

Discussion Paper No. 743

MARKET FRICTIONS AND HIERARCHICAL
TRADING INSTITUTIONS

by

Birger Wernerfelt*
J.L. Kellogg Graduate School of Management
Northwestern University
Evanston, IL 60201
(312-491-3465)

September 1987

*Comments by David Kreps, Roger Myerson, Mark Satterthwaite and Oliver Williamson are gratefully acknowledged. As usual, the standard caveat applies and comments are most welcome.

MARKET FRICTIONS AND HIERARCHICAL TRADING INSTITUTIONS

by

Birger Wernerfelt

Abstract

We consider a model of repeated bilateral trade with time varying private information about benefits and costs. The traders can communicate and each may terminate the relationship at any time. Since there are frictions in the market for trading partners such action will subject both to switching costs. In a class of simple hierarchical trading institutions, where price is constant over time, we consider both "hierarchical institutions," in which one trader dictates whether a trade will take place or not in each period, and "market institutions," in which both traders have a say. We establish folk theorems for the induced games. The folk theorem holds for both institutions, and in the hierarchical case it holds regardless of who the dictator is. However, without communication only the hierarchical institution is possible, but no folk theorem can be obtained and the choice of dictator becomes critical. In particular, a player makes a better dictator if he is subject to higher switching costs and has more valuable information.

Because hierarchical trading institutions require less communication per period, they are cheaper to administer than market alternatives. Their asymptotic efficiency, therefore, makes them attractive economic organizations. The theory is applied to the employment relation and it is shown that its predictions are similar to those of Williamson (1979). However, the underlying mechanisms are quite different.

Key words: Folk theorem, market frictions, hierarchies, economic organization.

1. Introduction

This paper presents a theory of the employment relationship based on the following intuition:

During a typical day an employee will be asked to do several different things, some less enjoyable than others. In principle, the employer and the employee could negotiate over the provision of each service. However, under such an arrangement, they would spend a lot of time negotiating. In practice, therefore, we have the institution called the employment relationship under which the employee has agreed to obey dictates. Under this institution, the power of the employee derives from the fact that he can quit and thereby normally subject the employer to some (perhaps small) costs. Because many different tasks have to be accomplished on an average day, the employee will not quit over one unpleasant hour but will keep score over a large period. Only if he is exploited "too much" will he quit. Similarly, the employer will be careful to avoid this by asking only for a "fair" amount of unpleasant work.

Formally, in a hierarchical trading institution, an incomplete contract is negotiated ex ante and one trader is given the right to dictate the residual aspects of trade over a given period. In contrast, market institutions require that both parties approve of any aspects of trade. The employment relationship, where an employer tells an employee what he wants done for the wage, is a prime example of a hierarchical institution. But, more generally, the entire gestalt of modern society, where production is organized by huge hierarchies, is dependent upon the efficiency of hierarchical trading.

In the following analysis, we compare hierarchical and market trading when

there are frictions in the market for trading partners. Formally, this gives us two games of repeated bilateral trade, both of which are nested in a market with switching costs. To keep things simple, we assume that price is agreed upon ex ante in both games. In the hierarchical setting the dictator has the right to stipulate whether or not trade will take place in each period. Each period may consist of five stages. In the hierarchical institution: (a) The players get private information about benefits and costs; (b) the dictator may receive a message from the other player; (c) the dictator announces whether or not he wants to trade; (d) either player may terminate the relationship; (e) the players trade. In the market alternative, stage (b) consists of two-way communication and (c) has the players agree whether to trade or not.

The first main result, given in Section 2, is a pair of folk theorems for these games. We establish that both institutions may achieve asymptotic efficiency. Intuitively, the underlying logic is that a player will be willing to take a loss on an individual trade if he can expect to make it up later. In the hierarchical setting, the folk theorem holds regardless of who the dictator is, although players with higher switching costs can be benevolent dictators for higher discount rates. However, without communication, the market institution is not possible and no folk theorem can be obtained for the hierarchical case. Further, the choice of dictator becomes critical. A second main result, given in Section 3, is that a player makes a better dictator if he is subject to higher switching costs and has more valuable information. Our interpretation of the folk theorem, given in Section 4, is that hierarchical trading institutions are cheaper to administer than market institutions, since the latter require more communication at each stage. Hierarchical trading should therefore be preferred as long as it is reasonably efficient. Our predictions are similar to those of Williamson

(1979). However, the underlying mechanisms are quite different. In particular, by focussing on administrative costs, we are able to sketch a theory of economic organization without reference to bounded rationality.

