



WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 10-28/R
INDUCING AGENTS TO REPORT HIDDEN TRADES:
A THEORY OF AN INTERMEDIARY**

Yaron Leitner
Federal Reserve Bank of Philadelphia

Revision forthcoming in the
Review of Finance

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

Inducing Agents to Report Hidden Trades: A Theory of an Intermediary

Yaron Leitner*

Federal Reserve Bank of Philadelphia

Revision Forthcoming in the Review of Finance.

Abstract

When contracts are unobserved (and nonexclusive), agents can promise the same asset to multiple counterparties and subsequently default. I show that a central mechanism can extract all relevant information about contracts that agents enter by inducing them to report one another. The mechanism sets position limits and reveals the names of agents who hit the limits according to (voluntary) reports from their counterparties. This holds even if sending reports is costly and agents can collude. In some cases, an agent's position limit must be nonbinding in equilibrium. The mechanism has some features of a clearinghouse.

JEL Classification: D02, D82, D86, G20, G28

1. Introduction

Consider a bank that buys a credit default swap from AIG. The bank knows that AIG has enough capital to honor the swap agreement if the bank is its only client. However, AIG might sell credit default swaps to many banks and thus create a liability that it cannot honor should it need to make a payment. A central mechanism might help by monitoring AIG's and other agents' positions so that no agent enters into too many liabilities relative to his capital. Indeed, the recent financial crisis has led banks and regulators to work toward the establishment of a clearinghouse for credit default swaps.

If a central mechanism could observe all the contracts that agents enter, it could ensure that agents don't enter into too many contracts by setting position limits. However, observing every contract that an agent can enter may be too costly, as agents may attempt to hide

*Comments from Franklin Allen (the editor) and two anonymous referees significantly improved the exposition of this paper. An earlier draft based on the first chapter of my dissertation was circulated as "Nonexclusive contracts, collateralized trade, and a theory of an exchange." I am grateful to Michael Fishman (main advisor), Philip Bond, Arvind Krishnamurthy, Adriano Rampini, and Asher Wolinsky for helpful discussions and comments. I also thank Mitchell Berlin, Ronel Elul, Robert Hunt, Cyril Monnet, Leonard Nakamura, Christine Parlour, Brooke Stanley, and John Weinberg, as well as seminar participants at Northwestern University, the Philadelphia Fed, and numerous conferences. The views expressed here are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at <http://www.philadelphiafed.org/research-and-data/publications/working-papers/>.

their transactions or make them too complicated to understand. It turns out that it might be enough that the central mechanism observes only the contracts agents *choose* to report. In particular, I show that by allowing agents to report the identities of their counterparties and revealing the names of those who hit prespecified position limits, a central mechanism can achieve the same outcome that would be achieved if agents could not enter into secret contracts. This is true even if sending a report to the mechanism involves some small cost and even if agents can collude. I also show that in some cases it is possible to prevent collusion only if the mechanism allows each agent to enter more contracts than he actually enters in equilibrium; that is, the position limit must be nonbinding in equilibrium.

The central mechanism can extract all relevant information about the contracts that agents enter by inducing them to report one another. In particular, the optimal position limit is such that if an agent enters a contract and reports it to the mechanism, his counterparty has the incentive to pay what he promised, but if an agent enters a contract without reporting it, his counterparty is induced to plan a strategic default in which he enters as many contracts as he can with the intention to default on all of them. In equilibrium, every agent can enter contracts until he hits the position limit, but the position limit is not sufficiently high to make default profitable; thus, every agent ends up with only one contract, which is enough to achieve an efficient allocation. However, if an agent enters a contract without reporting it, he allows his counterparty to enter more contracts than the position limit, and this increases the counterparty's gain from a strategic default to the point at which default is profitable.

To induce reporting, the position limit cannot be too high, nor can it be too low. If the position limit is too high, the counterparty will always default whether or not the agent reports him; but then the two agents will not enter a contract to begin with. If the position limit is too low, the counterparty will never default; but then the two agents will not report to save the reporting fee. In some cases, "not too low" means that the position limit must be higher than the number of contracts an agent enters in equilibrium.

To implement a position limit that is nonbinding in equilibrium, the mechanism must not reveal the exact number of contracts that an agent has entered. It should only reveal whether an agent has reached the position limit. The idea that the mechanism should not reveal too much information is well known in multistage games with private information and hidden actions: Too much information makes it easier for agents to manipulate the mechanism.¹ In my paper, an agent who learns that his counterparty has already entered a contract believes that his counterparty plans a strategic default. Hence, the agent cannot precommit to enter a contract with such a counterparty. But then the effective position limit is only one.

If agents cannot send reports to a central mechanism, they must put up cash as collateral. Collateral acts like a position limit because an agent may not have enough cash to enter the number of contracts needed to make default profitable. However, using collateral has an opportunity cost, as agents forgo investing in their positive net present value (NPV) projects. Since agents can divert money from investment to consumption, the optimal mechanism with reports (i.e., the position limit mechanism) may also require collateral, but less than the amount needed when agents cannot send reports.

Empirical predictions. According to the model, the gain from allowing agents to send reports to a central mechanism increases when the fixed cost per trade falls and/or the

¹See Myerson (1986), Bester and Strausz (2000, 2001, 2007), and the discussion in Subsection 4.4.

probability of finding a trading counterparty rises – both are features of a more liquid market.² The model also provides closed-form solutions and some comparative statics for the optimal amount of collateral (with and without reporting), the amount of investment, and the optimal position limits.

While this paper does not attempt to model any particular intermediary, the optimal mechanism has some features of a clearinghouse. The clearinghouse may be part of a futures exchange or a stand-alone institution; it can clear exchange-traded contracts as well as over-the-counter products.³ Clearinghouses deploy a number of safeguards to protect their members and customers against the consequences of default by a clearinghouse participant. In addition to requiring collateral, the clearinghouse monitors and controls the positions of its members (at least daily) and the financial statements, internal controls, and other indicators of financial strength (periodically). Some clearinghouses (e.g., in Sydney and Hong Kong) also set capital-based position limits.⁴ These safeguards, which reduce the amount of collateral that clearinghouse members must post, are more effective when clearinghouse members do not enter contracts secretly.⁵ In practice, the incentive to default may depend on activities in more than one market. Indeed, clearinghouses have recently moved toward more central clearing.⁶

Bernanke (1990) distinguishes between two roles of a clearinghouse: reducing the transactions cost of consummating agreed-upon trades (analogous to a bank that clears checks), and standardizing contracts by setting terms and format and by guaranteeing performance to both sides of trade (analogous to an insurance company).⁷ The optimal mechanism in my paper has a more minimal role, but the results remain even if we add other roles, such as guaranteeing performance. In addition, the model does not rule out multiple intermediaries.

Although this is not a model of regulation, in one interpretation the mechanism can be interpreted as a regulator (e.g., a central bank). My theory suggests that, in some cases, to induce banks to report all their transactions voluntarily, the regulator may need to commit to keep these reports private.⁸ The theory also illustrates a connection between regulation and private-sector incentives to discipline. The regulator, who sets position limits, relies on firms in the private sector to discipline one another; that is, each firm makes sure that its trading partner reports the trade to the regulator. The theory implies that regulations

²This seems consistent with the observation that the London Clearing House started clearing over-the-counter interest rate swaps only after they became a standardized and liquid product. Central clearing for credit default swaps also became relevant after a tremendous growth in market size (and the bad consequences during the recent financial crisis).

³For example, the London Clearing House clears over-the-counter interest rate swaps without being involved in the matching and bargaining processes.

⁴Capital-based position limits, whose purpose is to make sure that members maintain positions within their financial capability, are different from speculative position limits. The latter are set by exchanges and regulators to prevent speculators from manipulating spot prices.

⁵Netting may also reduce collateral. However, some clearinghouses (e.g., the Hong Kong Futures Exchange Clearing Corporation) calculate margin on a gross basis rather than a net basis. Also, while marking to market reduces the risk of default, when a counterparty's position is closed, there is a risk of not finding a new counterparty and remaining unhedged.

⁶For example, in 2004, the Chicago Mercantile Exchange (CME) fully integrated the clearing of all trades of the Chicago Board of Trade in addition to those of the CME. The CME has also developed cross-margin arrangements with other clearinghouses so that margins can be calculated based on the total position.

⁷See also Telser and Higinbotham (1977) and Edwards (1983).

⁸While my model provides a novel rationale for regulatory secrecy, I do not present a full discussion of the costs and benefits of regulatory secrecy.

that are too stringent may be counterproductive because they undermine private-sector incentives for agents to discipline one another.

Related literature. The paper contributes to the literature on financial intermediation by illustrating a very minimal role of an intermediary. The intermediary in my paper provides a cost-effective way to monitor agents' positions, not only because it saves on the cost of duplicate monitoring, but also because it relies on voluntary reports. Existing literature has focused on problems that arise because of asymmetric information regarding cash flows and on the role of intermediaries in enhancing liquidity. In contrast, I start with markets that are already liquid and show how an intermediary can help.⁹ The main problem in my paper is that an agent's history of transactions is private information. Unlike Diamond (1984), I do not rely on diversification, and unlike in Townsend (1978), the intermediary arises when the fixed cost per trade is low rather than high.¹⁰

The idea that a mechanism can extract information by inducing agents to report one another is related to the literature on information extraction in teams (e.g., Ma, 1988). This literature shows that when a principal interacts with many agents, he can implement his desired outcome if agents observe each other's choice of effort.¹¹ An agent is induced to take the desired amount of effort for fear of being reported to the principal by a teammate; however, as pointed out by many (e.g., Itoh, 1993; Bisin and Guaitoli, 2010), these information extraction mechanisms may not work when agents can enter side contracts. In my paper a mechanism can extract all relevant information about agents' positions by inducing agents to report the contracts they enter, including their counterparties' identities. An agent is induced to report because otherwise his counterparty can enter more contracts than the position limit, which leads to a strategic default. Crucially, the mechanism in my paper is collusion proof, and agents cannot undo the mechanism by entering side agreements.

