Limited Disclosure and Hidden Orders in Asset Markets

Cyril Monnet*
University of Bern and SZ Gerzensee

Erwan Quintin
Wisconsin School of Business

May 9, 2016

Abstract

Opacity assumes at least two prominent forms in asset markets. Dark exchanges and OTC markets enable expert investors to hide their orders while originators carefully control the disclosure of fundamental information about the assets they source. We describe a simple model where both forms of opacity – hidden orders and limited disclosure – complement one another. Costly investor expertise gives originators incentives to deliver assets of good quality. Keeping expert orders hidden generates the rents investors need to justify investing in expertise in the first place. Limiting disclosure mitigates the resulting adverse selection issues. Originators prefer to restrict the information they can convey to experts because it encourages the participation of non-experts. This optimal organization of asset markets can be decentralized using standard financial arrangements.

Keywords: market design, opacity, asymmetric information

JEL codes: D47; D82

*We wish to thank Morris Davis, Todd Keister, Vincent Maurin, Sangeeta Pratap, and Mark Ready for useful comments. We also thank seminar and session participants at the 2014 Midwest Macro meetings, the 2014 SED meeting, the University of Wisconsin and the Federal Reserve Bank of Dallas for their input.
Corresponding author: Cyril Monnet, Study Center Gerzensee, Dorfstr. 2, CH-3115 Gerzensee, Switzerland. cyril.monnet@szgerzensee.ch.

1. Introduction

Some financial markets are highly opaque. Dark exchanges and most over-the-counter (OTC) markets enable institutional investors to keep the details of their orders hidden from other investors, a form of pre-trade opacity emphasized by Bolton, Santos, and Scheinkman (2016). In addition and as discussed for instance by Pagano and Volpin (2012) in the context of asset-backed securitization, originators and underwriting intermediaries often withhold fundamental information from all investors, whether institutional or retail.

In this paper we make the case that the co-existence of hidden orders and limited disclosure is essential as long as expertise is necessary in financial markets. The idea is simple. Expertise is costly for investors to acquire but it is necessary to give originators incentives to deliver quality assets. Transparent venues would allow non-experts to free-ride on the investment of experts. The ability to hide orders allows experts to recoup their cost.[1] But non-expert investors may then have cold feet, as they legitimately fear that experts will exploit their ignorance, thus limiting the amount of liquidity at origination. Originators – or a financial intermediary representing them – can encourage the participation of non-experts by curtailing the flow of fundamental information to level the playing field between investors.

In our model, risk-neutral agents – originators – generate productive assets at a cost. Assets are of heterogeneous quality drawn from a known distribution. Investors are endowed with funds which they can either store safely or invest in the risky assets originators create. Investors can choose to become expert at a cost which means that they can understand the fundamental asset information originators convey to them. When originator effort is unobservable, the presence of expert investors rewards originators for producing quality assets.

[1] This aspect of our model is closely related to Grossman and Stiglitz (1980). They argue that when it is costly for investors to become informed markets cannot be informationally efficient since otherwise returns to information acquisition would be zero.

In such a context, limiting the disclosure of fundamental information to all investors can be strictly welfare-enhancing. Intuitively, non-expert investors are reluctant to participate in asset markets because originators with good projects prefer to deal with experts who can recognize the quality of their projects. Hiding fundamental information mitigates these adverse selection issues and leads to additional, welfare-improving asset origination.

Implementing the optimal informational arrangement via regulations would be difficult both on practical and legal grounds. But we show that contracts that are standard in financial markets can implement the desired information design. First, originators can sell their project forward before asset quality is revealed. Second, originators can hire an underwriter and design their compensation so that they will opt for the optimal disclosure policy. Third, originators can enter into a blank-check underwriting arrangement with investors.

The set of results we develop in this paper builds on a large literature devoted to the optimal level of information in asset markets. Ever since the seminal work of Hirshleifer (1971, e.g.), it is well understood that in second best environments the optimal level of disclosure is typically not full. This idea has been revived in recent work by Dang, Gorton, Holmstrom, and Ordonez (2013), Andolfatto, Berentsen, and Waller (2014), Monnet and Quintin (2013), Fishman and Hagerty (2003), and Pagano and Volpin (2012) among many others.

Both Dang, Gorton, Holmstrom, and Ordonez (2013) and Monnet and Quintin (2013) argue that limited disclosure preserves the liquidity of risky investments. Dang, Gorton, Holmstrom, and Ordonez (2013) implement the optimal disclosure level by delegating investments to an intermediary they interpret as a bank whereas Monnet and Quintin (2013) focus on the impact of information releases on secondary market liquidity.

By focusing on adverse selection issues this paper is more closely related to Pagano and Volpin (2012). They show that in a world with differently sophisticated investors, partial disclosure can serve to mitigate the winner's curse that arises when less knowledgeable investors must bid for assets alongside expert investors with a superior ability to interpret fundamental information.[2] The same idea is at the core of this paper but our goal is to fully characterize the optimal disclosure design in primary markets with differently skilled investors. Whereas Pagano and Volpin (2012) study the information-control problem from the point of view of a given issuer, discussing the social role of opacity as we do requires a model where the volume of origination is fully endogenous.[3] Among other key insights, such a model reveals that primary asset markets should optimally feature a juxtaposition of trading platforms with different ratio of experts to non-expert investors and different levels of disclosure.

Bolton, Santos, and Scheinkman (2016) also study the social value of costly information acquisition in a model with experts and boeotian investors. They interpret expert-only exchanges as OTC markets and use their framework to ask whether we should expect suboptimal investments in expertise in asset markets. Just like in our model, absent moral hazard the presence of experts leads to "cream-skimming" and inefficient origination levels but introducing moral hazard creates a socially useful role for costly information acquisition. These commonalities notwithstanding, our focus and main message are quite different from theirs. Their main point is that opacity in asset markets generically lead to excessive expert rents and suboptimal origination volumes. We are arguing essentially the contrary: opacity, carefully designed, can mitigate the negative welfare impact caused by adverse selection is-

[2] The idea that optimal disclosure is not full when some buyers cannot process fundamental information is also in Fishman and Hagerty (2003). They take the size of the market and the fraction of sophisticated buyers as exogenous whereas we ask whether limiting disclosure can help support additional, welfare-enhancing origination.

[3] Pagano and Volpin (2012) also discuss the socially optimal level of disclosure but one issuer at a time and from a completely different point of view. They study the trade-off between liquidity in primary markets and liquidity in secondary markets in a version of their model where sophisticated investors can invest in information acquisition after the primary trading stage is complete. They show that lack of transparency in primary markets can exacerbate adverse selection issues in secondary markets. If those secondary liquidity issues are associated with a greater risk of dead-weight losses caused e.g. by fire sales, issuers do no select the socially optimal level of transparency. We focus on the amount of primary disclosure that leads to the socially optimal level of productive origination.

sues in markets where experts play an essential role. Expert-only exchanges play a similar role in our model as in theirs – they enable experts to hide their orders hence to generate rents. But we argue that the resulting adverse selection issues should be mitigated by optimally controlling the disclosure of fundamental asset information to all investors. Rents always adjust in our model to exactly offset expertise acquisition cost and opacity, optimally designed, actually serves to make aggregate rents as small a fraction of origination volumes as possible in decentralized markets.[4]

Kurlat (2015) describes a model where informed and non-informed agents coexist and, like us, describes the trading arrangements that should emerge given this juxtaposition. His focus however is on trading and rationing protocols whereas we require that markets clear in a model where agents take other agents' willingness to pay for assets as given. We focus on the transparency characteristics of the trading venues that emerge in the resulting market environment with a fully endogenous fraction of experts and non-experts.

Our theory naturally applies to markets for initial-public-offerings (IPOs). There institutional investors usually receive detailed communications about target assets from underwriting intermediaries before formulating their bids whereas retail investors are often barred from these pre-IPO exchanges and must rely mostly on public disclosures. In fact, a significant share of retail investors who choose to participate in primary markets rely on blank-check underwriting vehicles that raise funds before a specific target has been located. Beyond easily rationalizing this segmentation, our model makes a host of predictions that are consistent with several well documented empirical regularities in IPO markets. As the survey provided by Ljungvist (2007) explains, IPOs tend to be "underpriced" in the sense that the

[4]Bolton, Santos, and Scheinkman (2016)'s other primary objective is to propose a mechanism that could explain why expert rents/compensation in Finance keeps increasing even as the number of experts is rising because this corresponds to their reading of the evidence presented in Philippon and Resheff (2012). While this is not our primary concern in this paper, we will explain in section 4. that our model makes the same prediction when expertise-acquisition costs are heterogeneous across investors, without altering any of our other results. See, in particular, footnote 12. Other papers that focus on expertise rents in the financial industry include Philippon (2008) and Biais et al. (2010).

offering price is below the price at which shares subsequently trade in secondary markets. In addition, IPOs that feature more representation from institutional investors tend to be more underpriced, (see e.g. Michaely and Shaw, 1994, and Caseres and Lowry, 2009) and institutional investors earn higher returns on IPOs than do retail investors (see Aggarwal, Prabhala, and Puri, 2002). These predictions are all borne by our model as they are by most primary asset market models with adverse selection.

The typical model in the vast literature on optimal disclosure in IPO market[5] follows Rock (1986) and takes the number of skilled investors as given. In such an environment, IPO underpricing constitutes pure rents for investors that happen to be informed and typically results in underinvestment vis-a-vis the full information benchmark. Among other stark predictions this class of environments generates, making information available to everyone reduces underpricing, increases the volume of transactions, and is clearly optimal.

