The Social Value of Financial Expertise

Pablo Kurlat*
Stanford University

February 2016

I study expertise acquisition in a model of trading under asymmetric information. I propose and implement a method to estimate the ratio of social to private marginal value of expertise. This can be decomposed into three sufficient statistics: traders' average profits, the fraction of bad assets among traded assets and the elasticity of good assets traded with respect to capital inflows. For venture capital, the ratio is between 0.64 and 0.83 and for junk bond underwriting, it is between 0.09 and 0.26. In both cases this is less than one so at the margin financial expertise destroys surplus.

Keywords: Financial industry, expertise, asymmetric information, sufficient statistic

JEL codes: D53, D82, G14, G20

*I am grateful to Sebastian Di Tella, Bob Hall Pete Klenow, Alejandro Martins, Monika Piazzesi, Martin Schneider, Johannes Stroebel, Victoria Vanasco and Iván Werning for helpful comments. Correspondence: Department of Economics, Stanford University. 579 Serra Mall, Stanford, CA, 94305. Email: pkurlat@stanford.edu

1 Introduction

The financial industry has been heavily criticized in recent years. One criticism often made is that it has simply become too large. Tobin (1984) worried that "we are throwing more and more of our resources, including the cream of our youth, into financial activities remote from the production of goods and services, into activities that generate high private rewards disproportionate to their social productivity". In the decades since Tobin's remark, the financial industry has become much larger. Philippon and Reshef (2012) and Philippon (2014) document that the share of value added of financial services in GDP has risen from about 5% in 1980 to about 8% in recent years.

While 8% of GDP is certainly a large number, it doesn't necessarily follow that it's excessive. In order to reach this conclusion one needs to have a framework for assessing how the size of the financial industry compares with the social optimum. Underlying the concern about the excessive size of the financial industry is a view that finance is a largely rent-seeking industry and that the resources it attracts would be better employed elsewhere. A converse point of view holds that the social value of the financial industry (fostering risk sharing, loosening credit constraints, increasing the informativeness of prices, etc.) may even exceed the income it obtains. Indeed, many of these benefits could be side-effects of activities that look a lot like rent-seeking: financial firms seeking out profitable trades end up reducing mispricing and making markets more liquid.

Several policies that have recently been under discussion would probably lead to reductions in the size of the financial industry, and in some cases that is their explicit purpose. These include special taxes on bank bonuses, higher capital requirements for banks and taxes on financial transactions. If indeed it is the case that the financial industry is too large relative to the social optimum, then the case for these policies is much stronger than otherwise.

In this paper I propose and implement a method to estimate r, the ratio of the marginal social value to the marginal private value of dedicating resources to the financial industry. If $r > 1$, then the marginal social value exceeds the marginal private value; under the assumption that marginal private value equals marginal cost, this implies that marginal social value exceeds marginal cost and a social planner would want the financial industry to expand from its current size. Conversely, if $r < 1$, the financial industry is too large.

The estimation is based on a particular model of what the financial industry does. I assume that financial firms earn income because they have expertise to trade in markets with asymmetric information: banks assess the creditworthiness of borrowers, venture capitalists

decide which startups are worth investing in, insurance companies evaluate risks, etc. Acquiring this expertise requires using productive resources that might be employed elsewhere: talented workers develop valuation models, IT equipment processes financial data, etc.

I formalize this in a model with the following elements. There is a group of households who own heterogeneous assets, either good or bad. Each household can keep its asset or sell it to a bank. Due to differential productivity or discount factors, selling assets creates gains from trade, which differ by household. Each household is privately informed about the quality of its own asset, while banks only observe imperfect signals about them. Each bank may, at a cost, acquire expertise. Having more expertise means receiving more accurate signals about the quality of the assets on sale.

I model trading using the competitive equilibrium concept proposed by Kurlat (2016). I assume markets at every possible price coexist and any asset can in principle be traded in any market. Households choose in what market (or markets, as there is no exclusivity) to put their asset on sale and banks choose what markets to buy assets from. Banks who want to buy may be selective, refusing to buy some of the assets that are on sale, but how selective they can be depends on their expertise. They can only discriminate between assets that their own signals allow them to tell apart. I do not impose market clearing. Assets may be offered on sale in a given market but not traded because there are not enough buyers who are willing to accept them. As in Gale (1996) and Guerrieri et al. (2010), rationing may and indeed does emerge as an equilibrium outcome.

In equilibrium, it turns out that all assets trade at the same price; owners of good assets can sell as many units as they choose at that price but owners of bad assets face rationing. Bad assets that are more likely to be mistaken for good assets face less rationing than easily detectable ones, and some assets cannot be traded at all. Only banks that are sufficiently expert choose to trade, while the rest stay out of the market. The price reflects the pool of assets acceptable to the marginal bank. Because this pool includes bad assets, households that sell good assets do so at a discount. Therefore, as in Akerlof (1970), households who have insufficiently large gains from trade choose not to sell, leading to a loss of surplus.

In this model, expertise is privately valuable to the individual bank because it enables it to better select which assets to acquire, improving returns. It is also socially valuable because it reduces overall information asymmetry, changing equilibrium prices and allocations and creating gains from trade. However, there is no reason for private and social values to be equal, i.e. no reason to believe $r = 1$. The private value depends on how expertise improves an individual trader's portfolio while the social value depends on how it shifts the entire

equilibrium.

It is possible to derive an analytical expression for r but it turns out to be quite complicated because it depends on various possible feedback effects. However, I show that it's possible to decompose the formula for r into sufficient statistics: measurable quantities that, combined, capture all the effects that are relevant for r without the need to separately estimate all the parameters of the model. In particular, I show that

$$r = \eta \left(1 - \frac{1-f}{\alpha}\right) \tag{1}$$

where η is the elasticity of the volume of good assets that are traded with respect to capital inflows, f is the proportion of bad assets among the assets that are traded and α is the average NPV per dollar invested earned by banks. α and f enter formula (1) because they measure the value of marginal trades: if banks make high profits despite acquiring a high fraction of bad assets, the adverse selection discount suffered by the marginal seller must be high, indicating large gains from trade at the margin. η enters formula (1) because an inflow of funds and an increase in the expertise of an individual bank affect the equilibrium through the same channel: by increasing the demand for good assets. Therefore η is informative about how many additional trades would take place if a bank increased its expertise at the margin.

I implement formula (1) empirically for two applications: venture capital and junk bond underwriting. For venture capital I rely on existing empirical studies while for junk bonds I rely on a combination of existing studies and new estimates.

Gompers and Lerner (2000) report elasticities of prices and outcomes of venture-backed firms with respect to inflows of capital into venture funds, which can be used to estimate η. Hall and Woodward (2007) estimate how the value of a venture-backed firm is split between founders, venture investors and general partners of venture funds. These estimates can be used to measure α. Both of these studies also report the distribution of outcomes for venture-backed firms, which can be used to get a value of f. Using these empirical estimates, I obtain values of r between between 0.64 and 0.83.

For the junk bond market, I use historical data on default rates in order to measure f. I then use variation in the the volume of new issues, prices and default rates around the time of the collapse of the investment bank Drexel Burnham Lambert in order to estimate η. Finally, I rely on studies of underwriting fees by Datta et al. (1997), Jewell and Livingston (1998) and Gande et al. (1999) to obtain measures of α. I obtain values of r between 0.09 and 0.26.

The estimates imply that out of the last dollar earned by venture capitalists, between 64 and 83 cents is value added and the rest is captured rents. Out of the last dollar earned by junk bond underwriters, between 14 and 26 cents is value added and the rest is captured rents. By these estimates, the venture industry and especially the junk bond underwriting industry are too large relative to the social optimum.

The argument that much of finance involves socially wasteful rent-seeking has been the subject of a large literature, surveyed by Cochrane (2013) and Greenwood and Scharfstein (2013). Bolton et al. (2011), Glode et al. (2012), Shakhnov (2014) and Fishman and Parker (2015) describe theoretical environments where over-investment in financial expertise emerges as an equilibrium outcome. In the context of an endogenous growth model, Philippon (2010) shows that optimal subsidies for innovation may be enough to prevent over-expansion of finance.

The empirical evidence based on aggregate cross-coutry data is somewhat mixed. Murphy et al. (1991) find that the proportions of university graduates in law (negatively) and engineering (positively) are correlated with economic growth, and argue that this roughly correponds to the distinction between financial and productive activities. Levine (1997, 2005) surveys cross country evidence that finds a positive correlation between economic growth and the size of the financial sector.

