Optimal Credit Scores Under Adverse Selection

Nicole Immorlica, Andre Sztutman, and Robert Townsend

May 19, 2022

Abstract

The increasing availability of data in credit markets may appear to make adverse selection concerns less relevant. However, when there is adverse selection, more information does not necessarily increase welfare. We provide tools for making better use of the data that is collected from potential borrowers, formulating and solving the optimal disclosure problem of an intermediary with commitment that seeks to maximize the probability of successful transactions, weighted by the size of the gains of these transactions. We show that any optimal disclosure policy needs to satisfy some simple conditions in terms of local sufficient statistics. These conditions relate prices to the price elasticities of the expected value of the loans for the investors. Empirically, we apply our method to the data from the Townsend Thai Project, which is a long panel dataset with rich information on credit histories, balance sheets, and income statements, to evaluate whether it can help develop the particularly thin formal rural credit markets in Thailand, finding economically meaningful gains from adopting limited information disclosure policies.

1 Introduction

Data is becoming increasingly available and more easily processed. New data and methods are useful across many economic sectors and applications, including credit markets. There is a large population of potential borrowers who have short credit histories and thus are unable to receive credit (Bricker et al., 2017). The new methods would allow banks to identify the creditworthy among these potential borrowers, giving credit to those who perhaps need it the most (Jagtiani and Lemieux, 2019).

Because of this increased capacity to identify creditworthy individuals, one may hope that the inefficiencies coming from information asymmetries would progressively disappear. However, a key reason that makes information asymmetries generate inefficiencies in these thin credit market segments is adverse selection: as the price of the loans decreases (or interest rates increase), the pool of borrowers can get progressively worse. Those who would be more likely to repay are only willing to borrow at higher prices. The credit market unravels, resulting in too few or no transactions happening. As long as there is some information asymmetry, some heterogeneity in expected repayment rates that lenders cannot observe, there can still be adverse selection problems.

Data owners, such as data-intensive firms and platforms, may hope that by making their data available to financial providers they will improve credit access. However, this hope lacks a theoretical justification. The inefficiencies arising from adverse selection do not necessarily get better with more information, and indeed, may as well get worse. As shown in Levin (2001), more information does not necessarily increase the number of transactions and the realized gains from trade. More information can prevent implicit cross-subsidization between different types, making the previously subsidized types leave the market. Hence, as more data and improved technologies for processing data arrive, there remain key issues concerning how much data to share.

To answer the question of how much data to share, we build on the literature on information design and formulate the optimal disclosure problem of a partially informed intermediary with commitment, maximizing the probability of successful transactions weighted by the size of gains from trade. This formulation allows us to answer the question of which variables in a dataset should be shared with financial providers – for example, whether geographic information should be shared and at which level of granularity, or whether only an index that combines different pieces of information should be shared. We construct an optimal disclosure system and derive new conditions for the optimality of a disclosure system in terms of local sufficient statistics. The optimal policy should satisfy three simple properties: i) generically messages should combine at most two signals; ii) there should be an increasing relationship between the price elasticities of the value of the loans to investors and the prices of these loans; and iii) when different signals are combined into a single message, there should be a decreasing relationship between these elasticities and the prices these loans would have if the signals were unbundled.

We apply our results to the rural credit markets in Thailand. This is a particularly fitting setup for at least four different reasons. First, these credit markets are thin and there is not much risk sharing across villages, so the potential welfare gains are large. Second, there is evidence of intensive risk sharing within villages, which makes us think that they, through a platform acting on their behalf, are able to organize and commit to an optimal disclosure policy. Third, a unique feature of this setup benefits us from an identification perspective. There is a main lender, the Bank of Agriculture and Agricultural Cooperatives, a government-owned bank, holding a significant fraction of the market for agricultural loans. This bank uses a rigid set of rules to set interest rates. We explore variation in these rules as a source of identification for slopes of supply and average value curves. These slopes are key ingredients in the computation of the optimal credit scores and appear as sufficient statistics in the necessary conditions we derive for the optimality of disclosure systems.[1] Fourth, we benefit from rich data from Townsend Thai Project, including detailed information on consumption, income and its different sources, crops, livestock, loans, and interest rates. This allows the construction of detailed balance sheets, income and cash flow statements for each household, as well as their credit histories. Assuming that the platform has access to the detailed information in this dataset, while investors do not, we show what pieces of information should be made available to investors and how, effectively constructing "optimal credit scores."

We find that the optimal disclosure policy substantially improves the gains from trade relative to a simple full disclosure policy, with the size of gains being of the order of 0.45% the size of a typical loan per household per month. Moreover, we find that the optimal policy puts higher weight than full disclosure credit scores on variables seemingly related to the solvency of farmers relative to variables that are informative about their current liquidity. Our findings can be instrumental in improving credit access in places where it is most needed by making better use of data.

Outline of the paper: The remainder of the paper is structured as follows. Section 2 discusses the related literature, and Section 3 presents the model and a simple motivating example, Section 4 presents our theoretical results. Section 5 presents the data, followed by the discussion of the empirical strategy and empirical results in Sections 6 and 7. Section 8 concludes.

2 Related literature

This paper is related to at least three different strands of the literature. There is a long line of research that has studied adversely selected markets, starting with Akerlof

[1] In Appendix C we complement this approach with an alternative identification strategy that leverages the evidence that there is intensive risk sharing within villages in Thailand.

(1970), including Glaeser and Kallal (1997) and Attar et al. (2011). More recently, a large literature has appeared on empirical tests for adverse selection and estimation of supply and demand in adversely selected markets (Finkelstein and Poterba, 2004; Cohen and Einav, 2007; Einav et al., 2010; Hendren, 2013, 2017; Handel and Kolstad, 2015; Finkelstein and Notowidigdo, 2019; Cabral et al., 2019). Relative to the empirical literature on adverse selection, our contribution is to combine machine learning methods to estimate slopes of supply and demand for many different markets simultaneously. To that goal, we borrow methods from the literature on the estimation of heterogeneous elasticities (Athey and Wager, 2019; Athey et al., 2019; Davis and Heller, 2017; Wager and Athey, 2018).

Another strand in the literature has studied information design problems (Lerner and Tirole, 2006; Ostrovsky and Schwarz, 2010; Rayo and Segal, 2010; Bergemann and Morris, 2013; Bergemann et al., 2015; Kamenica and Gentzkow, 2011). Our main contributions to this literature include relating the optimal policy to sufficient statistics that can be estimated, clarifying the economic mechanisms behind simple disclosure policies, and allowing for a fairly general multidimensional distribution of types.

Within this literature, we share the linear programming approach from Kolotilin (2018), and extend and derive new results that are closest to their conditions and to the pairwise signals condition found in Kolotilin and Wolitzky (2020)[2]. In the context of adversely selected markets as in Akerlof (1970), Levin (2001) has analyzed conditions under which more information increases trade volume, besides providing examples where welfare is not monotonic in the amount of information. In the context of insurance markets, an optimal disclosure algorithm assuming a fully informed intermediary, has been developed by Garcia et al. (2018). We in contrast do not suppose the intermediary is fully informed, but rather sees a signal and then decides on an information disclosure policy. Kartik and Zhong (2019) have characterized the set of feasible payoff vectors for a buyer and a seller across all possible information structures when the seller posts the price. Besides featuring a different market structure – perfect competition on the buyers (investors) side – we focus on a different question, which is what is the best information structure an intermediary can design, when it is constrained to a limited information set.

A recent literature on information design, including some of the articles cited above, has studied in detail what is called "linear" persuasion models, where the payoffs of senders (in our case, the platform) and receivers (investors) are linear in a single dimensional state variable (Dworczak and Martini, 2019; Arieli et al., 2020; Kolotilin et al., 2017; Dizdar and Kováč, 2020). The problem we analyze does not fit these assumptions, as we allow for more general payoff structures and multidimensional types.

Finally, this paper is related to the literature that has linked consumption patterns to risk sharing and insurance, including Gruber (1997); Ahlin and Townsend (2007); Giné

[2]The relationship between their results and ours is explained in more details in Section 4

(2011), and Townsend (1994, 1995). In the context of Thai villages, Chiappori et al. (2014), in particular, have documented intensive consumption risk sharing at the village level.

3 Model

The model features three types of agents: there are potential borrowers, investors, and a platform. Borrowers sell up to a one dollar claim to investors in the present, with a promise of repaying it in the future. The platform mediates these sales and specifies the information available to the other agents. If a borrower sells a claim for a price of $x \leq 1$, then the borrower receives a loan of x dollars today and promises to repay \$1 in the future. They can sell fractional shares of this claim to multiple investors, and investors can buy claims from multiple borrowers.

Borrowers have private types $\omega \in \Omega$. The types could be, for example, the borrower's default probability. The borrower's type ω determines the borrower's opportunity cost of selling a claim $b(\omega)$. More concretely, in the Appendix C we present a setup where $b(\omega)$ reflects the implicit option value that selling the claim offers to the borrower, coming from the possibility that the borrower would not repay the loan in the states of the world where the marginal value of consumption is high.

Investors are homogenous, each having value $a(\omega)$ for a unit claim from borrowers of type ω. Investors share a common prior over borrower types but do not observe the type of any given borrower. Assuming investors are risk neutral, this value is simply the discounted expected repayment conditional on the type ω of the borrower.

The platform has access to signals $z(\omega)$ that are partially informative about the borrowers' types, and commits to an information disclosure policy $m(\cdot)$ which is a randomized mapping from signal realizations $z(\omega)$ to an arbitrary message. For succinctness, we will write $m(\omega) \equiv m(z(\omega))$ and note that it must be measureable with respect to the signal realizations $z(\omega)$.

After observing the signal realization, the platform sends a message $m(\omega)$ to the investors, and the investors respond by offering contracts to borrowers. As shown in Appendix A, under a few technical assumptions, it is without loss to assume contracts are specified by a price $x(m(\omega))$ offered by investors to borrowers conditional on receiving message m. We focus on a competitive equilibrium in which investors make zero profits. If there are multiple prices that would guarantee investors break even, we take the highest price, which is the best price for the borrowers. More formally, we adopt the following definition.

