Econometrica, Vol. 70, No. 2 (March, 2002), 547–582

THE TIME STRUCTURE OF SELF-ENFORCING AGREEMENTS

BY DEBRAJ RAY[1]

A principal and an agent enter into a sequence of agreements. The principal faces an interim participation constraint at each date, but can commit to the current agreement; in contrast, the agent has the opportunity to renege on the current agreement. We study the time structure of agreement sequences that satisfy participation and no-deviation constraints and are (constrained) efficient. We show that every such sequence must, after a finite number of dates, exhibit a continuation that maximizes the *agent's* payoff over all such efficient, self-enforcing sequences. Additional results are provided for situations with transferable payoffs.

KEYWORDS: Contracts, constrained efficiency, principal-agent theory, limited enforcement, dynamic incentives.

1. INTRODUCTION

CONSIDER A REPEATED RELATIONSHIP between two individuals, a principal and an agent. At each date, an action vector (from some compact metric space) is agreed upon, and past actions are fully observed. (Though it is unnecessary to make this explicit, some components of this vector describe actions to be taken by the principal, and the remainder are actions to be taken by the agent.) Both principal and agent derive payoffs from actions, and intertemporal payoffs are evaluated using a common discount factor.

We restrict the space of agreements in two ways. First, we assume that one component in any agreement is a monetary transfer between principal and agent. However, we do not suppose that payoffs are necessarily linear in money, so all sorts of nontransferable payoff problems can be arbitrarily closely approximated. Second, we assume that agreements that lie within the Pareto-frontier of the one-shot game can always be improved by a small amount. This rules out "local maxima" or inflections and also implies that the space of agreements is rich enough to allow for continuous changes in the nontransfer components.

The fundamental (and only) asymmetry that we impose is that *the principal is committed to honoring the terms of the current agreement, while the agent is not.* To be more specific, while the principal may refuse to deal with an agent at the beginning of a date, he must adhere to any current agreement. On the other hand, the agent may take advantage of the going agreement and shirk his

[1] I am grateful to Parikshit Ghosh for numerous conversations on the subject matter of this paper, and for permitting me to record these results even though they are part of a larger joint project. I also thank Rui Albuquerque, Venkataraman Bhaskar, Mark Gertler, Michihiro Kandori, Jonathan Thomas, a co-editor, and three anonymous referees for extremely useful comments and suggestions.

responsibilities, or repudiate his obligations. In that case the agent receives some continuation value, which may sometimes be interpreted as an outside option.

A sequence of agreements that is proof against such deviations (and respects the principal's participation constraint at each date) is defined to be *self-enforcing*. A self-enforcing sequence that is Pareto-optimal (in the class of all self-enforcing sequences) will be called an *efficient self-enforcing sequence*, or simply, an ESE sequence. I analyze the time structure of ESE sequences, assuming that the no-deviation constraint does bite for some date for each such sequence.[2] To do so as starkly as possible, I abstract from any other form of incomplete or imperfect information.

The main result of this paper (Theorem 1) is that *every* ESE agreement sequence must, after a finite set of dates, exhibit a continuation that maximizes the *agent's* payoff from the class of all ESE sequences.

At this level of generality, I am unable to say much about the sequence that maximizes agent payoff, except that there must exist such a sequence that displays constant periodicity in actions and payoffs. Therefore—Theorem 2—there exists a cyclical (possibly stationary) path such that for any pair of ESE payoffs, there exists an ESE sequence with those payoffs that merges with the special path after at most a finite number of periods.

Specialization of this model to the case of transferable payoffs at each date yields further results. First, under a mild additional restriction, there is a unique ESE sequence that maximizes agent payoffs, and it is constant over time. It follows—Theorem 3—that the continuation tails of each ESE sequence must be constant (after a finite number of periods) and equal across all ESE sequences, *even in the magnitude of the transfer*. Second—Theorem 4—a description of the "initial phase" of an ESE sequence can be provided when payoffs are transferable. The agent's present-value payoff steadily rises through this phase before settling into the "tail contract" where she earns all the surplus. To be sure, the principal's present-value payoff must then fall through this phase. However, his *current* payoff steadily increases before plunging back to zero at the tail.

Thus, roughly speaking, a relationship can be divided into an "initial phase" and a "mature phase." In the mature phase, the structure of the relationship looks the same irrespective of the bargaining power that initially defined the relationship. All of that information is tucked away in the initial phase. The significance of the initial phase depends on the particular point on the (constrained) efficient frontier from which the relationship is presumed to start. If that point maximizes the agent's payoff, then there is no backloading at all: the mature phase is started up right away. If, on the other hand, that point maximizes the principal's payoff, then a significant initial phase will exist in which the principal picks up a disproportionate share of the gains. Once the mature phase is reached, the continuation path will be no different from that of any other efficient point.

Section 2 begins with simple examples that illustrate these results. The examples are chosen to highlight the potential usefulness of a general analysis—they

[2] As discussed below, this is equivalent to the assertion that discount factors are not too close to unity.

show that familar economic phenomena in diverse contexts may stem from identical formal considerations. The examples also show that the basic idea of this paper appears in different forms in the literature. For instance, Lazear (1981) and Harris and Holmstrom (1982) have argued—in the context of labor markets—that "backloading" of contractual payoff may be optimal when individual actions are subject to moral hazard. Thomas and Worrall (1994) show in the context of a repeated foreign investment model (with threat of expropriation) that continuation payoffs move over time in the direction of the agent that has the power to expropriate (and therefore faces a no-deviation constraint). Albuquerque and Hopenhayn (1997) study a repeated relationship in which working capital is periodically advanced from a lender to a borrower. The lender is assumed to maximize payoff. They observe that borrowing constraints become less severe across time. Finally, there is a (weaker) connection with the results of Mookherjee (1997), who has noted that in a static transferable-utility context, greater bargaining power accorded to an incentive-constrained agent can raise social surplus.

My results are related to the work of these authors. At the same time, I wish to highlight the remarkable generality of the propositions established here. First, they are stated for *every* constrained efficient or ESE sequence, and because of this they highlight the close connection between bargaining power and the existence of initial/mature phases. Second, the framework in which these results are established allows for arbitrary degrees of nontransferability (of payoffs) between principal and agent. In fact, as a technical aside, no curvature assumptions are made on payoff functions, and the arguments do not rely on first-order conditions or Euler equations. Third, by making the connection explicitly with repeated game theory, my findings provide a stark contrast to the usual folk theorems, which suggest an enormous lack of predictive power.[3] Finally, by allowing for several models (including dynamic contracting in labor and credit markets) as special cases, the framework here potentially unites seemingly disparate findings for the different markets.

2. EXAMPLES

I begin with simple examples to motivate the theoretical framework and the main result. The first example is a repeated employment model, in the spirit of Lazear (1981), Harris and Holmstrom (1982), Shapiro and Stiglitz (1984), and Thomas and Worrall (1988).

A REPEATED EMPLOYMENT RELATIONSHIP: At each date an employer demands a certain number of hours of work h from an employee, and pays out

[3] There is a broad correspondence between these findings and the "carrot-and-stick" structure unearthed by Abreu (1986, 1988) in his study of discounted repeated games. Recall that "worst punishment paths" often have the property that all the punishment is inflicted early on, with rewards (carrots) forthcoming along the punishment path itself. The main difference is that Abreu finds this to be a characteristic of *worst* payoff paths. In contrast, we find this to be a property of *efficient* paths, the characterization of which is an entirely different problem (if for no other reason than the fact that in one case, payoffs are being minimized, while in the other, they are being maximized).

a total income of w. The employee suffers disutility $e(h)$ from complying with this request; assume $e(h)$ satisfies standard assumptions.[4] Write the employee's (current) payoff as

$$w - e(h).$$

Suppose that h hours of work produces h units of monetary output. Then the employer's (current) payoff is

$$h - w.$$

Intertemporal payoffs are obtained by summing these current payoffs using a common discount factor of β.

Now, at any date, the employee is free to enjoy his current income but shirk his duties. Assume that shirking is detected at the end of the period. Denote the per-period equivalent of the punishment payoff thereafter by v. (This may involve firing, so that v represents some outside option, or some bad continuation equilibrium within the relationship, which is unnecessary to model for our purposes.)

Finally, suppose that at each date, the employer must make some nonnegative "present-value profit" from his dealings with the employee.

To see the structure of ESE agreements in this example, notice first that an agreement sequence $\{w_t, h_t\}$ is *self-enforcing* if at every date t,

$$\sum_{s=t}^{\infty} \beta^{s-t}[w_s - e(h_s)] \geq w_t + \beta(1-\beta)^{-1}v, \tag{1}$$

which captures the employee's incentive not to shirk, and

$$\sum_{s=t}^{\infty} \beta^{s-t}[h_s - w_s] \geq 0, \tag{2}$$

which captures the employer's interim participation constraint.

Suppose that of these self-enforcing agreements, we are interested in the one that maximizes *employee* payoffs. It can be shown (we will derive it later as a consequence of a more general model) that such an agreement sequence *must* be stationary. We may therefore write the agent's best self-enforcing agreement as one that solves

$$\max_{(w,h)} w - e(h),$$

subject to

$$w - e(h) \geq (1-\beta)w + \beta v, \tag{3}$$

[4] That is, e is increasing, smooth, and strictly convex, with $e'(0) = 0$ and $e(h) \to \infty$ as h approaches some upper bound H on hours.

and

(4) $\qquad h - w \geq 0,$

these constraints simply being stationary versions of (1) and (2). It is obvious that w should be set equal to h, so that the agent's best self-enforcing agreement involves the maximization of $h - e(h)$, subject to the no-deviation constraint

$$\beta h - e(h) \geq \beta v,$$

which is just (3) rewritten for $w = h$. Now let $\hat{h}$ be the unconstrained maximizer of $h - e(h)$; if this automatically satisfies the no-deviation constraint, we are done. In this paper I shall be interested in cases where the no-deviation constraint *does* bite, which is tantamount to saying that β isn't too high. For the problem to be interesting, however, β can't be too low either; otherwise no self-enforcing agreement will exist. Thus suppose that

(5) $\qquad \beta h - e(h) \geq \beta v \quad \text{for some } h, \text{ but that} \quad \beta \hat{h} - e(\hat{h}) < \beta v.$

(Figure 1 illustrates this condition.)

Then it is easy to see that the agent's best agreement involves hours h^*, characterized as the largest solution to the inequality $\beta h - e(h) \geq \beta v$. See Figure 1. As already discussed, $w^* = h^*$.

Now, this stationary agreement (w^*, h^*) is a special ESE agreement sequence: it maximizes employee payoffs. What about the other ESE sequences?

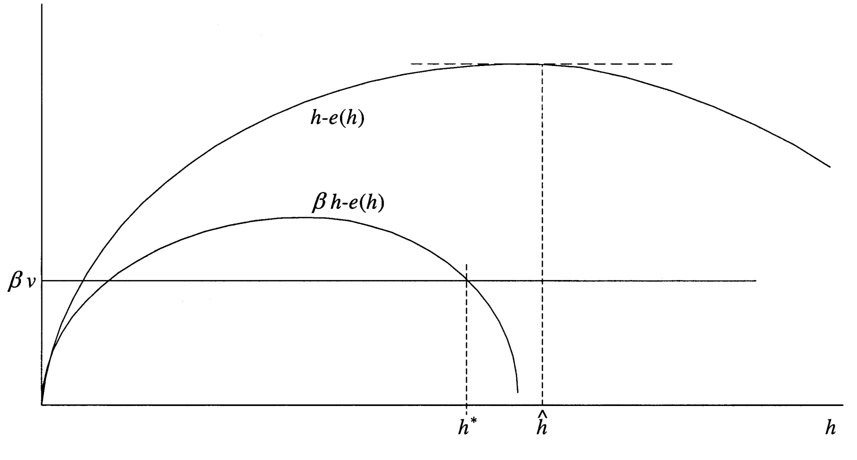

FIGURE 1.—Restrictions on the discount factor in Example 1.

