



Networked Markets and Simple Relational Contracts

Mathieu V Leduc

► To cite this version:

| Mathieu V Leduc. Networked Markets and Simple Relational Contracts. 2023. halshs-03948513

HAL Id: halshs-03948513

<https://shs.hal.science/halshs-03948513>

Preprint submitted on 20 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



WORKING PAPER N° 2023 – 05

Networked Markets and Simple Relational Contracts

Mathieu V. Leduc

JEL Codes: D00, D8, C7, L14

Keywords: Simple Relational Contracts, Shock Persistence, Principal-Agent Problems, Repeated Games, Social Capital.



Networked Markets and Simple Relational Contracts*

Mathieu V. Leduc[†]
Paris School of Economics
& Université Paris 1

January 2023

Abstract

Empirical studies of commercial relationships between firms reveal that (i) suppliers encounter situations in which they can gain in the short run by acting opportunistically and (ii) good conduct is sustained not exclusively by formal contracts but through informal relationships and the expectation of future business. In such relationships, the need to offer each supplier a large enough share of future business to deter cheating limits the number of supply relationships each buyer can sustain. The market thus becomes networked, with trade restricted to durable relationships. This article proposes a dynamic model to examine the structure of such *simple* overlapping relational contracts in equilibrium. These informal relational contracts are simple in the sense that they respect a set of realistic properties that can be easily understood and audited by both parties. Due to exogenous stochastic shocks, suppliers are nevertheless not always able to make good on their promises even if they wish to and thus relationships are constantly dissolving, leading to a destruction of social capital. New relationships are forming to take their place, but restoring social capital takes time due to search frictions. This creates a novel way for shocks to be persistent and also suggests new connections between the theory of relational contracting, on the one hand, and the macroeconomic analysis of recessions, on the other.

Keywords: Simple Relational Contracts, Shock Persistence, Principal-Agent Problems, Repeated Games, Social Capital

JEL codes: D00, D8, C7, L14

*I thank seminar participants at Cambridge University, the Paris School of Economics, the University of Paris 1, the Hungarian Academy of Science (Budapest) and WINE 2017: The 13th Conference on Web and Internet Economics (Bangalore, India). I thank Matthew Elliott (Cambridge) and Benjamin Golub (Northwestern) for very useful discussions on an earlier version of this project. I thank the European Research Council (ERC) (grant agreement number #757229) and the Cambridge-INET Institute. I also thank Grant ANR-17-EURE-0001 of the Agence Nationale de la Recherche (ANR).

[†](Corresponding author) Email: mattvleduc@gmail.com. Mailing address: Paris Jourdan Sciences Economiques, 48 boulevard Jourdan, 75014 Paris, France. Tel.: (+33) 01 80 52 16 60

1 Introduction

Relationships play a crucial role in the functioning of many markets. Insurers and other market participants tend to form a small set of long-term relationships with dealers in the over-the-counter corporate bond market (Hendershott, Li, Livdan, and Schürhoff (2020)). This allows them to trade off the value of repeated relationships with fewer dealers and the benefits of faster execution with more dealers. In the overnight lending market that banks rely on for liquidity, the same banks trade repeatedly with each other (Bräuning and Fecht (2012)). Agricultural traders in Madagascar rely heavily on social relationships to conduct business (Fafchamps and Minten (1999)). Medieval Maghribi traders utilized their social network to trade over long distances (Greif (1993)). Manufacturers in the Japanese electronic industry rely on durable relationships with their suppliers (Nishiguchi (1994)), and so on.

Building on the pioneering work of Granovetter (1985), an extensive literature in sociology shows that a primary function of these relationships is to help overcome short-term incentives to act opportunistically (see, for example, Uzzi (1997)). Common understandings of acceptable actions seem to limit opportunistic behavior and stabilize trade (e.g., McMillan and Woodruff (1999))¹. The relational contracting literature in economics has developed tools for modeling these kinds of informal understandings. However, with some important exceptions that we will discuss below, these models are typically limited to analyzing a single relationship or else the relationships of a single individual (say, buyer), and have stopped short of providing a theory of networked markets. This paper proposes model of market relationships made up of overlapping relational contracts.

Viewing business relationships as social capital, and formalizing the role they play in facilitating trade, raises new questions. What is the dynamic impact of shocks on the stock of social capital? Does the mistrust associated with turbulent economic times destroy social capital, and if so, how does the structure of the relational contracts relate to the extent of destruction? The typical approach economists have taken in modeling the formation of networked markets does not permit such questions to be addressed. For example, Kranton and Minehart (2001) and Elliott (2015) model relationships as partner-specific costly investments, abstracting from the roles relationships play in sustaining cooperative behavior. This limits the ability to consider the dynamic response of the network to shocks and policy interventions. Taking a more dynamic approach, based on relational contracting, complements the parsimonious static models.

We now describe the essentials of the model. Each of many principals (buyers) needs the services of one agent (supplier) every period. A buyer observes who has the capacity to fulfill his order and chooses one agent to trade with. The model features moral hazard: because contracts are incomplete and cannot cover every contingency, agents can act opportunistically (“defect”) and seize surplus at the buyer’s expense. In an incentive-compatible relational contract, to prevent such behavior, an agent’s good conduct must yield an expectation of trading sufficiently frequently in the future. That is, the principal must promise future business after successful trade, and the expected value of that business must exceed the short-term gain of opportunism. The need to promise a large enough share of future surplus limits the number of agents with whom a principal can sustain successful trade. Thus, the market becomes networked, with each supplier having a subset of *insiders* who currently enjoy favored status. The final key ingredient of the model is that exogenous shocks may cause a

¹Measuring trust by the size of trade credit, McMillan and Woodruff (1999) find evidence for relationships replacing formal contracts in Vietnam.

trade to fail through no fault of the supplier's. When a shock hits, even if the agent acted in good faith and did not defect, the outcome observed by the principal is as if the supplier had defected. Because the principal cannot distinguish such shocks from opportunistic cheating, the shocks must also lead to a disruption in the supplier's favored status; otherwise there would be no deterrent to defection. Thus, insider relationships are sometimes destroyed. We posit a frictional process of principals' matching to new agents as a channel that can replace the lost relationships.

In general, the best relational contract for a principal will be complex, rely on the fine details of the history, and involve randomizations, thus treating identical situations quite differently. Such relational contracts have two disadvantages. First, they stretch to an implausible extent the important practical constraint that contracts must be underpinned by common understandings of expected conduct. Second, they require the principal to be able to commit to intricate credible randomizing devices—something that is difficult to audit and rarely observed in practice. Instead of characterizing optimal contracts with these characteristics, we define a class of *simple contracts* that are easy for agents to understand and audit. The conditions we use to do this are intuitive and, when the principal optimizes within this class, lead to natural relational contracts. At the same time, the restrictions placed on the contracting space make the search for an optimal contract tractable, and allow us to study its characteristics.

The contracts that are found to be optimal within this restricted class are characterized by three main features. First, principals maintain relationships with a set of insiders, who are treated equally and receive priority as trading partners. Second, when no such insider is available for trade, the principal instead searches for an outsider to trade with as long as the current insider set is not too large. Finally, successful trade by an agent is always rewarded by insider status, while unsuccessful trade is always punished with outsider status. This simple structure is consistent with stylized facts. Given these contracts, new relationships are formed and existing relationships lost in a dynamic process creating a dynamic networked market, which we study. Formally, we analyze a Markov process on a set of networks that the relational contracts induce. We show how this random bipartite graph depends on the parameters of the problem.

The main questions concern how the behavior of the dynamic market depends on exogenous disruption. There are two ways in which such disruptions matter. First, when shocks are realized, there is a destruction of social capital, and restoring it takes time due to search frictions. Beyond the mechanical delay involved in meeting new agents, maintaining incentive-compatibility also implies prioritizing the remaining relationships, which delays the replacement of terminated relationships. This creates a novel source of persistence for shocks. Second, if the probability of shocks increases, the value of future trade decreases, and buyers are able to maintain fewer relationships, because they have less overall future surplus to divide. These forces suggest new connections between the theory of relational contracting, on the one hand, and the macroeconomic analysis of recessions, on the other. Relational contracts disrupted by shocks can then be viewed as a sort of “social” or network capital—an important input into production that is lost suddenly and replaced gradually.

There is a large related empirical literature in economics and sociology documenting the importance of relationships in markets. Beyond the examples already mentioned, Kirman and Vriend (2001) document that restaurateurs, particularly those from higher-end restaurants, purchase repeatedly from a small number of vendors. Petersen and Rajan (1994) and Petersen and Rajan (1995) find that the relationships between small firms and institutional creditors are crucial for gaining

access to financing and that firms attempting to access many creditors face higher interest rates and end up having less access to credit. Similarly, Cocco, Gomes, and Martins (2009) document the importance of relationships for accessing interbank lending markets and find that these relationships are especially important for distressed and smaller banks.

We now discuss the most closely related theoretical work. A few papers in the relational contracting literature have considered the relationships between a producer and his suppliers. The most closely related among these are Board (2011) and Andrews and Barron (2016). Board (2011) finds that an organization may prefer to maintain a limited number of strong relationships with suppliers rather than trading with the lowest-cost supplier in each period. This is the inspiration for the notion of insiders presented here². Calzolari and Spagnolo (2009) also develop a dynamic model of recurrent exchange of a non-contractible task between a principal and a population of agents and note the optimality of limiting the number of trading partners, although in a very different model. Andrews and Barron (2016) show that it can be optimal for an organization to favor suppliers that it has more recently traded with over others. In all these cases, the focus is on the interaction of a single principal (e.g., an organization) with several agents (e.g., suppliers). Moreover, the optimal contract may be complex and depend intricately on the history of play.

This article is also related to the static mechanism design literature on competing mechanisms and to the dynamic games literature. Samuelson and Stacchetti (2017) consider repeated relationships and use of transfers in sustaining them. Several recent papers also consider cooperation in networks, e.g. Ali and Miller (2013), Ali and Miller (2016), Wolitzky (2013) and Nava and Piccione (2014). For a survey, see Nava (2015). A key difference between these papers and the present one is that the networks considered here are dynamic and endogenously determined by relational contracts rather than exogenous constraints on interactions that are taken as given.

The need to study a simple, clearly understandable contracting setting is a departure from the standard contracting literature. The importance of simple strategic settings to sustain cooperation has been noted in both empirical and theoretical literature. For example, Kandori and Obayashi (2014) use a restricted framework, based on simple heuristics, to explain how cooperation is sustained. They show that it fits quite well empirical surveys of cooperation among labor union members in Japan. The importance of simple and understandable mechanisms, that can be easily understood by players (even cognitively limited), is also made clear namely in Li (2017). Indeed, human beings tend to play their equilibrium strategies in simple mechanisms, but often make mistakes in complex mechanisms. This can affect the performance of complex optimal mechanisms in ways that are difficult to predict.

Finally, there is a literature on macroeconomics and social capital (e.g. Sobel (2002), Dasgupta and Serageldin (2001)). This literature does not provide microfoundations for social capital, while this article does. Moreover, the microfoundations provided here are crucial for understanding the aggregate dynamic impact of shocks on social capital.

The paper is structured as follows. Section 2 introduces the model, focusing on the interaction setting and the timing of decisions. Payoffs and incentive-compatibility are analyzed in a general contracting space. In Section 3, simpler contracts are defined and their basic characteristics are

²Note that there are many key differences between the model presented in this paper and that of Board (2011). These will be made clear throughout the paper. Namely, there are on-path errors due to exogenous shocks, meaning that agents can lose insider status on the equilibrium path, whereas in Board (2011) agents can lose insider status only off-path.

described. Section 4 provides an equilibrium characterization in this simpler contracting space and derives comparative statics and welfare results. Section 5 extends the model to a regime-switching setting, where we can examine the rapid destruction of social capital during a crisis and the slow rebuilding of market relationships during the recovery. Section 6 provides some robustness checks, while the proofs are relegated to the Appendix.

2 Model

There is a continuum of principals and a continuum of agents, also called *suppliers*. We normalize the measure of principals to 1 and let x be the measure of agents. There are infinitely many discrete time periods $t \in \{0, 1, 2, \dots\}$. The principal and all agents have a utility equal to the discounted sum (with discount rate $\delta \in (0, 1)$) of period payoffs, which we are about to define.

At each time, each principal j is associated with a set of agents \mathcal{I}_t^j called *insiders*, and if we interpret this set as the neighborhood of j , the neighborhoods define a bipartite network of trading relationships.

2.1 Timing within a period and payoffs

2.1.1 Timing and primitives

We now present the timing of the model within each period t .

- (i) Each principal is matched to one *outsider* uniformly at random³.
- (ii) Each agent is either *available* to trade or not. The probability that any agent is available is α . The availability realizations are independent across agents and observable by the principal, all insiders, and the outsider.
- (iii) Simultaneously, each principal chooses at most one agent to trade with: this may be an available insider, the matched outsider, if available, or nobody. The indicator variable corresponding to principal j selecting agent i at time t is denoted by $Q_{it}^j \in \{0, 1\}$, where $\sum_i Q_{it}^j \leq 1$ for each j . If j selected i , we also say that j *initiated* trade with i .
- (iv) For each principal-agent pair (j, i) with $Q_{it}^j = 1$ —i.e., where the principal has selected the agent—a *shock* is realized with probability q . In case the shock is realized, the agent cannot produce for the principal. The surplus \tilde{v}_t available in that relationship in that period is thus

$$\tilde{v}_{it}^j = \begin{cases} 0 & \text{w.p. } q \\ v & \text{w.p. } 1 - q, \end{cases}$$

where $v > 0$ is a parameter of the model. The realization of \tilde{v}_{it}^j is observed by the selected agent, i , at time t , but not by the principal j .

- (v) The selected agent chooses whether to *defect* in period t on principal j or not. If the agent defects, he receives \tilde{v}_{it}^j , leaving the principal with nothing. If the agent does not defect (the

³There is no limit on how many principals an agent may be matched to through this process.

agent *cooperates*), he receives

$$\tilde{p}_{it}^j = \begin{cases} 0 & \text{if } \tilde{v}_{it}^j = 0 \\ p & \text{if } \tilde{v}_{it}^j = v, \end{cases}$$

where p is a parameter of the model. We assume $p < v$.

- (vi) The principal observes the amount left to him by the agent, $\tilde{v}_{it}^j - \tilde{p}_{it}^j$, and if either the shock has occurred or if the agent has defected (i.e., if $\tilde{v}_{it}^j - \tilde{p}_{it}^j \neq v - p$), the principal incurs a cost $c > 0$. The principal then selects his insider set for next period, \mathcal{I}_{t+1}^j . This consists of an arbitrary subset of \mathcal{I}_t^j and can also include the outsider, if the principal traded with the outsider.

At the end of each period, the actions of the principal and success of trade are observed by the principal, the insiders, and the outsider. Interactions within a principal's neighborhood are not observed by anyone outside it, and there is no cheap talk technology⁴.

2.1.2 Discussion

The model features moral hazard: because contracts are incomplete—not every contingency can be described in the contract—there is an opportunity for the agent to act to benefit himself at the principal's expense. This is reflected in the model by the ability of an agent to ‘hold up’ the principal, i.e. he can keep the produced good of value $v > p$ for himself and the principal is not able to directly observe whether this is due to cheating or to the production shock hitting with probability q .

