, information increases secondary market liquidity as information not released by issuers creates scope for private information acquisition and hence leads to adverse selection. The second channel is also present in our model. While in Pagano and Volpin (2012) information is of an all-or-nothing nature, in our model information is continuous. This allows us to show that the value of a unit of information can vary with the asset's level of opacity, which is the source of the opacity benefit in our paper.

Carlin et al. (2013) focus on an issue similar to the differential information processing in Pagano and Volpin (2012). They consider an experimental setting in which the complexity of an asset is varied. Complexity relates to the computational difficulty required to obtain information about the asset's payoff. Carlin et al. (2013) find that when subjects are aware that other subjects are more adept at performing the required calculations, adverse selection becomes pronounced. This is consistent with agents anticipating a lower degree of common information present in markets.

While in our setting there is no social benefit to information, recent papers by Monnet and Quintin (2015) and Dang et al. (2013b) have shown that transparency (i.e. more information) can lead to more efficient interim decisions.[6] However, there is also a cost, as investors may be forced to liquidate their positions in response to negative information. In the presence of secondary markets that are not always liquid, the benefits of good interim

[6]Boot and Thakor (2001) provide an analysis of disclosure of various types of information that are all beneficial (as it reveals agent's types). They show that in equilibrium firms find it beneficial to disclose all types of information.

information cannot be fully capitalized by investors. Transparency is shown to mitigate this problem, at the cost of allocative efficiency.

The remainder of the paper is organized as follows. Section 2 sets up the baseline model for the analysis of the link between opacity and illiquidity. Section 3 examines the cross-section of U.S. firms to see whether it exhibits a hump-shape relationship. In Section 4 we consider the incentives of asset originators. Section 5 discusses some policy implications. Section 6 concludes.

2 The Model

We consider a model in which an investor can "learn" about an asset of varying degrees of opacity. This learning is not about the asset's expected pay-off (which is the focus of Grossman and Stiglitz (1980) and several other papers) but about how it pays in different states of the world. This can be likened to an investor (or the risk manager of a financial institution) exerting effort in analyzing how an asset performs under several scenarios (e.g., an oil price shock, deflation or an economic downturn). For opaque assets, this will be inherently more difficult than for transparent ones.

Take for instance the stock of Coca-Cola (or a community bank) versus the stock of JP Morgan. The business models of Coca-Cola and the community bank are simple and transparent; it is hence easy to predict how their stock will perform in a set of circumstances. By contrast, the operations of JP Morgan are extremely complex, involving a wide set of activities (such as trading in derivatives, or holdings of securitization products) which are often difficult to understand even on an individual basis. Learning about how JP Morgan's business will perform under different circumstances is hence difficult and requires substantial effort by investors. Another example is credit products. A mezzanine tranche formed from a portfolio of credits, for instance, is much more opaque than an exposure to a single name. As a result, learning about its pay-offs in different states (for instance, its dependence on a clustering of default events in the economy) is more demanding.

The economy consists of an investor I and an agent M, representing the market. There are two dates, $t = 1, 2$. The preferences of both agents are linear and given as follows:

- The investor's utility depends on whether she is patient or impatient. If patient (occurring with probability $\pi \in (0, 1)$), the investor can consume at both dates: $U^I = C_1^I + C_2^I$. If impatient, the investor derives only utility from consumption at date 1: $U^I = C_1^I$. The investor privately learns her type (patient or not) at $t = 1$ and this information is not verifiable.

- The market agent consumes at both dates: $U^M = C_1^M + C_2^M$.

The endowments of the agents are as follows. At $t = 1$, the investor holds an asset which pays off at date 2. This asset returns one in a subset L (of mass $l \in (0, 1)$) of uniformly distributed states of the world $s \in S = [0, 1]$ and zero otherwise. Given the uniform distribution, the unconditional value of the asset is hence l. The market agent has a cash endowment of $w^M > 1$ at date 1.[7] The agents hold no other endowments.

Given the allocation of endowments, it is natural that gains from trade can be realized. If the investor turns out to be impatient at date 1, she can sell the asset to M. However, reaping these gains is complicated by the opportunity for the investor to acquire private information about the asset prior to trading: Acquisition of private information results in adverse selection when trading with the market. The incentives to acquire information, in turn, are affected by the asset's opacity.

Opacity is modeled as follows. There is a set of states O containing the payoff states ($L \in O$). This set is publicly known. We refer to the mass of this set, o ($\in [l, 1]$), as the asset's opacity. Maximum opacity ($o = 1$) arises when there is no information about the set of payoff states. At the other extreme, if $o = l$, the precise set of payoff states is common knowledge and there is no scope for private information acquisition – the asset is

[7] We abstract from issues that can arise because endowments are constrained (as in Dang et al. (2013b)), which potentially lead to cash-in-the-market pricing.

transparent. For $o \in (l, 1)$, opacity is of an intermediate degree and there is incomplete knowledge about payoff states. The more transparent the asset, the smaller o and the more precise is the public information about the location of the payoff states, that is, the circumstances under which the asset pays off. Note that opacity is distinct from the asset's ex-ante return and risk: the expected payoff is l and variance of the asset is $l(1-l)$.

At the beginning of date 1, the investor has the option to acquire private information. Specifically, she decides on an amount a ($\in [0, o-l]$) of information to acquire. Following this, nature reveals a random subset of states A of mass a from O for which the asset does not pay off. Private information acquisition reduces the size of the set containing the payoff states from o to $o-a$. There are proportional costs of acquiring information $k_I \cdot a$, where $k_I > 0$. We assume that these costs take intermediate values:

Assumption 1

$$\frac{\pi l(1-\pi)}{1-\pi l} < k_I < \pi.$$

This assumption ensures that information acquisition is nontrivial.

The choice of a as well as the realization of the subset A are private to the investor and are not verifiable. Following the investor's information acquisition decision, the state of the world s becomes available. Subsequently, the investor can sell the asset to the market. For this we assume that the market posts a competitive price for the asset and the investor decides whether or not to sell at this price.[8]

To focus the analysis, it is convenient to rearrange the states s of the world. Specifically, we reorder states such that the payoff states are on $[0, l]$, the public set of potential payoff states is on $[0, o]$, and the set of potential payoff states privately known to the investor is $[0, o-a]$. In addition, agents no longer observe the exact state, but only the set in which the state falls. If $s > o$, both investor and the market learn that the state of the world falls outside the public set O, and hence that the asset does not pay off. If $s \in (o-a, o]$, the state

[8]This avoids the use of price as a signal about the asset's quality or the investor's type. A competitive price may, for example, arise if market participants compete by posting bid prices for the asset.

of the world is in the public set of possible states of the world O, but not in the investor's private set. The investor privately learns that the asset does not pay, while the market only learns that s is within the public set of potential payoff states of mass o. If $s \leq o - a$, both investor and market have incomplete knowledge about the payoff. The investor knows that s is within the private set of payoff states, while the market only observes that the state is within the public set.

The timing of the model can be summarized as follows:

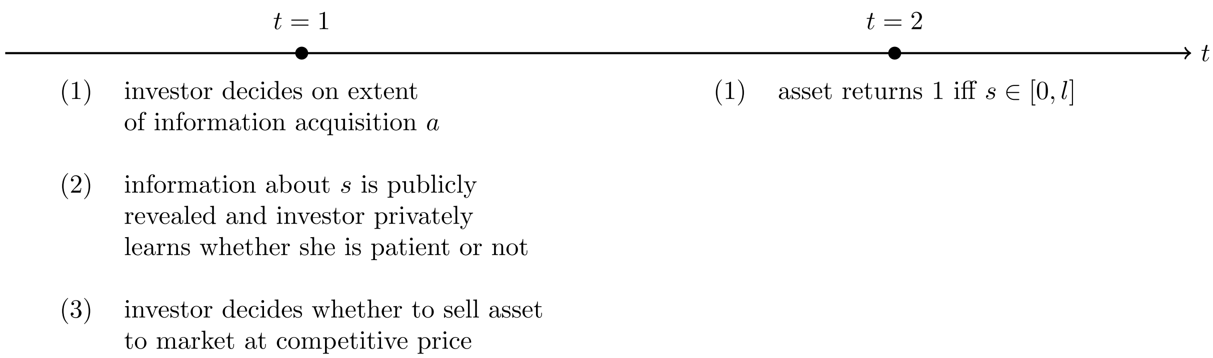

Figure 2: Timeline of the Baseline Model

2.1 Trade with the Market

To solve for an equilibrium of the game, we first analyze the final stage in which the investor has the opportunity to sell to the market. At this stage, public information about the state s has been revealed. The public set of payoff states depends on the asset's opacity level and is given by $[0, o]$. Furthermore, the investor has potentially acquired information a; her private set of payoff states is thus $[0, o - a]$. Denote by $\widetilde{a}$ the market's beliefs about how much information the investor has acquired, and by $p(\widetilde{a}, o)$ the competitive price given these beliefs and the opacity level o.

We first analyze the investor's selling decision for a given price p. To rule out no-trade equilibria, we assume that the investor has a weak preference for selling when she is impatient, and a weak preference for not selling when she is patient. We focus on pure strategy equilibria.

The following cases arise depending on the realization of s. First, there is the trivial case of s being outside the public set ($s > o$). Both the investor and the market know that the asset does not pay off and trade is irrelevant. We can ignore this case for the analysis of the trade equilibrium as trade, if it takes place, occurs at a price of zero.

Consider next the case of s being inside the public set ($s \leq o$). If an investor is impatient, she will sell regardless of price (given her weak preference for selling) since there is zero utility from holding on to the asset. For a patient investor, the decision to sell depends on whether the signal is in the private set. If the signal is outside the private set ($s > o - a$), the investor knows that the asset is worthless. She will hence sell at any positive price. If the signal is inside the private set ($s \leq o - a$), the investor's expected utility of keeping the asset is $\frac{l}{o-a}$. Taking into account the weak preference for holding on to the asset, she will hence sell the asset if and only if the price p is larger than $\frac{l}{o-a}$. Such a price, however, is inconsistent with market rationality. To see this, note that $\frac{l}{o-a}$ is higher than the value of the asset even without adverse selection (that is, when the investor only sells when she is impatient). The market can hence never break even at this price and such a price cannot prevail in equilibrium. It follows that when $s \leq o - a$ the patient investor does not sell the asset.

We can summarize these results as follows:

Lemma 1 *In equilibrium, the asset is offered to the market (when $s \leq o$) iff*

i) the investor is impatient, or

ii) the investor is patient and s is in her public set but outside her private set ($s \in (o - a, o]$).

We next solve for the price at which an asset is sold (in the case of $s \leq o$). Since the price is set competitively, the market breaks even in expectation. The price hence has to be equal to the asset's expected value (given beliefs $\widetilde{a}$) *conditional on being sold.* According to Lemma 1, the asset is sold either when the investor is impatient, or when she is patient and the state is outside her private set. The first case occurs with probability $1-\pi$ and the likelihood of the asset paying off in this case is $\frac{l}{o}$, i.e., the ratio of the size of the payoff set l to the size of the public set o. The second case, a patient investor with s outside her private set, is perceived by the market to occur with probability $\pi \cdot \frac{\widetilde{a}}{o}$ ($\frac{\widetilde{a}}{o}$ is the likelihood of the state being outside the private set given beliefs $\widetilde{a}$ about the extent of information acquisition). The asset is worthless in this case. The expected value of the asset (conditional on being sold) is hence $\frac{(1-\pi)\frac{l}{o}}{1-\pi+\pi\frac{\widetilde{a}}{o}}$. Rearranging yields the competitive price $p(\widetilde{a}, o)$ given beliefs $\widetilde{a}$ and opacity level o:

$$p(\widetilde{a}, o) = \frac{1-\pi}{o-\pi(o-\widetilde{a})} l. \tag{1}$$

Note that for $\widetilde{a} = 0$ (that is, if the market believes there is no private information) we have $p(0, o) = \frac{l}{o}$. Furthermore, $\frac{\partial p}{\partial \widetilde{a}} < 0$ because if the market believes that the investor privately acquired more information, it prices in more adverse selection as it becomes more likely that a worthless asset is offered.

2.2 Information Acquisition

Consider a candidate for information acquisition a^*, and corresponding market beliefs $\widetilde{a}$. For a^* to constitute an equilibrium amount of information acquisition, it has to be the case that a^* maximizes the investor's utility given that the market believes $\widetilde{a} = a^*$. We thus have for a^* that

$$a^* = \underset{a \in [0, o-l]}{\operatorname{argmax}}\, u(a, a^*),$$

where $u(a, \widetilde{a})$ denotes the investor's expected utility given that she chooses a level of information acquisition a and the market holds beliefs $\widetilde{a}$.