Definition 1. *A competitive equilibrium is a set of prices $x^*(m(\omega))$ and allocations such that:*

Table 1: Joint distribution of values in the example

	L	M	H
value for investors (as buyers) $a(\omega)$	0.12	0.3	0.84
value for borrower (as sellers) $b(\omega)$	0.06	0.24	0.36
probability $\rho(\omega)$	1/3	1/3	1/3

- *investors break even and prices are borrower-optimal:*

$$x^*(m(\omega)) = sup(x|x = E[a(\omega)|x \geq b(\omega), m(\omega)]),$$

- *borrower types for which $b(\omega) \leq x^*(m(\omega))$ borrow up to the borrowing limit of* \$1*,*

- *and borrower types for which $b(\omega) > x^*(m(\omega))$ do not borrow.*

In general, there can be multiple solutions to the fixed-point equation $x = E[a(\omega)|x \geq b(\omega), m(\omega)])$. Regions where the expectation is decreasing as price decreases correspond to adverse selection: the borrowers willing to sell claims at this lower price are worse borrowers (and there are fewer of them), so the value of the investor goes down and the market may unravel resulting in no trade. Similarly, regions where the expectation is increasing as the price decreases correspond to advantageous selection and can result in unraveling in which all trades, even inefficient ones, happen.

The key idea behind Definition 1 is that, i) investors need to take into account that the price affects the pool of the types that are willing to borrow; and, ii) if there are multiple prices that make the investors break even, and if the equilibrium were not the highest among them, some investor could offer a price in a neighborhood of the highest, attracting the borrowers and generating positive profits.

We illustrate our model and implications for optimal disclosure in the following example adapted from Levin (2001). In particular, this example demonstrates that information can decrease welfare.

There are three types of borrowers $\omega \in \Omega$, associated with their repayment probabilities: low (L), medium (M) or high (H). The three types occur with equal probability $\rho(\omega) = 1/3$. The platform has access to binary signal realizations. If $\omega = H$, the signal $z(\omega)$ is $\{H\}$. Else, if $\omega = L$ or $\omega = M$, the signal $z(\omega)$ is $\{M, L\}$. In other words, the platform can differentiate H borrowers from M and L borrowers, but can not differentiate M and L borrowers from each other. The borrowers' and investor values $b(\omega)$ and $a(\omega)$ are as in Table 1.

Consider the following disclosure policies.

- In the full disclosure policy, the platform sends message $m(\omega) = z(\omega)$. In this case, in the H market there would be no information asymmetry. Competition between the investors would drive the price up to their willingness to pay, 0.84, and these gains from trade would be realized by the borrowers. In the L or M market there is asymmetric information. There cannot be a pooling equilibrium in this market, where a bundle of the two types would be sold together, because investors' willingness to pay is less than the reservation value of a borrower of type M, that is, $E[a(\omega)|\omega \in \{M, L\}] = 0.21 < 0.24$. Thus, M's would not be willing to sell, and the bundle unravels to L. So there will be a separating equilibrium, and only the L will types trade, at a price of 0.12, which we can call the full disclosure price of the bundle $\{M, L\}$. The total surplus generated by the sales of L and H borrower claims would be $(0.06+0.48)/3 = 0.18$, and not all gains from trade would be realized, as the M types are left out of the market.

- In the no-disclosure policy, the platform sends a null message. Then there is a pooling equilibrium because the investor willingness to pay for the whole bundle is higher than the highest reservation value for the borrowers — $E[a(\omega)] = 0.42 > 0.36$, and we can call analogously 0.42 the no-disclosure price of the market. The welfare now increases to $(0.06 + 0.06 + 0.48)/3 = 0.2$, and all the gains from trade are realized.

Thus, in this example welfare decreases when the platform reveals more information, which is illustrative of why more information is not always better.[3]

4 Optimal Disclosure Policies

In this section, we formulate the optimal disclosure problem of the intermediary as a linear programming problem. Then we derive necessary conditions for the optimality of garbling and separating signals. These conditions give economic content to what optimal disclosure policies do and tell us simple properties that they need to satisfy. These properties are summarized in three rules-of-thumb, relating prices and elasticities of the value for investors: i) generically messages should combine at most two signals; ii) there should be an increasing relationship between the price elasticities of the value of the loans to investors and the prices of these loans; and iii) when different signals are combined into a single message, there should be a decreasing relationship between these elasticities and the prices these loans would have if the signals were unbundled. We illustrate these properties with an example at the end of this section.

[3]Moreover, there is a non-monotonic relationship between information and welfare: suppose the platform has full information, i.e., $z(\omega) = \omega$. Then welfare would increase again with disclosure policy $m(\omega) = z(\omega) = \omega$. Hence full disclosure would be as good as no disclosure which, in turn, is better than partial disclosure of the form described above.

4.1 Characterization

As demonstrated by Example 3, the choice of disclosure policy impacts welfare. The optimal disclosure policy can be described by a linear program whose constraints characterize the competitive equilibria of Definition 1 and whose objective maximizes welfare.

This linear program can be simplified by noticing two facts: First, messages can be identified by their equilibrium prices. This is because, if investors break even at the same price for two messages, they still break even if these two messages are combined into a single message. Further, no higher price generates positive profits. Therefore we can label each message by its equilibrium price. Second, because welfare is increasing in prices, the supremum from Definition 1 can be dropped. That is, we can simply confine $x = E[a(\omega)|x \geq b(\omega), m(\omega)]$ and the objective will guarantee we select the highest such x.

Thus, for an arbitrary distribution of signals $G(z)$[4], we can formulate the optimal informational disclosure problem as follows. Note the choice variables are the conditional distribution of messages (equivalently, equilibrium prices) given signals. By assigning different probabilities of the signals to different messages, the platform affects the objective function because for each of these messages there will be a different equilibrium price and a different pool of borrowers signing a loan contract.

Proposition 1. *The optimal informational disclosure problem can be formulated as:*

$$\begin{aligned} \max_{\phi(x|z)} & \int_{X \times Z} \sigma(x|z) d\phi(x|z) dG(z) \\ st. & \int_{\tilde{X} \times Z} \pi(x|z) d\phi(x|z) dG(z) = 0 \quad \textit{for any measurable set } \tilde{X} \subset X \\ & \int_X d\phi(x|z) = 1 \quad \textit{almost surely } \forall z \in Z \end{aligned}$$

$$\begin{aligned} \textit{where } \sigma(x|z) &\equiv \mathbb{E}\left[\mathbb{1}(x \geq b)(x-b)|z\right] \quad \textit{Borrowers' surplus / total surplus} \\ \textit{and } \pi(x|z) &\equiv \mathbb{E}\left[\mathbb{1}(x \geq b)(a-x)|z\right] \quad \textit{Investors' surplus} \end{aligned}$$

Where we have denoted the expected borrowers' surplus for a prevailing price x and conditional on the signal z by $\sigma(x|z)$. Similarly, we denote the expected investors profits for a prevailing price x and conditional on signal z by $\pi(x|z)$.

4.2 Conditions for optimality of separating signals

By starting from a given information policy, and considering revealing less information, we can arrive at simple conditions, presented in Proposition 2[5] and Proposition 3.

[4] To make the notation cleaner, in what follows, we drop the explicit dependence of z, a, b on ω.

[5] The conditions in Propostion 2 are analogous to the full disclosure conditions derived in Kolotilin (2018) under a similar but different set of assumptions. In particular, we do not assume a single-crossing

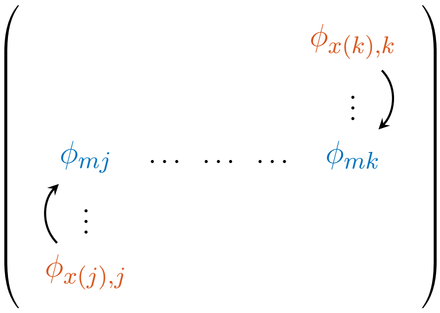

Figure 1: Deviations from full disclosure

To state these conditions, it will be convenient to discretize the model. Let m index a price, and j and k index signals. Consider a discrete approximation of the infinite linear programming problem, with a finite number of signals and prices, and such that for each signal z_j, its full disclosure price $x^*(z_j)$ is included in the discretization of the prices. Denote its index by $x(j)$. Let $\sigma_{mj} = \mathbb{E}[(x_m - b)\mathbb{1}(x_m \geq b)|z_j]\hat{g}(z_j)$ and $\pi_{mj} = \mathbb{E}[(a - x_m)\mathbb{1}(x_m \geq b)|z_j]\hat{g}(z_j)$, where $\hat{g}(z_j)$ is the probability of signal z_j. Let the decision variables $\phi_{m,j}$ denote the conditional probability of price m given signal j.

The first proposition says that full disclosure, which fully separates all signals, is optimal if and only if there is no benefit to pooling pairs of signals. We state this proposition as a characterization for the optimality of full-disclosure, it should be noted that it applies more generally. For any policy to be optimal, it must be the case that garbling two messages when feasible does not increase welfare. The sufficiency of this statement only holds for full disclosure because for full disclosure garbling is the only feasible perturbation.

Proposition 2. *Full disclosure is optimal if and only if*

$$\left(\sigma_{mj} - \sigma_{x(j),j} - \frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk} - \sigma_{x(k),k})\right) \leq 0 \quad \textit{for all } m, j, k \textit{ such that } \pi_{mj} > 0 \textit{ and } \pi_{mk} < 0.$$

The necessity part of the proof is based on the following argument. The expression above is the change in welfare from the following deviation: move some probability of signal z_j from being assigned to its full disclosure price to another price x_m, and make sure investors break even by moving some probability of another signal z_k to the same price x_m in the right proportion. This deviation is illustrated in matrix form in Figure 1. If the policy is optimal, then this change in welfare must be negative. To show that this condition is also sufficient we notice that any feasible direction can be written as conical combinations of feasible directions of the form above. The formal proof appears in Appendix B.

To gain further economic insight into the content of this proposition and the properties

condition on types.

of optimal disclosure policies, we further specialize to deviations of full disclosure where signals that have nearby full disclosure prices are combined.[6] This results in a simple monotonicity condition, relating prices and the price elasticity of the average value to investors. Recall that the price elasticity of value is defined as the percent change in value in response to a percent change in price, i.e.:

$$\epsilon_{V,x}(x^*(z), z) \equiv \frac{\partial E[a|b \leq x, z]}{\partial x} \frac{x}{E[a|b \leq x, z]}\Big|_{x=x^*(z)}.$$

Proposition 3. *Suppose there is an interval $[\underline{x}, \overline{x}]$ such that the full disclosure price $x^*(z)$ is dense over it. If full disclosure is optimal over this interval, then for z and z' such that $x(z) \in [\underline{x}, \overline{x}]$ and $x(z') \in [\underline{x}, \overline{x}]$, the signal with the higher price must have higher elasticity:*

$$x^*(z) > x^*(z') \Rightarrow \epsilon_{V,x}(x^*(z), z) \geq \epsilon_{V,x}(x^*(z'), z')$$

where $\epsilon_{V,x}(x^*(z), z)$ is the price elasticity of the value for investors.