For the moment, the notion of insiders simply captures the idea that the principal is not able to contact every potential agent, and can trade at most with those he has met before⁵. In the class of contracts we study, the insiders correspond to the suppliers who have priority in being considered to trade at a given time. The properties defining the restricted class of contracts we study will give additional operational content to this notion.

The key tool a principal has for incentivizing a selected agent not to defect is the choice of whether to retain that agent as an insider for the next period. This captures the idea that buyers can choose to fire a preferred supplier when trade with that supplier is unsuccessful.

The notion of availability corresponds to the idiosyncratic realization of whether a given supplier has the skills and resources to fulfill the buyer's needs. When there is no available supplier among a principal's insiders, the principal can search for an outsider supplier. To capture the idea that this search is not frictionless, we assume that in a given period the principal is only able to sample one outsider. But this outsider also faces the same kind of availability draw that the insiders do: he may not be able to take the buyer's order.

The cost c incurred by the principal when either the shock or defection occurs might reflect lost upfront payments made to the agent that cannot be recovered (another aspect of incomplete contracts), revenues that the principal loses, or costs the principal incurs because his suppliers do not deliver.

⁴Our continuum assumptions effectively rule out the prospect of interesting “communication” or “reporting” strategies of any kind, even through actions. No player in a given interaction will meet, with positive probability, anyone that his current partners will ever meet.

⁵Note the principal may retain every agent ever encountered in his insider set, which simply acts as a business contacts list in this case.

2.1.3 Efficiency benchmark

The constrained-efficiency benchmark we will consider is one in which the agent cannot defect and thus trade will occur if there is no shock. In this case, it is optimal for a principal never to fire any agents, to add every outsider he meets to his insider set and thus eventually to have an arbitrarily large insider set. This makes it overwhelmingly likely, in the long run, that at least one agent is available and the expected surplus per principal achieved per period is $v(1 - q) - qc$.

2.2 General contracts and the agent's problem

Next, we formalize the incentives facing the agents and the principal. Here and at many points throughout, we will focus on a particular principal, dropping the superscript j . Recall Q_{it} is an indicator variable specifying whether the principal selects agent i in period t or not. The principal's history at the end of period t is denoted by $h^t = \{h_\tau\}_{\tau \leq t}$. Because the principal's actions and the success of trade are publicly observed, $h_\tau = (\mathcal{I}_\tau, A_\tau, i, \tilde{v}_\tau - \tilde{p}_\tau)$, where \mathcal{I}_τ is the set of insiders of the principal in period τ , the set A_τ contains the agents available to the principal in that period, the index i indicates the identity of the agent selected by the principal (and may be equal to \emptyset , corresponding to nobody), and $\tilde{v}_\tau - \tilde{p}_\tau \in \{v - p, 0\}$ reveals the success of the trade⁶.

A *relational contract* Γ is a pair (γ, η) along with a strategy for the agent. Here, γ describes the selection strategy: it is a mapping from $h^{t-1} \cup \{\mathcal{I}_t, A_t\}$ into the space of probability distributions over available agents, $\Delta(A_t)$. The other component, η , describes the retention strategy, and is a mapping from h^t into (i) distributions over subsets of $\mathcal{I}_t \cup \{O\}$, when trade in period t was with the outsider, or (ii) distributions over subsets of \mathcal{I}_t otherwise. This identifies which agents will be insiders for the next period. This completes the description of the principal's strategy. A strategy for the agent, which is also part of a relational contract, specifies an action for every agent history. We will sometimes omit the description of an agent's behavior when it is clear what it should be.

Suppose an agent i is selected to trade with a particular principal in period t . To make the agent's problem separable across principals, we make our first assumption:

Assumption 1 (Infinite supplier capacity) *Each supplier (agent) can simultaneously supply any number of principals, and payoffs within a period are additively separable across principals.*

This infinite supply assumption simplifies the analysis when agents are selected by multiple principals at once. It means that principals need not care about the employment status of the agents, and whom else they are matched to as outsiders. From the supplier's perspective, the combination of unlimited capacity and additively separable payoffs means that each agent can optimize separately in every relationship with a principal. There are interesting implications of relaxing this assumption and they are explored in Section 6.

Now, returning to our selected agent i , let $V_i(h^t)$ be i 's continuation value, given a relational contract, after history h^t . If no shock hits the agent (which happens with probability $1 - q$), then that agent's value from attempting to trade with the principal is

$$p + \delta V_i(h^t),$$

⁶Here an outsider who ends up unavailable is considered unobserved; this makes no difference to any of the results.

where $h^t = h^{t-1} \cup \{\mathcal{I}_t, A_t, i, v - p\}$ and $\delta < 1$ is the discount factor representing time preferences. If agent i instead defects after being selected to trade with the principal, the agent receives

$$v + \delta V_i(h^t),$$

where $h^t = h^{t-1} \cup \{\mathcal{I}_t, A_t, i, 0\}$.

Likewise, if a shock hits the agent (which occurs with probability q), then the agent receives

$$0 + \delta V_i(h^t),$$

where, again, $h^t = h^{t-1} \cup \{\mathcal{I}_t, A_t, i, 0\}$.

2.3 A first look at incentive compatibility

We will now focus on incentive-compatible relational contracts. To describe the incentive compatibility condition, fix a relational contract and consider an interaction with a particular principal. The continuation value of a cooperating agent i at the beginning of period t after history h^{t-1} is

$$V_i(h^{t-1}) = \mathbb{E}[Q_{it}(1 - q)p \mid h^{t-1}] + \delta \mathbb{E}[V_i(h^t) \mid h^{t-1}]. \quad (1)$$

Conditional on trading in period t with the principal, the relational contract with that principal induces a probability distribution over whether agent i will be selected to trade in each subsequent period. Denote the probability that a cooperating agent selected to trade in period t will also be selected to trade in period s by $\rho(s \mid t, \Gamma)$. We then have that for an agent selected to trade in period t ,

$$V_i(h^{t-1} \cup \{\mathcal{I}_t, A_t, i\}) = (1 - q)p + \sum_{s=t+1}^{\infty} \delta^{s-t} \rho(s \mid t, \Gamma)(1 - q)p. \quad (2)$$

For purposes of illustration, suppose now that the contract prescribes that the principal fire the agent if the trade fails. Incentive compatibility can then be represented by

$$v \leq p + \sum_{s=t+1}^{\infty} \delta^{s-t} \rho(s \mid t, \Gamma)(1 - q)p, \quad (3)$$

where the right-hand side is the value of present and future trades when cooperating with that principal, while the left-hand side is the value of cheating. Without loss of generality, we can set $p = \lambda v$ for some $\lambda \in [0, 1)$ and we get

$$\frac{(1 - \lambda)}{(1 - q)\lambda} \leq \sum_{s=t+1}^{\infty} \delta^{s-t} \rho(s \mid t, \Gamma). \quad (4)$$

The formulation emphasizes that incentive-compatibility is determined entirely by the frequency with which an agent is selected to trade in the future. Moreover, in order for an agreement to be incentive-compatible for all agents, the principal must commit to trade with all of them sufficiently often in the future. This will place a limit on the number of different agents a principal can credibly trade with.

The principal’s expected payoff from an incentive-compatible relational contract is simply

$$\pi(h^{t-1}) = \sum_{s=t}^{\infty} \delta^{s-t} \sum_i \mathbb{E}[Q_{is}] [(1-q)(v-p) - qc]. \quad (5)$$

We assume that $(1-q)(v-p) > qc$ so that we are in the interesting case in which trade is at least sometimes optimal for the principal. Thus the principal’s expected payoff is increasing in the probability the principal selects an available agent who is incentivized to complete the trade. Maintaining more insiders reduces the probability that there will be no agent available for trade in a given period. On the other hand, we have seen that a principal has to commit to sufficiently frequent future trade with a given agent to incentivize the agent not to defect and this places a limit on how many well-incentivized agents the principal can maintain as insiders. This is the key trade-off we study.

3 Simple Contracts

3.1 Simple contracts: Properties

3.1.1 The need to study a simple class of contracts

The purpose of this discussion is to explain why the optimal contracts within the general class we have defined so far are likely to be complex. Indeed, the arguably unrealistic features of truly optimal contracts motivate us to study the tradeoffs within a simpler and more plausible class of contracts.

First, suppose that the principal punishes failed trade with probability less than 1. If there is no public randomization device that is available, the principal is unlikely to be able to commit to actions that are unobservable and he would like to renege on. In this case the principal would like to always pretend that the realization of the randomization device does not require punishment. If the principal always punishes failed trade with probability 1, then some incentive-compatibility constraints are likely to be slack and the principal will do better by punishing with probability less than 1.

Second, there are likely to be further intricacies regarding which insider is selected for trade and when. For example, as in Andrews and Barron (2016), the principal may want to prioritize trade with insiders who recently traded (in order to strengthen incentives at the time of trade). This can lead to a complex form for γ , the probability distribution with which a principal selects an agent to trade. Indeed, γ can then depend intricately on the history of play, h^t .

These observations are in tension with another feature of implicit contracts. These are often tacit agreements where common understanding is crucial (e.g. Uzzi (1997), Fafchamps and Minten (1999) or Greif (1993)). To avoid misunderstanding (sincere or strategic), intuition—supported by casual empiricism—suggests relational contracts cannot be too complex.

We therefore focus on a considerably simpler, and possibly more realistic, contracting space by requiring that a number of properties constrain the principal.

3.1.2 Properties for simple contracts

First we require that the contract is simple in terms of the mutual understanding of what is required from the agents: They are expected to never defect. Clearly a contract in which agents frequently defect will be suboptimal because of the cost c that defection imposes on the principal⁷. However, as there are no transfers, defection could be occasionally used as a means of transferring rents to the agents. As such defections are occasional they must depend on the history of play in a non-trivial way, creating opportunities for misunderstandings. This motivates the following Property.

Property 1 (No Defection on Path) *A contract never permits defection of an agent.*

Thus in any history reached on the equilibrium path of play, an agent's strategy will prescribe defection with probability 0.

Second, we require that insiders are treated equally. What it means to be an insider depends on the difference in the way the relational contract treats them compared with outsiders. To give insider status a simple and easily understood meaning, we require that all insiders must have the status of a principal's most favored supplier. Technically, this requirement ensures the continuation values of all insiders are the same at the start of a period.

Property 2 (Equal Treatment of Insiders) *After any history, the treatment of all current insiders is the same and cannot depend on the insider's identity. Specifically, (i) under the measure $\gamma(h^{t-1} \cup \{\mathcal{I}_t, A_t\})$, any $i, j \in A_t \cap \mathcal{I}_t$ are chosen with equal probability; (ii) moreover letting i denote the selected agent, then under the measure $\eta(h^t)$ any other insiders $j, k \in \mathcal{I}_t$ have equal probabilities of being in the retained set \mathcal{I}_{t+1} ; and (iii) the probability with which i is in \mathcal{I}_{t+1} depends only on the outcome of the trade $\tilde{v}_t - \tilde{p}_t$.*

Specifically, this will imply that each insider will have to be selected with the same probability for a trade and the treatment of a selected insider after trade, in terms of being granted insider status for the next period, will not depend on the identity of the insider selected.

Next, to give insider status further meaning, we will require that insiders are prioritized over an outsider. A principal can never select an outsider if an insider is available. This requirement considerably simplifies the analysis. For a principal with a very small insider set, absent this constraint, it may be optimal to initially favor outsiders in order to build up a larger insider set more quickly, thereby increasing the chance that he will be able to trade with an available agent. While this would tighten the insiders' incentive-compatibility constraints, it may increase the principal's expected profit. However, eventually incentive compatibility concerns will force the principal to prioritize trade with insiders. In any case, the issue considerably complicates the task of finding an optimal contract. While we expect this requirement to be restrictive in the present model, it can be motivated by robustness considerations: In a model where there is limited monitoring of the principal, he may be tempted to prioritize outsiders more often than he is supposed to. It is easier to verify compliance with simpler contracts.

Property 3 (Insiders' Priority) *An available outsider can only be selected if there are no available insiders. Specifically, under the measure $\gamma(h^{t-1} \cup \{\mathcal{I}_t, A_t\})$, the outsider is chosen with probability 0 whenever $\mathcal{I}_t \cap A_t \neq \emptyset$.*

⁷In practice there are often such costs representing a real loss of social surplus: for example, the failure to generate surplus with one of the principal's customers caused by the principal's lack of a key input.

In this model, a principal can punish defection by making the selected agent an outsider (i.e., by “firing” the selected agent). Given the equal treatment of insiders property, the principal cannot target punishment at a specific agent while retaining that agent as an insider. Making the agent an outsider implies that with probability 1 the agent will never be selected again by this principal, since there is a continuum of agents. Mixing when punishing would allow the agent’s incentive-compatibility constraint to bind and thus may be preferable for the principal since he would in some cases have more insiders to trade with. On the other hand, a mixed strategy is difficult to interpret in reality. Moreover, without a public randomization device, it may be hard to maintain the credibility of this mixed punishment strategy. The principal would like to convince the agents they will be punished with sufficiently high probability that they never defect. But then, upon observing a failed trade, the principal would infer that the shock must have hit and would want to renege and not punish. To better formalize this, we now state the following definition.

Definition 1 (Principal-Credible Contract) *We say that a deviation is non-detectable if it is on the equilibrium path, i.e. if it is consistent with the strategy prescribed by the relational contract. A relational contract is principal credible if all non-detectable deviations are weakly unprofitable.*

To prevent the kind of situations described previously, we require that contracts fulfill the following property:

Property 4 (Credibility) *A contract must be principal credible.*

Thus, whenever a relational contract gives the principal latitude in the actions that he can take, then the action prescribed by this contract must be weakly preferred to all others. For example, if the relational contract requires the principal to mix between actions, then the principal must be indifferent between those actions.

The next remark follows from Property 4.

Remark 1 (No Mixing when Punishing) *To maintain credibility, the principal cannot mix when punishing an unsuccessful trade (i.e. when setting the status of the agent in question to “outsider” for the next period). That is, he punishes either with probability 0 or 1.*

Property 4 limits the ability of the principal to randomize in different ways that he would be tempted to game. While the principal generally likes to have more insiders, as having more of them increases the probability of there being an available insider to trade with each period, insiders will not typically like to share future trades with more other insiders. Thus, in order to maximize the expected number of insiders the principal has, while respecting existing insiders’ incentive-compatibility constraints, the principal might want to mix over whether to take on an additional insider or not. However, if this mixing probability cannot be observed, the principal would always like to cheat and take on the extra insider for sure, while the current insiders believe the decision is random.

Moreover, to avoid intricate conditioning on the history of play that could otherwise proxy for mixing and would make the relational contract complex, we have the following property.

Property 5 (Simple Conditioning on History when Hiring a New Agent) *When $A_t \cap \mathcal{I}_t = \emptyset$ and $O \in A_t$, then the selection strategy γ can be conditioned only on \mathcal{I}_t .*

An implication of both Properties 4 and 5 is summarized in the following remark.

Remark 2 (No Mixing when Hiring a New Insider) *If after a history h^t there are no available insiders and the outsider is available, the principal must initiate trade with the outsider with probability 1 or 0 and can condition that decision on the history h^t only through the current number of insiders. Further, if trade is successful, the principal must grant insider status to the outsider with probability 1 or 0.*

3.2 Simple contracts: Characterization

Before stating our main characterization of optimal contracts within the restricted space we have laid out, we define a k -insider contract.