We can derive $u(a, \widetilde{a})$ as follows. With probability $1-o$, the state of the world falls outside the public set ($s > o$). In this case, the investor does not derive any utility from owning the asset as it is common knowledge that the asset is worthless. With probability o, the state of the world falls inside the public set ($s \leq o$). The investor then sells whenever she is impatient or when she is patient and the state is outside her private set ($s \in (o-a, o]$). The combined probability for this is $1 - \pi + \pi \cdot \frac{a}{o}$ and she obtains $p(\widetilde{a}, o)$ from selling the asset. When she is patient and the state is inside the private set ($s \in [0, o-a]$) she holds onto the asset. This happens with probability $\pi \cdot \frac{o-a}{o}$ and she receives (in expectation) $\frac{l}{o-a}$ from the date-2 return. Together with the information costs $k_I \cdot a$, her utility is thus

$$u(a, \widetilde{a}) = o\left((1 - \pi + \pi\frac{a}{o})p(\widetilde{a}, o) + \pi\frac{o-a}{o}\frac{l}{o-a}\right) - k_I \cdot a. \tag{2}$$

Note that when beliefs are consistent with actual information acquisition ($\widetilde{a} = a$), the above simplifies to $l - k_I \cdot a$.

Differentiating with respect to a, we obtain the marginal gain from acquiring information:

$$\frac{\partial u(a, \widetilde{a})}{\partial a} = \pi p(\widetilde{a}, o) - k_I. \tag{3}$$

Equation (3) shows that information acquisition trades off marginal benefits $\pi p(\widetilde{a}, o)$ with information acquisition costs k_I. The benefits are derived as follows: By acquiring one additional unit of information, the investor reduces her private set by one state. If this state realizes, she knows that the asset is worthless. If she turns out to be patient, she will hence sell and obtain $p(\widetilde{a}, o)$, while before she would have held a worthless asset. Note that the incentives to acquire information increase in the asset's price.

The marginal benefits in (3) are constant as they do not depend on the amount of information acquired (a). There are hence three cases to consider. If $\pi p(\widetilde{a}, o) - k_I < 0$ (or rearranging, if $p(\widetilde{a}, o) < \frac{k_I}{\pi}$), the marginal benefits are always outweighed by the marginal costs. Zero information ($a^* = 0$) thus maximizes investor utility. Likewise, if $\pi p(\widetilde{a}, o) - k_I > 0$ ($p(\widetilde{a}, o) > \frac{k_I}{\pi}$), the marginal benefits outweigh the marginal costs and the highest possible level of information acquisition ($a^* = o - l$) maximizes utility. Finally, if $p(\widetilde{a}, o) = \frac{k_I}{\pi}$, the investor is indifferent as to which level of information acquisition to choose. We can hence summarize for the investor's choice of information given beliefs $\widetilde{a}$:

$$\underset{a\in[0,o-l]}{\operatorname{argmax}}\, u(a,\widetilde{a}) = \begin{cases} 0 & \text{if} \quad p(\widetilde{a},o) < \frac{k_I}{\pi} \\ [0,o-l] & \text{if} \quad p(\widetilde{a},o) = \frac{k_I}{\pi} \\ o-l & \text{if} \quad p(\widetilde{a},o) > \frac{k_I}{\pi}. \end{cases} \tag{4}$$

This allows us to solve for equilibrium information acquisition. Note that higher opacity reduces the price p for a given belief $\widetilde{a}$ (equation (1)) and hence the incentives to acquire information. Define $\underline{o}$ as the critical opacity level which just leads to full information acquisition ($a^* = o - l$). Recall that in equilibrium, we have that $a^* = \widetilde{a}$. Inserting $\widetilde{a} = \underline{o} - l$ into $p(\widetilde{a}, \underline{o}) = \frac{k_I}{\pi}$, we obtain after rearranging: $\underline{o} = \pi l + \frac{\pi(1-\pi)l}{k_I}$. Likewise, define $\bar{o}$ as the critical opacity which deters acquisition of any information. We obtain $\bar{o} = \frac{\pi l}{k_I}$ by rearranging $p(0, \bar{o}) = \frac{k_I}{\pi}$. For intermediate values of o, an interior equilibrium arises. By solving for $\widetilde{a}$ in the condition $p(\widetilde{a}, o) = \frac{k_I}{\pi}$, we obtain for the interior equilibrium that $a^* = \widetilde{a} = \frac{(1-\pi)}{\pi}(\frac{\pi l}{k_I} - o)$.

Note that Assumption 1 ensures $\underline{o} < \min\{\bar{o}, 1\}$, which allows to summarize

Proposition 1 *The equilibrium level of information acquisition a^* is*

$$a^*(o) = \begin{cases} o-l & \textit{if} \quad o \leq \underline{o} \\ \frac{(1-\pi)}{\pi}(\frac{\pi l}{k_I} - o) & \textit{if} \quad o \in (\underline{o}, \bar{o}) \\ 0 & \textit{if} \quad o \geq \bar{o} \end{cases} \tag{5}$$

with $\underline{o} = \pi l + \frac{(1-\pi)\pi l}{k_I}$ *and* $\bar{o} = \frac{\pi l}{k_I}$.

Figure 3 shows equilibrium information acquisition $a^*(o)$ as a function of an asset's opacity o. At $o = l$, the asset is fully transparent and it is not possible to acquire information ($a^* = 0$). For values of o between l and $\underline{o}$, the maximum feasible amount of information is acquired ($a^* = o - l$). In this range, opacity increases information acquisition, as higher opacity increases the feasible amount. Beyond $\underline{o}$, however, opacity reduces information acquisition. This is until $\bar{o}$ is reached, at which point no information is acquired. Note that while in the figure we have that $\bar{o} < 1$, this is not necessarily always the case. If not,

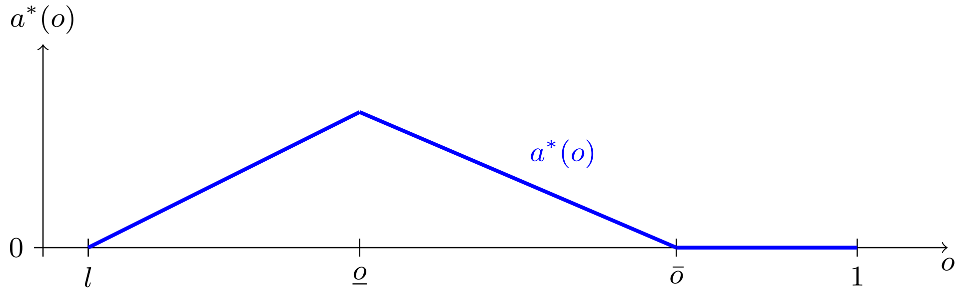

Figure 3: Information Acquisition as a Function of o

information will be acquired even at full opacity.

What is the reason why opacity can deter information acquisition? When opacity is high, the public set is large and hence a realization of s in this set becomes less informative about payoffs. This can be appreciated from the fact that p (for given $\widetilde{a}$) is declining in opacity (see equation (1)). A lower p in turn means that learning about a given number of states in the public set becomes less valuable for the investor as the investor benefits from private information by selling to the market in cases where the asset is worthless.

Note that the non-monotonic impact of opacity on information acquisition translates also into a non-monotonic impact on liquidity as well as welfare. This is, first, because information acquisition always lowers liquidity, and second, because information acquisition is the only source of welfare losses in our setting.

2.3 Robustness

In this section, we discuss several modifications of the model.

2.3.1 Random Discovery of the Payoff Interval

We have considered an information acquisition technology which is deterministic: The investor eliminates non-paying states from the set of potential payoff states with certainty. This has significantly simplified the analysis. Alternatively, the outcome of information acquisition may be random. In particular, the investor may decide to analyze a specific state and then find out whether the asset pays off in this state or not. Following this, she may decide to acquire information about more states. A consequence is that the extent to which opacity is eliminated becomes random: the investor may be lucky and discover the payoff interval early on or she may be unsuccessful and decide to stop after having acquired a certain amount of information. Another consequence is that the amount of information acquired (and hence also the deadweight loss from information acquisition) becomes random is well.

Appendix A analyzes an information technology with random discovery of states. The results from the baseline model carry over in that information acquisition is first increasing and then decreasing in opacity. There can also be interior equilibrium amounts of information acquisition, where there is a threshold for information acquisition such that an investor acquires information until this threshold is reached or until the pay-off interval is discovered.

2.3.2 Learning about Loss States

What happens when information acquisition allows us to learn about the set of states where the asset does not pay off? Suppose that – the exact opposite of the baseline model – the asset pays off on $[l, 1]$ but not on $[0, l]$. In addition, suppose that reducing opacity and information acquisition also work in the opposite way: the owner's opacity choice narrows down the set of loss states to $[0, o]$, while information acquisition further narrows it to $[0, o - a]$.

A difference to the baseline model is that the investor now benefits from states in which he has *positive* private information about the asset. The intuition for this observation is as follows (Appendix B contains the full analysis). Suppose that selling the asset yields a given price p. Suppose a state s realizes in which the investor knows that the asset pays off but the market does not ($s \in [o - a, o]$). A patient investor will then not sell the asset and thus realize a return of 1, whereas she would have realized p without information acquisition. Suppose next that a state of the world realizes where both investor and market are uncertain about whether there is a pay-out ($s \in [l, o - a]$). Since the investor also observes that this state is not within her private set of payoff-states $[o - a, o]$, she perceives a higher chance that the asset will not pay than the market. This will cause her to sell the asset when patient. However, under symmetric information, the investor would have been indifferent between selling and not selling so that no additional gains are incurred.

Information acquisition makes it more likely that a state realizes where the investor has positive information about the asset. In such a state the investor will refrain from selling the asset, while prior to information acquisition she would have sold the asset. A consequence is that the gains from information acquisition are decreasing in the market price p, the opposite to the case in the baseline model (equation (3)). This eliminates the possibility for interior choices of information acquisition. However, as shown in the appendix, it is still the case that opacity lowers information acquisition at low opacity levels and that sufficiently high opacity prevents information acquisition.

2.3.3 State-Dependent Information Acquisition Costs

The baseline model assumes that the cost of acquiring information is proportional to the number of states which are analyzed. Implicit to this is that states have equal information costs. Alternatively, information costs may differ across states. For example, it might be easier to ascertain the value of an asset in the case of an inflationary shock, than for instance in the event of a financial crisis. If costs are state-dependent, it becomes optimal for an investor to first analyze cheaper states, resulting in increasing marginal information acquisition costs. Intuitively, increasing costs make it more likely that we obtain an interior equilibrium. Appendix C contains the analysis of increasing marginal costs, showing that the quantitative results are the same as in the baseline model. In particular, the relationship between opacity and information acquisition still follows a hump-shape.

It is critical, however, that the *total* cost of gathering information is higher when more states are analyzed. To see this, suppose to the contrary that any level of information acquisition incurs a fixed cost, independently of how many states are analyzed. The marginal benefit from information once some information has been acquired ($a > 0$) is then strictly positive ($\frac{\partial u(a,\widetilde{a})}{\partial a} = \pi p(\widetilde{a}, o) > 0$ from equation (3) for $k_I = 0$). Hence there can no longer be an interior equilibrium. The investor will hence either acquire no information or all information. This case corresponds to the technology of information acquisition considered in Dang et al. (2013a) and Dang et al. (2013b).

2.3.4 Alternative Mechanisms

Opacity discourages information acquisition by reducing the price at which an asset can be sold in the case of negative private information. An opaque asset is worth less in such a sale for the following reason: when the public domain is wider, it is less likely that a state in this domain belongs to the payoff interval. The main mechanism, however, does not rely on this aspect of the model. There are various other reasons for why opaque assets generally will command lower prices. For instance, market participants may have heterogenous preferences for assets (arising, for example, because they have different endowment profiles). For transparent assets there is more knowledge about when the asset pays out, allowing market participants with the highest valuation for these states to bid for the asset. Alternatively, more transparency may allow for more efficient hedging in the transparent region, decreasing hedging costs.

In appendix E we develop a model of endogenous valuation of opaque assets based on inefficient liquidations. We consider a setting in which a market agent, after having bought the asset, may decide to scrap the asset before $t = 2$. Efficiency requires to scrap whenever the asset is worthless at $t = 2$. If the asset is opaque, it is more likely that the scrapping decision is inefficient as there are then less states where the market is informed about the future payout. Anticipating this, the market's valuation of an opaque asset will be lower at $t = 1$. This in turn makes information acquisition less attractive for investors. The appendix shows that in this richer setup it still holds that information is only acquired when opacity is not large. For sufficiently large opacity, there is no information acquisition.