In a different framework, Bizer and DeMarzo (1992) and Parlour and Rajan (2001) study the effect of nonexclusivity on equilibrium interest rates and competition in credit markets.¹² In Bizer and DeMarzo, contracts entered in the past are observable and have priority. Thus, a reporting mechanism, as in my paper, cannot improve welfare.¹³ In Parlour and Rajan, intermediaries offer contracts simultaneously, and then a single borrower can accept any subset of these contracts. As in my paper, agents who strategically default do so on all the contracts they entered. In their model, this can rule out entry even though competing lenders make positive profits. In my paper, this helps to sustain an equilibrium in which agents do not enter contracts secretly.

Paper outline. In Section 2, I present the economic environment, and in Section 3, I

⁹Since my paper illustrates a negative aspect of liquidity, it relates to Myers and Rajan (1998). In their model, greater asset liquidity reduces a firm's capacity to raise external finance because it reduces the firm's ability to commit to a specific course of action.

¹⁰Madhavan (2000) summarizes the extensive literature on the effects of different trading mechanisms on liquidity provision. Gorton and Winton (2003) summarize the extensive literature on the role of banks. Many other papers focus on the role of an intermediary. For example, Brusco and Jackson (1999) show how a market maker can economize on the fixed costs of trading across periods; and Rubinstein and Wolinsky (1987) show how middlemen can reduce search costs.

¹¹See also Ma et al (1988), in which agents observe correlated state variables before they take actions.

¹²See also Kahn and Mookherjee (1998), who study insurance contracts; Bisin and Rampini (2006), who study bankruptcy; and Bisin and Guaitoli (2004), who show that intermediaries can make positive profits by offering contracts that are not traded in equilibrium.

¹³The problem in their paper is that additional contracts impose a negative externality on existing contracts because the agent's hidden effort affects his future income.

solve for the optimal contract when contracts are exclusive (second best). Section 4 solves for an optimal mechanism when contracts are nonexclusive and shows that the second best outcome can be achieved; this is the paper's main result. I also discuss the role of nonbinding position limits and show that, in general, the mechanism should not reveal the information it has. In section 5, I show that the ability to send reports to the central mechanism is crucial and I solve for the best outcome that can be achieved when agents cannot send reports (third best). Section 6 discusses some robustness issues, and Section 7 concludes. The appendix contains proofs and omitted details.

2. The Model

The model has a continuum of agents who enter bilateral contracts for mutual insurance purpose. Half of the agents are type 1 and half are type 2. Each contract is between a type-1 agent and a type-2 agent. Contracts are entered at date 0, and they specify payments at date 1 contingent on the realized state. An agent can default strategically, but if he defaults, he loses his future income. I first solve the case in which agents can precommit to enter exclusive contracts (second best) and show that to achieve an efficient allocation, it is sufficient that each agent enters only one contract. Then I solve for an optimal mechanism when contracts are assumed to be nonexclusive and agents cannot observe the set of contracts their counterparties have entered or will enter.

In more detail: There are three dates, $t = 0, 1, 2$, and one divisible good, called dollars, or simply cash. Uncertainty is modeled by assuming two states of nature, state 1 and state 2, one of which is realized at date 1. Agents are risk-neutral and obtain an expected utility of $E(c_0 + c_1 + c_2)$ from consuming c_0, c_1 , and c_2 dollars at dates 0, 1, and 2, respectively. Agents are protected by limited liability, so $c_t \geq 0$ at each date.

At date 0, each agent has one dollar and an investment opportunity (project) that requires his human capital. Each project lasts for two periods and yields nothing if transferred to another agent; thus, a bank cannot invest on behalf of agents, as in Diamond and Dybvig (1983). Projects' cash flows are in Figure 1. Start with the project of a type-1 agent. The agent invests $I_1 \in [0, 1]$ at date 0. At date 1, the project yields εI_1 in state 1 ($\varepsilon > 0$) but requires an additional investment εI_1 in state 2. The additional investment must be made in full for the project to continue to date 2 and is called a negative cash flow. If the project continues to maturity (because it had a positive cash flow at date 1 or it had a negative cash flow and the additional investment was made in full), the project yields RI_1 dollars at date 2. Similarly, the project of a type-2 agent yields εI_2 in state 2 but requires εI_2 in state 1. If the project continues to maturity, it yields RI_2 dollars at date 2. Note that there is no aggregate uncertainty at date 1: Half of the projects have positive cash flows and half have negative cash flows. Liquidation values are zero at every date, consistent with the assumption that projects require human capital.

The risk-free rate is normalized to be zero percent (i.e., there is a storage technology that gives one dollar at date $t + 1$ for every dollar invested at date t), and it is assumed that $R > 1 > \varepsilon$. The assumption $R > \varepsilon$ implies that it is efficient to make the additional investment at date 1 if cash is available, and $R > 1$ implies that in a world without frictions, each project has a positive NPV; the NPV of i 's project is $(R - 1)I_i$. The assumption $\varepsilon < 1$ ensures that entering bilateral contracts is preferred to autarky despite the moral hazard problems below (see Section 3).

Contracts can be contingent on the state that is realized at date 1. However, agents cannot commit to make payments:

Assumption 1. An agent cannot commit to pay out of the project’s final cash flows (RI_i).

Assumption 2. An agent cannot commit to pay out of the project’s interim cash flows (εI_i).

Assumption 1 implies that an agent with a negative cash flow cannot borrow at date 1 against the future cash flows from his project. This is why agents enter insurance contracts at date 0, such that an agent with a positive cash flow transfers cash to an agent with a negative cash flow. Assumption 2 implies that agents may default on these contracts. Both assumptions can be motivated by assuming that cash flows are not verifiable in court. The first assumption can also be motivated by assuming moral hazard as in Holmström and Tirole (1998), or by assuming that final cash flows are unobservable.

Enforcement technology. If an agent defaults (i.e., does not pay what he promised in full), his project is terminated. In this case the agent keeps current cash flows but loses future cash flows. It is optimal to terminate the project of a defaulting agent with probability one, even if it is possible to choose a probability less than one, and it is assumed that it is possible to commit to this closure policy. Allowing for additional penalties for default, such as spending time in prison, losing one’s reputation, or losing other sources of future income, does not alter the nature of the results. What’s important is that penalties impose a finite cost, rather than an infinite cost.

Collateral. Agents cannot post the projects’ assets as collateral, but they can post cash as collateral. Specifically, a pair of agents can open an escrow account; they can store cash through a third party who can commit not to divert it. Money placed in escrow is observable to both agents and can be contracted upon. However, it is unobservable to agents who are not part of the bilateral contract (or to the central planner). That is, agents can open secret escrow accounts.

The last assumption in the benchmark model is that at date 0 agents can divert cash invested in their projects.

Assumption 3. The amount that an agent invests in his project (I_i) and the amount that an agent consumes are private information.

Assumption 3 introduces the risk of strategic default via “asset substitution”: An agent can consume his initial endowment, instead of investing it, and subsequently default at date 1, as he has no cash flows to pay from. Think of it as an agent diverting cash from one project (the original project) to another project (“consumption”) that yields some unobservable cash flow at date 0 and nothing afterward. Assumption 3 is the reason agents may need to post collateral, even in the benchmark case in which contracts are exclusive (see Section 3).¹⁴ Note that while everyone can observe whether a project operates (it can be terminated upon default), the level of investment (I_i) is private information. In addition, a project can operate even if $I_i = 0$. For example, an agent can go to work and keep his business open but effectively do nothing (e.g., Bernie Madoff).

Trading game. Agents can enter contracts, as described below. The game captures the idea that a counterparty may have entered contracts in the past and may enter additional

¹⁴Assumption 3 is needed because without it one could infer how many contracts an agent has entered by observing I_i , the amount left for investment after posting collateral on all contracts. If posting collateral did not reduce the amount available for investment, Assumption 3 would not be necessary.

contracts in the future, with none of these contracts being observable. The game also captures the idea of a large market in which a deviation by one agent (or a finite number of agents) does not affect bargaining power and contract terms. I do not focus on search frictions.¹⁵

There are N trading rounds; all happen at date 0 before agents post collateral, invest in their projects, and/or consume (see Figure 2). In each trading round, a fraction $\frac{1}{N}$ of agents, chosen randomly, arrive to trade for the first time, with an equal mass of both types. Agents who are present in each trading round are pairwise matched according to their types, and each pair can enter a contract. Then each agent decides whether to enter additional contracts or stick with the contracts he has entered so far. An agent who wants to enter additional contracts stays for the next trading round to be matched with another counterparty. If there are no more trading rounds, the agent leaves the trading game and moves to the next stage, in which he posts collateral. An agent who does not want to enter additional contracts also leaves the trading game.¹⁶

The sequence of events for an individual agent is in Figure 3. For simplicity, I assume that the contract is set by a planner and that it is entered only if both agents agree to enter (they choose whether to enter simultaneously). The results remain even if a pair of agents can enter a contract that is different from the one suggested by the planner (Subsection 6.1), and even if we assume other bargaining processes (e.g., one agent makes an offer and the other agent can either accept or reject it).

The main assumption is as follows:

Assumption 4. An agent cannot observe contracts that other pairs of agents enter (either in the past or in the future).

Assumption 4 has a few interpretations: Agents can enter contracts secretly¹⁷; trading is too fast for agents to keep track of a counterparty's history of transactions; or existing contracts are observable but not understood. An example is the complex derivative positions and off-balance-sheet transactions made by many hedge funds. As noted earlier, an agent cannot observe the amount of collateral that his counterparty posts with other agents; an agent who enters nonexclusive contracts opens a different escrow account with each counterparty.

Finally, the assumption that agents have the same initial endowment is made for simplicity. The results extend to the case in which agents have different endowments, which are private information (see Subsection 6.2).

¹⁵The effects of search and bargaining frictions on valuation and bid-ask spreads in over-the-counter markets are analyzed in Duffie et al. (2005, 2007).

¹⁶In equilibrium, collateral requirements and/or position limits put an upper bound on the number of contracts that an agent can enter. Thus, the mass of agents present in each round is finite, even if a continuum of agents decides to deviate by staying for more than one round. In addition, each type has the same proportion. In the out-of-equilibrium event in which the mass of type-1 agents does not equal the mass of type-2 agents, some agents remain unmatched.