In terms of related literature, there are really two bodies of relevant work. On the technical side, other folk theorems for games with private information have been proved by several authors. Fudenberg and Tirole (1986) looked at the case with time invariant private information; Radner (1985) and Abreu, Pearce and Stachetti (1986) considered various combinations of moral hazard and time varying private information; and Townsend (1982) looked at one-sided time varying information. Our topic, two-sided time varying private information with no moral hazard, has not been covered.

On the economic side there is a widespread acknowledgement of the importance of hierarchical trading but not yet a complete analysis of its properties. In Simon's (1951) original paper, one trader is allowed to specify the details of trade when the other trader is relatively indifferent. While this has considerable intuitive appeal, it rests on some degree of bounded rationality. Similarly, Williamson (1979) presented hierarchical trading as a solution supported by bounded rationality or contracting impossible. Despite the fact that nearly all actual incidents of hierarchical trading are of a longer lasting nature, neither of these models make use of the power of repetition. The idea that repetition can discipline the dictator has only recently been introduced by Kreps (1984), in the context of trades where the dictator faces a sequence of players (e.g., a plumber sending a "fair" bill to protect his reputation).¹ The present model is, in many ways, an elaboration and reformulation of Kreps' ideas. In particular, his formulation has been changed in two ways. First, we respecify the game to get the normally observed situation in which two traders form an ongoing

trading relationship. Second, the relationship is nested in a market with frictions such that the threat of termination may work even if this market is ex ante competitive market.

2. The Folk Theorem

a. Hierarchical Version

We consider a possibly infinitely repeated game between two players, A and B. The stage game consists of a sequence of five steps: (a) A and B privately observe the realizations of the random variables α_t and β_t , respectively; (b) player B may make a statement to player A, S_{Bt} , reporting what he just saw; (c) player A decides whether the variable a_t takes the value zero or one; (d) either player may terminate the game (set T_{At} or T_{Bt} equal to one); (e) unless the game is terminated, payoffs $u_A = \alpha_t a_t$ and $u_B = \beta_t a_t$ are realized. If the game is terminated, the players return to a competitive market in which they find new trading partners. Because of frictions in the this market (search time, partner specific investments, etc.), this subjects them to switching costs Δ_A and Δ_B , respectively. As the game is specified, they do not become "lemons" in that market.

For simplicity we first assume that $\alpha_t \in \{\underline{\alpha}, \bar{\alpha}\}$, $\beta_t \in \{\underline{\beta}, \bar{\beta}\}$ such that the commonly known probabilities of $\bar{\alpha}$ and $\bar{\beta}$ are p_A and p_B , respectively. We further assume that $\underline{\beta}$ and $\underline{\beta} + \underline{\alpha}$ are negative, while $\underline{\alpha}$, $\bar{\alpha}$, $\bar{\beta}$, $\underline{\beta} + \bar{\alpha}$ and $\bar{\beta} + \underline{\alpha}$ are positive. So A always wants to trade, while B only sometimes finds it attractive. Both players use $r > 0$ as their interperiod discount rate. (Under the assumption that the underlying market is competitive, we can make statements about the expected equilibrium payoffs. However, except for the postulate that the individual rationality constraints are satisfied, we make no use of this.)

In the repeated game, we use the notation $x^t \equiv (x_1, x_2, \dots, x_t)$. A

strategy for player A is then given by three sequences of functions:

$S_{At}(\alpha^t, S_A^{t-1}, S_B^{t-1}, a^{t-1})$, $a_t(\alpha^t, S_A^t, S_B^t, a^{t-1})$, $T_{At}(\alpha^t, S_A^t, S_B^t, a^t)$. Similarly, a

strategy for player B is given by the two function sequences

$S_{Bt}(\beta^t, S_A^{t-1}, S_B^{t-1}, a^{t-1})$, $T_{Bt}(\beta^t, S_A^t, S_B^t, a^t)$.