To analyze these, suppose, first, that no lower bound on employee payments or payoffs need to be honored at any particular date. Then the following observation is true:

OBSERVATION 1: *If there are no restrictions on the values of w or h, then every ESE agreement sequence must coincide with the agent's best self-enforcing agreement from the second period onwards.*

We shall presently be providing formal arguments for a more general case, so let us see informally why Observation 1 must be true. The following considerations are relevant:

(i) Under our assumptions, along any ESE sequence, the initial working hours must be strictly less than the unconstrained efficient point $\hat{h}$. (While our restriction on the discount factor was only used to achieve this for the agent's best agreement, the same restriction can be shown to have the wider implication as well.)

(ii) If the ESE agreement does not coincide with the agent's best agreement from the second period onwards, then the agent's continuation payoff must be less than the best possible self-enforcing payoff to the agent, and consequently, by efficiency, the *principal's* continuation payoff must be strictly positive.

(iii) By (ii), it is possible to modify the agreement sequence by having the principal transfer some extra income to the agent in the second period without violating the principal's participation constraint.

(iv) This act relaxes the no-deviation constraint at date 0. The agent can be made to work harder (and can also be paid less) at date 0 to compensate the principal for his promised future transfer. By consideration (i), as long as *some* extra work is done, it can compensate for the principal's proposed transfer while still leaving the agent better off. We have therefore Pareto-improved on the original agreement sequence while maintaining all the self-enforcement constraints.

(v) This contradicts the presumption that we had an ESE sequence to start with. Therefore the starting premise of consideration (ii) is invalid, and the observation is proved.

The reasoning is in line with Lazear (1981), Harris and Holmstrom (1982), and Thomas and Worrall (1988), all of whom observe that a worker's compensation would generally be backloaded in a repeated relationship. To see how one might begin to go further, consider now the case in which there are additional restrictions on the contract at every date. These could be driven, for instance, by considerations of limited liability (on the part of the employee) or by other legal or even social considerations. Suppose, for example, that net payoff to the employee must be nonnegative at every date, so that w and h are jointly restricted by the condition that $w - e(h) \geq 0$. In this case, one still obtains convergence to the agent's best self-enforcing agreement, but not in one period:

OBSERVATION 2: *If w and h are restricted as described, and if the ESE agreement sequence in question is not the agent's best self-enforcing sequence, then there exists $T \geq 1$ such that*

$$h_s < h_{s+1} \tag{6}$$

for all dates $s = 0, \ldots, T-1$, and (provided $T \geq 2$),

$$w_s = e(h_s) \tag{7}$$

for all dates $s = 0, \ldots, T-2$. Thereafter work hours are constant at a value h^* that is the largest root of the equation

$$\beta h - e(h) = \beta v,$$

and wages are increased to h^*, which exceeds $e(h^*)$.

Observation 2 is harder to establish than Observation 1, but the basic logic is the same. An ESE agreement sequence appears to divide itself into two parts: an "initial phase" in which the worker obtains no surplus but is progressively made to work harder and harder,[5] and a "final phase" in which the worker's best ESE agreement is implemented. This sort of analysis has potentially interesting implications, as we now illustrate with the use of this example.

First, note that while the agent is made to work progressively harder, and indeed, is working at his hardest during the final phase, he feels "overworked" during the initial phase and "underworked" during the final phase, *relative* to the income that he receives.[6] One could, of course, make this sort of argument very simply by assuming that there is some fixed number of hours to be worked in each date, but our arguments show that the phenomenon is at once more subtle and more robust (hours are endogenously chosen here). It is true that senior employees feel underworked in this model, but they also work longer hours. Whether this holds up in suitably controlled empirical analysis appears to be an open question.[7]

[5] This statement is slightly imprecise, as a careful reading of (6) and (7) will reveal. At the very end of the initial phase (period $T-1$), the worker may receive some surplus even though we are not yet at the agent's best continuation.

[6] To see this, define the *hourly rate* by $\nu \equiv w/h$. If the employee could choose hours in response to the hourly rate, he would select h to maximize $\nu h - e(h)$. This leads to a preferred choice of hours that is lower than what he is being asked to do in the initial phase, and higher than the hours stipulation in the final phase.

[7] Meanwhile, the employer's payoffs change as well. They rise steadily through the initial phase, but drop in the final phase to zero as the employee's best self-enforcing agreement takes over. This variation raises issues about employer credibility in upholding the agreement sequence. Because this question is of wider applicability, I take it up after discussing the general model.

Second, the relative importance of the initial—overworked—phase is linked to the bargaining power of the employer. The larger this power, the larger the length of the initial phase and the more prevalent will "overworking" be. At the other extreme, if the agent picks up all the surplus, the initial phase disappears altogether.[8] If one provisionally accepts the notion that in times of high unemployment, employers have greater bargaining power (over and above their shifting outside options), then high unemployment societies are more likely to be correlated with the presence of overworked employees. In contrast, in regimes where the worker has power, we would accept a preponderance of the second phase in our observations, and workers would largely feel underworked.[9]

Finally—and most important from the point of view of motivating a general analysis—there are formal similarities between the labor-market characteristics that I have discussed in some detail here, and apparently different phenomena in different markets or relationships. To this end, I briefly sketch a second example involving self-enforcing credit agreements.

A Repeated Credit Relationship: We base this example on Thomas and Worrall (1994) and Albuquerque and Hopenhayn (1997). At each date a lender provides a (working capital) loan ℓ to a borrower, and demands a repayment R at the end of the period. The loan produces an output $F(\ell)$, where F has the usual properties.[10] At any date the borrower may refuse to repay, in which case the credit relationship is terminated. The borrower obtains an outside value with a per-period equivalent of v from the next date onwards. The lender's next best opportunity at any date is to earn a safe rate of interest of r on his loan.

Thus current payoffs on an agreement (ℓ, R) (that's complied with) are

$$R-(1+r)\ell$$

to the lender, and

$$F(\ell)-R$$

to the borrower.

Once again assume a common discount factor β. Then an agreement sequence $\{\ell_t, R_t\}$ is self-enforcing if at every date t,

$$\sum_{s=t}^{\infty}\beta^{s-t}[F(\ell_s)-R_s]\geq F(\ell_t)+\beta(1-\beta)^{-1}v, \tag{8}$$

[8] There are two differences in the present setup (compared, say, to Lazear (1981)) that precipitate this additional structure. First, capital markets are imperfect—the agent cannot borrow from future earnings. Second, the principal (or employer, in this example) cannot be driven to negative present-value earnings at any date. As I have shown, these two assumptions are sufficient to guarantee that *no* backloading occurs when the agent has all the bargaining power. The initial, nonstationary phase is linked to the extent of the surplus that the principal can capture.

[9] Authors such as Lang and Kahn (2001) document the persistence of "underworking" in relatively high-employment U.S. and Canada, while noting the pressures to reduce the work week in continental Europe, where unemployment rates are higher.

[10] Specifically, assume that F is increasing, strictly concave, smooth, and satisfies the Inada conditions.

which captures the borrower's incentive not to default, and

$$\sum_{s=t}^{\infty} \beta^s [R_s - (1+r)\ell_s] \geq 0, \tag{9}$$

which describes the lender's interim participation constraint.

Impose the limited liability constraint on the borrower that net payoffs $F(\ell) - R$ be nonnegative. Then analogues of Observations 1 and 2 hold for this case. When the borrower's best (self-enforcing) agreement sequence is considered, the solution is stationary, just as in the employment relationship. If we assume—in parallel fashion to the previous example—that the unconstrained efficient loan size cannot be self-enforcing, the borrower's best loan size ℓ^* is the largest solution to the equation

$$\beta F(\ell) - (1+r)\ell = \beta v,$$

and repayment R^* is set equal to $(1+r)\ell^*$.

For *no* other ESE sequence are the agreements time-invariant. In general, there will exist T such that

$$\ell_s < \ell_{s+1} \tag{10}$$

for all dates $s = 0, \ldots, T-1$, and such that

$$R_s = F(\ell_s) \tag{11}$$

for all dates $s = 0, \ldots, T-2$ (provided $T \geq 2$). Thereafter the loan size—and indeed the entire agreement sequence—settles down to the stationary terms under the borrower's best sequence.

Construct implicit rates of interest corresponding to the two phases. It is very easy to see that in the initial phase, there must be *loan pushing* relative to the implicit rate of interest at any date. In the final phase, in contrast, there is credit rationing (relative to the rate of interest in that phase), even though the loan is now higher. Observe the close parallel between loan pushing and overworking on the one hand, and credit rationing and underworking on the other. Parallels such as these motivate and justify an examination of the generality of the results, to which we now turn.

3. A GENERAL MODEL

3.1. *Agreements and Payoffs*

There are two actors in this model, a principal and an agent. At every date an agreement ξ is specified from some feasible set Ξ. An agreement may be a complicated object: it could specify an entire list of inputs to be supplied, actions to be taken, various responsibilities as well as monetary transfers. Some components of these are the responsibility of the principal (e.g., the supply of capital or land), and others are the responsibility of the agent (e.g., hours worked).[11] I assume throughout that all components of an agreement are fully observable ex post.

Both principal and agent have payoff functions defined on Ξ, which we denote by P and A respectively.

I am going to put some structure on the set of feasible agreements. Specifically, I suppose that each agreement $\xi \in \Xi$ can be written as a pair (c, m), where c is drawn from some compact metric space, and m is some scalar variable ("money"), to be interpreted as a transfer made from the principal to the agent. In principle, m could take on negative values in which case the opposite interpretation applies.

For instance, in the employment example c simply denotes hours worked, while the wage paid is m. In the credit example c denotes loan size, while the repayment is to be interpreted as the negative of m.

This structure allows me to introduce a function $L(\xi)$, increasing in m, such that an agreement ξ lies in Ξ if and only if $L(\xi) \geq 0$. Think of this as the "limited liability" function. With continuity of the payoff functions (to be assumed in a moment), one can mimic unlimited liability by simply shifting L downwards by a large constant. Thus our formulation includes all sorts of liability constraints (or their absence). For instance, L could be made to coincide with the function A, or some shift of it.

I shall make the following assumption.

(A.1) *P, A, and L are continuous functions. They are continuously differentiable in m, with A and L strictly increasing in m and P strictly decreasing, for each c. Moreover, the partial derivatives of A, L, and P with respect to m are uniformly bounded above and below—in absolute value—by strictly positive numbers.*

The substantive concern of (A.1) is to formalize the interpretation of the money variable m as a transfer from principal to agent. *Note that linearity is not assumed*, and that arbitrarily high degrees of nontransferability may be approximated.[12]

Say that an agreement ξ is *improvable* if there exists $\xi' \in \Xi$ such that $A(\xi') > A(\xi)$ and $P(\xi') > P(\xi)$. Otherwise, it is *unimprovable*. (I avoid the descriptions "inefficient" and "efficient" here because I reserve such terms for dynamic payoff comparisons.) An agreement ξ is *ϵ-improvable* if the ϵ-ball in Ξ around ξ contains only improvable agreements. An improvable agreement ξ can be δ-improved if for every $\nu > 0$ and sufficiently small, there exists a feasible agreement ξ' within the ν-ball of ξ such that $\min\{A(\xi') - A(\xi), P(\xi') - P(\xi)\} \geq \delta\nu$. I make the following assumption.

