Joint Berkeley-Stanford Finance Seminar
Wednesday, October 31, 2012
3:40-4:30 pm
Room F320, Koret Classroom
Paper #2

Impatience vs. Incentives*

Marcus M. Opp† John Y. Zhu ‡

October 26, 2012

Abstract

This paper studies the long-run dynamics of Pareto-optimal self-enforcing contracts in a repeated principal-agent framework with differential discounting. In such a setting, impatience concerns encourage contracts to favor the patient player in the long run and incentive concerns encourage contracts to favor the agent in the long run. When these two forces are aligned or one is relatively strong, we show that every Pareto-optimal self-enforcing contract converges to a steady state in the long run in a well-behaved way. In particular, the results of Ray (2002) and Lehrer and Pauzner (1999) can be recovered as limiting cases. However, when the forces are opposed and of comparable strength, our main result is that some Pareto-optimal self-enforcing contracts deterministically oscillate between agent-preferred and principal-preferred states. Oscillation may be damped - so that in the long run a steady state is reached; or explosive until the allocation reaches one of the participation constraints and then alternates between a pair of states forever.

*We thank Peter DeMarzo, Mike Fishman, Willie Fuchs, Drew Fudenberg, Nicolae Garleanu, Milton Harris, Rich Kihlstrom, Natalia Kovrijnykh, Jonathan Levin, Andrey Malenko, John Morgan and Bilge Yilmaz for helpful comments.

†University of California, Berkeley (Haas), email: mopp@haas.berkeley.edu.

‡University of Pennsylvania, Wharton, e-mail: zhuyiran@wharton.upenn.edu.

1 INTRODUCTION

In this paper, we study a simple, broadly applicable repeated principal-agent framework with differential discounting. We focus on how impatience and incentives affect the long-run dynamics of Pareto-optimal self-enforcing contracts. Implied is the idea that impatience and incentives are distinct forces that, depending on the parameters, may either reinforce or go against each other. The impatience force implies that contracts should favor the more patient player in the long-run in order to exploit trading gains across time. Lehrer and Pauzner (1999) validates this dynamic in repeated games when players are arbitrarily patient. The incentives force implies that contracts should favor the agent in the long-run in order to efficiently provide incentives. This is the main thesis of Ray (2002), who in a similar model shows that when the agent and principal are equally patient, all Pareto-optimal self-enforcing contracts favor the agent in the long-run.

When the agent is at least as patient as the principal, the two forces are aligned and we show that every Pareto-optimal payoff can be supported by a self-enforcing contract that converges to a steady state in the long run in a well-behaved way. Here, steady state denotes a stationary Pareto-optimal self-enforcing contract. We also show that this result holds when either one of the forces is sufficiently strong relative to the other. This result can greatly simplify the study of long-run dynamics by focusing attention onto easily characterized steady states.

We then consider the alternate case, when the impatience and incentive forces are comparable but opposite. This is the case when the principal is strictly more patient and the agent's incentive problem is non-trivial. Here, as Ray (2002) notes, the two forces "tug in different directions [and] it may be worth exploring if one of the two factors always dominates." In our paper, we provide a complete characterization of Pareto-optimal self-enforcing contracts in an important model within the general framework. Many special cases of the model have appeared in the previous literature including Ray (2002). One of the most surprising findings is that almost all Pareto-optimal self-enforcing contracts oscillate.

Here, oscillation refers to the deterministic non-monotonic dynamics of the agent's payoff around the unique steady state. The oscillation may be damped - so that in the long run the steady state is still reached, but stable or even explosive oscillation also occur. By always cycling back to a principal-preferred state, an oscillating self-enforcing contract ensures that the gains from trade across time between patient principal and impatient agent are at least partially recouped. By following each such trip with a promise to return to an agent-preferred state, such an oscillating self-enforcing contract also ensures that the agent does not have sufficient incentives to deviate.

Among applied papers, the assumption of relative impatience of the agent is commonly made to make constraints bind in the long run. In particular, Acemoglu, Golosov, and Tsyvinski (2008) and Aguiar, Amador, and Gopinath (2009) establish long run investment distortions in these settings. Opp (2012) shows that the "distortionary" steady state is characterized by the trade off between the static relationship gains (effort) with the gains from trading across time (credit). None of these papers, however, uncover the feature of non-trivial oscillation dynamics even in the absence of exogenous uncertainty.