In our model we obtain the same prediction: making information cheaper to analyze for instance by establishing strict accounting standards reduces underpricing. But, in an environment where information is costly to process, imposing more disclosure than would be optimal can also increase average underpricing, decrease the volume of primary transaction and reduce welfare. This is because in our model expertise acquisition and participation by experts are both endogenous. Releasing more information without reducing the cost of processing it can exacerbate the adverse selection problem and reduce the willingness of non-experts to participate in asset markets. Hence, the ratio of experts to non-experts and average underpricing can increase.

The key distinction between our model (or the model of Pagano and Volpin, 2012) and the typical model used in optimal disclosure investigations is that being informed in our environment means being able to process and interpret fundamental information. Availability of information is only part of the problem if, as seems a natural assumption, non-expert

[5] See Kumar et al. (2013) for a detailed review.

investors are not able to process that information anyway. In agency mortgage-backed security (MBS) markets for instance, providing detailed information on hundreds of diverse mortgages to retail investors does not mean that they can refine their prepayment forecasts in the same way that a mortgage Real Estate Investment Trust (REIT) and its army of quantitative analysts can. This means that hiding some information is the only way to level the playing field between experts and non-experts.

The remainder of the paper is organized as follows. Section 2. describes the physical environment. Section 3. characterizes the solution of a benevolent social planner problem with and without moral hazard in origination. It shows, in particular, that the socially optimal fraction of experts is interior when moral hazard frictions are present. Section 4. considers a decentralized, Walrasian environment with full disclosure where originators can choose to sell either in expert-only markets or in a market open to all investors. This introduces the adverse selection issues emphasized by Bolton, Santos, and Scheinkman (2016) and as a result, decentralized markets cannot implement the second-best allocation. Section 5. shows that controlling the disclosure of fundamental information to all investors mitigates these issues and, therefore, strictly increases origination volumes and welfare. Section 6. explains how the resulting optimal disclosure design can be implemented using contracts that are standard in financial markets. Section 7. concludes.

2. The environment

Consider an environment with two dates $t = 0, 1$, a measure one of investors and a large mass of originators. All agents are risk neutral. Investors care for consumption at $t = 1$, while originators care for consumption at $t = 0$. Investors are endowed with one unit of the unique good at $t = 0$. They can individually invest all or part of that endowment in a risk-free storage technology with a zero net return.

Originators have no endowment but can each create/find a project provided they incur a utility cost at date 0. This utility cost is distributed among originators according to a distribution function H with a continuous derivative. That is, $H(k_0)$ denotes the mass of originators whose creation cost is below k_0.

Once a project is created it only becomes productive if the originator performs due diligence at an additional cost k. This cost is the same for all originators.[6] A diligent originator draws a probability $\lambda \in [0, 1]$ from a continuous cumulative distribution function F with continuous density f. Projects of type λ yield a payoff r with probability λ at date 1 and nothing otherwise. In other words λ measures the quality of a project. A law of large numbers holds: the distribution of quality is always F among diligent originators. A productive project does not require any investment beyond the originator's creation and diligence costs, k_0 and k.

Investors do not directly observe due diligence. But at date zero any investor can choose to suffer a utility cost $\gamma \geq 0$ to become an expert. Experts know whether or not a given originator was diligent and originators can convey to them the quality λ of a project once it is created. Investors who choose to remain non-experts cannot tell whether an originator was diligent and originators cannot communicate the quality of their project to them.

Originators cannot signal their project type by retaining ownership, as they only consume at date 0.[7] Therefore, they only engage in project creation and in due diligence if this enables them to consume part of the date -0 endowment.

We make two assumptions on the distribution F of project quality. First, we restrict parameters so that absent any frictions it is optimal (in a sense we make precise below) to invest a positive amount of resources in risky projects, that is we assume

[6]The heterogeneity in project creation cost k_0 produces a non-degenerate supply curve of projects. Much of this paper is about the socially optimal quantity of origination. We use the additional effort cost k to introduce moral hazard frictions in a tractable fashion. We could make that aspect heterogeneous as well at some notational cost without changing the nature of any of our results.

[7]This contrasts with models in the tradition of Leland and Pyle (1977).

Assumption 1. $\int \lambda r dF > k$.

When assumption 1 is violated projects are too unproductive in expectation to cover the due diligence cost. Not even the most talented originators – originators whose $k_0 = 0$ – would create one. Second and to simplify the analysis, we also assume that it is never optimal for all resources to be directed towards risky projects,

Assumption 2. $H\left(\int \lambda r dF - k\right) \int \lambda r dF < 1$.

As we will show, assumption 2 guarantees that the storage technology is used at any optimal solution and, a fortiori, at any market solution.

3. The social role of expertise

In this paper as in Bolton, Santos, and Scheinkman (2016) or Pagano and Volpin (2012), opacity only plays an essential role provided differently skilled investors populate asset markets. In this section, we consider the problem of a benevolent social planner who can directly allocate resources among agents. We do not think of this planner as a government or regulator in a literal sense but as a device to characterize socially optimal allocations.

The planner is benevolent in the sense that her goal is to maximize average hence total utility. In addition to the economy's fundamental characteristics, she knows whether a given originator created a project. But the planner does not know the cost type k_0 of a given originator and she cannot force investors to become experts so that any allocation has to satisfy standard participation constraints. As for the due diligence decision of originators, we will consider two benchmark cases.

3.1. When due diligence is public information

Assume first that due diligence choices by originators are public information. Denote by $c^n \geq 0$, $c^e \geq 0$ and $c^o \geq 0$ the consumption of non-experts (n), experts (e) and active originators (o), respectively, while s^n and s^e are the resources stored by non-experts and experts. Without loss of generality, we assume that only *active* originators consume positive amounts. Since agents are risk neutral, giving consumption to inactive originators is a pure transfer from investors and does not raise total surplus. In addition to consumption levels, the planner decides a threshold ability $\bar{k}_0$ below which originators are active, and the fraction of experts $\mu \in [0, 1]$ in the economy. Naturally, the allocation has to be feasible and satisfy incentive constraints.

Resource constraints impose that storage s^n and s^e are in $[0, 1]$ and limit the consumption of active originators at date 0:

$$\int_0^{\bar{k}_0} c^o dH(k_0) + (1-\mu)s^n + \mu s^e \leq 1, \tag{1}$$

while at date 1 we must have:

$$(1-\mu)c^n + \mu c^e \leq (1-\mu)s^n + \mu s^e + H(\bar{k}_0) \int \lambda r dF. \tag{2}$$

In order to implement the desired skill level $\bar{k}_0$, active originators must be compensated for the overall effort they put in. The associated participation constraint is

$$c^o \geq \bar{k}_0 + k. \tag{3}$$

Furthermore, experts and non-experts must receive at least 1 in expected utility in order to participate in the social plan. But the planner cannot force investors to become experts. Therefore, if there are experts, investors have to be indifferent between the two investor types,[8]

$$c^e - \gamma = c^n \geq 1. \tag{4}$$

The planner maximizes a social objective function that puts equal weight on all investors and originators,

$$SP1 = \max_{\bar{k}_0, c^n, c^e \geq 0, (s^n, s^e, \mu) \in [0,1]^3,} (1-\mu)c^n + \mu(c^e - \gamma) + \int_0^{\bar{k}_0} (c^o - k_0 - k) dH(k)$$

subject to (1)-(4).

When due diligence is public information the planner has no reason to spend resources on training experts. Hence, it is optimal to set $\mu = 0$. Substituting the resource constraints in the objective, we can write the planner's problem as maximizing date 1 consumption net of the required utility cost for originators:

$$SP1 = \max_{\bar{k}_0 \geq 0} 1 + H(\bar{k}_0) \int \lambda r dF - \int_0^{\bar{k}_0} (k_0 + k) dH(k_0)$$

The first order condition gives the optimal k_0^* which equates the marginal benefit of an additional project to its marginal cost,

$$\int \lambda r dF = k_0^* + k. \tag{5}$$

Assumptions 1 and 2 guarantee that this solution is compatible with all resource constraints.

[8]We could write an apparently more general set of participation conditions as follows:

$$\text{If } \mu > 0 \text{ then } c^e - \gamma \geq \max(c^n, 1) \text{ and if } \mu < 1 \text{ then } c^n \geq \max(c^e - \gamma, 1).$$

But there is no loss of generality in immediately requiring that investors be indifferent between the two possible types. Indeed, take any allocation that satisfies the more general set of constraints. Then the alternative consumption plan $(\tilde{c}^n, \tilde{c}^e)$ defined by $\tilde{c}^n = \tilde{c}^e - \gamma = max(c^n, c^e - \gamma)$ supports the same fraction μ of experts without violating any feasibility constraint. Indeed, the two conditions stated above imply that $c^e - \gamma = c^n$ unless $\mu = 0$ or $\mu = 1$ and in those two extreme cases one can raise either consumption level until $c^e - \gamma = c^n$ without affecting feasibility.

In summary,

Proposition 1. The first best allocation satisfies $\mu = 0$, $s^n = 1 - H(k_0^*)(k_0^* + k)$, $c^n = 1$, and $c^o = k + k_0^*$ where

$$\int \lambda r dF = k_0^* + k.$$

All investors should remain non-experts because expertise is costly and serves no purpose when originators' due diligence is public information. By assumption there are more than enough resources to support the optimal activation level of originators and, as k_0 is not observable, originators are able to extract all the surplus from the investment. Note that assumption 2 guarantees that there is positive storage at the first best, i.e. that $s^n > 0$.[9] This is also a property of all the allocations we study later in this paper. On the other hand, assumption 1 implies that a strictly positive mass of risky projects is activated at the first best so that not all resources are stored.