There exists a second, independent, reason for why opaque assets may lower information acquisition. It arises when the informational gain from a given amount of information depends on opacity. A unit of information is conceivably less informative if an asset is very opaque as there will then be large uncertainty even after the unit has been acquired. To demonstrate, consider a situation where an asset is valuable to an agent only if it meets a criterion in every state of the world. For instance, agents may have a subsistence requirement $\overline{c}$; reaching this level of consumption gives a utility of one, if it is not reached, utility will be zero. Suppose an agent can acquire information about an asset (=project) before deciding whether to undertake it (the alternative to investment being to store funds to meet subsistence requirements). Clearly, the agent will only choose the asset if the subsistence requirement is fulfilled in every state. This requires the agent to investigate all states. Suppose that there are a discrete number o of opaque states and that in all transparent states it is known that the asset pays at least $\overline{c}$.[9] Let the probability of a payoff in an individual state meeting the subsistence requirement be $q \in (0, 1)$ and let payoffs be independent across states. The likelihood of all states meeting the criterion is then q^o. The expected benefit from acquiring full information is given by $q^o \overline{c} \cdot 1 - k_I o$, where k_I is the per-state cost of information acquisition. Dividing by o yields a benefit $\frac{q^o}{o}\overline{c} - k_I$ of acquiring a unit of information. This value of information is decreasing in opacity for two reasons: First, for higher opacity o, the likelihood that the asset will eventually meet the criteria is lower (q^o is lower). Second, for higher o more states have to be inspected, hence the gain per state is lower.

[9] If the latter is not the case, the asset is known to be worthless and hence there is also no incentive to acquire information about it.

3 The Empirical Relationship between Opacity and Liquidity

The model's key prediction is that opacity encourages private information only up to a point. Beyond this point, the relationship inverts, and opacity makes information acquisition less attractive (see Figure 3). In this section we analyze firm-level data to see whether such a pattern is consistent with the data.[10]

Following the theoretical contributions of Glosten and Milgrom (1985) and Kyle (1985), private information leads to higher bid-ask spreads on a firm's stock. Market makers need to be compensated for the risk of trading with an informed party, leading them to widen the bid-ask spread when they expect private information to be more prevalent. Similarly, in our model, the market bids less when information asymmetries are higher. Empirically, it is well documented that adverse selection is an important determinant of bid-ask spreads (see e.g. Stoll (1989), Huang and Stoll (1997)). We hence proxy the amount of private information on a firm using its stock's bid-ask spread.

The opacity of a firm is measured by the extent of disagreement among analysts about future earnings (following Flannery et al. (2004) and others). The idea is that opaque firms exhibit large potential for divergence among analysts, while disagreement is naturally limited for transparent firms. The literature has suggested alternatives to the dispersion proxy, which are however less appropriate for our purpose. For example, Morgan (2002) uses rating splits as measure of firm opacity. While conceptually similar to analyst dispersion, rating splits are not ideal in our context because testing our theory requires an opacity measure that varies over a sufficiently large interval in order to be able to identify a non-monotonic relationship (rating splits, in their simplest form, are a binary measure). Another

[10] Agarwal (2007) finds a hump-shape relationship between institutional ownership and liquidity. This is interpreted as the presence of two offsetting effects. On the one hand, higher institutional ownership leads to more informational assymmetries and hence lower liquidity. On the other hand, it leads to more competition (among institutions) which should result in pricing better reflecting information and hence higher liquidity. This exercise differs from ours in that we vary asset characteristics (opacity) rather than characteristics of the holders of the asset.

measure of opacity that is used in the literature is the number of analysts following a firm, see e.g. Roulstone (2003). This measure does not measure underlying firm opacity itself, but rather the extent to which analyst activity alleviates this opacity. We will account for this in our analysis by including the number of analysts following a stock as a control variable.

3.1 Data

We conduct an analysis of firms listed in the U.S. by relating their bid-ask spreads to the dispersion in analyst forecasts. In our analysis, we control for factors that may affect bid-ask spreads and which are unrelated to adverse selection.

We use the universe of firms contained in the CRSP database. Our measure of the bid-ask spread is the average of a firm's bid-ask spread in CRSP during the last three months of 2013 (October 2013-December 2013). In addition, we obtain various controls from CRSP: the (log) of the market capitalization as a measure of firm size and the stock price itself (both as of October 2013); the standard deviation of stock price returns, trading volume and the standard deviation of trading volume averaged over the two years prior to October 2013.

The dispersion measure is obtained from the monthly summary statistics of the I/B/E/S database. Specifically, we calculate dispersion as the average standard deviation of monthly one-year ahead earnings forecasts. We calculate this average over the two years prior to October 2013 (October 2011 until September 2013) and scale the standard deviation by the mean earnings forecast. We deliberately choose a long horizon to capture structural disagreement among analysts and to mitigate the impact of any short term events that may cause analysts to diverge or converge in their forecasts, e.g. earnings announcements. We also obtain the average number of analysts submitting forecasts following a firm from I/B/E/S.

To be included in our final data set, we require firms to have complete information during each month over which averages are computed and to have at least two analyst forecasts in any month (results are robust to requiring a higher number of analysts). As there are outliers in both the spread and dispersion measure, we exclude observations in the 1% tail in either variable. We also drop stocks with an average price of less than $5 (a common practice in the literature) because such stocks tend to trade infrequently. We arrive at a final sample of 2067 observations. Table 1 in Appendix F contains the summary statistics for all variables.

3.2 Results

We first summarize the relationship between spreads and analyst dispersion using rolling windows which sort on dispersion. Figure 4 depicts the results for a window size of 500 (the first datapoint is the mean spread of the sample of firms with the 500 smallest dispersion measure, the second datapoint is the mean of the firms with a dispersion rank between 2 and 501, ...). Up to around window 800, there is a clear positive relationship between dispersion and bid-ask spreads. However, for subsequent windows, the bid-ask spread drops significantly. There is thus a non-monotonic relationship between the two variables – as predicted by theory. Interestingly, for the extreme dispersion portfolios (starting at around window 1300) the negative relationship ceases. Spreads no longer fall (and even increase somewhat). A potential explanation for this is that at this point the degree of opacity that prevents information acquisition is reached ($\bar{o}$ in our model). We also note that the range of spreads implied by dispersion changes is large. While portfolios with intermediate dispersion have an average spread of around 0.028, spreads at the low and high end of the dispersion spectrum around are at 0.024 and 0.023 (the standard deviation over the entire sample is 0.03).

The rolling window analysis of Figure 4 is based on the raw data and is subject to the disadvantage that for any window information from the observations outside the window are completely ignored. This is an inefficient use of data and, among others, results in a more

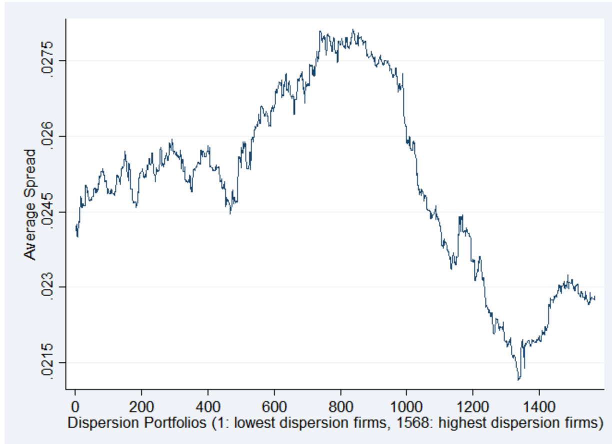

Figure 4: Rolling Windows Analysis of Bid-Ask Spreads

variable relationship in the figure. In addition, it does not allow inferences for individual specific dispersion levels, as each datapoint equally summarizes 500 data points. As an alternative, we analyze the relationship using Lowess smoothing.[11] Figure 9 in Appendix F presents the results, which confirm the rolling windows analysis. In particular, there is a monotonically increasing relationship between dispersion and spreads up to the 36th dispersion percentile of firms, after which a negative relationship obtains. Around the 86th percentile, the lowest bid-ask spread and hence highest liquidity is reached. Beyond, only smaller fluctuations in the liquidity proxy occur, consistent with opacity levels that preclude information acquisition being reached.

Previous research has indicated that bid-ask spreads reflect other factors besides adverse selection costs. It is thus important to control for these factors in the analysis. Addressing this, we analyze bid-ask spreads which are net of these factors. For this, we first regress bid-ask spreads on a set of controls and obtain residuals from this regression. We proceed to analyze the spread residuals using rolling window portfolios and Lowess smoothing (Cleveland (1979)).

As a first control, we use size (the log of market capitalization) as larger firms are

[11]Lowess smoothing (Cleveland (1979)) is based on a series of local regressions which are combined using non-parametric smoothing.

expected to have smaller spreads independent of adverse selection considerations. Second, we include the stock price, as higher price firms have a tendency to have larger spreads. We also include proxies for inventory costs, as prior literature has emphasized that such costs should result in wider spreads by market makers. A first (and inverse) proxy is trading volume (scaled by market capitalization). Higher trading volume makes it easier for market makers to adjust their inventory and should hence lead to lower spreads (see e.g. Chordia et al. (2000)). A second proxy is the standard deviation of the stock return. This variable captures firm risk, which has the effect of increasing the cost of holding inventory and results in larger spreads. We also include the number of analysts following a stock. This is because the presence of more analysts is considered to lead to more efficient information transmission to other market participants, effectively reducing private information (see e.g. Roulstone (2003)) and thus the spread. Finally, we include the standard deviation of the daily trading volume (scaled by market capitalization) as firms with volatile trading volumes require more market depth to provide smooth pricing (Roulstone (2003)).

Table 2 in Appendix F summarizes the results of a regression of the spread on these controls. Several controls are significant: market capitalization, price, volatility of returns and number of analysts. In each case, the coefficient of the significant variable has the expected sign. We calculate the residuals from this regression to separate the components that do not relate to adverse selection. Figure 5 depicts the rolling window analysis of the residuals, using the same approach as in Figure 4. The pattern looks fairly similar to the previous analysis, one noticeable difference being that there is now a more pronounced peak in residual spreads (at around window 820). Figure 10 in Appendix F presents a locally smoothed graph based on Lowess regressions which plots residual spread against dispersion rank, again showing the hump-shaped relationship.

The basic result of a non-monotonic relationship between spreads and dispersion is robust to various considerations. First, the length of the rolling window can be modified within reasonable ranges without fundamentally modifying the observed pattern (similarly,

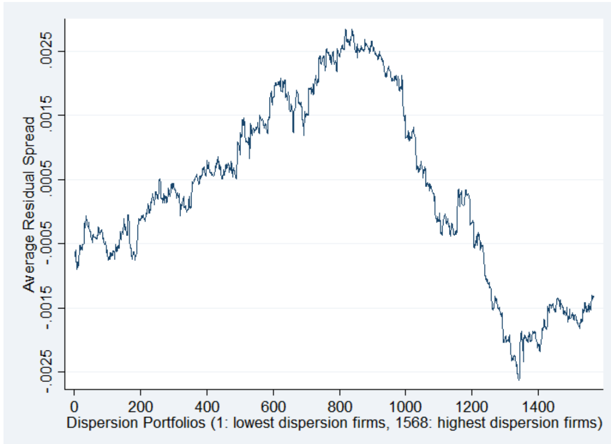

Figure 5: Rolling Windows Analysis of Bid-Ask Spreads Residuals

results are robust to variations in the bandwidth for the Lowess regression). Second, trading volume (which we use here as a control) arguably can be considered as a measure of liquidity itself. We hence re-run Figure 5 excluding trading volume in stage 1, with results unchanged. Furthermore, the results are robust to exclusion of the number of analysts as control. This is potentially important since the number of analysts may itself be a function of firm opacity, which may obscure the analysis. The analysis is also robust to requiring a high number of analysts following a firm (i.e., a minimum requirement of 5, 7 or 10 analysts following the firm). Finally, results are robust to different outlier treatments, such as including stocks with a price of less than \$5 and extending the tail cut-off for the spread and dispersion measures.

4 The Incentives of Asset Originators

In this section, we endogenize several characteristics of the asset held by the investor. For this, we consider an original owner of the asset who can influence an asset's characteristics before selling it on to the investor. We first analyze the question of how much information an owner wants to release about an asset prior to the sale. Following this, we consider implications for which assets should be sold and how to sell them. Often originators (which may for instance be banks) have several assets for sale. In this case, they can decide whether to sell them together or separately, and when they sell them together, which assets to include in the bundle.

To analyze these questions, let us assume that there is an original owner of the asset, O. Prior to selling the asset to the investor, the owner can choose (some) characteristics of the asset. The choice of these characteristics affects future information acquisition by the investor, and through this, the price at which the owner can sell in the primary market.