2 Theory

Consider an infinite horizon, discrete time repeated principal-agent relationship with perfect public information in which time is indexed by $t \in \mathbb{Z}_+$. There is a compact, convex space $(\mathscr{A}, \geq)$ that serves as the action space. The two players - the agent and the principal - are denoted by A and P, respectively. Each player i has a continuous utility function u_i defined over $\mathscr{A}$, a discount factor δ_i and an outside option worth $O_i = \frac{o_i}{1-\delta_i}$ in present value terms.[1] There is a continuous deviation function $D : \mathscr{A} \longrightarrow \mathbb{R}$ characterizing the agent's deviation incentives.

Given a sequence of actions $\{a_t\}$, introduce the time t discounted payoff of the sequence:

$$(U_{A,t}, U_{P,t}) = \left(\sum_{s=t}^{\infty} \delta_A^{s-t} u_A(a_s), \sum_{s=t}^{\infty} \delta_P^{s-t} u_P(a_s) \right) \tag{1}$$

A (self-enforcing) *contract* is a sequence of actions $\{a_t\}$ such that for each date t, the participation constraint $(U_{A,t}, U_{P,t}) \geq (O_A, O_P)$ and the incentive-compatibility constraint $\delta_A U_{A,t+1} \geq D(a_t)$ are both satisfied.

Definition. *D is semi-convex over $(\mathscr{A}, u_A, u_P)$ if for any $a_1, a_2 \in \mathscr{A}$ and $\lambda \in [0,1]$ there exists an $a_\lambda \in A$ such that $u_i(a_\lambda) \geq \lambda u_i(a_1) + (1-\lambda)u_i(a_2)$ for $i = A, P$, and $D(a_\lambda) \leq \lambda D(a_1) + (1-\lambda)D(a_2)$. D is strictly semi-convex if for at least one player i, $u_i(a_\lambda) > \lambda u_i(a_1) + (1-\lambda)u_i(a_2)$.*

For example, if D is convex then it is semi-convex. Throughout the rest of the paper, we will assume semi-convexity of D. Let C denote the set of contract payoffs. Let $F \subset C$ denote the set of stationary contract payoffs. So, let V denote the set of all Pareto-optimal payoffs of C. As an abuse of notation, we will also think of $V = V(U_A)$ as the principal's value function over agent payoffs. It is easy to show that C is compact and convex.[2] Therefore, V-contracts are randomization-proof and as a function,

Lemma 1. *$V(U_A)$ is concave with compact domain.*

Our goal is to establish some facts about the long-run dynamics of V-contracts with a focus on the role played by discount factors. The following fundamental result holds for all pairs of discount factors:

Lemma 2. *There exists at least one steady state (stationary V-contract.) It is unique if D is strictly semi-convex. In general, the continuation contract of any V-contract is a V-contract.*

Lemma 2 offers the tantalizing possibility of reducing an exhaustive study of the long-run dynamics of V-contracts down to characterizing steady states. The following two results provide conditions under which this possibility is realized:

[1]Note that without imposing incentive/liquidity constraints, the assumption of differential discounting would lead to unbounded payoffs.

[2]Convexity is an immediate consequence of the semi-convexity of D. Compactness means bounded and closed. Boundedness follows from the participation constraints. Closedness follows from an easy application of the diagonalization trick.

Lemma 3. *If the only steady state is V_R, the rightmost payoff of V, then every V-contract converges monotonically to the steady state. A mirror result holds with V_R replaced with V_L, the leftmost payoff of V.*

Both Lemma 2 and Lemma 3 follow from Kakutani's Fixed Point Theorem. While the condition in Lemma 3 is not on fundamentals, it is still of practical value. Indeed, it is often straightforward to impose conditions on fundamentals so as to rule out, a priori, any non-corner stationary contracts as Pareto-optimal. For example, consider Ray (2002) which analyzes a repeated employment model that is an example of our model. The author imposes an upper bound on the discount factor, which effectively guarantees that in the steady state, the principal's participation constraint must bind. This implies the steady state is V_R. More generally, if the agency problem is sufficiently severe relative to the principal's participation constraint, payoffs to the agent must be backloaded (regardless of differential discounting).