3.2. When due diligence is private information

Assume instead that the planner cannot observe originators' choice of due diligence. Expertise now plays a useful role in distinguishing bad projects (low λ) from good ones and it becomes worthwhile spending resources on training some investors. Then the planner can instruct experts to transfer $q^e(\lambda)$ to originators with projects of quality λ at date 0. Non-experts, for their part, cannot tell good from bad projects, hence the planner can only instruct these agents to transfer a uniform amount q^n to each active originators. Under the proposed allocation, the consumption of an originator whose project happens to be of quality

[9]Indeed, $k_0^* = \int \lambda r dF - k$ so that

$$s^n = 1 - H(k_0^*)(k_0^* + k) = 1 - H\left(\int \lambda r dF - k\right) \int \lambda r dF > 0.$$

λ is:

$$c^o(\lambda) = q^n + q^e(\lambda). \tag{6}$$

This set-up embodies two implicit assumptions. First, experts are physically separated from non-experts, as non-experts cannot observe and mimic the actions of experts. In other words, their transfers to a specific originator cannot depend on what experts do. Second, non-experts cannot entrust their endowment to experts. As is well known, the information problem would otherwise become trivial.[10]

Since originators need incentives to perform due diligence, it must be that active originators are better off exercising due diligence given c^o than just selling a bad project, or

$$\int c^o(\lambda)dF - k \geq c^o(0). \tag{7}$$

To support skill threshold $\bar{k}_0$, the expected payoff of originators must also satisfy

$$\int c^o(\lambda)dF \geq \bar{k}_0 + k. \tag{8}$$

Since it is clearly optimal to induce due diligence at the lowest possible cost, we can impose $q^e(0) = 0$. Feasibility requires that experts and non-experts have enough resources in order to pay q^n and $q^e(\lambda)$ to each active originator, i.e. we will also require

$$H(\bar{k}_0)q^n \leq (1-\mu)(1-s^n), \tag{9}$$

and

$$H(\bar{k}_0)\int q^e(\lambda)dF \leq \mu(1-s^e). \tag{10}$$

[10] See assumptions A2 and A3 in Rock, 1986, or p. 2426 in Pagano and Volpin, 2012, for a discussion. In our case, the planner would simply train a set of negligible mass of experts and have them make fully informed transfers on behalf of all agents. One specific friction that rules out this option is to assume that the planner cannot enforce ex-post transfers across investors, as in Koeppl et al. (2013) or Sappington (1983).

Constraints (9) and (10) formalize the assumption that experts cannot borrow resources from non-experts. It should be clear that they hold with equality at any optimum and, therefore the aggregate expected consumption of active originators at date 0 equals what investors do not save,

$$H(\bar{k}_0)\int c^o(\lambda)dF = H(\bar{k}_0)\int (q^n + q^e(\lambda))dF = 1-(1-\mu)s^n - \mu s^e. \tag{11}$$

In this version of the model with moral hazard, the planner solves:

$$SP2 = \max_{c^n, c^e \geq 0, (s^n, s^e, \mu)\in[0,1]^3, \bar{k}_0 \geq 0} (1-\mu)c^n + \mu(c^e - \gamma) + \int_0^{\bar{k}_0}\int (c^o(\lambda) - k_0 - k)\, dF(\lambda) dH(k_0)$$

subject to (2), (4), (7), (8), (10), and (11).

The introduction of moral hazard causes aggregate investment and consumption to fall because it is costly to train experts. To economize on the number of experts the planner will pick μ so that there is just enough experts to pay the due diligence cost of all active originators, that is $\mu = kH(k_0^{SB})$ where k_0^{SB} denote the originator cut-off in this second-best problem. These considerations imply:

Proposition 2. With unobservable effort, the optimal allocation satisfies $\mu = kH(k_0^{SB})$ and $s^e = 0$, where k_0^{SB} solves

$$\int \lambda r dF = k_0^{SB} + k + \gamma k.$$

All proofs are in the appendix. Notice that when $\gamma = 0$ this solution coincides with the first best allocation. The planner internalizes the cost of expertise by reducing the number of active originators from $H(k_0^*)$ to $H\left(k_0^{SB}\right)$. Furthermore, it is optimal that experts spend all their endowment on diligent originators, and that originators get as much surplus as possible, as this reduces the number of experts necessary to finance originators. So at the optimum

$s^e = 0$, $c^e = 1 + \gamma$ and $c^n = 1$.[11] In other words, investors, whether experts or not, get the same utility they would if they invested their entire endowment in the storage technology leaving the entire surplus generated by projects for originators. This reduces the number of experts to the bare minimum while still providing the right incentives to active originators.

4. Market equilibrium with full disclosure

We will now compare the second-best allocation we just described to the allocations decentralized markets can deliver under full disclosure. Similarly to Bolton, Santos, and Scheinkman (2016) we formalize the segmentation between experts and non-experts by assuming that the economy comprises two trading stages at date 0. In the first trading stage, originators can sell all or part of their projects to any investor. The second trading stage is only accessible to experts. This segmentation is meant to formalize the assumption that non-experts cannot learn from the orders or mimic the behavior of experts.

Bolton, Santos, and Scheinkman (2016) think of the segment of the market only accessible to experts as OTC markets. They assume that all trades in these markets take place in bilateral meetings. Instead, we adopt a Walrasian trading protocol in the sense that investors take all prices as given and we will require that demand and supply coincide for each project type on the expert market. Rents arise in equilibrium not due to bilateral monopoly power but because entry into expert markets is costly. This difference in trading protocol on the

[11] As $\mu = kH(k_0^{SB})$, $c^e = c^n + \gamma$ and $s^n = \frac{1-H(k_0^{SB})(k_0^{SB}+k)}{1-k_0^{SB}H(k_0^{SB})}$, we can use the resource constraint (2) to obtain

$$\begin{aligned} c^n + kH(k_0^{SB})\gamma &= 1 - H(k_0^{SB})(k_0^{SB} + k) + H(k_0^{SB}) \int \lambda r dF \\ \iff c^n &= 1 + H(k_0^{SB}) \int \lambda r dF - H(k_0^{SB})(k_0^{SB} + k + \gamma k) \\ \iff c^n &= 1. \end{aligned}$$

Since, in addition, $\mu \in (0, 1)$ at the optimum, we must have $c^e = 1 + \gamma$.

expert market is not essential and we could easily establish this paper's main results under a matching arrangement.[12]

The key assumption both in our model and in Bolton, Santos, and Scheinkman (2016) is that expert markets are dark: it is not possible for non-experts to infer project types from the willingness of experts to pay for specific projects, i.e. from their orders. Otherwise no positive rents could exist in equilibrium since investors would simply free-ride on the experts' knowledge. Exactly as in Bolton, Santos, and Scheinkman (2016) then, segmented expert markets are essential to justify costly investment in expertise in the first place.[13]

In the second trading stage, originators convey the quality of their project to experts. So there are effectively as many expert submarkets as there are λ's. We conjecture and will later verify that experts will choose to participate only in the second trading stage. Under that conjecture, projects are indistinguishable to investors in the first stage and so they sell for the same price. Let p_n denote the price of a claim to the entire cash flow of a project selling on the non-expert market, while $p_e(\lambda)$ is the price of an entire project of quality λ on the expert market.

Consider first the problem of an originator who chose to incur cost k_0 to become active – a choice which is publicly observable – and chose to exert due diligence at cost k – a choice which is private information. Given prices p_n and $p_e(\lambda)$, the problem of an active and

[12] Bolton, Santos, and Scheinkman (2016) use the bilateral matching aspect of their environment to propose a mechanism that could explain why expert rents have grown even as the number of experts in the Finance industry has risen. This happens in their model because more experts implies more cream-skimming by OTC markets hence a lower threat point for asset sellers. A version of our model with heterogeneous and unobservable expertise-acquisition cost γ across investors yields the same prediction. In that case, investors with the lowest expertise cost will become experts first and the size of rents is pinned down by the marginal expert. Any increase in volume – caused say by a decrease in origination cost, a decrease in storage returns, or an increase in investor endowments – will increase the number of experts in equilibrium, hence the training cost level of the marginal expert, hence the size of rents. In that simple variation on our model, our other results go through unchanged. The main change is that investors who have a low expertise cost now receive pure economic rents.

[13] Our market environment can also be mapped into a setting where originators conduct uniform price auctions as in Pagano and Volpin (2012). In that case, originators know that investors' willingness-to-pay for project is as defined in lemma 1 and sell their projects to the highest bidders.

diligent originator is to maximize his payoff $V(\lambda)$ by choosing the proportion $\theta(\lambda) \in [0,1]$ of his project's cash flow to sell on the expert market, or

$$V(\lambda) = \max_{\theta(\lambda)\in[0,1]} (1-\theta(\lambda))\, p_n + \theta(\lambda) p_e(\lambda).$$

An originator with project type λ will sell his entire project in the expert market whenever $p_e(\lambda) > p_n$. An active originator will perform due diligence whenever the expected value of the resulting project exceeds what he would get from selling a project of type $\lambda = 0$ to non-experts, which requires no effort:

$$\int V(\lambda) dF(\lambda) - k \geq p_n.$$

An originator activates a project whenever its expected payoff is higher than the activation cost,

$$\max\left\{\int V(\lambda) dF(\lambda) - k, p_n\right\} \geq k_0. \tag{12}$$

Since the left hand side of (12) is independent of k_0, only originators with low enough k_0 will activate their projects. Letting $\bar{k}_0$ be such that (12) holds with equality, the measure of active originators is $H(\bar{k}_0)$.