(A.2) *For every $\epsilon > 0$, there exists $\delta > 0$ such that if an agreement is ϵ-improvable, then it can be δ-improved.*

[11] A more conventional game-theoretic formulation would separate Ξ into two action spaces, but for our purposes it is unnecessary to do so.

[12] To be sure, perfect nontransferability is ruled out.

This assumption essentially states that there are no "local maxima" or inflection points in the interior of the set of feasible payoffs.

3.2. *Repeated Agreements*

The relationship between principal and agent takes place over infinite time, unless the agent breaks the relationship (see below). Consider a sequence of agreements $\{\xi_t\}$ and assume for the moment that the agreements are honored. Suppose that both individuals have a common discount factor β. Then we may write the continuation payoff to the agent from any date t onwards as

$$\alpha_t \equiv \sum_{s=t}^{\infty} \beta^{s-t} A(\xi_s), \tag{12}$$

and the continuation payoff to the principal from any date t onwards as

$$\pi_t \equiv \sum_{s=t}^{\infty} \beta^{s-t} P(\xi_s). \tag{13}$$

When we say that an agreement sequence $\{\xi_t\}$ yields a payoff of α to the agent and π to the principal, we mean that $(\alpha_0, \pi_0) = (\alpha, \pi)$.

3.3. *Self-Enforcing Agreement Sequences*

At any date, the principal is bound to his responsibilities as described by the going agreement ξ, which includes any promised transfer to the agent before the latter's actions are observed. But we assume that the agent can renege on the agreement. In case of defection the agent obtains some value, which we denote by $V(\xi)$.

V includes a one-time "deviation payoff," followed by some continuation payoff. As already indicated, this may be viewed as an outside option or some punishment payoff within the same relationship.[13] To be sure, the sum of the deviation and continuation payoffs will generally depend on the "going agreement" from which the deviation occurred, which is why we write V as a function of ξ.

The following assumption on V will be maintained throughout:

(A.3) *V is continuous in ξ and differentiable in m. Moreover, $V - A$ is Lipschitz,*[14] *and*

$$0 \le \frac{\partial V}{\partial m}(\xi) \le \frac{\partial A}{\partial m}(\xi). \tag{14}$$

[13] Under the "outside-option" interpretation, it should be observed that firing a deviating agent is easily sustained as a subgame perfect equilibrium of the repeated relationship (even if the hunt for a new agent is costly). In applications, one would simply choose a continuation equilibrium—between the principal and the "unfired" agent—that is well within the constrained efficient frontier. This will work as long as the cost of replacement is not too large.

[14] Recall that a real-valued function f on some metric space $(\mathscr{X}, \rho)$ is *Lipschitz* if there exists $\lambda < \infty$ such that for all $x, x' \in \mathscr{X}$, $|f(x) - f(x')| \le \lambda\rho(x, x')$. We will only use the upper Lipschitz bound on $V - A$.

The substantive part of (A.3) is the restriction (14), which needs some discussion. The first inequality is harmless given our interpretation of m: it is reasonable to assume that an increased transfer does not lower the outside value for the agent. (To be sure, it might leave the value unchanged if the agent cannot take the transfer with him in the event of a default.) The second condition is stronger. Informally, it states that the agent cannot do any better with an increase in m if he defaults, than he can do "on the job." This assumption is a real restriction and will need to be checked in applications; my conjecture is that the results in this paper will not generally survive the dropping of (A.3).

Notice that in the employment game, $V(\xi) = V(h, w) = w + \beta(1-\beta)^{-1}v$, and $A(\xi) = A(h, w) = w - e(h)$. The derivatives with respect to the transfer are both 1, so that (14) is satisfied. In the credit game, $V(\xi) = V(\ell, -R) = F(\ell) + \beta(1-\beta)^{-1}v$, while $A(\xi) = A(\ell, -R) = F(\ell) - R$. Noting that $m = -R$, we see that the derivative of V with respect to m is zero, while the corresponding derivative for A is 1, so that once again (14) is satisfied. Therefore (A.3) is met in both the examples.

Let $\{\xi_t\}$ be a sequence of agreements in Ξ (with associated payoff sequence $\{\alpha_t, \pi_t\}$). Say that such a sequence is *self-enforcing* if for every $t \geq 0$,

$$\alpha_t \geq V(\xi_t) \tag{15}$$

and

$$\pi_t \geq 0. \tag{16}$$

Inequality (15) requires that at no date should the agent find it profitable to seize his (post-deviation) continuation value and default on the agreement. Inequality (16) is a participation constraint for the principal. It states that the principal must find it profitable—at each date—to continue with the relationship (that is, the present value of his outside option has been normalized to zero).

The presumption that a self-enforcing agreement must satisfy a participation constraint for the principal but an incentive constraint for the agent is fundamental. So as to explore the implications in a sharp way, we assume no other asymmetry between principal and agent.

3.4. *Efficient Self-Enforcing Agreement Sequences*

A self-enforcing agreement sequence yielding payoffs (α, π) is *weakly efficient* (WESE) if there is no other self-enforcing sequence with payoffs $(\alpha', \pi') \gg (\alpha, \pi)$. It is *efficient* (ESE) if, given π, α is the maximum payoff available among all self-enforcing sequences that yield at least π to the principal.

Weak efficiency and efficiency typically coincide, except possibly at jump points in the efficient frontier of self-enforcing payoffs.

I make the following additional assumptions.

(A.4) *No WESE sequence is wholly composed of unimprovable agreements.*

(A.5) *Along no WESE sequence is the limited liability condition binding at all dates*.

I know these assumptions are somewhat unsatisfactory, in that they are not placed on the primitives of the problem. But they are compelling and intuitive enough that I feel comfortable presenting them in this way. Certainly, if a condition such as (A.4) does not hold, incentive constraints are of no special import, and the problem is uninteresting.

Moreover, in examples, (A.4) and (A.5) are particularly easy to verify. (A.4) typically boils down to an upper bound on the discount factor, so that the interesting incentive problems do not entirely disappear. (A.5) is needed because the limited liability constraint is described in a very abstract way. Essentially, all that (A.5) says is that nontrivial self-enforcing agreement schemes do exist (if the only ones were trivial, involving no agreement, then the limited liability condition may well bite throughout). When nontrivial schemes do exist, (A.5) is satisfied in all the examples I have considered. The reason is simple: the post-deviation continuation value V (which includes the immediate deviation payoff) typically dominates any sequence of agreements in which the agent is pushed to the limits of her liability *at every date*. So no agent would conform to such a path of agreements.

In the context of the examples in Section 2, see condition (5) as well as its illustration in Figure 1.

An ESE sequence of particular interest is one that maximizes the *agent's* payoff among the class of all ESE sequences. Call such a sequence an *agent's best self-enforcing sequence* (at this level of generality there may be more than one such sequence).

3.5. *The Structure of ESE Sequences*

Our principal result is the following theorem.

THEOREM 1: *Let $\{\xi_t\}$ be an ESE agreement sequence. Then there exists a finite date such that an agent's best self-enforcing sequence is followed thereafter.*

That is, regardless of bargaining power, the agreement sequence "ultimately" looks like one that maximizes agent payoff. This is a substantial generalization of the main observation in the examples of Section 2, and one that may be worth some intuitive discussion.

We will prove Theorem 1 by showing that along *any* ESE sequence, there must exist some date at which the principal's continuation payoff is zero. Given the simple observation that the continuation of any ESE sequence is itself ESE, this suffices to establish that the agent's best self-enforcing sequence must be followed from some finite date onwards.

The idea is to try and imitate the familiar backloading argument up to a point: if the principal's continuation payoffs are *always* positive, then shift some payoff for the agent from the present into the future. One could do this by cutting transfers to the agent today and then reimbursing her at some future date (with continuation payoffs positive for the principal, this should be possible to do). Why is such a transfer (if feasible) a good idea? By shifting payoff into the future, it raises agent continuation payoffs and therefore improves incentives, allowing for an increase in efficiency.

Two relatively minor problems stand in the way. First, to guarantee an improvement in incentives, the post-deviation continuation value cannot change in arbitrary ways as transfers are made; this is the role of assumption (A.3). Moreover, there must be a rich set of dates for which the incentive constraint is binding; this is guaranteed by Lemma 3 below.

The major problem is that such a transfer may be prohibitively costly, because payoff functions are generally nonlinear in the transfer. It turns out that (when contemplating small changes), the appropriate variable to look at is the ratio of the agent's marginal utility to that of the principal's marginal utility in the transfer—for any date t, call this $\psi_t \equiv -(\partial A(\xi_t)/\partial m)/(\partial P(\xi_t)/\partial m)$. With transferable payoffs this ratio is independent of time. With nonlinearities, its movement cannot be controlled a priori.

The solution is a proof by contradiction: suppose that the principal's continuation payoffs remain strictly positive at every date. Then the ratio ψ_t cannot rise over any two points of time.[15] For if it could, a small current sacrifice by the agent can be more than paid back in the future, and the "standard" backloading argument goes through with no difficulty. This contradicts the presumption that the original agreement sequence is ESE.

So ψ_t cannot increase, but this observation only holds for a subsequence of periods, namely for those in which the limited liability condition is not binding for the agent, so that the required current sacrifice can be made. In the formal proof, Lemma 4 guarantees that there is indeed such a set of periods, and Lemmas 5 and 6 show that ψ_t is indeed declining over this set of dates.

With this insight, we can travel far enough out on the subsequence of the previous paragraph so that ψ_t is almost at its limit point. Here the transferable payoff argument can be mimicked. There are possible losses but these are of second order in the transfer. The subtlety of the argument now lies in showing that the incentive gains are of first order in the transfer, and the proof is complete.

A small technical query: is the theorem true for all WESE sequences? Without further assumptions, the answer is generally no. However, the difference is truly mild. For instance, if the ESE payoff frontier (viewed as a map from π to α) is continuous at $\pi = 0$, then the theorem extends without difficulty (footnote 23 in Section 7 fills in the details).

Figure 2 summarizes the assertion of Theorem 1 in graphical form. Notice, however, that the theorem does *not* assert that all ESE sequences settle down to some unique, infinitely repeated agreement. In fact, whether reasonable additional assumptions can be made (while still permitting some nontransferability)

[15] This is not entirely precise. One also needs to ensure that the incentive constraints are binding, which is taken care of.