The paper is organized as follows. Section 2 presents the model, defines and characterizes the equilibrium and derives an expression for r. Section 3 derives the sufficient statistics needed to estimate r and presents the estimates for venture capital and junk bonds. Section 4 discusses the implications of the findings and some of their limitations.

2 The Model

2.1 Agents, Preferences and Technology

The economy is populated by households and banks, all of whom are risk neutral.

Banks are indexed by $j \in [0,1]$. Bank j has an endowment $w(j)$ of goods that it may use to buy assets from households. It is best to think of this endowment as including both the bank's equity and its maximum debt capacity, i.e. the maximum amount of funds it can invest.

Households are indexed by $s \in [0,1]$. Each household is endowed with a single divisible asset $i \in [0,1]$, which it may keep or sell to a bank. The household's type s and the index of

its asset i are independent. If sold to a bank, asset i will produce a dividend of

$$q(i) = \mathbb{I}(i \geq \lambda) V$$

This means a fraction λ of assets are bad and yield 0 and a fraction $1-\lambda$ are good and yield V. If instead household s keeps asset i, it will produce a dividend of $\beta(s) q(i)$. Therefore $(1-\beta(s)) V$ are the gains produced if a household of type s sells a good asset to a bank. Assume w.l.o.g. that $\beta(\cdot)$ is weakly increasing, so higher types get more dividends out of good assets. There is no need to assume that $\beta(s) < 1$ for all s, the model can allow for households for whom there are no gains from trade.

Several applications fit this general framework. In an application to household borrowing, $q(i)$ represents future income and $\beta(s)$ is the household's discount factor. In an application to venture capital, households represent startup companies, banks represent venture capital funds and $\beta(s)$ is the fraction of the startup's potential value that can be realized without obtaining venture financing. In an application to insurance, $q(i)$ is the household's expected income net of any losses and $\beta(s) q(i)$ is its certainty-equivalent.

2.2 Information and Expertise

The household knows the index i of its asset and therefore its quality $q(i)$. Banks do not observe i directly but instead observe signals that depend on their individual expertise. A bank with expertise $\theta \in [0,1]$ will observe a signal

$$x(i,\theta) = \mathbb{I}(i \geq \lambda\theta) V \tag{2}$$

whenever he analyzes asset i, as illustrated in Figure 1.[1] Higher-θ banks are more expert because they make fewer mistakes: they are more likely to observe signals whose value coincides with the true quality of the asset.

The level of expertise θ is endogenously chosen by each bank. The cost for bank j of acquiring expertise θ is given by $c_j(\theta)$. The function $c_j(\cdot)$ is allowed to be different for different banks.

[1]The information structure implied by equation (2) is special in that banks only make mistakes in one direction. Kurlat (2016) analyzes other possible cases.

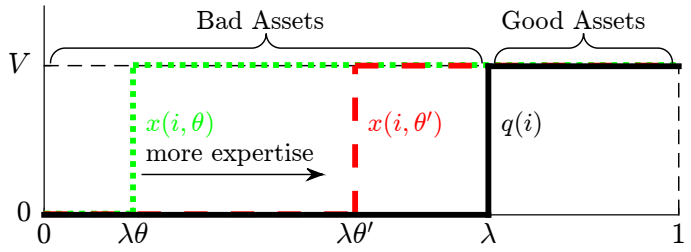


Figure 1: Asset qualities and signals

2.3 Equilibrium Definition

I define equilibrium using the definition of competitive equilibrium definition from Kurlat (2016). Each possible price $p \in [0, V]$ defines a market and any asset can in principle be traded in any market. Markets need not clear: assets that are offered for sale in market p may remain totally or partially unsold.

Households trade by choosing at what prices to put their asset on sale. Markets are non-exclusive: households are allowed to offer their asset for sale at as many prices as they want. This implies that a household of type s who owns asset i will simply choose a reservation price $p^R(i,s)$ and put its asset on sale at every $p \geq p^R(i,s)$ and not at any price below that.[2] From the household's point of view, the only thing that matters about the equilibrium is at what price it's possible to sell its asset i, i.e. the extent to which it will face rationing at each price. Formally, this is captured by a "rationing function" $\mu : [0,V] \times [0,1] \to \mathbb{R}$. $\mu(p,i)$ is the number of assets that a household would end up selling if it offers one unit of asset i on sale with the reservation price p (thereby offering it on sale at every price in $[p,V]$). Implicit in this formulation is the assumption that assets are perfectly divisible, so there is exact pro-rata rationing rather than a probability of selling an indivisible unit.

A household of type s who owns asset i solves:

$$\max_{p^R} \quad \int_{p^R}^{V} p d\mu(p,i) + \left[1 - \mu\left(p^R, i\right)\right] \beta(s)\, q(i) \tag{3}$$

$$s.t. \quad \mu\left(p^R, i\right) \leq 1 \tag{4}$$

The first term in (3) represents the proceeds from selling the asset, possibly fractionally and across many prices. The second term represents the dividends obtained from whatever

[2]There is an extra assumption involved in this. There will be many prices at which it's impossible to sell assets so the household is indifferent between offering its asset on sale in them or not. A reservation price is the only optimal strategy that is robust to a small chance of selling at every price.

fraction of the asset the household retains. Constraint (4) limits the household to not sell more than one unit in total.

This problem as a simple solution. Define

$$p^L(i) \equiv \max\{\inf\{p : \mu(p,i) < 1\}, 0\}$$

$p^L(i)$ is the highest reservation price that a household can set and still be sure to sell its entire asset; if there is no positive price that guarantees selling the entire asset, then $p^L(i) = 0$. It's immediate that the solution to program (3) is:

$$p^R(i,s) = \max\left\{p^L(i), \beta(s) V\right\} \tag{5}$$

If it's possible to sell the entire asset at a price above the household's own valuation, then the household sets the reservation price at the level that guarantees selling; otherwise the reservation price is the household's own valuation.

Turn now to the bank's problem. It has two stages: first the bank chooses a level of expertise and then it trades assets. In the second stage, the bank trades by choosing a quantity δ, a price p and an acceptance rule χ. An acceptance rule is a function $\chi : [0,1] \to \{0,1\}$ from the set of assets to $\{0,1\}$, where $\chi(i) = 1$ means that the bank is willing to accept asset i and $\chi(i) = 0$ means it is not. By trading in market p with acceptance rule χ, the bank obtains χ-acceptable assets in proportion to the quantities that offered on sale at price p. A bank may only impose acceptance rules that are informationally feasible given the expertise it has acquired, so it cannot discriminate between assets that it cannot tell apart, i.e. $\chi(i) = \chi(i')$ whenever $x(i,\theta) = x(i',\theta)$.

From the point of view of banks, the only thing that matters about the equilibrium is what distribution of assets it will obtain for each possible combination of price and acceptance rule it could choose. Formally, this is captured by a measure $A(\cdot;\chi,p)$ on the set of assets $[0,1]$ for each χ, p. For any subset $I \subseteq [0,1]$, $A(I;\chi,p)$ is the measure of assets $i \in I$ that a bank will end up with if it demands one unit at price p with acceptance rule χ.

Therefore in the trading stage, a bank with expertise θ and wealth w solves:

$$\max_{\delta,p,\chi} \quad \delta \left[\int_{[0,1]} q(i)\, dA(i;\chi,p) - pA([0,1];\chi,p) \right] \tag{6}$$

$$s.t. \quad \delta p A([0,1];\chi,p) \leq w \tag{7}$$

$$\chi(i) = \chi(i') \quad \text{whenever} \quad x(i,\theta) = x(i',\theta) \tag{8}$$

(6) adds all the dividends $q(i)$ of the assets the bank buys, subtracts what it pays per unit and multiplies by total demand δ; (7) is the budget constraint and (8) imposes that the bank use an informationally feasible acceptance rule.

Notice that w enters the problem only in the budget constraint, which is linear. This implies that δ will be linear in w and p and χ will not depend on w. Let $\delta(\theta)$, $p(\theta)$ and $\chi(\theta)$ denote the solution to the bank's problem for a bank with $w = 1$ and expertise θ, and let $\tau(\theta)$ be the maximized value of (6) for $w = 1$.

The first stage of the bank's problem is straightforward. Bank j chooses expertise $\theta(j)$ by solving:[3]

$$\max_{\theta} w(j)\, \tau(\theta) - c_j(\theta) \tag{9}$$

Let $W(\theta)$ be the total wealth of banks that choose expertise at most θ, i.e.

$$W(\theta) \equiv \int w(j)\, \mathbb{I}(\theta(j) \leq \theta)\, dj \tag{10}$$

and let $w(\theta) \equiv \frac{\partial W(\theta)}{\partial \theta}$. Nothing depends on $W(\theta)$ being differentiable but it simplifies the exposition.