This proposition has a simple graphical interpretation. Presented in Figure 2 is a case where the state necessary condition fails. That is, $x^*(z) > x^*(z')$ but $\epsilon_{V,x}(x^*(z), z) < \epsilon_{V,x}(x^*(z'), z')$. We will show that we can garble these signals an increase the planners' objective.

First note our equilibrium condition requires that investors break even, i.e., $x = E[a|b \leq x, z]$, and so the price elasticity of investor value reduces to:

$$\epsilon_{V,x}(x^*(z), z) = \frac{\partial E[a|b \leq x, z]}{\partial x}\Big|_{x=x^*(z)}$$

for equilibrium prices $x^*(z)$. That is, the elasticity at an equilibrium price is simply the slope of the average value curve at that point. In the figure, we denote this average value curve by $AVG_V(x, z) = E[a|b \leq x, z]$. In general, this can be an arbitrary function; in the figure we draw it linearly as we imagine the full disclosure prices of the signals are close and hence a linear curve is a good approximation. Points in the forty-five degree line (dashed) correspond to prices where investors break-even, that is where they pay exactly what that claim is worth for them. In the figure there are two signals, z and z'. The solid lines correspond to the average value for investors as a function of prices given those corresponding signals. The average value curve for z crosses the forty-five degree line at a higher price than the average value curve for z', that is, the full disclosure price of z is higher than z'.

Now, starting from the full disclosure configuration, let's consider the consequences of creating a message that partially garbles the two signals. In particular, let's combine

[6] Again, this result applies more generally than only full disclosure policies, by instead of combining signals the platform considers combining messages with nearby prices.

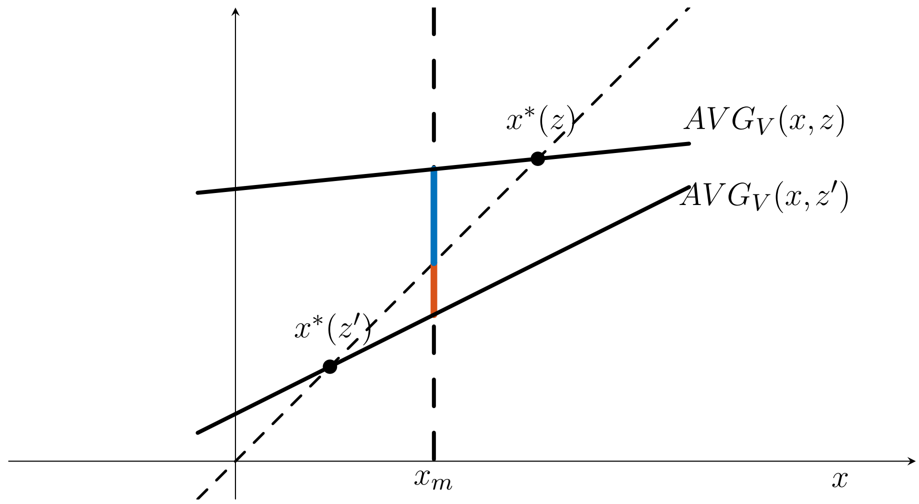

Figure 2: Example where condition for full disclosure to be optimal fails

units of these two signals into a new message, such that the resulting price (x_m) is the midpoint of the two full disclosure prices.

In order to do that, we need to combine them at the right proportions to make sure the investors still break even. When we lower the price of a unit of z, investors are going to face a profit of the size of the blue bar, which is the difference between the average value and the price. The average value changes as the composition of borrowers that take up the loan changes. When we do the same for a unit of z', raising their prices, investor are going to face a loss of the size of the red segment. For investors to break even, we take a number of units of z and z' that is inversely proportional to the size of these segments [7]. Because the condition fails in the proposition fails (the signal with the higher price has the flatter curve), the blue segment is larger than the red segment. Therefore, we will be lowering the price of fewer units of z than we are raising the price of units of z'. This means that on average we have increased prices. Because to a first order, the change in welfare is the change in prices times the number of units for which we are changing prices, this garbling procedure has increased welfare.

In contraposition, if the condition in Proposition 3 holds (with strict inequality), it cannot be optimal to pool two signals z and z' into a price $\bar{x}$ where $x^*(z) > \bar{x} > x^*(z')$, as depicted in Figure 3.

4.3 Conditions for optimality of pooling signals

Instead of starting from a full disclosure configuration and analyzing the consequences of garbling signals, Proposition 4 below starts from a no disclosure configuration, i.e., one in which all signals are pooled, and checks the consequences of revealing more information.

[7] Because we are near the region where the average value of investors is equal to prices, we can ignore the change in profits that would come from the change in quantities, because they are to a first order zero.

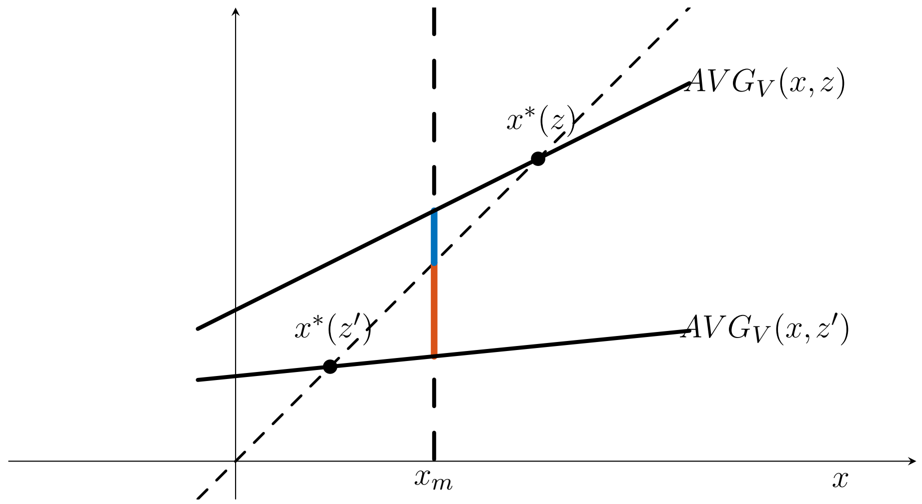

Figure 3: Example where condition for full disclosure to be optimal holds

Let x_0 be the no disclosure price, and 0 its index. Let m, m' index messages, and i, j, i', j' index signals. For notational convenience, let

$$\Delta_\pi(m,j,k) \equiv \frac{\pi_{mj}}{\pi_{mk}}\pi_{0k} - \pi_{0j}.$$

Proposition 4. *No disclosure is optimal if and only if:*

$$\left(\sigma_{mj} - \sigma_{0j} - \frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk} - \sigma_{0k})\right) - \frac{\Delta_\pi(m,j,k)}{\Delta_\pi(m',j',k')}\left(\sigma_{m'j'} - \sigma_{0j'} - \frac{\pi_{m'j'}}{\pi_{m'k'}}(\sigma_{m'k'} - \sigma_{0k'})\right) \leq 0$$

for all m, j, k, m', j', k' such that $\pi_{mj}, \pi_{m'j'}, \Delta_\pi(m,j,k) \geq 0$, and $\pi_{mk}, \pi_{m'k'}, \Delta_\pi(m',j',k') < 0$. If there are no m, j, k, m', j', k' such that $\pi_{mk}, \pi_{m'k'}, \Delta_\pi(m',j',k') < 0$, then no disclosure is optimal if and only if:

$$\left(\sigma_{mj} - \sigma_{0j} - \frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk} - \sigma_{0k})\right) \leq 0$$

for all m, j, k, m', j', k', such that $\pi_{mj}, \pi_{m'j'}, \Delta_\pi(m,j,k) = 0$.

This is the change in welfare from the following deviation: move some probability of signal z_j away from the no disclosure price x_0 and to some other price x_m, make sure investors break even by moving some probability of another signal z_k to the same price x_m, in the right proportion. But now, investors do not break even at the price x_0; to fix this, repeat the procedure above for a price x'_m, and signals z'_j,z'_k and combine the two procedures in the right proportion. This deviation in illustrated in Figure 4. The quantity $\Delta_\pi(m,j,k)$ plays an analogous role to π_{mj}. It says when two signals j and k are assigned to price m in a proportion that makes investors break even at m, what is the size of the profit that investors will be making at the original price x_0.

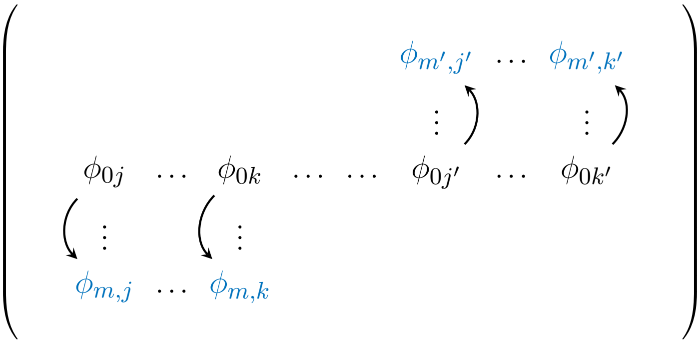

Figure 4: Deviations from no disclosure

To gain further economic insight into under what conditions revealing more increases the total value of transactions, we can specialize Proposition 4 above to the cases where disclosing information results in infinitesimally small changes in prices. Define $\pi(x,z) = \mathbb{E}[(a-x)\mathbb{1}(x \geq b)|z]g(z)$, where $g(z)$ is the density of z or the probability mass of z.