While the general model of Ray (2002) does not completely fit within our framework, it is worth noting that both of the paper's motivating examples do. In fact, in both examples, one can easily recover Ray's main result - that every V-contract must converge to the V_R-contract - by showing that the conditions for Lemma 3 hold.

While monotonic dynamics are common, they are not the only type of robust dynamics observed in V-contracts. The following definition introduces some useful terminology:

Definition. *A sequence $\{x_t\}$ is damped if for every t, $x_{t+1} > x_t \Rightarrow x_{t+2} > x_t$ and $x_{t+1} < x_t \Rightarrow x_{t+2} < x_t$. A sequence $\{x_t\}$ oscillates if for every t, either $x_{t+1} > x_t, x_{t+2}$ or $x_{t+1} < x_t, x_{t+2}$. A sequence damped-oscillates if it oscillates and is damped. A sequence stable-oscillates if it oscillates and for every t, $x_t = x_{t+2}$. A sequence explosive-oscillates if it oscillates and for every t, either $x_{t+2} < x_t < x_{t+1}$ or $x_{t+1} < x_t < x_{t+2}$.*

Lemma 4. *If there does not exist an stable-oscillating V-contract, then every V-contract damped-converges to a steady state.*

Once again, the condition is not on fundamentals. However, we will show in Theorem 1 that the condition is satisfied whenever the agent is at least as patient as the principal. Lemma 4 will follow from the proof of Theorem 1. Of course, Lemma 4 also hints at something more intriguing: if it is the case that some model does not exhibit convergence to steady state, then oscillation may be a major feature of constrained-optimal behavior.

2.1 DYNAMICS: $\delta_A \geq \delta_P$

Theorem 1. *If $\delta_A \geq \delta_P$, then every V-payoff can be supported by a contract that damped converges to a steady state.*

Theorem 1 does not rule out Pareto optimal contracts that exhibit damped-oscillations. One may wonder how likely it is that a simple, convex and economically sensible model of self-enforcing contracting might lead to such seemingly exotic behavior.

In the next subsection, we consider a family of such models. Indeed, some members of the family are quite well-known and have been considered in numerous previous papers. Yet the myriad examples of non-monotonic, specifically oscillating, behavior that emerges from a full analysis of this family has, until now, gone largely undocumented.

2.2 Dynamics: $\delta_A < \delta_P$

When the agent is strictly more impatient, a tension arises between incentive provision and efficiency. Incentive provision typically involves backloading payments to the agent. Conversely, efficiency, which requires exploiting gains from trade across time, involves backloading payments to the more patient player, which in this case is the principal (see e.g., Lehrer and Pauzner (1999)).

The impatience versus incentives conflict generates a notable second-best optimal arrangement: oscillation. We now explore this phenomenon by characterizing optimality in a specific but important model that fits within the previously developed general framework.

Let $\mathscr{A} = \{(w,e)|w,e \in \mathbb{R}\}$, $u_A = m - c(e)$, $u_P = r(e) - m$ and $D = c(e) + d(e) - (1-\theta)m + \delta_A O_A$.[3] We assume c is increasing and strictly convex, r is increasing and concave, $r - c$ achieves an interior maximum at e^*, d is increasing and concave and $\theta \in [0,1]$. The constant θ is the fraction of money the agent can take with him should he decide to deviate. Inspection of the deviation function illustrates that θ governs the (in)efficiency of monetary transfers in reducing the current period incentive problem.[4]

This model includes as special cases those considered in Ray (2002), Opp (2012), as well as simplified versions of those found in Acemoglu, Golosov, and Tsyvinski (2008), and Thomas and Worrall (1988, 1994).

In light of Lemma 3, the following analysis will restrict attention to when there is at least one interior steady state, i.e., both incentives and impatience matter. A sufficient condition for an interior steady state, is to require existence of a stationary contract with the static efficient effort level e^* and some constant transfer m^s such that $(U_A^s, U_P^s) > (O_A, O_P)$. In the following analysis, we will denote (U_A^s, U_P^s) as any stationary contract whereas $(\bar{U}_A, \bar{U}_P)$, refers to the (unique) steady state, i.e., $\bar{V}$.