The two key equilibrium objects are the rationing function $\mu(p,i)$ and the allocation measures $A(\cdot;\chi,p)$. Informally, A is consistent with equilibrium if, for any χ, p, the distribution $A(\cdot;\chi,p)$ is a representative sample of the χ-acceptable assets that are on sale at price p. μ is consistent with equilibrium if, for any i, p, $\mu(i,p)$ is equal to the total fraction of supply that is bought by banks who each buy representative samples. The Appendix spells this out formally.

I define equilibrium in two steps. First I define a conditional equilibrium, i.e. an equilibrium given the first-stage choices by banks that result in $W(\theta)$.

Definition 1. Taking $W(\theta)$ as given, a conditional equilibrium is given by reservation

[3]For simplicity, the cost $c_j(\theta)$ is expressed directly in utility terms and does not enter (7).

prices $p^R(i,s)$, buying plans $\{\delta(\theta), p(\theta), \chi(\theta)\}$, rationing measures $\mu(\cdot; i)$ and allocation measures $A(\cdot; \chi, p)$ such that: $p^R(i,s)$ solves the household's problem for all i, s, taking $\mu(\cdot, i)$ as given; $\{\delta(\theta), p(\theta), \chi(\theta)\}$ solves the bank's second stage problem for all θ, taking $A(\cdot; \chi, p)$ as given and $\mu(\cdot; i)$ and $A(\cdot; \chi, p)$ satisfy the consistency conditions (40) and (41) (derived in the Appendix).

Using this, I now define a full equilibrium. The usefulness of this two-step definition is that it is possible to focus on characterizing the conditional equilibrium without fully specifying the cost functions c_j that govern the banks' first-stage decisions.

Definition 2. An equilibrium is given by expertise choices $\theta(j)$, a wealth distribution $W(\theta)$ and a conditional equilibrium $\left\{p^R, \delta, p, \chi, \mu, A\right\}$ such that: $\theta(j)$ solves the bank's first stage problem for all j, taking the conditional equilibrium as given; $W(\theta)$ is defined by (10) and $\left\{p^R, \delta, p, \chi, \mu, A\right\}$ is a conditional equilibrium given $W(\theta)$.

2.4 Equilibrium Characterization

Taking $W(\theta)$ as given, let p^*, θ^* and s^* be the highest-p^* solution to the following system of equations:

$$p^* = \beta(s^*) V \tag{11}$$

$$p^* = \frac{s^*(1-\lambda)}{s^*(1-\lambda) + \lambda(1-\theta^*)} V \tag{12}$$

$$p^* = \int_{\theta^*}^{1} \frac{1}{s^*(1-\lambda) + \lambda(1-\theta)} dW(\theta) \tag{13}$$

Furthermore, assume the following:

Assumption 1. $\frac{1}{p} \frac{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)}{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta^*)} V < 1$ *for all* $p > p^*$

The role of Assumption 1 is discussed below.

Proposition 1. *If Assumption 1 holds, there is a unique conditional equilibrium, where:*

1. Reservation prices are:

$$p^R(i,s) = \begin{cases} \max\{p^*, \beta(s) V\} & \text{if } i \geq \lambda \\ 0 & \text{if } i < \lambda \end{cases} \tag{14}$$

2. The solution to the banks' problem is:

$$\{\delta(\theta), p(\theta), \chi(\theta)\} = \begin{cases} \left\{\frac{1}{p^*}, p^*, \mathbb{I}(i \geq \lambda\theta)\right\} & \text{if } \theta \geq \theta^* \\ \{0,0,0\} & \text{if } \theta < \theta^* \end{cases} \tag{15}$$

3. At price p^*, bank θ, who applies acceptance rule $\chi(\theta) = \mathbb{I}(i \geq \lambda\theta)$, obtains the following density over assets:

$$a(i; \chi(\theta), p^*) = \begin{cases} \frac{s^*}{\lambda(1-\theta)+s^*(1-\lambda)} & \text{if } i \geq \lambda \\ \frac{1}{\lambda(1-\theta)+s^*(1-\lambda)} & \text{if } i \in [\lambda\theta, \lambda) \\ 0 & \text{if } i < \lambda\theta \end{cases} \tag{16}$$

4. The rationing function at price p^* is:

$$\mu(p^*, i) = \begin{cases} 1 & \text{if } i \geq \lambda \\ \int_{\theta^*}^{\frac{i}{\lambda}} \frac{1}{\lambda(1-\theta)+s^*(1-\lambda)} \frac{1}{p^*} dW(\theta) & \text{if } i \in [\lambda\theta, \lambda) \\ 0 & \text{if } i < \lambda\theta \end{cases} \tag{17}$$

See the Appendix for a full statement of the equilibrium objects (in particular A and μ at other prices).

In equilibrium, all trades take place at the same price p^*. Condition (17) says that households who offer good assets on sale at p^* are able to sell them. Therefore they set their reservation price according to (14): the highest of either their valuation $\beta(s)V$ or the price p^* at which they know they will be able to sell the asset. This defines a cutoff type s^* who is just indifferent between selling the asset or keeping it (equation (11)). Conversely, condition (17) says that households who own a bad asset cannot sell all of it at p^*; since they don't value it at all, they set a reservation price of 0 and offer it on sale at every price.

A bank who buys at price p^* faces a supply which consists of 1 unit of each $i \in [0, \lambda)$ and s^* units of each $i \in [\lambda, 1]$. If it has expertise θ it will impose the acceptance rule $\mathbb{I}(i \geq \lambda\theta)$ (i.e. only accept assets for which it observes a good signal). This filters out some, but not all, the bad assets. Hence it will obtain assets distributed according to (16). This implies it will obtain a surplus of

$$\tau(\theta) = \frac{1}{p^*}\left[\frac{s^*(1-\lambda)}{s^*(1-\lambda)+\lambda(1-\theta)}V - p^*\right] \tag{18}$$

per unit of wealth that it dedicates to buying assets. Notice that $\tau(\theta)$ is increasing in θ. More expert banks are able to filter out more bad assets and therefore obtain higher returns. There is a cutoff value θ^* such that $\tau(\theta)$ is positive if and only if $\theta > \theta^*$. Rearranging leads to equation (12). Banks with expertise above θ^* spend all their wealth buying assets while banks with expertise below θ^* choose not to buy at all. This gives equation (15).

Banks also have the option to buy assets at prices other than p^*. Buying at lower prices is clearly worse than buying at p^* because the reservation price for good assets is at least p^* so no good assets are on sale at lower prices. Assumption 1 ensures that buying at higher price is not preferred either. Given the reservation prices (14), the surplus per unit of wealth for bank θ^* if it buys at price $p > p^*$ is:

$$\frac{1}{p}\left[\frac{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)V}{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta^*)}-p\right]$$

In principle, the bank faces a tradeoff: better selection (because β^{-1} is an increasing function) but a higher price. Assumption 1 ensures that the direct higher-price effect dominates and a bank with expertise θ^* has no incentive to pay higher prices to ensure better selecion. It is then possible to show that if this is true for the marginal bank θ^*, it is true for all banks: higher-θ banks care even less about selection because they can filter assets themselves and lower-θ banks can never earn surplus in a market where θ^* would not. One can still solve for equilibria where Assumption 1 does not hold, but they are somewhat more complicated. Wilson (1980), Stiglitz and Weiss (1981) and Arnold and Riley (2009) analyze the implications of models where an analogue of Assumption 1 doesn't hold.

Condition (13) is a market clearing condition. The total supply of good assets is $s^*(1-\lambda)$. Equation (16) implies that, per unit of wealth, a bank with expertise θ obtains $\frac{1}{p^*}\frac{s^*(1-\lambda)}{s^*(1-\lambda)+\lambda(1-\theta^*)}$ good assets. Adding up across all banks and imposing that all good assets end up being sold results in (13).

Recall that market clearing is not imposed as an equilibrium condition. Indeed, (17) implies that the market for good assets at price p^* clears but that for bad assets does not. How do we know that the market for good assets must clear? If it didn't, (5) implies that households $s < s^*$ who own good assets would choose a lower reservation price. But since banks would still find it optimal to buy assets, these would run out, violating condition (4).