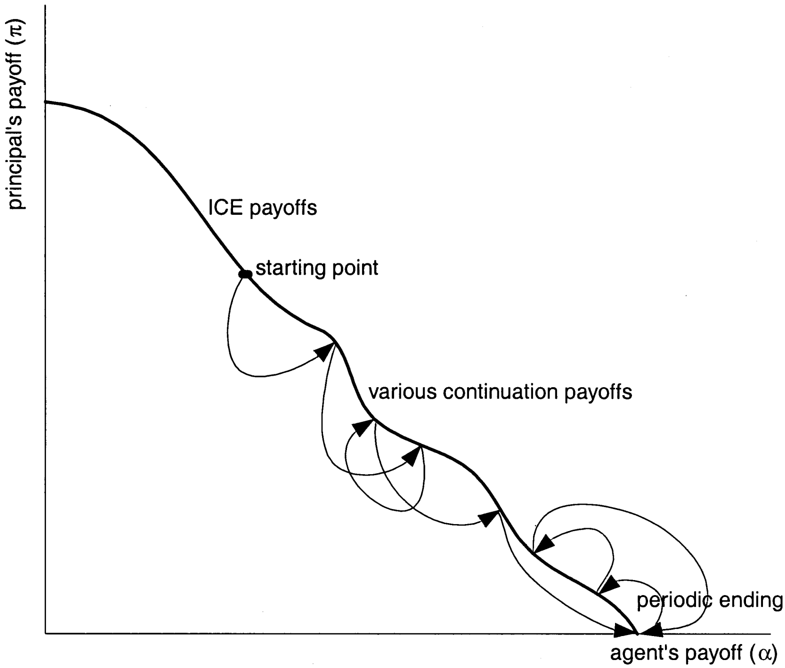

FIGURE 2.—ESE sequences converge to the agent's best self-enforcing sequence.

that will guarantee this, remains an open question. I am, however, able to say a bit more about the structure of the agent's best self-enforcing sequence.

THEOREM 2: *There exists an agent's best self-enforcing sequence* $\{\xi_t^*\}$ *that is periodic*; *that is*, *there exists a date* N *and a finite collection of agreements* $\{\xi_0, \ldots, \xi_N\}$ *such that for all* t, $\xi_t^* = \xi_i$, *where* $i = \text{Remainder}(t/N+1)$.

Moreover, for every ESE payoff vector, there exists an ESE payoff sequence such that after some finite number of dates, *the above periodic sequence is started up*.

Figure 2 summarizes this as well. The intuitive discussion following Theorem 1 outlines the approach to a proof; we showed there that every ESE sequence must—after some finite date—hit a point where the principal's continuation payoff is zero. This holds, in particular, for the agent's best self-enforcing sequence, which is also ESE. That there exists some best sequence that is periodic is now an easy consequence of this last observation.

4. TRANSFERABLE PAYOFFS

At this point, it is an open question whether any additional structure can be gleaned from the general framework. I therefore turn to the case of transferable payoffs, and place more restrictions on the model.

We specialize the payoff functions by imposing quasi-linearity. That is, we assume that $A(\xi) = A(c, m) = A(c) + m$ and $P(\xi) = P(c, m) = P(c) - m$. I realize that there is some abuse of notation by using the same functional notation A and P, but this is perhaps compensated for by mnemonics and continuity.

In this section, we regard limited liability as tantamount to some lower bound on the agent's total payoff. That is, $L(c, m) = A(c) + m - \lambda$, so that $A(c, m) \geq \lambda$ for feasibility.

A minor technical assumption on the post-deviation continuation value simplifies the analysis significantly, though it is possible to make progress without it (details excluded).

(A.6) *There are no two agreements* ξ *and* ξ' *such that* $\partial V(\xi)/\partial m = 0$ *and simultaneously* $\partial V(\xi')/\partial m = 1$.

If V is also assumed to be linear in m—and given transferable payoffs this may not be entirely objectionable—then (A.7) is automatically satisfied.[16]

Further assumptions will be made below.

4.1. *The Agent's Best Stationary Problem*

The following artificial problem plays a key role. This is the problem of finding the agent's best solution in the class of *stationary* self-enforcing agreements. That is,

$$\text{maximize}_{(c,m)} A(c) + m \tag{17}$$

subject to

$$V(c, m) \leq (1-\beta)^{-1}[A(c) + m] \tag{18}$$

and

$$P(c) - m \geq 0. \tag{19}$$

We assume the following.

(A.7) *There is a unique solution* (c^*, m^*) *to the agent's best stationary problem, and the agent's per-period payoff at the solution strictly exceeds* λ (*the lower bound on agent payoff given by limited liability*).

Obviously, the first part of (A.7) can be easily derived from more primitive assumptions on the curvature of the A, P, and V functions. The second part is self-explanatory and constitutes an obviously weak restriction.

The examples in Section 2 satisfy all the assumptions made so far.

[16] In general, something stronger than linear utility is needed to guarantee that V will also be linear; for V summarizes not just payoffs, but also incorporates the agent's ability to walk away with transfers promised her by the principal (or to evade transfers that she may have promised to make). There is no reason why this should automatically precipitate linearity even if overall payoffs are linear.

4.2. *More on the Structure of ESE Sequences*

The following results can now be proved.

THEOREM 3: *Let $\{\xi_t\}$ be an ESE agreement sequence. Then there exists a finite date T such that for all $t \geq T$,*

$$\xi_t = (c^*, m^*), \tag{20}$$

where (c^*, m^*) is the solution to the agent's best stationary problem.

That is, *every* ESE sequence must end in the same way. If we interpret these different sequences as emanating from the different bargaining powers of the two individuals, the result shows that "in the long run," different bargaining power has no effect. Thus our prediction is very sharp indeed.

Transferable payoffs also permit a tighter characterization of the nonstationary initial phase of an ESE sequence.

THEOREM 4: *For any ESE sequence, consider the "initial phase" up to date $T-1$, before the agent's best self-enforcing stationary sequence (c^*, m^*) is applied. Then the limited liability constraint binds throughout this phase (except possibly at $T-1$), and the agent's present-value payoff rises throughout:*

$$\alpha_t < \alpha_{t+1} \tag{21}$$

for all $t = 0, \ldots, T-1$.

At the same time, the principal's current payoff increases through most of this phase, before it drops to zero at the onset of the agent's best self-enforcing sequence:

$$P(c_t) - m_t < P(c_{t+1}) - m_{t+1} \tag{22}$$

for $t = 0, \ldots, T-2$ (but not necessarily $T-1$).

Given the intuition of the analysis so far, it is not surprising that, with transferable utility, the agent passes through a phase where the limited liability condition is binding. In other words, the ESE agreements are backloaded to the maximal extent possible, and this is borne out in (21). Perhaps a bit more subtle is the fact that the *principal's* one-shot payoff must climb steadily through this phase, before finally dropping to zero once the tail sequence takes over.[17] This finding also implies that *total* additive surplus rises through the initial phase.

Conditions (21) and (22), taken together, seemingly contradict ESE. But this is illusory: despite (22), the *continuation* value to the principal, π_t, must steadily fall through the initial phase.

[17] That these results do not necessarily hold at the penultimate point $T-1$, just before the agent's best self-enforcing agreement takes over, stems from an integer problem with discrete time.

I end this section by indicating how the theorems can be applied to derive the intertemporal structure asserted for the examples in Section 2. Of course, Theorem 3 immediately implies the main observation for both these examples: ultimately the agent's best *stationary* self-enforcing sequence must be followed.

To apply Theorem 4, note the following points. By assumption, the first-best number of hours (in Example 1) or the first-best loan size (in Example 2) cannot be supported. In particular, these variables must lie below their first-best values, for higher values are doubly cursed with lower payoffs *and* higher deviation gains.

Now it is easy to see why hours (and loan size) must steadily rise over the initial phase. As discussed, Theorem 4 tells us that *total* additive surplus—a well-defined concept in the transferable payoffs case—must climb through the initial phase. Because the two examples have a convex structure, and given the observations of the previous paragraph, this can only happen if work hours (and loan size) rise steadily over the initial phase.

5. COMMENTS, CONNECTIONS AND EXTENSIONS

I make several remarks on these results, some in the form of possible extensions and generalizations.

5.1. *Credibility*

Suppose that the principal may have the opportunity to contract with several identical agents, starting in each case from the same point on the constrained frontier. As soon as continuation payoffs depart in an adverse way from this starting point (which they *must*, by our findings), why can't the principal simply fire this agent and replace him with another?

It may as well be stated up front that if this option is available *without cost* to the principal, then all ESE sequences must be stationary. However, the time drift identified in this paper makes an immediate reappearance as soon as one or more of the following conditions are satisfied.

(i) *The principal's actions, and those of the agent, are observed by other agents.* In this case a broader notion of equilibrium, which permits future agents to condition on this history, would lead to the same conclusions as the ones obtained here.[18] To be sure, new equilibria may emerge in which there is no drift (no-drift equilibria exist in the original model as well). But the *efficient* self-enforcing outcomes must exhibit exactly the same characteristics as those identified in this paper.

One might argue—correctly—that (perfect) observability of actions is too strong an assumption. In that case, the equilibria to be constructed are more

[18] In such equilibria, future agents must punish the principal for deviating from the "agreed-upon" path. Notice, by the way, that the assumption of observability does not necessarily imply the ability to commit. No court of law may be available, or even if one is, verifiability is different from observability.

complex; the efficient among these will surely involve statistical tests to figure out whether the principal is systematically deviating, or whether it is the agent who is responsible for breaking the relationship. I conjecture that such equilibria will also exhibit intertemporal drift towards the agent's best self-enforcing outcome, though the drift may stop short of actually getting to that outcome.

(ii) *It is costly for the principal to find or train a new agent.* In this case, a certain degree of agent-directed drift will be automatically "tolerated" by the principal because it is simply too costly to find a replacement. The larger the cost, the closer we come to a full replication of the results here. Both search and training costs may be significant (I return to this issue in a later subsection).

5.2. *Repeated Bargaining*

The careful reader will have noticed that the same bargaining power is not applied to equilibrium selection from *subgames*. Bargaining power is only used to pin down date-zero payoffs. It may be argued that this approach falls prey to renegotiation, or something that resembles it in ex-post equilibrium selection.

Let us get to the bottom line of this argument: If equilibrium selection demands that the same bargaining power be applied period after period, then only equilibrium outcomes are stationary, and there can be no drift of continuation values in the direction of the person who faces the incentive constraint. In particular, long-term payoffs *do* depend on bargaining power, and the results of this paper are invalidated.

Whether or not this form of continued negotiation *must* be imposed on the problem is, however, another matter. Notice, in passing, that our schemes are always constrained efficient, so that they are immune to the standard form of renegotiation that looks for (constrained) *Pareto*-improvements. Thus repeated renegotiation really refers to ex-post movements along the Pareto frontier—in the direction of the principal—and away from agreed-upon continuations. This relates the "repeated bargaining" issue to the credibility problems raised in the previous section. A principal who deals with many agents—either sequentially or simultaneously—may face serious reputational consequences from engaging in such behavior.

5.3. *Nonbinding Enforcement Constraints*

A fundamental assumption, (A.4), of this paper is that for every ESE agreement sequence, incentives do matter at some date (and consequently for an infinite subsequence of dates). This assumption may well be violated for certain games when the discount factor is close enough to unity. In that case it is easy enough to see that after a finite number of dates, an ESE agreement sequence must be fully (i.e. unconstrained) efficient. In this case we will not, in general, obtain convergence of all ESE sequences to the agent's best self-enforcing outcome.

For instance, in the case of transferable payoffs, suppose that the no-deviation constraint does not bite at the agent's best self-enforcing payoff. Then a small shift of bargaining power to the principal will not change the "nonmonetary terms" of the agreement, which is fully efficient. It will only change the monetary transfer at each date. Thus the full agreement (with transfer included) is distinct, and remains so over time.

This suggests a more general—yet weaker—description of the main result: every ESE sequence is, after a finite number of dates, *either* unconstrained efficient, *or* follows the agent's best self-enforcing sequence.[19] In this more qualified characterization, the assumption (A.4) may be dropped. Because I wish to shift the focus away from folk-theorem-like unpredictability, I choose to emphasize the stronger conclusion, obtained with stronger assumptions.