Let $\hat{\sigma}$ be a pair of subgame perfect equilibrium supergame strategies, and let \hat{u}_A and \hat{u}_B be the expected discounted average period payoffs in this

equilibrium. The cooperative solution, which entails setting $a_t = 0$ iff

$(\alpha_t, \beta_t) = (\underline{\alpha}, \underline{\beta})$, gives the expected per stage payoffs $u_A^* \equiv p_A \bar{\alpha} + (1 - p_A) p_B \underline{\alpha}$,

$u_B^* \equiv (1 - p_B) p_A \underline{\beta} + p_B \bar{\beta}$. For this model, we will now prove the following folk

theorem:

Proposition 1H: $\forall \epsilon > 0 \exists \bar{r} > 0 \forall r < \bar{r} \exists \hat{\sigma}: \hat{u}_A + \epsilon > u_A^*$ and $\hat{u}_B + \epsilon > u_B^*$.

Proof: We will proceed by construction. Divide time into a sequence of non-overlapping blocks of time each consisting of τ periods. We will consider strategies where termination results from B sending the message $\underline{\beta}$ more than $(1 - p_B)\tau$ times in a block or A responding to this message with $a = 1$ more than a fraction p_A of these times.² The proposition will be established in two steps. First, we show that the discounted average period payoffs, given this block structure, have the desired limiting properties. Second, we find conditions under which the block structure is compatible with subgame perfection.

Given the block structure, the players face conceptually simple Markov decision problems: each player wants to take/claim his benefits early and when they are worth the most, but they do not know ex ante when these times will be. While it is computationally difficult to solve these problems, their solutions clearly exist.

Let us first take a look at player A. From the perspective of A, the

worst thing B can do is to claim his $\underline{\beta}$'s on the first $(1 - p_B)\tau$ periods of each block. In this case, we will say that B plays his "greedy" strategy. Suppose now that A responds with the following, not necessarily optimal, "friendly" strategy: take $a_t = 0$ iff $S_{Bt} = \underline{\beta}$ and $\alpha_t = \underline{\alpha}$ until the last few periods in which a series of $a_t = 0$ or $a_t = 1$ will be taken to use exactly (up to an integer) $p_A(1 - p_B)\tau$ instances of $a_t = 1$ after $S_{Bt} = \underline{\beta}$. The expected discounted average per period payoff to A, resulting from this scenario, is a lower bound to that obtained in a perfect equilibrium within the block structure. Call this \underline{u}_A . We can then prove the following lemma.

Lemma 1: $\lim_{\tau \rightarrow \infty} \lim_{r \rightarrow 0} \underline{u}_A = u_A^*$.

Proof: See Appendix A.

To prove the analog result for B, we assume that A plays greedily and B plays friendly. That is, we assume that A takes $a_t = 1$ after the first $p_A(1 - p_B)\tau$ instance of $S_{Bt} = \underline{\beta}$. Similarly, we assume that B sends truthful messages until the last few periods when he claims $\underline{\beta}$ or $\bar{\beta}$ constantly such that he just fills his quota of $(1 - p_B)\tau$ $\underline{\beta}$'s. Again, here, B's expected discounted average per period payoff is a lower bound to that he will get in a perfect equilibrium under the block structure. By another trivial but tedious calculation we can see that $\lim_{\tau \rightarrow \infty} \lim_{r \rightarrow 0}$ of this lower bound equals u_B^* .

It remains to be shown that the block structure can be a subgame perfect equilibrium.

Let us first establish that it can be a subgame perfect equilibrium. Note that individual rationality is guaranteed by assumption. So, it is sufficient to demonstrate that the players in no circumstances will go beyond their quotas.