Incorporating the owner, the economy now consists of three agents: an owner O, an investor I, and the market M. There are three dates ($t = 0, 1, 2$) of which date 1 and 2 are identical to the baseline model. The preferences of agents are as follows:

- The owner derives utility from consumption at date 0 only: $U^0 = C_0^O$.
- The investor can now also consume at date 0. His utility when patient is hence $U^I = C_0^I + C_1^I + C_2^I$ and $U^I = C_0^I + C_1^I$ when impatient.
- The utility of the market is unchanged: $U^M = C_1^M + C_2^M$.

At date 0, the owner is endowed with the asset. The owner has no other endowment besides the asset. The investor has an endowment of w^I $(> l)$ at date 0; the market still has an endowment of w^M (> 1) at date 1.

The owner first decides on the characteristics of the asset. Following this, he can sell the asset to the investor. For this, we assume that the owner and the investor bargain and that the owner captures a fraction $\delta \in (0, 1]$ of the investor's surplus. Following this, actions proceed as in the baseline model. Figure 6 depicts the timeline.

4.1 Opacity

We first analyze the owner's choice of opacity. We assume that the owner is perfectly informed about the states in which the asset pays off. Before selling to the investor, he decides how much of this information to release. Specifically, he discloses a set of states of

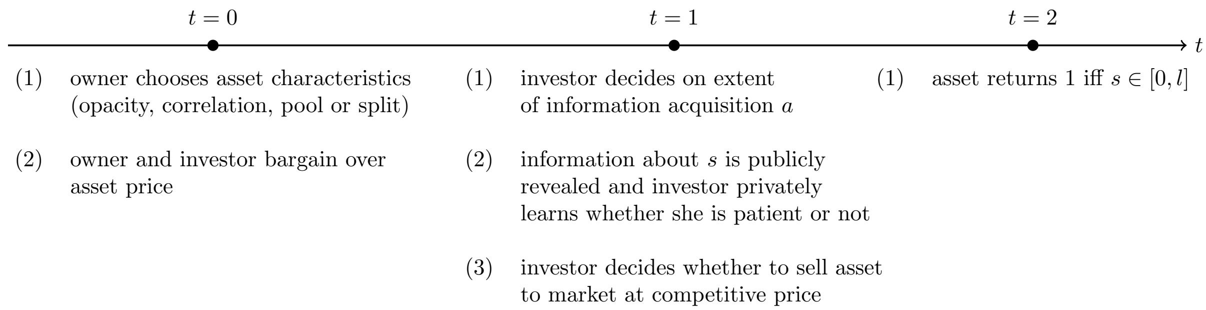

Figure 6: Timeline of Extended Model

measure o which contain the payoff states. Releasing information comes at a cost for the owner: reducing opacity from 1 to o incurs a proportional cost of $k_O \cdot (1 - o)$ $(k_O > 0)$.[12] Such costs arise because it is costly to collect information about an asset and to convey it credibly to the other agents in the economy.

4.1.1 Efficient Opacity

Since the owner does not capture the full surplus whenever $\delta < 1$, his choice of opacity may differ from the welfare maximizing one. We first solve for the welfare-maximizing opacity level and subsequently contrast it with the owner's opacity choice.

From date 1 onwards, the setup is identical to the model of fixed opacity; trading and information acquisition are still characterized by Lemma 1 and Proposition 1. We now analyze the level of opacity that maximizes welfare. Given linearity of utility, (utilitarian) welfare is simply the expected sum of resources in the economy that are available for consumption. Welfare thus consists of the endowments, $w^I + w^M$, the asset's expected payoff, l, minus the cost of reducing opacity $k_O \cdot (1 - o)$, minus the cost of acquiring information $k_I \cdot a^*(o)$:

$$W(o, a^*(o)) = w^I + w^M + l - k_O \cdot (1 - o) - k_I \cdot a^*(o). \tag{6}$$

[12] A richer model could allow assets to differ with respect to *fundamental* opacity, that is, the level of opacity before any efforts by the owner to reduce opacity.

Welfare is hence maximized by minimizing the sum of the two costs in the economy.

The opacity choice has two effects on welfare. There is the direct cost of opacity reduction $k_O \cdot (1-o)$ incurred by the owner. Furthermore, opacity affects date-1 information acquisition $a^*(o)$ and hence the information acquisition costs. Two cases arise. If $\bar{o} \leq 1$, information acquisition can be deterred by leaving the asset fully opaque, that is setting $o=1$ (see Proposition 1). As this induces neither information acquisition (and associated costs) nor costs of opacity reduction; the first best is reached.

If $\bar{o} > 1$, this is not possible. In this case, the problem can be broken down as follows. First, choosing an opacity level that leads to partial information acquisition (that is, choosing an o on $[\underline{o}, 1))$ such that $a^*(o) \in (0, o-l)$ is never optimal. A completely opaque asset ($o=1$) would dominate this choice as it would entail less information acquisition (recall that information acquisition is decreasing in opacity in the interior range) and also no opacity reduction cost. Second, when an opacity level of $[l, \underline{o})$ is chosen, all possible information is acquired ($a^*(o) = o - l$) and welfare is given by

$$W(o, a^*(o)) = w^I + w^M + l - k_O \cdot (1-o) - k_I \cdot (o-l). \tag{7}$$

Equation (7) shows that optimal opacity depends on which cost parameter is larger. If information is more costly ($k_I > k_O$), welfare is maximized by choosing the smallest opacity in the range: $o = l$. If this is not the case ($k_I \leq k_O$), the optimal choice would be to choose the largest opacity in the range: $o = \underline{o}$. However, as previously discussed, $\underline{o}$ is dominated by a completely opaque asset ($o = 1$).

It follows that to find the optimal opacity level o whenever $\bar{o} > 1$, one has to compare welfare for a fully transparent and a fully opaque asset ($o = l$ versus $o = 1$). This boils down to comparing the cost of fully eliminating opacity, $k_O \cdot (1-l)$, with the cost of investor information acquisition that arises for an entirely opaque asset, $k_I \cdot a^*(1)$.

Summarizing:

Proposition 2 *Selling an opaque asset ($o^* = 1$) maximizes welfare if*

(i) this deters information acquisition ($\bar{o} \leq 1$), or

(ii) $k_I \cdot a^*(1) < k_O \cdot (1 - l)$.

Otherwise, selling a fully transparent asset maximizes welfare ($o^* = l$).

There are three important messages. First, it can be optimal to sell a fully opaque asset – independent of the magnitude of opacity reduction costs k_O. This is because under certain conditions, full opacity prevents any information acquisition by the investor. Second, intermediate degrees of opacity are undesirable as such opacity levels induce the investor to acquire costly information. Third, if the costs of opacity reduction are sufficiently small, it can be optimal for the owner to sell a fully transparent asset, which precludes information acquisition.

Adverse selection costs: Even though there is adverse selection at the trading stage (since a patient investor sells when he has negative private information), there are no direct welfare consequences of this in our model. This is because the impatient investor and the market have identical marginal utilities of consumption. A lower market price resulting from adverse selection thus does not affect the gains from trade (the equation for welfare does not contain the price). If an impatient investor were to have higher marginal utility than the market, this neutrality no longer obtains. Appendix D analyzes this case, showing that information acquisition then has an additional, negative, effect on welfare through its effect on the equilibrium price. This, however, does not affect the key results. In particular, the hump-shaped relationship between opacity and information acquisition is still obtained.

4.1.2 The Owner's Choice of Opacity

The owner maximizes the price at which he can sell the asset to the investor, minus any cost incurred by him. Given that the investor's surplus is $l - k_I \cdot a^*(o)$, the owner maximizes

$$W_O(o, a^*(o)) = \delta\,(l - k_I \cdot a^*(o)) - k_O \cdot (1 - o). \tag{8}$$

The owner thus minimizes a combination of costs of opacity reduction and information acquisition costs. However, his objective function is not identical to the social one as he only internalizes a fraction δ of the investor's information acquisition costs.

Similar to the previous section, the solution can be derived as:

Proposition 3 *The owner sells a fully opaque asset ($o = 1$) if*

(i) this deters information acquisition ($\bar{o} \leq 1$), or

(ii) $\delta k_I \cdot a^*(1) < k_O \cdot (1 - l)$.

Otherwise, he sells a fully transparent asset ($o^* = l$).

Proof. *The owner's opacity choice mirrors the one in the baseline model. If $1 > \bar{o}$, information acquisition can be deterred and the owner can avoid costs entirely by choosing full opacity ($o = 1$). If this is not the case, he chooses either full opacity or full transparency. The respective utilities from these choices are $W_O(1, a^*(1)) = \delta\,(l - k_I \cdot a^*(1))$ and $W_O(l, a^*(l)) = \delta_O l - k_O \cdot (1 - l)$. He hence chooses full opacity iff $\delta k_I \cdot a^*(1) < k_O \cdot (1 - l)$.* ■

This yields the following corollary

Corollary 1 *The owner chooses an opacity level that is inefficiently high if and only if $\delta k_I \cdot a^*(1) < k_O \cdot (1 - l) < k_I \cdot a^*(1)$. Otherwise his choice of opacity is efficient.*

Proof. *Follows from comparing condition (ii) in Proposition 2 and 3.* ■

The intuition is clear. Since the owner incurs transparency costs fully but only internalizes a fraction of the investor's information acquisition costs, he has comparatively lower benefits from outcomes where transparency is high and information acquisition low. He hence may not sell a transparent asset even when transparency maximizes welfare.[13]

[13]Note that a sharing of rents between the investor and the market in the secondary trading stage does not lead to any bias in the owner's opacity choice as it does not create a wedge with the welfare maximizing level of opacity.

Fundamental and effective opacity: The effective opacity of an asset (that is, the opacity of an asset when sold to the investor) will in practice consist of two factors. First it consists of the fundamental opacity of the asset, determined by its business characteristics. This was the focus of the analysis in the baseline model. For example, firms in certain industries are intrinsically more opaque. Large and complex firms will also have a fundamental tendency towards higher opacity. Second, there is the opacity choice of the owner (which we focused on in this section). This choice can be understood as efforts by the owner to reduce opacity below its fundamental opacity. In cases where such efforts are not taking place, effective opacity may approximate fundamental opacity. In addition, reaching a certain level of transparency will be more costly when initial opacity is high. In practice we would thus expect to observe a wide range of levels of effective opacity, with higher levels generally corresponding to a higher fundamental opacity.

4.2 Correlation

Suppose an owner wants to sell a number of assets, for instance, through a securitization. Should he include correlated or uncorrelated assets in the sale? And does this decision depend on the characteristics of the assets available?

To analyze this, we consider the following modification of the model. At $t = 0$, the owner is endowed with two pools of assets, each containing x ($x \geq 2$) assets of fixed opacity o. The assets in each pool are individually identical to the that of the baseline model: an asset pays 1 in a mass l states of the world and zero otherwise. The investor can narrow down the set of pay-off states for each individual asset by incurring cost a. The only difference between the two pools is that in the first pool, assets pay off in exactly the same states of the world. In the second pool, the payoff-states are independently distributed across assets.

At $t = 0$ the owner decides which pool of assets to sell. This choice is public information. At $t = 1$ the investor can acquire information about each asset in the pool and subsequently sell assets to the market. We assume that assets are sold individually to market participants and that each market participant cannot observe how many assets in total the investor is selling. To focus the analysis, we analyze in the following the case of $\delta = 1$, in which case there is no conflict between the owner's incentives and the welfare maximizing outcome.

Suppose first that the owner chooses to sell the pool consisting of correlated assets. At the trading stage, the investor has to decide for each individual asset whether or not to sell it. The market has formed beliefs about information acquisition and since assets are identical, these beliefs boil down to a single parameter $\widetilde{a}$ about the investor's private information set $[o - \widetilde{a}, o]$. The decision whether or not to sell is identical to the baseline model, but now applies to x-assets at the same time. That is, the investor will sell all assets whenever she is impatient or when she privately knows that the assets are worthless ($s \in [o - a, o]$).

At the beginning of $t = 1$, the investor decides how much information to acquire about each asset. Since assets are perfectly correlated, it is strictly optimal to acquire information about one asset only. The investor thus has a single choice a, as in the baseline model. However, acquiring information about one asset now provides additional benefits: Because of perfect correlation, the investor learns about several assets at the same time. Similar to equation (2), we can write the utility of the investor as

$$u(a, \widetilde{a}) = x \cdot o \left((1 - \pi + \pi \frac{a}{o}) p(\widetilde{a}, o) + \pi \frac{o - a}{o} \frac{l}{o - a} \right) - k_I \cdot a. \tag{9}$$

From this we can derive the investor's optimal information acquisition:

Proposition 4 *The equilibrium level of information acquisition for the correlated pool of assets is*

$$a_C^* = \begin{cases} o - l & \text{if} \quad o \leq \underline{o}_C \\ (1 - \pi)(x \frac{l}{k_I} - \frac{o}{\pi}) & \text{if} \quad o \in (\underline{o}_C, \bar{o}_C) \\ 0 & \text{if} \quad o \geq \bar{o}_C \end{cases} . \tag{10}$$

with $\underline{o}_C = \pi l + x \frac{\pi(1-\pi)l}{k_I}$ *and* $\bar{o}_C = x \frac{l\pi}{k_I}$.