¹⁷For example, according to the *Wall Street Journal* (August 25, 2005), "(hedge) funds sometimes move out of trades — 'assign' them — without telling the bank that sold them the credit-derivative contract that their counterparty has changed." Another example is the Nigerian barge deal between Enron and Merrill Lynch in which Enron allegedly arranged for Merrill Lynch to serve as a temporary buyer (of the barges) so as to make Enron appear more profitable than it was. According to a release by the Department of Justice (October 15, 2003), "Enron promised in a secret oral 'handshake' side-deal that Merrill Lynch would receive a return on its investment plus an agreed-upon profit..."

3. The Benchmark of Exclusive Contracts (Second Best)

In this section, I analyze the benchmark case in which every agent can enter at most one contract (i.e., $N = 1$). I characterize optimal contracts as a solution to a planning problem in which the planner sets a contract and recommends to each agent how much to invest.

A *contract* is a pair $(k_i, x_i)_{i=1,2}$, where k_i is the amount of cash that an agent of type i posts as collateral, and x_i is the amount of cash that he transfers to the other agent at date 1; an agent with a positive cash flow transfers cash to an agent with a negative cash flow (i.e., an agent of type i promises to pay in state i). An agent can default only on the amount $x_i - k_i$ ($k_i \leq x_i$). Assuming that there are no cash transfers at date 2 and that all transfers at date 0 are in the form of collateral is without loss of generality.

The triple $(I_i, k_i, x_i)_{i=1,2}$, which includes the contract and the planner's recommended level of investment, is referred to as the *agreement* and is denoted by ψ . An agreement induces the following consumption stream. At date 0, an agent of type i consumes $1 - I_i - k_i$, which is his initial endowment minus the amounts he invests and posts as collateral (without loss of generality, agents do not store on their own). The agent's consumption at date 1 depends on the state. In state i , the agent consumes $k_i + \varepsilon I_i - x_i$, which is the amount left after paying what he promised using the collateral he posted and the project's cash flows. In the other state, denoted by $-i$, the agent consumes $k_i + x_{-i} - \varepsilon I_i$, which is the amount left after making the additional investment using the collateral he posted and the payment received from his counterparty. (We can assume, without loss of generality, that the agreement is such that each agent has enough cash to make the additional investment.) At date 2, the agent consumes RI_i . The agent's expected utility is

$$\begin{aligned} U_i(\psi) &\equiv 1 - I_i - k_i + \frac{1}{2}(k_i + \varepsilon I_i - x_i) + \frac{1}{2}(k_i + x_{-i} - \varepsilon I_i) + RI_i \\ &= 1 + (R - 1)I_i + \frac{1}{2}(x_{-i} - x_i). \end{aligned} \quad (1)$$

When an agreement is symmetric, I sometimes drop the index i . The agreement is *feasible* if: (i) $I_i \geq 0$; (ii) $x_i \geq k_i \geq 0$; and (iii) the amount that every agent consumes at each date and state is nonnegative. That is,

$$1 - I_i - k_i \geq 0, \quad \text{for } i \in \{1, 2\}, \quad (2)$$

$$k_i + \varepsilon I_i - x_i \geq 0, \quad \text{for } i \in \{1, 2\}, \quad (3)$$

$$k_i + x_{-i} - \varepsilon I_i \geq 0, \quad \text{for } i \in \{1, 2\}. \quad (4)$$

Since the two types of agents are identical ex-ante and have equal proportion, it is natural to assume that the planner's objective is to maximize the unweighted sum of agents' utilities, $U_1(\psi) + U_2(\psi)$. This is equivalent to maximizing $I_1 + I_2$. The solution to the unconstrained problem (*first best*) is $I_1 = I_2 = 1$, $k_1 = k_2 = 0$, and $x_1 = x_2 = \varepsilon$. In the first best, agents do not post collateral, and the utility for each agent is R .

In the second best, we need to ensure that (i) agents have the incentives to invest and make the transfers suggested by the planner (incentive compatibility); and (ii) each agent prefers the proposed agreement to autarky (participation).

Participation. Denote by U_A an agent's utility in autarky. The participation constraint is

$$U_i(\psi) \geq U_A, \quad \text{for } i \in \{1, 2\}. \quad (5)$$

In autarky, an agent can self-insure by investing I and storing $s = 1 - I$ so that $s = \varepsilon I$. In this case, the agent continues his project in both states and obtains $s + RI = \frac{R+\varepsilon}{1+\varepsilon}I$ ($I = \frac{1}{1+\varepsilon}$). Alternatively, the agent can invest $I = 1$ and store nothing. In this case, the agent cannot continue his project when he realizes a negative shock, and his expected utility is $\frac{R+\varepsilon}{2}$. Since $\varepsilon < 1$, self-insuring is preferred. Thus, $U_A = \frac{R+\varepsilon}{1+\varepsilon}$. Entering bilateral contracts is preferred to autarky because a pair of agents can allocate all the cash stored to the agent with the negative cash flow so that each agent can invest more and store less.¹⁸

Incentives to make payments. Suppose an agent has entered the contract (k, x) and invested I' . When $k + \varepsilon I' < x$, the agent does not have enough cash to deliver the full amount, and it is optimal for him to pay nothing because if he makes a partial payment he still loses his project. When $k + \varepsilon I' \geq x$, the agent can pay the full amount, and it is optimal for him to do so because otherwise he gains $x - k \leq \varepsilon I'$ but loses $RI' > \varepsilon I'$. Denote by d whether an agent delivers ($d = 1$) or not ($d = 0$). The optimal delivery rule is $d(I', k, x) = 1$, if $k + \varepsilon I' \geq x$, and $d(I', k, x) = 0$, otherwise.

Incentives to invest. Denote by $U_i(I'_i|\psi)$ the utility for an agent of type i , if he deviates from the agreement ψ by investing $I'_i \in [0, 1 - k_i]$ instead of I_i . Denote by $\beta_i(I'_i, \psi)$ whether the agent has enough cash to make the additional investment; that is, $\beta_i(I'_i, \psi) = 1$, if $k_i + x_{-i} \geq \varepsilon I'_i$, and $\beta_i(I'_i, \psi) = 0$, otherwise. Then

$$\begin{aligned} U_i(I'_i|\psi) &= 1 - k_i - I'_i & (6) \\ &+ \frac{1}{2}[\varepsilon I'_i - d(I'_i, k_i, x_i)(x_i - k_i) + d(I'_i, k_i, x_i)RI'_i] \\ &+ \frac{1}{2}[k_i + x_{-i} + \beta_i(I'_i, \psi)(R - \varepsilon)I'_i]. \end{aligned}$$

The first line in (6) is the amount the agent consumes at date 0, the second is the amount he consumes at dates 1 and 2 after a positive cash flow, and the third line is the amount consumed after a negative cash flow. Observe that $U_i(I_i|\psi) = U_i(\psi)$.

The incentive constraint is that for $i \in \{1, 2\}$,

$$U_i(I_i|\psi) \geq U_i(I'_i|\psi), \text{ for every } I'_i \in [0, 1 - k_i]. \quad (7)$$

The second-best problem is to find a feasible agreement that maximizes $I_1 + I_2$ subject to the participation constraint and the incentive constraint.

Equation (7) can be replaced with $U_i(I_i|\psi) \geq U_i(0|\psi)$. In other words, it is enough to focus on deviations in which an agent invests nothing in his project and then defaults when he needs to make a payment. Intuitively, an agent who plans to default is better off consuming his initial endowment, rather than investing it and losing it upon default.¹⁹

Hence, the incentive constraint reduces to

$$\frac{1}{2}(x_i - k_i) \leq (R - 1)I_i, \text{ for } i \in \{1, 2\}. \quad (8)$$

Intuitively, the expected gain from not delivering the promised amount (left-hand side) must be less than or equal to the expected loss from not investing in one's project (right-hand side).

¹⁸In other words, the symmetric agreement (I, x, k) that satisfies $\varepsilon I = 2k$, and $x = k = 1 - I$, is strictly preferred to autarky.

¹⁹Lemma 1 in the appendix contains a proof.

The problem reduces to finding a feasible agreement that maximizes $I_1 + I_2$ subject to Equations (5) and (8). This is a linear programming problem. When $R \geq 1 + \frac{1}{2}\varepsilon$, the first-best agreement satisfies all the constraints and is a unique solution. In this case, the incentive constraint is not binding. In contrast, when $R < 1 + \frac{1}{2}\varepsilon$, the incentive constraint binds and the optimal agreement is such that: agents do not consume at date 0; each agent has exactly what he needs to make the additional investment but not more; and each agent is indifferent between following the agreement and planning a strategic default. Collateral is needed to prevent a strategic default in which an agent consumes his initial endowment instead of investing it.

Proposition 1. (Second best) *If $R \geq 1 + \frac{1}{2}\varepsilon$, the second-best agreement equals the first best. Otherwise, the second-best agreement is given (uniquely) by $k_1 = k_2 = k$, $x_1 = x_2 = \varepsilon - (1 + \varepsilon)k$, and $I_1 = I_2 = 1 - k$, where $k = \frac{\varepsilon - 2(R-1)}{\varepsilon - 2(R-1) + 2}$.*

Denote the second-best agreement by $\psi_{sb} \equiv (I_{sb}, k_{sb}, x_{sb})$. Then $U_i(\psi_{sb}) = R - (R-1)k_{sb}$ (from Equation (1) and Proposition 1). The first term (R) is the agent's first-best utility, and the second term is the opportunity cost of collateral: By posting collateral, agents forgo investing in their positive NPV projects.

The optimal amount of collateral decreases in R but increases in ε . An increase in R reduces the gain from a strategic default because the agent has more to lose. In contrast, an increase in ε increases the gain because the liquidity need is higher and the agent defaults on a larger amount.