Proposition 1. *There is a unique steady state V-contract with $\bar{U}_A = d(\bar{e}) + \theta\bar{m} + \delta_A O_A$. If $\theta = 0$ or 1, every V-contract is monotone convergent. If $\theta \in (0,1)$, every non-stationary V-contract is an oscillation contract.*

Proposition 1 shows the important role played by oscillation. For any non-corner θ, any Pareto-optimal contract that is not the steady state oscillates.[5] The following lemma is an important step in the proof of Proposition 1 that helps explain why oscillation is important.

Lemma 5. *Let $\theta \in (0,1)$. Fix a stationary contract with action (m^s, e^s) and payoff (U_A^s, U_P^s). Let $\rho := \frac{1}{\delta_A}\frac{1-\theta}{\theta}$. For every $m \in \mathbb{R}$, the action sequence $C(m^s, e^s, m) := \{(m^s + (-\rho)^t m, e^s)\}_{t=0}^{\infty}$ satisfies the IC-constraint at every date and has payoff $U(m^s, e^s, m) := (U_A^s + \frac{m}{1+\delta_A\rho}, U_P^s - \frac{m}{1+\delta_P\rho})$. The set of payoffs $\{U(m^s, e^s, m) \mid m \in \mathbb{R}\}$ is linear.*

[3]It is often economically sensible to impose a non-negativity constraint on m. We ignore this constraint purely for expositional reasons.

[4]Technically, this model doesn't fit the general framework since $\mathscr{A}$ is not compact. However, the outside options guarantee that only a compact subset of $\mathscr{A}$ is ever used in any contract.

[5]Part of the reason this behavior seems to have been missed by the literature is that previous papers have focused on the $\theta = 0$ and 1 cases.

Recall V must be concave. So suppose the stationary contract of Lemma 5 was the unique steady state V-contract of Proposition 1. Then, were it not for participation constraints, the Lemma 5 action sequences would actually be the set of Pareto-optimal V-contracts (because of linearity).

Indeed, when $\theta \geq \frac{1}{1+\delta_A}$ so that $\rho \leq 1$, then starting at date $t = 1$ (at the latest), the action sequence of every V-contract becomes a Lemma 5 action sequence. This is not surprising since every non-stationary Lemma 5 action sequence oscillates or damped-oscillates around the stationary contract. If such a sequence satisfies participation constraints for the first two dates then it will satisfy participation constraints for all dates in the future. In this case, the agent value evolves according to a simple non-homogeneous first-order difference equation.

$$U_{A,t+1} = -\rho U_{A,t} + k(\bar{e}) \tag{2}$$

with $k(e) = \frac{c(e)}{\delta_A} + \frac{d(e)+\delta_A O_A}{\theta \delta_A}$ so that $\bar{U}_A = \frac{k(\bar{e})}{1+\rho}$.

In general, any non-stationary Lemma 5 action sequence continually cycles the monetary transfer below the stationary level. This assures that the more patient principal repeatedly receives a continuation payoff strictly higher than the stationary level. The arrangement represents a form of second-best backloading of payments to the more patient player - the principal. Of course, whenever the sequence is in such a principal-preferred state, the deviation incentives of the agent are relatively high. To combat this, the monetary transfer needs to swing back to a greater than stationary level so that the agent will be rewarded tomorrow for choosing to stay in the relationship today. This creates the ubiquitous oscillation dynamic and assures that the IC-condition holds when the action sequence is in the principal-preferred state.

When the action sequence is in the agent-preferred state, there is also a strong incentive to deviate - this time because the continuation payoff will be low for the agent. However, the agent understands that if he deviates now, he is losing a portion of an above average monetary transfer. This is the $-(1-\theta)m$ component of the deviation function If the oscillation is calibrated correctly and the transfer is sufficiently above average, it creates an equally strong disincentive to deviate and the IC-condition will again hold.

Of course, the higher the θ, the smaller the cost of deviation. Consequently, higher θ means it is harder to use the oscillation arrangement to discourage deviation when the action sequence is in the agent-preferred state. This is why as $\theta \to 1$, the oscillation feature disappears and the unique long-term dynamic is to immediately sink into the steady state.