Proof. *Analogous to Proposition 1.* ■

Compared to the sale of a single asset, information acquisition now tends to be higher. First, the threshold at which the investor starts to acquire information is higher ($\bar{o}_C > \bar{o}$). Second, information acquisition in the interior cases is always higher ($a_C^* > a^*$ for given o). The reason for this is that since information can be applied to several assets, it becomes more attractive to acquire information.

Suppose next that the owner has sold uncorrelated assets. At the trading stage the market will again have beliefs $\widetilde{a}$ about the level of private information for each asset. These beliefs will be asset-independent due to symmetry of the setup. The trading stage for each asset is hence the same as in the baseline case. Consequently, information acquisition for each individual asset is also unchanged and given by a as laid out in Proposition 1. Total information acquisition, however is $x \cdot a^*$.

We can now turn to the owner's choice of which assets to sell. Since the owner only consumes at $t = 0$, he does not care about the assets that are retained.[14] He will hence sell the pool that obtains the highest price, which will be the one with the lowest information cost. The owner's problem is thus to identify the pool that induces the lowest amount of private information. This choice will be subject to a basic trade-off. The incentives to acquire information for an individual asset are stronger in the correlated pool, as shown above. This speaks for the uncorrelated pool. However, for a given amount of information acquired about an asset, total costs are higher in the uncorrelated pool because information is then acquired about each asset individually.[15]

The consequences for the owner's decision are as follows. When information acquisition is sufficiently unattractive ($o \geq \bar{o}_C$), there will be no information acquisition for either

[14] If the owner could also consume at $t = 2$, he would still be indifferent as to which assets are retained as both pools have the same expected payoff.

[15] Dang et al. (2013b) also analyze the impact of diversification on the incentives for information acquisition. They show that selling a diversified portfolio discourages private information acquisition by hiding private information. They do this in a setting where the cost of acquiring information is independent of the security design. Our model of endogenous information acquisition shows that while incentives to acquire information are indeed lower in the case of uncorrelated assets, correlated pools avoid duplicating private information production and may thus be preferred.

pool and the owner is indifferent between the pools. When $\bar{o} < o < \bar{o}_C$, there will be information acquisition in the correlated pool only; hence the uncorrelated pool is preferred. For lower levels of opacity ($o < \bar{o}$), information is acquired in both pools. In this case the above trade-off comes into play. If $o > \underline{o}_C$ (that is, there is incomplete information acquisition in the correlated pool), an uncorrelated pool still maximizes welfare. This can be seen by noting that interior information acquisition in the correlated pool, $a_C^* = (1-\pi)(x\frac{l}{k_I} - \frac{o}{\pi})$, is always higher than in the uncorrelated pool, $xa^* = x(1-\pi)(\frac{l}{k_I} - \frac{o}{\pi})$. However, for o that is sufficiently below $\underline{o}_C$, information costs in the uncorrelated pool dominate (information acquisition in the correlated pool even decline because they are then already at their maximum feasible level, $o - l$). The critical opacity level at which this happens is determined by the condition $o - l = xa^*(o)$. Rearranging yields:

$$\widehat{o} = \frac{x(1-\pi)\frac{\pi}{k_I} + 1}{x(1-\pi) + 1} l. \tag{11}$$

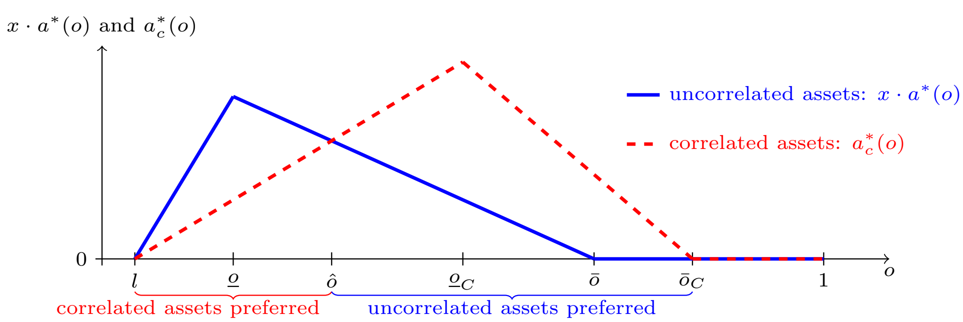

Figure 7: Information Acquisition as a Function of o

Figure 7 illustrates the different cases. We can summarize:

Proposition 5 *Consider the owner's choice to sell a correlated or uncorrelated pool of assets.*

1. *If $o \leq \widehat{o}$, the owner prefers to sell a correlated pool of assets.*

2. *If $o \in (\widehat{o}, \bar{o}_C)$, the owner prefers to sells an uncorrelated pool of assets.*

3. If $o \geq \bar{o}_C$, the owner is indifferent between both pools.

In their review of securitization practices, Gorton and Metrick (2012) identify the lack of diversification as one of the main puzzles: "The choice of loans to pool and sell to the SPV also remains a puzzle. Existing theories cannot address why securitized-loan pools are homogeneous – all credit cards or all prime mortgages, for example. The existing theory suggests that credit card receivables, auto receivables, mortgages, and so on should be in the same pool – for diversification, but this never happens." Proposition 5 shows that selling homogenous (or correlated assets) can be optimal. The reason is that this lowers the total cost of private information acquisition because information acquisition costs do not need to be spend on each individual asset.

4.3 Splitting and Pooling

Information acquisition also has consequences for whether an owner should sell cash-flows individually, or in a pool. To this end, consider that the owner has at date 0 an asset that pays x in l states and zero otherwise. The owner has the option to sell this asset in its entirety. Alternatively, he can split the asset into x smaller assets (each paying 1 in l states) and sell them to x separate investors. Assume that per-state information costs are k_I regardless of the size of the asset. In addition, assume that investors cannot credibly reveal information to each other (otherwise, one investor could obtain the information and sell them to all other investors) and that the market cannot observe how many investors are selling assets (this would reveal the private information of investors).

Consider first the sale of the asset in one piece. This case is identical to that of a correlated pool in the previous section. While for a correlated pool information acquisition for one asset applied to x assets of size one, it now applies to one asset of size x. Information acquisition is hence a_C^* as given by Proposition 4. Consider next the sale of split assets to different investors. Each investor is in the same situation as in the baseline model: he can decide to acquire information about an asset of size 1. Thus, the results from the baseline model apply. However, since there are now x investors in total, overall information costs are $x \cdot a^*$, identical to the case of an uncorrelated pool.

The decision whether or not to split the asset thus creates the same trade-off as the decision whether to sell a correlated pool. We can conclude:

Proposition 6 *Consider the owner's choice to split an asset for sale.*

1. *If* $o \leq \widehat{o}$, *the owner prefers not to split.*

2. *If* $o \in (\widehat{o}, \bar{o}_C)$, *the owner prefers to split.*

3. *If* $o \geq \bar{o}_C$, *the owner is indifferent.*

The intuition behind the trade-off is as follows. On the one hand, the incentives to acquire information for an investor who has bought the entire asset are high because private information can then be applied to an asset that pays off $x > 1$. On the other hand, when investors who have bought the split assets acquire information, information acquisition is duplicated because each individual investor will acquire information. This means that in cases where there are large incentives to acquire information, the owner should sell the entire asset in order to avoid duplication of a large amount of information.

5 Public Policy

Regulation of information disclosure by firms has a long tradition and takes many forms. Examples are requirements for listed companies to publish certified accounts at specified intervals or to disclose material information in a timely fashion. Prior to the crisis of 2007-2009, disclosure policies were predominantly targeted at protecting investors in standard securities (debt and equity). Following the breakdown of trade in various classes of asset-backed securities, a new focus of regulation is on the transparency of assets issued by financial institutions. For example, the Dodd-Frank act requires disclosure of information about asset-backed securities.

Transparency policies typically take the form of minimum standards. Issuers are obliged to follows these standards, but are free to implement higher standards of transparency. The non-monotonic nature of opacity suggests that a (uniform) minimum standard is not a desirable approach to regulation. We have shown that transparency reduces adverse selection only when transparency is sufficiently large, while increasing it otherwise. Consider Figure 4 and 9, which depict the (smoothed) cross-sectional relationship between opacity and bid-ask spreads at the firm level. The turning point at which transparency reduces asset liquidity is around the 40th percentile in both figures, suggesting that a mandated increase in transparency may increase bid-ask spreads for a large share of the population of firms. Since higher transparency brings about costs for issuers, the net effect of uniformly higher transparency may hence easily be negative.[16] Note that this does not imply that transparency regulation per se is undesirable as actual opacity levels already reflect existing efforts to enhance transparency.

Nonetheless, our analysis provides a clear rationale for regulation: issuers do not internalize the full cost of opacity for other agents in the economy and may hence choose inefficiently low amounts of transparency. Firm-specific disclosure standards that take into account that optimal opacity is heterogenous are in principle welfare-enhancing. However, as the analysis in Section 4.1 has shown, the extent to which transparency is optimal depends on deep parameters such as the cost of information to firms and investors. Regulation that conditions on these parameters seems practically infeasible.

A less demanding approach is to provide subsidies (implicit or explicit) to issuers for reducing transparency. From previous analysis we know that issuers sometimes choose inefficient opacity since they only take into account a fraction $\delta < 1$ of the full cost of opacity, $k_I \cdot a^*(1)$. A subsidy of $(1-\delta)k_I \cdot a^*(1)$ for each issuer can hence implement efficiency (essentially, a negative Pigouvian tax). And when the regulator has incomplete knowledge

[16]Kurlat and Veldkamp (2012) provide an alternative reason for why disclosure can reduce welfare. The channel is based on a general equilibrium effect. Disclosure makes assets less risky. This, in turn, will result in assets commanding a lower return in equilibrium.

about the size of externalities posed by individual issuers, he can still implement a welfare-improving policy through a subsidy that is equal to the minimum of $(1-\delta)k_I \cdot a^*(1)$ across all firms (in this case, transparency will be optimally increased at some firms – without leading to any increases in transparency that are welfare-reducing at other firms).

A subsidy could, for example, take the form of a government-sponsored rating agency that allows issuers (at their discretion) to obtain free ratings. In addition, publicly-run information repositories could help reduce the costs of providing transparency to issuers. It is crucial, however, that participation is left to the discretion of the issuers – compulsory participation suffers from the same problem as mandatory disclosure requirements.

Our analysis also speaks to the current discussion on asset correlation in the financial system. There are good reasons to believe that at the system-level financial institutions inefficiently invest in assets with high correlation, due to the externalities associated from creating systemic risk (see e.g. Acharya (2009) or Wagner (2011)). Correlation is hence typically viewed with suspicion by policy-makers. On the micro-level, there is the puzzle that securitization pools predominantly consist of similar assets, which flies in the face of diversification (Gorton and Metrick (2012)). Our analysis suggests that there is also a benefit to issuing assets with high correlation: it minimizes information costs in the economy as information about one asset can be applied to other assets as well. Thus even though adverse selection problems for correlated assets were observed during the recent crisis, this does not necessarily indicate inefficiency, as issuing correlated assets can be optimal under certain circumstances. Similarly to asset correlation, opacity is also often associated with inefficiencies. Commentators have been especially alarmed by the opacity of complex securitization structures. Since opacity also comes with benefits – it may discourage people from obtaining private information – one has to be careful in deriving direct implications from this.

6 Conclusion

How does opacity affect liquidity when investors can acquire information about an asset? This paper has suggested that the link between the two is non-monotonic. Both very transparent and very opaque assets preserve commonality of information. While full transparency directly precludes information asymmetries, sufficiently large opacity deters acquisition of private information by making learning about an asset more costly. Assets with either very low or very high opacity can hence be expected to be liquid. Assets which display intermediate degrees of opacity, in contrast, are prone to information acquisition. These assets may suffer from adverse selection problems when they need to be traded. An empirical analysis of the cross-section of listed U.S. firms strongly supported this hump-shape relationship between opacity and illiquidity.

Our analysis points to a significant benefit to opacity, which may help understand the phenomenon that issuers often choose to sell surprisingly opaque assets, as for instance observed in the case of securitization products. Policy makers thus have to be careful in equating opacity with inefficiencies. The results also have implications for transparency regulation. In particular, our analysis suggests that uniform transparency requirements are not desirable. This is simply because they may increase adverse selection for the more opaque assets in the economy. Rather, a more appropriate policy is to subsidize the provision of information by issuers. This can help internalizing the externalities associated with opacity, while allowing issuers to optimally preserve heterogeneous transparency levels.