Consider first A. His incentives to go beyond his quota are largest if

he has used his $p_A(1 - p_B)\tau$ instances of " $a_t = 1$ after $S_{Bt} = \underline{\beta}$ " in the initial $p_A(1 - p_B)\tau$ periods of a block. At that point his expected net present value from cheating is given by:

$$\sum_{t=1}^{[1-p_A(1-p_B)]\tau} \left(\frac{1}{1+r}\right)^{t-1} [p_A \bar{\alpha} + (1 - p_A)\underline{\alpha}] + (V_A - \Delta_A) \left(\frac{1}{1+r}\right)^{[1-p_A(1-p_B)]\tau}$$

Conversely, his expected net present value from staying within his quota is bounded from below by

$$\sum_{t=(1-p_A)(1-p_B)\tau+1}^{[1-p_A(1-p_B)]\tau} \left(\frac{1}{1+r}\right)^{t-1} [p_A \bar{\alpha} + (1 - p_A)\underline{\alpha}] + V_A \left(\frac{1}{1+r}\right)^{[1-p_A(1-p_B)]\tau}$$

where V_A is A's expected net present value from the equilibrium, evaluated at the start of a block. From this, a sufficient condition for A not violating his quota is:

$$(1) \quad \sum_{t=1}^{(1-p_A)(1-p_B)\tau} \left(\frac{1}{1+r}\right)^{t-1} [p_A \bar{\alpha} + (1 - p_A)\underline{\alpha}] < \Delta_A \left(\frac{1}{1+r}\right)^{[1-p_A(1-p_B)]\tau}$$

To derive an analog condition for B, we focus on the situation where he has used his $(1 - p_B)\tau$ messages claiming $\underline{\beta}$ in the first $(1 - p_B)\tau$ periods of a block. From that point his expected net present value from cheating is bounded from above by

$$\sum_{t=1}^{p_B\tau} \left(\frac{1}{1+r}\right)^{t-1} p_A [p_B \bar{\beta} + (1 - p_B)\underline{\beta}] + (V_B - \Delta_B) \left(\frac{1}{1+r}\right)^{p_B\tau}$$

Conversely, his expected net present value from staying within his quota is

$$\sum_{t=1}^{p_B\tau} \left(\frac{1}{1+r}\right)^{t-1} [p_B \bar{\beta} + (1 - p_B)\underline{\beta}] + V_B \left(\frac{1}{1+r}\right)^{p_B\tau}$$

where V_B is B's expected net present value from the equilibrium, evaluated at the start of a block. So a sufficient condition for B to stay within his quota is:

$$(2) \quad \sum_{t=1}^{P_B \tau} \left(\frac{1}{1+r}\right)^{t-1} (1 - p_A) [p_B \bar{\beta} + (1 - p_B) \underline{\beta}] < \Delta_B \left(\frac{1}{1+r}\right)^{P_B \tau}$$

For a given r , (1) and (2) give an upper bound on τ , call $\bar{\tau}(r)$. If τ is greater than this, the temptation to cheat will become overwhelming. We can show:

Lemma 2: $\bar{\tau}(r) \rightarrow \infty$ as $r \rightarrow 0$.

Proof: See Appendix B.

So the limiting arguments from the first part of the proof remain valid under (1) and (2).

To show subgame perfection we need to make sure that termination is a rational response to violation of a quota. The myopic or one-shot equilibrium in which A sets $a = 1$ and B claims $\bar{\beta}$ at all times is clearly a subgame perfect equilibrium of the supergame. Suppose that the players switch to this equilibrium after any violation until the game is terminated. In this case B will terminate the game if

$$(C) \quad [p_B \bar{\beta} + (1 - p_B) \underline{\beta}] / r < -\Delta_B + \underline{u}_B / r$$

Note that this constraint only binds for sufficiently high values of r .³ So the suggested equilibrium is perfect if (C) holds or if r is sufficiently low. Q.E.D.

Inspection of (1) and (C) reveals the role of the switching costs Δ_A and Δ_B . For any equilibrium and any parameter configuration, higher Δ_A 's increase $\bar{r}(r)$. So increased switching costs will, ceteris paribus, discipline the dictator and allow more efficient equilibria to be sustained. On the other hand, for high values of r , it is necessary that Δ_B be low enough to make the threat of termination credible.

Let us now consider the implications of making player B the dictator. That is, suppose that a_t is set by player B. In this case an identical folk theorem can be established from a construction where A sends messages, subject to a quota system, and B takes action subject to another quota. Thus, we have found a limiting form of the Coase conjecture (Coase, 1960), that the allocation of decision rights is irrelevant.