Definition 2 (k -Insider Contract) *A contract with the principal's strategy given by $\Gamma = (\gamma, \eta)$ is a k -insider contract if it has the following structure:*

- (i) *If $\mathcal{I}_t \cap A_t \neq \emptyset$, then each available insider is selected for trade with probability $1/|\mathcal{I}_t \cap A_t|$.*
- (ii) *If $\mathcal{I}_t \cap A_t = \emptyset$ and the outsider is available, the outsider is selected (with probability 1) if and only if $|\mathcal{I}_t| < k$.*
- (iii) *Successful trade is rewarded by insider status, while unsuccessful trade is punished with outsider status.*

We say a relational contract is *constrained-efficient* when it is incentive-compatible and maximizes the expected payoff of a principal among the incentive-compatible contracts that satisfy Properties 1, 2, 3, 4 and 5 (and thus Remarks 1 and 2).

Proposition 1 *A relational contract is constrained-efficient if and only if it is a k^* -insider contract with k^* given by*

$$k^* = \max\{k : V_i(k) \geq (v - p)/\delta\}$$

where $V_i(k)$ is an agent's expected payoff when there are k insiders and the principal commits to a k -insider contract.

The simple contract properties do a lot of the work of pinning down the decisions of the principal. Proposition 1 is an exact characterization of what the simplicity conditions imply about principal behavior when combined with optimality within this class of simple contracts.

Even though the contracting space is much simplified, this characterization requires several lemmas to be proved, which is done in the following section.

3.3 Simple contracts: Proof of characterization

We prove Proposition 1 through a series of lemmas, the proofs of which are relegated to the appendix.

First we show that an agent i is removed from the insider set only when trade is initiated with that agent and fails.

Lemma 1 *In all constrained-efficient relational contracts, if $i \in \mathcal{I}_t$ and trade with i was either not initiated in period t , or else trade with i was initiated and succeeded, then $i \in \mathcal{I}_{t+1}$.*

We next show that, given the insider priority property (Property 3), the principal will initiate trade with an insider whenever one is available.

Lemma 2 *In all constrained-efficient relational contracts, the principal initiates trade with an insider whenever there is an available insider.*

We now show that there is an upper bound on the number of insiders the principal can maintain.

Lemma 3 *There exists a \bar{k} such that the principal cannot maintain more than \bar{k} insiders while satisfying incentive-compatibility.*

Lemma 3 shows that if the principal always takes on available outsiders, eventually incentive-compatibility will be violated. We now use this lemma to show that the principal takes on an available outsider whenever there is no available insider and trade with the outsider is incentive-compatible. It will turn out to be incentive-compatible precisely when there are fewer than k^* insiders.

Lemma 4 *In all constrained-efficient relational contracts, there exists an integer k such that if there are no available insiders the principal trades with an available outsider if and only if there are fewer than k insiders and rewards successful trade with insider status.*

Together, Lemmas 1 and 4, along with Property 4 (Remark 1), establish that part (iii) of Definition 2 must hold in all constrained-efficient relational contracts: Failed trade is always punished with outsider status (else, there would be no way to incentivize agents not to cheat) while successful trade is always rewarded with insider status.

We now turn to part (i) of Definition 2. Lemma 2, Properties 2, 3 and 4 establish that all constrained-efficient relational contracts satisfy part (i) of Definition 2. When at least one insider is available, each available insider is chosen with equal probability for trade. Note that mixing here is credible since all available insiders are homogenous and thus the principal has no incentive to favor one in particular. Property 4 is thus not violated.

Part (ii) of Definition 2 holds in all constrained-efficient relational contracts by Lemma 4. This proves that Definition 2 establishes necessary characteristics of a constrained-efficient relational contract. Moreover, as under Property 1 there is no defection on path, Eq. (5) characterizes the principal's expected profits in a constrained-efficient relational contract. This shows that the principal's expected profit is increasing in the frequency with which the principal trades and so is increasing in k , for a k -insider contract. Letting $V_i(k)$ be an agent's expected payoff when there are k insiders and the principal commits to a k -insider contract, the largest k than can be sustained without violating incentive compatibility is

$$k^* = \max\{k : V_i(k) \geq \frac{v - p}{\delta}\}.$$

Thus all constrained-efficient relational contracts are k^* -insider contracts. Moreover, the set of k^* -insider contracts is nonempty and, as this pins down all the principal's actions after every possible history, Definition 2 with $k = k^*$ also provides a sufficient condition for a constrained-efficient relational contract. ■

4 Dynamics of Optimal Insider Contracts

Now that we have established the optimality of insider contracts within the space of simple contracts, we can calculate a principal's expected profit and an agent's expected utility under a k -insider contract and study in detail the conditions for incentive-compatibility. This will then allow us to find the equilibrium behavior of principals and analyze the dynamics of relations being created and destroyed.

4.1 Principal's expected profit

Even within the simple class of contracts we have described, the principal's expected discounted profits in any state involve an expectation over the stochastic process governing the size of the insider set. Fix a principal j and a k -insider contract, i.e. an insider contract in which j holds a maximum of k insiders, in which agents cooperate whenever they are available (which is justified by the incentive-compatibility of such contracts when k is low enough). Then, recalling Eq. (5), the principal's expected profit vector can be written as

$$\mathbf{\Pi}_j(k) = \mathbb{E}[\mathbf{Q}^j(k)]((1-q)(v-p) - qc) + \delta \mathbf{M}_P(k) \mathbf{\Pi}_j(k) \quad (6)$$

where $\mathbf{\Pi}_j(k) = [\Pi_j(0, k), \Pi_j(1, k), \dots, \Pi_j(k, k)]^\top$ is the vector of expected profits when the number of insiders is $0, 1, \dots, k$. The matrix $\mathbf{M}_P(k)$ describes transitions among states and $\mathbf{Q}^j(k)$ is a vector we are about to define.

We let $Q^j(d, k) = \sum_{i \in \mathcal{I} \cup \{O\}} Q_i^j(d, k)$ denote the indicator variable the principal chooses *some* agent to trade with at a given time, given that he currently has d agents in his insider set \mathcal{I} and aims to keep a maximum of k such insiders. Given the form of the contract, $\mathbb{E}[Q^j(d, k)]$ is the probability that an agent is available in state (d, k) , and so

$$\mathbb{E}[\mathbf{Q}^j(k)] = \begin{bmatrix} \mathbb{E}[Q^j(0, k)] \\ \mathbb{E}[Q^j(1, k)] \\ \vdots \\ \mathbb{E}[Q^j(k-2, k)] \\ \mathbb{E}[Q^j(k-1, k)] \\ \mathbb{E}[Q^j(k, k)] \end{bmatrix} = \begin{bmatrix} 1 - (1 - \alpha)^1 \\ 1 - (1 - \alpha)^2 \\ \vdots \\ 1 - (1 - \alpha)^{k-1} \\ 1 - (1 - \alpha)^k \\ 1 - (1 - \alpha)^k \end{bmatrix}. \quad (7)$$

$\mathbf{M}_P(k)$ is the Markov transition matrix among those different states (d, k) :

$$\mathbf{M}_P(k) = \begin{bmatrix} a_{0,0} & a_{0,1} & \cdots & a_{0,k} \\ a_{1,0} & a_{1,1} & \cdots & a_{1,k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k,0} & a_{k,1} & \cdots & a_{k,k} \end{bmatrix}, \quad (8)$$

where, for $0 \leq d \leq k$:

$$\begin{aligned} a_{d,d-1} &= (1 - (1 - \alpha)^d)q \\ a_{d,d} &= (1 - (1 - \alpha)^d)(1 - q) + (1 - \alpha)^d(1 - \alpha) + (1 - \alpha)^d q \alpha \\ a_{d,d+1} &= (1 - \alpha)^d \alpha (1 - q). \end{aligned}$$

with the exception that $a_{k,k} = (1 - (1 - \alpha)^k)(1 - q) + (1 - \alpha)^k$. All other entries are 0.

Thus we have the following characterization:

Lemma 5 *In any k -insider contract where agents cooperate when they are available,*

$$\Pi_j(k) = (\mathbf{I} - \delta \mathbf{M}_P(k))^{-1} \mathbb{E}[\mathbf{Q}^j(k)] ((1 - q)(v - p) - qc). \quad (9)$$

The matrix $(\mathbf{I} - \delta \mathbf{M}_P(k))$ is a non-singular M -matrix, so its inverse is nonnegative. Moreover, for any k , it holds that $\Pi_j(0, k) < \Pi_j(1, k) < \dots < \Pi_j(k, k)$.

The principal is thus better off if he currently has more insiders. Proofs not in the main text are in the appendix.

4.2 Agent's expected utility

An agent i 's expected utility vector from a relationship with a principal j , when this principal commits to a k -insider contract and the agent cooperates whenever able, is

$$\mathbf{U}_i^j(k) = (1 - q)p\mathbb{E}[\mathbf{Q}_i^j(k)] + \delta \mathbf{M}_A(k) \mathbf{U}_i^j(k) \quad (10)$$

where $\mathbf{U}_i^j(k) = [U_i^j(O, k), U_i^j(I, 1, k), \dots, U_i^j(I, k, k)]^\top$ is the vector of utilities in all possible states the agent can be in and

$$\mathbb{E}[\mathbf{Q}_i^j(k)] = \begin{bmatrix} \mathbb{E}[Q_i^j(O, k)] \\ \mathbb{E}[Q_i^j(I, 1, k)] \\ \vdots \\ \mathbb{E}[Q_i^j(I, k, k)] \end{bmatrix} \quad (11)$$

is the probability that agent i will be chosen to trade with principal j in each possible state. The state (I, d, k) means agent i is an insider to principal j , the latter has currently d insiders and follows a k -insider contract. The state (O, k) means agent i has been fired (made an outsider) by principal j , which follows a k -insider contract, and thus will never be chosen again by principal j as there is a continuum of principals. Note that $\mathbb{E}[Q_i^j(I, d, k)] = \alpha \mathbb{E}[\frac{1}{1+n_a^d}]$ where $n_a^d \sim \text{bin}(\alpha, d - 1)$. On the other hand, $\mathbb{E}[Q_i^j(O, k)] = 0$ is the probability the agent will be chosen again by principal j after being fired.

The Markov transition matrix $\mathbf{M}_A(k)$ is:

$$\mathbf{M}_A(k) = \begin{bmatrix} a_{O,O} & a_{O,1} & \cdots & a_{O,k} \\ a_{1,O} & a_{1,1} & \cdots & a_{1,k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k,O} & a_{k,1} & \cdots & a_{k,k} \end{bmatrix} \quad (12)$$

where the entries are, for $1 \leq d \leq k$:

$$\begin{aligned} a_{d,d-1} &= \frac{(d-1)(1-(1-\alpha)^d)}{d}q \\ a_{d,d} &= (1-(1-\alpha)^d)(1-q) + (1-\alpha)^d q \alpha + (1-\alpha)^d (1-\alpha) \\ a_{d,d+1} &= (1-\alpha)^d \alpha (1-q), \end{aligned}$$

with the exception that $a_{k,k} = (1-(1-\alpha)^k)(1-q) + (1-\alpha)^k$. Also,

$$\begin{aligned} a_{d,O} &= \frac{(1-(1-\alpha)^d)}{d}q \\ a_{O,d} &= 0 \\ a_{O,O} &= 1. \end{aligned}$$

All other entries are 0.

Thus we have

Lemma 6 *In any k -insider contract in which agent i cooperates when available, his expected utility is*

$$U_i^j(k) = (\mathbf{I} - \delta \mathbf{M}_A(k))^{-1} \mathbb{E}[\mathbf{Q}_i^j(k)](1-q)p, \quad (13)$$

where $(\mathbf{I} - \delta \mathbf{M}_A(k))$ is a non-singular M -matrix. Moreover, for any k , it holds that $U_i^j(I, 1, k) > U_i^j(I, 2, k) > \cdots > U_i^j(I, k, k)$.

4.3 Incentive-compatibility

Lemma 7 *Given a maximum number of insiders k , a contract is incentive-compatible if, for any $d \leq k$,*

$$p + \delta U_i^j(I, d, k) \geq v + \delta U_i^j(O, k). \quad (14)$$

It follows from Lemma 6 that if a contract is incentive-compatible with k insiders, it is also incentive-compatible with k' insiders, for any $k' < k$. By Lemma 5, a principal's expected profit is increasing in the number of insiders he keeps. In a principal-optimal contract, the latter will choose the maximum number of insiders k such that the incentive-compatibility condition holds, or equivalently

$$U_i^j(I, k, k) - U_i^j(O, k) \geq \frac{v-p}{\delta}. \quad (15)$$




Figure 1: A realization of the equilibrium random bipartite graph linking the set of principals \mathcal{P} to the set of agents \mathcal{A} in steady state. The number of edges d incident on a principal follows distribution $f_{k^*}(d)$. In this example, $k^* = 3$ and thus a principal cannot have degree higher than 3. An agent, on the other hand, can have any degree. The first-best configuration—when a principal does not have to maintain incentive-compatibility—would be the complete bipartite graph.

Note also that from Eq. (10), it follows that $U_i^j(O, k) = \delta U_i^j(O, k)$ and thus that⁸ $U_i^j(O, k) = 0$. Thus this condition reduces nicely to $U_i^j(I, k, k) \geq \frac{v-p}{\delta}$.

4.4 Equilibrium

To analyze the behavior of the entire market, we define the insider contract game by allowing all principals to simultaneously choose, and commit to, any insider relational contracts. We look for a Nash equilibrium of this game.

Proposition 2 (Existence and Uniqueness) *There exists a unique equilibrium of the insider contract game in which all principals choose an insider contract with k^* insiders. To this k^* , there corresponds a random bipartite graph with stationary principal degree distribution $f_{k^*}(d)$, which is the stationary distribution of the Markov transition matrix $\mathbf{M}_P(k^*)$.*

The steady-state distribution of the number of agents in a principal's insider set is thus given by $f_{k^*}(d)$. Therefore, in steady state, the equilibrium strategy k^* yields a random bipartite graph linking the set of principals to the set of agents and $f_{k^*}(d)$ is the distribution of the number of edges incident on a principal. Fig. 1 illustrates an equilibrium network in steady state.

4.5 Comparative statics and welfare

4.5.1 Welfare

Given a symmetric strategy k followed by principals and agents always cooperating, the welfare (social surplus) per period is:

$$W(k) = ((1-q)v - qc) \left(\sum_{d=0}^{k-1} f_k(d)(1 - (1-\alpha)^{d+1}) + f_k(k)(1 - (1-\alpha)^k) \right). \quad (16)$$

The upper bound on welfare (first-best) is $\bar{W} = \lim_{k \rightarrow \infty} W(k) = (1-q)v - qc$, which is also reached when $\alpha = 1$. This corresponds to the efficiency benchmark.

⁸Indeed, after being made an outsider (i.e. being fired) by principal j , agent i will never be chosen again by principal j and thus can obtain no further expected utility from principal j .

It can be shown that social surplus is greater when principals maintain a larger number of insiders who cooperate. This follows from the intuitive fact that when each principal seeks a larger number of insiders, in equilibrium each principal tends to have a higher degree. This is summarized in the next proposition.