Now, when $\theta < \frac{1}{1+\delta_A}$, every non-stationary Lemma 5 action sequence explosively-oscillates around the stationary sequence. Such action sequences will eventually permanently break both participation constraints. Thus, Pareto-optimal contracts will exhibit persistent distortions and will not become exactly like Lemma 5 action sequences over time.

This will imply that V is no longer linear around the steady state (see left panel of Figure 1). In this case, we can exploit the changing slope of V to provide an implicit characterization of both transfer and effort dynamics.

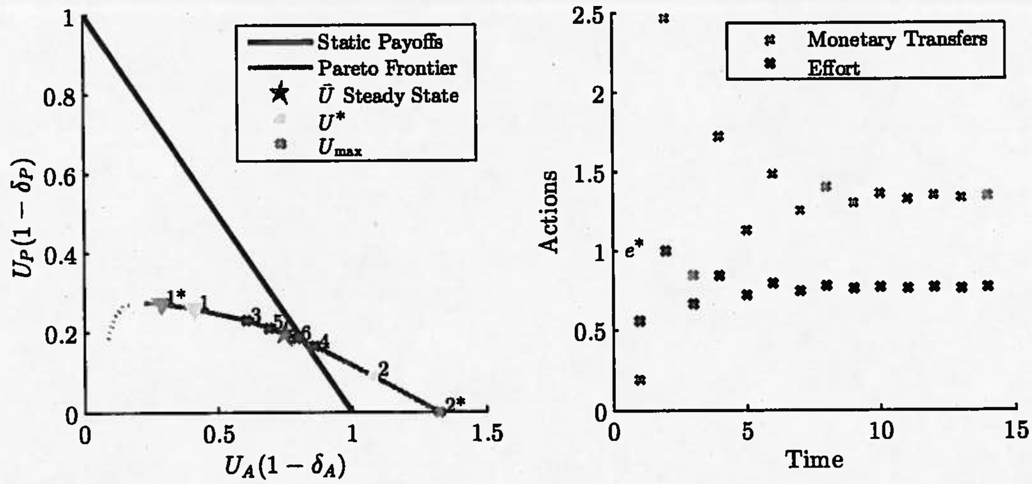

Figure 1: **Pareto Frontier and Dynamics.** The left panel plots the Pareto Frontier in the normalized payoff space $U_i(1-\delta_i)$ for the case $\theta < \frac{\delta_P}{\delta_P+\delta_A}$. Dynamic trading gains make allocations outside the static payoff space feasible. Two sequences of optimal dynamic allocations starting with 1 and 1* are plotted. From period 3 onwards, both sequences coincide. Any allocation between U^* and $U_{\max}$ features investment at the static optimum $e^* = 1$. The right panel plots the sequence of actions as a function of time.

Lemma 6. *For any V-contract with payoff $(U_{A,t}, U_{P,t} = V(U_{A,t}))$ and continuation payoff $(U_{A,t+1}, U_{P,t+1} = V(U_{A,t+1}))$, the following differential conditions must be satisfied:*

$$d^-V(U_{A,t+1}) \geq -\kappa - \frac{1}{\rho\delta_P} d^+V(U_{A,t}) \tag{3}$$

$$d^+V(U_{A,t+1}) \leq -\kappa - \frac{1}{\rho\delta_P} d^-V(U_{A,t}) \tag{4}$$

where $\kappa = \frac{\delta_A}{\delta_P}\frac{1}{1-\theta}$.

Lemma 6 implies that (U_A, U_P) is the steady state if and only if

$$d^-V(\bar{U}_A) \geq -\frac{\delta_A}{(1-\theta)\delta_P + \theta\delta_A} \quad \text{and} \quad d^+V(\bar{U}_A) \leq -\frac{\delta_A}{(1-\theta)\delta_P + \theta\delta_A}.$$