The very special support for α and β obviously limit the direct value of Proposition 1H. In particular, given the impossibility result of Myerson and Satterthwaite (1983), it would be desirable to look at the case where α and β have more general distributions F_α and F_β on $[0,1]$. As one would expect, this is no problem.

Proposition 1H': Proposition 1H holds for any F_α, F_β .

Proof: Suppose first that the F 's have no mass points. Consider the same strategies as in the proof of 1H with the following difference. Fix an integer $n \geq 2$. If in any block player A announces "too many" values in any interval $[0, 1/n), [1/n, 2/n), \dots, [(n-1)/n, 1]$, player B will terminate the game. Similarly, if player B makes "too many" dictates which place his value in any interval $[1/n, 1], [2/n, 1], \dots, [(n-1)/n, 1]$, player A will terminate the game. In such a setting, "truthtelling" consists in announcing the actual intervals in which α and β are realized. Clearly for fixed n we can proceed

as in 1H to show that the players almost surely will tell the truth as $\tau \rightarrow \infty$ after $r \rightarrow 0$. The average efficiency loss from this is

$$\sum_{i=0}^{n-1} [F_{\alpha}(\frac{i+1}{n}) - F_{\alpha}(\frac{i}{n})][F_{\beta}(\frac{i+1}{n}) - F_{\beta}(\frac{i}{n})](1/2)$$

If we let n go to infinity, this goes to zero for differentiable F_{α}, F_{β} .

If one or both F 's have mass points, the proof is modified by letting each mass point be its own "interval." Q.E.D.

b. Market Version

In this institution the price is still agreed upon ex ante, but the game has the following stages each period: (a) A and B privately observe α_t and β_t , respectively; (b) they may make statements to each other, S_{At} and S_{Bt} , reporting what they just saw; (c) they negotiate about whether or not to trade; (d) either player may terminate the game; (e) unless the game is terminated, payoffs are realized. The differences between this and the hierarchical version are in stages (b) and (c): because both parties need to approve each trade, we now have two-way communication, instead of one-way communication in these two stages.

Not surprisingly, we can prove a folk theorem here also:

Proposition 1M: $\forall \epsilon > 0 \exists \bar{r} > 0 \forall r < \bar{r} \exists \hat{\sigma}: \hat{u}_A + \epsilon > u_A^*$ and $\hat{u}_B + \epsilon > u_B^*$.

Proof: We use the same strategies as in the proof of 1H with the following change. If in any block either player announces "too many" high or low values, his opponent will terminate the game. Further, if a player refuses to trade when the announced values indicate that it is efficient, his opponent will terminate the game.

In the proof of 1H we showed that A would converge to true reporting as

$\tau \rightarrow \infty$ after $r \rightarrow 0$. This argument now applies to both players. Similarly, the argument from 1H that B will converge to "fair" dictates now applies to both players. Q.E.D.

We state without proof:

Proposition 1M': Proposition 1M holds for any F_α, F_β .

We used the same strategies to construct the proofs of Propositions 1H and 1M. Asymptotically, the two institutions perform identically, except for communication costs. For given parameter values, however, the market institution can perform better (again except for communication costs) since more information can be exchanged. The folk theorems show that the efficiency differences vanish as the frequency of trading goes up. At the same time, the weight of the communication costs will increase in that case.

3. Choice of Dictator

If the players cannot communicate, or find it too costly to do so, the market institution is no longer possible. For the hierarchical institution, matters may change greatly. In this case, the dictator has no way of finding out when his action is more costly to the other player. The most efficient equilibria are the ones where the dictator sets a_t equal to zero or one at all times, or lets a_t reflect only his own private information. In the latter case, the choice of dictator matters.

Suppose first that A is the dictator. The highest attainable joint average payoff is $p_A \bar{\alpha} + p_A p_B \bar{\beta} + p_A(1 - p_B) \underline{\beta}$. However, this upper bound is only attainable if A can be restrained from cheating. That is, we need a sufficiently low r and/or a high Δ_A , analogous to (1). Similarly, if B is the dictator, the upper bound on the joint average payoff is

$p_B p_A \bar{\alpha} + p_B (1 - p_A) \underline{\alpha} + p_B \bar{\beta}$ and feasibility requires that r is low and/or that Δ_B is high, as in (2).