Assets $i < \lambda$ will not be accepted by all banks. Assets $i < \lambda\theta^*$ are rejected by all banks that choose to buy while assets $i \in [\lambda\theta, \lambda)$ will be accepted by some banks but not others. The fraction of each asset that will be sold in equilibrium depends on how many units are

bought by banks willing to accept them. This gives equation (17).

2.5 Welfare

I measure welfare as the total surplus that is generated by trading assets, ignoring the distribution of gains. When a household of type s sells a good asset, this creates $(1-\beta(s))V$ social surplus. Integrating over all households that sell yields a total surplus of:

$$S = (1-\lambda)\int_0^{s^*}(1-\beta(s))\,V ds \tag{19}$$

Consider an individual bank j that in equilibrium chooses to acquire expertise θ_j. Holding the expertise choices of all other banks constant, let $S_j(\theta)$ be the social surplus that would result if instead bank j were to acquire expertise θ. Define

$$r_j \equiv \frac{S_j'(\theta_j)}{w(j)\,\tau'(\theta_j)} \tag{20}$$

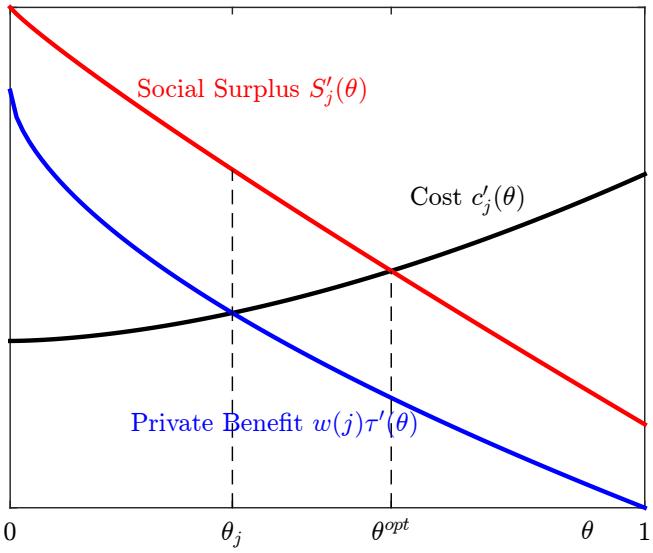


Figure 2: Example of marginal social surplus, private benefit and cost of additional investments in expertise. Bank j will choose expertise θ_j, equating marginal private benefit and marginal cost. The socially optimal level of expertise would be θ^{opt}.

Why is r_j an object of interest? The logic is illustrated in Figure 2. The first order

condition for problem (9) is:

$$w(j)\tau'(\theta_j) = c_j'(\theta_j)$$

and therefore

$$r_j = \frac{S_j'(\theta_j)}{c_j'(\theta_j)}$$

Hence r_j is a measure of the amount of value created per unit of marginal resources that bank j invests in acquiring expertise. In the example in Figure 2, at the equilibrium level of expertise θ_j, we have $S_j'(\theta_j) > w(j)\tau'(\theta_j)$ so $r_j > 1$, which means that at the margin investing more in expertise increases the net social surplus.

It is worth noting that if it were possible to redistribute banks' endowments, then investing in expertise would always be socially wasteful. Rather than having many banks invest independently in acquiring the same expertise, the efficient thing to do would be to have a single bank acquire expertise and manage everyone's endowment. The maintained assumption is that for unmodeled moral hazard or span-of-control reasons this is not possible. Studying r_j answers the question of what is the marginal social value of investments in expertise taking as given the duplicative nature of these investments.

3 Estimating r

3.1 Solving for r_j

Using (19), the marginal social surplus is

$$S_j'(\theta) = (1-\lambda)(1-\beta(s^*))V\frac{ds^*}{d\theta_j} \tag{21}$$

A change in bank j's expertise increases the social surplus if the change in the equilibrium that it brings about induces marginal households to sell their asset, creating gains from trade. In equation (21), $(1-\lambda)(1-\beta(s^*))V$ are the gains from trade by the marginal household s^* and $\frac{ds^*}{d\theta_j}$ is the shift in s^* when bank j increases its expertise.

Using (18), private marginal utility is

$$w(j)\tau'(\theta_j) = \frac{w(j)}{p^*}V\frac{\lambda(1-\lambda)s^*}{[(1-\lambda)s^* + \lambda(1-\theta_j)]^2} \tag{22}$$

In formula (22), $\frac{w(j)}{p^*}$ is the number of assets the bank can afford to acquire, V is the value of each good asset and $\frac{\lambda(1-\lambda)s^*}{[(1-\lambda)s^*+\lambda(1-\theta_j)]^2}$ is how the fraction of good assets in the bank's portfolio changes when the bank acquires additional expertise.

Replacing (21) and (22) in (20):

$$r_j = \frac{(1-\lambda)\left(1-\beta\left(s^*\right)\right)V}{w\left(j\right)\frac{V}{p^*}\frac{\lambda(1-\lambda)s^*}{[(1-\lambda)s^*+\lambda(1-\theta)]^2}}\frac{ds^*}{d\theta_j} \tag{23}$$

A key ingredient of equation (23) is $\frac{ds^*}{d\theta_j}$, how many additional households sell good assets when the expertise of bank j changes. In order to compute this, rewrite equations (11)-(13) compactly as:

$$K\left(p^*, \theta^*, s^*\right) = 0 \tag{24}$$

where

$$K\left(p^*, \theta^*, s^*\right) = \begin{pmatrix} p^* - \beta\left(s^*\right)V \\ p^* - \frac{(1-\lambda)s^*}{(1-\lambda)s^*+\lambda(1-\theta^*)}V \\ p^* - \int_{\theta^*}^{1}\frac{1}{(1-\lambda)s^*+\lambda(1-\theta)}dW\left(\theta\right) \end{pmatrix}$$

Let K_i denote the i_{th} dimension of the function K and $D = \nabla K$ denote the matrix of derivatives of K.

Using the implicit function theorem, (24) implies:

$$\frac{ds^*}{d\theta_j} = -D_{33}^{-1}\frac{\partial K_3}{\partial \theta_j} \tag{25}$$

where

$$D_{33}^{-1} = -\frac{1}{|D|}\frac{\lambda\left(1-\lambda\right)s^*}{\left[\left(1-\lambda\right)s^* + \lambda\left(1-\theta^*\right)\right]^2}V \tag{26}$$

$$|D| = \frac{V}{\left[\left(1-\lambda\right)s^* + \lambda\left(1-\theta^*\right)\right]^2}\begin{bmatrix} -\frac{\lambda(1-\lambda)(1-\theta^*)}{(1-\lambda)s^*+\lambda(1-\theta^*)}w\left(\theta^*\right) \\ +\left[\lambda\left(1-\lambda\right)s^*V + \left(\left(1-\lambda\right)s^* + \lambda\left(1-\theta^*\right)\right)w\left(\theta^*\right)\right]\beta'\left(s^*\right) \\ +\lambda\left(1-\lambda\right)s^*\int_{\theta^*}^{1}\frac{(1-\lambda)}{[(1-\lambda)s^*+\lambda(1-\theta)]^2}dW\left(\theta\right) \end{bmatrix} \tag{27}$$

$$\frac{\partial K_3^*}{\partial \theta} = -w_j\frac{\lambda}{\left[\left(1-\lambda\right)s^* + \lambda\left(1-\theta\right)\right]^2} \tag{28}$$

Equation (28) captures the direct effect of an increase in bank j's expertise. More expertise implies rejecting more bad assets and therefore buying more good assets. This shifts

the market clearing condition. Other things being equal, prices would have to rise to restore equation (13). But, of course, all the endogenous variables respond: higher prices attract marginal sellers of good assets and repel marginal banks, so both s^* and θ^* respond as well. The term D_{33}^{-1} measures how shifts in the market clearing condition translate, through all the feedback channels in the model, into a change in the marginal seller. Equation (28) implies this is always positive: more expert banks lead to a higher equilibrium price and this induces marginal households to sell good assets.

Replacing equations (25)-(28) into equation (23) and simplifying:

$$r_j = \frac{1}{|D|} \frac{\lambda (1-\lambda) (1-\beta (s^*)) p^* V}{\left[(1-\lambda) s^* + \lambda (1-\theta^*)\right]^2} \tag{29}$$

Formula (29) immediately implies the following result.

Proposition 2. *r_j does not depend on θ_j or w_j*

One might have conjectured that the misalignment of social and private returns to expertise might be different for banks with different wealth or for banks that (for instance due to different cost functions) choose different levels of θ. That turns out not to be the case. This means that if the financial industry has incentives to either over- or under-invest in expertise, this will be true across the board, and any corrective policies don't need to be applied selectively.