References

Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability*, 5(3):224–255.

Agarwal, P. (2007). Institutional ownership and stock liquidity. *working paper*.

Boot, A. W. and Thakor, A. V. (2001). The many faces of information disclosure. *Review of Financial Studies*, 14(4):1021–1057.

Carlin, B., Kogan, S., and Lowery, R. (2013). Trading complex assets. *Journal of Finance*, 68(1):1937–1960.

Chordia, T., Roll, R., and Subrahmanyam, A. (2000). Commonality in liquidity. *Journal of Financial Economics*, 56(1):3–28.

Cleveland, W. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(1):829–836.

Dang, T., Gorton, G., and Holmström, B. (2013a). The information sensitivity of a security. *working paper*.

Dang, T., Gorton, G., Holmström, B., and Ordoñez, G. (2013b). Banks as secret keepers. *working paper*.

Farhi, E. and Tirole, J. (2014). Liquid bundles. *Journal of Economic Theory*.

Flannery, M., Kwan, S., and Nimalendran, M. (2004). Market evidence on the opaqueness of banking firms' assets. *Journal of Financial Economics*, 71(3):419–460.

Glosten, L. and Milgrom, P. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100.

Gorton, G. and Metrick, A. (2012). Securitization. *working paper*.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, pages 393–408.

Huang, R. D. and Stoll, H. R. (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4):995–1034.

Kaplan, T. (2006). Why banks should keep secrets. *Economic Theory*, 27(2):341–357.

Kurlat, P. D. and Veldkamp, L. (2012). Should we regulate financial information? *working paper*.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, pages 1315–1335.

Monnet, C. and Quintin, E. (2015). Rational opacity in private equity markets. *working paper*.

Morgan, D. (2002). Rating banks: Risk and uncertainty in an opaque industry. *American Economic Review*, 92(4):874–888.

Pagano, M. and Volpin, P. (2012). Securitization, transparency, and liquidity. *Review of Financial Studies*, 25(8):2417–2453.

Roulstone, D. T. (2003). Analyst following and market liquidity. *Contemporary Accounting Research*, 20(3):552–578.

Sato, Y. (2014). Opacity in financial markets. *Review of Financial Studies*, 56(1):3–28.

Stoll, H. R. (1989). Inferring the components of the bid-ask spread: Theory and empirical tests. *The Journal of Finance*, 44(1):115–134.

Wagner, W. (2011). Systemic liquidation risk and the diversity diversification trade-off. *Journal of Finance*, 66(4):1141–1175.

Yang, M. (2012). Optimality of securitized debt with endogenous and flexible information acquisition. *working paper*.

Appendix

A Random Discovery of the Payoff Interval

This analyzes a stochastic information acquisition technology. While in the baseline model information acquisition started at the upper end of the interval $[l, o]$, we now consider a random starting point. Specifically, we denote the *starting state* for information acquisition with y and assume that it is uniformly distributed on $[l, o]$. The distribution of the starting state is known by the investor, but not its realization.

As before, the investor learns about an interval of mass a when choosing a level of information acquisition a. For given starting state y, the investor thus learns about the interval $[y-a, y])$. If a is such that $y-a > l$, she learns that the interval $[y-a, y]$ does not contain payoff states, as in the baseline model. If a is sufficiently large such that $y-a \leq l$, she "discovers" the payoff interval. In this case, she ends up with complete knowledge about the distribution of payoff states.

We allow information acquisition to take place sequentially, that is, the investor can first decide to obtain information about a certain mass of states, and following this decide whether to analyze more states (and so on). Note that since the investor does not know the realization of y, she does not know in advance whether a certain amount of information acquisition will lead to discovery of the payoff interval.

It is easy to see that the modification in the information technology does not alter the investor's incentives to sell to the market at date 1 (Lemma 1): she will offer the asset if impatient; otherwise she will offer the asset only if she knows that the asset is worthless. The price of the asset will again depend on the market's belief about information acquisition. These beliefs, however, are no longer necessarily characterized by a single parameter since information acquisition can become stochastic (for instance, depending on y, investor may discover the payoff interval early on and stop). Let us denote the market price with $\widetilde{p}$ to indicate its dependence on beliefs.

We start with the analysis of the investor's incentives to acquire information. When deciding about information, the investor takes as given the price $\widetilde{p}$ at which she can sell to the market. We consider information acquisition that takes place by acquiring knowledge about (small) intervals of size $b > 0$ (we later consider the limit of b tending to zero).

Consider first that the investor has already discovered the payoff interval. She then has complete information about the asset, and hence will not acquire any further information.

Consider next the decision of an investor to acquire information about an interval b given that she has already acquired an amount $a \geq 0$ of information and has not yet discovered the payoff interval. Two cases arise. First, if a is sufficiently large such that $o - a - b \leq l$, the investor knows that the payoff interval will be discovered with certainty with the next information acquisition. The discovery will benefit the investor when a state of nature s materializes that falls in the interval $[l, o - a]$ and when she is impatient. The probability of this is $(o - a - l)\pi$, in which case she is able to sell at price $\widetilde{p}$ rather than holding onto a worthless asset. Her expected gains from additional information acquisition are thus

$$u(l, \widetilde{p}) - u(a, \widetilde{p}) = (o - a - l)\pi\widetilde{p} - bk_I. \tag{12}$$

These gains are identical to equation (3) in the baseline model - except that an interval of size $o-a-l$ is discovered by incurring costs for b ($\geq o-a-l$) states. Equation (12) shows that information acquisition is beneficial whenever $(o-a-l)\pi\widetilde{p} > bk_I$. We can hence define the *option value* of information acquisition in this case as $\max\{(o - a - l)\pi\widetilde{p} - bk_I, 0\}$.

Second, we have the case of $o - a - b > l$. In this case the investor does not know whether the next information acquisition will discover the payoff interval – it depends on the starting state y. While the realization of y is unknown to the investor, she infers from not having discovered the payoff interval up to now that $y \in [l + a, o]$. The impact of information acquisition in this case is as follows. When $y > l + a + b$, she does not discover the payoff interval. In this case, she can rule out an interval of mass b as containing payoff-states. When $y \leq l + a + b$, she discovers the payoff interval. She then rules out in total a mass of $o-a-l$ states. The likelihood of non-discovery and discovery is $1-\frac{b}{o-a-l}$ and $\frac{b}{o-a-l}$, respectively. We hence have for the total expected mass of loss states discovered $b(1+\frac{o-l-a-b}{o-a-l})$. Recalling that the investor benefits from knowledge about loss states when impatient, we obtain for the total expected gains from acquiring information

$$u(a+b,\widetilde{p})-u(a,\widetilde{p})=b\left((1+\frac{o-l-a-b}{o-a-l})\pi\widetilde{p}-k_I\right)+\left(1-\frac{b}{o-a-l}\right)V(a+b), \quad (13)$$

where $V(a+b)$ is the option value from acquiring further information when the payoff interval has not been discovered.

The value of information acquisition can hence be recursively defined as

$$V(a)=\begin{cases}\max\{b\left((1+\frac{o-l-a-b}{o-a-l})\pi\widetilde{p}-k_I\right)+\left(1-\frac{b}{o-a-l}\right)V(a+b),0\} & \text{if} \quad a<o-l-b\\ \max\{(o-a-l)\pi\widetilde{p}-bk_I,0\} & \text{if} \quad a\in[o-l-b,o-l)\\ 0 & \text{if} \quad a\geq o-l\end{cases}. \quad (14)$$

Note that $f(a):=b\left((1+\frac{o-l-a-b}{o-a-l})\pi\widetilde{p}-k_I\right)$ is decreasing in a. This implies that the value of acquiring information about an interval of size b is declining in the amount of information already acquired. The reason is as follows. While the likelihood of discovering the payoff interval $(\frac{b}{o-a-l})$ is increasing in a, the expected gains conditional on discovery are decreasing. The latter is because the mass of states ruled out by discovery, $o-l-a-b$, falls in a. Because of this latter effect, the gains from information acquisition are ultimately decreasing.

It follows that $f(a)\leq 0$ implies $f(a+b)<0$. In addition, we can conclude that when $a\in[o-l-b,o-l)$ (that is, when the next information acquisition discovers the payoff interval with certainty) we have $f(a-b)>(o-a-l)\pi\widetilde{p}-bk_I$. From this it follows that whenever $f(a)\leq 0$, the option value of information acquisition beyond the next interval is zero ($V(a+b)=0$). Thus, $V(a)=0$ whenever $f(a)\leq 0$. The consequence is that an investor will acquire information as long as $f(a)>0$, and will stop when $f(a)\leq 0$ or when the payoff interval is discovered.

An equilibrium strategy for information acquisition is hence defined by a threshold $a^* \in (0, o-l)$ such that $f(a^*) \leq 0$, but $f(a^*+b) > 0$. For arbitrarily small intervals of information acquisition ($b \to 0$), we find that $f(a) = 0$ precisely when

$$\widetilde{p} = \frac{k_I}{2\pi}. \tag{15}$$

This condition is almost identical to the condition for an interior equilibrium in the baseline model ($\widetilde{p} = \frac{k_I}{\pi}$). The difference arises because information acquisition is now more effective as it can result in the discovery of the payoff interval, in which case the entire distribution becomes known (in the baseline model, it only allowed us to proportionally narrow down the set of payoff states). In order for the gains from information acquisition to be identical to the costs k_I, the price at which the asset can be sold when information is of use to the investor hence has to be lower.

We next derive the break-even market price $\widetilde{p}$ as a function of beliefs about information acquisition. Recall that the investor's strategy can be summarized by a threshold value a^*. The market's beliefs can hence be summarized by a single parameter $\widetilde{a}$. Note that even though information discovery is stochastic, it only has two possible outcomes: either the investor finds the payoff interval or she reaches $\widetilde{a}$ and stops. Given that the starting point y is distributed on $[l, o]$, the probability of the payoff interval being discovered is simply

$$\pi_0 = \frac{\widetilde{a}}{o-l}. \tag{16}$$

The investor will offer the asset if either she is impatient or if she is patient and privately knows the asset will not pay out. The probability of the latter is $\frac{o-l}{o}$ when she has discovered the payoff interval and $\frac{\widetilde{a}}{o}$ when she has not discovered the payoff interval. The total probability of offering is thus

$$1 - \pi + \pi\left(\pi_0 \frac{o-l}{o} + (1-\pi_0)\frac{\widetilde{a}}{o}\right). \tag{17}$$

An offered asset only has a positive expected value if the investor is impatient (occurring with probability $1-\pi$), in which case the expected value to the market is $\frac{l}{o}$. We can then use (16) and (17) to express the expected value (and hence the price) of the asset conditional on being offered as

$$p(\widetilde{a}, o) = \frac{1-\pi}{(1-\pi)o + \pi(\widetilde{a}(2 - \frac{\widetilde{a}}{o-l}))} l. \tag{18}$$

Combining (15) and (18) to eliminate $p(\widetilde{a}, o)$, and solving for $a^* = \widetilde{a}$ yields:

$$a^* = (o-l) - \sqrt{(o-l)\left((o-l) - (1-\pi)(\frac{2l}{k_I} - \frac{o}{\pi})\right)}. \tag{19}$$

Differentiating with respect to o gives

$$\frac{\partial a^*}{\partial o} = 1 - \frac{(o-l)\left(2 - \frac{1-\pi}{\pi}\right) - (1-\pi)(\frac{2l}{k_I} - \frac{o}{\pi})}{2\sqrt{(o-l)\left((o-l) - (1-\pi)(\frac{2l}{k_I} - \frac{o}{\pi})\right)}} < 0. \tag{20}$$

Information acquisition (in an interior equilibrium) is hence declining in opacity o, as in the baseline model.

The cases of no and full information acquisition are straightforward to analyze. No information acquisition results if at $a = 0$ we have $f(a) \leq 0$. Noting that zero information acquisition implies $p = \frac{l}{o}$, we can obtain from $f(a) = 0$ a critical threshold opacity of $\bar{o} = 2\frac{l\pi}{k_I}$, such that an opacity level of $o \geq \bar{o}$ deters information acquisition. Full information acquisition arises when $f(o-l) \geq 0$ (as $b \to 0$, we can ignore the case of $a \in [o-l, o-l+b]$). Equation (18) yields for $a^* = o - l$ that $p = \frac{(1-\pi)l}{o-\pi l}$. Combining with $f(o-l) = 0$ and rearranging gives a critical threshold $\underline{o} = \pi l + \frac{(1-\pi)\pi l}{2k_I}$. For $o \leq \underline{o}$ we hence have a full information acquisition equilibrium.