4. Optimal Contracts with Nonexclusivity

In this section — which contains the main results — I analyze the case in which agents cannot precommit to enter exclusive contracts. For ease of exposition, I focus on the limit case in which there is an infinite number of trading rounds ($N = \infty$). Thus, an agent assigns a probability of zero to the event that he or his counterparty will not be able to enter additional contracts should either of them decide to do so. The nature of the results remains even if N is finite.²⁰

4.1 The mechanism design problem

Using a mechanism design approach, I extend the trading game from Section 2 by allowing agents to communicate with a central planner, who designs the rules of communication to maximize the sum of agents' utilities.

Agents here have both private information (e.g., the identities of their counterparties) and private actions (e.g., whether to stay for additional trading rounds). By the revelation principle (Myerson, 1982, 1986), we can restrict attention to direct communication mechanisms that are incentive compatible. A direct mechanism means that in each stage of the game every agent reports his new private information to the planner, and in return, the planner recommends an action to him, using some prespecified *recommendation rule* that maps reports to recommendations. Incentive compatibility means that it is optimal for

²⁰In particular, we can assume that N is a random variable with a geometric distribution. More details are available upon request.

each agent to be truthful and obedient (follow the planner’s recommendation), given that all other agents are.

Hence, the sequence of events in each trading round is as follows:

1. Every agent who is present in the trading round reports the identity of his counterparty to the planner (an agent can lie).
2. The planner tells every agent whether to agree to enter a contract, and every agent chooses whether to follow this recommendation. A contract is entered only if both the agent and his counterparty agree to enter.
3. Every agent tells the planner whether he has entered a contract in the current round (an agent can lie).
4. The planner tells every agent whether to leave or stay for the next round, and every agent chooses whether to follow this recommendation.

The problem reduces to finding a contract and a recommendation rule that maximize the sum of agents’ utilities subject to incentive compatibility.

Our main result is that the direct mechanism above can implement the second best. The result holds under the standard assumptions of the revelation principle, but it continues to hold even if we violate two assumptions. In particular, I assume that sending a report to the planner involves some small cost and that agents can collude.

In my setting, it is natural to focus on collusion between a pair of agents who were matched, but the results extend to a group of more than two agents. Collusion is modeled by assuming that a side planner can recommend to a pair of agents what to do. The implementation is *collusion proof* if the side planner cannot come up with a recommendation that increases the sum of the two agents’ utilities, such that it is optimal for each of them to follow the recommendation if the other agent does and if all other agents are truthful and obedient to the central planner.²¹

For ease of exposition, I assume that the cost of sending a report to the planner approaches zero and exclude it from the expressions below. I also assume, for simplicity, that the planner can observe the identity of agents who send reports; this assumption can be relaxed without affecting the results. Finally, assume that an agent agrees to enter a contract only if he believes that there is a positive probability that his counterparty will also agree. This can be motivated by assuming there is some small cost involved in making an offer.²²

4.2 The main result

Suppose the planner sets the second-best contract. The planner wants to implement an outcome in which every agent enters exactly one contract and follows it. The utility for each agent is then $U(\psi_{sb})$.

A possible deviation is that an agent enters more than one contract, invests nothing in his project, and subsequently defaults on all contracts. The maximum number of contracts

²¹The side planner is a modeling device that captures what a pair of agents can achieve via direct communication in an environment of asymmetric information; see, for example, Laffont and Martimort (1997). The side planner knows that the two agents were matched with one another, but he cannot observe anything else about the history of the game.

²²The last assumption is used in the proof of Proposition 4.

that a deviating agent can enter depends on the recommendation rule used by the planner and on the equilibrium strategies of all other agents. Suppose that if all other agents are truthful and obedient, a deviating agent can enter L contracts (assume the deviating agent remains truthful). I refer to L as the *position limit*. The deviating agent's utility is then

$$\begin{aligned}\bar{U}(\psi_{sb}, L) &\equiv (1 - Lk_{sb}) + \frac{1}{2}L(k_{sb} + x_{sb}) + \frac{1}{2}(0) \\ &= 1 + \frac{1}{2}L(x_{sb} - k_{sb}).\end{aligned}\tag{9}$$

The first expression $(1 - Lk_{sb})$ is what the agent consumes at date 0 after posting collateral. The other two expressions represent the expected amount consumed at date 1. In one state, the agent receives back all the collateral he posted plus a payment from each of his counterparties. In the other state, the agent needs to deliver, but he defaults; thus, he loses his collateral and ends up with nothing (limited liability). Observe that $\bar{U}(\psi_{sb}, L)$ is increasing in L .

To prevent the deviation above, we must have:

$$\bar{U}(\psi_{sb}, L) \leq U(\psi_{sb}).\tag{10}$$

Denote $L^* = \lfloor \max(1, \frac{2(R-1)}{\varepsilon}) \rfloor$, where the function $\lfloor z \rfloor$ denotes the largest integer less than or equal to z . Equation (10) reduces to $L \leq L^*$.

Another deviation is that a pair of agents enters a contract secretly. Specifically, a side planner can recommend that the two agents enter a contract without reporting each other's identities and without reporting the fact that they entered a contract; he can further recommend that the two agents do not enter additional contracts afterward. If both agents follow the side planner's recommendations, each of them obtains the second-best utility without incurring the reporting cost.

For the implementation to be collusion proof, we need to make sure that if one agent follows the side planner's recommendation, it is optimal for the other agent to cheat. An agent can cheat by entering additional contracts and defaulting on all contracts. Since the planner does not observe that the two agents entered a contract, he continues to allow each of them to enter L additional contracts, according to the position limit. Hence, together with the contract that was entered secretly, each of the two agents can enter a total of $L + 1$ contracts. Thus, an agent will cheat if and only if

$$\bar{U}(\psi_{sb}, L + 1) > U(\psi_{sb}).\tag{11}$$

Equation (11) reduces to $L \geq L^*$.

Combining the two results above, we obtain that $L = L^*$.

Proposition 2. (Optimal position limits) *To implement the second best in a collusion-proof way, the planner must allow each agent to enter L^* contracts, i.e., he must set a position limit $L = L^*$.*

Intuitively, the position limit cannot be too high nor can it be too low. If the position limit is too high, agents are induced to enter too many contracts with the intention to default on all of them; if the position limit is too low, agents can collude by entering a contact without reporting it to the planner. Agents have the incentive to report one

another only if not reporting triggers a counterparty’s default. When a contract is not reported, a counterparty can enter more contracts than the position limit. However, if the position limit is too low, a counterparty will never find it optimal to do so, whether the contract is reported to the planner or not. In this case, a pair of agents can benefit by not reporting and saving the reporting cost.

Implementation. Denote by $n(i, \tau)$ the number of times that agent i (here i denotes an agent’s identity) reported entering a contract by the end of round τ , and denote by $n'(i, \tau)$ the number of agents who reported entering a contract with agent i (an agent reported entering a contract with agent i if he reported agent i ’s identity and then reported entering a contract).

Consider the following rule, which we call a *position limit rule*: If agent i reports the identity of agent j in round τ , the planner recommends that agent i agree to enter a contract if and only if $n'(j, \tau) < L^*$. The planner recommends that agent i leave if and only if $n(i, \tau) = n'(i, \tau) = 1$ or $n'(i, \tau) = L^*$ (i.e., if the agent entered a contract and both the agent and his counterparty reported it, or if the agent has hit the position limit).

One can show that the game induced by the recommendation rule above has a perfect Bayesian equilibrium in which every agent enters the second-best contract exactly once. In addition, the implementation is collusion proof.

Proposition 3. (Main Result) *A position limit rule with a position limit L^* can implement the second best. In addition, the implementation is collusion proof.*

Under a position limit rule, the planner sets a position limit L^* and counts the number of times that an agent has entered a contract, according to reports from his counterparties. The planner recommends that an agent enter a contract only if his counterparty has not hit the position limit. Thus, if everyone is truthful and obedient, a deviating agent can enter at most L^* contracts. Along the equilibrium path (when everyone is truthful), the planner recommends that an agent leave after he enters his first contract. From our previous analysis, we know that it is optimal to follow this recommendation and that a pair of agents cannot gain by entering a contract secretly. If an agent reports entering a contract, and his counterparty does not, the planner does not count the contract and recommends that the agent who reported should default. This rules out a unilateral deviation in which one agent reports and his counterparty does not. Lying to the planner has the same effect as not reporting and is therefore suboptimal.

Essentially, the planner can extract all the relevant information about contracts that agents enter by inducing them to “police” one another: If one agent does not report entering a contract (and/or the identity of his counterparty), the counterparty is induced to plan a strategic default. Note that the planner can extract information that is shared by two agents (e.g., the fact that they entered a contract), but he cannot extract information that is held by only one agent (e.g., the amount that an agent invests).

As I mentioned in the introduction, the idea that agents police one another relates to the literature on information extraction in teams. However, in contrast to this literature, the planner in my setting can extract all relevant information even when agents can collude. The reason is that in my setting, a pair of agents cannot extract surplus from other agents, because when an agent defaults, he defaults on all the contracts he entered. In contrast, in the problem of a principal who deals with multiple agents, a group of agents can extract surplus from the principal by putting low effort and reporting high effort. Hence, whenever

the principal offers contracts that induce high effort, agents have the incentive to undo what the principal does by entering a side contract.

The optimal position limit, L^* , increases in R but decreases in ε . Intuitively, when the gain from strategic default falls (rises), the optimal position limit is higher (lower). It can also be shown that the optimal position limit is lower when markets become more liquid, as defined in Section 5 below.

4.3 The role of nonbinding position limits

When $R > 1 + \varepsilon$, it follows that $L^* > 1$. Thus, the planner must allow each agent to enter more than one contract, even though in equilibrium every agent enters only one contract.

Corollary 1. If $R > 1 + \varepsilon$, the position limit in Propositions 2 and 3 must be nonbinding in equilibrium.