It can also be used to characterize non-stationary V-contracts for $\theta < 1/(1+\delta_A)$ much in the same way Lemma 5 characterizes V-contracts for $\theta \geq 1/(1+\delta_A)$. For $\theta \in (0, \frac{1}{1+\delta_A})$ every non-stationary contract oscillates around the steady state. When $\theta \in (\frac{\delta_P}{\delta_P+\delta_A}, \frac{1}{1+\delta_A})$, i.e., $\rho\delta_P < 1$, these oscillation contracts are explosive and all converge to a unique maximal oscillating contract. Here, maximal means that (at least) one of the continuation payoffs is a corner payoff. Interestingly, even arbitrarily low participation constraints will bind in the long-run. When $\theta = \frac{\delta_P}{\delta_P+\delta_A}$, there is a continuum of stable oscillating contracts around the steady state, not unlike the $\frac{1}{1+\delta_A}$ case. Lastly, when $\theta < \frac{\delta_P}{\delta_P+\delta_A}$, all oscillating contracts are damped and converge to the steady state (see Figure 1).

A key implication of Lemma 6 that differentiates the $\theta < \frac{1}{1+\delta_A}$ case from the $\theta \geq \frac{1}{1+\delta_A}$ case is that effort will now also vary over time due to the permanent distortions induced by the participation constraints. For example, if c, r and V are differentiable, then Lemma 6 implies that effort satisfies

$$\frac{r'(e) - c'(e)}{c'(e)\,\theta + d'(e)} = \frac{1 + dV(U_A)}{1 - \theta} \tag{5}$$

In general, Lemma 6 predicts that effort and monetary transfers will exhibit co-movement (see right panel of Figure 1). When transfers are high and deviation entails an above average loss in transfer, a Pareto-optimal contract can credibly demand a higher effort from the agent, potentially even at the static first best effort level (such as in period 2 of the sequence plotted in Figure 1). When transfers are low and deviation is cheap, a Pareto-optimal contract can only demand a below steady state level of effort.

It is worth highlighting that if close the gap between δ_A and δ_P by, say, reducing δ_P, then eventually, oscillation and investment distortions will disappear. Indeed, when $\delta_P = \delta_A$, there exists a continuum of steady states (where the slope of $V = -1$) and any V-contract with $d^+V(U_A) > -1$ will move rightwards to a steady state with static efficient investment e^* in the next period. This is a simple example of the idea that if you shut down the impatience force, then the incentive force dominates and Ray (2002)'s result is recovered: V-contracts will favor the agent in the long-run.

Similarly, Lehrer and Pauzner (1999) can be recovered by shutting down the incentive force, say, by making both parties arbitrarily patient:

Corollary 1. *Fix $\frac{\log(\delta_P)}{\log(\delta_A)}$ and take the limit as δ_P approaches one, then the normalized payoff of the agent in the steady state, $\bar{U}_A(1-\delta_A)$, converges to the per-period value of the outside option o_A.*

Since the steady state approaches V_L, monotone convergence favoring the principal over time is implied by Lemma 3.

3 Conclusion

In this paper, we explore the implications for self-enforcing contracting when the principal and agent have potentially different discount factors. When the agent is at least as patient as the principal, we showed that every Pareto-optimal self-enforcing contract converges to a steady state. This helps put into context Ray (2002) and reduces the study of long-run dynamics down to characterizing steady states.

However, when the agent is strictly more impatient, there is a conflict between incentive provision which implies backloading payments to the agent and efficiency which implies backloading rewards to the principal. We find that a robust compromise between these two forces is oscillation. By always switching between principal-preferred and agent-preferred states, such arrangements find an interesting solution to the incentive-compatibility problem: the agent has too much to gain tomorrow from staying when the current state is bad and has too much to lose today from deviating when the current state is good.

The "strict" optimality of oscillation also implies that when the agent is relatively impatient, simple stationary contracts likes those of Levin (2003) can be insufficient to characterize optimal allocations even in seemingly well-behaved, non-stochastic models.

While oscillation arrangements are a natural theoretical implication of relative impatience of the agent, its connection with features of real-life contracts seems unclear. We hope our findings serve as a catalyst for further research.