The main difficulty with estimating (29) is that the expression for the determinant $|D|$ is quite complicated. This is because $|D|$ captures the magnitude of all the various feedback effects in the model: how selection depends on prices, the extensive margin of bank participation, etc. The key to the sufficient statistic approach is that it is not necessary to estimate all the elements of $|D|$ separately. $|D|$ measures the strength of feedback effects with respect to any driving force; therefore it enters the formula for any elasticity that one could measure.

3.2 Sufficient Statistics

Let α be the average net present value per dollar invested that banks obtain. In the model:

$$\alpha = \frac{(1-\lambda) s^* V}{\int_{\theta^*}^1 dW(\theta)} \tag{30}$$

The numerator represents the total dividends obtained from assets acquired by banks and the denominator is the total funds they spend.

Let f be the fraction of assets traded that turn out to be bad. In consumer loans, this would correspond to the default rate; in venture capital it would correspond to the fraction of ventures that fail, etc. If N is the total number of assets that are traded and G is the number of good assets that are traded, then:

$$f \equiv 1 - \frac{G}{N}$$

In the model we have:

$$N = \frac{\int_{\theta^*}^{1} dW(\theta)}{p^*} \tag{31}$$

$$G = (1-\lambda)\, s^* \tag{32}$$

The numerator in (31) is the total funds spent by banks who choose to trade and the denominator is the price they pay per asset. Therefore:

$$f = 1 - \frac{(1-\lambda)\, s^* p^*}{\int_{\theta^*}^{1} dW(\theta)} \tag{33}$$

Notice that measuring f only requires tracking failures among assets that actually trade, not among all projects, which would be harder to measure. It is not necessary, for instance, to measure counterfactual default rates among applicants that are denied credit.

Suppose there is an exogenous capital inflow into banks that increases all banks' endowments by Δ, from $w(j)$ to $(1+\Delta)\, w(j)$. For instance, this could be the result of a relaxation in leverage limits that lets banks manage larger portfolios with the same net worth. According to the model, the elasticity of G with respect to this increase is

$$\begin{aligned}
\eta &\equiv \frac{d\log(G)}{d\Delta} \\
&= \frac{d\log(s^*)}{d\Delta} \\
&= -D_{33}^{-1}\frac{\partial K_3}{\partial \Delta}\frac{1}{s^*} \\
&= \frac{1}{|D|}\frac{\lambda(1-\lambda)}{\left[(1-\lambda)\, s^* + \lambda(1-\theta^*)\right]^2} p^* V
\end{aligned} \tag{34}$$

Replacing (30), (33) and (34) into (29) and rearranging results in equation (1):

$$r = \eta \left(1 - \frac{1-f}{\alpha}\right)$$

Formula (1) has the following interpretation. If $\frac{1-f}{\alpha}$ is low, this means that banks obtain high returns despite the fact that only a small fraction of the assets they buy are good. For this to be true it must be that $\frac{p^*}{V}$ is low, i.e. they must be making very high profits on the good assets that they do buy, which means that the marginal household s^* is preventing large gains from trade by not selling.[4] When this is the case, marginal trades create high surplus, which makes r large. η enters the formula because it is a way to measure the strength of the extensive margin $\frac{ds^*}{d\theta}$. An increase in the expertise of one bank affects the equilibrium through the same channel than an inflow of funds for all banks: through the market clearing condition (13). An inflow of funds means that the more expert banks can afford to buy more assets; prices must rise to restore equilibrium and s^* responds to this. An increase in expertise means that the same bank will reject more bad assets and therefore buy more good ones. Again, prices must rise to restore equilibrium and s^* responds. Both effects involve the same mechanism and the same feedback channels.

The quantities α and f can be measured relatively straightforwardly because they are simple averages. η is more challenging because it requires identifying a plausibly exogenous capital inflow or outflow and measuring its consequences. If such identifying assumptions are satisfied, there are a few different ways to measure η depending on what outcomes are easier to measure. The first, if the number of good assets traded can be measured, is simply to measure $\eta = \frac{d\log(G)}{d\Delta}$ directly. The second is almost as simple: if one can measure total number of assets traded and failure rates, then relying on (33) one gets:

$$\eta = \frac{d\log(1-f)}{d\Delta} + \frac{d\log N}{d\Delta} \tag{35}$$

A third option, if one measures failure rates, prices and total funds invested, is to use (31) to further decompose:

$$\eta = \frac{d\log(1-f)}{d\Delta} + \frac{d\log\left(\int_{\theta^*}^1 dW(\theta)\right)}{d\Delta} - \frac{d\log(p^*)}{d\Delta} \tag{36}$$

In all cases, measuring elasticities with respect to Δ requires measuring Δ itself, i.e. how

[4] This is the standard type of inefficiency in models based on Akerlof (1970)

much banks' endowments change. In some cases it might be possible to do this directly, for instance if there is an increase in leverage limits that expands maximum balance sheets by a known factor. In other cases one might have to rely on measured changes in the the total number of funds actually invested in buying assets, which is not exactly the same. One of the things that can happen when Δ increases is that, because prices rise, marginal banks exit. Therefore the measured proportional change in total funds spent buying assets could be an underestimate of Δ. Formally:

$$\frac{d\log\left(\int_{\theta^*}^{1} dW\left(\theta\right)\right)}{d\Delta} = 1 - \frac{\frac{d\theta^*}{d\Delta} w\left(\theta^*\right)}{\int_{\theta^*}^{1} dW\left(\theta\right)} \leq 1$$

However, it is not unreasonable to assume that $w\left(\theta^*\right) = 0$. Choosing $\theta = \theta^*$ means that a bank would earn $\tau\left(\theta^*\right) = 0$ despite having invested a strictly positive amount of resources in acquiring expertise. Assuming $w\left(\theta^*\right) = 0$ means assuming that no banks choose to do this. Under this assumption, measuring an elasticity with respect to measured capital flows and with respect to Δ is equivalent, i.e.

$$\frac{d\log\left(\int_{\theta^*}^{1} dW\left(\theta\right)\right)}{d\Delta} = 1 \tag{37}$$

and therefore $d\Delta$ and be replaced with $d\log\left(\int_{\theta^*}^{1} dW\left(\theta\right)\right)$ in formulas (35) or (36).

3.3 Application to Venture Capital

Hall and Woodward (2007) estimate how the value of venture-backed firms is, on average, split between the firm's founders and the general and limited partners of venture funds. I map these participants to the model as follows. The firm's founders are like the households in the model. They own an asset (the firm) and there are possible gains from trade in transferring part of the ownership of the firm to the venture fund. The general partners of venture funds are like the banks in the model. They have expertise in determining which firms are valuable. The limited partners are absent from the model. Hall and Woodward find that limited partners, who provide capital to venture funds but are not directly involved in decision-making, get almost no risk-adjusted excess returns from venture investments. I assume that general partners commit to deliver zero excess returns to limited partners and keep all excess returns for themselves in the form of fees. If this is true, the incentives to acquire expertise are proportional to the capital that the general partners administer. Hence

$w(j)$ in the model corresponds to the total capital administered by a venture fund, including the capital supplied by limited partners. Hall and Woodward find that general partners, on average, earn 26% of funds invested. This suggests a value of $\alpha = 1.26$. This is probably an upper bound on α (and therefore an upper bound on r) since the rewards to venture capitalists compensate them for other services the provide firms besides screening them.

Gompers and Lerner (2000) estimate the elasticity of valuations for venture investments with respect to inflows of capital into venture funds. They estimate $\frac{\partial \log(p^*)}{\partial \Delta} \in [0.12, 0.22]$ depending on the specification used. They don't report estimates of $\frac{\partial \log(1-f)}{\partial \Delta}$ directly but it's possible to reconstruct them on the basis of the time series of f that they do report. Based on this data, $\frac{\partial \log(1-f)}{\partial \Delta} \in [0.11, 0.21]$ depending on the exact definition of a successful venture that is used. Using these estimates, in formula (36) and assuming $w(\theta^*) = 0$ so that (37) holds, we can assign a value of $\eta \in [0.89, 1.14]$.