Proposition 5. *All signals z that are pooled together in x_0 must satisfy:*

$$\frac{x_0}{E[a|b \leq x_0, z] - x_0} = \kappa(x_0)\epsilon_{P,x}(x_0,z) + \gamma(x_0)$$

for some constants $\kappa(x_0)$ and $\gamma(x_0)$, and where $\epsilon_{P,x}(x_0,z) = \frac{\partial \pi(x_0,z)}{\partial x}\frac{x_0}{\pi(x_0,z)}$

The proposition implies that at any x_0, few signals should be pooled together. Most of the time, complete no disclosure is not optimal. The result arises from the following observation: in general, if there are three or more signals combined into the same price, picking two signals and increasing their prices, and picking another pair and decreasing their prices (making sure the investors break even) either decreases or increases welfare. These possibilities are illustrated in Figure 5, where three signals are assigned to the same price x_0, and one tries deviate from the original configuration by raising the prices of a pair and lowering the prices of another [8]. Flipping the pairs, the planner can in general increase welfare, unless it happens that for any pair of signals the change in welfare from feasible price increases is exactly the same. This indifference can only be met if the affine relationship is satisfied. For single-dimensional or discretely distributed types, given an arbitrary x_0, this is generically satisfied only for pairs of types. For multidimensional types in $\mathbb{R}^d$ this defines a subspace of dimension $d-1$. [9]

[8] Although we plot the average values, the change in profits in this case also depends on the change in quantities, which are not plotted.

[9] A similar result appears in Kolotilin and Wolitzky (2020) stating that is generically optimal to use pairwise signal structures. In their case, payoffs are more general, but there is a single dimension of heterogeneity. Our result additionally states that signals should satisfy this affine relationship, which for multidimensional types in $\mathbb{R}^d$ defines a subspace of dimension $d-1$.

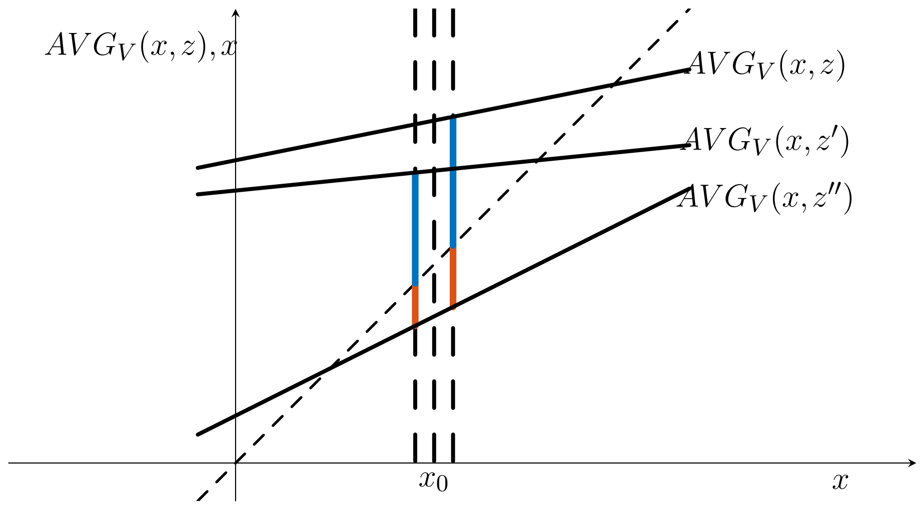

Figure 5: Example where the condition for no disclosure to be optimal fails

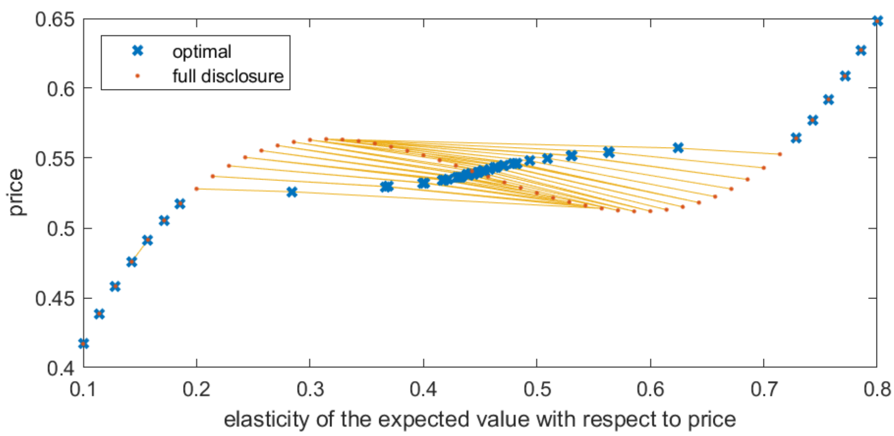

Figure 6: Example - Optimal information disclosure policy

4.4 Economic implications

The conclusions from Propositions 3 and 5 can be summarized in three simple rules of thumb. To arrive at those, we also notice that Proposition 3 applies more generally, not only for signals that the platform sees, but also to messages that combines multiple signals. Thus, these rules of thumb state that:

1. Markets with slightly higher prices should have higher price elasticities of the expected value for investors;

2. If two signals with nearby full disclosure prices are combined into a single message, the signal with a higher full disclosure price should feature a lower elasticity. In the example, whenever signals are combined, the signal with the high full disclosure price has a low elasticity;

3. Generically, each message should combine few signals. In the discrete case, one or two signals.

Figure 6 illustrates these rules in a hypothetical example. In this example, we postulate that conditional on each signal, supply and demand are such that: i) conditional on each signal, elasticities are constant as a function of prices; ii) elasticities of supply are equal to one for every signal; and iii) the signals are uniformly distributed, that is $\rho(z) = 1/\#Z$.

In Figure 6, each red dot is a different z, with their full disclosure price displayed on the vertical axis, and the elasticity of the expected value for the investors on the horizontal axis. Blue crosses represent the optimal policy, each cross is a different message x, and on the horizontal axis is their average elasticity.

In line with the rules of thumb, one can see that the arrangement of the blue crosses are such that there is an increasing relationship between prices and elasticities of the value for the investors, as implied by rule (1); the yellow lines connecting the messages to the signals that are assigned to them with positive probability are downward sloping, as implied by rule (2); and each blue cross is connected to at most two red dots, as implied by rule (3).

5 Data

The Townsend Thai monthly survey is an intensive monthly survey initiated in 1998. The analysis presented in this paper is based on 156 months from January 1999 (month 5) to December 2011 (month 160), which coincides with 13 calendar years. The four provinces of Thailand from which the sample is drawn are Chachoengsao and Lopburi in a more developed central region and Buriram and Srisaket located in the less developed northeastern region. The sampled townships (counties) of these provinces were part of an initial larger baseline initiative in 1997. The data utilized here are the continuously sampled households, those present in the survey throughout the 156 months. As we are concerned in this paper with rural credit markets, we include only households that generated income from farm and nonfarm business activities and drop the households whose income was almost exclusively from wage earnings. In the end, there are 541 households in the sample: 129 from Chachoengsao, 140 from Lopburi, 131 from Buriram, and 141 from Srisaket.

Notably, the monthly data have been used to create for each of these business households complete consistent financial accounts: Income Statement, including revenue, expenses and disposition of income (e.g., consumption as if dividends from profits and internal investment); Balance Sheet, including assets and liabilities, both real and financial, with net worth as the residual; and Statement of Cash Flow, with flows for production, consumption, investment, and financing [10]. The data set also has a loan form each loan. When the loan is initiated it specifies the lender and term, then the loan is placed on a roster and tracked each month until it is fully repaid, if ever.

A key assumption of this paper is that the platform has access to all the variables of the Townsend Thai monthly survey including the financial accounts, including disaggregated to the sectoral categories: fish/shrimp, farm, business, and livestock, as well as calendar time and geographic data.

6 Empirical Strategy

In this section, we present our main empirical strategy using the Thai Data, which uses exogenous variation in interest rates to identify slopes of supply and demand in these adversely selected markets. In Appendix C we present an alternative identification strategy, leveraging the documented fact that there is intensive within-village risk sharing, to estimate a structural model where consumption risk is shared at the village level while default decisions are idiosyncratic.

The Bank of Agriculture and Agricultural Cooperatives (BAAC) uses rigid rules to set interest rates. We can use changes in the rules as a source of exogenous variation in interest rates, to estimate the slopes of supply and average value curves as in Einav et al. (2010), Cabral et al. (2019), among others. The key idea is that the variation in interest, rates together with data on take-up, allows us to infer the supply of bonds from the borrowers. Similarly, using data on repayment rates, and assuming that investors are risk neutral, as interest rates change we can trace how repayment rates and thus the average value of the new loans for investors change.

[10] See Samphantharak and Townsend (2010) for details

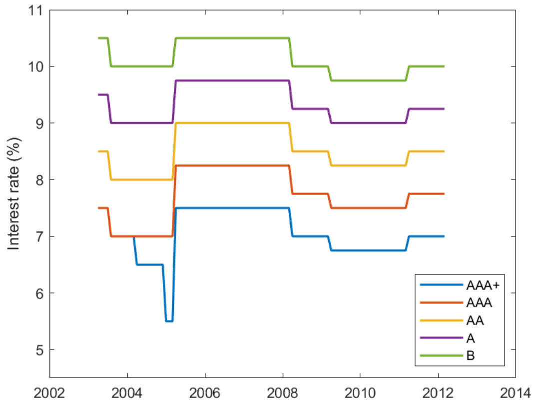

Symbol	Record of loan repayment
AAA+	full and in time for the past 4 years
AAA	full and in time for the past 3 years
AA	full and in time for the past 2 years
A	full and in time in the previous fiscal year
B	overdue in the previous year or a new client

Figure 7: BAAC Interest rate structure for farmers

Because we are not only interested in the average slopes across the different households, but also in the observable heterogeneity in these slopes, we use causal forests (Wager and Athey, 2018; Athey et al., 2019) to estimate the levels of take-up and expected returns and, importantly, the heterogenous elasticities of take-up and of the average value for investors.

With these estimates we can, first, test whether full disclosure is optimal (or if no disclosure or the current scoring system would be optimal under a competitive equilibrium) and, second, compute the optimal policy.

The econometric model we will estimate can be described by the following set of equations:

$$Q_i = X_i \cdot \gamma_i + \xi_i \qquad R_i = X_i \cdot \beta_i + \zeta_i$$

where Q_i denotes whether the household i takes up a loan at a particular month, and R_i is the repayment rate conditional on taking the loan. We allow for different individuals to respond differently to changes in prices X_i. We are interested in $\gamma(z) = \mathbb{E}[\gamma_i | Z_i = z]$ and $\beta(z) = \mathbb{E}[\beta_i | Z_i = z]$, average slopes of take-up and repayment rates with respect to price x, conditional on observable characteristics Z (which correspond to what we refer to as signals in the model), and the levels $\mathbb{E}[Q_i | Z_i = z, X_i = x]$, and $\mathbb{E}[R_i | Z_i = z, X_i = x]$. These are the average slopes (with respect to price) of take-up and repayment rates, respectively, over all households in the sample who generate signal z. The interpretation of these average slopes more concretely is in how repayment rates, in percentage points, change when there is a \$1 increase in the price of a \$100 bond ($\beta(z)$); and how take-up rates change, in percentage points, when there is a \$1 increase in the price of a \$100 bond ($\gamma(z)$).