A Proofs

Proof of Lemma 2. For each payoff $v \in V$, define the set $\kappa(v) \subset V$ to consist of all continuation payoffs of contracts with payoff v. κ is a convex valued correspondence with closed graph. The result then follows from Kakutani's Fixed Point Theorem. □

Proof of Lemma 3. Recall κ from the proof of Lemma 2. For each set $\kappa(v)$, let $\kappa(v)_L$ and $\kappa(V)_R$ denote the leftmost and rightmost points of $\kappa(v)$. Suppose for all $v \in V - V_R$, $\kappa(v)_L$ is strictly to the right of v. Then the lemma holds. Otherwise, let v' be a point such that $\kappa(v')_L$ is to the left of v'. If $\kappa(v')_R$ is to the right of v', then v' is a fixed point. Contradiction. So suppose $\kappa(v')_R$ is also to the left of v'. Consider the subset of V: $V|_{[V_L,v']}$ and the correspondence $\kappa' := \eta \circ \kappa$ where $\eta : V \to V$ is the continuous map that is the identity on $[V_L, v']$ and maps $(v', V_R]$ to v'. By Kakutani's Fixed Point Theorem, κ' has a fixed point $v^* \in V|_{[V_L,v']}$. Since $\kappa(v')_R$ is to the left of v', $v^* \neq v'$. But then that means v^* must also be a fixed point of κ. Contradiction. □

Lemma 7. *Let F be a continuous function from $[0,1]$ to itself and consider the set of all iteration sequences generated by F. Either every iteration sequence converges stably or at least one of them is an oscillation.*

Proof. Suppose there exists a sequence $\{F^n(x)\}$ which is not stable convergent. If $\{F^n(x)\}$ is stable, then $F(\limsup\{F^n(x)\}) = \liminf\{F^n(x)\}$ and $F(\liminf\{F^n(x)\}) = \limsup\{F^n(x)\}$ is an oscillation. Otherwise, there exists some $y = F^N(x)$ such that $F^2(y) < y < F(y)$ or $F(y) < y < F^2(y)$. WLOG assume $F(y) < y < F^2(y)$. This implies the existence of a fixed point z^* of F in the interval $(F(y), y)$. Since $F^2(y) - y > 0$ and $F^2(1) \leq 1$, there exists a value $z^{**} \in (y, 1]$ such that $F^2(z^{**}) = z^{**}$. If $F < z^*$ on $(y, 1]$ then $\{F(z^{**}), z^{**}\}$ is a 2-period cycle. Otherwise, let z be the smallest value in $(y, 1]$ such that $F(z) = z^*$. Then $F^2(z) - z < 0$. Therefore, z^{**} may be chosen to be in (y, z). Again, $\{F(z^{**}), z^{**}\}$ is an oscillation. □

Proof of Lemma 4. Assume D is strictly semi-convex over $\mathscr{A}$.[6] Then the κ from the proof of Lemma 2 is a continuous function. The result follows immediately from Lemma 7. □

Proof of Theorem 1. Suppose the theorem is false. Then Lemma 7 implies there exists an oscillating V-contract $a_1 \circlearrowleft a_2$. Let $u_1 := (u_A(a_1), u_P(a_1))$, $u_2 := (u_A(a_2), u_P(a_2))$ and

[6]The general semi-convex case can be deduced from the strict semi-convex case by using a limits argument.

$U_{12} = (U_{12,A}, U_{12,P}) :=$ payoff of $a_1 \circlearrowleft a_2$. Of course $a_2 \circlearrowleft a_1$ is also an V-contract with payoff $U_{21} = (U_{21,A}, U_{21,P})$.

Semi-convexity implies that there exists an action a_h with utility $u_h := (u_A(a_h), u_P(a_h))$ such that $u_h \geq \frac{u_1+u_2}{2}$ and $D(a_h) \leq \frac{D(a_1)+D(a_2)}{2} \leq \frac{\delta_A(U_{12,A}+U_{21,A})}{2} = \frac{\delta_A(u_A(a_1)+u_A(a_2))}{2}$. Therefore $u_h \in F$. Also, $u_2 \in F$. Thus there exists a λ such that $\lambda u_h + (1-\lambda)u_2 \geq U_{21}$. Semi-convexity implies there exists an $\tilde{a}$ with utility $\tilde{u} := (u_A(\tilde{a}), u_P(\tilde{a}))$ such that $\tilde{u} \geq \lambda u_h + (1-\lambda)u_2$ and $D(\tilde{a}) \leq \delta_A(\lambda u_A(a_h) + (1-\lambda)u_A(a_2))$. Therefore $\tilde{u} \in F$ and either $\tilde{u} > U_{21}$ or $\tilde{u} = U_{21} \in F$. Both are contradictions.