We can summarize

Proposition A.1 *The equilibrium threshold for information acquisition $a^*(o)$ is given by*

$$a^*(o) = \begin{cases} o - l & \text{if} \quad o \leq \bar{o} \\ (o-l) - \sqrt{(o-l)\left((o-l) - (1-\pi)(\frac{2l}{k_I} - \frac{o}{\pi})\right)} & \text{if} \quad o \in (\underline{o}, \bar{o}) \\ 0 & \text{if} \quad o \geq \underline{o} \end{cases} \tag{21}$$

with $\underline{o} = \pi l + \frac{(1-\pi)\pi l}{2k_I}$ *and* $\bar{o} = 2\frac{\pi l}{k_I}$.

B Learning about Loss States

Assume that the asset pays 1 if nature selects a state $s \in [l, 1]$ and zero otherwise. Increasing transparency and information acquisition each narrow down the potential set of states where the asset does not pay off. In particular, for transparency choice o and information acquisition a, the public knows the set of non-paying (loss) states to be on the interval $[0, o]$, while the investor knows that the loss states are distributed on $[0, o - a]$.

For given beliefs about private information acquisition, $\widetilde{a}$, the trading decision of the investor is as follows. When $s \geq o$, both investor and market know that the asset will certainly pay . Its price will hence be 1. An impatient investor will sell the asset, while a patient investor will not sell given the assumptions we made about the investor's actions whenever indifferent. When $s \in [o - a, o]$, the investor knows that the asset will certainly pay off but the market only has imperfect knowledge about the payoff. The investor has thus positive private information about the asset. If she is patient, she will hence not sell. If impatient, the investor will still sell. Finally, when $s \in [0, o - a]$, both investor and market are uncertain about the payoff. However, the investor observes that s is not in her private set of payoff states $[o - a, o]$. She thus has negative private information. She will hence sell, regardless of whether she is patient or not (the expected value of the asset is $\frac{o-a-l}{o-a}$ in this case). The market price of the asset (conditional on $s < o$) can hence be derived as

$$p(\widetilde{a}, o) = \frac{o - l - \widetilde{a}\pi}{o - \widetilde{a}\pi}. \tag{22}$$

Note that for $\widetilde{a} = 0$, this simplifies to $p = \frac{o-l}{o}$, which is the expected value of the asset conditional on $s < o$. Note also that $\frac{\partial p(\widetilde{a},o)}{\partial \widetilde{a}} < 0$, because of adverse selection.

Similar to equation (2) we can derive the investor's utility given market beliefs $\widetilde{a}$:

$$u(a, \widetilde{a}) = 1 - o + o\left(\frac{a}{o}(\pi + (1-\pi)p(\widetilde{a}, o) + \frac{o-a}{o}p(\widetilde{a}, o)\right) - k_I \cdot a. \tag{23}$$

The derivative with respect to a is

$$\frac{\partial u(a, \widetilde{a})}{\partial a} = \pi(1 - p(\widetilde{a}, o)) - k_I. \tag{24}$$

It is useful to contrast this with the marginal benefit of information acquisition in the baseline model ($\frac{\partial u(a,\widetilde{a})}{\partial a} = \pi p(\widetilde{a}, o) - k_I$). Private information benefits the investor whenever it causes her to modify her selling decision. In the baseline model, the investor learns that the asset will *not* pay off in certain states. A patient investor will then sell the asset if such a state materializes; and hence benefits from a higher market price. In the extension considered here, the investor learns about states in which the asset *does* pay off. She thus does not sell the asset if such a state materializes. Her gains hence decline in the market price (which she would otherwise obtain by selling the asset). This has a consequence: because more information acquisition leads to lower prices in equilibrium, the gains from information will now be increasing in the amount of information acquired (formally, we have that $\left.\frac{\partial u(a,\widetilde{a})}{\partial a}\right|_{a=\widetilde{a}}$ is increasing in a).

Two cases arise. Consider first that $\left.\frac{\partial u(a,\widetilde{a})}{\partial a}\right|_{a=\widetilde{a}} > 0$ at $a = 0$. This implies that at a conjectured equilibrium with no information acquisition, the marginal gains from information acquisition are positive. Since we know that $\left.\frac{\partial u(a,\widetilde{a})}{\partial a}\right|_{a=\widetilde{a}}$ is increasing in a, the marginal gains from information acquisition are hence also positive for any $a > 0$. The unique equilibrium is hence full information acquisition: $a^* = o - l$. Consider next that $\frac{\partial u(a,\widetilde{a})}{\partial a}\Big|_{a=\widetilde{a}} < 0$ at $a = 0$. In this case, the gains from information acquisition at an equilibrium with no information acquisition are negative. Hence, no information acquisition is an equilibrium ($a^* = 0$). Since $\frac{\partial u(a,\widetilde{a})}{\partial a}\Big|_{a=\widetilde{a}}$ is increasing in a, there might also be a second equilibrium with positive information acquisition. However, this equilibrium would be pareto-dominated by no information acquisition (which involves no information cost) and we hence rule it out.

Whether full or no opacity is chosen thus depends on the sign of $\frac{\partial u(a,\widetilde{a})}{\partial a}\Big|_{a=\widetilde{a}}$. Using equation (24) one can find that this derivative is zero when $o = \frac{\pi l}{k_I}$. We can hence state

Proposition B.1 *When the investor learns about loss states, the equilibrium level of information acquisition a^* is*

$$a^*(o) = \begin{cases} o - l & \text{if} \quad o < \bar{o} \\ 0 & \text{if} \quad o \geq \bar{o} \end{cases} \tag{25}$$

with $\bar{o} = \frac{\pi l}{k_I}$.

C Increasing Cost of Information Acquisition

To analyze increasing costs of information acquisition, let the total cost of acquiring information about a mass of a states be $K_I(a)$ with $K_I(0) = 0$, $K_I'(a) > 0$, $K_I''(a) > 0$.

Differentiating the investor's utility $u(a, \widetilde{a})$ (equation (2), after replacing $k_I \cdot a$ with the new information cost function) with respect to a yields

$$\frac{\partial u(a,\widetilde{a})}{\partial a} = \pi p(\widetilde{a}, o) - K_I'(a). \tag{26}$$

Equation (26) determines a new threshold for zero information acquisition $\bar{o}$. Rearranging $\pi p(0, o) - K_I'(0) = \pi \frac{l}{o} - K_I'(0) = 0$ gives $\bar{o} = \frac{\pi l}{K_I'(0)}$. Since $K_I'(0) > 0$, there is hence a unique $\bar{o}$ above which no information is acquired. Likewise, $\underline{o}$ is uniquely pinned down by the condition: $\pi p(o-l,o)-K_I'(o-l)=\pi\frac{1-\pi}{1-\pi\frac{l}{o}}\frac{l}{o}-K_I'(o-l)=0$. This yields $\underline{o}=\pi l+\frac{(1-\pi)\pi l}{K_I'(o-l)}$. Finally, we can write down the condition for the interior equilibrium: $\pi p(a^*,o)-K_I'(a^*)=\frac{\pi(1-\pi)l}{o-\pi(o-a^*)}-K_I'(a^*)=0$. Totally differentiating with respect to o and rearranging gives

$$a^{*\prime}(o)=-\frac{(1-\pi)K_I'(a^*(o))}{\pi K_I'(a^*(o))+((1-\pi)o+\pi a^*(o)))K_I''(a^*(o))}<0. \tag{27}$$

Thus, opacity reduces information acquisition in an interior equilibrium. We hence have the same properties as in the baseline model. For $o\leq\underline{o}$ we have full information acquisition ($a^*=o-l$). Between $\underline{o}$ and $\bar{o}$ there is an interior degree of information acquisition which is declining in opacity. For opacity larger than $\bar{o}$, no information is acquired.

D Adverse Selection Costs

Modify the baseline model by assuming that the utility of the impatient investor is

$$U^I=C_0^I+qC_1^I,\text{ with }q\geq 1. \tag{28}$$

This modification does not affect trading with the market: an impatient investor will always sell while the patient investor only sells when she know the asset is worthless. Consider next the investor's incentives to acquire information. Similar to equation (2), utility is now

$$u(a,\widetilde{a})=o\left(((1-\pi)q+\pi\frac{a}{o})p(\widetilde{a},o)+\pi\frac{o-a}{o}\frac{l}{o-a}\right)-k_I\cdot a. \tag{29}$$

The derivative with respect to a is $\pi p(\widetilde{a},o)-k_I$ – the same as in the baseline model (equation (3)). The incentives to acquire information are hence unchanged and Proposition 1 still applies. The reason is that information acquisition only benefits the investor if she turns out to be patient, thus the fact that q may be larger than one does not matter.

The expression for welfare is now as follows. Whenever the investor is impatient and sells to the market, there is an additional welfare gain of $(q-1)p(a^*(o), o)$ compared to the baseline model. We thus have for welfare that

$$W(o, a^*(o)) = w^I + w^M + l + \pi(q-1)p(a^*(o), o) - k_O \cdot (1-o) - k_I \cdot a^*(o). \tag{30}$$

Opacity now has a new effect, arising because it can affect the price p through a change in information acquisition. Since p is declining in information acquisition (equation (1)), opacity-induced increases in information acquisition now have two effects. First, they directly lead to costs k_I. Second, they reduce the gains for the impatient investor by lowering the price at which she can sell to the market (when $q > 1$, these losses are not completely offset by gains for the market).

Optimal opacity is determined analogous to the baseline model. For $1 \leq \bar{o}$, full opacity maximizes welfare. For $1 > \bar{o}$, one needs to compare welfare under full and no opacity (because of the dependence on p, welfare is now non-linear in the region $[l, \bar{o}]$, but this does not affect the optimal decision). Full opacity is optimal if and only if $W(1, a^*(1)) > W(l, 0)$, which is when $\pi(q-1)p(a^*(1), 1) - k_I \cdot a^*(1) > \pi(q-1)p(0, o) - k_O \cdot (1-l)$. Otherwise full transparency should be chosen.

E Endogenous Discounting of Opaque Assets

We present a model in which low pricing of opaque assets does not follow from the assumed information acquisition technology (in the baseline model, a state in the public domain is less likely to pay-out when opacity is higher, resulting in a lower market price). Rather, the market values opaque assets less because it does not allow him to make efficient scrapping decisions.

We modify the baseline model as follows. First, we assume a distribution of asset returns and a type of information acquisition for which the expected value of an asset in the public domain is independent of opacity. As in the baseline model, the asset's pay-off states are randomly distributed. We assume that each state has the same likelihood of paying out, given by l. An asset is said to have opacity o when the returns on $(o, 1]$ are publicly known (there is thus a mass $(1-o)l$ of known pay-off states and mass $(1-o)(1-l)$ of known non-payoff states) but the returns in the states $[0, o]$ are not. An amount of information acquisition a results in the investor learning the returns in the interval $[o-a, o]$. She will thus know about a mass $a \cdot l$ of pay-off states and a mass $a \cdot (1-l)$ of states where the asset does not pay off.

Second, we assume that at $t = 1$, after potential information acquisition by the investor, a public signal s' about the state of the world arrives (in the baseline model, the state itself was revealed at this stage). The signal is imperfect in that only with probability σ ($\sigma \in (0, 1)$) it points to the actual state of the world ($s' = s$). With probability $1-\sigma$, the signal is wrong, in which case the distribution of states is identical to the ex-ante distribution (each state of the world $s \in [0, 1]$ is then equally likely). At $t = 1.5$, the state of the world is publicly revealed. Following the revelation of state, a market agent who has bought the asset can decide to scrap the asset. Scrapping the asset yields a sure pay-off of γ ($< l$). We assume in addition that $\gamma < \frac{1-\sigma}{\sigma}$.

All other actions are as in the baseline model. Figure 8 summarizes the timing of the modified model.