Towards estimating these slopes, we want to make the following comparison: For households with observable characteristics z, how are repayment and take-up rates differ-

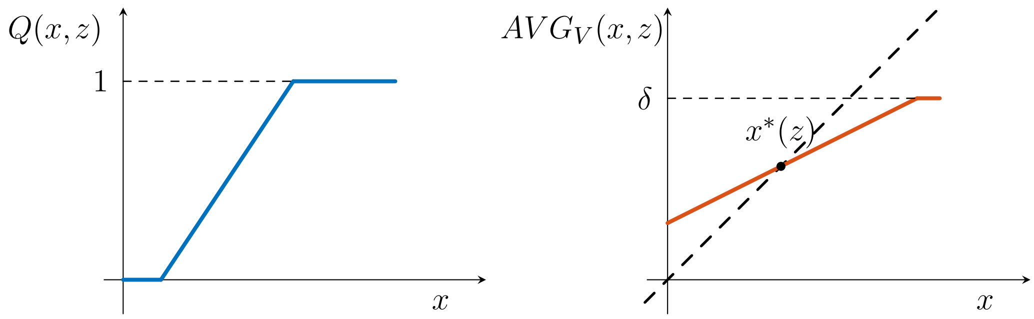

Figure 8: Supply and average value curves

ent when price is x versus when the price is x'? No two households have exactly the same observable characteristics, so we need to find a way to define who is similar to whom. Fortunately, causal forests (Athey and Wager, 2019; Wager and Athey, 2018) provide us with a solution to this question. The method uses forest-based weights to group together as a function of the observables households with similar slopes, separating households with distinctly different slopes. The approach solves a 'curse of dimensionality' problem, efficiently grouping households to estimate conditional elasticities. In doing that the algorithm effectively redefines the signals z in an economically meaningfully way.

The key identification assumption that will allows to use the method is that conditional on observables, X_i is exogenous, $\{\beta_i, \zeta_i, \gamma_i, \xi_i\} \perp\!\!\!\perp X_i | Z_i$. Figure 7 provides justification for this assumption. The BAAC uses very coarse rules for setting interest rates, which are seldom updated. Roughly, the BAAC considers the cost of funds in the Bangcok money market but, because it receives government subsidies, it can wait wait before making changes. It is also reluctant for public relations reasons as chartered development bank to increase interest rates. Thus, the updates in those coarse rules are unlikely to track changes in the pool of borrowers in these particular sample of villages. Therefore, we use the time variation in interest rates to estimate how take-up and repayment rates change in response to interest rate changes[11].

After estimating how repayment rates change with prices, we use additional assumptions to map those into reservation values for the investor. Assuming investors are risk neutral and given a discount factor, the average value for the investors of a unit bond they can buy is given by the simple relationship $AVG_V(x,z) = \mathbb{E}[a|b \leq x, z] =$ repayment rate $\times$ discount rate. Furthermore, because take-up and repayment rates above one hundred percent or below zero would be nonsensical, we postulate that take-up and value curves are piece-wise linear, bounded by zero and one, implying that the reservation values for investors are bounded by zero and the discount rate δ, as depicted in Figure 8.

The estimated slopes are presented in Figure 9. On the left side of the figure is a

[11]Because we can construct from the monthly surveys the individual credit histories, we can create a proxy for the BAAC rating, and conditioning on this rating, use only the time variation in interest rates to estimate the slopes of supply and repayment rates.

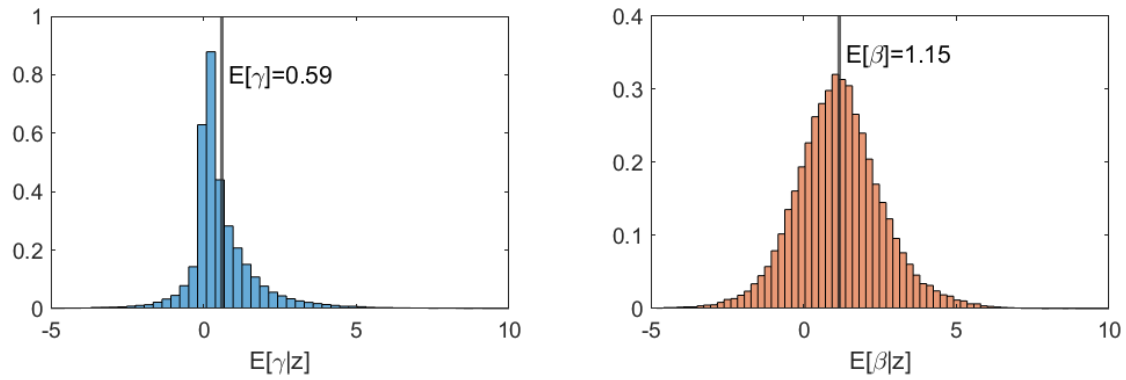

Figure 9: Histograms of slope estimates for take-up and repayment curves

histogram of the estimates of slopes of take-up $\mathbb{E}[\gamma|z]$, evaluated at each z_i in the dataset. The right hand side figure is the histogram for the estimates of slopes of repayment rates $\mathbb{E}[\beta|z]$. We can see that estimated take-up elasticities are mostly above zero, with a mean of 0.59, and most of the estimates lying between zero and five. The slopes of repayment rates are much more spread out, showing evidence that in these rural credit markets, if all information were used, there would be both adverse selection ($\beta(z) > 0$), with higher prices leading to higher repayment rates, and advantageous selection ($\beta(z) < 0$), with higher prices leading to lower repayment rates.

To be able to compute the optimal credit scores, all we need are estimates of the joint distribution of values for investors and borrowers. Even though the estimated slopes and levels can be readily translated into this joint distribution, we take an additional step that is meant to make the problem computationally easier and interpretable. Given estimates of slopes and intercepts for supply and average value curves, we cluster them into k different groups, using a k-means algorithm[12]. The results of this estimation procedure will be used in the next section to give empirical substance to the optimal credit scores that were theoretically discussed above.

7 Empirical results

In this section, we combine the estimates for the joint distribution of values we presented in the previous section to arrive at the theoretically optimal credit scores. As shown in Section 3, it is without loss to identify messages with prices. We interchangeably refer to these messages or prices as credit scores.

Figure 10 illustrates the three rules of thumb we established for the optimal disclosure policy. We have shown that: markets with higher prices should have higher elasticities of the expected value for investors; in general, at most two signals should be combined in a

[12] Alternatively, we could for example have used the empirical distribution of z for the joint distribution. However, this would require later solving a very large linear programming problem.

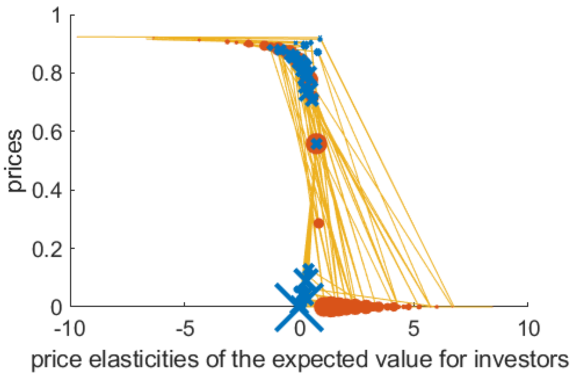

Figure 10: Optimal and full disclosure credit scores

single message; and when these two signals are combined into the same message, the one that has the higher full disclosure price should feature the lowest price elasticity of the value for investors.

These three rules of thumb imply that, in Figure 10, i) downward sloping yellow lines should connect two red dots to a blue cross, reflecting the second and third rules-of-thumb; ii) blue crosses should be upward sloping, reflecting the first rule-of-thumb. These rules-of-thumb, however, rely on local comparisons between adjacent prices and signals that do not need to hold when dots and crosses are further apart, and that can get blurred by the nature of the discrete approximation and the mechanics of linear programming algorithms [13]. In spite of it, we can see that most of yellow lines are downward sloping, they connect two red dots to a blue cross, and blue crosses are (mostly) upward sloping.

To compute welfare gains from the optimal policy relative to full disclosure, while avoiding counting the gains from overfitting, we use a simple form of sample splitting. We split the sample in two halves, repeating the procedure twenty times. At each time, in one half of the sample we compute the credit scoring rule, while in the second half, we apply this scoring rule and evaluate the welfare gains. This procedure shows that there would be monthly welfare gains from moving from full disclosure to the optimal policy of 0.45 % the size of a typical loan per household. This is approximately $3 per household per month.

To get a better sense of how optimal credit scores differ in practice from standard scores that would be exclusively targeted at predicting repayment probabilities, in Figure 11, we rank the variables in the dataset by their relevance in predicting repayment probabilities (red, on the left) and optimal credit scores (blue, on the right).

The relevance measure is a weighted sum of how many times each variable was used to create a split on the trees of the random forest, with weights proportional to the depth where the split was created. Interestingly, only two variables appear in the top ten most

[13]Moreover, the assumption that average value curves are piecewise linear implies that elasticities can change abruptly as prices change.