Gompers and Lerner's estimates are based on exploiting time-series variation in inflows to venture funds, which raises questions about identification. Possibly, funds flow into venture funds attracted by better prospects for firms, which leads to higher prices and lower failure rates. Gompers and Lerner control for the most plausible channels of reverse-causality by including measures of stock market valuation as controls and by using inflows into leveraged buyout funds as instruments. Furthermore, they argue that regulatory changes like the clarification of the "prudent man" rule that allowed pension funds to invest in venture capital and changes in the capital gains tax rate account for much of the variation. Still, it's possible that the estimates of elasticity have omitted variable bias. This would bias both $\frac{\partial \log(p^*)}{\partial \Delta}$ and $\frac{\partial \log(1-f)}{\partial \Delta}$ upwards, with an uncertain net effect on η.

Asset payoffs in the model are binary, either 0 or V. Payoffs from venture-backed firms are far from binary. Many fail and pay close to zero while among the successful ones there is a long right tail of extremely successful ones. This can be reconciled with the binary-payoff model by assuming that the value of successful firms is a random variable $\tilde{V}$ with expected value V. If we assume that entrepreneurs are not privately informed about the realization of $\tilde{V}$, then the fact that it's random makes no difference.

The question remains of how to measure f (the fraction of outright failures) empirically. Both Hall and Woodward and Gompers and Lerner discuss this issue. Gompers and Lerner propose using the failure to either conduct an IPO or be acquired at twice the original valuation as a definition of failure (that definition is implicitly used in the measured elasticity above). Under this definition, in their data, $f = 0.66$. Hall and Woodward report similar figures. In their sample, the fraction of venture-backed firms that have not been acquired nor

undergone an IPO is $f = 0.65$. This is not fully satisfactory since some firms in the sample will conduct IPOs or be acquired later on, or simply continue as privately held firms and produce positive (though rarely large) dividends. Because of this, these estimates should probably be regarded as an upper bound on f (and therefore an upper bound on r).

Replacing the range of empirical estimates of α, η and f into formula (1) gives range of r between 0.64 to 0.83. This means that for the last dollar that general partners of venture funds earn by being good at selecting which firms to invest in, between 64 and 83 cents are value added and the remainder is captured rents. Compared to the social optimum, the venture capital industry is too large.

3.4 Application to Junk Bond Underwriting

I map the junk bond market to the model following the "certification" view of underwriting proposed by Booth and Smith (1986). The company issuing the bonds is like the households in the model and the asset is a stream of cashflows. Investment banks that underwrite bonds are like the banks in the model. In exchange for a fee, they certify that they observed a good signal from a bond, after which it is acquired by ordinary investors (not in the model), who make zero profits.

Gande et al. (1999) report underwriting fees averaging 2.76% for bonds rated between Caa and Ba3 between 1985 and 1996; Jewell and Livingston (1998) rerport similar figures. Furthermore, Datta et al. (1997) report an average initial-day return of 1.86% for low-grade bonds. Arguably, this is also part of the underwriter's compensation since it allows the underwriter to place the bonds with favored clients or bolster its reputation. Accordingly, I add these two fees and set $\alpha = 1.046$.

In order to obtain measures of η and f, I focus on the period around 1990. The reason for this is that the investment bank Drexel Burnham Lambert filed for bankruptcy in Februrary 1990 following an SEC investigation for various forms of wrongdoing. Drexel was a major participant in the junk bond market, with a market share above 40%, and its demise had a major impact on the market (Brewer and Jackson 2000). I exploit the variation in volumes of bonds issued, bond prices and default rates around 1990 in order to obtain an estimate of η. Clearly, this is not ideal because the bankruptcy of Drexel is not the only thing that happened around that time (the economy was undergoing a recession) and it's also not exogenous to other developments in the junk bond market. This could in principle bias the estimate of r in either direction. The main channel of reverse causation is likely to be from a fall in the qualities of bond issuers to a fall in volume and price. This would mean that

both the elasticity of $1-f$ and the elasticity of p to capital flows are overestimated, with an uncertain net effect on η and r.

My baseline sample includes (subject to data availability) all the corporate bonds denominated in US dollars, issued between 1987 and 1990 and rated below investment grade by either S&P, Moody's or Fitch, a total of 585 individual bonds. The source is the Bloomberg database. For each bond I observe the total dollar amount issued. its coupon rate, its maturity, the yield spread against treasuries of comparable maturity and whether it subsequently defaulted.[5] I measure f simply as the dollar-weighted fraction of bonds that defaulted, and obtain $f = 0.09$, which is somewhat lower than the numbers reported in previous studies (Altman 1989, 1992, Asquith et al. 1989, McDonald and Van de Gucht 1999, Zhou 2001).

I estimate η using formula (36) and assuming $w(\theta^*) = 0$. I normalize the cash flows of each bond by dscounting the coupon and principal payments at treasury rate of the corresponding maturity at the time of issuance. I then compute a normalized price for each bond by discounting the same cash flows at the bond's actual yield constructing by adding the bond's spread to the treasury rate. Then, for each year t of the sample, I compute p_t as the dollar-weighted average p, f_t as the dollar-weighted fraction of bonds that default and $\int_{\theta^*}^{1} dW_t(\theta)$ by adding the dollar amount of all the bonds issued. I then separately regress $\log(p_t)$ and $\log(f_t)$ on $\log\left(\int_{\theta^*}^{1} dW_t(\theta)\right)$ to obtain elasticities and apply formula (36) to obtain an estimate of η. I find $\frac{\partial \log(p)}{\partial \Delta} = 0.004$, $\frac{\partial \log(1-f)}{\partial \Delta} = 0.14$, which results in $\eta = 1.13$

Replacing the estimates of α, η and f into formula (1) gives $r = 0.15$. Different estimation windows around 1990 produce estimates of $r \in [0.09, 0.26]$. This means that out of the last dollar that junk bond underwriters earn by being good at certifying the quality of bond issuers, between 9 and 26 cents are value added and the remainder is captured rents. Compared to the social optimum, the junk bond underwriting industry is too large, and the wedge between the private and social value is estimated to be quite large. The reason for this is that, given relatively low values of f and α, the value of the marginal trade is not estimated to be very large. This could be because junk bond issuers have alternative sources of funding (for instance, bank loans) or because they are close to indifferent between obtaining financing or not.

[5] I don't observe all of these measures for all the bonds. In particular, data on spreads is missing for many of them, so I exclude them from measures of p, though not from measures of total volume.

4 Discussion

The method I use to measure r has both advantages and limitations, some of which have to do with the method itself and others with the particular applications.

One advantage is that it does not require estimating or making assumptions about the nature of the cost function $c_j(\theta)$ ("how many physicists with PhDs does it take to value a mortgage-backed security?"). Simply assuming that θ is chosen optimally makes it possible to sidestep this question. Another advantage, common to methods based on sufficient statistics, is that the ingredients of r can be estimated without estimating all the structural parameters of the model. Chetty (2008) offers a discussion of this type of approach.

One disdvantage, also common to sufficient statistics methods, is that r is a purely local measure at the equilibrium. If some policy were to result in a different equilibrium, then r at the new equilibrium might be different. If one wanted calculate the optimal rate of a simple Pigouvian tax to align private and social incentives it would be necessary to know r at the new equilbrium rather than at the original equilibrium.

Another limitation is that r measures the size of the wedge between $S_j'(\theta_j)$ and $w(j)\tau'(\theta_j)$ but not the distance between the equilibrium θ_j and the social optimum θ^{opt} in Figure 2. In order to assess this, it would be necessary to know more about the cost function. For instance, if the marginal cost of expertise increased very steeply, then even a large wedge between r and 1 would imply a small difference between θ_j and θ^{opt}.

In interpreting the estimates of r, it's important to bear in mind that evaluating trades in environments with asymmetric information is just one of the many things that financial firms do. Therefore the measured r is informative about the net social value of dedicating resources to these types of activities within finance and not necessarily about the industry as a whole. Indeed, the method for estimating r could be applied to businesses that are not usually classified as finance but also involve expertise for trading under asymmetric information, such as used car dealerships.

Within the literature on venture capital, there is some debate about whether asymmetric information is a major issue at all. Of course, the estimates of r for venture capital only make sense if one takes the view that indeed asymmetric information prevents gains from trade. In particular, one must believe that there are entrepreneurs with good projects who refuse to seek venture capital financing (or choose not to become entrepreneurs at all) because venture capital funds offer terms that are too onerous. Gompers (1995), Amit et al. (1998) and Ueda (2004) find evidence consistent with models in which this sort of effect is present.

The typical venture transaction differs from the simple outright sales that take place in

the model: the venture capitalist's funds are invested in the firm rather that paid in cash to the founders. This is an important distiction but it need not change the basic force at play: venture capitalists demand a higher stake in the companies they finance than they otherwise would in order to compensate for investing in the firms that end up failing.