Position limits that are nonbinding in equilibrium are essential if agents can collude, but they are not essential if agents cannot collude. When agents cannot collude, the planner can implement the second best by allowing each agent to enter only one contract and punishing an agent who lied or did not send a report to the planner by allowing his counterparty — and only his counterparty — to enter L^* additional contracts so that the counterparty will have the incentive to default.²³

In contrast, when agents can collude, the planner must rely on nonbinding position limits to induce agents to punish one another. Since the planner cannot detect a joint deviation, he must give each agent enough latitude to cheat on his counterparty. Then if a pair of agents attempts to enter a contract without reporting it to the planner, it is optimal for each of them to enter additional contracts and default.

The planner's commitment to use the prespecified recommendation rule (which is one of the assumptions behind the revelation principle) is crucial. When $L^* > 1$, the planner must allow an agent to enter L^* contracts, although an agent will enter these contracts only if he plans a strategic default. Ex ante, the threat of default is optimal because it induces agents to reveal information to the planner. However, ex post, once an agent attempts to enter more than one contract, it is suboptimal to let him do so.

4.4 How much information should the planner reveal?

In the implementation above, an agent can infer from the planner's recommendations whether his counterparty has reached the position limit and whether his counterparty reported him. Instead of recommending an action, the planner can simply reveal this information. For example, at the beginning of each round, the planner can make a public announcement (available at no cost to agents in the given round) of the identities of agents who have reached the position limit, and after an agent reports entering a contract (and the identity of his counterparty), the planner can let him know confidentially whether his counterparty reported him.

²³ Alternatively, a pair of agents can include a clause that voids a contract if the planner does not certify that both agents reported entering it.

Is it possible to implement the second best if the planner reveals more information? Under some restrictions on out-of-equilibrium beliefs, the answer is no. In other words, the assumption that reports are confidential (which is one of the assumptions behind the revelation principle) is crucial, and the planner cannot be replaced by a “bulletin board.”

The restriction used in the proposition below is that an agent who learns that his counterparty was reported entering a contract believes that his counterparty indeed entered a contract. This holds if agents assign a zero probability to the event that other agents have lied to the planner; for example, if an agent who reports entering a contract sends a copy of the signed contract.²⁴

Proposition 4. *The planner must not reveal the exact number of times that a counterparty was reported entering a contract. The planner should only reveal whether the counterparty has reached the position limit or not.*

The logic behind Proposition 4 is as follows: Unlike the planner, who can precommit to use the prespecified recommendation rule, agents cannot precommit to take actions that are optimal ex ante but are suboptimal ex post. An agent who believes that his counterparty has already entered a contract, or more than one contract, will not agree to enter another contract with him, since he will expect that his counterparty will continue to enter contracts afterward and default on all contracts. However, to implement the second best, an agent must be able to find counterparties to enter contracts with until he hits the position limit.

The result relates to Bester and Strausz (2007), who study principal-agent problems when the principal cannot precommit to take prespecified actions. They show that, without loss of generality, the principal can restrict himself to communication devices such that the agent reports his type honestly to the device, and the device garbles this information when sending a message to the principal. The purpose of this noisy device is to “fine tune” the amount of information that the principal has.²⁵

More generally, we know from the revelation principle for multistage games with adverse selection and moral hazard (Myerson, 1986) that we can focus, without loss of generality, on communication mechanisms in which the planner tells agents what to do without giving them extra pieces of information. This is because more information makes it easier for agents to manipulate the planner by lying to him or by disobeying his recommendations. The same intuition applies in my setting. If an agent knows that his counterparty has already entered a contract (but has not hit the position limit), he will not follow the planner’s recommendation to enter an additional contract with him.

²⁴ Another assumption is that upon observing an out-of-equilibrium event, an agent updates his beliefs about his counterparty’s past actions but continues to believe that all other agents in the current trading round have just showed up for trade and will follow their equilibrium strategies (to enter one contract and leave). Thus, an agent assigns a zero probability to the event that in the next round he will be matched with an agent from a previous round. The assumption is analogous to the notion of passive beliefs (McAfee and Schwartz, 1994), which is common in the contract theory literature.

²⁵ Bester and Strausz (2000, 2001) also analyze a situation in which agents send messages directly to the principal without using a noisy device. They show that if there is only one agent, the principal can restrict attention to a direct mechanism in which (i) the agent’s message space is the set of his types, and (ii) it is an optimal strategy for the agent to report his true type with a positive probability. However, when the principal deals with more than two agents, the message space may need to include more messages than types. Intuitively, the principal can add noise – just like a noisy device does – by implementing an equilibrium in which an agent sometimes lies, or by including more messages than types.

4.5 The planner as a clearinghouse

The central planner can be interpreted as an intermediary who sets position limits and lets a pair of agents register their contract, as long as none of them have reached the limit. The closest real-world example is a clearinghouse, as discussed in the introduction.

The results (including nonbinding position limits) remain even if, in addition to the minimal role above, the intermediary becomes a central counterparty that guarantees payments. Since default never happens in equilibrium, the intermediary does not need to have any capital to make this guarantee credible. In the out-of-equilibrium event in which an agent enters more than one contract and defaults, the intermediary defaults as well. The intermediary can prevent this type of default by setting aside some capital, but this is not necessary in our model.²⁶

5. What if Agents Cannot Send Reports to a Central Planner?

In the analysis above, I showed that the second best can be achieved if agents can send reports to a central planner. Below I show that the ability to send reports is crucial: If agents cannot send reports to a central planner (e.g., if there is no clearinghouse), the second best cannot be achieved.

To see why, suppose agents are unable to send reports to the central planner, and suppose, by contradiction, that there is an equilibrium in which every agent enters the second-best contract exactly once and follows it.²⁷ The only belief consistent with the equilibrium path is that all the agents who are present in the current round have just showed up for trade. Given this belief, an agent expects that each of his counterparties will enter one contract and deliver on it. A necessary condition is that if all other agents follow their equilibrium strategies, an agent cannot gain by entering more than one contract. However, the second-best contract does not satisfy this condition, as follows: If $R \geq 1 + \frac{1}{2}\varepsilon$, the second-best contract does not require collateral and a deviating agent can enter $L^* + 1$ contracts, thereby obtaining more than what he obtains if he enters only one contract. If $R \leq 1 + \frac{1}{2}\varepsilon$, the second-best contract requires collateral. Since $k_{sb} < \frac{1}{2}$ (by simple algebra), a deviating agent can enter at least two contracts and obtain $\bar{U}(\psi, 2) > U(\psi)$. Thus, it is optimal to deviate.

Can agents benefit from bilateral trade? Yes. However, rather than relying on position limits, agents must rely on collateral to limit the number of contracts that a deviating agent can enter. If the collateral is k (per contract), a deviating agent can enter at most $\frac{1}{k}$ contracts, since he has only one dollar to begin with.

²⁶If the intermediary sets aside some capital, Proposition 4 no longer holds, since an agent is guaranteed to obtain what is promised to him, even if his counterparty defaults. However, Proposition 4 continues to hold if a large group of agents can collude. Suppose these agents learn that their counterparties have already entered contracts, and suppose these agents share this information among themselves. Then the agents may rationally expect that the capital set by the clearinghouse may not be enough to guarantee payments to all of them.

²⁷I restrict attention to symmetric perfect Bayesian equilibria, in which agents of the same type follow the same (pure) strategy. I also assume that strategies do not depend on the index of the trading round in which an agent is present; this can be motivated by assuming that an agent does not know in what round he showed up.

Formally, denote by $\bar{U}_i(\psi, n_i)$ the utility for an agent of type i who enters $n_i \leq \frac{1}{k_i}$ contracts, planning to default on all of them, if each of his counterparties delivers. Similar to Equation (9), we obtain that

$$\bar{U}_i(\psi, n_i) = 1 + \frac{1}{2}n_i(x_{-i} - k_i). \quad (12)$$

Thus, the equilibrium contract must satisfy $U_i(\psi) \geq \bar{U}_i(\psi, n_i)$, for $i \in \{1, 2\}$ and for every integer $n_i \leq \frac{1}{k_i}$. This reduces to the following incentive constraint:

$$\frac{1}{2}(x_i - k_i) + \frac{1}{2}(n_i - 1)(x_{-i} - k_i) \leq (R - 1)I_i, \quad (13)$$

for $i \in \{1, 2\}$ and for every integer $n_i \leq \frac{1}{k_i}$. In the special case $n_i = 1$, (13) reduces to (8). The extra term when $n_i > 1$ is the expected net payoff from entering additional contracts and not delivering on them: In one state the agent collects payments from each of his additional counterparties; in the other state the agent loses his collateral.

The optimal agreement (and related contract) when agents cannot communicate with a central planner is referred to as *third best*. The third-best problem is to find a feasible agreement that maximizes $I_1 + I_2$, subject to the participation constraint, (5), and the incentive constraint, (13). To ensure that a solution exists, I drop the restriction that n_i be an integer; the appendix contains a micro foundation for this.²⁸ The optimal agreement is obtained by setting $n_i = \frac{1}{k_i}$ and solving (2), (4), and (13) with equalities. The next proposition characterizes the unique solution.

Proposition 5. (Third best) *The optimal agreement when agents cannot send reports to a central planner involves collateral, and more than in the second best. The agreement is given (uniquely) by $k_1 = k_2 = k^*$, $x_1 = x_2 = \varepsilon - (1 + \varepsilon)k^*$, and $I_1 = I_2 = 1 - k^*$, where $k^* \equiv \frac{1}{4r} \left(b - \sqrt{b^2 - 8r\varepsilon} \right)$, $r = R - 1$, and $b = 2 + \varepsilon + 2r$.*

Under the third-best agreement, an agent promises more than the amount of cash that he posts as collateral, i.e., $x > k$. Specifically, it follows from Equation (13) that $x = k + \frac{2(R-1)I}{n}$, where $n = \frac{1}{k}$. The first term (k) captures the idea that an agent cannot default on the amount of cash that he posts as collateral. The second term captures the idea that requiring collateral limits the number of contracts that an agent can enter, and this makes the threat of losing future cash flows valuable in backing promises.