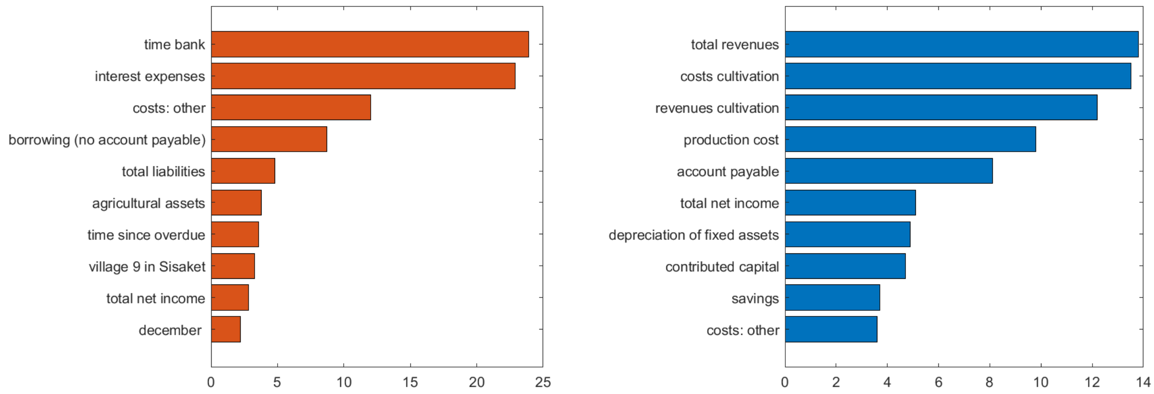

Figure 11: Variable importance comparison: Full disclosure vs optimal scores

important variables at both of these two rankings: "total net income" and "costs: other". The most important variables for predicting repayment probabilities include for how long the individual has been a client from the BAAC, and how long since the last repayment was overdue – which are the two variables that the BAAC uses to create its own rating for the farmers. These two variables do not appear on the top ten variables of the optimal credit score. Instead, variables from the balance sheet and cash flow statements of the farmers appear more prominently, with total revenues, and cultivation costs and revenues being the three most important. Surprisingly, the comparison of these two lists of variables seems to suggest that the optimal credit score would give higher weight to the solvency of farmers, instead of their short term liquidity conditions.

8 Conclusion

In this paper, we presented results on how to compute optimal information disclosure policies and, more generally, what they look like, focusing in on credit markets with asymmetric information.

We presented simple rules of thumb that describe the solution of the optimal disclosure problem in terms of local sufficient statistics that can be estimated. Additionally, we presented a portable empirical strategy to estimate these sufficient statistics relying on exogenous variation in interest rates[14]. We found that there are economically meaningful welfare gains from pursuing optimal disclosure policies in rural credit markets. Our estimates indicate that the monthly gains from moving from full disclosure to the optimal policy are of the order of 0.45% the size of a typical loan, or approximately $3 per household per month.

Our framework can be applied to other setups, where in principle one can think that

[14] A second empirical strategy, presented in Appendix C leverages the evidence of intensive risk sharing within villages in rural Thailand.

more information could ameliorate adverse selection. Besides credit scores, the method can be applied to markets ranging from health insurance to unemployment, disability insurance and workers compensation.

References

Christian Ahlin and Robert M Townsend. 2007. Using repayment data to test across models of joint liability lending. *The Economic Journal* 117, 517 (2007), F11–F51.

George A Akerlof. 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84, 3 (1970), 488–500.

Itai Arieli, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita. 2020. Optimal Persuasion via Bi-Pooling. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 641–641.

Susan Athey, Julie Tibshirani, and Stefan Wager. 2019. Generalized random forests. *The Annals of Statistics* 47, 2 (2019), 1148–1178.

Susan Athey and Stefan Wager. 2019. Estimating treatment effects with causal forests: An application. *Observational Studies* 5, 2 (2019), 37–51.

Andrea Attar, Thomas Mariotti, and François Salanié. 2011. Nonexclusive competition in the market for lemons. *Econometrica* 79, 6 (2011), 1869–1918.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2015. The limits of price discrimination. *American Economic Review* 105, 3 (2015), 921–57.

Dirk Bergemann and Stephen Morris. 2013. Robust predictions in games with incomplete information. *Econometrica* 81, 4 (2013), 1251–1308.

Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.

Jesse Bricker, Lisa J Dettling, Alice Henriques, Joanne W Hsu, Lindsay Jacobs, Kevin B Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A Windle. 2017. Changes in US family finances from 2013 to 2016: Evidence from the Survey of Consumer Finances. *Fed. Res. Bull.* 103 (2017), 1.

Marika Cabral, Can Cui, and Michael Dworsky. 2019. What is the Rationale for an Insurance Coverage Mandate? Evidence from Workers' Compensation Insurance. (2019).

Pierre-André Chiappori, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M Townsend. 2014. Heterogeneity and risk sharing in village economies. *Quantitative economics* 5, 1 (2014), 1–27.

Alma Cohen and Liran Einav. 2007. Estimating risk preferences from deductible choice. *American economic review* 97, 3 (2007), 745–788.

Jonathan Davis and Sara B Heller. 2017. Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review* 107, 5 (2017), 546–50.

Deniz Dizdar and Eugen Kováč. 2020. A simple proof of strong duality in the linear persuasion problem. *Games and Economic Behavior* 122 (2020), 407–412.

Piotr Dworczak and Giorgio Martini. 2019. The simple economics of optimal persuasion. *Journal of Political Economy* 127, 5 (2019), 1993–2048.

Liran Einav, Amy Finkelstein, and Mark R Cullen. 2010. Estimating welfare in insurance markets using variation in prices. *The quarterly journal of economics* 125, 3 (2010), 877–921.

Amy Finkelstein and Matthew J Notowidigdo. 2019. Take-up and targeting: Experimental evidence from SNAP. *The Quarterly Journal of Economics* 134, 3 (2019), 1505–1556.

Amy Finkelstein and James Poterba. 2004. Adverse selection in insurance markets: Policyholder evidence from the UK annuity market. *Journal of Political Economy* 112, 1 (2004), 183–208.

Daniel Garcia, Roee Teper, and Matan Tsur. 2018. Information design in insurance markets: Selling peaches in a market for lemons. *Available at SSRN 3024500* (2018).

Xavier Giné. 2011. Access to capital in rural Thailand: An estimated model of formal vs. informal credit. *Journal of Development Economics* 96, 1 (2011), 16–29.

Edward L Glaeser and Hedi D Kallal. 1997. Thin markets, asymmetric information, and mortgage-backed securities. *Journal of Financial Intermediation* 6, 1 (1997), 64–86.

Jonathan Gruber. 1997. The Consumption Smoothing Benefits of Unemployment Insurance. *The American Economic Review* 87, 1 (1997), 192–205.

Benjamin R Handel and Jonathan T Kolstad. 2015. Health insurance for" humans": Information frictions, plan choice, and consumer welfare. *American Economic Review* 105, 8 (2015), 2449–2500.

Nathaniel Hendren. 2013. Private information and insurance rejections. *Econometrica* 81, 5 (2013), 1713–1762.

Nathaniel Hendren. 2017. Knowledge of future job loss and implications for unemployment insurance. *American Economic Review* 107, 7 (2017), 1778–1823.

Julapa Jagtiani and Catharine Lemieux. 2019. The roles of alternative data and machine learning in fintech lending: evidence from the LendingClub consumer platform. *Financial Management* 48, 4 (2019), 1009–1029.

Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.

Navin Kartik and Weijie Zhong. 2019. Lemonade from Lemons: Information Design and Adverse Selection. (2019).

Anton Kolotilin. 2018. Optimal information disclosure: A linear programming approach. *Theoretical Economics* 13, 2 (2018), 607–635.

Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. Persuasion of a privately informed receiver. *Econometrica* 85, 6 (2017), 1949–1964.

Anton Kolotilin and Alexander Wolitzky. 2020. Assortative information disclosure. (2020).

Josh Lerner and Jean Tirole. 2006. A model of forum shopping. *American Economic Review* 96, 4 (2006), 1091–1113.

Jonathan Levin. 2001. Information and the Market for Lemons. *RAND Journal of Economics* (2001), 657–666.

Michael Ostrovsky and Michael Schwarz. 2010. Information disclosure and unraveling in matching markets. *American Economic Journal: Microeconomics* 2, 2 (2010), 34–63.

Luis Rayo and Ilya Segal. 2010. Optimal information disclosure. *Journal of political Economy* 118, 5 (2010), 949–987.

Krislert Samphantharak and Robert M Townsend. 2010. *Households as corporate firms: an analysis of household finance using integrated household surveys and corporate financial accounting.* Number 46. Cambridge University Press.

Robert M Townsend. 1994. Risk and Insurance in Village India. *Econometrica* 62, 3 (1994), 539–591.

Robert M Townsend. 1995. Consumption insurance: An evaluation of risk-bearing systems in low-income economies. *Journal of Economic perspectives* 9, 3 (1995), 83–102.

Stefan Wager and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* 113, 523 (2018), 1228–1242.

A Omitted definitions and propositions

As discussed in the main text, definition 1 corresponds to the outcome of a game where the investors offer arbitrary menus of contracts to the potential borrowers without exclusivity, under the assumption 1 below (Attar et al., 2011).

Assumption 1.

- *The distribution P_m of b, conditional on any signal m, has bounded support.*
- *$E[a|b,m]$ exists for every b and m and it is finite.*
- *If b is an atom of P_m, then $E[a|b,m] \geq b$*
- *If $x > x^*(m)$ then $\pi(x;m) < 0$, where $\pi(x;m) = \int_{b \leq x}(a-x)\ d\pi(a,b|m)$. In other words, at any higher price than the equilibrium price for the message m, investors would make a loss.*

Denote a contract investor i offers by a vector (t_i, q_i), where t_i is how much the borrower would receive in the first period, and q_i is how much the borrower would pay the investor back in the second period.

Proposition 6. *Under Assumption 1, the non-exclusive competition between the investors result in all contracts that are traded being of the form $t_i = x^*(m)\ q_i$, where $x^*(m) = sup(x|x = E[a|x \leq b,m])$ and investors exactly break even.*

Proof. Assumption 1 is exactly parallel to the assumptions in Attar et al. (2011) and the result follows from their Propositions 1 and 2. □

B Omitted proofs

Proposition 2. *Full disclosure is optimal if and only if*

$$\left(\sigma_{mj} - \sigma_{x(j),j} - \frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk} - \sigma_{x(k),k})\right) \leq 0 \quad \textit{for all } m,j,k \textit{ such that } \pi_{mj} > 0 \textit{ and } \pi_{mk} < 0.$$

Proof. A full disclosure configuration p is optimal if and only if for any direction $\Delta\overrightarrow{p} \in C$, where

$$C = \{\Delta\overrightarrow{p} \in \mathbb{R}^{n\times m} | \Delta p_{x(j),j} \leq 0,\ \Delta p_{m\neq x(j),j} \geq 0,\ \Delta p_{x(j),j} = -\sum_{m\neq x(j)} \Delta p_{mj} \text{ and } 0 = \sum_j \pi_{mj}\Delta p_{mj}\}$$

, the resulting change in the objective function should be non positive, that is, $\Delta W = d \cdot \Delta\overrightarrow{p} \leq 0$.