Proposition 3 *Let k and k' be two symmetric strategies played by all principals and suppose all agents cooperate when available. If $k' > k$, then $f_{k'}(d) \succeq f_k(d)$, where \succeq indicates first-order stochastic dominance. Moreover, $W(k)$ is weakly increasing in k .*

4.5.2 Comparative statics in q

Suppose that market conditions deteriorate, making it more difficult for suppliers to live up to their commitments. As a first pass, we can model this as a reduction in q . How does k^* , the principal-optimal and efficient number of insiders, depend on q ?

Proposition 4 *Let $k^*(q)$ be the principal-optimal number of insiders under q . For all \hat{q} such that $k^*(\hat{q}) > 0$, there exists $\underline{q} > \hat{q}$ and $\epsilon > 0$ such that $k^*(q) \leq k^*(\hat{q}) - 1$ for all $q > \underline{q}$ and $k^*(q) = k^*(\hat{q}) - 1$ for all $q \in (\underline{q}, \underline{q} + \epsilon)$.*

Proposition 4 states that when q increases sufficiently, the corresponding k^* is bounded by a strictly decreasing upper bound. Moreover, this upper bound decreases by steps of size one. This decrease in the number of relationships across the market is associated with a corresponding decrease in welfare, due to fewer trade opportunities.

Corollary 1 *Let $k'^* < k^*$ be the principal-optimal and efficient number of insiders corresponding to $q' > \underline{q}$ in Proposition 4. Then $W(k'^*) \leq W(k^*)$.*

Fig. 2 illustrates Proposition 4 and Corollary 1. We see that an increase in q leads to a rapid decrease in the number of insiders k^* a principal can keep while maintaining incentive-compatibility. These lost relationships can be viewed as a destruction of social capital and are associated with large welfare losses. Welfare is impacted by two effects: (i) the direct and continuous effect of an increase in q on $(1 - q)v - qc$, seen in the slope of the $W(k^*)$ curve; (ii) the indirect effect of a decreasing $k^*(q)$ as q increases, which is seen in the downward jumps of $W(k^*)$. As a comparison, the first best welfare—which assumes a complete bipartite graph and also that the principal does not have to maintain incentive-compatibility—only suffers a continuous decrease due to an increasing q .

Because of search frictions, the model involves an asymmetry in how it responds to unexpected positive and negative changes in q . For example, for the parameters shown in Fig. 2, if the economy enters a crisis and q suddenly rises from 0.09 to 0.12, then a principal would have to immediately adjust the relational contracts⁹ in order to maintain incentive-compatibility. This would imply k^* suddenly dropping from 9 to 2 (as per Fig. 2). This implies an immediate destruction of social capital, quantified here by a sharp drop in welfare $W(k^*)$ of more than 50%. However, were the economy to unexpectedly enter a post-crisis state in which q returns to its initial value of 0.09, rebuilding the lost social capital takes a considerable amount of time. Indeed, the principal would

⁹We assume that the principal would be allowed to renegotiate contracts following the transition to an unanticipated crisis and the associated change in q .




Figure 2: Principal-optimal and efficient number of insiders k^* in equilibrium, corresponding welfare $W(k^*)$ and first best welfare (i.e. $(1 - q)v - qc$) versus the shock probability q . Other model parameters are $\delta = 0.99$, $c = 0.2$, $v = 3$, $p = 0.5$, $x = 1$ and $\alpha = 0.2$.

$ \mathcal{I}_t $										
0	1	2	3	4	5	6	7	8	9	
0	5.49	13.04	23.73	39.49	64.15	105.93	185.04	358.66	817.15	

Table 1: Expected first-passage times (in periods) from a starting insider set size $|\mathcal{I}_t|$ of 0 to various other values of $|\mathcal{I}_t|$ (columns).

now rewrite optimal contracts in which he aims to return to a maximum of $k^* = 9$ insiders, but search frictions would delay the rebuilding of his set of insiders. In Table 1, we show the expected first-passage times (in periods) from an insider set size of 0 to a higher insider set size (columns). We see that building up the insider set from 2 insiders (the median number of insiders in the steady state when $k^* = 2$) to 5 insiders (the median number of insiders in the steady state when $k^* = 9$) takes more 50 periods¹⁰. Search frictions limit the speed of the transition and are exacerbated by the relational contracts, which prioritize insiders thereby relaxing their incentive compatibility constraints.

5 Regime Switching

5.1 Generalizing equilibrium results

The previous analysis essentially assumes that changes in q are observable but unanticipated. We can now naturally extend the model to permit these changes to be anticipated. Suppose that at any time t , the economy can be either in a ‘crisis’ regime or in a ‘non-crisis’ regime. We let the variable $\theta \in \{0, 1\}$ represent that regime of the economy (i.e. $\theta = 0$ means ‘non-crisis’ and $\theta = 1$ means

¹⁰As the process is a Markov chain, to get from 0 insiders to 5 insiders, we must first pass by 2 insiders and thus from Table 1, the expected first-passage time from 2 insiders to 5 insiders is $64.15 - 13.04 \approx 51$.

‘crisis’) and assume this is commonly observable. The probability of an agent being hit by a shock should be greater in a time of crisis than in normal times. Thus, we can let $q = q_\theta$ depend on θ , such that $q_1 > q_0$.

The transitions between the two different regimes of the economy can be represented by the following Markov transition table:

θ	0	1
0	s_{00}	s_{01}
1	s_{10}	s_{11}

It would be natural to assume that $s_{00} > s_{01}$ and $s_{11} > s_{10}$ so that the regimes are persistent. Naturally, $s_{01} = 1 - s_{00}$ and $s_{10} = 1 - s_{11}$. The transition in θ occurs at the very beginning of a period (just before step (i) in Section 2.1.1) and is publicly observable.

For trading in each regime to be incentive compatible, a different incentive compatibility constraint will have to be satisfied for each θ . Let $U_i^j(I, d, k_0, k_1, \theta)$ be an agent i ’s utility from interacting with principal j , in regime θ when the principal-optimal number of insiders kept by principals are k_0 and k_1 in each regime and when principal j currently has d insiders. There are now two incentive compatibility constraints that must be satisfied (one in each regime, $\theta = 0$ and $\theta = 1$):

$$U_i^j(I, k_0, k_0, k_1, 0) \geq \frac{v - p}{\delta} \quad (17)$$

$$U_i^j(I, k_1, k_0, k_1, 1) \geq \frac{v - p}{\delta}. \quad (18)$$

From Proposition 4, we can foresee that k_1^* may be smaller than k_0^* . Indeed, the only way to satisfy these insiders’ incentive compatibility constraint is to divide the smaller value of future trades among fewer insiders. This will be formalized in Proposition 5. This could also imply that, following a transition to the crisis regime, the principal may need to ‘fire’ as many insiders as necessary to satisfy the incentive compatibility constraints of the remaining insiders (i.e. Eq. (18)). As mentioned earlier, the transition in θ occurs at the very beginning of a period (just before step (i) in Section 2.1.1) and is publicly observable. We may therefore add an additional step (0) in the timing of the model, in which θ is observed and in which the principal can select \mathcal{I}_t^j to reflect that new information. To guarantee equal treatment of insiders, we thus add an additional subpart to Property 2:

Property 2(iv) (Equal Treatment of Insiders Following a Transition to a Crisis) *At time t , upon observing a transition to the crisis regime ($\theta = 1$), any subset of insiders of size k_1 currently in \mathcal{I}_t is chosen to remain in \mathcal{I}_t with the same probability.*

Following a return to the non-crisis regime, it may take time to rebuild the insider set. This can limit the ability of the principal to find an available insider and thus decrease his expected profit. The principal may thus find it optimal to halt trading in the crisis regime and then resume trading with a larger insider set when the crisis ends. The following proposition gives conditions under which the principal finds it optimal to pursue trading with a reduced insider set during a crisis. The crisis regime essentially needs to be persistent enough and the probability of shocks in the crisis regime, q_1 , must not be so high that the incentive compatibility condition would always be violated.

Proposition 5 (Trading in crisis and non-crisis regimes) *Let $1 < k^*(q_1) < k^*(q_0)$ in the single regime model. Then, in the two-regime model:*

- (i) *There exists $\underline{s}_{11} < 1$ such that the principal always chooses to continue trading in the crisis regime if $s_{11} > \underline{s}_{11}$.*
- (ii) *Moreover, there also exists $\underline{s}_{00} < 1$ such that when $s_{00} > \underline{s}_{00}$, $k_1^* < k_0^*$. A principal thus retains a maximum of k_1^* randomly-chosen insiders following the transition to the crisis regime.*

This proposition shows that the coming of a (sufficiently persistent) crisis implies a rapid destruction of social capital. Indeed, the expected value of future trade being lower, the principal cannot sustain as many incentive-compatible relationships when the crisis regime hits. He must thus fire insiders to maintain incentive compatibility, i.e. $k_1^* < k_0^*$ when q_1 is high enough. This is true even before an actual production shock hits any agent and therefore exacerbates the destruction of social capital already caused by the higher shock probability $q_1 > q_0$. Upon returning to the normal regime, this insider set will then take some time to rebuild, since an insider is prioritized over an outsider. Thus $f_{k_0^*,t}(d) \preceq f_{k_0^*}(d)$, where $f_{k_0^*,t}(d)$ is the transient principal degree distribution when exiting the crisis regime, while $f_{k_0^*}(d)$ is the steady state principal degree distribution in the non-crisis regime. This thus decreases the probability of a trade compared to steady state and therefore decreases welfare, i.e. $W_t(k_0^*) \leq W(k_0^*)$, from Proposition 3. It is to be noted that the mere possibility of a transition to a crisis regime may imply a lower number of insiders in the non-crisis regime, because of the lower continuation value faced by any agent. That is, $k^*(q_0) \geq k_0^*$, where $k^*(q_0)$ is the optimal number of insiders in the single-regime model presented earlier in the paper, while k_0^* is the optimal number of insiders in the non-crisis regime of the two-regime model.

Fig. 3 illustrates this proposition in a setting where a transition into the crisis regime occurs at time $t = 40$ and a transition back into the non-crisis regime occurs at $t = 100$. Here, the optimal maximum number of insiders a principal can keep in each regime is $(k_0^*, k_1^*) = (9, 5)$. The upper panel describes the dynamics of the average number of insiders principals keep at any time. We see that the coming of the crisis regime implies an immediate firing of insiders in order to preserve incentive compatibility (since $k_1^* < k_0^*$). The occurrence of shocks then further decreases the average number of insiders principals have as time passes and this average tends to a new steady state. When the crisis ends, principals can afford to hire up to $k_0^* > k_1^*$ insiders, but it takes them time to rebuild their insider set, since current insiders must be prioritized *and* search frictions also limit the hiring of new insiders.

The lower panel describes the dynamics of welfare. The dotted line represents the first-best welfare (if a principal could keep an arbitrarily high number of insiders). In this case, the drop in welfare in the crisis regime is simply due to the higher probability that an agent be hit with a shock ($q_1 > q_0$). The full red line represents the actual welfare $W_t(k_{\theta_t}^*)$. Here, incentive compatibility concerns further decrease welfare since a principal can only keep a maximum of k_0^* agents in the non-crisis regime and k_1^* in the crisis regime. Insider prioritization and search frictions also slow down the rebuilding of the welfare after the return to the non-crisis regime.

5.2 Computing the equilibrium

Here we explain in greater details how to find an equilibrium as in Fig. 3. We must find a pair (k_0^*, k_1^*) such that an agent's incentive compatibility constraints in Eqs. (17)-(18) are satisfied and at least one of them would be violated by increasing k_0^* or k_1^* .




Figure 3: Transition into and out of a crisis: (i) upper panel describes the dynamics of $\mathbb{E}_t[d]$, the market-wise average number of insiders a principal has; (ii) lower panel describes the dynamics of the welfare $W_t(k_{\theta_t}^*)$ (full line) and the first-best welfare (dotted line). Insiders must be ‘fired’ to maintain incentive compatibility when the crisis hits (at $t = 40$). Upon returning to the non-crisis regime (at $t = 100$), insider prioritization and search frictions limit the speed at which relationships can be rebuilt. Parameter values are $v = 5$, $p = 2$, $c = 0.2$, $\delta = 0.95$, $\alpha = 0.2$, $q_0 = 0.01$, $q_1 = 0.2$, $s_{00} = s_{11} = 0.98$, $s_{01} = s_{10} = 0.02$. The equilibrium maximum numbers of insiders a principal can keep in each regime are respectively $k_0^* = 9$ and $k_1^* = 5$.

5.2.1 Principal's problem

To give a formal expression for a principal's profit, we first need to describe the Markov chain governing the states in which a principal j can be. This Markov chain is now defined by a transition matrix $\mathbf{M}_P(k_0, k_1)$. It describes not only the variation of a principal j 's insider set due to shocks, but also due to regime switches:

$$\mathbf{M}_P(k_0, k_1) = \begin{bmatrix} s_{00}\mathbf{M}_P^{q_0}(k_0) & s_{01}\bar{\mathbf{M}}_P(k_0, k_1) \\ s_{10}\hat{\mathbf{M}}_P(k_0, k_1) & s_{11}\mathbf{M}_P^{q_1}(k_1) \end{bmatrix} \quad (19)$$

where $\mathbf{M}_P^{q_0}(k_0)$ (resp. $\mathbf{M}_P^{q_1}(k_1)$) is the one-period state transition probability matrix after it has been observed that we remain in the non-crisis regime $\theta = 0$ (resp. in the crisis regime $\theta = 1$). $\mathbf{M}_P^{q_0}(k_0)$ (resp. $\mathbf{M}_P^{q_1}(k_1)$) is defined as in Eq. (8), with $k = k_0$ and $q = q_0$ (resp. with $k = k_1$ and $q = q_1$).

On the other hand, $\bar{\mathbf{M}}_P(k_0, k_1)$ (resp. $\hat{\mathbf{M}}_P(k_0, k_1)$) is the one-period state transition probability matrix after it has been observed that we just entered the crisis regime (resp. the non-crisis regime), i.e. that θ just changed from 0 to 1 (resp. from 1 to 0). Such a matrix again captures the immediate firing of agents upon entering a crisis (resp. the search frictions causing the slow rebuilding of relationships upon exiting the crisis). The expressions, although relatively simple, are given in the Appendix (Section 7.1) in order to save space.

The vector of the profit $\Pi_j(d, k_0, k_1, \theta)$ in each possible state a principal j can be in is then written as

$$\boldsymbol{\Pi}_j(k_0, k_1) = \begin{bmatrix} \Pi_j(0, k_0, k_1, 0) \\ \Pi_j(1, k_0, k_1, 0) \\ \Pi_j(2, k_0, k_1, 0) \\ \vdots \\ \Pi_j(k_0, k_0, k_1, 0) \\ \Pi_j(0, k_0, k_1, 1) \\ \Pi_j(1, k_0, k_1, 1) \\ \Pi_j(2, k_0, k_1, 1) \\ \vdots \\ \Pi_j(k_1, k_0, k_1, 1) \end{bmatrix}. \quad (20)$$

We can express $\boldsymbol{\Pi}_j(k_0, k_1)$ as

$$\boldsymbol{\Pi}_j(k_0, k_1) = \mathbf{R}_P + \delta \mathbf{M}_P(k_0, k_1) \cdot \boldsymbol{\Pi}_j(k_0, k_1) \quad (21)$$

where \mathbf{R}_P is, for each state a principal can be in, the expected profit in the current period. Its expression is given in the Appendix (Section 7.1).