5.4. *Several Individuals*

I sketch one generalization of the results in this paper that perhaps opens other avenues of research. First suppose that there are n individuals instead of two; and assume that of these n agents, a subset m face incentive constraints due to limited enforcement. I conjecture that every constrained efficient path must have the property that after some finite date, it must be constrained efficient for the subset of m agents. The same argument, with some variations, will go through.

This sort of generalization draws our attention to the possibility that "facing or not facing an incentive constraint" may not be a 0-1 feature. *Any* of the individuals may face a binding incentive constraint if pushed too far. One simple way to model this is to actually posit a perfectly symmetric game with n agents, but include agent-specific additive costs for breaking a relationship. How does the extent of time-drift relate to the variation in these costs?

These research directions also serve to tie up our model with a recent literature on mutual insurance with enforcement constraints. Because economic reciprocity is of particular importance in developing countries (see, e.g., Townsend (1993) and Udry (1994)), both development economists and theorists have made recent efforts to model this phenomenon (see, e.g., Coate and Ravallion (1993), Fafchamps (1996), Kletzer and Wright (1996), Kocherlakota (1996), and Ligon, Thomas, and Worrall (2001)). One insight that emerges from (some of) these papers is that provided the first-best insurance system cannot be maintained, the second-best scheme has credit-like characteristics. That is, a person for whom the incentive constraint is binding today is "rewarded" with a higher continuation payoff. Because a binding incentive constraint today arises from the fact that the person has to make a transfer to others, and because a higher continuation payoff means that (on average) he receives more transfers at a later date, credit-like features enter the scheme.

Notice that the incentive constraints here are assigned stochastically—to the person who has received good output relative to others and so must pay up. At

[19] I am indebted to an anonymous referee for suggesting this observation.

the instance of that shock, future continuation values drift in the direction of the person concerned. If all individuals are perfectly symmetric, however, there is no long-term drift in any preassigned direction.

The results of our paper—and the n-person generalizations thereof—suggest that the observations for mutual insurance can be extended to the case in which there is asymmetry among individuals. For instance, some people in an insurance group may find it easier to migrate than others. Out approach suggests that long-term drift towards certain individuals will now appear—perhaps stochastic in nature. These are precisely those individuals who find it easiest *not* to comply with a scheme of reciprocity.

5.5. *Heterogeneous Discount Factors*

Lehrer and Pauzner (1999) consider repeated games in which the players possess different discount factors. In their paper, both efficiency and constrained efficiency require that the patient player be rewarded later, while the impatient player is rewarded earlier. I have not explored an extension to this case, but it is clear that there are interesting connections. Recall my informal discussion of the arguments underlying Theorem 1. Two factors were crucial: the incentive constraint, of course, and the (relative) marginal evaluations of monetary transfers across principal and agent. When discount factors are different, the latter must be weighted to allow for variable discounting, and the interplay between these two factors becomes nontrivial and interesting. The interest is heightened by the following interpretation: suppose that the agent is poor relative to the principal (she is an employee, or tenant, or borrower). Then it is not unreasonable to postulate that the agent is more impatient than the principal. But in that case, the Lehrer-Pauzner findings and the results of this paper tug in different directions. It may be worth exploring if one of the two factors always dominates.

5.6. *Endogenous Outside Options*

Consider a large number of identical principal-agent relationships, in which a particular principal-agent pair is created by random matching. Suppose, moreover, that on the termination of an ongoing relationship, both principal and agent are returned to their pools for future rematching. With rematching, assume that a fresh relationship is started up with no knowledge of the past.[20]

Now, the costs of termination of a particular relationship can be modeled in a variety of ways. One simple way to do so is to postulate an additive cost for the agent, which proxies some one-time sanction or other punishment that can

[20] To make this precise requires further description. The idea is to avoid previous matches. One can do this in several ways. For instance, suppose that there is a continuum of each type of agent, and that there is a fresh inflow of a positive measure of individuals at each date. Then rematching will involve a new partner with probability one. Or one can assume that existing relationships break apart for exogenous reasons, which (along with the continuum postulate) will also guarantee that the same partner is not encountered twice.

be imposed on him before rematching occurs. Under this interpretation, we can proceed as follows.

Let $\{\hat{c}_t, \hat{m}_t\}$ be the "going" ESE sequence in a large number of identical bilateral relationships. Then for any *particular* relationship, we may write the outside option for the agent as

$$V(c, m) \equiv D(c, m) + \beta\hat{\alpha}_0 - \kappa_a, \tag{23}$$

where $D(c, m)$ represents the one-time payoff from a deviation, κ_a a present-value cost of deviation (e.g., sanctions), and $\beta\hat{\alpha}_0$ the continuation payoff from rematching. Notice that $\hat{\alpha}_0$ is just the starting payoff from the sequence $\{\hat{c}_t, \hat{m}_t\}$, which we assume pervades the economy at large.

To be sure, the possibility of rematching also affects the principal's participation constraint. We must dispense with the zero-normalization and rewrite this as

$$\pi_t \geq \hat{\pi}_0 - \kappa_p, \tag{24}$$

where κ_p is the cost of acquiring a new relationship for the principal, and $\hat{\pi}_0$ is the payoff if she does so. As in the case of the agent, $\hat{\pi}_0$ is the starting payoff from the "going" sequence $\{\hat{c}_t, \hat{m}_t\}$.

We may now think of some agreement sequence $\{\hat{c}_t, \hat{m}_t\}$ as a *general equilibrium sequence* if it is the chosen sequence in any specific bilateral problem, when the outside options and participation constraints are described by (23) and (24).

This extension is of interest because it has definite predictions regarding what might be called "market selection" in the face of a changing informational environment. For instance, with economic development, mobility increases significantly, leading to a decline in information of the sort needed to sustain repeated transactions. It may be harder to place sanctions on a deviating agent. And it may be harder to find an agent that the principal considers trustworthy.

The penultimate sentence in the previous paragraph may be translated thus: κ_a falls with the increased mobility that development brings.[21] It is harder to translate the last sentence into our model, for it is essentially based on incomplete information regarding agent characteristics.[22] But if we permit a certain degree of expository license, this can be proxied by an increase in κ_p (recall equation (24)).

What are the effects of these changes? Imagine that there are a large number of markets, each peopled by a large number of principal-agent pairs. While the allocation of bargaining power is unchanged *within* a market, allow it to vary fully *across* markets. Now, if κ_a falls (and κ_p rises), markets in which the agent has all the bargaining power (or close to it) are likely to collapse. The reason is that enough punishment cannot be imposed on a deviating agent—the continuation payoff is very close to the payoff from which the agent is deviating. If, on the other hand, the principal has all the bargaining power, we have seen

[21] To be sure, this cannot be true permanently. For instance, in developed countries, credit histories can be effectively tracked on computer networks.

[22] See Ghosh and Ray (1996), Kranton (1996), and Watson (1999) for analyses along these lines.

that the agreement path involves increasing continuation value to the agent, as in Theorem 4. This means that a defaulting agent will have to return to the very beginning of an improving path when a new relationship is sought. If follows that such markets are more likely to be robust to a fall in κ_a.

This extension suggests that as informal markets cease to function under the progressive decay of information, only those markets with relatively high bargaining power to the principal are likely to survive. This sort of “market-selection” argument may be worth exploring in future research.

5.7. *The Effects of Changed Bargaining Power*

Finally, my implicit focus on variations in bargaining power has its roots in literature on economic development. That two parties enter into a relationship unequally matched has often been the subject of inquiry in development economics (whereas in other subjects it is more customary to impose the zero-profit condition on one side or the other, implying certain “corner” selections from the constrained frontier). Recently, this sort of analysis has received some attention in the work of Mookherjee (1997), Banerjee, Gertler, and Ghatak (1997), Banerjee, Mookherjee, Munshi, and Ray (2001), Mookherjee and Ray (2001), and others. The main insight is that a change in bargaining power does affect the total surplus (and certainly the chosen contract) generated by a relationship (of course, concepts such as total surplus having meaning only in a transferable-utility situation). The general result of this model states—in contrast—that in dynamic situations variations in bargaining power may not be observable once the relationship has travelled far enough in time. But in obtaining this result it relies on the observations of Mookherjee (1997) and others that total surplus increases as bargaining power shifts to the agent who is incentive-constrained. It is precisely because of this that dynamically efficient contracts shift (over time) in favor of the agent, though the fact that we obtain this result for general nontransferable-payoff situations suggests that there are even more basic forces at work (which do not rely on concepts of total surplus).

6. FINAL REMARKS

The results in this paper appear to be significant for the following reasons:

(i) They reveal the striking generality of the observation that long-term relationships ultimately shift in favor of the individual whose incentive constraints are “binding.” This occurs even when payoffs cannot be costlessly transferred across agents (that is, even when the payoff functions are nonlinear in some transfer variable such as money).

(ii) At the same time, our findings imply that the extent of this shift (in terms of the observable movement of agreements) relies strongly on the allocation of bargaining power. If the agent has all the power to begin with, the structure of agreements has no discernable time trend. In the most general scenario, one can always prove that when the agent has all the bargaining power, there exists a periodic equilibrium sequence. With transferable payoffs, this sequence *must* be stationary. Thus the time-shifts are entirely due to the fact that the principal has some of the bargaining power to start with.

(iii) These results also have empirical implications over and above the predicted time paths. They suggest that if one wishes to empirically study the allocation of "power" in relationships such as repeated credit, tenancy, or employment, it is important to have data over the duration of the relationship. In particular, observation of the "mature phase" alone will say little or nothing to answer this question.

7. PROOFS

In what follows, the assumptions (A.1), (A.2), (A.3), (A.4), and (A.5) are taken to hold throughout. In the proofs that follow, the notations ξ and (c, m) are used interchangeably, the latter being employed whenever I feel it is useful to stress the two components of the agreement.

LEMMA 1: *There exists* $-\infty < m < M < +\infty$ *such that if* $\{c_t, m_t\}$ *is any WESE agreement sequence, then* $m \leq m_t \leq M$ *for all* t.

PROOF: Given that $L(c_t, m_t)$ must be nonnegative, noting that c lies in some compact set, and recalling that L is continuous in (c, m) with positive derivative bounded away from zero in m, it is easy to see that there exists $m > -\infty$ such that $m_t \geq m$ for all t. Noting that the derivative of P with respect to m is negative, this places an upper bound on the possible value of $P(c_t, m_t)$ at any date; call it $\overline{P}$. Using (16), it must be the case that at any date, $P(c_t, m_t) \geq -\beta(1-\beta)^{-1}\overline{P}$. That is, P must be uniformly bounded below. Using this information along with the assumptions that P is continuous in (c, m) with negative derivative bounded away from zero in m, we may conclude that there exists $M < \infty$ such that $m_t \leq M$ for all t. *Q.E.D.*

In what follows, define $G(\xi) \equiv V(\xi) - A(\xi)$ for each $\xi \in \Xi$.

Fix some feasible agreement ξ and some $\theta > 0$. We shall say that ξ can be θ-*relaxed* if for each $\nu > 0$ and sufficiently small, there is a feasible agreement ξ' such that

$$G(\xi') - G(\xi) \leq \nu, \tag{25}$$

and

$$\min\{A(\xi') - A(\xi), P(\xi') - P(\xi)\} \geq \theta\nu. \tag{26}$$

LEMMA 2: *For each* $\theta > 0$, *there exists* $\delta > 0$ *such that if a feasible agreement* ξ *cannot be* θ*-relaxed, it cannot be* δ*-improved. Moreover,* $\delta \to 0$ *as* $\theta \to 0$.