Summarizing, in the limit, player A is a better dictator if Δ_A is high, Δ_B is low and

$$(3) \quad p_A (1 - p_B) (\bar{\alpha} + \underline{\beta}) > (1 - p_A) p_B (\underline{\alpha} + \bar{\beta})$$

To interpret (3), note that $p_A (1 - p_B) (\bar{\alpha} + \underline{\beta})$ is the minimum foregone utility if B is the dictator, while at least $(1 - p_A) p_B (\underline{\alpha} + \bar{\beta})$ is lost if A is the dictator. Thus, we find that a player is a better dictator if he is subject to higher switching costs and has more valuable information.⁴ (See also Farrell, 1987.)

4. Interpretation and Extensions

Proposition 1H and 1M apply to a very special economic structure patterned after a bargaining problem. It should, however, be quite clear that the technique used in the proof can be adapted to yield much more general results. Such games may, of course, fit the description of an employment relationship more closely.⁵

The main message of Proposition 1H is that a hierarchical trading institution is asymptotically efficient when traders engage in repeated exchange under time varying incomplete information. The market institution is more efficient, as are institutions with time varying price menus or arbitrators (Farrell, 1986). The problem is that the market institution is costly to administer. In fact, it is particularly unattractive when trades are frequent and small: one does not want to spend too much time talking about a \$.50 deal. However, frequent trading is exactly the case in which the hierarchical alternative has good efficiency properties.

At this point it is instructive to compare the assumptions of our model with those of alternative explanations of hierarchical trading. If one takes the mechanism design literature on face value, it is surprising how few complex contracting relationships one actually observes. The extant literature on hierarchical trading (Williamson, 1979) has solved this dilemma by assuming that complete contracting is impossible, ultimately because of bounded rationality. Thus the emergence of hierarchies as an inefficient, but feasible, alternative. In contrast, we have shown that hierarchies may be asymptotically efficient, such that small differences in the costs of creating and administering alternative institutions may be the decisive factor. So the crux of our argument is administrative costs, not bounded rationality.

If we compare the present model to Williamson's (1979) theory of economic organization, the two make very similar predictions. In the absence of market frictions, spot markets suffice. With market frictions, infrequent trading demand contracts and frequent trading leads to hierarchies. However, the results do depend on very different premises. For Williamson, hierarchies are expensive to create and high frequency helps spread the costs over many trades. In contrast, we here look at hierarchies as relatively cheaper to administer, with high frequency giving them approximate efficiency.

5. Conclusion

On a fundamental level, the present theory suggests a refinement of the central premise of the Coaseian research program: the idea that the transaction should be the unit of analysis. In our model, individual transactions need not balance out; it is only over the life of the trading relationship that incentives are important and efficiency is achieved.

More directly, and perhaps more importantly, the paper breaks away from tradition by emphasizing the role of implicit contracts and the costs of

administering them. So far, the literature has had a tendency to pay only lip service to contracting costs and neglect the costs of administering ongoing trading relationships. To evaluate the merits of this, it is useful to think of the model without communication as it applies to the employment relation. We have a tendency to think of communication as "cheap" relative to the losses given in equation (3), and yet it is very rare that such communication actually is observed. Further, the fact that firms commonly delegate hierarchical authority shows that even simple dictatorship entails a nontrivial administrative burden.

Summarizing, the purpose of this paper has been to suggest a new theory for characterizing the class of transactions for which hierarchical trading institutions are useful. Even though the theory is based on very different premises, its predictions are essentially identical to those of Williamson's (1979) transaction cost theory. At this stage it is therefore difficult to envision a practical and yet discriminating empirical test. Perhaps further work on both theories, including applications to property rights, could help in this regard. Overall, however, we suspect that the two theories are complementary. Either way, we do not dispute the validity of the extant explanations of hierarchical trading, with its focus on noncontractability. However, by putting emphasis on the costs of administering trading relationships, we can explain hierarchical trading institutions without reference to bounded rationality. Even in such a setting, there is obviously more to this than efficient trade and administrative costs. In particular, the influence costs analyzed by Milgrom and Roberts (1986) and the incentive decay stressed by Williamson (1984) seem important.