Now consider the problem of non-experts in the first trading stage. Given prices, non-experts form expectations about the quality of the projects sold in their market. In any equilibrium, they understand that only projects with quality $\lambda \in \Lambda_n = \{\lambda : p_n \geq p_e(\lambda)\}$ are sold to them. Letting θ^n denote the investment of non-experts in risky projects, they solve

$$V_n = \max_{s^n\in[0,1],\theta^n\geq 0} s^n + \theta^n \int_{\Lambda_n} \lambda r \frac{dF(\lambda)}{F(\Lambda_n)}$$

subject to

$$1 - s^n = \theta^n p_n.$$

Replacing θ^n using the budget constraint, the problem becomes

$$V_n = \max_{s^n \in [0,1]} s^n + (1 - s^n) \int_{\Lambda_n} \frac{\lambda r}{p_n} \frac{dF(\lambda)}{F(\Lambda_n)} \tag{13}$$

As for the second stage, given the set of project types $\lambda \in \Lambda_e = \{\lambda : p_n < p_e(\lambda)\}$ available to experts and given prices $p_e(\lambda)$ for each type $\lambda \in \Lambda_e$, experts choose an investment $\theta^e(\lambda)$ in each sub-market $\lambda \in \Lambda_e$ to solve

$$V_e = \max_{s^e \in [0,1], \{\theta^e(\lambda) \geq 0 : \lambda \in \Lambda_e\}} s^e + \int_{\Lambda_e} \theta^e(\lambda) \lambda r d\lambda$$

subject to

$$1 - s^e = \int_{\Lambda_e} \theta^e(\lambda) p_e(\lambda) d\lambda.$$

Therefore, once again,

$$V_e = \max_{\{\theta^e(\lambda) \geq 0 : \lambda \in \Lambda_e\}} 1 + \int_{\Lambda_e} \theta^e(\lambda) p_e(\lambda) \left(\frac{\lambda r}{p_e(\lambda)} - 1 \right) d\lambda. \tag{14}$$

Finally, agents decide whether or not to become experts given V_e and V_n. Denote by μ the measure of experts, so that, $\mu = 1$ whenever $V_e - \gamma > V_n$, $\mu = 0$ whenever $V_e - \gamma < V_n$, and $\mu \in [0, 1]$ otherwise. We can now define an equilibrium.

Definition 3. An equilibrium is a list $\left\{\Lambda_e, \Lambda_n, p_n, \theta^n, s^n, s^e, [p_e(\lambda), \theta^e(\lambda) : \lambda \in \Lambda_e], \bar{k}_0, \mu\right\}$ such that given p_e and p_n,

1. the decisions of originators and investors are optimal (i.e. solve the problems described above),

2. these decisions are consistent with investor beliefs, i.e. for almost all λ,

$$\lambda \in \Lambda_e \quad \Leftrightarrow \quad \theta(\lambda) > 0,$$
$$\lambda \in \Lambda_n \quad \Leftrightarrow \quad \theta(\lambda) < 1,$$

3. markets clear,

$$\mu\theta^e(\lambda) \;=\; f(\lambda)\theta(\lambda)H(\bar{k}_0) \qquad \text{for almost all } \lambda \in \Lambda_e$$
$$\text{and } (1-\mu)\theta^n \;=\; H(\bar{k}_0) - H(\bar{k}_0)\int_{\Lambda_e}\theta(\lambda)dF(\lambda).$$

Experts will spend resources in the second trading stage whenever the return on a positive mass of projects is higher than the return on storage. If experts buy projects with different quality then it must be that the return on these projects is the same, i.e. $R_e \equiv \frac{\lambda r}{p_e(\lambda)}$ for all $\lambda \in \Lambda^e$ for some constant $R_e \geq 1$. Without loss of generality then, we can rewrite (14) as

$$V_e = \max_{s^e \in [0,1]} 1 + (1 - s^e)(R_e - 1). \tag{15}$$

We will now argue that in any equilibrium where non-experts use the storage technology their return from investment must be one, while the return for experts must be just enough to cover the utility cost they must bear to become experts, namely $1 + \gamma$.

Lemma 1. *In any equilibrium in which a positive mass of projects is activated, prices satisfy* $p_n \geq E(\lambda r|\lambda \in \Lambda_n)$ *and* $R_e = \frac{\lambda r}{p_e(\lambda)} = 1 + \gamma$ *for all* $\lambda \in \Lambda_e$. *Furthermore,* $p_n = E(\lambda r|\lambda \in \Lambda_n)$ *whenever* $s^n < 1$.

All investors must earn at least the return they would expect from storage. Experts must also be compensated for the training cost they bear. This implies lower bounds on the returns of each type of investor. On the other hand, investors would spend all their resources on buying projects if their expected net returns (net of the utility cost for experts) strictly exceeded storage returns. But assumption 2 implies that this cannot be an equilibrium since some originators would otherwise expect negative net returns from activating their project. In addition, expert markets must be active if projects are sold in equilibrium since otherwise originators have no incentives to be diligent. It follows that in any trading equilibrium, $R_e = 1 + \gamma$. Still it is possible that in equilibrium non-experts only use storage and do not buy any project. Thus $p_n = E(\lambda|\lambda \in \Lambda_n)r$ can only be guaranteed if non-experts do participate, i.e. if $s^n < 1$.

One consequence of this result is that, as conjectured, in any equilibrium experts only participate in the second market, where they earn a higher return. Another key consequence of this result is that when both market types are active, there is a pivotal originator who is indifferent between the two markets:

Corollary 4. If both markets are active, all projects with quality $\lambda \leq \tilde{\lambda}^M$ will be sold in the non-experts market while projects with quality $\lambda > \tilde{\lambda}^M$ will be sold in the expert market, where $\tilde{\lambda}^M$ solve:

$$E(\lambda|\lambda \leq \tilde{\lambda}^M) = \frac{\tilde{\lambda}^M}{1+\gamma} \tag{16}$$

This equation must hold to guarantee that originators with project quality $\lambda \leq \tilde{\lambda}^M$ do not want to sell their project in the expert market and, vice-versa, that originators who sell in expert markets would not be better off selling in non-expert markets. The number of solutions to (16) depends on the properties of $E(\lambda|\lambda \leq \tilde{\lambda}^M) - \frac{\tilde{\lambda}^M}{1+\gamma}$. The case with $\gamma = 0$ has been extensively studied since it arises in many search models and signaling games.[14] It is known for instance that $E(\lambda|\lambda \leq \tilde{\lambda}) - \tilde{\lambda}$ is (strictly) monotone and log-concave if F is

[14]See Bagnolli and Bergstrom (2005) for a survey or Burdett (1996). In general, there could be several solutions to (16) but a standard application of Sard's lemma implies that as long as the distribution F of project quality is continuous, the set of solutions to this equation generically has measure zero, a fact we will use in the next section.

(strictly) log-concave, which holds for most distributions typically used in economics. While there is no such general results for $E(\lambda|\lambda \leq \tilde{\lambda}^M) - \frac{\tilde{\lambda}^M}{1+\gamma}$ where $\gamma > 0$, monotonicity will hold at least approximately for γ small.

Figure 1 illustrates this for the case where F is assumed log-normal. The two sides of equation (16) obviously agree at $\tilde{\lambda} = 0$ which means, as we will explain below, that equilibria with only experts always exist. The question is whether a market equilibrium with both expert and non-expert transactions also exists. A necessary condition for this is that equation (16) admit at least one strictly interior solution, as happens in the case drawn in figure 1. But as we will discuss below, it is also necessary that the resulting expected rents for originators be sufficiently high to induce effort on their part. Graphically, this means that the $F-$weighted surplus past the intersection of the two curves on figure 1 must exceed the due diligence cost k. This second requirement turns out to imply that all market equilibria are inferior to the second-best solution, and this is the main result of this section:

Proposition 1. *At a market equilibrium, fewer projects are activated than at the second best and the ratio of experts to projects is higher.*

The allocation that prevails in a market-based economy is suboptimal because markets impose a condition that the social planner could ignore: originators know the willingness-to-pay for projects of both types of investors and, therefore, choose to sell their projects to the highest bidder. Non-experts, as a result, know that they are getting the worst projects and this affects their willingness to participate in the market in the first place. The result above says that market outcomes are always suboptimal. In fact, the adverse selection problem may be so severe that non-experts may prefer to just rely on the storage technology,

Remark 1. There are market equilibria in which only experts buy projects.