PROOF: By (A.3), we have that $G = V - A$ is Lipschitz; let $\lambda < \infty$ be the Lipschitz bound. Given $\theta > 0$, define $\delta \equiv \theta\lambda$. Suppose that a feasible agreement ξ cannot be θ-relaxed but it can be δ-improved. Then for all $\nu' > 0$ and sufficiently small, there exists feasible ξ' within the ν'-ball of (ξ) such that

$$\min\{A(\xi') - A(\xi), P(\xi') - P(\xi)\} \geq \delta\nu' = \theta\nu,$$

where $\nu \equiv \nu'\lambda$. At the same time,

$$G(\xi') - G(\xi) \leq \lambda\nu' = \nu.$$

These arguments imply that ξ can be θ-relaxed, a contradiction.

Note that the way we have defined δ in this proof allows us to immediately infer the second part of the lemma. *Q.E.D.*

LEMMA 3: *Let $\{\xi_t\}$ be a WESE agreement sequence. Then there exists $\theta > 0$ and an infinite set of dates t such that ξ_t can be uniformly θ-relaxed.*

PROOF: Note that the continuation of a WESE sequence is trivially WESE. With this in mind, define an infinite collection of WESE agreement sequences $\{\xi_t^n\}$ by

$$\xi_t^n = \xi_{n+t}$$

for all nonnegative integers n and t. Given Lemma 1 and the fact that the component c of an agreement is drawn from a compact space, we see that ξ_t^n lies in some compact set for each n and t. Consequently, by a standard diagonal argument, we may assert the existence of a sequence ξ_t^* and a subsequence $n(k)$ of n such that

$$\lim_{k\to\infty} \xi_t^{n(k)} = \xi_t^*$$

for each date t.

I claim that $\{\xi_t^*\}$ is a WESE agreement sequence.

Notice that by the continuity of L, $L(\xi_t^*) \geq 0$ for all t. Next—because ξ_t^n lies in some compact set that's uniform in n and t—it is easy to see that continuation payoffs converge over the subsequence of agreement sequences:

$$\alpha_t^{n(k)} \to \alpha_t^* \quad \text{and} \quad \pi_t^{n(k)} \to \pi_t^*$$

for each t, as $k \to \infty$. From these observations and the assumed continuity of the post-deviation continuation value V, we may conclude that $\{\xi_t^*\}$ satisfies (15) and (16) at all dates. Finally, because the agreement sequences $\{\xi_t^{n(k)}\}$ are weakly efficient for all k, it is easy to check that the same must be true of the limit $\{\xi_t^*\}$.

Suppose, now, that the lemma is false. Then for every $\theta > 0$, there is a date T such that for all $t \geq T$, no ξ_t can be θ-relaxed. Applying this information to our construction above, we may deduce that for each t, there exists a sequence $\theta_k \downarrow 0$ such that for all k, $\xi_t^{n(k)}$ cannot be θ_k-relaxed. It follows from Lemma 2 that there exists a sequence $\delta_k \downarrow 0$ such that for each t and k, $\xi_t^{n(k)}$ cannot be δ_k-improved. Using (A.2), we may conclude that there is some sequence $\epsilon_k \downarrow 0$ such that for all k, $\xi_t^{n(k)}$ is *not* ϵ_k-improvable. Therefore the limit point ξ_t^* cannot be ϵ-improvable for any $\epsilon > 0$. Consequently ξ_t^* must be unimprovable.

We have therefore shown that $\{\xi_t^*\}$ is WESE *and* wholly composed of unimprovable agreements. This contradicts (A.4). *Q.E.D.*

LEMMA 4: *Let $\{\xi_t\}$ be a WESE agreement sequence. Then the collection of all dates k such that $L(\xi_k) > 0$ is an infinite set (call it K). Moreover, there is a uniform bound Γ on the distance between adjacent dates in K.*

PROOF: Recall that the continuation of any WESE sequence must itself be WESE. Applying (A.5), it follows that K (as defined in the statement of the lemma) must be an infinite set.

We now prove the second part of the lemma. Suppose that there is no uniform bound as claimed between adjacent members of K. Then there exists an infinite ordered subset of K, call it $\{k(n)\}$, $n = 1, 2, \ldots$, such that $k'(n) - k(n) \uparrow \infty$, where $k'(n)$ is the smallest index in K larger than $k(n)$.

Now define an infinite collection of WESE agreement sequences $\{\xi_t^n\}$ by

$$\xi_t^n = \xi_{k(n)+t}$$

for all nonnegative integers n and t. By the same diagonal argument as in the proof of Lemma 3, we may assert the existence of a sequence $\{\xi_t^*\}$ and a subsequence $n(s)$ of n such that

$$\lim_{s\to\infty} \xi_t^{n(s)} = \xi_t^*$$

for each date t. By the same arguments as in the proof of Lemma 3, $\{\xi_t^*\}$ is a WESE agreement sequence.

Pick any date t. Pick an index N such that $k'(n) - k(n) > t$ for all $n \geq N$. Then for such n, we have

$$L(\xi_t^n) = 0.$$

It follows from the continuity of L that at the limit point ξ_t^*,

$$L(\xi_t^*) = 0.$$

Since t is arbitrary and $\{\xi_t^*\}$ is WESE, we contradict (A.5). *Q.E.D.*

Let $\{\xi_t\}$ be a sequence of agreements. For each t, define

$$\psi_t \equiv -\frac{\partial A(\xi_t)/\partial m}{\partial P(\xi_t)/\partial m}. \tag{27}$$

Notice that by assumption (A.1), ψ_t is well-defined and positive for all t, and bounded away from zero in t.

In what follows, I will use the generic form $o(z)$ to denote any functional expression such that $o(z)/z \to 0$ as $z \to 0$. I will also use the notation $A'(\xi)$ (resp. $P'(\xi)$) to denote the partial derivative of A (resp. P) with respect to the m component.

LEMMA 5: *Let $\{c_\tau, m_\tau\}$ be a sequence of agreements. Consider any date s for which $L(\xi_s) > 0$ and any date $t > s$ such that $\pi_k > 0$ for all dates $k = s+1, \dots, t$. Then for all $\epsilon > 0$ and sufficiently small, there exists $\epsilon' > 0$ such that*

$$L(c_s, m_s - \epsilon) > 0, \tag{28}$$

$$\pi_k > \beta^{t-k}[P(c_t, m_t) - P(c_t, m_t + \epsilon')] \tag{29}$$

for all $k = s+1, \dots, t$,

$$P(c_s, m_s - \epsilon) + \beta^{t-s} P(c_t, m_t + \epsilon') = P(c_s, m_s) + \beta^{t-s} P(c_t, m_t), \tag{30}$$

and

$$\begin{aligned}[A(c_s, m_s - \epsilon) + \beta^{t-s} A(c_t, m_t + \epsilon')] - [A(c_s, m_s) + \beta^{t-s} A(c_t, m_t)] \\ = (\psi_s - \psi_t) P'(c_s, m_s)\epsilon + o(\epsilon).\end{aligned} \tag{31}$$

PROOF: For $\epsilon > 0$ and sufficiently small, it is trivial to see that there exists $\epsilon' > 0$ such that (28), (29), and (30) are satisfied. Note that

$$\begin{aligned}&[A(c_s, m_s - \epsilon) + \beta^{t-s} A(c_t, m_t + \epsilon')] - [A(c_s, m_s) + \beta^{t-s} A(c_t, m_t)] \\ &= -A'(c_s, m_s)\epsilon + \beta^{t-s} A'(c_t, m_t)\epsilon' + o(\epsilon) + o(\epsilon') \\ &= \psi_s P'(c_s, m_s)\epsilon - \beta^{t-s}\psi_t P'(c_t, m_t)\epsilon' + o(\epsilon) + o(\epsilon') \\ &= \psi_t[P'(c_s, m_s)\epsilon - \beta^{t-s} P'(c_t, m_t)\epsilon'] \\ &\quad + (\psi_s - \psi_t) P'(c_s, m_s)\epsilon + o(\epsilon) + o(\epsilon').\end{aligned} \tag{32}$$

Moreover, using (30),

$$P'(c_s, m_s)\epsilon - \beta^{t-s} P'(c_t, m_t)\epsilon' = o(\epsilon) + o(\epsilon'). \tag{33}$$

Combining (32) and (33), we may conclude that

$$\begin{aligned}[A(c_s, m_s - \epsilon) + \beta^{t-s} A(c_t, m_t + \epsilon')] - [A(c_s, m_s) + \beta^{t-s} A(c_t, m_t)] \\ = (\psi_s - \psi_t) P'(c_s, m_s)\epsilon + o(\epsilon) + o(\epsilon').\end{aligned} \tag{34}$$

Now observe that ϵ and ϵ' are of the same order (because the derivatives of P with respect to m are bounded above and below), so that $o(\epsilon) \simeq o(\epsilon')$. Using this information in (34), we obtain (31). *Q.E.D.*

LEMMA 6: *Let $\{\xi_\tau\}$ be a WESE sequence of agreements. Consider any date s for which $L(\xi_s) > 0$ and any date $t > s$ such that $\pi_k > 0$ for all dates $k = s+1, \ldots, t$. Then $\psi_s \geq \psi_t$.*

PROOF: Suppose not. Then for some s and t as described in the statement of the lemma, we have $\psi_s < \psi_t$. Pick $\epsilon > 0$ and $\epsilon' > 0$ so that all the conditions of Lemma 5 are satisfied.

Construct a new sequence of agreements $\{\xi'_\tau\}$ by only changing the monetary transfers at dates s and t: specifically, $m'_s = m_s - \epsilon$, $m'_t = m_t + \epsilon'$, and m_τ unchanged elsewhere. Do this in a way such that (28)–(31) are satisfied. I claim that this sequence is self-enforcing. By (28), limited liability is satisfied. By (29) and (30), the principal's participation constraint (16) is satisfied.

To check the agent's constraint (15), apply the presumption that $\psi_s < \psi_t$ to (31), keeping in mind that $P'(c_s, m_s) < 0$. It is easy to see (when ϵ is small enough) that at all dates, the agent's continuation payoffs have either gone up (for instance, strictly so at date t) or have remain unchanged. Thus we only need to examine the (possibly changed) post-deviation continuation value at dates s and t. At date s, using (A.3),

$$V(c_s, m_s - \epsilon) \leq V(c_s, m_s) \leq \alpha_s < \alpha'_s \tag{35}$$

so (15) is met (and is in fact met strictly, because $\alpha'_t > \alpha_t$). At date t, using (A.3) again,

$$V(c_t, m_t + \epsilon') - V(c_t, m_t) \leq A(c_t, m_t + \epsilon') - A(c_t, m_t),$$

so (15) must be satisfied once again. Therefore $\{\xi'_\tau\}$ is self-enforcing. Notice that

$$\alpha'_0 > \alpha_0 \quad \text{while} \quad \pi'_0 = \pi_0. \tag{36}$$

Now perturb the new sequence slightly by reducing m'_t further by a tiny amount. Because (35) holds strictly, it will continue to hold following the perturbation. Thus the new perturbation also creates a self-enforcing sequence. But now, using (36), *both* principal and agent can be made strictly better off relative to (α_0, π_0). This contradicts our starting point that $\{\xi_\tau\}$ is WESE. *Q.E.D.*

PROOF OF THEOREM 1: Let $\{\xi_\tau\}$ be an ESE agreement sequence. Obviously, it suffices to show that $\pi_t = 0$ at some date t, because an agent's best self-enforcing sequence must be followed thereafter.[23]

[23] This is the point at which the theorem may fail to extend to all WESE sequences. However, if the ESE payoff frontier is continuous at $\pi = 0$, then the fact that $\pi_t = 0$ at some date t once again implies that an agent's best self-enforcing sequence must be followed thereafter. For if not, continuation payoffs at this point are of the form $(\alpha_t, 0)$, where $\alpha_t < \alpha^*$, the agent's best self-enforcing payoff. At the same time, the continuity assumption implies that there is some self-enforcing payoff $(\alpha, \pi) \gg (\alpha_t, 0)$, which in turn means that $(\alpha_t, 0)$ cannot be a WESE payoff. But this is a contradiction, for the continuation of a WESE sequence is trivially WESE.