For all pair of indexes except those of the form $x(m), m$, divide the indexes in two groups, one such that $\pi_{mj} > 0$. Then construct vectors v_{mjk} such that for each index in the first group each vector v_{mjk} has a one entry at the index mj, a $-\frac{\pi_{mj}}{\pi_{mk}}$ entry for an index mk in the second group and an index -1 in the index $x(j), j$ and $\frac{\pi_{mj}}{\pi_{mk}}$ entry in the diagonal index $x(k), k$.

Notice that by construction $V \subset C$, that is, deviations in any direction of these vectors $v \in V$ are feasible. Thus if p is optimal, then any of these deviations should generate a non positive change in welfare, which is exactly the condition in the claim. This proves that it is necessary.

To prove sufficiency, first define S as the set of vectors v such that for each rs with $\pi_{rs} = 0$, there is vector $v_{rs} \in S$ with a plus one entry at rs and a minus one entry at $x(s), s$. Notice that the definition of competitive price implies that for the vectors in S, the resulting change in welfare is non positive. We are going to show that conical combinations of the vectors in $V \cup S$ generate all the feasible directions $d \in C$, starting from the full disclosure benchmark.

Notice that any direction $d \in C$, with a positive index mj with $\pi_{mj} > 0$ must have another entry mk with $\pi_{mk} < 0$, otherwise it is not a feasible direction.

The proof now proceeds by induction. Given a vector d with N positive entries r, s such that $\pi_{rs} \neq 0$, take a pair of entries mj and mk with $\pi_{mj} > 0$ and $\pi_{mk} < 0$ if there are any, and the corresponding vector $v_{mjk} \in V$ which has these positive entries. Define $\alpha_{mjk} = min(\sigma_{mj}, -\frac{\sigma_{mj}}{\pi_{mk}}\sigma_{mk})$, and build a new vector $\tilde{d} = d - \alpha_{mjk} \cdot v_{mjk}$. This new vector now has at least one zero entry less and it is still lies in C. Inductively, one can repeat the procedure until all entries rs with $\pi_{rs} \neq 0$ are zero. For the remaining entries with $\pi_{rs} = 0$ we can use the simpler vectors in S which have a plus one entry at rs and a minus one entry at $x(s), s$. We concluded implies that we can write d as a conical combination of vectors in $V \cup S$. Thus we conclude that the stated condition is sufficient.

□

Proposition 3. *Suppose there is an interval $[\underline{x}, \overline{x}]$ such that the full disclosure price $x^*(z)$ is dense over it. If full disclosure is optimal over this interval, then for z and z' such that $x(z) \in [\underline{x}, \overline{x}]$ and $x(z') \in [\underline{x}, \overline{x}]$, the signal with the higher price must have higher elasticity:*

$$x^*(z) > x^*(z') \Rightarrow \epsilon_{V,x}(x^*(z), z) \geq \epsilon_{V,x}(x^*(z'), z')$$

where $\epsilon_{V,x}(x^*(z), z)$ is the price elasticity of the value for investors.

Proof. For two signals z and z' with $x(z)$ close to $x(z')$ and $x(z) > x(z')$, the expression in proposition 1 can be rearranged as:

$$\frac{\frac{\partial \pi(x(z),z)}{dx}}{\frac{\partial \sigma(x(z),z)}{\partial x}} \geq \frac{\frac{\partial \pi(x(z'),z')}{dx}}{\frac{\partial \sigma(x(z'),z')}{\partial x}}$$

$$\text{Notice that: } \frac{\partial \sigma(x(z), z)}{\partial x} = P(z, b \leq x(z))$$

and $\dfrac{\partial\pi(x(z),z)}{dx} = P(z,b\leq x(z))\left[\dfrac{\partial E[a|b\leq x,z]}{\partial x}-1\right]+\dfrac{\partial P(z,b\leq x(z))}{\partial x}\left(E[a|b\leq x,z]-x\right)$

$$= P(z,b\leq x(z))\left[\frac{\partial E[a|b\leq x,z]}{\partial x}-1\right]$$

$$\Rightarrow \epsilon_V^z(x(z))\geq \epsilon_V^{z'}(x(z'))$$

□

Proposition 4. *No disclosure is optimal if and only if:*

$$\left(\sigma_{mj}-\sigma_{0j}-\frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk}-\sigma_{0k})\right)-\frac{\Delta_\pi(m,j,k)}{\Delta_\pi(m',j',k')}\left(\sigma_{m'j'}-\sigma_{0j'}-\frac{\pi_{m'j'}}{\pi_{m'k'}}(\sigma_{m'k'}-\sigma_{0k'})\right)\leq 0$$

for all m,j,k,m',j',k' *such that* $\pi_{mj},\pi_{m'j'},\Delta_\pi(m,j,k)\geq 0$, *and* $\pi_{mk},\pi_{m'k'},\Delta_\pi(m',j',k')<0$.

If there are no m,j,k,m',j',k' *such that* $\pi_{mk},\pi_{m'k'},\Delta_\pi(m',j',k')<0$, *then no disclosure is optimal if and only if:*

$$\left(\sigma_{mj}-\sigma_{0j}-\frac{\pi_{mj}}{\pi_{mk}}(\sigma_{mk}-\sigma_{0k})\right)\leq 0$$

for all m,j,k,m',j',k', *such that* $\pi_{mj},\pi_{m'j'},\Delta_\pi(m,j,k)=0$.

Proof. The proof follows the same logic of Proposition 3. The expression is the change in welfare that results from deviating from no disclosure to a feasible direction. This vector combines two other generally infeasible vectors of the same form of the previous proposition. The first has has an one entry in the mj position, a minus one entry in the $0j$ position, a $-\frac{\pi_{mj}}{\pi_{mk}}$ entry in the mk position and $\frac{\pi_{mj}}{\pi_{mk}}$ entry in the $0k$ position. The second has analogous entries in the $m'j'$, $0j'$, $m'k'$ and $0k'$ positions. They are generally infeasible because π_{0s} is generally different than zero, while in the full disclosure case $\pi_{x(s),s}$ is zero by definition. By definition this infeasible vector should increase the expected value for investors conditional on the no disclosure price, while the second infeasible vector should decrease it. Then those are combined in the right proportion so that the zero profit constraint of the investors holds with equality, that is, at the ratio $-\frac{\frac{\pi_{mj}}{\pi_{mk}}\pi_{0k}-\pi_{0j}}{\frac{\pi_{m'j'}}{\pi_{m'k'}}\pi_{0k'}-\pi_{0j'}}$. This shows that this direction is feasible, and therefore the condition is necessary. The qualification says that if any of these ratios turn to be zero, then it is not necessary to find another entry or vector that would compensate for the violation in the zero profit condition. Observe that if on the other hand there is an index with a strictly positive π_{mj} but for this price there is no signal such that $\pi_{mk}<0$, then it is not feasible to increase mj.

Now, to prove sufficiency, notice that any feasible direction can be decomposed as conical combinations of these directions, using the same argument in the proof of Proposition 2.

□

Proposition 5. *All signals z that are pooled together in x_0 must satisfy:*

$$\frac{x_0}{E[a|b \leq x_0, z] - x_0} = \kappa(x_0)\epsilon_{P,x}(x_0, z) + \gamma(x_0)$$

for some constants $\kappa(x_0)$ and $\gamma(x_0)$, and where $\epsilon_{P,x}(x_0, z) = \frac{\partial \pi(x_0,z)}{\partial x}\frac{x_0}{\pi(x_0,z)}$

Proof. The condition on proposition 4 can be rewritten as:

$$\frac{\left(\frac{\sigma_{mj}-\sigma_{0j}}{\pi_{mj}} - \frac{\sigma_{mk}-\sigma_{0k}}{\pi_{mk}}\right)}{\frac{\pi_{0k}}{\pi_{mk}} - \frac{\pi_{0j}}{\pi_{mj}}} \leq \frac{\left(\frac{\sigma_{m'j'}-\sigma_{0j'}}{\pi_{m'j'}} - \frac{\sigma_{m'k'}-\sigma_{0k'}}{\pi_{m'k'}}\right)}{\frac{\pi_{0k'}}{\pi_{m'k'}} - \frac{\pi_{0j'}}{\pi_{m'j'}}}$$

By assumption, $\pi_{mj} > 0$, and $\pi_{mk} < 0$. If we take x_m to be close to x_0, then $\pi_{0j} > 0$ and $\pi_{0k} < 0$. At the same time, taking this limit with $x_m > x_0$, implies that the denominator on the left side of the inequality is proportional to $\epsilon_{P,x}(x_0, z) - \epsilon_{P,x}(x_0, z')$, which by assumption is positive. Notice, however, that taking the under a sequence where $x_m < x_0$ would flip the sign of the denominator, which implies that the same pair of signals could play the role of j' and k' as long as x'_m approaches x_0 from below.

With that in mind, and taking the limit in both sides of the inequality we conclude that the condition can be written as:

$$\frac{\frac{x_0}{E[a|b\leq x_0,z_j]-x_0} - \frac{x_0}{E[a|b\leq x_0,z_k]-x_0}}{\epsilon_{P,x}(x_0, z_j) - \epsilon_{P,x}(x_0, z_k)} = \frac{\frac{x_0}{E[a|b\leq x_0,z'_j]-x_0} - \frac{x_0}{E[a|b\leq x_0,z'_k]-x_0}}{\epsilon_{P,x}(x_0, z'_j) - \epsilon_{P,x}(x_0, z'_k)}$$

Which implies that

$$\frac{\frac{x_0}{E[a|b\leq x_0,z]-x_0} - \frac{1}{E[a|b\leq x_0,z']-x_0}}{\epsilon_{P,x}(x_0, z) - \epsilon_{P,x}(x_0, z')} = k_{x_0}$$

for some constant k_{x_0}. Further, we can rewrite this expression as:

$$\frac{x_0}{E[a|b \leq x_0, z] - x_0} = \kappa(x_0)\epsilon_{P,x}(x_0, z) + \gamma(x_0)$$

□

C Alternative identification strategy: Within village risk-sharing

In this section, we present our alternative empirical strategy to estimate the joint distribution of values for borrowers and investors. The empirical strategy leverages the evidence that there is intensive consumption risk sharing at the village level (Chiappori et al., 2014). We explore this idea, with a simple, more structural, model: individuals share consumption risk, but each has its own cut-off for which if the marginal utility of consumption at the second period is above it, she defaults.

Concretely, we use data on the time series of consumption to estimate the distribution of marginal utilities, [15]. We will assume that loans are infinitesimal, so that marginal utilities of consumption do not change with borrowing decisions, both at the individual level and at the village level. Then we turn to data on default and take-up to to estimate distribution of cut-offs.