A maintained assumption is that $(1-\beta(s))V$ represents the social value of the gains from trade. If the trade itself generates externalities then the social gains from trade should be adjusted accordingly. A firm that expands thanks to venture capital financing could generate positive externalities through technological spillovers or negative ones through business-stealing. A firm that finances a buyout by issuing junk bonds could bring about new management techniques that other firms learn from or could be destroying value to take advantage of tax benefits. Taking this into account could make the social value of financial expertise higher or lower than estimated.

Another assumption in the model is that trading bad assets neither produces nor destroys social value. If giving funding to bad firms means wasting resources then these trades destroy social value. Fishman and Parker (2015) analyze a model of strategic information acquisition where this is an important effect. Instead, if funding accelerates the development and thus the failure of a business model, this can liberate resources and the trades create social value. Since an increase in expertise leads to fewer trades of bad assets, making either of these alternative assumptions would result in a different estimate of the social value of expertise.

References

Akerlof, G. A.: 1970, The market for "lemons": Quality uncertainty and the market mechanism, *The Quarterly Journal of Economics* **84**(3), 488–500.

Altman, E. I.: 1989, Measuring corporate bond mortality and performance, *The Journal of Finance* **44**(4), 909–922.

Altman, E. I.: 1992, Revisiting the high-yield bond market, *Financial Management* pp. 78–92.

Amit, R., Brander, J. and Zott, C.: 1998, Why do venture capital firms exist? theory and canadian evidence, *Journal of business Venturing* **13**(6), 441–466.

Arnold, L. G. and Riley, J. G.: 2009, On the possibility of credit rationing in the stiglitz-weiss model, *American Economic Review* **99**(5), 2012–21.

Asquith, P., Mullins, D. W. and Wolff, E. D.: 1989, Original issue high yield bonds: Aging analyses of defaults, exchanges, and calls, *The Journal of Finance* **44**(4), 923–952.

Bolton, P., Santos, T. and Scheinkman, J. A.: 2011, Cream skimming in financial markets. NBER working paper No. 16804.

Booth, J. R. and Smith, II, R. L.: 1986, Capital raising, underwriting and the certification hypothesis, *Journal of Financial Economics* **15**(1-2), 261 – 281.

Brewer, E. and Jackson, W. E.: 2000, Requiem for a market maker: the case of drexel burnham lambert and junk bonds, *Journal of Financial Services Research* **17**(3), 209–235.

Chetty, R.: 2008, Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. NBER working paper No. 14399.

Cochrane, J. H.: 2013, Finance: Function Matters, Not Size, *Journal of Economic Perspectives* **27**(2), 29–50.

Datta, S., Iskandar-Datta, M. and Patel, A.: 1997, The pricing of initial public offers of corporate straight debt, *the Journal of Finance* **52**(1), 379–396.

Fishman, M. J. and Parker, J. A.: 2015, Valuation, adverse selection, and market collapses, *Review of Financial Studies* .

Gale, D.: 1996, Equilibria and pareto optima of markets with adverse selection, *Economic Theory* **7**(2), 207–235.

Gande, A., Puri, M. and Saunders, A.: 1999, Bank entry, competition, and the market for corporate securities underwriting, *Journal of Financial Economics* **54**(2), 165–195.

Glode, V., Green, R. C. and Lowery, R.: 2012, Financial expertise as an arms race, *The Journal of Finance* **67**(5), 1723–1759.

Gompers, P. A.: 1995, Optimal investment, monitoring, and the staging of venture capital, *The Journal of Finance* **50**(5), 1461–1489.

Gompers, P. and Lerner, J.: 2000, Money chasing deals? the impact of fund inflows on private equity valuation, *Journal of financial economics* **55**(2), 281–325.

Greenwood, R. and Scharfstein, D.: 2013, The growth of finance, *The Journal of Economic Perspectives* pp. 3–28.

Guerrieri, V., Shimer, R. and Wright, R.: 2010, Adverse selection in competitive search equilibrium, *Econometrica* **78**(6), 1823–1862.

Hall, R. E. and Woodward, S. E.: 2007, The incentives to start new companies: Evidence from venture capital. NBER working paper No. 13056.

Jewell, J. and Livingston, M.: 1998, Split ratings, bond yields, and underwriter spreads, *Journal of Financial Research* **21**(2), 185–204.

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

Levine, R.: 1997, Financial development and economic growth: Views and agenda, *Journal of Economic Literature* **35**(2), 688–726.

Levine, R.: 2005, Finance and growth: theory and evidence, *Handbook of economic growth* **1**, 865–934.

McDonald, C. G. and Van de Gucht, L. M.: 1999, High-yield bond default and call risks, *Review of Economics and Statistics* **81**(3), 409–419.

Murphy, K. M., Shleifer, A. and Vishny, R. W.: 1991, The allocation of talent: Implications for growth, *The Quarterly Journal of Economics* **106**(2), 503–30.

Philippon, T.: 2010, Financiers versus engineers: Should the financial sector be taxed or subsidized?, *American Economic Journal: Macroeconomics* **2**(3), 158–82.

Philippon, T.: 2014, Has the us finance industry become less efficient? on the theory and measurement of financial intermediation. NYU working paper.

Philippon, T. and Reshef, A.: 2012, Wages and human capital in the u.s. finance industry: 1909-2006, *The Quarterly Journal of Economics* .

Shakhnov, K.: 2014, The allocation of talent: Finance versus entrepreneurship. EUI working paper.

Stiglitz, J. E. and Weiss, A.: 1981, Credit rationing in markets with imperfect information, *American Economic Review* **71**(3), 393–410.

Tobin, J.: 1984, On the efficiency of the financial system, *Lloyds Bank Review* .

Ueda, M.: 2004, Banks versus venture capital: Project evaluation, screening, and expropriation, *The Journal of Finance* **59**(2), 601–621.

Wilson, C.: 1980, The nature of equilibrium in markets with adverse selection, *The Bell Journal of Economics* **11**(1), 108–130.

Zhou, C.: 2001, Credit rating and corporate defaults, *The Journal of Fixed Income* **11**(3), 30–40.

Appendix

Deriving A and μ

The allocation measures $A(\cdot;\chi,p)$ formalize the notion that banks obtain representative samples from the assets on sale that they find acceptable. The rationing function μ formalizes the notion that whether assets that are put on sale are actually sold depends on how many units are demanded by banks who finde them acceptable. To compute A and μ, first define supply and demand.

The supply of asset i at price p is:

$$S(i;p) = \int_s \mathbb{I}\left(p^R(i,s) \leq p\right) \tag{38}$$

(38) is just aggregating all the supply from households whose reservation prices are below p.

Demand is defined as a measure. Suppose X is some set of possible acceptance rules. Define

$$\Theta(X,p) \equiv \{\theta : \chi(\theta) \in X, p(\theta) \geq p\}$$

$\Theta(X,p)$ is the set of bank types who choose to buy at prices above p using acceptance rules in the set X. Aggregating $\delta(\theta)$ over this set gives demand:

$$D(X,p) = \int_{\theta\in\Theta(X,p)} \delta(\theta)\, dW(\theta) \tag{39}$$

One complication is that if different banks impose different acceptance rules in the same market, the allocation will depend on the order in which they execute their trades because each successive bank will alter the sample from which the following banks draw assets. Kurlat

(2016) shows that if one allows markets for each of the possible orderings and lets traders self-select, then in equilibrium trades will take place in a market where the less restrictive banks execute their trades first.[6] Less-restricive banks' trades do not alter the relative proportions of acceptable assets available for the more-restrictive banks who follow them so, as long as acceptable assets don't run out, all bankers obtain assets as though they were drawing from the original sample. This means that (as long as acceptable assets don't run out before a bank with rule acceptance rule χ trades, which does not happen in equilibrium) the density of measure $A(\cdot;\chi,p)$ is:

$$a(i;\chi,p)=\begin{cases}\frac{\chi(i)S(i;p)}{\int\chi(i)S(i;p)di} & \text{if } \int\chi(i)S(i;p)di>0\\ 0 & \text{otherwise}\end{cases} \tag{40}$$

Knowing A, the rationing faced by an asset i depends on the the ratio of the total demand that gets satisfied (added across all χ) to supply, so

$$\mu(p,i)=\int_{\substack{\tilde{p}\geq p\\ \text{all }\chi}}\frac{a(i;\chi,\tilde{p})}{S(i;\tilde{p})}dD(\chi,\tilde{p}) \tag{41}$$

Proof of Proposition 1

Full statement of the equilibrium.