Suppose, on the contrary, that $\pi_\tau > 0$ for all τ. Let the set K be given as in Lemma 4. Because $\pi_\tau > 0$ for all τ, Lemma 6 is applicable to the entire sequence of dates in K, so that

(37) $$\psi_s \geq \psi_t$$

for all dates s and t such that $s < t$ and $s, t \in K$. Lemma 4 provides us with a uniform bound Γ on the gap between adjacent dates in K. Pick $\theta > 0$ such that Lemma 3 is applicable to the sequence $\{\xi_\tau\}$, and define

(38) $$\mu \equiv \beta^{2\Gamma}\theta(q^2/Q),$$

where q and Q are positive lower and upper bounds respectively on the absolute values of the derivatives of A and P with respect to m. Choose $\eta > 0$ such that

(39) $$\mu - \eta Q > 0.$$

Because Lemma 3 applies, ξ_τ can be θ-relaxed for an infinite set of dates. This allows us to pick two dates s and t in K satisfying the following conditions:

(40) $$\psi_t \geq \psi_s - \eta$$

and

(41) $$T - \Gamma \leq s < T < t \leq T + \Gamma$$

for some date T such that ξ_T can be θ-relaxed.

These choices are made possible by (37)—so that ψ_k converges over K—and the existence of the uniform bound Γ asserted in Lemma 4.

Now choose $\epsilon > 0$ and $\epsilon' > 0$ so that all the conditions of Lemma 5 are satisfied. We will create a new agreement sequence $\{\xi''_\tau\}$ in two stages. First, we move to a sequence $\{\xi'_\tau\}$ by imitating the construction in the proof of Lemma 6, which is to only alter monetary transfers at dates s and t: specifically, $m'_s = m_s - \epsilon$, $m'_t = m_t + \epsilon'$, and m_τ unchanged elsewhere.

After this first set of changes, note that (agent) continuation payoff viewed *at* date T has increased by

(42) $$\begin{aligned} \alpha'_T - \alpha_T &= \beta^{t-T}[A(c_t, m_t + \epsilon') - A(c_t, m_t)] \\ &\geq \beta^\Gamma q\epsilon' \\ &\geq \beta^\Gamma (q^2/Q)\epsilon, \end{aligned}$$

where the last inequality follows from the fact that $\epsilon' \geq (q/Q)\epsilon$ is a necessary condition for (30).

Moreover, since $\xi_T = \xi'_T$ can be θ-relaxed, there exists feasible ξ''_T such that (for ϵ positive and sufficiently small),

(43) $$G(\xi''_T) - G(\xi'_T) \leq \beta^\Gamma (q^2/Q)\epsilon,$$

and

$$(44)\qquad \min\{A(\xi_T'')-A(\xi_T'), P(\xi_T'')-P(\xi_T')\} \geq \theta\beta^{\Gamma}(q^2/Q)\epsilon.$$

Let the move from ξ_T' to ξ_T'', as described above, represent the entire second set of changes to bring us to our final sequence $\{\xi_\tau''\}$. I claim that this sequence payoff-dominates $\{\xi_\tau\}$ in date-0 payoffs, and moreover, that it is self-enforcing.

That the principal's payoffs are higher follows from the fact that $\pi_0' = \pi_0$ (by (30), and the fact that $\pi_0'' > \pi_0'$ by (44)). To see that the agent's payoffs are higher, notice that at date s, using (31),

$$(45)\qquad \begin{aligned}\alpha_s' - \alpha_s &= [A(c_s, m_s-\epsilon) + \beta^{t-s}A(c_t, m_t+\epsilon')] \\ &\quad - [A(c_s, m_s) + \beta^{t-s}A(c_t, m_t)] \\ &= (\psi_s - \psi_t)P'(c_s, m_s)\epsilon + o(\epsilon) \\ &\geq -\eta Q\epsilon + o(\epsilon),\end{aligned}$$

where the second equality uses (31) and the last inequality follows from (40) and the fact that $P'(c_s, m_s) \geq -Q$.

At the same time,

$$(46)\qquad \begin{aligned}\alpha_s'' - \alpha_s' &= \beta^{T-s}[A(\xi_T'') - A(\xi_T')] \\ &\geq \beta^{\Gamma}[A(\xi_T'') - A(\xi_T')] \\ &\geq \beta^{2\Gamma}\theta(q^2/Q)\epsilon = \mu\epsilon,\end{aligned}$$

where the first inequality comes from (41), the second inequality from (44), and the last equality from the definition of μ (see (38)). Combining (45) and (46), we see that

$$(47)\qquad \alpha_s'' - \alpha_s \geq [\mu - \eta Q]\epsilon + o(\epsilon),$$

which is strictly positive for small ϵ, by (39). Because the two contract sequences are identical before period s, this proves that $\alpha_0'' > \alpha_0$, completing the proof of payoff-dominance.

Next, I show that $\{\xi_\tau''\}$ is self-enforcing. Note that by our choice of ϵ and ϵ', the sequence $\{\xi_\tau'\}$—and a fortiori our sequence $\{\xi_\tau''\}$—must satisfy limited liability, as well as the principal's participation constraint (16) at every date. To complete the demonstration of self-enforcement, I show that (15) is satisfied. The three critical dates to consider are dates s, T, and t (at all other dates $V(\xi_\tau) = V(\xi_\tau'')$, while agent continuation payoffs are no lower). At date s, invoking (A.3) and (47),

$$V(c_s'', m_s'') = V(c_s, m_s-\epsilon) \leq V(c_s, m_s) \leq \alpha_s < \alpha_s'',$$

so (15) is met. At date t, using (A.3) again,

$$\begin{aligned}V(c_t'', m_t'') - V(c_t, m_t) &= V(c_t, m_t+\epsilon') - V(c_t, m_t) \\ &\leq A(c_t, m_t+\epsilon') - A(c_t, m_t) = \alpha_t'' - \alpha_t,\end{aligned}$$

so (15) must be satisfied once again. Finally, at date T, combine (42) and (43) to conclude that

$$\begin{aligned} G(\xi''_T) - G(\xi_T) = G(\xi''_T) - G(\xi'_T) &\leq \beta^{\Gamma}(q^2/Q)\epsilon \\ &\leq \alpha'_T - \alpha_T \\ &= \beta[\alpha'_{T+1} - \alpha_{T+1}] \\ &= \beta[\alpha''_{T+1} - \alpha_{T+1}], \end{aligned}$$

where the second-last equality follows from the fact that the agreement at date T is unaltered when moving to the sequence $\{\xi'_\tau\}$, and the last equality follows from the fact that no agreements beyond date T are altered in the further move to $\{\xi''_\tau\}$. Using the definition of G to rearrange the first and last terms of this string of inequalities, we see that

$$V(\xi''_T) - \alpha''_T \leq V(\xi_T) - \alpha_T \leq 0,$$

which verifies (15) at date T.

So $\{\xi''_\tau\}$ not only payoff dominates $\{\xi_\tau\}$, it is self-enforcing. But this contradicts the fact that $\{\xi_\tau\}$ is ESE. *Q.E.D.*

PROOF OF THEOREM 2: Let $\{\xi_t\}$ be an agent's best self-enforcing agreement sequence. Observe that

$$\pi_0 = 0. \tag{48}$$

(For if π_0 were positive, we could increase m_0, thus raising α_0. The increase would not violate feasibility, because L is increasing in m, and it would not violate (15), because $V'(c, m) \leq A'(c, m)$ by assumption (A.3).)

Note that the continuation sequence of this sequence from date 1 is also ESE. By the proof of Theorem 1, we know that there is a first date $T \geq 1$ for which $\pi_T = 0$ along the continuation sequence. By the same argument as in the first paragraph of the proof of Theorem 1,

$$\alpha_0 = \alpha_T. \tag{49}$$

Let $N \equiv T - 1$. Define, for each $t = 0, \ldots, N$, $\hat{\xi}_t \equiv \xi_t$. Consider the periodic agreement sequence generated by this collection in the obvious way: $\xi^*_t = \hat{\xi}_i$, where $i =$ Remainder$(t/N+1)$. It is easy to see, using (49), that $\alpha^*_0 = \alpha_0$. The self-enforcement constraints (15) and (16) are easy to verify for this new sequence.

So if the original self-enforcing agreement sequence were best for the agent, this one must be so as well. *Q.E.D.*

We now turn to the proofs for the transferable utility case. Thus assumptions (A.1)–(A.7) will all be assumed to hold in what follows.

Let S^* denote the maximum value of $A(c) + P(c)$.

LEMMA 7: *In no self-enforcing contract can the agent's payoff exceed* $S^*(1-\beta)^{-1}$.

PROOF: If α and π denote the payoffs to agent and principal under $\{\xi_t\}$, then

$$\alpha+\pi=\sum_{t=0}^{\infty}\beta^t[A(c_t)+P(c_t)]\leq S^*(1-\beta)^{-1}.$$

The lemma now follows from the fact that $\pi\geq 0$.

LEMMA 8: *An agreement* $\xi=(c,m)$ *is unimprovable if and only if* c *maximizes* $A(c)+P(c)$.

PROOF: Obvious.

LEMMA 9: *If* $\{\xi_t\}$ *is self-enforcing, then* ξ_t *must be improvable at every date* t.

PROOF: Suppose that at some date, $\xi_t=(c_t,m_t)$ is unimprovable. Then, by Lemma 8, c_t maximizes $A(c)+P(c)$ over the set of feasible agreements. Consider the stationary agreement (c,m) where $c=c_t$ and $m=P(c_t)$. Notice that the agent's payoff under this stationary agreement is exactly $S^*(1-\beta)^{-1}$. I claim this is self-enforcing.

To see this, note that if $m\leq m_t$, then

$$V(c,m)\leq V(c_t,m_t)\leq\alpha_t\leq S^*(1-\beta)^{-1}.$$

Otherwise, if $m>m_t$, then using (A.3),

$$V(c,m)-A(c,m)\leq V(c_t,m_t)-A(c_t,m_t)\leq\beta\alpha_{t+1}\leq\beta S^*(1-\beta)^{-1}.$$

In either case, we see that the stationary agreement satisfies the agent's no-deviation constraint (15). Clearly, (16) is trivially met. Therefore the stationary agreement with unimprovable (c,m) is self-enforcing. By Lemma 7, it must be efficient. Therefore it is ESE. But now we contradict (A.4). *Q.E.D.*

LEMMA 10: *If* $\{\xi_t\}$ *is WESE, then* $L(\xi_t)=0$ *whenever* $\pi_{t+1}>0$.