When both markets coexist, the willingness to pay of each type of investors is pinned down by lemma 1 so that, in turn, the project quality cutoff must solve equation (16). But at

Figure 1: Project quality cutoff in a market equilibrium

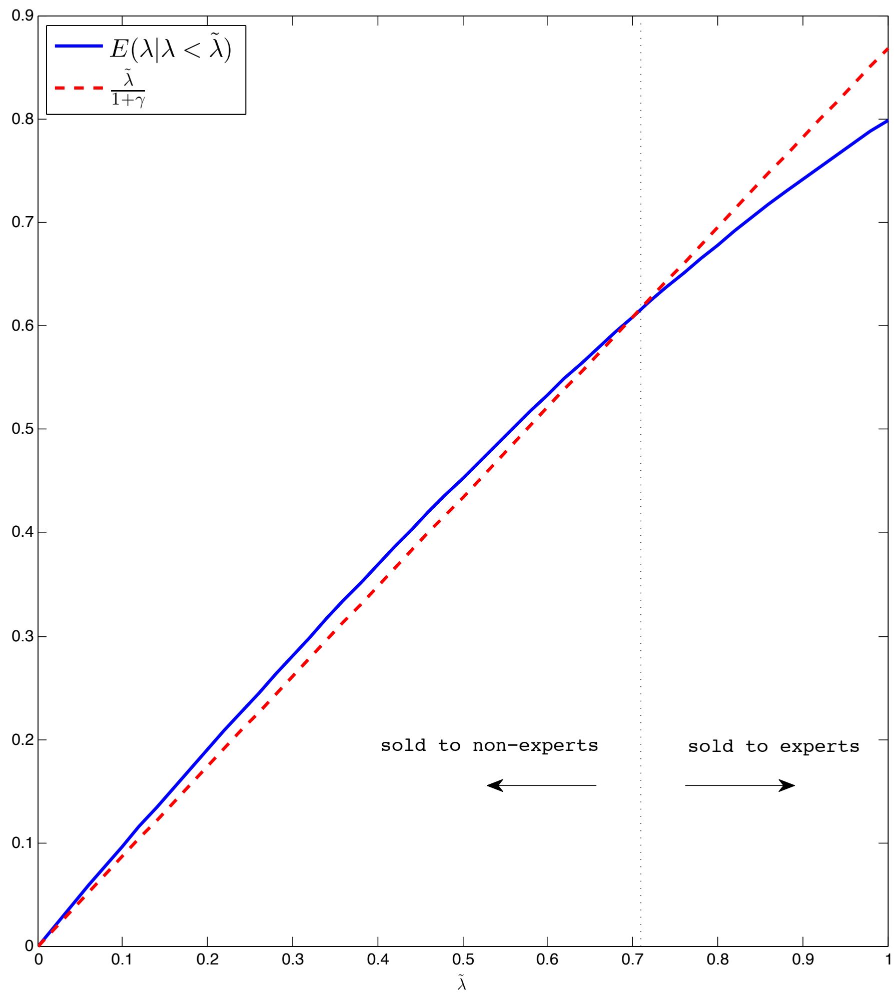

Notes: The figure is drawn for the case where F is log-normal with location parameter 0 and dispersion parameter 0.3, truncated to the unit interval, and $\gamma = 15\%$ so that experts buy projects in primary markets at a 15% discount to fundamental value.

this cut-off the co-existence of expert and non-expert investors may be such that originators do not have incentives to be diligent. As illustrated in figure 1, F may be such that only one interior solution to equation (16) exists. If that one solution does not justify due diligence by originators, no projects would be created. In that case, non-experts do not participate at all and markets for projects are small. Indeed, the originator talent cutoff must solve:

$$\int_0^1 \frac{\lambda r}{1+\gamma} dF = \bar{k}_0 + k.$$

To support the resulting allocation as an equilibrium, non-experts must believe that $\Lambda_n = \{0\}$ and those degenerate beliefs are trivially consistent with the resulting equilibrium choices. Worse still, there could be equilibria where no projects are activated at all, i.e. it could be that all market equilibria feature $\Lambda_n = \Lambda_e = \{0\}$. This is the case when

$$\int_0^1 \frac{\lambda r}{1+\gamma} dF < k. \tag{17}$$

In that situation, the incentive compatibility constraint cannot hold for any originator since $\int_0^1 \frac{\lambda r}{1+\gamma} dF$ is the highest possible reward they can expect for exerting effort.

Market equilibria with both investor types, when they exist, are inefficient because resources are wasted on making originators diligent. To understand why, recall that the quality threshold under full information must solve equation (16). Furthermore, originators choose to be diligent only if

$$\int_{\tilde{\lambda}^M}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda | \lambda \leq \tilde{\lambda}^M) r \right) dF \geq k. \tag{18}$$

Among the solutions to equation (16) one must satisfy the above inequality for an equilibrium with two active markets to exist. But there is no economic force that guarantees this inequality is tight. Resources are wasted on making sure that it is incentive compatible for originators to be diligent which requires informing more agents than necessary. This causes the equilibrium quantity of projects to be low given the rents experts must earn to offset their date-0 investment. In the next section we show that erasing fundamental information about assets, in a way we make precise, can improve market outcomes.

5. Optimal disclosure

We have already argued that when costly investments in expertise are essential, the ability of experts to hide their orders plays an essential role in supporting an equilibrium in which some risky projects are activated. But we will now argue that a second form of opacity – limited disclosure - can also create social value. Withholding information about the quality of projects is potentially beneficial because it can encourage the participation of non-experts in primary markets hence lead to origination volumes closer to the second-best level.

To see this, consider a planner who is equipped with a technology that can coarsen information in the same sense as in Monnet and Quintin (2013). Specifically, the planner can set up a disclosure technology that sends a public message m once λ is realized on a project. The message is fully understood by originators and experts alike, but not by non-experts. The planner can choose any message function, but to keep things simple and without loss of generality for our purposes, we consider messages $m(\lambda)$ that are subsets of $[0, 1]$ where $\lambda \in m(\lambda)$. This allows for full disclosure – $m(\lambda) = \lambda$ for all λ –, full opacity –$m(\lambda) = [0, 1]$ for all λ –, as well as any partial disclosure policy in between these two extremes.

All agents know the disclosure technology, but only experts can fully process the information contained in the message. For clarity, we emphasize that the planner must still rely on markets. In other words, the planner gets to choose the information environment under which agents will trade but once she has selected an information design, decentralized trade as described in the previous section dictates the allocation that prevails. Put simply, the planner chooses what information to provide before trading for projects begins.

The market equilibria we discussed in the previous section are those that obtain when there is full disclosure, i.e. when $m(\lambda) = \lambda$ for all $\lambda \in [0, 1]$. But the planner can improve welfare by selecting a more limited disclosure design. Indeed, with control over information the planner can make sure that (18) is tight. Better yet, even when (18) is violated at all interior solutions to equation (16), the planner can induce participation by non-experts by limiting disclosure. To do this, the planner simply needs to provide just enough information to direct projects to the desired market – i.e announce the quality subset to which a project belongs –, provide sufficient information for experts to verify the expected quality of the project, and make sure that due diligence holds. It follows that the optimal information solution assumes a simple form:

Proposition 5. Assume that $\int_0^1 \frac{\lambda r}{1+\gamma} dF > k$. A message function that partitions the project quality interval in two subsets $\Lambda_n^* = [0, \lambda^B]$ and $\Lambda_e^* = [\lambda^B, 1]$ is optimal, where

$$\int_{\lambda^B}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda | \lambda \leq \lambda^B) r \right) dF = k.$$

Projects in Λ_n^* are sold to non-experts while projects in Λ_e^* are sold to experts.

Generically, this optimal disclosure design generates strictly more social surplus than all market equilibria under full information. More projects are activated under the optimal information design than under full information.

The assumption $\int_0^1 \lambda r dF > k(1 + \gamma)$ is necessary to justify due-diligence by originators. This remains true even under the best possible disclosure design. Several additional comments are in order. While full disclosure is usually not optimal, neither is full opacity. The planner needs to reveal enough information to induce originators to perform due diligence. In addition, while the design above is best for a planner who must rely on markets once a disclosure policy is set, it still does not achieve the second best. In the second best environment, the planner only faces basic participation and incentive compatibility constraints. She does not need to worry about originators with high-quality projects trading with investors once consumption profiles have been proposed.

Finally, although weakly optimal, it is typically not necessary to suppress all information within the expert subset. Take for instance the case depicted in figure 1. Since $E(\lambda|\lambda \leq \tilde{\lambda})$ cuts $\frac{\tilde{\lambda}}{1+\gamma}$ from above, there is no desire for any originator in $\Lambda_e^* = [\lambda^B, 1]$ to defect to the non-expert market. In that case, the planner can provide finer information, even full disclosure for all quality past λ^B. Expected payoffs for experts are the same whether there is full disclosure in their market or they only know that $\lambda \geq \lambda^B$. Suppressing information on the non-expert market, on the other hand, is necessary. Indeed, since the objective is to increase the quality threshold between the two markets beyond the limits imposed by equation (16), some originators in Λ_n^* would defect to the expert market if they knew exactly their project type. So the market for non-experts must be opaque under the optimal disclosure design.

6. Decentralization of the optimal disclosure design

The previous section establishes that a benevolent planner who has to depend on markets would opt for a restricted information design. This invites a natural implementation question: what mechanisms can deliver the optimal level of information in financial markets? In other words, how can the optimal information solution be operationalized? One obvious candidate is a regulatory approach: well-meaning governments could in principle mandate the desired information design. This section will begin by questioning the practicality of such an approach. We will then argue that a decentralized approach that relies on arrangements that are routine in financial markets can implement the desired solution.

6.1. Regulation

Ex ante, originators want to commit to the optimal disclosure structure described in the previous section. Once project qualities are revealed however, originators with good projects want to share that information with experts. That is incompatible with the optimal information structure. So originators need to find a way to tie their hands ex-ante.

One way to deliver the optimal information solution would be for financial authorities – the Securities and Exchange Commission (SEC), e.g. – to directly control the availability of information in asset markets. While trivial in principle, this approach is impractical for several reasons. First, financial authorities would have to understand each and every specific asset markets fully to arrive at the right information prescription. In practice, they would likely get it wrong and may in fact propose a policy that is worse than full transparency. With a permanently fixed distribution of fundamentals one could argue that authorities could converge to the right solution via experimentation but such an assumption seems hard to defend in the context of primary asset markets.

Even if those informational issues can be overcome, financial authorities would need to control communication between originators and investors. While it seems feasible to mandate more transparency, it seems much more difficult to require opacity, as the regulator would need to verify that originators and investors do not share any information, even in private meetings. This would require a level of surveillance that seems infeasible both on practical and legal grounds.