As in the second best, the optimal amount of collateral decreases in R but increases in ε . The proof in the appendix applies to a more general case in which entering a contract involves some fixed cost δ , which can represent the time and effort involved in entering a contract. I show that when δ falls, the third-best agreement requires more collateral. A lower δ represents a more liquid market. Liquid markets present a problem in this model because they create more opportunities for strategic default.²⁹ The following conclusion then follows:

²⁸If n_i is restricted to be an integer, a solution may not exist because the set of feasible agreements that satisfy Equation (13) may be open (since n_i is not a continuous function of k_i).

²⁹In an appendix available upon request, I focus on another feature of a liquid market: the probability of not finding a counterparty. The effects of reducing this probability are similar to the effects of reducing the fixed cost per trade.

Corollary 2. The gain from allowing agents to send reports to a central planner increases when the fixed cost per trade (δ) falls.

6. Robustness

6.1 What if agents can choose the contract terms?

In the analysis above I assumed that the contract is set by the central planner. The main result remains even if a pair of agents can choose to enter a contract that is different from the one suggested by the central planner. To see why, suppose a side planner suggests that a pair of agents enter $\psi \neq \psi_{sb}$ without reporting. The two agents save the reporting cost, but to prevent default they must post more collateral than in the second best. If the cost of sending a report to the central planner is sufficiently low, the extra cost of collateral outweighs the benefits of not reporting.

Proposition 5 (third best) also remains. In fact, the third-best agreement is the *only* feasible agreement satisfying Equations (5) and (13) (participation and incentive) that is both symmetric and renegotiation proof.³⁰

6.2 Unobservable endowments

In the analysis above I assumed that endowments are observable. Under some conditions, the results extend to the case in which endowments differ across agents and are private information. In particular, assume that an agent who claims having an endowment \hat{e} must prove it by showing it to the planner. That is, an agent who has e can report $\hat{e} \in [0, e]$. Then agents will report their endowments truthfully and the second best will be achieved.

The extra step in which agents reports their endowments to the planner will be before trading begins. If an agent reports \hat{e} , his position limit will be $\hat{e}L^*$, and the planner will recommend that he enter the scaled contract $(\hat{e}k_{sb}, \hat{e}x_{sb})$ and invest $\hat{e}I_{sb}$. To abstract from search and matching issues, assume that there are many trading venues and that in each trading venue the planner recommends that agents enter a different contract. Each agent will select a trading venue according to the contract that he wants to enter in the given round; the appendix contains more details.

Essentially, when endowments are observable, the planner's role is very simple because he only needs to rely on voluntary reports regarding the contracts that agents enter. When endowments are private information, the planner can still rely on these voluntary reports to learn whether a pair of agents has entered into a contractual relationship; however, to set position limits appropriately, the planner must verify that an agent who claims to have a certain endowment actually has it. The need to verify endowments also arises when agents set optimal collateral levels in the decentralized environment

What if we add deep-pocket lenders to our model? Nonexclusivity rules out a situation in which agents borrow money secretly to make their endowments look better than what they truly are. Lenders will report all the loans they make to the planner because otherwise

³⁰A proof is available upon request. (An agreement ψ is renegotiation proof if a side planner cannot improve the agents' utilities by giving them the opportunity to enter the agreement $\psi' \neq \psi$ instead of ψ , such that it is optimal for each agent to follow ψ' if the other agent does, and if all other agents enter and follow ψ .)

an agent could borrow from multiple lenders, divert the borrowed funds for consumption, and subsequently default on all loans.

6.3 Multiple intermediaries

The analysis above shows that one intermediary can implement the second best, but Proposition 2 does not rule out multiple intermediaries. For example, a position limit of four can be implemented by four intermediaries, each setting a position limit of one.³¹ To see that, adjust the trading game by assuming more than one location, such that each location has its own intermediary (planner) who can observe only the information that is reported to him. Each agent shows up for trade in a randomly chosen location. Initially, an agent must trade in the location where he shows up, but if an agent decides to stay for additional rounds, he can switch back and forth among the different locations. Pairwise matching in each location is as in the original game, and an agent can communicate only with the intermediary in the location where he is. The game above has an equilibrium in which every agent enters one contract and reports his information truthfully to the intermediary in his original location.

7. Conclusion

The paper constructs a mechanism that induces agents to voluntarily reveal to it all the contracts they enter. The mechanism allows each agent to report every contract he enters, and it reveals the names of agents who have reached some prespecified position limit. The main result is that such a mechanism can implement the same outcome that could be implemented if agents could not enter contracts secretly. This is true even if reporting involves some small cost, and even if agents can collude. If agents' endowments are private information, the mechanism must also verify that an agent who claims to have a certain endowment, actually has it. In general, the mechanism should not reveal the information it has and so it cannot be replaced by a bulletin board. The paper also provides a closed-form solution for the best outcome when agents cannot send reports to a central mechanism and shows that the gain from allowing agents to send reports increases when markets become more liquid.

Appendix

Unobservable Endowments. Assume that instead of one group of agents, there are an infinite number of groups corresponding to the interval $(0, 1]$. Agents in group $e \in (0, 1]$ have an initial endowment e . There are also an infinite number of trading venues, corresponding to the interval $(0, 1]$. In trading venue $\mu \in (0, 1]$, the planner recommends entering the scaled contract $(\mu k_i, \mu x_i)_{i=1,2}$ and investing $(\mu I_i)_{i=1,2}$. In each trading round, the same mass of agents from each group and type shows up for trade for the first time. An agent can switch back and forth among different trading venues.

The optimal mechanism includes an initial stage in which agents report their endowments to the planner. An agent with an initial endowment e can report $\hat{e} \in [0, e]$. In

³¹More generally, L^* can be implemented by n intermediaries, such that intermediary j sets a position limit $l_j \geq 1$, and $\sum_{j=1}^n l_j = L^*$.

addition, an agent also announces the trading venue in which he chooses to trade. The planner recommends entering a contract, and he later counts a contract as being entered, only if the counterparty's endowment is consistent with the trading venue.

Following a similar logic as in the single-endowment case, it is possible to show that the second best can be implemented if the planner sets a position limit of $\widehat{e}L^*$ for an agent who reports \widehat{e} . Given this position limit, an agent cannot gain by saying that he has less than he truly has, and by assumption, he cannot say that he has more.

Dropping the Restriction That n_i in Equation (13) Be an Integer. Consider the extension for unobservable endowments, and assume that an agent can trade in venue μ only if $e \geq \mu$ (i.e., the agent must reveal his endowment or a portion of it to his counterparty). When agents cannot send reports to the planner, an agent can deviate by entering multiple contracts in the venue that corresponds to his initial endowments, but he can also trade in venues that correspond to less than his initial endowment. Since choosing a venue $\mu < e$ is like entering a fraction of a contract, the number of contracts that a deviating agent can enter is no longer restricted to be an integer.

Lemma 1. *Equation (7) can be replaced with $U_i(I_i|\psi) \geq U_i(0|\psi)$.*

Proof. First, observe that Equation (2) must be binding; that is, $I_i = 1 - k_i$. Otherwise, we can increase the value of the objective function without violating the constraints by increasing I_i and k_i by Δ and $\varepsilon\Delta$, respectively, where Δ is small enough.

An agent can deviate by choosing $I'_i < I_i$. He can pay what he promised if and only if $I'_i \geq \frac{x_i - k_i}{\varepsilon}$. The result follows because $U(I'_i|\psi)$ is linear on $[0, \frac{x_i - k_i}{\varepsilon}]$, has a positive jump at $I'_i = \frac{x_i - k_i}{\varepsilon}$, and is increasing on $[\frac{x_i - k_i}{\varepsilon}, I_i]$ (since $R > 1$). ■

Proof of Proposition 1. The proof applies to a more general case in which there is a penalty $M \geq 0$ upon default (measured in terms of utility), so $U_i(0|\psi)$ is replaced with $U_i(0|\psi) - \frac{1}{2}M$, and the incentive constraint becomes $\frac{1}{2}(x_i - k_i) - \frac{1}{2}M \leq (R - 1)I_i$. When $R + \frac{1}{2}M \geq 1 + \frac{1}{2}\varepsilon$, the second best equals the first best. Otherwise, Equation (3) follows from the incentive constraint, and the solution is obtained by solving the incentive constraint and Equations (2), (4) with equalities. From (2), $I_i = 1 - k_i$. Substituting this in (4) and in the incentive constraint, and rearranging terms, we obtain: $x_{-i} = \varepsilon - (1 + \varepsilon)k_i$ and $x_i = M + 2(R - 1) + (3 - 2R)k_i$, for $i \in \{1, 2\}$. Thus, $x_2 - x_1 = (1 + \varepsilon)(k_2 - k_1)$ and $x_2 - x_1 = (3 - 2R)(k_2 - k_1)$. Since $2(1 - R) < 0 < \varepsilon$, it follows that $1 + \varepsilon \neq 3 - 2R$. Thus, $k_1 = k_2$. Denoting $k_i = k$, we obtain $x_1 = x_2 = \varepsilon - (1 + \varepsilon)k = M + 2(R - 1) + (3 - 2R)k$. Solving for k , we obtain, $k = \frac{\varepsilon - 2(R - 1) - M}{\varepsilon - 2(R - 1) + 2}$. Observe that $sign(\frac{\partial k}{\partial R}) = -2(2 + M) < 0$, and $sign(\frac{\partial k}{\partial \varepsilon}) = 2 + M > 0$. Thus, the optimal amount of collateral decreases in R but increases in ε . ■

Proof of Proposition 2. As explained in the text, to prevent a deviation in which an agent enters more than one contract and defaults, we must have $L \leq L^*$. To prevent collusion, we must have $L \geq L^*$. Thus, $L = L^*$ is a necessary condition. ■