As before, we can then express $\boldsymbol{\Pi}_j(k_0, k_1)$ explicitly:

$$\boldsymbol{\Pi}_j(k_0, k_1) = (\mathbf{I} - \delta \mathbf{M}_P(k_0, k_1))^{-1} \mathbf{R}_P. \quad (22)$$

5.2.2 Agent's problem

To give a formal expression for an agent's utility, we first need to describe the Markov chain governing the states in which an agent i interacting with a principal j can be. This Markov chain is now defined by a transition matrix $\mathbf{M}_A(k_0, k_1)$. It describes not only the variation of a principal j 's insider set due to shocks, but also due to regime switches:

$$\mathbf{M}_A(k_0, k_1) = \begin{bmatrix} s_{00}\mathbf{M}_A^{q_0}(k_0) & s_{01}\bar{\mathbf{M}}_A(k_0, k_1) \\ s_{10}\hat{\mathbf{M}}_A(k_0, k_1) & s_{11}\mathbf{M}_A^{q_1}(k_1) \end{bmatrix} \quad (23)$$

where $\mathbf{M}_A^{q_0}(k_0)$ (resp. $\mathbf{M}_A^{q_1}(k_1)$) is the one-period state transition probability matrix after it has been observed that we remain in the non-crisis regime $\theta = 0$ (resp. in the crisis regime $\theta = 1$). $\mathbf{M}_A^{q_0}(k_0)$ (resp. $\mathbf{M}_A^{q_1}(k_1)$) is defined as in Eq. (12), with $k = k_0$ and $q = q_0$ (resp. with $k = k_1$ and $q = q_1$).

On the other hand, $\bar{\mathbf{M}}_A(k_0, k_1)$ (resp. $\hat{\mathbf{M}}_A(k_0, k_1)$) is the one-period state transition probability matrix after it has been observed that we just entered the crisis regime (resp. the non-crisis regime), i.e. that θ just changed from 0 to 1 (resp. from 1 to 0). Such a matrix again captures the immediate firing of agents upon entering a crisis (resp. the search frictions causing the slow rebuilding of relationships upon exiting the crisis). The expressions, although relatively simple, are also given in the Appendix (Section 7.1) in order to save space.

The vector of expected utilities in every possible state an agent i can be in when dealing with principal j can be written as

$$\mathbf{U}_i^j(k_0, k_1) = \begin{bmatrix} U_i^j(O, k_0, k_1, 0) \\ U_i^j(I, 1, k_0, k_1, 0) \\ U_i^j(I, 2, k_0, k_1, 0) \\ \vdots \\ U_i^j(I, k_0, k_0, k_1, 0) \\ U_i^j(O, k_0, k_1, 1) \\ U_i^j(I, 1, k_0, k_1, 1) \\ U_i^j(I, 2, k_0, k_1, 1) \\ \vdots \\ U_i^j(I, k_1, k_0, k_1, 1) \end{bmatrix}. \quad (24)$$

We can express $\mathbf{U}_i^j(k_0, k_1)$ as

$$\mathbf{U}_i^j(k_0, k_1) = \mathbf{R}_A \cdot p + \delta \mathbf{M}_A(k_0, k_1) \cdot \mathbf{U}_i^j(k_0, k_1) \quad (25)$$

where \mathbf{R}_A is, for each state an agent can be in, the probability of being chosen and not failing in the current period. Its expression is given in the Appendix (Section 7.1).

As before, we can then express $\mathbf{U}_i^j(k_0, k_1)$ explicitly:

$$\mathbf{U}_i^j(k_0, k_1) = (\mathbf{I} - \delta \mathbf{M}_A(k_0, k_1))^{-1} \mathbf{R}_A \cdot p. \quad (26)$$

5.2.3 Equilibrium dynamics

We now study the equilibrium dynamics when the parameters are $v = 5$, $p = 2$, $c = 0.2$, $\delta = 0.95$ and $\alpha = 0.2$. We let $q_0 = 0.01$ in the non-crisis regime while $q_1 = 0.2$ in the crisis regime and we let the transition probabilities between the two regimes be as follows:

θ	0	1
0	0.98	0.02
1	0.02	0.98

In this case, there is a unique pair $(k_0^*, k_1^*) = (9, 5)$ satisfying the equilibrium conditions.

As we saw earlier, Fig. 3 illustrates this setting when a transition into the crisis regime occurs at time $t = 40$ and a transition back into the non-crisis regime occurs at $t = 100$.

It is easy to track the dynamics of the degree distribution of a principal, $\mathbf{f}_{k_{\theta_t}^*, t}$ using the Markov transition matrices. For $t < 40$, we assume that we are at the steady-state of the non-crisis regime, i.e. $\mathbf{f}_{k_{\theta_t}^*, t} = \mathbf{f}_{k_0^*}$, where

$$\mathbf{f}_{k_0^*} = \mathbf{f}_{k_0^*} \mathbf{M}_P^{q_0}(k_0^*).$$

For $t = 40$, we transit from the non-crisis regime into the crisis regime and thus

$$\mathbf{f}_{k_1^*, 40} = \mathbf{f}_{k_0^*, 39} \bar{\mathbf{M}}_P(k_0^*, k_1^*)$$

For $t \in [41, 99]$, we are in the crisis regime and thus

$$\mathbf{f}_{k_1^*, t+1} = \mathbf{f}_{k_1^*, t} \mathbf{M}_P^{q_1}(k_1^*)$$

For $t = 100$, we transit from the crisis regime back into the non-crisis regime and thus

$$\mathbf{f}_{k_0^*, 100} = \mathbf{f}_{k_1^*, 99} \hat{\mathbf{M}}_P(k_0^*, k_1^*)$$

Finally, for $t \geq 100$, we are in the non-crisis regime and thus

$$\mathbf{f}_{k_0^*, t+1} = \mathbf{f}_{k_0^*, t} \mathbf{M}_P^{q_0}(k_0^*)$$

Moreover, since we assume a unit mass of principals, $\mathbf{f}_{k_{\theta_t}^*, t}$ captures the actual distribution of the number of insiders that principals have at any given time t . We can then easily compute the welfare $W_t(k_{\theta_t})$ across the market using Eq. (16) and also compute the market-wise average number of insiders a principal has at any time, which is defined as $\mathbb{E}_t[d] = \sum_d \mathbf{f}_{k_{\theta_t}^*, t}(d) \cdot d$.

6 Robustness check: Limited Supply Capacity

6.1 Incentive compatibility

We will now examine the case where an agent has supply capacity L and thus can only service L principals at a time. In order to make the analysis clearer, we will deal with the model of Section 4, ignoring anticipated regime switches, although this will not affect the conclusions.

To model this, we will assume that the number of outsider draws are limited to $L - d^a$ for each agent, where d^a is the agent's degree (the number of principals to which he is an insider). This implies that $d^a \leq L$ at any time. This can model a decreasing search ability on the part of the agent as he accumulates relationships. An agent thus has L 'stubs' to which principals can be connected. Then to derive an incentive compatibility condition, we must consider all possibilities in which an agent could cheat. Say that, at time t , an agent of degree $d^a \leq L$ is solicited by $n \leq L$ principals on a subset $\mathcal{N} \subset [1, 2, \dots, L]$ of his stubs (as an insider and/or as an outsider). Then for any $m \leq n$ principals, on a subset $\mathcal{M} \subseteq \mathcal{N}$ of stubs on which the agent might consider cheating, incentive compatibility requires that

$$(mv + \sum_{l \in \mathcal{M}} \delta U_i^l(O, k)) \leq \min_{\{d_l\}_{l \in \mathcal{M}}} (mp + \sum_{l \in \mathcal{M}} \delta U_i^l(I, d_l, k))$$

Canceling out the equal terms on both sides, we get

$$\begin{aligned} m \cdot (v + \delta U_i^l(O, k)) &\leq \min_{\{d_l\}_{l \in \mathcal{M}}} (mp + \sum_{l=1}^m \delta U_i^l(I, d_l, k)) \\ &= m \cdot (p + \delta U_i^l(I, k, k)) \end{aligned}$$

which reduces nicely to

$$v + \delta U_i^l(O, k) \leq p + \delta U_i^l(I, k, k)$$

or

$$U_i^l(I, k, k) \geq \frac{v - p}{\delta} + U_i^l(O, k)$$

which indicates that k^* may be lower than in the infinite capacity case, because here cheating 'frees up' space on one of the L slots and thus has higher option value than when $L = \infty$.

6.2 Agents' utilities

We let the mass of principals be normalized to 1 and the mass of agents be x . At each time, a principal looking for an outsider is matched randomly to an agent's free 'stub'. To guarantee that this matching process can take place, x will now have to be large enough. We may thus let $x \geq \frac{\bar{k}^* + 1}{L}$, where $\bar{k}^* = \max_{L \in \mathbb{N}_+} k_L^*$ and k_L^* is the equilibrium maximum number of insiders when each agent has supply capacity L . Such an upper bound \bar{k}^* exists and is finite since k_L^* is bounded. The above choice of x thus guarantees that a principal is always matched to a free agent stub, since $xL \geq \bar{k}^* + 1$ (i.e. the total supply of agent stubs is larger than the total principals' demand for agents stubs) and, as before, the agent (or stub) will be 'available' with exogenous probability α . Then the calculation of the transition matrix $\mathbf{M}_P(k)$ and a principal's profit $\Pi(k)$ are exactly as in the infinite capacity case.

An agent i 's expected utility vector from a relationship with a principal on stub l , when this principal commits to a k -insider contract and the agent cooperates whenever available, is then

$$U_i^l(k) = (1 - q)p\mathbb{E}[Q_i^l(k)] + \delta \mathbf{M}_A(k)U_i^l(k)$$

where $\mathbf{U}_i^l(k) = [U_i^l(O, k), U_i^l(I, 1, k), \dots, U_i^l(I, k, k)]^\top$ is the vector of utilities in all possible states the agent can be on stub l and

$$\mathbb{E}[\mathbf{Q}_i^l(k)] = \begin{bmatrix} \mathbb{E}[Q_i^l(O, k)] \\ \mathbb{E}[Q_i^l(I, 1, k)] \\ \vdots \\ \mathbb{E}[Q_i^l(I, k, k)] \end{bmatrix}$$

is the probability that agent i will be chosen to trade with a principal on stub l in each possible state. Note that $\mathbb{E}[Q_i^l(I, d, k)] = \alpha \mathbb{E}[\frac{1}{1+n_d^d}]$ where $n_\alpha^d \sim \text{bin}(\alpha, d-1)$, as in the infinite capacity case. On the other hand, $\mathbb{E}[Q_i^l(O, k)] = \sum_{d=1}^k \frac{a_{O,d}}{1-q} > 0$ is now the probability the agent will be solicited for a trade as an outsider by some new principal on stub l , after having been fired by the previous principal connected to that stub. Note that this probability is now strictly positive.

The transition matrix $\mathbf{M}_A(k)$ is now

$$\mathbf{M}_A(k) = \begin{bmatrix} a_{O,O} & a_{O,1} & \cdots & a_{O,k} \\ a_{1,O} & a_{1,1} & \cdots & a_{1,k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k,O} & a_{k,1} & \cdots & a_{k,k} \end{bmatrix} \quad (27)$$

where,

$$\begin{aligned} a_{d,O} &= \frac{(1 - (1 - \alpha)^d)}{d} q \\ a_{O,d} &= \alpha \frac{1 - f_k(k)}{xL - \sum_{d=0}^k f_k(d)d} f_k(d-1)(1 - \alpha)^{d-1}(1 - q) \\ a_{O,O} &= 1 - \sum_{d=1}^k a_{O,d} \end{aligned}$$

$a_{O,d}$ now reflects the probability of being selected as an outsider on a given stub by a new principal with $d-1$ insiders. Fixing k , we see that this probability decreases in L (which will ultimately drive $U_i^l(O, k)$ to 0, as in the infinite capacity case).

The entries for $1 \leq d \leq k$ are as in the infinite capacity case:

$$\begin{aligned} a_{d,d-1} &= \frac{(d-1)(1 - (1 - \alpha)^d)}{d} q \\ a_{d,d} &= (1 - (1 - \alpha)^d)(1 - q) + (1 - \alpha)^d \alpha q + (1 - \alpha)^d (1 - \alpha) \\ a_{d,d+1} &= (1 - \alpha)^d \alpha (1 - q), \end{aligned}$$

and $a_{k,k} = (1 - (1 - \alpha)^k)(1 - q) + (1 - \alpha)^k$. All other entries are 0.

6.3 Equilibrium existence and characterization

Note that the population feedback through $U_i^l(O, k)$ can lead to a situation where

$$U_i^l(I, k, k) - U_i^l(O, k) \geq \frac{v-p}{\delta} \quad \text{and} \quad U_i^l(I, k+1, k+1) - U_i^l(O, k+1) < \frac{v-p}{\delta}$$




Figure 4: Equilibrium $k^* + \mu^*$ and corresponding welfare versus L . L is varied from 1 to 40. Other parameters are $q = 0.1$, $v = 3$, $p = 0.5$, $\delta = 0.99$, $\alpha = 0.2$, $c = 0.2$, $x = 7$.

but

$$U_i^l(I, k+1, k+1) - U_i^l(O, k) \geq \frac{v-p}{\delta}$$

and thus there may be cases where a symmetric strategy equilibrium k may not exist.

We will therefore look for a potentially non-symmetric equilibrium. When all model primitives are homogeneous across principals, then the profit functions and incentive compatibility conditions are the same across principals. A principal will thus always prefer deviating from a non-symmetric equilibrium and choosing a contract with the highest possible k . To guarantee existence of a non-symmetric equilibrium, we can perturb the principals' payoffs by letting $v_j \sim \text{unif}(v-r, v+r)$ for some $r > 0$ and then look for an equilibrium in the limit as $r \rightarrow 0$.

We now have the following proposition.

Proposition 6 (Equilibrium with limited supply capacity) *Consider the setting where payoffs are perturbed by allowing the value of the good produced for a principal j to be $v_j \sim \text{unif}(v-r, v+r)$ for some $r > 0$. Then, in the limit as $r \rightarrow 0$, there exists an equilibrium in which a fraction $\mu^* \in [0, 1)$ of principals choose a (k^*+1) -insider contract while the rest choose a k^* -insider contract, where $k^* = \max\{k : U_i^l(I, k, k) - U_i^l(O, k) \geq \frac{v-p}{\delta}\}$.*

Note that $\mu^* = 0$ occurs with positive probability and corresponds to a symmetric equilibrium, as the one studied so far.

Fig. 4 shows equilibria for various values of L . Between discrete values of k^* , the non-symmetric equilibrium is represented by $k^* + \mu^*$. We see that when L decreases, the corresponding k^* decreases, reflecting the fact that principals must keep fewer insider agents in order to preserve incentive compatibility. This further decreases welfare, as fewer trade opportunities necessarily occur. As L grows large, k^* converges to its value in the infinite capacity model.