Notes

1. Gilson and Mnookin (1985) informally use a folk theorem in their analysis of sharing among partners of law firms.
2. There is no reason to believe that this equilibrium is second best. One would expect that constructions similar to those of Holmstrom and Milgrom (1987) and Radner (1985) dominate it.
3. This statement depends on Δ being independent of r . If we interpret Δ as sunk costs and nest the trading relationship in an ex ante competitive market, then $\Delta_A \rightarrow u_A^*/r$ as $r \rightarrow 0$ and (C) degenerates to $p_B \beta + (1 - p_B) \beta < 0$ in the limit.
4. In principle we could find a single inequality condition for the choice of dictator by comparing payoffs from the optimal equilibria for given parameter values. However, such an exercise is very difficult to carry out, and it is not clear that the insights would extend beyond the specifics of the model.
5. To interpret these more general versions of Proposition 1, note that the semantic meaning of different levels of a_t is immaterial. We do not need ex ante knowledge of them, nor do they have to be time invariant. All we need is a constant distribution over their payoff implications.

References

- Abreu, D., D. Pearce and E. Stachetti, "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," mimeo, Harvard University, 1986.
- Coase, R., "The Problem of Social Cost," Journal of Law and Economics 3 (1960), 1-44.
- Farrell, J., "Rigidity vs. License," American Economic Review, 77 (1987), 195-77.
- Farrell, J., "Rights and Efficiency," mimeo, GTE Labs, 1986.
- Fudenberg, D. and E. Maskin, "The Folk Theorem in Repeated Games with Incomplete Information," Econometrica, 54 (1986), 533-54.
- Gilson, R. J. and R. H. Mnookin, "Sharing Among the Human Capitalists," Stanford Law Review, 37 (1985), 313-92.
- Holmstrom, B. and P. Milgrom, "Aggregation and Linearity in the Provision of Intertemporal Incentives," Econometrica, 55 (1987), 303-28.
- Kreps, D. M., "Corporate Culture and Economic Theory," mimeo, Stanford University (1984).
- Myerson, R. B. and M. A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," Journal of Economic Theory, 29 (1983), 265-81.
- Milgrom, P. and J. Roberts, "Bargaining and Influence Costs and the Organization of Economic Activity," mimeo, Stanford University, 1987.
- Radner, R., "Repeated Principal-Agent Games with Discounting," Econometrica 53 (1985), 1173-98.
- Simon, H. A., "A Formal Theory of the Employment Relationship," Econometrica 19 (1951), 293-305.
- Williamson, O. E., "Transaction-cost Economics: The Governance of Contractual

Relations," Journal of Law and Economics 22 (1979), 223-61.

Williamson, O. E., "The Incentive Limits of Firms," Weltwirtschaftliches Archiv, 120 (1984), 736-62.

Appendix A

Proof of Lemma 1: $\lim_{\tau \rightarrow \infty} \lim_{\rho \rightarrow 0} \underline{u}_A = u_A^*$.

We look at the average discounted per period payoff to player A if he plays friendly and B plays his greedy strategy. \underline{u}_A denotes A's average discounted per period payoff within a block. We can write \underline{u}_A as:

$$\begin{aligned}
 (A.1) \quad \tau \underline{u}_A &= \sum_{s=0}^{(1-P_A)(1-P_B)\tau} [P_A^{P_A(1-P_B)\tau} (1-P_A)^s \binom{P_A(1-P_B)\tau+s}{s}] \\
 &\times \sum_{t=1}^{P_A(1-P_B)\tau+s} \left(\frac{1}{1+r}\right)^{t-1} \bar{\alpha} \frac{P_A(1-P_B)\tau}{P_A(1-P_B)\tau+s} \\
 &+ \sum_{s=0}^{P_A(1-P_B)t} [P_A^s (1-P_A)^{(1-P_A)(1-P_B)\tau} \binom{(1-P_A)(1-P_B)\tau+s}{s}] \\
 &\times \left(\sum_{t=1}^{(1-P_A)(1-P_B)\tau+s} \left(\frac{1}{1+r}\right)^{t-1} \bar{\alpha} \frac{s}{(1-P_A)(1-P_B)\tau+s} \right. \\
 &\left. + \sum_{t=(1-P_A)(1-P_B)\tau+s+1}^{(1-P_B)\tau} \left(\frac{1}{1+r}\right)^{t-1} \bar{\alpha} \right) \\
 &+ \sum_{t=(1-P_B)\tau+1}^{\tau} \left(\frac{1}{1+r}\right)^{t-1} [P_A \bar{\alpha} + (1-P_A) \underline{\alpha}]
 \end{aligned}$$