Proof of Proposition 3. If everyone is truthful and obedient, the outcome is that every agent enters one contract (the second best) and leaves. It is optimal to follow the planner's recommendations as follows: If the planner recommends that an agent agree to

enter a contract, the agent believes that his counterparty has just showed up for trade and will follow the contract. If the planner recommends that an agent should not agree to enter a contract, the agent believes that his counterparty has already entered L^* contracts and will default if he enters an additional contract. If an agent has entered one contract, and the planner recommends that he leave, the agent believes that his counterparty reported him, and since $\bar{U}(\psi_{sb}, L^*) \leq U(\psi_{sb})$, it is optimal to follow the recommendation. If the planner recommends that the agent stay, the agent believes that his (first) counterparty did not report, and since $\bar{U}(\psi_{sb}, L^* + 1) > U(\psi_{sb})$, it is optimal to stay for L^* additional rounds (and default). It is optimal to report the counterparty's identity (truthfully) and the fact that an agent has entered a contract because otherwise the planner acts as if the counterparty did not enter a contract and recommends that he stay for additional rounds and default. Collusion proofness follows from Equation (11), as explained in the text. ■

Proof of Proposition 4. The relevant case is $L^* \geq 2$. Suppose an agent learns that his counterparty has entered $n \geq 1$ contracts. Since the counterparty is expected to use his equilibrium strategy from now on, he will either stick with the contracts he has already entered and not agree to enter a contract in the current round, or else he will agree to enter a contract in the current round and continue to enter as many contracts as he can, planning to default on all contracts. In the first case, it is suboptimal to agree to enter a contract, since there is zero probability that the counterparty will also agree. In the second case, it is suboptimal to enter a contract, since the counterparty will default. But then the effective position limit is less than L^* , and the implementation is not collusion proof. ■

Proof of Proposition 5. The proof applies to a more general case in which $U_i(\psi) = 1 + (R - 1)I_i + \frac{1}{2}(x_{-i} - x_i) - \delta$ and $\bar{U}_i(\psi, n_i) = 1 + \frac{1}{2}n_i(x_{-i} - k_i) - \frac{1}{2}M - n_i\delta$. The parameter $\delta \geq 0$ represents a fixed cost per trade, and $M \geq 0$ represents a penalty upon default, both measured in terms of utility. Assume that δ is sufficiently small so that the participation constraint is satisfied, and the third best does not equal the second best. The incentive constraint becomes $\frac{1}{2}(x_i - x_{-i}) + \frac{1}{2}n_i(x_{-i} - k_i) - (n_i - 1)\delta \leq (R - 1)I_i + \frac{1}{2}M$. As in Lemma 1, $I_i = 1 - k_i$. Substituting this and $n_i = \frac{1}{k_i}$ in the incentive constraint, summing over $i = 1, 2$, and rearranging terms, we obtain $g_1(\psi) \equiv \frac{1}{2} \sum_{i=1}^2 (\frac{x_{-i}}{k_i} - 1) - (R - 1) \sum_{i=1}^2 (1 - k_i) - M - \delta \sum_{i=1}^2 (\frac{1}{k_i} - 1) \leq 0$. Similarly, from Equation (4) we obtain $g_2(\psi) \equiv \varepsilon \sum_{i=1}^2 (1 - k_i) - \sum_{i=1}^2 (k_i + x_i) \leq 0$.

Denote $\varepsilon' \equiv \varepsilon - 2\delta$; $r \equiv R - 1$; $b \equiv 2 + \varepsilon' + 2r + M$; and $k^* = \frac{1}{4r}(b - \sqrt{b^2 - 8r\varepsilon'})$. (The statement of the proposition in the text is for the special case, $\delta = M = 0$.) To prove the proposition, it is enough to show that the agreement $(k, x, I) = (k^*, \varepsilon - (1 + \varepsilon)k^*, 1 - k^*)$ is feasible, satisfies the participation and the incentive constraint, and is a unique solution to $\min(k_1 + k_2)$ subject to $g_i(\psi) \leq 0$ for every $i \in \{1, 2\}$. I prove the last part below. The rest follows easily.

First, $g_i(\psi) \leq 0$ is binding ($i = 1, 2$), as follows: Since the second best is not achieved, $k_1 > 0$ and/or $k_2 > 0$. Without loss of generality, $k_1 > 0$. If $g_1(\psi) < 0$ (by contradiction), we can increase the value of the objective function without violating the constraints by replacing k_1 and x_1 with $k_1 - \Delta$ and $x_1 + (1 + \varepsilon)\Delta$, where Δ is sufficiently small. If $g_2(\psi) < 0$, we can increase the value of the objective function without violating the constraints by replacing k_1 and x_2 with $(1 - \Delta)k_1$ and $(1 - \Delta)x_2$. Hence, $g_1(\psi) = g_2(\psi) = 0$.

Denote the Lagrange multiplier of $g_i(\psi)$ by λ_i , and let $L(\psi) = k_1 + k_2 + \sum_{j=1}^2 \lambda_j g_j(\psi)$. An (optimal) solution satisfies $\frac{\delta L}{\delta x_i} = \frac{\lambda_1}{2k_{-i}} - \lambda_2 = 0$ and $\frac{\delta L}{\delta k_i} = 1 - \frac{x_{-i}}{2k_i^2} \lambda_1 + (R-1)\lambda_1 + \delta \frac{1}{k_i^2} \lambda_1 - \varepsilon \lambda_2 - \lambda_2 = 0$ ($i = 1, 2$). Thus, $k_1 = k_2 \equiv k$ and $x_1 = x_2 \equiv x$. From $g_2(\psi) = 0$, we obtain $x = \varepsilon - (1 + \varepsilon)k$. Thus, $g_1(\psi) = 0$ reduces to $\frac{\varepsilon - (1 + \varepsilon)k}{k} - 1 = 2r(1 - k) + M + 2\delta(\frac{1}{k} - 1)$. This further reduces to $2rk^2 - bk + \varepsilon' = 0$, which has two roots, $k = \frac{1}{4r}(b \pm \sqrt{b^2 - 8r\varepsilon'})$. The smallest root (k^*) is the unique solution, since it gives a lower value for the objective function than the other root.

It is easy to verify that $k^* \in (0, 1)$. To see that $k^* < 1$, note that $k^* < 1$ is equivalent to $b - 4r < \sqrt{b^2 - 8r\varepsilon'}$. If $b < 4r$, the result follows since the left-hand side is negative and the right-hand side is positive. Otherwise, we need to show that $(b - 4r)^2 < b^2 - 8r\varepsilon'$, which is equivalent to $b > 2r + \varepsilon'$. The last inequality follows from the definition of b . To see that $k^* > 0$, note that $\sqrt{b^2 - 8r\varepsilon'} < \sqrt{b^2} = b$.

To do comparative statics, define $H = 2rk^2 - bk + \varepsilon'$. Observe that $\frac{\partial H}{\partial k}|_{k=k^*} = 4rk^* - b = -\sqrt{b^2 - 8r\varepsilon'} < 0$. Thus, $\text{sign}(\frac{\partial k^*}{\partial \varepsilon}) = \text{sign}(\frac{\partial H}{\partial \varepsilon}|_{k=k^*}) = \text{sign}(1 - k^*) > 0$; $\text{sign}(\frac{\partial k^*}{\partial R}) = \text{sign}(\frac{\partial H}{\partial R}|_{k=k^*}) = \text{sign}(2k(k-1)) < 0$; and $\text{sign}(\frac{\partial k^*}{\partial \delta}) = \text{sign}(\frac{\partial H}{\partial \delta}|_{k=k^*}) = \text{sign}(2(k-1)) < 0$. ■

Proof of Corollary 2. Denote the third-best agreement by $\psi_{sb} = (k_{sb}, x_{sb}, I_{sb})$. The gain from allowing agents to send reports to a central planner is $U(\psi_{sb}) - U(\psi_{tb}) = (R - 1)(k_{tb} - k_{sb})$. Since k_{sb} does not depend on δ , but k_{tb} does ($\frac{\partial k_{tb}}{\partial \delta} < 0$), the gain increases when δ falls. ■

References

- [1] Bester, H. and Strausz, R. (2000) Imperfect commitment and the revelation principle: the multi-agent case, *Economics Letters* 69, 165-171.
- [2] Bester, H. and Strausz, R. (2001) Contracting with imperfect commitment and the revelation principle: the single agent case, *Econometrica* 69, 1077-1098.
- [3] Bester, H. and Strausz, R. (2007) Contracting with imperfect commitment and noisy communication, *Journal of Economic Theory* 136, 236-259.
- [4] Bernanke, B. S. (1990) Clearing and settlement during the crash, *Review of Financial Studies* 3, 133-151
- [5] Bisin, A. and Guaitoli, D. (2004) Moral hazard and nonexclusive contracts, *RAND Journal of Economics* 35, 306-328.
- [6] Bisin, A. and Guaitoli, D. (2010) Information extraction and norms of mutual protection, *Journal of Economic Behavior and Organization*, forthcoming.
- [7] Bisin, A. and Rampini, A. (2006) Exclusive contracts and the institution of bankruptcy, *Economic Theory* 27, 277-304.
- [8] Bizer, D. S. and DeMarzo, P. M. (1992) Sequential banking, *Journal of Political Economy* 100, 41-61.

- [9] Brusco, S. and Jackson, M. O. (1999) The optimal design of a market, *Journal of Economic Theory* 88, 1-39.
- [10] Diamond, D. W. (1984) Financial intermediation and delegated monitoring, *Review of Economic Studies* 51, 393-414.
- [11] Diamond, D. W. and Dybvig, P. H. (1983) Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401-419.
- [12] Duffie, D., Gârleanu, N. and Pedersen, L. H. (2005) Over-the-counter markets, *Econometrica* 73, 1815-1847.
- [13] Duffie, D., Gârleanu, N. and Pedersen, L. H. (2007) Valuation in over-the-counter markets, *Review of Financial Studies* 20, 1865-1900.
- [14] Edwards, F. (1983) The clearing association in futures markets: guarantor and regulator, *Journal of Futures Markets* 3, 369-392.
- [15] Gorton, G. and Winton, A. (2003). Financial intermediation, in: G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbook of the Economics of Finance, Corporate Finance*, Vol. 1A. North-Holland, Amsterdam, pp. 431-552.
- [16] Holmström, B. and Tirole, J. (1998) Private and public supply of liquidity, *Journal of Political Economy* 106, 1-40.
- [17] Itoh, H. (1993) Coalitions, incentives, and risk sharing, *Journal of Economic Theory* 60, 410-427.
- [18] Kahn, C. M. and Mookherjee, D. (1998) Competition and incentives with nonexclusive contracts, *RAND Journal of Economics* 29, 443-465.
- [19] Laffont, J. and Martimort, D. (1997) Collusion under asymmetric information, *Econometrica* 65, 875-911.
- [20] Ma, C. (1988) Unique implementation of incentive contracts with many agents, *Review of Economic Studies* 55, 555-572.
- [21] Ma, C., Moore, J. and Turnbull, S. (1988) Stopping agents from cheating, *Journal of Economic Theory* 46, 355-372.
- [22] Madhavan, A. (2000) Market microstructure: a survey, *Journal of Financial Markets* 3, 205-258.
- [23] McAfee, R. P. and Schwartz, M. (1994) Opportunism in multilateral vertical contracting: nondiscrimination, exclusivity, and uniformity, *American Economic Review* 84, 210-230.
- [24] Myers, S. C. and Rajan, R. (1998) The paradox of liquidity, *Quarterly Journal of Economics* 113, 733-771.
- [25] Myerson, R. B. (1982) Optimal coordination mechanisms in generalized principal-agent problems, *Journal of Mathematical Economics* 10, 67-81.