Given informational and practical constraints therefore, regulation is unlikely to bring about the optimal informational design. A decentralized approach where originators choose the information policy solves both issues in principle. It only requires asset originators to understand the uncertainty they face in their area of specialization and to control their own disclosures. But it presents a fundamental commitment problem: originators, as a group, are willing to commit to the optimal solution before they discover the quality of the asset they produced. But upon receiving good news, some originators want to renege on that arrangement. We will now argue that this commitment problem can be overcome using simple contracts that are routinely written in financial markets.

6.2. Forward contracts

A simple way for a coalition of originators to implement the desired allocation is to sell pre-set quantities of projects in forward markets. To see this, assume that originators of ability $k_0 < k_0^B$ enter at date zero in a forward contract that promises the delivery of projects ex-post for a pre-set price $p_n^F = E(\lambda|\lambda \leq \lambda^B)r$.[15] Ex-ante, the coalition agrees to deliver a quantity $H\left(k_0^B\right)\int_0^{\lambda^B} dF$ of projects for that price. Ex-post, originators use the lowest value projects to fulfill their forward obligations – this is the standard *cheapest-to-deliver* result – and sell the remaining projects to experts.[16] In this context, the only role left for a regulator to play is that of enforcing contracts.

In the expert market, originators sell projects forward for $p_e^F = \frac{E(\lambda|\lambda \geq \lambda^B)r}{1+\gamma}$ but they must show that the projects delivered in that market satisfy the condition $\lambda \geq \lambda^B$, i.e. that they meet certain underwriting standards, which experts can authenticate. Originators agree to deliver a quantity $H\left(k_0^B\right)\int_{\lambda^B}^1 dF$ of projects that meet those standards. As we pointed out at the end of section 5., this second forward market and the associated opacity may not be necessary. If fundamentals look as they do in figure 1 then there is no need for any opacity at the top end and it is enough in that case to have a simple, full-information spot market alongside the opaque bottom-of-the-quality-distribution forward market.

Each originator in the coalition gets p_e^F if they end up contributing a project of quality $\lambda \geq$

[15] There is no need for all originators to form one grand coalition. Smaller representative coalitions of originators of ability $k_0 < k_0^B$ could enter into the same arrangement, the argument we present in this section can easily be scaled down.

[16] The same bundling outcome could be accomplished via the block-booking arrangements described by Kavajecz and Kelm (2005).

λ^B and p_n^F otherwise. This arrangement replicates exactly the ex-ante payoff of originators under the optimal disclosure policy. And since pre-diligence payoffs are unchanged, due-diligence is guaranteed as well. In other words, selling a fraction of projects before they are created optimally solves the disclosure problem.

This implementation closely resembles for instance the way agency mortgage backed-securities trade in the United States. As discussed *inter alia* by Vickery and Wright (2013) the government-sponsored enterprises (GSEs) build and market mortgage pass-through securities in two distinct ways. The majority of deals are floated in a forward market termed the TBA market. Investors buy pools that have yet to be delivered a few weeks in advance with the GSE having committed only to abiding by a few average pool-wide characteristics. Other pools are marketed in the "specified", spot market. By hiding individual loan characteristics, TBA market enhance liquidity by commoditizing transactions. The specified market, for its part, creates an avenue to market loans which are not TBA eligible.

But a non-negligible fraction of TBA-eligible loans are also sold in the specified market. Interestingly given the purpose of this paper, the majority of these loans are loans whose observable characteristics make them low prepayment risk. GSEs provide insurance against default which is implicitly backed by the full faith of the US government and make Agency securities effectively as safe as US Treasuries. Unlike Treasuries however, pass-throughs are subject to prepayment risk and the corresponding yield degradation. GSEs thus recognize that loans with characteristics such as low balances have a higher market value hence finds it profitable to convey that information to investors such as investment banks and mortgage REITs that know how to process that information. On the flip side and just like in our model, TBA investors understand that they face a "cheapest-to-deliver-problem" which affects their willingness to pay for the pools they purchase.

6.3. *Proportional-fee underwriting*

Another way for originators to implement the ex-ante optimal disclosure design is to delegate market decisions to an underwriter. Assume that, before project quality levels are drawn, originators of ability $k_0 < k_0^B$ hire one of the inactive prospectors (the underwriter) and delegate to this underwriter the market selection decision. The contract gives the underwriter the choice of a quality threshold λ^U below which projects are sold to non-experts while other projects are sold to experts. In addition, the contract stipulates that the underwriter will be paid a fraction $\delta > 0$ of project sales proceeds. Given this contract, the underwriter solves

$$\max_{\lambda^U \in [0,1]} \delta \left(\int_0^{\lambda^U} E(\lambda|\lambda \leq \lambda^U) r dF + \int_{\lambda^U}^1 \frac{\lambda r}{1+\gamma} dF \right)$$

subject to:

$$\int_{\lambda^U}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda|\lambda \leq \lambda^U) r \right) dF \geq k.$$

Like the planner, the underwriter has to provide enough incentives to originators to exert due diligence, which brings in the constraint above.[17] The underwriter takes as given that non-experts require a zero net return on their purchases while experts demand a return of $1+\gamma$. Furthermore, as it is the case in practice where underwriters usually control all communication with potential investors during the IPO process, the underwriter in this construction controls the disclosure policy. Exactly like the planner in the previous section therefore, the underwriter maximizes originator surplus subject to participation by investors – implicitly

[17] Since there are many prospectors and one underwriter by assumption, the underwriter compensation has negligible mass. The assumption is that originators can share underwriting services. If underwriting is associated with net transaction costs for originators, the implementation can only approximate the optimal disclosure solution we derived in the previous section, which is not surprising since that problem assumed no transaction costs of any sort. See appendix 8.5. for a full discussion. The same caveat applies to our other implementations. Introducing transaction costs make those implementations approximations of the optimal arrangement.

reflected in the willingness to pay of investors – and subject to a due diligence constraint. The key assumption is that the underwriter can erase the information content on $[0, \lambda^U]$. And, once again like the planner, the underwriter will take full advantage of that flexibility until the due diligence condition binds. It follows that under an underwriting arrangement with proportional fees the threshold quality is $\lambda^U = \lambda^B$, so that the proportional-fee underwriter selects the optimal disclosure design.

The underwriter cannot deviate from the announced disclosure policy without being detected. The set of projects he handles is representative of the project-type population. Therefore, the mass of projects directed to either market is predictable under the informational arrangement. Any non-negligible deviation by the underwriter from the pre-announced policy becomes public information. The contract can specify zero compensation in that case.

This implementation scheme accords well with the different way in which underwriting intermediaries treat different investor types in practice. Institutional investors receive direct information from management via pre-IPO "road-shows" and other events from which most retail investors are barred.

In the context of this IPO interpretation, our model suggests that the connection between disclosure and underpricing is more complex than models with exogenous masses of investor types may suggest. Improving the quality and reliability of accounting disclosures, for instance, is probably best interpreted in our model as a reduction in γ since it makes the evaluation of asset quality easier, and leads to less underpricing. Releasing complex information that only experts can interpret, on the other hand, can lead to less participation by non-experts and more underpricing in our model.[18]

[18] The empirical literature (see e.g. Ljungqvist 2004 and the references therein) on the relationship between disclosure and underpricing has found that disclosure policies that reduce asymmetric information between experts and non-experts lead to less underpricing while others – typically the release of soft information such as risk-factors or use-of-funds information – can lead to more underpricing. The second type of finding is usually interpreted as a consequence of increased uncertainty or ambiguity (see Arnold, Fishe, and North 2010). Our model is in line with this second interpretation. If the information provided is difficult to interpret by non-specialists, they will tend to shy away from the IPO, which increases the ratio of experts to

6.4. *Blank-check underwriting*

The commitment device provided by forward markets can also be replicated via blank-check underwriting. Under a pure blank-check agreement, a corporation raises funds to purchase, at a later date, unspecified assets in a broadly defined class. In our model, a coalition of originators of ability $k_0 < k_0^B$ can raise funds $E(\lambda|\lambda \leq \lambda^B)r$ with the commitment to deliver an agreed-upon number $H\left(k_0^B\right)\int_0^{\lambda^B} dF$ of active projects. Ex-post once again originators will direct the lowest quality projects to the blank-check corporation and save better projects for expert markets. Because of the presence of these experts, non-experts know that originators have sufficient incentives to perform due diligence ex-ante. Therefore, once again, this arrangement replicates the desired disclosure design and allocation.

The 1990 Penny Stock Reform Act stipulates a right of rescission for investors, requiring that at least 80% of original investors approve targets, and imposes a time-limit of no more than two years within which funds must be either deployed or returned to investors. In the context of our model, one interpretation of these rescission and approval provisions is that they provide enough safeguards for investors to make sure that originators who enter into blank-check contracts do in fact incur the creation cost k_0 required to produce a project. Otherwise, there would be nothing preventing originators from selling empty-shells forward at no cost to them.

Blank-check vehicles accounted for 20% of funds raised for IPOs in 2007 according to Jenkinson and Sousa (2011). The fact that blank-check underwriting has thrived despite increasing regulatory hurdles is evidence that it serves an important role in financial markets. On the investor side, it enables non-institutional investors to participate in IPO markets. On

non-experts and increases underpricing. Boone et. al. (2016) point out that almost 40 percent of IPO firms redact some information from their SEC filings and that those firms tend to experience greater underpricing. They also find that firms that are younger and more research-and-development-intensive are more likely to control information releases. These are once again firms whose releases are more difficult to interpret by non-specialists.

the target side, it enables some companies to go public, while being excluded from standard IPO routes because they are small or young.