To interpret this, it is helpful to look at Figure 1. The first term

accounts for the realizations where the cumulative number of times in which $a_t = 1$ hits the horizontal line; the second term accounts for the realizations where this total hits the sloping line; and the third term gives the expected net present value for the rest of the block.

Insert Figure 1 About Here

Now fix a τ . As $r \rightarrow 0$ the different temporal positions of the three terms becomes immaterial and (A.1) degenerates to

$$\begin{aligned}
 (A.2) \quad \tau \underline{u}_A = & \sum_{s=0}^{(1-P_A)(1-P_B)\tau} P_A^{(1-P_B)\tau} (1-P_A)^s P_A^{(1-P_B)\tau+s} \bar{\alpha} P_A^{(1-P_B)\tau} \\
 & + \sum_{s=0}^{P_A(1-P_B)\tau} P_A^s (1-P_A)^{(1-P_A)(1-P_B)\tau} (1-P_A)^{(1-P_B)\tau+s} \\
 & \times (\bar{\alpha}s + \bar{\alpha}[(1-P_B)\tau - s - (1-P_A)(1-P_B)\tau]) + \tau P_B [P_A \bar{\alpha} + (1-P_A)\underline{\alpha}]
 \end{aligned}$$

If we let $\tau \rightarrow \infty$ the value of the first two terms converge to that realized when the process moves along a "straight" line from (0,0) to $((1-P_B)\tau, P_A(1-P_B)\tau)$. So as $\tau \rightarrow \infty$, (A.2) degenerates to

$$\begin{aligned}
 (A.3) \quad \underline{u}_A = & (1-P_B)P_A \bar{\alpha} + P_B [P_A \bar{\alpha} + (1-P_A)\underline{\alpha}] \\
 = & [P_A \bar{\alpha} + P_B(1-P_A)\underline{\alpha}] = u_A^*
 \end{aligned}$$

So the average discounted per period payoff goes to u_A^* as first $r \rightarrow 0$ and then $\tau \rightarrow \infty$. Q.E.D.

Appendix B

Proof of Lemma 2: $\bar{\tau}(r) \rightarrow \infty$ as $r \rightarrow 0$.

We here look at the maximum block length $\bar{\tau}(r)$, under which player A will refrain from violating his quota. If we hold r and V_A constant, $\bar{\tau}(r)$ is given by

$$(B.1) \quad \bar{\tau}(r) = \max\{\tau(r) \mid \begin{aligned} & (1 - P_A)(1 - P_B)\tau(r) \\ & \sum_{t=1}^{\tau(r)} \left(\frac{1}{1+r}\right)^{t-1} [P_A \bar{\alpha} + (1 - P_A)\underline{\alpha}] \\ & < \Delta_A \left(\frac{1}{1+r}\right)^{[1 - P_A(1 - P_B)]\tau(r)} \end{aligned} \}$$

from equation (1).

We can rewrite (B.1) as

$$(B.2) \quad \bar{\tau}(r) = \max\{\tau(r) \mid \begin{aligned} & \sum_{t=1 - [1 - P_A(1 - P_B)]\tau(r)}^{\tau(r)} \left(\frac{1}{1+r}\right)^{t-1} \\ & < \Delta_A [P_A \bar{\alpha} + (1 - P_A)\underline{\alpha}]^{-1} \end{aligned} \}$$

If we let $r \rightarrow 0$, the left side of this inequality goes to $(1 - P_B)(1 - P_A)\tau(r) - 1$ while the right side goes to positive infinity. So the maximum block length under which A will abstain from violating his quota goes to infinity as r goes to zero.

The arguments for player B, using equation (2), can be made by analogous methods.

Q.E.D.

Figure 1

Possible Realizations of A's Friendly Strategy