References

- ALI, S. N. AND D. A. MILLER (2013): “Enforcing cooperation in networked societies,” *Unpublished paper*. [282].
- (2016): “Ostracism and forgiveness,” *The American Economic Review*, 106, 2329–2348.
- ANDREWS, I. AND D. BARRON (2016): “The allocation of future business: Dynamic relational contracts with multiple agents,” *The American Economic Review*, 106, 2742–2759.
- BOARD, S. (2011): “Relational contracts and the value of loyalty,” *The American Economic Review*, 101, 3349–3367.
- BRÄUNING, F. AND F. FECHT (2012): “Relationship lending and peer monitoring: Evidence from interbank payment data,” *Available at SSRN 2020171*.
- CALZOLARI, G. AND G. SPAGNOLO (2009): “Relational contracts and competitive screening,” *CEPR discussion paper no. DP7434*.
- COCCO, J. F., F. J. GOMES, AND N. C. MARTINS (2009): “Lending relationships in the interbank market,” *Journal of Financial Intermediation*, 18, 24–48.
- DASGUPTA, P. AND I. SERAGELDIN (2001): *Social capital: a multifaceted perspective*, World Bank Publications.
- ELLIOTT, M. (2015): “Inefficiencies in networked markets,” *American Economic Journal: Microeconomics*, 7, 43–82.
- FAFCHAMPS, M. AND B. MINTEN (1999): “Relationships and traders in Madagascar,” *The Journal of Development Studies*, 35, 1–35.
- GRANOVETTER, M. (1985): “Economic action and social structure: The problem of embeddedness,” *American journal of sociology*, 481–510.
- GREIF, A. (1993): “Contract enforceability and economic institutions in early trade: The Maghribi traders’ coalition,” *The American economic review*, 525–548.
- HENDERSHOTT, T., D. LI, D. LIVDAN, AND N. SCHÜRHOFF (2020): “Relationship trading in over-the-counter markets,” *The Journal of Finance*, 75, 683–734.
- KANDORI, M. AND S. OBAYASHI (2014): “Labor union members play an OLG repeated game,” *Proceedings of the National Academy of Sciences*, 111, 10802–10809.
- KIRMAN, A. P. AND N. J. VRIEND (2001): “Evolving market structure: An ACE model of price dispersion and loyalty,” *Journal of Economic Dynamics and Control*, 25, 459–502.
- KRANTON, R. AND D. F. MINEHART (2001): “A Theory of Buyer-Seller Networks,” *American Economic Review*, 91, 485–508.
- LI, S. (2017): “Obviously strategy-proof mechanisms,” *American Economic Review*, 107, 3257–87.

- McMILLAN, J. AND C. WOODRUFF (1999): “Interfirm relationships and informal credit in Vietnam,” *The Quarterly Journal of Economics*, 114, 1285–1320.
- NAVA, F. (2015): “Efficiency in decentralized oligopolistic markets,” *Journal of Economic Theory*, 157, 315–348.
- NAVA, F. AND M. PICCIONE (2014): “Efficiency in repeated games with local interaction and uncertain local monitoring,” *Theoretical Economics*, 9, 279–312.
- NISHIGUCHI, T. (1994): *Strategic industrial sourcing: The Japanese advantage*, Oxford University Press on Demand.
- PETERSEN, M. A. AND R. G. RAJAN (1994): “The benefits of lending relationships: Evidence from small business data,” *The journal of finance*, 49, 3–37.
- (1995): “The effect of credit market competition on lending relationships,” *The Quarterly Journal of Economics*, 407–443.
- SAMUELSON, L. AND E. STACCHETTI (2017): “Even up: Maintaining relationships,” *Journal of Economic Theory*, 169, 170–217.
- SOBEL, J. (2002): “Can we trust social capital?” *Journal of economic literature*, 40, 139–154.
- UZZI, B. (1997): “Social structure and competition in interfirm networks: The paradox of embeddedness,” *Administrative science quarterly*, 35–67.
- WOLITZKY, A. (2013): “Cooperation with network monitoring,” *The Review of Economic Studies*, 80, 395–427.

7 Appendix

7.1 More details on the regime-switching model

7.1.1 More details on the expressions in Section 5.2.1 (Principal’s problem)

We now give the expressions of the submatrices $\bar{\mathbf{M}}_P(k_0, k_1)$ and $\hat{\mathbf{M}}_P(k_0, k_1)$ in Eq. (19). We let $k_0 > k_1$.

More specifically,

$$\bar{\mathbf{M}}_P(k_0, k_1) = \begin{bmatrix} \mathbf{M}_P^{q_1}(k_1) \\ \mathbf{N} \end{bmatrix} \quad (28)$$

where \mathbf{N} captures the transitions after the occurrence of a switch to the crisis regime, when a number of randomly chosen insiders must be fired:

$$\mathbf{N} = \begin{bmatrix} a_{1,0} & a_{1,1} & \cdots & a_{1,k_1} \\ \vdots & \vdots & \ddots & \vdots \\ a_{(k_0-k_1),0} & a_{(k_0-k_1),1} & \cdots & a_{(k_0-k_1),k_1} \end{bmatrix} \quad (29)$$

is a $(k_0 - k_1) \times (k_1 + 1)$ matrix where for any $1 \leq d \leq (k_0 - k_1)$

$a_{d,k_1-1} = (1 - (1 - \alpha)^{k_1})q_1$
 $a_{d,k_1} = (1 - (1 - \alpha)^{k_1})(1 - q_1) + (1 - \alpha)^{k_1}$
 and all other entries are 0.

Finally,

$$\hat{\mathbf{M}}_P(k_0, k_1) = \begin{bmatrix} \tilde{\mathbf{M}}_P & \mathbf{0} \end{bmatrix} \quad (30)$$

where $\tilde{\mathbf{M}}_P$ is a $(k_1 + 1) \times (k_1 + 2)$ matrix corresponding to the first $k_1 + 1$ rows of $\mathbf{M}_P^{q_0}(k_1 + 1)$, while $\mathbf{0}$ is a $(k_1 + 1) \times (k_0 - k_1 - 1)$ matrix of zeros.

In Eq. (22), \mathbf{R}_P is, for each state a principal can be in, the expected profit in the current period. It can be expressed as

$$\mathbf{R}_P = \begin{bmatrix} s_{00}\alpha((1 - q_0)(v - p) - q_0c) + s_{01}\alpha((1 - q_1)(v - p) - q_1c) \\ \vdots \\ s_{00}(1 - (1 - \alpha)^{d+1})((1 - q_0)(v - p) - q_0c) + s_{01}(1 - (1 - \alpha)^{\min(d+1, k_1)})((1 - q_1)(v - p) - q_1c) \\ \vdots \\ s_{00}(1 - (1 - \alpha)^{k_0})((1 - q_0)(v - p) - q_0c) + s_{01}(1 - (1 - \alpha)^{k_1})((1 - q_1)(v - p) - q_1c) \\ s_{11}\alpha((1 - q_1)(v - p) - q_1c) + s_{10}\alpha((1 - q_0)(v - p) - q_0c) \\ \vdots \\ s_{11}(1 - (1 - \alpha)^{d+1})((1 - q_1)(v - p) - q_1c) + s_{10}(1 - (1 - \alpha)^{d+1})((1 - q_0)(v - p) - q_0c) \\ \vdots \\ s_{11}(1 - (1 - \alpha)^{k_1})((1 - q_1)(v - p) - q_1c) + s_{10}(1 - (1 - \alpha)^{\min(d+1, k_0)})((1 - q_0)(v - p) - q_0c) \end{bmatrix}. \quad (31)$$

7.1.2 More details on the expressions in Section 5.2.2 (Agent's problem)

We now give the expressions of the submatrices $\tilde{\mathbf{M}}_A(k_0, k_1)$ and $\hat{\mathbf{M}}_A(k_0, k_1)$ in Eq. (23). We let $k_0 > k_1$.

More specifically,

$$\tilde{\mathbf{M}}_A(k_0, k_1) = \begin{bmatrix} \mathbf{M}_A^{q_1}(k_1) \\ \mathbf{L} \end{bmatrix} \quad (32)$$

where \mathbf{L} captures the transitions after the occurrence of a switch to the crisis regime, when a number of randomly chosen insiders must be fired:

$$\mathbf{L} = \begin{bmatrix} a_{1,0} & a_{1,1} & \cdots & a_{1,k_1} \\ \vdots & \vdots & \ddots & \vdots \\ a_{(k_0-k_1),0} & a_{(k_0-k_1),1} & \cdots & a_{(k_0-k_1),k_1} \end{bmatrix} \quad (33)$$

is a $(k_0 - k_1) \times (k_1 + 1)$ matrix where for any $1 \leq d \leq (k_0 - k_1)$

$$\begin{aligned}
 a_{d,0} &= \frac{d}{k_1+d} + \left(1 - \frac{d}{k_1+d}\right) \frac{1-(1-\alpha)^{k_1}}{k_1} q_1 \\
 a_{d,k_1-1} &= \left(1 - \frac{d}{k_1+d}\right) \frac{k_1-1}{k_1} (1 - (1 - \alpha)^{k_1}) q_1 \\
 a_{d,k_1} &= \left(1 - \frac{d}{k_1+d}\right) ((1 - (1 - \alpha)^{k_1})(1 - q_1) + (1 - \alpha)^{k_1})
 \end{aligned}$$

and all other entries are 0.

Finally,

$$\hat{\mathbf{M}}_A(k_0, k_1) = \begin{bmatrix} \tilde{\mathbf{M}}_A & \mathbf{0} \end{bmatrix} \quad (34)$$

where $\tilde{\mathbf{M}}_A$ is a $(k_1 + 1) \times (k_1 + 2)$ matrix corresponding to the first $k_1 + 1$ rows of $\mathbf{M}_A^{q_0}(k_1 + 1)$ while $\mathbf{0}$ is a $(k_1 + 1) \times (k_0 - k_1 - 1)$ matrix of zeros.

In Eq. (26), \mathbf{R}_A is, for each state an agent can be in, the probability of being chosen and not failing in the current period. It can be expressed as

$$\mathbf{R}_A = \begin{bmatrix} 0 \\ \vdots \\ s_{00}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^d}\}(1-q_0) + s_{01}\min(\frac{k_1}{d}, 1)\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^d}\}(1-q_1) \\ \vdots \\ s_{00}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^{k_0}}\}(1-q_0) + s_{01}\min(\frac{k_1}{k_0}, 1)\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^{k_0}}\}(1-q_1) \\ 0 \\ \vdots \\ s_{11}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^d}\}(1-q_1) + s_{10}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^d}\}(1-q_0) \\ \vdots \\ s_{11}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^{k_1}}\}(1-q_1) + s_{10}\alpha\mathbb{E}\{\frac{1}{1+n_\alpha^{k_1}}\}(1-q_0) \end{bmatrix} \quad (35)$$

where d runs from 1 to k_θ and where $n_\alpha^d \sim \text{Bin}(\alpha, d - 1)$, as seen before.

7.2 Proofs

Proof of Remark 1 (No mixing when punishing). Suppose the contract Γ prescribes that the principal punishes a failed trade with probability $\mu \in (0, 1)$. Then punishing with a probability smaller than μ is a strictly profitable deviation for the principal since he can keep more insiders than he would under μ and thus may increase the frequency of future trades. Moreover, such a deviation is non detectable, as defined in Definition 1. Indeed, agents cannot detect it by observing any history h^t . The contract Γ is thus not principal credible and this violates Property 4. On the other hand, if a contract prescribes a pure punishment strategy $\mu = 0$ or $\mu = 1$, then any deviation is detectable and thus there are no strictly profitable non-detectable deviations. Such a contract is thus principal credible. ■

Proof of Remark 2 (No mixing when hiring a new insider). Suppose the contract Γ prescribes that, when there are no available insiders, the principal hires an outsider with probability $\tilde{\mu} \in (0, 1)$. Then hiring the outsider with a probability larger than $\tilde{\mu}$ is a strictly profitable deviation for the principal since he can keep more insiders than he would under $\tilde{\mu}$ and thus may increase the frequency of future trades. As in the previous proof, such a deviation is non-detectable. This contract is thus not principal credible, and this violates Property 4. On the other hand, if a contract prescribes a pure hiring strategy $\tilde{\mu} = 0$ or $\tilde{\mu} = 1$, then any deviation is detectable and thus there are no strictly profitable non-detectable deviations. Such a contract is thus principal credible.

That the decision to hire an outsider can only be conditioned on the current number of insiders follows directly from Property 5.

That following a successful trade, insider status must be granted to the outsider with probability 0 or 1 follows from a similar credibility argument: if the contract specified that the principal granted insider status with probability $\mu \in (0, 1)$, then the principal would always prefer deviating and granting it with a higher probability since he would then be able to trade with more insiders in

future periods. Since this strictly-profitable deviation is non-detectable, then the contract is not principal credible, whereas a contract prescribing a pure strategy $\mu = 0$ or $\mu = 1$ would be credible since any deviation would be immediately detectable. ■

Proof of Lemma 1. Suppose, towards a contradiction, there is a constrained-efficient relational contract Γ in which after some history h^{t-1} , an insider who either did not trade or who traded successfully is made an outsider in period t with positive probability $\mu > 0$. Equal treatment of insiders, Property 2, then implies that at the start of period $t - 1$ all insiders are equally likely to be made outsiders following no trade or a successful trade. Consider now an alternative contract Γ' that is identical to Γ , except that $\mu = 0$. In some random realizations, immediately after h^{t-1} , there will be an additional insider under Γ' that the principal can trade with. Moreover, trade with an outsider that would be incentive-compatible under Γ may not be incentive-compatible under Γ' , in which case it is not attempted under Γ' . Γ' then is associated with a larger insider set than Γ until either: (i) a trade that would not have occurred under Γ , but occurs under Γ' , fails and the agent is fired, or (ii) a trade that would have been successful with an outsider under Γ and resulted in that outsider becoming an insider is not undertaken under Γ' . After such an event there is the same number of insiders under Γ' as there is under Γ and henceforth Γ' follows Γ exactly.

The principal prefers Γ' to Γ . By construction the principal has the same sequence of trades except there are now, in expectation, additional trades. The expected value of these additional trades is positive $(1 - q)(v - p) - qc > 0$ and thus the principal is better off.

The contract Γ' is incentive-compatible. The initial contract Γ cannot, therefore, have been constrained-efficient. ■

Proof of Lemma 2. Without loss of generality, suppose that nature initially draws an infinite sequence of binary random variables, where each random variable is drawn independently from each other and is 1 with probability $1 - q$ and 0 with probability q . We then let the success of initiated trade conditional on no defection follow this sequence, where 1 represents a success and 0 represents a failure.

Suppose, towards a contradiction, there exists a constrained-efficient relational contract Γ in which, following some history h^t , the principal does not trade with an insider in period $t + 1$ despite there being an available insider. Property 3 (Insider's priority) then implies that the principal does not trade in period $t + 1$. Further, by Property 1 (No defection on path), all trade under this contract is incentive-compatible.

Consider now the relational contract Γ' that is almost identical to Γ . We will assume and then verify that Γ' is also incentive-compatible. Let Γ' differ from Γ in (only) the following way: After history h^t , in period $t + 1$ the principal trades with an insider.

If trade is successful in period $t + 1$, Γ' follows Γ henceforth. The expected value to the principal from period $t + 2$ onwards is therefore the same under Γ' and Γ . Moreover, as Γ was incentive-compatible from period $t + 2$ onwards, so is Γ' . Also, an additional successful trade in period $t + 1$ yields profits for the principal of $v - p$ and additional rents for the trading agent i of p .

If instead the extra trade in period $t + 1$ under Γ' is unsuccessful, then the principal suffers a loss of c . Following a given path of availabilities, let the next time the principal would initiate trade under Γ be in period $s > t + 1$. The only other difference between Γ' and Γ is that in period s trade may not be initiated under Γ' (due to the failing agent having been potentially fired).