- [26] Myerson, R. B. (1986) Multistage games with communication, *Econometrica* 54, 323-358.
- [27] Parlour, C. A. and Rajan, U. (2001) Competition in loan contracts, *American Economic Review* 91, 1311-1328
- [28] Rubinstein, A. and Wolinsky, A. (1987) Middlemen, *Quarterly Journal of Economics* 102, 581-593.
- [29] Telser, L. G. and Higinbotham, H. N. (1977) Organized futures markets: costs and benefits, *Journal of Political Economy* 85, 969-1000.
- [30] Townsend, R. M. (1978) Intermediation with costly bilateral exchange, *Review of Economic Studies*, 45 417-425.

Figures




Figure 1. Project's cash flows for an agent of type i (i=1,2) if the agent makes the additional investment at date 1.

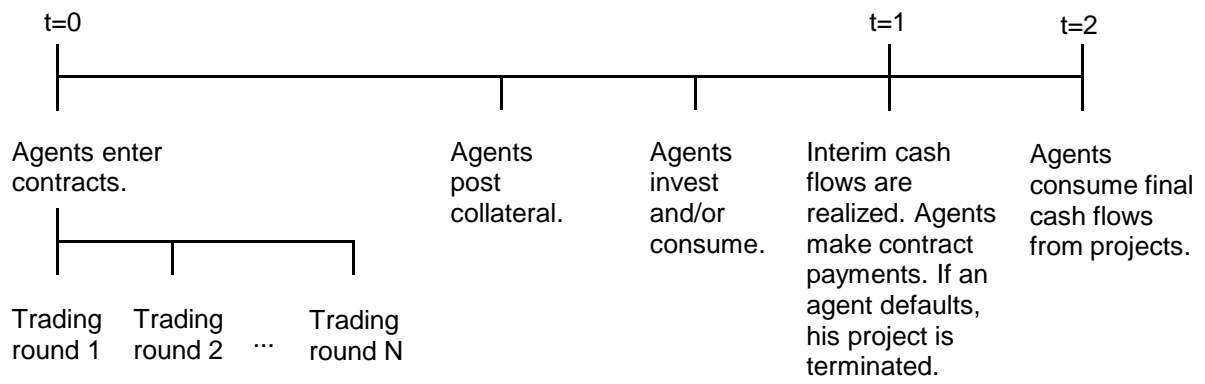


Figure 2. Time Line

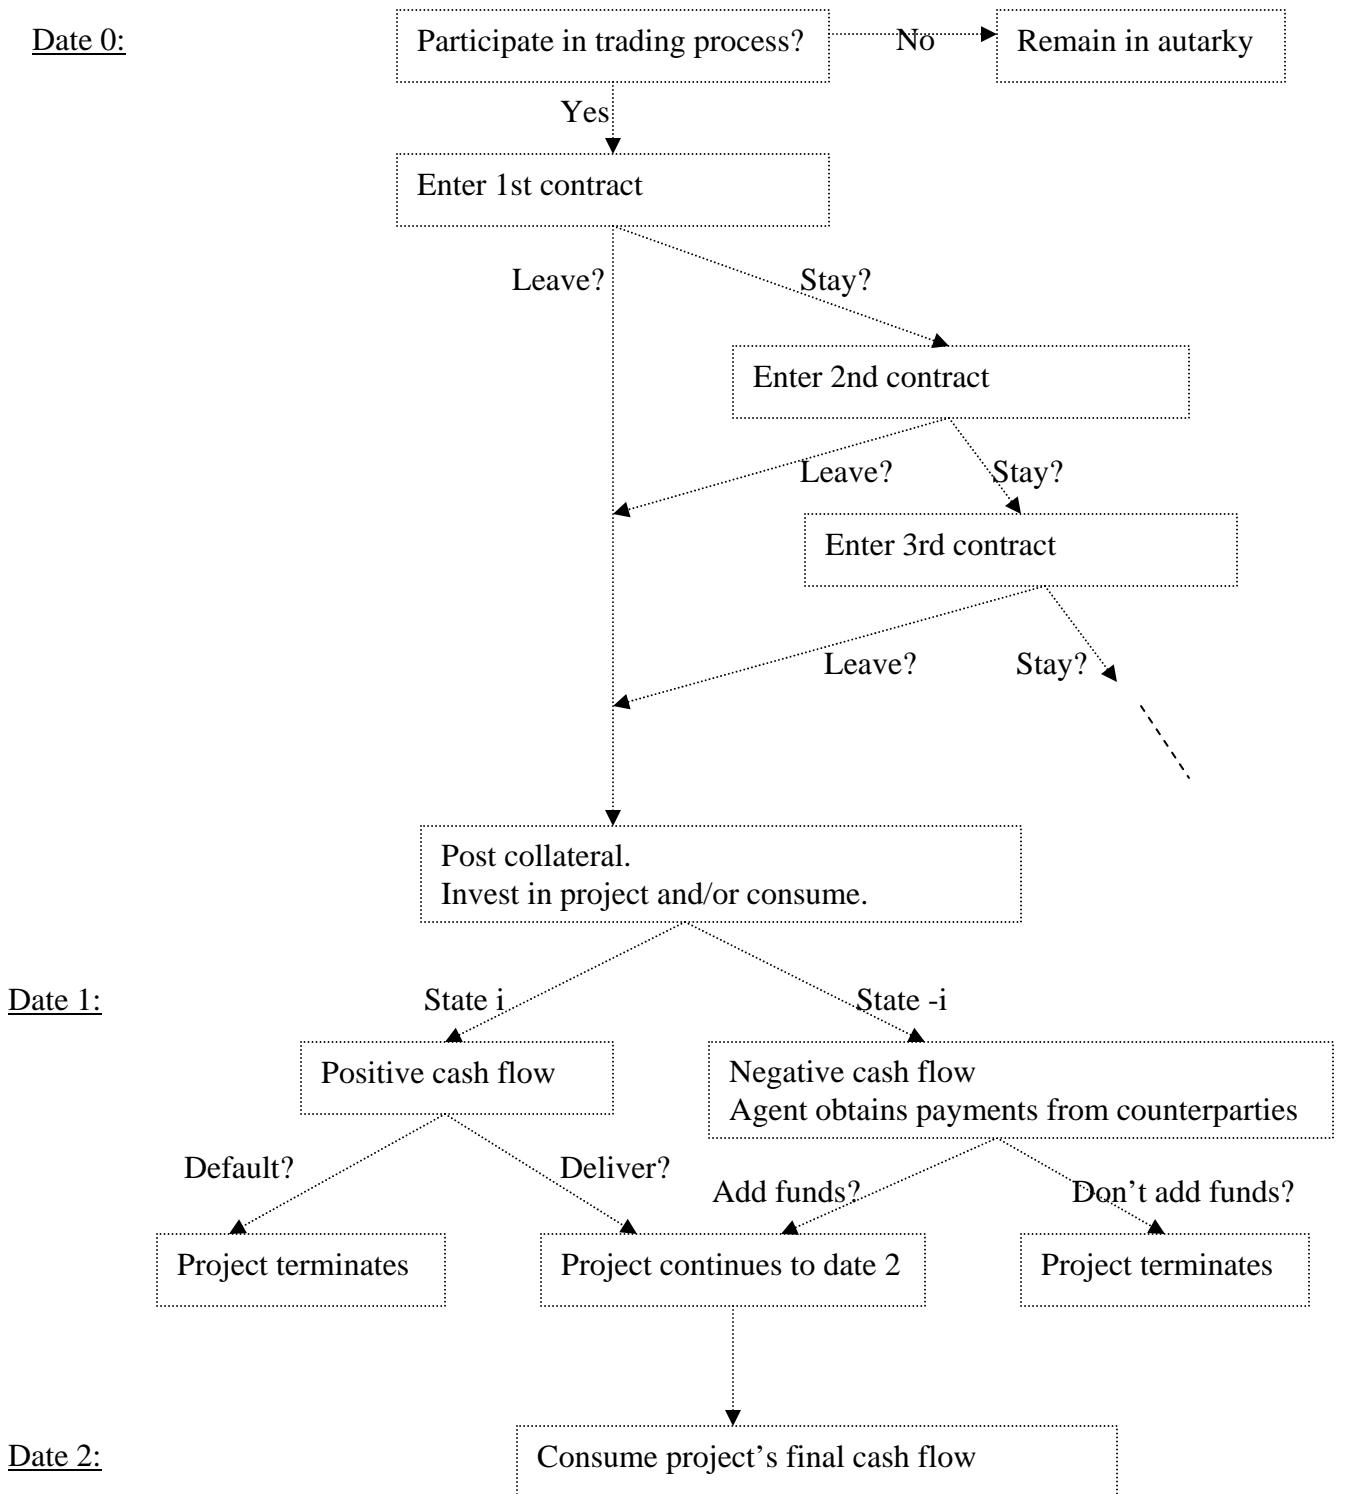


Figure 3. Sequence of events for an agent of type i.