7. Conclusion

This paper is founded on a simple premise: the ubiquity of dark trading platforms and partial disclosure in asset markets suggests that both forms of opacity serve an essential purpose. We show that two simple frictions – moral hazard on the part of asset suppliers and the fact that acquiring expertise is costly for investors – suffice to produce an environment where, optimally, expert and non-expert investors coexist and platforms of varying transparency level emerge to cater to those different types of investors. Compared to models where the distribution of investor types is exogenous, this point of view leads to a different set of predictions: mandating disclosure of information that is difficult to interpret may exacerbate average underpricing in primary asset markets and reduce welfare.

8. Appendix

8.1. *Proof of proposition 2*

Substituting feasibility constraints into the objective function simplifies the planner's problem to

$$SP2 = \max_{\bar{k}_0, \mu \in [0,1],} 1 + H(\bar{k}_0) \int \lambda r dF - \mu\gamma - \int_0^{\bar{k}_0} (k_0 + k)\, dH(k)$$

subject to (7), (9) and (10). The planner clearly wants μ as small as possible while still satisfying (10). Therefore, it is optimal for the planner to require $s^e = 0$ and from (10) we obtain

$$\mu = H(\bar{k}_0) \int q^e(\lambda) dF.$$

We can then use (6) to substitute for $q^e(\lambda)$

$$\mu = H(\bar{k}_0) \int (c^o(\lambda) - q^n)\, dF$$

and using (7) with $c^p(0) = q^n$,

$$\mu = kH(\bar{k}_0).$$

Hence the planner is maximizing

$$SP2 = \max_{\bar{k}_0} 1 + H(\bar{k}_0) \int \lambda r dF - \bar{k} H(\bar{k}_0)\gamma - \int_0^{\bar{k}_0} (k_0 + k)\, dH(k_0)$$

with first order condition

$$\int \lambda r dF = \bar{k}_0 + k + \gamma k.$$

Condition (7) then implies that

$$q^n = \bar{k}_0$$

so that s^n is

$$s^n = \frac{1 - H(\bar{k}_0)(\bar{k}_0 + k)}{1 - kH(\bar{k}_0)}$$

The optimal originator ability threshold is in fact $\bar{k}_0$ as long as $H(\bar{k}_0)(\bar{k}_0+k) < 1$. Assumption (2) implies that this inequality holds.

8.2. *Proof of lemma 1*

As we argued in the text, the return to investing in projects must be constant on Λ_e. If it were the case that $p_e(\lambda)(1 + \gamma) > \lambda r$ on Λ_e then no investor would choose to become an expert. In this case, no originator would exert due diligence and all projects in the non-expert market have a zero payoff. Hence, no investors would purchase any project in the non-expert market. If the opposite inequality held then all agents would invest their entire endowment in projects. But this cannot be in equilibrium since, in that case, assumption (2) would be violated.

To see this, denote by q the expected hence average price per project so that $H(q - k)$ is the mass of active originators. Assume by way of contradiction that storage is not used by any investor. Then we must have

$$H(q - k)q = 1. \tag{19}$$

Investors must expect at least the same return as they would from the storage technology so we must have $q \leq \int \lambda r dF$. But this inequality together with condition (19) implies

$$H\left(\int \lambda r dF - k\right) \int \lambda r dF \geq 1,$$

a direct contradiction of assumption (2).

Therefore, in any equilibrium with positive investment in projects, $R_e = 1 + \gamma$. This immediately implies that $s^e = 0$ but, once again, should $s^n = 0$ as well then assumption (2) would be violated. Hence it must be that $p_n \geq E(\lambda|\lambda \in \Lambda_n)r$, with equality if $s^n < 1$.

8.3. Proof of proposition 1

Recall that at the second best the cutoff satisfies:

$$\int \lambda r dF = k_0^{SB} + k(1+\gamma).$$

Now consider a market equilibrium with both experts and non-experts – in market equilibria with expert transactions only, the result is obvious. From Lemma 1, the overall rate of return of non-experts has to be 1, while that of the experts has to be $1+\gamma$, as in the second best. (Here we assume that $s^n < 1$, but once we reach the conclusion of the proof, that will be shown as well since the market volume of projects is below the second-best volume.) This implies a cutoff $\tilde{\lambda}^M$ past which originators opt for expert markets. This cutoff must be one of the possibly multiple solutions to equation (16), the argument here does not depend in any way on whether that cut-off point is unique.

In turn, (12) as well as the functional forms for p_n and p_e imply that the cutoff participation level k_0^M under any such market equilibrium solves:

$$\int_0^{\tilde{\lambda}^M} E(\lambda|\lambda \leq \tilde{\lambda}^M) r dF + \int_{\tilde{\lambda}^M}^1 \frac{\lambda r}{1+\gamma} dF = k_0^M + k. \tag{20}$$

To prove the proposition, we just need to compare k_0^M and $\bar{k}_0$. First, observe that[19]:

$$\int_0^{\tilde{\lambda}^M} E(\lambda|\lambda \leq \tilde{\lambda}^M) r dF + \int_{\tilde{\lambda}^M}^1 \frac{\lambda r}{1+\gamma} dF = \int \lambda r dF - \int_{\tilde{\lambda}^M}^1 \frac{\gamma \lambda r}{1+\gamma} dF \tag{21}$$

Also, in any equilibrium, there has to be due diligence. Therefore:

$$\int_{\tilde{\lambda}^M}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda|\lambda \leq \tilde{\lambda}^M) r \right) dF \geq k, \tag{22}$$

from which it directly follows that

$$\int_{\tilde{\lambda}^M}^1 \frac{\gamma \lambda r}{1+\gamma} dF > k\gamma. \tag{23}$$

Combining (20)-(23), it follows that

$$\int \lambda r dF > k_0^M + k(1+\gamma).$$

Hence $k_0^{SB} > k_0^M$ so that fewer projects are activated.

Now recall from proposition 2 that in the second best the planner makes μ as small as possible by setting $\mu^{SB} = kH(k_0^{SB})$ which is the smallest set of resources compatible with incentive compatibility for originators. Define

$$q_e^M = \int_{\tilde{\lambda}^M}^1 \frac{\lambda r}{1+\gamma} dF.$$

[19] For the algebraic details, note first that:

$$\int \lambda r dF = \int_0^{\tilde{\lambda}^M} E(\lambda|\lambda \leq \tilde{\lambda}^M) r dF + \int_{\tilde{\lambda}^M}^1 \lambda r dF$$

But

$$\int_{\tilde{\lambda}^M}^1 \lambda r dF = \int_{\tilde{\lambda}^M}^1 \frac{\lambda r}{1+\gamma} dF + \int_{\tilde{\lambda}^M}^1 \frac{\gamma \lambda r}{1+\gamma} dF.$$

Substituting the second expression for $\int_{\tilde{\lambda}^M}^1 \lambda r dF$ in the first equation gives the desired result.

In words, q_e^M is the payoff originators expect from expert markets when they do exert effort. When $\mu < 1$, the due diligence constraint implies that

$$q_e^M > k.$$

But, exactly as in the proof of proposition 2, multiplying both sides by $H\left(k_0^M\right)$ and using feasibility implies

$$\mu^M > kH\left(k_0^M\right),$$

which implies

$$\frac{\mu^M}{H\left(k_0^M\right)} > \frac{\mu^{SB}}{H(k_0^{SB})}$$

as claimed.

8.4. *Proof of proposition 5*

Assume first that under full disclosure (when $m(\lambda) = \lambda$ a.s) a market equilibrium exists with $\tilde{\lambda}^M \in (0,1)$. A comparison of (22) and the definition of λ^B shows that $\tilde{\lambda}^M \leq \lambda^B$. Furthermore, under disclosure design $(\Lambda_n^*, \Lambda_e^*)$, all projects in Λ_n^* are observably the same. Since the willingness of non-experts to pay for expected cash-flows is higher than that of experts, originators are better off selling projects in Λ_n^* to non-experts than to experts, as needed for equilibrium. Given the definition of λ^B and conversely, originators sell projects in $[\lambda^B, 1]$ to experts. Therefore, under the proposed disclosure design, there is an equilibrium where $[0, \lambda^B]$ is the uninformed market while $[\lambda^B, 1]$ is the expert market. Furthermore and in the generic case where (22) does not hold exactly average utility is strictly higher with blind trading than without.

This suffices to show that limiting disclosure raises welfare strictly. But we will also show that the disclosure strategy $(\Lambda_n^*, \Lambda_e^*)$ cannot be improved upon by any other disclosure policy.

This will require several steps.