If the trading agent j fails under Γ in period s , then the principal under Γ suffers a loss of c and

treats agent j exactly as the failing agent i would have been treated under Γ' . Then, from period $s + 1$ onwards, Γ' again exactly mimics Γ in terms of how each insider is treated. If instead the trade succeeds under Γ in period s , the principal obtains $v - p$, the agents obtains p and from period $s + 1$ onwards, Γ' again exactly mimics Γ in terms of how each insider is treated.

Thus, under Γ , a trade with an expected value to the principal of $(1 - q)(v - p) - qc > 0$ has simply been moved forward from period $t + 1$ to period $s > t + 1$. Given that the discount factor is $\delta < 1$, the principal strictly prefers Γ' to Γ .

Finally, the contract Γ' is still incentive-compatible. Indeed, relative to Γ , a trade is moved backwards in time and since $\delta < 1$, this increases the continuation value of all the insiders.

■

Proof of Lemma 3. By Lemma 1 and Property 2 (equal treatment of insiders), an insider's expected payoff is highest when the principal never again trades with an outsider. In this case, the expected payoff of all remaining insiders in a given period, when there are $k > 1$ insiders, is:

$$V_{\text{tot}}(k) \equiv (1 - (1 - \alpha)^k)[(1 - q)(p + \delta V_{\text{tot}}(k)) + q\delta V_{\text{tot}}(k - 1)] + (1 - \alpha)^k \delta V_{\text{tot}}(k).$$

When there is 1 insider

$$V_{\text{tot}}(1) = \alpha[(1 - q)(p + \delta V_{\text{tot}}(1))]$$

and thus

$$V_{\text{tot}}(1) = \frac{\alpha(1 - q)p}{1 - \alpha(1 - q)\delta}.$$

We can therefore solve this system to find an expression for $V_{\text{tot}}(k)$.

An upper bound on the continuation value of any agent will then, by the equal treatment Property, be $V_{\text{tot}}(k)/k$, which is decreasing in k and converges to 0 as $k \rightarrow \infty$. Denoting by $V_i(k)$ any agent's expected payoff when there are k insiders, any agent's incentive-compatibility constraint is

$$p + \delta V_i(k) \geq v$$

or

$$V_i(k) \geq \frac{v - p}{\delta}.$$

There thus exists a \bar{k} such that the IC constraint of the agent is not satisfied, i.e.

$$V_i(k) \leq \frac{V_{\text{tot}}(k)}{k} \leq \frac{v - p}{\delta}.$$

This completes the proof. ■

Proof of Lemma 4. By Remark 2 the principal cannot mix over trading with the outsider and must grant insider status following successful trade (otherwise, by Remark 2, the outsider must not be granted insider status for sure and would choose to defect violating Property 1). Moreover, Remark 2 implies that d insiders can be obtained only if the principal always trades with an available outsider when there are $d - 1$ insiders, all of whom are unavailable and the outsider is available. Consider two such contracts and suppose both satisfy Property 1 and Remark 2. Let Γ_d be such a contract in which the principal keeps at most d insiders. Let Γ_{d+1} be a similar contract, but in which the principal keeps at most $d + 1$ insiders. Under both Γ_d and Γ_{d+1} , the incentive-compatibility constraint of trading agents are always satisfied by Property 1.

Under these two contracts, the principal trades more often under Γ_{d+1} : On all paths where there are strictly less than d insiders, the two contracts are identical. On paths where there are d insiders and none is available and the outsider is not available, the two contracts are identical. On paths where there are d unavailable insiders and the outsider is available, the principal trades and earns an expected period payoff of $(1 - q)(v - p) - qc > 0$ under Γ_{d+1} , whereas he does not trade under Γ_d . On paths where there are $d + 1$ insider and only one is available, the principal trades and earns a positive expected payoff under Γ_{d+1} , whereas he may or may not trade under Γ_d . Therefore Γ_{d+1} does strictly better than Γ_d on some paths and the same on other paths. It is thus strictly preferred by the principal. This is valid for all d and $d + 1$ such that the trading agents' IC constraints are always satisfied. By Lemma 3 there is a maximum such d . ■

Proof of Lemma 5.

Eq. (9) follows directly from Eq. (6).

Now to show monotonicity, note that for any k , we may write

$$\Pi_j(d, k) = \left(\mathbb{E}[Q^j(d, k)] + \sum_{s=t+1}^{\infty} \delta^{s-t} \mathbf{M}_{P,d}^{s-t}(k) \mathbb{E}[Q^j(k)] \right) ((1 - q)(v - p) - qc) \quad (36)$$

where $\mathbf{M}_{P,d}^{s-t}(k)$ is the d^{th} row in the $(s - t)$ -period transition matrix $\mathbf{M}_P^{s-t}(k)$. Note that the dot product $\mathbf{M}_{P,d}^{s-t}(k) \mathbb{E}[Q^j(k)]$ is the probability that the principal j will be able to select an agent to trade with at time $s > t$, given that he has d insiders at time t and aims to keep a maximum of k .

Let $\mathbb{P}(\tilde{d} \mid d, k)$ be the one-period transition probability to \tilde{d} insiders, given that the principal currently has d insiders. $\mathbb{P}(\tilde{d} \mid d, k)$ corresponds to the d^{th} row $\mathbf{M}_{P,d}(k)$ of the transition matrix $\mathbf{M}_P(k)$. It is easy to show by inspection that $\mathbb{P}(\tilde{d} \mid d + 1, k) \succeq \mathbb{P}(\tilde{d} \mid d, k)$, where \succeq means first-order stochastic dominance. Equivalently, we may write $\mathbf{M}_{P,d+1}(k) \succeq \mathbf{M}_{P,d}(k)$. Since $\mathbb{E}[Q^j(d, k)]$ is an increasing function of d (see Eq. (7)), it follows that $\mathbf{M}_{P,d+1}(k) \mathbb{E}[Q^j(k)] > \mathbf{M}_{P,d}(k) \mathbb{E}[Q^j(k)]$.

Now suppose that for any d and for some $l \geq 1$, $\mathbf{M}_{P,d+1}^l(k) \mathbb{E}[Q^j(k)] > \mathbf{M}_{P,d}^l(k) \mathbb{E}[Q^j(k)]$. Then it follows that

$$\mathbf{M}_{P,d+1}^{l+1}(k) \mathbb{E}[Q^j(k)] = \mathbf{M}_{P,d+1}(k) (\mathbf{M}_P^l(k) \mathbb{E}[Q^j(k)]) \quad (37)$$

$$> \mathbf{M}_{P,d}(k) (\mathbf{M}_P^l(k) \mathbb{E}[Q^j(k)]) \quad (38)$$

$$= \mathbf{M}_{P,d}^{l+1}(k) \mathbb{E}[Q^j(k)] \quad (39)$$

Thus, it follows by induction that for any $s - t \geq 1$, $\mathbf{M}_{P,d+1}^{s-t}(k) \mathbb{E}[Q^j(k)] > \mathbf{M}_{P,d}^{s-t}(k) \mathbb{E}[Q^j(k)]$ and thus from Eq. (36), we have that $\Pi_j(d, k) < \Pi_j(d + 1, k)$. ■

Proof of Lemma 6.

Eq. (13) follows directly from Eq. (10).

Now to show monotonicity, note that for any k , we may write

$$U(I, d, k) = \mathbb{E}[Q_i^j(I, d, k)](1 - q)p + \sum_{s=t+1}^{\infty} \delta^{s-t} \mathbf{M}_{A,d}^{s-t}(k) \mathbb{E}[Q_i^j(k)](1 - q)p \quad (40)$$

where $\mathbf{M}_{A,d}^{s-t}(k)$ is the d^{th} row in the $(s - t)$ -period transition matrix $\mathbf{M}_A^{s-t}(k)$. Note that the dot product $\mathbf{M}_{A,d}^{s-t}(k) \mathbb{E}[Q_i^j(k)]$ is the probability that agent i will be selected for trading by principal j

at time s , given that he is in state (I, d, k) at time t .

Let $\mathbb{P}((I, \tilde{d}) \mid d, k)$ be the one-period transition probability to state (I, \tilde{d}) , given that the principal currently has d insiders (i.e. $\mathbb{P}((I, \tilde{d}) \mid d, k)$ corresponds to the d^{th} row $\mathbf{M}_{A,d}(k)$ of the transition matrix $\mathbf{M}_A(k)$). It is easy to show by inspection that $\mathbb{P}((I, \tilde{d}) \mid d+1, k) \succeq \mathbb{P}((I, \tilde{d}) \mid d, k)$, where \succeq indicates first-order stochastic dominance. Equivalently, we may write $\mathbf{M}_{A,d+1}(k) \succeq \mathbf{M}_{A,d}(k)$. Since $\mathbb{E}[Q_i^j(I, d, k)]$ is a decreasing function of d (from Eq. (11)), it follows that $\mathbf{M}_{A,d+1}(k)\mathbb{E}[Q_i^j(k)] < \mathbf{M}_{A,d}(k)\mathbb{E}[Q_i^j(k)]$.

Now suppose that for any d and for some $l \geq 1$, $\mathbf{M}_{A,d+1}^l(k)\mathbb{E}[Q_i^j(k)] < \mathbf{M}_{A,d}^l(k)\mathbb{E}[Q_i^j(k)]$. Then it follows that

$$\mathbf{M}_{A,d+1}^{l+1}(k)\mathbb{E}[Q_i^j(k)] = \mathbf{M}_{A,d+1}(k)(\mathbf{M}_{A,d+1}^l(k)\mathbb{E}[Q_i^j(k)]) \quad (41)$$

$$< \mathbf{M}_{A,d}(k)(\mathbf{M}_{A,d}^l(k)\mathbb{E}[Q_i^j(k)]) \quad (42)$$

$$= \mathbf{M}_{A,d}^{l+1}(k)\mathbb{E}[Q_i^j(k)] \quad (43)$$

It thus follows by induction that for any $s - t > 1$, $\mathbf{M}_{A,d+1}^{s-t}(k)\mathbb{E}[Q_i^j(k)] < \mathbf{M}_{A,d}^{s-t}(k)\mathbb{E}[Q_i^j(k)]$ and thus from Eq. (40), we have that $U_i^j(I, d, k) > U_i^j(I, d+1, k)$. ■

Proof of Lemma 7. Whenever the agent is selected, his payoff from not deviating in state (I, d, k) is $p + \delta U_i^j(I, d, k)$, i.e. the price he obtains from the transaction plus his continuation value after the trade. His payoff from deviating and then reverting back to on-path play is $v + \delta U_i^j(O, k)$, i.e. the value of the good plus his continuation value when fired by the principal. Thus i will not have a profitable one-period deviation if and only if $p + \delta U_i^j(I, d, k) \geq v + \delta U_i^j(O, k)$. Invoking the one-shot deviation principle completes the proof. ■

Proof of Proposition 2.

We first show that any equilibrium must be symmetric. Let $\Pi_j(d, k_j)$ denote principal j 's expected profit in state (d, k_j) . It is easy to show, by an argument similar to that of Lemma 5, that $\Pi_j(d, k_j)$ is increasing in k_j . Principal j 's optimal strategy is thus necessarily the largest k such that $U_i^j(I, k, k) \geq \frac{v-p}{\delta}$. From Lemma 6, it then follows that $U_i^j(I, d, k) \geq \frac{v-p}{\delta}$, $\forall d \leq k$ and thus the incentive-compatibility constraints of all insiders are satisfied.

Any k^* such that $U_i^j(I, k^*, k^*) \geq \frac{v-p}{\delta}$ and $U_i^j(I, k^*+1, k^*+1) < \frac{v-p}{\delta}$ is an equilibrium. Note that all principals face the same decision problem. It follows that any equilibrium must be symmetric.

Any k^* that just precedes the crossing of $U_i^j(I, k^*, k^*)$ and $\frac{v-p}{\delta}$ is thus an equilibrium. There exists an equilibrium since $U_i^j(I, 1, 1) > \frac{v-p}{\delta}$ and $\lim_{k \rightarrow \infty} U_i^j(I, k, k) = 0 < \frac{v-p}{\delta}$ and thus there exists a crossing. Moreover, the equilibrium is unique since $U_i^j(I, k, k)$ is strictly decreasing in k and thus there is a single crossing of $U_i^j(I, k, k)$ and $\frac{v-p}{\delta}$. In this case, the unique, symmetric equilibrium strategy followed by all principals is $k^* = \max\{k \mid U_i^j(I, k, k) \geq \frac{v-p}{\delta}\}$.

When all principals follow strategy k^* , then in steady state, any principal's degree (the number of insiders he can potentially trade with) follows a distribution $f_{k^*}(d)$, which is the stationary distribution of the Markov transition matrix $\mathbf{M}_P(k^*)$. Since the Markov chain described by $\mathbf{M}_P(k^*)$ is time-homogenous, irreducible and all its states are positive recurrent, then this stationary distribution necessarily exists and is unique. ■

Proof of Proposition 3. Let x' and x be distribution vectors over the support $\{0, 1, 2, \dots, k\}$. We will first show that for any $x' \succeq x$, then $x'^\top \mathbf{M}_P(k) \succeq x^\top \mathbf{M}_P(k-1)$. Here $\mathbf{M}_P(k)$ is the

$(k+1) \times (k+1)$ Markov transition matrix in Eq. (8) when the principal keeps a maximum of k insiders. $\mathbf{M}_P(k-1)$ is the Markov transition matrix when the principal keeps a maximum of $k-1$ insiders with an additional last row and last column consisting of zeros so that $\mathbf{M}_P(k-1)$ is a $(k+1) \times (k+1)$ matrix.

First note that $x'^\top \mathbf{M}_P(k)$ and $x^\top \mathbf{M}_P(k-1)$ can be expressed as mixture distributions. Indeed, let $\mathbf{M}_{P,d}(k)$ be the d^{th} row of the transition matrix $\mathbf{M}_P(k)$. This row corresponds to the transition probability distribution $\mathbb{P}(\tilde{d} \mid d, k)$. We may then write $x'^\top \mathbf{M}_P(k) = \sum_{d=0}^k x'(d) \mathbf{M}_{P,d}(k)$ and $x^\top \mathbf{M}_P(k-1) = \sum_{d=0}^k x(d) \mathbf{M}_{P,d}(k-1)$. Moreover, $\mathbf{M}_{P,d}(k) \succeq \mathbf{M}_{P,d}(k-1)$ for all d .

If $x' \succeq x$, then it follows that $\sum_{d=0}^k x'(d) \mathbf{M}_{P,d}(k) \succeq \sum_{d=0}^k x(d) \mathbf{M}_{P,d}(k-1)$ and thus that $x'^\top \mathbf{M}_P(k) \succeq x^\top \mathbf{M}_P(k-1)$. Thus, if for some $l > 0$, $x'^\top \mathbf{M}_P^l(k) \succeq x^\top \mathbf{M}_P^l(k-1)$, then it follows that

$$\begin{aligned} x'^\top \mathbf{M}_P^{l+1}(k) &= (x'^\top \mathbf{M}_P^l(k)) \mathbf{M}_P(k) \\ &\succeq (x^\top \mathbf{M}_P^l(k-1)) \mathbf{M}_P(k-1) \\ &= x^\top \mathbf{M}_P^{l+1}(k-1) \end{aligned}$$

Thus, by induction we conclude that $x'^\top \mathbf{M}_P^l(k) \succeq x^\top \mathbf{M}_P^l(k-1)$ for any $l \geq 1$ and $x' \succeq x$. It thus follows that the stationary distributions satisfy

$$f_k(d) = \lim_{l \rightarrow \infty} x'^\top \mathbf{M}_P^l(k) \succeq \lim_{l \rightarrow \infty} x^\top \mathbf{M}_P^l(k-1) = f_{k-1}(d)$$

By transitivity, $f_{k'}(d) \succeq f_k(d)$ for any $k' > k$.