With some convenient parametric assumptions, the model is just identified without exogenous variation in interest rates[16]. As such, to the extent that we may worry about the changes in interest rates by the BAAC not being completely exogenous to the demand and creditworthiness of the borrowers, this approach provides an alternative source of identification.

The model will generate two key equations defining the reservation values for borrowers and for investors in terms of types (that is the idiosyncratic cutoffs that trigger default) and the empirical processes for marginal utilities of consumption.

At the village-level, a household type ω is associated with their their idiossyncratic cutoff in terms of marginal utilities of consumption that would trigger default $l(\omega)$. Writing this cutoff in units of current period marginal utilities of consumption we can express the reservation value of a potential borrower as:

$$b(\omega) = \delta E_y \left[min \left\{ y, \tilde{l}(\omega) \right\} \right]$$

where $\tilde{l}(\omega)$ is the idiosyncratic penalty and the cutoff that triggers default, y is common at the village-level ratio of marginal utilities of consumption between period 1 and 2, and δ is the discount rate. Intuitively, because the household does not repay the loan when the marginal utility is of consumption is higher, preferring instead to incur the non-pecuniary penalty, the loan works as an insurance device. It insures them against the states of the world where the village-level income is low, and marginal utility of consumption are high.

Assuming the investors are risk neutral, their reservation value, conditional on ω can be written as:

[15] We will assume, for simplicity, a common power utility, with risk aversion coefficient of three

[16] With exogenous variation in interest rates, as in the previous section, the model is non-parametrically identified

$$a(\omega) = \delta P_y\left(y < \tilde{l}(\omega)\right)$$

That is, for investors, the value of a unit claim to be paid in the future is the discount factor, times the probability of repayment. The probability of repayment is just the probability that the marginal utility of consumption is lower than the idiosyncratic cutoff that triggers default.

More formally, these expressions are derived from the following model: There are two periods. In the first period, each household decides whether to borrow from investors. After these decisions, the resources are pooled and a centralized decision is made with respect to consumption. In the second period, uncertainty in wealth and preferences are realized, and households decide whether to default or not, and if they default, they suffer a non-pecuniary idiosyncratic penalty. Then, within village risk-sharing takes place, that is, a centralized planner makes the second round of consumption decisions subject to a feasibility constraint.

That is, in each period $t = 1, 2$, the planner solves

$$\max_{c_t^h} \sum \lambda^h u^h(c_t^h) \quad \text{s.t.} \ \sum c_t^h = W_t$$

The solution of this problem implies that $\lambda^h u^{h'}(c_t^h) = \mu_t$ and therefore $c_t^h = u^{h'^{-1}}\left(\frac{\mu_t}{\lambda^h}\right) = f^h(W_t)$, where μ_t is the Lagrange multiplier on the village resource constraint, and we can define $U^h(W_t) = u^h(f^h(W_t))$. For all households and every state of nature the rate of marginal utilities at time t and t' is the same, that is, there is full consumption risk sharing.

Households individually decide whether to default or not, facing a household specific cost k^h per unit of loan. Default decisions happen before consumption is realized, and the resources the household did or did not pay are added to the village level wealth W_t. Therefore the indebted household with debt b solves at the second period:

$$\max_{repay,\ default} \{U^h(W_2 - q), U^h(W_2) - l^h q\}$$

Where W_2 incorporates the equilibrium responses for the other households. This formulation implies that the equilibrium strategy is such that each household has a particular cutoff for village wealth below which she defaults. In the limit of a small loan q, the household will default if $U'^h(W_t) > l^h$. Moreover, if the total volume of loans at the village level also goes to zero, then the process for marginal utilities of consumption do not depend on the disclosure policy. If we know the distribution of penalties l^h, and the distribution of marginal utilities we can compute the distribution of default probabilities and reservation values for investors. Assuming they are risk neutral and that they have a common discount factor δ, the reservation value for investors, conditional on the household type ω, is the probability of repayment times the discount factor, that is $a(\omega) = \delta P(U'^{h,\omega}(W_t) \leq l^{h,\omega})$.

Likewise, we can find the reservation value for the potential borrowers. A borrower selling claims q and receiving t has an utility function:

$$V^h(b,t;\omega) = U_1^{h,\omega}(W_1+t) + \delta E\left[max\{U^{h,\omega}(W_2-q), U^{h,\omega}(W_2) - l^{h,\omega}q\}\right]$$

Which in the limit of a small loan can be written as:

$$\tilde{V}^h(b,t;\omega) = t - \delta E\left[min\left\{\frac{U'^{h,\omega}(W_2)}{U'^{h,\omega}(W_1)}, \tilde{l}^{h,\omega}\right\}\right] q \quad \text{where } \tilde{l}^{h,\omega} = \frac{l^{h,\omega}}{U'^{h,\omega}(W_1)}$$

Thus, the reservation value for a borrower of type ω is $b(\omega) = \delta E\left[min\left\{\frac{U'^{h,\omega}(W_2)}{U'^{h,\omega}(W_1)}, \tilde{l}^{h,\omega}\right\}\right]$

Given this formulation we can relate moments in the data to the moments in the model by the following set of equations.

First, the observed price x can be related to type that is just indifferent $\bar{l}$ between borrowing or not by the equation:

$$x(z) = \delta E_y\left[min\left\{y, \bar{l}\right\}|z\right]$$

Second, all types below this cutoff $\bar{l}$ have a lower reservation value, and therefore will sell the claim. Thus the observed takeup rates ($takeup(z)$), are related to the model moments by the following equation:

$$takeup(z) = P_l(l \leq \bar{l}|z)$$

Third, the average repayment rate is given by the average probability that marginal utility of consumption is below the idiosyncratic cutoff of those who decided to borrow:

$$repayment(z) = P_{l,y}(y \leq l | l \leq \bar{l}, z)$$

With additional parametric assumptions, these three equations are going to be sufficient to identify the joint distribution of values. [17] We will assume y and l are each log normally distributed. Thus, conditional on the distribution of marginal utilities of consumption, for each z and x, we will have three equations and three unknowns ($\bar{l}(x|z)$, $\mu_{l|z}$, $\sigma_{l|z}$).

The whole model has five moments and five parameters for each z. Two moments and parameters are associated with the data on consumption – the mean and variance of

[17] If we jointly consider the model equations and assume there is exogenous variation in prices, than we would be overidentified. Indeed, without the structural assumptions, we are already non-parametrically identified.

marginal utility of consumption. Then there are the mean and variance of cutoffs which are associated with take-up and repayment rates, and an incidental parameter which is the type who is just indifferent $\bar{l}(x|z)$ at a price x. These last three are the moments we are referencing here.

Concretely, the first equation above, that defines the type who is just indifferent $\bar{l}(x|z)$ at a price x, becomes:

$$x(z) = \delta \int_{-\infty}^{\bar{l}(x|z)} exp(w)\phi\left(\frac{w-\mu_{u|z}}{\sigma_{u|z}}\right)dw + \delta exp(\bar{l}(x|z)) \cdot \left(1 - \Phi\left(\frac{\bar{l}(x|z)-\mu_{u|z}}{\sigma_{u|z}}\right)\right)$$

The second equation, describing the take-up rate, becomes:

$$takeup(z) = \Phi\left(\frac{\bar{l}(x|z)-\mu_{l|z}}{\sigma_{l|z}}\right)$$

And the third equation, describing the average repayment rate, becomes:

$$repayment(z) = \int_{-\infty}^{\bar{l}(x|z)} \Phi\left(\frac{w-\mu_{u|z}}{\sigma_{u|z}}\right)\frac{\phi\left(\frac{w-\mu_{l|z}}{\sigma_{l|z}}\right)}{\Phi\left(\frac{\bar{l}(x|z)-\mu_{l|z}}{\sigma_{l|z}}\right)}dw$$

Having related the theoretical and empirical moments, it remains to estimated the latter. That is, we need to estimate take-up, and repayment rates as functions of the observable characteristics z. To accomplish this task, we use random forests (Breiman, 2001; Athey and Wager, 2019). [18].

These moments are translated into estimated functions $\mu_l(z)$, $\sigma_l(z)$, $\mu_u(z)$, $\sigma_u(z)$, using the equations above. As in the previous estimation strategy, even though we could from these estimators infer the joint distributions of values and observable characteristics,we take an additional step that is meant to make the problem computationally easier and interpretable. To that goal, we use a k-means algorithm to classify households into different clusters, according to their estimated $\mu_l, \sigma_l, \mu_u, \sigma_u$.

Given the estimates for the joint distribution of values, we can compute optimal credit scores as in the section 7. Figure 12 compares these scores to the ones we derived using exogenous variation in interest rates. We can see that credit scores and elasticities of the value for investors are roughly centered around the means, but much less spread out.

[18]Random forests had the best out-of-sample performance among a variety of other machine learning methods, such as boosted trees, lasso, and linear random forests.

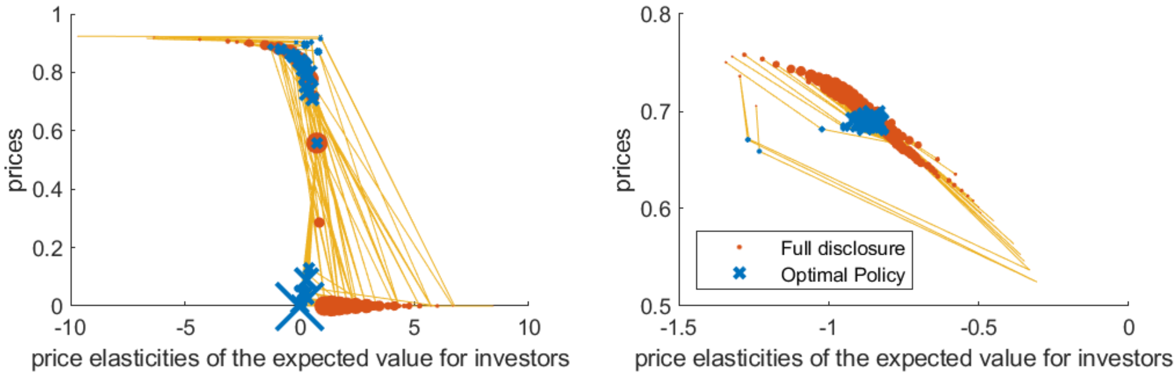

Figure 12: Reduced form and structural model