First, it is easy to see that there is no need for the planner to partition the quality interval in more than two subsets: expert and non-expert. Indeed, assuming otherwise, take all the messages that lead to expert sales and replace those messages by their union. The ex-ante payoff to originators and their incentives to perform due diligence has not changed. Neither has the average return of expert investors. A similar argument proves that non-expert messages can, likewise, be merged. Since the ex-ante payoff is left unaffected, binary messages suffice.[20]

Next consider any disclosure strategy such that is is **not** the case that, outside possibly of sets of F-measure zero, $\Lambda_n = [0, \lambda]$ for some $\lambda \geq 0$ and let $\Lambda_e = [0, 1] - \Lambda_n$. In other words, assume that there is a positive mass of projects sold to the non-expert market whose quality exceeds the quality of a positive mass of the projects sold to experts. We will show that such a strategy cannot be welfare maximizing. To be feasible, it must first be compatible with due-diligence:

$$\int_{\Lambda_e} \left(\frac{\lambda r}{1+\gamma} - E(\lambda | \lambda \in \Lambda_n) r \right) dF \geq k$$

The originator expected payoff under this messaging strategy or for that matter, any binary strategy, can be computed in two ways:

$$\int_{\Lambda_n} \lambda r dF + \int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF = \int_{\Lambda_n} E(\lambda | \lambda \in \Lambda_n) dF + \int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF \tag{24}$$

$$= \int \lambda r dF - \gamma \int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF \tag{25}$$

[20]This aspect of the argument builds on the strategic information transmission literature introduced by Crawford and Sobel (1982.) More specifically, this result closely resembles the "straightforward signal" results of Kamenica and Gentzkow (2011). The logic behind the optimality of simple messaging strategies is the same as in their paper. Since the planner knows what market originators will choose given a message and since this market choice is all that matters for ex-ante payoffs, the planner can effectively use messages that call for a specific market. Since there are two market options, two messages suffice.

Now pick $\tilde{\lambda} > 0$ such that

$$\int_{\tilde{\lambda}}^{1} \frac{\lambda r}{1+\gamma} dF = \int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF$$

and consider alternative messaging strategy $\left(\tilde{\Lambda}_n = [0, \tilde{\lambda}], \tilde{\Lambda}_e = [\tilde{\lambda}, 1]\right)$. In words, the alternative strategy picks the smallest interval at the top of the quality distribution that gives the same volume of expert transactions as under messaging strategy (Λ_n, Λ_e). Expression (25) implies that the expected payoff to originators is unchanged under this alternative, since the $\int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF$ term has not changed in value. But since the alternative strategy directs only the very top projects to experts, the average quality of projects sold to experts has strictly risen hence,

$$F\left(\tilde{\Lambda}_e\right) < F\left(\Lambda_e\right) \text{ while } F\left(\tilde{\Lambda}_n\right) > F\left(\Lambda_n\right).$$

Indeed, the average payoff per project on expert markets has increased. If the number of expert transactions had not decreased, the $\int_{\Lambda_e} \frac{\lambda r}{1+\gamma} dF$ part of the originator payoff would have risen strictly which we ruled out by construction of $\tilde{\lambda}$.

But now, since the left-hand of expression (24) is the same under $\left(\tilde{\Lambda}_n, \tilde{\Lambda}_e\right)$ than under (Λ_n, Λ_e), the right-hand side of that expression implies that $E(\lambda|\lambda \in \tilde{\Lambda}_n) < E(\lambda|\lambda \in \Lambda_n)$. In words, since we have moved high-quality projects from the non-expert set to the expert set, the average quality of non-expert transactions has fallen strictly. But this implies that the due-diligence constraint is slack under $\left(\tilde{\Lambda}_n, \tilde{\Lambda}_e\right)$ which means that increasing the quality cutoff strictly above $\tilde{\lambda}$ is feasible and compatible with due-diligence, which raises the expected payoff of originators strictly without changing the average consumption of investors, which increases welfare strictly.

This establishes that only messaging strategies such that $\Lambda_n^* = [0, \lambda]$ for some $\lambda \geq 0$ can be optimal. Within that class, the planner cannot go further than $\lambda = \lambda^B$ because that would violate the due diligence condition, and this establishes the desired result.

What if the only market equilibrium when $m(\lambda) = \lambda$ a.s. is such that only expert investors buy projects? Then the condition defining λ^B guarantees in one fell swoop that originators prefer selling "good" projects to expert agents and that non-expert agents are now willing to participate since due diligence is assured. And once again, pushing the boundary past λ^B would violate (22).

8.5. Underwriter compensation when the underwriter has positive mass

Assume that each originator is matched with one underwriter and that the payment to the underwriter reduces the eventual payoff to the originator in a one-to-one fashion. In other words, the ex-ante originator payoff under the optimal information arrangement is

$$(1-\delta)\left(\int_0^{\lambda^U} E(\lambda|\lambda \leq \lambda^U) r dF + \int_{\lambda^U}^1 \frac{\lambda r}{1+\gamma} dF\right).$$

In that context originators who do use an underwriter want to minimize the impact of underwriting fees on their payoff. In addition, the due diligence constraint becomes:

$$(1-\delta)\int_{\lambda^U}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda|\lambda \leq \lambda^U) r\right) dF \geq k.$$

Letting $\lambda^U(\delta)$ be the solution to this problem. By the same logic as in the text, that solution solves

$$\int_{\lambda^U(\delta)}^1 \left(\frac{\lambda r}{1+\gamma} - E(\lambda|\lambda \leq \lambda^U) r\right) dF = \frac{k}{(1-\delta)}$$

so that $\lambda^U(\delta) \nearrow \lambda^B$ as δ falls so that the solution is close to the optimal design solution when δ is low.

In the presence of contracting costs then, underwriting arrangement can only approximate the optimal disclosure design, which is not surprising since the optimal disclosure design problem we solved assumed no such costs. The same remark holds for the other implementations we develop in the text. In practice of course, underwriters perform services that would reduce origination costs k_0 and k. In the special case where those fall by a fraction of exactly δ, the underwriting solution coincides exactly with the optimal disclosure solution.

References

- Andolfatto, D., Berentsen, A., Waller, C., 2014. Optimal disclosure policy and undue diligence. Journal of Economic Theory 149, 128-152.

- Arnold, Tom, Fishe, R. P. H., North D., 2010. The effects of ambiguous information on initial and subsequent IPO returns. Financial Management 39, 1497-1519.

- Bagnoli, M., Bergstrom, T., 2005. Log-concave probability and its applications. Economic Theory 26, 445-469.

- Biais, B., Rochet J., Woolley, P., 2010. Innovations, rents, and risk. Unpublished working paper. Paul Woolley Centre Working Paper No 13, London School of Economics.

- Bolton, P., Santos, T., Scheinkman, J. A., 2016. Cream-skimming in financial markets. Journal of Finance. 17, 709-736.

- Boone, A. L., Floros, I. V., Johnson, S. A., 2016. Redacting Information at the Initial Public Offering. Journal of Financial Economics (forthcoming.)

- Burdett, K., 1996. Truncated means and variances. Economics Letters 52, 263–267.

- Caseres Field, L., Lowry, M., 2009. Institutional versus individual investment in IPOs: The importance of firm fundamentals. Journal of Financial and Quantitative Analysis 44(3), 489-516.

- Cazavan-Jeny, A., Jeanjean, T., 2007. Levels of voluntary disclosure in IPO prospectuses: An empirical analysis. Review of Accounting and Finance 6(2), 131-149.

- Crawford, V. P., Sobel, J., 1982. Strategic information transmission. Econometrica 50(6), 1431-1451.

- Dang, T. V., Gorton, G., Holmstrom, B., Ordonez, G., 2013. Banks as secret keepers. Unpublished working paper, University of Pennsylvania.
- Fishman, M., Hagerty, K., 2003. Mandatory versus voluntary disclosure in markets with informed and uninformed customers. Journal of Law, Economics and Organization 19, 45-63.
- Grossman, S. J., Stiglitz, J. E., 1980. On the impossibility of informationally efficient markets. American Economic Review 70, 573-585.
- Heyman, D.K., 2007. From blank check to SPAC: The regulator's response to the market, and the market's response to the regulation. Entrepreneurial Business Law Journal 2(1), 531-552.
- Hirshleifer, J., 1971. The private and social value of information and the reward to inventive activity," American Economic Review 61(4), 561–574.
- Jog, V., Sun, C., 2007. Blank check IPOs: A home run for management. Unpublished working paper, Carleton University.
- Jenkinson, T., Sousa, M., 2011. Why SPAC investors should listen to the market. Journal of Applied Finance 2(21), 38-57.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. American Economic Review 101, 2590-2615.
- Kavajecz, K., Keim, D., 2005. Packaging liquidity: Blind auctions and transaction efficiencies. Journal of Financial and Quantitative Analysis 40, 465 - 492.
- Koeppl, T., Monnet, C., Quintin, E., 2014. Efficient contract enforcement. Economic Theory 55, 161-183.

- Kurlat, P., 2015. Asset markets with heterogeneous information. Unpublished working paper. Stanford University.
- Leland, H., Pyle, D., 1977. Informational asymmetries, financial structure, and financial intermediation. Journal of Finance 32(2), 371-387.
- Ljungqvist, A., 2007. IPO underpricing: a survey. In: Eckbo, B.E. (Ed.), Handbook of Corporate Finance, North-Holland, Amsterdam.
- Monnet, C., Quintin, E., 2013. Rational opacity. Unpublished working paper, Wisconsin School of Business.
- Pagano, M., Volpin, P., 2012. Securitization, disclosure and liquidity. Review of Financial Studies 25, 2417-2453.
- Philippon, T., 2008. The evolution of the US financial industry from 1860 to 2007: Theory and evidence. Unpublished working paper, New York University.
- Philippon, T., Reshef, A., 2012. Wages and human capital in the U.S. finance industry: 1909–2006. Quarterly Journal of Economics 127(4), 1551-1609.
- Ritter, J. R., Welsh, I., 2002. A review of IPO activity, pricing, and allocations. Journal of Finance 42(4), 1795-1828.
- Rock, K., 1986. Why new issues are underpriced. Journal of Financial Economics 15 (1–2), 187–212.
- Sappington, D., 1983. Limited liability contracts between principal and agent. Journal of Economic Theory 29, 1-21.
- Vickery, J., Wright, J., 2013. TBA trading and liquidity in the Agency MBS market. Federal Reserve Bank of New York Economic Policy Review / May.