The equilibrium is given by equations (14)-(17) plus the statement of $A(\cdot;\chi,p)$ for other values of p and χ and $\mu(p,i)$ for other values of p:

[6] An acceptance rule $\tilde{\chi}$ is less restrictive than another rule χ if $\chi(i)=1$ implies $\tilde{\chi}(i)=1$ but there exists some i such that $\tilde{\chi}(i)=1$ and $\chi(i)=0$. Under the information structure (2), all feasible acceptance rules can be ranked by restrictiveness.

$$a(i;\chi,p)=\begin{cases}\frac{\beta^{-1}\left(\frac{p}{V}\right)\chi(i)}{\int_0^\lambda\chi(i)di+\int_\lambda^1\chi(i)\beta^{-1}\left(\frac{p}{V}\right)di} & \text{if } i\geq\lambda \text{ and } p\geq p^*\\ \frac{\chi(i)}{\int_0^\lambda\chi(i)di+\int_\lambda^1\chi(i)\beta^{-1}\left(\frac{p}{V}\right)di} & \text{if } i<\lambda \text{ and } p\geq p^*\\ 0 & \text{if } i\geq\lambda \text{ and } p<p^*\\ \frac{\chi(i)}{\int_0^\lambda\chi(i)di} & \text{if } i<\lambda \text{ and } p<p^*\end{cases} \tag{42}$$

$$\mu(p,i)=\begin{cases}0 & \text{if } p>p^*\\ \mu(p^*,i) & \text{if } p\leq p^*\end{cases} \tag{43}$$

Equations (14)**-**(17)**,** (42) **and** (43) **constitute an equilibrium.**

1. Household optimization. (43) and (17) imply that:

$$p^L(i)=\begin{cases}p^* & \text{if } i\geq\lambda\\ 0 & \text{if } i<\lambda\end{cases}$$

   This immediately implies that $p^R(i,s)$ from (14) solves the household's problem.

2. Bank optimization.

   (a) $\chi(\theta)$ is the optimal acceptance rule because, given (42), any other rule that satisfies (8) includes a higher proportion of bad assets.

   (b) At any $p<p^*$, there are no good assets on sale so it is not optimal for any bank to choose this. For any $p>p^*$:

$$\begin{aligned}\frac{1}{p}\frac{\beta^{-1}\left(\frac{p}{V}\right)}{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta^*)}&<\frac{1}{p^*}\frac{s^*}{s^*(1-\lambda)+\lambda(1-\theta^*)}\\ \frac{p^*}{p}\frac{\beta^{-1}\left(\frac{p}{V}\right)}{s^*}&<\frac{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta^*)}{s^*(1-\lambda)+\lambda(1-\theta^*)}\\ \frac{p^*}{p}\frac{\beta^{-1}\left(\frac{p}{V}\right)}{s^*}&<\frac{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta)}{s^*(1-\lambda)+\lambda(1-\theta)}\quad\text{for all }\theta\geq\theta^*\\ \frac{1}{p}\frac{\beta^{-1}\left(\frac{p}{V}\right)}{\beta^{-1}\left(\frac{p}{V}\right)(1-\lambda)+\lambda(1-\theta)}&<\frac{1}{p^*}\frac{s^*}{s^*(1-\lambda)+\lambda(1-\theta)}\quad\text{for all }\theta\geq\theta^*\end{aligned} \tag{44}$$

   The first step is Assumption (1); the second is just rearranging; the third follows because the right hand side is increasing in θ and the last is just rearranging.

Inequality (44) implies that all banks with $\theta \geq \theta^*$ prefer to buy at price p^* than at higher prices. Therefore if they buy at all they buy at price p^*.

(c) For $\theta > \theta^*$, $\tau(\theta) > 0$ so the budget constraint (7) binds; for $\theta < \theta^*$ there is no $\chi(\theta)$ that satisfies (8) and leads to a positive value for the objective (6). Therefore $\delta(\theta)$ is optimal .

3. Consistency of A and μ. Replacing reservation prices (14) into (38) and using this to replace $S(i;p)$ into (40) leads to (42). Adding up demand using(15) and (39) and replacing in (41) implies (43).

The equilibrium is unique

Note first that since no feasible acceptance rule has $\chi(i) \neq \chi(i')$ for $i, i' \geq \lambda$, this implies that $p^L(i) = p^L(\lambda)$ and $S(i,p) = S(\lambda,p)$ for all $i \geq \lambda$. Now proceed by contradiction.

Suppose there is another equilibrium with $p^L(\lambda) < p^*$. Households' optimization condition (5) and formula (38) for supply imply that for $p \in \left[p^L(\lambda), p^*\right]$:

$$S(i,p) = \begin{cases} \beta^{-1}\left(\frac{p}{V}\right) & \text{if } i \geq \lambda \\ 1 & \text{if } i < \lambda \end{cases} \tag{45}$$

(45) implies that all banks with $\theta > \theta^*$ can attain $\tau(\theta) > 0$ by choosing p^*. By (44), they prefer p^* to any $p' > p^*$ and therefore in equilibrium they all chose some $p(\theta) \in \left[p^L(\theta), p^*\right]$ and $\delta(\theta) = \frac{1}{p(\theta)}$. Using (40):

$$a(i, \chi(\theta), p(\theta)) = \frac{\beta^{-1}\left(\frac{p(\theta)}{V}\right)}{\beta^{-1}\left(\frac{p(\theta)}{V}\right) + \lambda(1-\theta)} \quad \text{for all } i \geq \lambda$$

Using (41), this implies that

$$\mu(p,\lambda) = \int_{\{\theta : p(\theta) \geq p\}} \frac{1}{\beta^{-1}\left(\frac{p(\theta)}{V}\right) + \lambda(1-\theta)} \frac{1}{p(\theta)} dW(\theta)$$

and therefore

$$\begin{aligned}\mu\left(p^{L}(\lambda),\lambda\right) &\geq \int_{\theta^{*}}^{1} \frac{1}{\beta^{-1}\left(\frac{p(\theta)}{V}\right)+\lambda(1-\theta)} \frac{1}{p(\theta)} dW(\theta) \\ &\geq \int_{\theta^{*}}^{1} \frac{1}{s^{*}+\lambda(1-\theta)} \frac{1}{p^{*}} dW(\theta) \\ &= 1 \end{aligned} \tag{46}$$

The first inequality follows because the set $\{\theta : p(\theta) \geq p(\lambda)\}$ includes $[\theta^{*}, 1]$; the second follows because $\beta^{-1}\left(\frac{p^{*}}{V}\right) = s^{*}$, β^{-1} is increasing and $p^{*} \geq p(\theta)$; the last equality is just the market clearing condition (13). Furthermore, if $p(\theta) < p^{*}$ for a positive measure of banks, then (46) is a strict inequality, which leads to a contradiction. Instead, if $p(\theta) = p^{*}$ for almost all banks, then $p^{L}(\lambda) = p^{*}$, which contradicts the premise.

Suppose instead that there is an equilibrium such that $p^{L}(\lambda) > p^{*}$. This implies that there is no supply of good assets at any price $p < p^{L}(\lambda)$ and therefore no bank with $\theta < \theta^{*}$ chooses $\delta(\theta) > 0$ and banks $\theta \in [\theta^{*}, 1]$ choose some price $p(\theta) \geq p^{L}(\lambda)$ and $\delta(\theta) \leq \frac{1}{p(\theta)}$. Therefore, using (40) and (41), we have

$$\begin{aligned}\mu\left(p^{L}(\lambda),\lambda\right) &\leq \int_{\theta^{*}}^{1} \frac{1}{\beta^{-1}\left(\frac{p(\theta)}{V}\right)+\lambda(1-\theta)} \frac{1}{p(\theta)} dW(\theta) \\ &< \int_{\theta^{*}}^{1} \frac{1}{s^{*}+\lambda(1-\theta)} \frac{1}{p^{*}} dW(\theta) \\ &= 1 \end{aligned}$$

The first inequality follows from $\delta(\theta) \leq \frac{1}{p(\theta)}$; the second follows because $\beta^{-1}\left(\frac{p^{*}}{V}\right) = s^{*}$, β^{-1} is increasing and $p^{*} < p(\theta)$; the last equality is just the market clearing condition (13). Again, this is a contradiction.

Therefore any equilibrium must have $p^{L}(\lambda) = p^{*}$. The rest of the equilibrium objects follow immediately.