PROOF: Suppose on the contrary that $\pi_{t+1}>0$ for some t but that $L(\xi_t)>0$. Pick $\epsilon>0$ such that $L(c_t,m_t-\beta\epsilon)\geq 0$ and $\pi_{t+1}-\epsilon\geq 0$. Define a sequence $\{m'_s\}$ such that $m'_t=m_t-\beta\epsilon$, $m'_{t+1}=m_{t+1}+\epsilon$, and $m'_s=m_s$ for all other dates s.

The corresponding contract sequence $\{c_s,m'_s\}$ is surely feasible. We first verify the agent's no-deviation constraint. It should be obvious that t and $t+1$ are the only dates to check. Notice that $\alpha'_{t+1}=\alpha_{t+1}+\epsilon$, so that using (A.3) and (15) for the original sequence,

$$\beta\alpha'_{t+1}=\beta\alpha_{t+1}+\beta\epsilon\geq G(c_t,m_t)+\beta\epsilon\geq G(c_t,m_t-\beta\epsilon)=G(c_t,m'_t). \tag{50}$$

This verifies (15) at date t. At date $t+1$, note that

$$\beta\alpha'_{t+2} = \beta\alpha_{t+2} \geq G(c_{t+1}, m_{t+1}) \geq G(c_{t+1}, m'_{t+1}), \quad (51)$$

using (A.3) and (15) for the original sequence once again.

Next, the principal's participation constraints are satisfied—the only change in π is at date $t+1$, but $\pi'_{t+1} = \pi_{t+1} - \epsilon \geq 0$, by choice of ϵ.

So $\{c_s, m'_s\}$ is self-enforcing. Since it yields the same payoffs as the original sequence, it must be WESE.

Our next step is to note that at least one of the following inequalities, obtained from (50) and (51), *must be strict*:

$$\beta\alpha'_{t+1} \geq G(c_t, m'_t) \quad (52)$$

or

$$\beta\alpha'_{t+2} \geq G(c_{t+1}, m'_{t+1}). \quad (53)$$

To show this, we use (A.6). Assume that (52) holds with equality. Then, using (50), we conclude that

$$G(c_t, m_t) + \beta\epsilon = G(c_t, m_t - \beta\epsilon),$$

which implies that $V'(c_t, m_t) = 0$. But then, using (A.3) and (A.6), it must be the case that $V'(c, m) < 1$ for all (c, m). Consequently,

$$G(c_{t+1}, m_{t+1}) > G(c_{t+1}, m'_{t+1}),$$

and using this information in (51), we conclude that (53) must hold strictly.

Without loss of generality let (52) hold strictly. Now, by Lemma 9, (c_t, m'_t) is improvable. Our assumptions (notably (A.1) and (A.2)) imply that (c_t, m_t) can be replaced by some feasible agreement (c''_t, m''_t) such that both

$$A(c''_t) + m''_t > A(c_t) + m'_t \quad \text{and} \quad P(c''_t) - m''_t > P(c_t) - m'_t$$

hold, and

$$\beta\alpha'_{t+1} \geq G(c''_t, m''_t).$$

But this is easily seen to contradict our conclusion that $\{c_s, m'_s\}$ is WESE. Therefore our original supposition is wrong, and $L(\xi_t) = 0$ whenever $\pi_{t+1} > 0$. *Q.E.D.*

LEMMA 11: *If* $\{\xi_t\}$ *is ESE, then* $\pi_{t+1} = 0$ *whenever* $\pi_t = 0$.

PROOF: Suppose, on the contrary, that $\pi_t = 0$ and $\pi_{t+1} > 0$ for some t. By Lemma 10, it follows that $L(\xi_t) = 0$.

Because $\pi_t = 0$, it must be the case that

$$\alpha_t \geq [A(c^*) + m^*](1-\beta)^{-1}, \quad (54)$$

where (c^*, m^*), it will be recalled, is the solution to the agent's best stationary problem. Relation (54) simply follows from the fact that the stationary sequence in problem (17)–(19) is self-enforcing.[24]

[24] Once again, to extend (54) to WESE sequences, we need an argument such as the one made in footnote 23.

By (A.7) and the fact that $L(c_t, m_t) = 0$, we know that $A(c_t) + m_t = \lambda < A(c^*) + m^*$. Combining this information with (54), we see that

(55) $$\alpha_{t+1} > \alpha_t.$$

Moreover, we already have that

(56) $$\pi_{t+1} > \pi_t = 0.$$

But these two inequalities contradict the fact that the tail sequence from date t onwards is ESE. *Q.E.D.*

PROOF OF THEOREM 3: By Theorem 1, there exists some first date T at which $\pi_T = 0$, and such that the agent obtains her best self-enforcing payoff from that date onwards.

By Lemma 11, $\pi_t = 0$ for all $t \geq T$. Notice that this implies

(57) $$P(c_t) - m_t = 0$$

for all $t \geq T$.

Notice also that

(58) $$\alpha_t = \alpha_{t+1}$$

for all $t \geq T$. For if not, let s be the *first* date (no less than T) at which the equality is violated. It cannot be that $\alpha_{s+1} > \alpha_s$, for then $\alpha_s = \alpha_T$ cannot be the agent's maximal ESE payoff. But neither can the opposite inequality hold. For if it did, we could replace the tail sequence at $s+1$ by the contract sequence starting from T (using the fact that $\pi_{s+1} = 0$). This would raise agent payoff at date T while still satisfying all the self-enforcement constraints, a contradiction.

Given (57) and (58), it is now easy to check that for each $t \geq T$ and each ξ_t, we can construct a *stationary* self-enforcing contract sequence in which ξ_t is repeated at every date, with exactly the same payoff—α^*—to the agent. But by assumption (A.7), there can be only one such stationary contract, (c^*, m^*). So $\xi_t = (c^*, m^*)$ for all $t \geq T$. *Q.E.D.*

PROOF OF THEOREM 4: Fix any ESE sequence, and let T be given as in Theorem 3. Then $\pi_t > 0$ for all $t < T$. Using Lemma 10, this immediately establishes that $L(c_t, m_t)$ equals zero for all $t = 0, \ldots, T-2$.

By (A.7), $\lambda < (1-\beta)\alpha^*$. Because $\alpha_{T-1} = [A(c_{T-1}) + m_{T-1}] + \beta\alpha^* \geq \lambda + \beta\alpha^*$, we may also conclude that $\lambda < (1-\beta)\alpha_{T-1}$. Consequently, for all $t = 0, \ldots, T-2$,

$$\alpha_{t+1} - \alpha_t = \beta^{T-t-2}[(1-\beta)\alpha_{T-1} - \lambda] > 0.$$

Because $\pi_{T-1} > \pi_T$, it is easy to see that $\alpha_{T-1} < \alpha_T$ as well.[25] Thus (21) is proved.

[25] Suppose, on the contrary, that $\alpha_{T-1} \geq \alpha_T$. By making a small monetary transfer from principal to agent—only at date $T-1$—none of the constraints are disturbed. So this perturbed sequence from $T-1$ is also self-enforcing. But the agent value of such a sequence must *strictly* exceed α_T, which contradicts the fact that α_T is the agent's best self-enforcing payoff.

We complete the proof by establishing (22). Suppose, contrary to our assertion, that (22) is false for some date $s = 0, \ldots, T-2$. Consider a new agreement sequence $\{\xi'_t\}$ that simply replaces the agreement at date $s+1$ of the old sequence by a single repetition of ξ_s. By (22), it is easy to see that the new sequence must be self-enforcing, Moreover, since the new sequence weakly dominates $\{\xi_t\}$, which is ESE to start with, $\{\xi'_t\}$ must be ESE. However,

$$G(c'_{s+1}, m'_{s+1}) = G(c_s, m_s) \leq \beta\alpha_{s+1} < \beta\alpha_{s+2} = \beta\alpha'_{s+2}, \tag{59}$$

which shows that the new incentive constraint for the agent must be slack at $s+1$. Because $\{\xi'_t\}$ is ESE, it follows that $\xi'_{s+1} = \xi_s$ cannot be θ-relaxed for any $\theta > 0$, so by Lemma 2, it is unimprovable. But this observation contradicts Lemma 9. Therefore (22) holds. *Q.E.D.*

Department of Economics, New York University, 269 Mercer Street, New York, NY 10003, U.S.A.; debraj.ray@nyu.edu

Manuscript received March, 2000; final revision received February, 2001.

REFERENCES

ABREU, D. (1986): "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39, 191–225.

——— (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383–396.

ALBUQUERQUE, R., AND H. HOPENHAYN (1997): "Optimal Dynamic Lending Contracts with Imperfect Enforceability," Mimeo., University of Rochester.

BANERJEE, A., P. GERTLER, AND M. GHATAK (2001): "Empowerment and Efficiency: The Economics of Tenancy Reform," forthcoming, *Journal of Political Economy*.

BANERJEE, A., D. MOOKHERJEE, K. MUNSHI, AND D. RAY (2001): "Inequality, Control Rights and Rent Seeking: Sugar Cooperatives in Maharashtra," *Journal of Political Economy*, 109, 138–190.

COATE, S., AND M. RAVALLION (1993): "Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements," *Journal of Development Economics*, 40, 1–24.

FAFCHAMPS, M. (1996): "Risk Sharing, Quasi-Credit, and the Enforcement of Informal Contracts," Mimeograph, Department of Economics, Stanford University.

GHOSH, P., AND D. RAY (1996): "Cooperation in Community Interaction without Information Flows," *Review of Economic Studies*, 63, 491–519.

HARRIS, M., AND B. HOLMSTROM (1982): "A Theory of Wage Dynamics," *Review of Economic Studies*, 49, 315–333.

KLETZER, K., AND B. WRIGHT (2000): "Sovereign Debt as Intertemporal Barter," *American Economic Review*, 90, 621–639.

KOCHERLAKOTA, N. (1996): "Implications of Efficient Risk-Sharing Without Commitment," *Review of Economic Studies*, 63, 595–610.

KRANTON, R. (1996): "The Formation of Cooperative Relationships," *Journal of Law, Economics and Organization*, 12, 214–233.

LANG, K., AND S. KAHN (2001): "Hours Constraints: Theory, Evidence and Policy Implications," in *Working Time in Comparative Perspective*, ed. by S. Houseman and G. Wong. Kalamazoo, MI: Upjohn Institute.

LAZEAR, E. (1981): "Agency, Earnings Profiles, Productivity and Hours Restrictions," *American Economic Review*, 71, 606–620.

LEHRER, E., AND A. PAUZNER (1999): "Repeated Games with Differential Time Preferences," *Econometrica*, 67, 393–412.

LIGON, E., J. THOMAS, AND T. WORRALL (2001): "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," forthcoming, *Review of Economic Studies*.

MOOKHERJEE, D. (1997): "Informational Rents and Property Rights in Land," in *Property Rights, Incentives and Welfare*, ed. by J. Roemer. New York: Macmillan Press.

MOOKHERJEE, D., AND D. RAY (2001): "Contractual Structure and Wealth Accumulation," Mimeo., Boston University.

SHAPIRO, C., AND J. STIGLITZ (1984): "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review*, 74, 433–444.

THOMAS, J., AND T. WORRALL (1988): "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55, 541–554.

——— (1994): "Foreign Direct Investment and the Risk of Expropriation," *Review of Economic Studies*, 61, 81–108.

TOWNSEND, R. (1993): "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.

UDRY, C. (1994): "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria," *Review of Economic Studies*, 61, 495–526.

WATSON, J. (1999): "Starting Small and Renegotiation," *Journal of Economic Theory*, 85, 52–90.