Now to show the last part, note that Eq. (16) can be rewritten as

$$W(k) = ((1-q)v - qc) \mathbb{E}[g(d)] \quad (44)$$

where $d \sim f_k(d)$ and

$$g(d) = \begin{cases} 1 - (1-\alpha)^{d+1} & \text{for } d \leq k-1 \\ 1 - (1-\alpha)^d & \text{for } d = k \end{cases}$$

is a non-decreasing function of d . When $k' > k$, then it follows from Proposition 3 that $f_{k'} \succeq f_k$. It thus follows in turn that $W(k') \geq W(k)$. ■

Proof of Proposition 4. Let $\underline{q} = \max\{\tilde{q} \in [q, 1] : U_{\tilde{q},i}^j(I, k^*, k^*) = \frac{v-p}{\delta}\}$, where k^* is the principal-optimal number of insiders under q . Such a \underline{q} exists since $U_{\tilde{q},i}^j(I, k, k)$ is continuous in \tilde{q} , since $U_{\tilde{q}=q,i}^j(I, k^*, k^*) \geq \frac{v-p}{\delta}$ and since $U_{\tilde{q}=1,i}^j(I, k^*, k^*) = 0$.

Then $\forall q' > \underline{q}$, we have that $U_{q',i}^j(I, k^*, k^*) < \frac{v-p}{\delta}$ and thus it must be that $k'^* \leq k^* - 1$, since this is the only way to potentially satisfy the IC constraint. Moreover, $U_{\underline{q},i}^j(I, k^* - 1, k^* - 1) > \frac{v-p}{\delta}$ and continuity in \tilde{q} implies that $U_{q',i}^j(I, k^* - 1, k^* - 1) > \frac{v-p}{\delta}$ for all $q' \in (\underline{q}, \underline{q} + \epsilon)$ for some $\epsilon > 0$. Thus, $k'^* = k^* - 1$ for all $q' \in (\underline{q}, \underline{q} + \epsilon)$, since this is the greatest number of insiders for which the incentive compatibility constraint is satisfied. ■

Proof of Corollary 1.

The result follows directly from Proposition 3. ■

Proof of Proposition 5.

Part (i):

We first show that trade can be incentive compatible in the crisis regime.

Let τ be the first transition time back to the non-crisis regime ($\theta = 0$). Suppose we are in the crisis regime ($\theta = 1$). We can write agent i 's utility from trading with principal j as

$$U_i^j(I, k_1, k_0, k_1, 1) = \mathbb{E}\left[\sum_{s=t}^{\tau} \delta^{s-t} \rho_t^i(s)(1 - q_1)p\right] + \mathbb{E}\left[\sum_{s=\tau}^{\infty} \delta^{s-t} \rho_{\tau}^i(s)(1 - q_{\theta_s})p\right]$$

where the expectation is taken over τ . $\rho_t^i(s)$ is the probability, at time t , that agent i will be selected to trade at time s , with $t \leq s < \tau$, before the first transition from regime $\theta = 1$ to regime $\theta = 0$. On the other hand, $\rho_{\tau}^i(s)$ is the probability, at time τ , that agent i will be selected to trade at time $s \geq \tau$, after the first transition from regime $\theta = 1$ to regime $\theta = 0$ and onwards into the future.

If $s'_{11} > s_{11}$, then $f'(\tau) \succeq f(\tau)$, where f denotes the pdf of τ and where \succeq means first-order stochastic dominance. Moreover as $s_{11} \rightarrow 1$, there exists $\underline{\tau}$ such that $f(\tau) \rightarrow 0, \forall \tau < \underline{\tau}$. It follows that $\mathbb{E}[\sum_{s=\tau}^{\infty} \delta^{s-t} \rho_{\tau}^i(s)(1 - q_{\theta_s})p] < B(s_{11})$ where $B(s_{11}) \rightarrow 0$ as $s_{11} \rightarrow 1$.

Now suppose that in the single regime model, $k(q_1) > 0$. Then, in the two-regime model, since $\mathbb{E}[\sum_{s=t}^{\tau} \delta^{s-t} \rho_t^i(s)(1 - q_1)p] \rightarrow U_{i,q_1}^j(I, k_1, k_1)$, there exists \underline{s}_{11} such that for $s_{11} > \underline{s}_{11}$, we have $k_1 > 0$.

It follows that, for $s_{11} > \underline{s}_{11}$, if the principal chooses to keep on trading in the crisis regime with k_1 insiders, he gets

$$\Pi^{trade}(k_1, k_0, k_1, 1) = \mathbb{E}\left[\sum_{s=t}^{\tau} \delta^{s-t} \rho_t(s)(1 - q_1)p - q_1c\right] + \mathbb{E}\left[\sum_{s=\tau}^{\infty} \delta^{s-t} \bar{\rho}_{\tau}(s)(1 - q_{\theta_s})p - q_{\theta_s}c\right] \quad (45)$$

where the first term is greater than 0 since $k_1 > 0$. In the above expression, $\rho_t(s)$ is the probability, at time t , that some agent will be selected to trade at time s , with $t \leq s < \tau$, before the first transition from regime $\theta = 1$ to regime $\theta = 0$. On the other hand, $\bar{\rho}_{\tau}(s)$ is the probability, at time τ , that some agent will be selected to trade at time $s \geq \tau$, after the first transition from regime $\theta = 1$ to regime $\theta = 0$ and onwards into the future, conditional on the fact that the principal has chosen to continue trading in the crisis regime from time t to time $\tau - 1$.

On the other hand, if he chooses to stop trading, for any number of insiders d that he currently holds, he gets

$$\Pi^{no\ trade}(d, k_0, k_1, 1) = 0 + \mathbb{E}\left[\sum_{s=\tau}^{\infty} \delta^{s-t} \hat{\rho}_{\tau}(s)(1 - q_{\theta_s})p - q_{\theta_s}c\right] \quad (46)$$

with $\hat{\rho}_{\tau}(s)$ being the probability, at time τ , that some agent will be selected to trade at time $s \geq \tau$, after the first transition from regime $\theta = 1$ to regime $\theta = 0$ and onwards into the future, conditional on the fact that the principal has chosen *not* to continue trading in the crisis regime from time t to time $\tau - 1$.

Since the second terms in both Eq. (45) and Eq. (46) can be bounded by $\bar{B}(s_{11})$ and $\hat{B}(s_{11})$ respectively, both bounds converging to 0 as $s_{11} \rightarrow 1$, then we can conclude that there exists \underline{s}_{11} such that for all $s_{11} > \underline{s}_{11}$, $\Pi^{trade}(k_1, k_0, k_1, 1) > \Pi^{no\ trade}(d, k_0, k_1, 1)$ for any $d \leq k_0$. It is thus optimal for a principal to keep on trading in the crisis regime with a reduced set of insiders, when the crisis regime is persistent enough.

Part (ii):

We can similarly express an agent's utility in the safe regime as

$$U_i^j(I, d, k_0, k_1, 0) = \mathbb{E}\left[\sum_{s=t}^{\tau} \delta^{s-t} \rho_t^i(s)(1 - q_0)p\right] + \mathbb{E}\left[\sum_{s=\tau}^{\infty} \delta^{s-t} \rho_{\tau}^i(s)(1 - q_{\theta_s})p\right]$$

where τ is now the first transition time to the crisis. We can now find a bound $B(s_{00})$ for the second term such that $B(s_{00}) \rightarrow 0$ as $s_{00} \rightarrow 1$ and $\mathbb{E}[\sum_{s=t}^{\tau} \delta^{s-t} \rho_t^i(s)(1 - q_0)p] \rightarrow U_{i,q_0}^j(I, d, k_0)$. Thus $k_0 \rightarrow k(q_0)$.

Since we supposed $k(q_1) < k(q_0)$ in the single-regime model, it thus follows that there exist \underline{s}_{00} and \underline{s}_{11} such that for all $s_{00} > \underline{s}_{00}$ and all $s_{11} > \underline{s}_{11}$, we have that $k_1 < k_0$.

Finally, satisfying Property 2(iv) (equal treatment of insiders after a transition to a crisis) means randomly selecting the subset of agents that will remain in the insider set after the transition to the crisis has been observed (step (0)). Property 4 (mixing must be credible) is satisfied since agents are homogenous and the principal has no incentive to favor one over another. ■

Proof of Proposition 6.

We first note that under such a non-symmetric strategy $(k, k+1, \mu)$, an agent's Markov transition matrix can be readily expressed as

$$\mathbf{M}_A(\mu) = \begin{bmatrix} \mathbf{F} & & \\ & \mathbf{M}_A(k)_{[1:k, 0:k]} & \mathbf{0}_{k \times (k+1)} \\ & \mathbf{M}_A(k+1)_{[1:k+1, 0]} & \mathbf{0}_{(k+1) \times k} & \mathbf{M}_A(k+1)_{[1:k+1, 1:k+1]} \end{bmatrix} \quad (47)$$

This is a $(2k+2) \times (2k+2)$ matrix formed by known submatrices, except for \mathbf{F} , which we will shortly define. For instance, $\mathbf{M}_A(k)_{[1:k, 0:k]}$ is the $k \times (k+1)$ submatrix formed by rows 1 to k and columns 0 to k of the $(k+1) \times (k+1)$ matrix $\mathbf{M}_A(k)$ in Eq. (27). $\mathbf{0}_{k \times (k+1)}$ is a $k \times (k+1)$ matrix of zeros and so on. On the other hand,

$$\mathbf{F} = [a_{O,O}, a_{O,1_k}, a_{O,2_k}, \dots, a_{O,k_k}, a_{O,1_{k+1}}, a_{O,2_{k+1}}, \dots, a_{O,k+1_{k+1}}]$$

are the probabilities of being recruited as an outsider by a principal following strategy k or $k+1$ and having degree $d-1$. These probabilities are found easily:

$$a_{O,d_k} = \alpha \frac{(1-\mu)(1-f_k(k))}{xL - ((1-\mu) \sum_{d=0}^k f_k(d)d + \mu \sum_{d=0}^{k+1} f_{k+1}(d)d)} f_k(d-1)(1-\alpha)^{d-1}(1-q)$$

$$a_{O,d_{k+1}} = \alpha \frac{\mu(1-f_{k+1}(k+1))}{xL - ((1-\mu) \sum_{d=0}^k f_k(d)d + \mu \sum_{d=0}^{k+1} f_{k+1}(d)d)} f_{k+1}(d-1)(1-\alpha)^{d-1}(1-q)$$

and $a_{O,O} = 1 - \sum_{d=1}^k a_{O,d_k} - \sum_{d=1}^{k+1} a_{O,d_{k+1}}$.

We now proceed to proving the proposition:

Consider first the case where all principals have $v_j = v$. We know that $U_i^l(I, k, k)$ converges to 0 as k grows large. Thus, there exists k^* such that

$$U_i^l(I, k^*, k^*) - U_i^l(O, k^*) \geq \frac{v-p}{\delta}$$

and

$$U_i^l(I, k, k) - U_i^l(O, k) < \frac{v-p}{\delta} \quad \text{for all } k > k^*.$$

If $U_i^l(I, k^* + 1, k^* + 1) - U_i^l(O, k^*) < \frac{v-p}{\delta}$, then all principals choosing a k^* -insider contract is a symmetric equilibrium. This corresponds to $\mu^* = 0$. On the other hand, if $U_i^l(I, k^* + 1, k^* + 1) - U_i^l(O, k^*) > \frac{v-p}{\delta}$, some principals will have an interest in deviating and choosing a $(k^* + 1)$ -insider contract. Note that there exists μ^* such that $U_i^l(I, k^* + 1, k^* + 1|\mu^*) - U_i^l(O|\mu^*) = \frac{v-p}{\delta}$, where $U_i^l(\cdot|\mu^*)$ denotes the expected utility of being in a given state when a fraction μ^* of the principals choose a $(k^* + 1)$ -contract and a fraction $1 - \mu^*$ choose a k^* -contract. Such a μ^* exists because $U(I, k^* + 1, k^* + 1|\mu = 0) - U_i^l(O|\mu = 0) > \frac{v-p}{\delta}$ and $U_i^l(I, k^* + 1, k^* + 1|\mu = 1) - U(O|\mu = 1) < \frac{v-p}{\delta}$, and $U_i^l(\cdot|\mu)$ is continuous in μ .

Now consider the case where $v_j \sim \text{unif}(v-r, v+r)$ for some $r > 0$, i.e. v_j is uniformly distributed over the interval $(v-r, v+r)$. Let a principal choose a $(k^* + 1)$ -insider contract if $v_j \leq \tilde{v}$ for some \tilde{v} and a k^* -insider contract if $v_j > \tilde{v}$. We will verify that this is an equilibrium in the limit as $r \rightarrow 0$.

Note that $\mathbb{P}(v_j \leq \tilde{v}) = \frac{\tilde{v} - (v-r)}{2r}$. Calling $\mu \in [0, 1]$ the fraction of principals choosing a $(k^* + 1)$ -contract, we have $\frac{\tilde{v} - (v-r)}{2r} = \mu$, which leads to $\tilde{v} = v + r(2\mu - 1)$.

Now let μ^* be such that

$$U_i^l(I, k^* + 1, k^* + 1|\mu^*) - U_i^l(O|\mu^*) = \frac{\tilde{v} - p}{\delta}$$

Then it follows that

$$U_i^l(I, k^* + 1, k^* + 1|\mu^*) - U_i^l(O|\mu^*) \geq \frac{v_l - p}{\delta} \quad \text{for all } v_l \leq \tilde{v}$$

and

$$U_i^l(I, k^* + 1, k^* + 1|\mu^*) - U_i^l(O|\mu^*) < \frac{v_l - p}{\delta} \quad \text{for all } v_l > \tilde{v}$$

and thus no principal with $v_j \leq \tilde{v}$ has an incentive to switch to a k^* -contract and likewise no principal with $v_j > \tilde{v}$ has an incentive to switch to a $(k^* + 1)$ -contract. Thus μ^* is an equilibrium fraction of principals choosing a $(k^* + 1)$ -contract.

Now note that

$$\lim_{r \rightarrow 0} \tilde{v} = \lim_{r \rightarrow 0} v + r(2\mu^* - 1) = v$$

and the limiting μ^* is such that

$$U_i^l(I, k^* + 1, k^* + 1|\mu^*) - U_i^l(O|\mu^*) = \frac{v - p}{\delta}.$$

Thus in the limit,

$$k^* = \max\{k' : U_i^l(I, k', k'|\mu^*) - U_i^l(O|\mu^*) \geq \frac{v_j - p}{\delta} \mid v_j > v\}$$

and

$$k^* + 1 = \max\{k' : U_i^l(I, k', k'|\mu^*) - U_i^l(O|\mu^*) \geq \frac{v_j - p}{\delta} \mid v_j \leq v\}$$

and these are the optimal contracts that can be chosen by such principals. ■