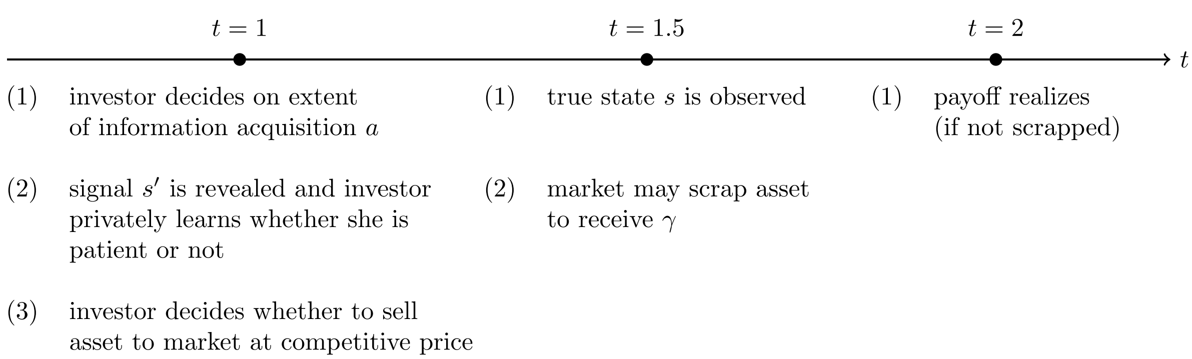

Figure 8: Timeline of Modified Model

The model is significantly more complicated than the baseline model in that multiple equilibria exist for a range of opacity levels and that small perturbations in the parameters can lead to discrete changes in equilibrium variables. In particular, there is no intuitive characterization of the behaviour of the equilibria with positive amounts of information acquisition. We hence focus the analysis on showing that for sufficiently high opacity, no information is acquired. By contrast, for assets of low or intermediate opacity, equilibria without private information acquisition are unsustainable.

We prove this by analyzing the investor's incentives to deviate from a conjectured equilibrium without information acquisition. We first derive the valuation of an asset sold by the investor presuming that the market believes no information has been acquired. Following this we solve for the investor's gains from deviating by acquiring information.

E.1 Market Price without Information Acquisition

Suppose that the market believes that no information has been acquired by the investor ($\widetilde{a} = 0$). In this case information asymmetry is absent and the selling decision of the investor does not provide any signal to the market.

Consider the scrapping decision of the market at $t = 1.5$. At this stage the market knows the initial signal s' as well as the actual state s. Optimal scrapping is as follows. If s is in the transparent region ($s \in (o, 1]$), the market knows whether the asset will pay off at $t = 2$. The market will thus scrap if and only if the asset does not pay off. If s is in the opaque region ($s \in [0, o]$), the market is not certain whether the asset will pay-off at date 2. Since the likelihood of pay-off is l, it is optimal not to scrap since $\gamma < l$.

Similar to the baseline model, we derive the price conditional on the signal being in the opaque region ($s' \in [0, o]$). Given the efficient scrapping decision, the expected value of an asset to the market, and hence the price, is given by

$$p(o, \widetilde{a} = 0) = \sigma l + (1 - \sigma)[ol + (1 - o)(l + (1 - l)\gamma)]. \tag{31}$$

This is derived as follows. With probability σ, the signal is true ($s = s'$). In this case, the actual state of the world is in the opaque region. The market then does not scrap and receives l in expectation. With probability $1-\sigma$, the signal is not true. In this case, the state falls in the opaque region with probability o. The market then again does not scrap and receives an expected value of l. With probability $1 - o$, the state falls in the transparent region and the pay-off is known to the investor. If it is a pay-off state (occurring with probability l), the investor receives 1. If it is not a pay-off state (occurring with probability $1 - l$), the investor scraps and receives γ.

Note that $p(o, \widetilde{a} = 0)$ simplifies to $l + (1 - \sigma)(1 - o)(1 - l)\gamma$. We hence have $p > l$ due to the option value from scrapping. Second, more opaque assets (high o) command lower prices. This is because for such assets there is a greater chance that a state realizes for which the investor does not know the date-2 payoff of the asset, in which case she cannot make efficient scrapping decisions.

E.2 Information Acquisition

We next analyze the investor's incentives to deviate by choosing a positive amount of information $a > 0$.

Consider the selling decision of an investor who has acquired information a. Information acquisition is only relevant in cases where the signal falls in the opaque region ($s' \in [0, o)$), so we can restrict attention to this region. Note that if the investor sells, she obtains $p(o, \widetilde{a} = 0)$ as the market believes there is no information acquisition.

If the investor is impatient, it is strictly dominating to sell to the market at $p(o, \widetilde{a} = 0)$. If the investor is patient, there are three cases to consider:

1. The signal s' falls in the private domain ($s' \in [o - a, o]$) and s' is not a pay-off state. The expected pay-off for the investor is then $(1-\sigma)l$, which is smaller than $p(o, \widetilde{a} = 0)$. It is hence optimal to sell.

2. The signal s' falls in the private domain ($s' \in [o - a, o]$) and s' is a pay-off state. The expected payoff from holding on to the asset is then given by $\sigma \cdot 1 + (1 - \sigma)l$. This is

less than $p(o, \widetilde{a} = 0)$ by our assumption that $\gamma < \frac{1-\sigma}{\sigma}$. It is hence optimal not to sell the asset.

3. The signal s' falls outside the private domain ($s' \in [0, o - a)$). The investor then does not know whether the asset pays off if the signal turns out to be true. Her expected payoff from holding on to the asset is then simply l. Since $l < p(o, \widetilde{a} = 0)$, she strictly prefers to sell.

The selling decision of the patient investor is thus similar to the baseline model: she continues to hold the asset unless she has negative private information about the realization of the asset at date 2.

It follows that the investor benefits from information acquisition *if and only if* she has negative private information (Case 1). The probability of this is $a \cdot l$. She benefits by receiving $p(o, \widetilde{a} = 0)$ from selling the asset, instead of the getting $(1 - \sigma)l$ in expectation from holding onto the asset. The gains from acquiring information are thus given by $a(p(o, \widetilde{a} = 0) - (1 - \sigma)l) - k_I a$. Simplifying gives:

$$u(a, \widetilde{a} = 0) - u(0, \widetilde{a} = 0) = a \cdot [\sigma l + (1 - \sigma)(1 - o)(1 - l)\gamma] - k_I a. \tag{32}$$

The marginal benefits from a unit of information are hence

$$\frac{\partial u(a, \widetilde{a} = 0)}{\partial a} = \sigma l + (1 - \sigma)(1 - o)(1 - l)\gamma - k_I, \tag{33}$$

and are constant. It follows that an equilibrium with no information acquisition can only arise when $\frac{\partial u(a, \widetilde{a}=0)}{\partial a} \leq 0$. Rearranging, this gives a critical opacity level $\bar{o}$:

$$\bar{o} = 1 - \frac{k_I - \sigma l}{(1 - \sigma)(1 - l)\gamma}, \tag{34}$$

such that for $o \geq \bar{o}$ no information is acquired. By contrast, for $o < \bar{o}$ investors always have incentives to acquire information when the market believes that no information has been acquired. In this case there is no equilibrium without information; in any equilibrium there hence has to be some level of information acquisition.

F Empirical Analysis

Table 1: Summary statistics

Variable	Description	Time Window	Mean	Std. Dev.	Min.	Max.
Spread	Average spread	Oct'13-Dec'13	0.025	0.03	0.01	0.244
Market capitalization	Market capitalization (log)	Oct'13	14.656	1.494	10.67	19.876
Share price	Share price	Oct'13	41.339	40.054	5.020	672.78
Volume	Avg monthly trading volume scaled by market cap	Oct'11-Sep'13	0.116	0.212	0.002	3.835
St.Dev. Volume	St.Dev. of daily returns	Oct'11-Sep'13	0.025	0.032	0.008	0.622
Volatility	St.Dev. of scaled monthly trading volume	Oct'11-Sep'13	0.538	1.153	0.008	19.586
Number of analysts	Average number of analyst forecasts	Oct'11-Sep'13	10.994	7.229	2	53.167
Dispersion	Average standard deviation of analyst forecasts	Oct'11-Sep'13	0.117	0.217	0.007	1.651

Table 2: Control factors in the determination of bis-ask spreads

This table summarizes results from an ordinary least squares regression where the bid-ask spread is the dependent variable. We report coefficients for the independent variables logged market capitalization, share price, average trading volume scaled by market capitalization, the standard deviation of past daily returns, the standard deviation of scaled past trading volume, and the number of analysts giving estimates.

Coefficient	Spread
Market capitalization	-0.00961***
	(0.000605)
Share price	0.000417***
	(0.000026)
Volume	-0.00960
	(0.00645)
St.Dev. volume	-0.00161
	(0.00122)
Volatility	0.11520***
	(0.04058)
Number of analysts	-0.000263***
	(0.00008)
Constant	0.15023***
	(0.00874)
Observations	2,067
R-squared	0.344

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

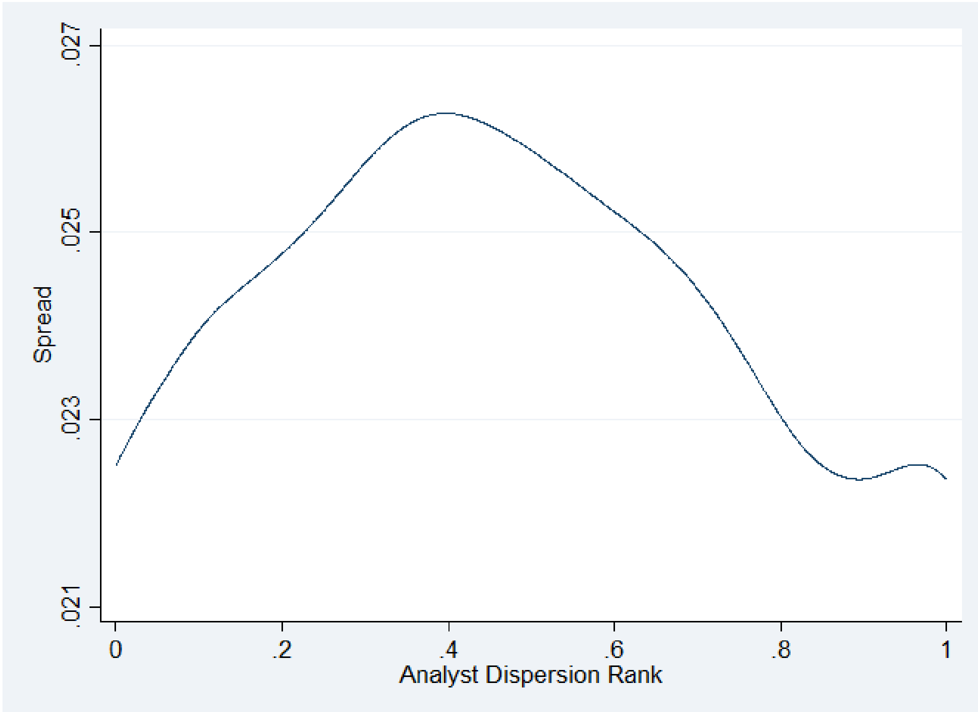

Figure 9: Lowess Regression for Bid-Ask Spread

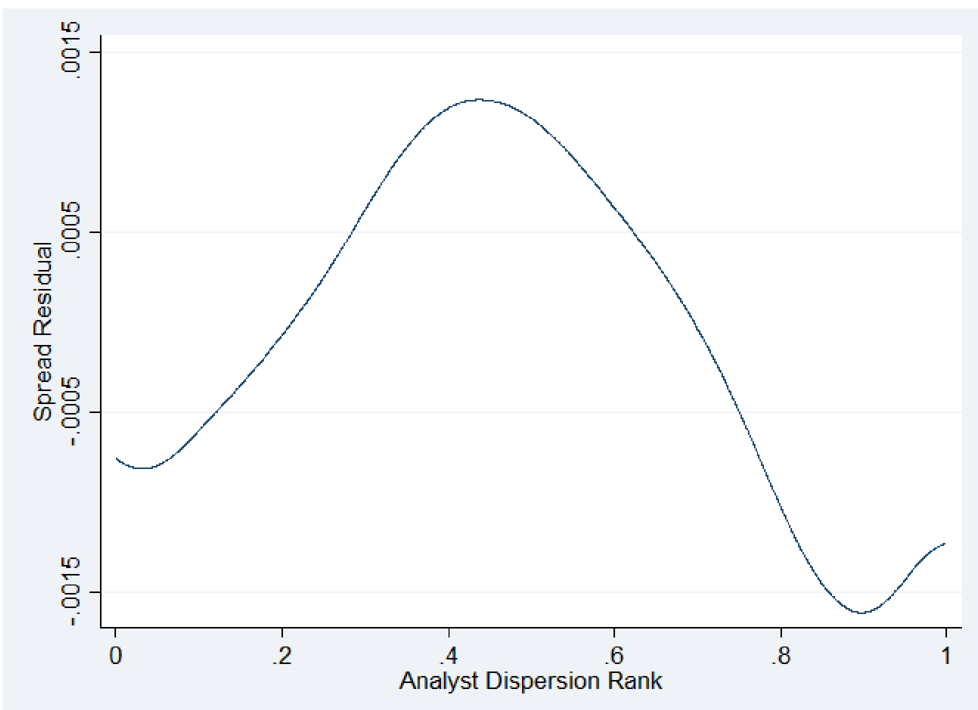

Figure 10: Lowess Regression of Bid-Ask Spread Residual