□

Proof of Lemma 6. Let (w_0, e_0) be the initial action. Incentive-compatibility requires $\delta_A U_{A,1} \leq c(e_0) + d(e_0) - (1-\theta)w_0 + \delta_A O_A$. For every $\epsilon \neq 0$, consider the perturbed contract C_ϵ with initial action $(e_0, w_0 + \epsilon)$ followed by a continuation V-contract with agent payoff $U_{A,1} - \frac{(1-\theta)\epsilon}{\delta_A}$. Contract C_ϵ's agent payoff is $U^\epsilon_{A,0} = U_{A,0} + \theta\epsilon$. Its principal payoff is

$$U^\epsilon_{P,0} = r(e) - (w+\epsilon) + \delta_P V\left(U_{A,1} - \frac{(1-\theta)\epsilon}{\delta_A}\right) =$$

$$V(U_{A,0}) - \epsilon + \delta_P \left(V\left(U_{A,1} - \frac{(1-\theta)\epsilon}{\delta_A}\right) - V(U_{A,1})\right)$$

The efficiency of C requires $V(U_{A,0} + \theta\epsilon) \geq U^\epsilon_{P,0}$, which implies

$$V\left(U_{A,1} - \frac{(1-\theta)\epsilon}{\delta_A}\right) - V(U_{A,1}) \leq \frac{\epsilon + V(U_{A,0}+\theta\epsilon) - V(U_{A,0})}{\delta_P}$$

Dividing both sides by ϵ and taking the limit as ϵ goes to zero from above and below, we have the two differential conditions □

Proof of Proposition 1.
Case 1: $\theta \geq \frac{1}{1+\delta_A}$
Let $\mathcal{S}$ denote the set of payoffs (U_A, U_P) in V satisfying $d^- V(U_A) = \frac{-\delta_A}{\theta\delta_A + (1-\theta)\delta_P}$ or $d^+ V(U_A) = \frac{-\delta_A}{\theta\delta_A + (1-\theta)\delta_P}$. Lemma 6 implies that any stationary V-payoff (i.e. any V-payoff $\in F$) lies in $\mathcal{S}$. Suppose there are two distinct stationary V-payoffs with actions (m^s, e^s) and (m', e'). Then $\mathcal{S}$ is a line segment with slope $\frac{-\delta_A}{\theta\delta_A+(1-\theta)\delta_P}$. Suppose $e^s \neq e'$. Since $c(e)$ is strictly convex, any point on $\mathcal{S}$ strictly between the two stationary payoffs is in the interior of F. Contradiction. This implies that $m^s \neq m'$. But then, the slope of $\mathcal{S}$ is $-\frac{1-\delta_A}{1-\delta_P} \neq \frac{-\delta_A}{\theta\delta_A+(1-\theta)\delta_P}$. Contradiction. Thus, there is a unique stationary V-payoff and therefore, a unique stationary V-contract. Call this contract C^s

Let (m^s, e^s) and (U^s_A, U^s_P) be the action and payoff of C^s. Consider an action sequence of the type considered in Lemma 5: $C(m^s, e^s, m) := \{(m^s + (-\rho)^t m, e^s)\}_{t=0}^\infty$ where $\rho := \frac{1}{\delta_A}\frac{1-\theta}{\theta}$. Since $\theta \geq \frac{1}{1+\delta_A}$, $C(m^s, e^s, m)$ exhibits stable dynamics. This implies that if $C(m^s, e^s, m)$'s payoff $(U^s_A + \frac{w}{1+\delta_A\rho}, U^s_P - \frac{w}{1+\delta_P\rho})$ and continuation payoff $(U^s_A - \frac{\rho w}{1+\delta_A\rho}, U^s_P + \frac{\rho w}{1+\delta_P\rho})$ satisfy the participation constraints, $C(m^s, e^s, m)$ is a contract.

Since the stationary contract is assumed to be in the interior of V, the participation constraints do not bind. Therefore m can be picked to be nonzero. Let $\overline{w} > 0$ be the largest m such that $C(m^s, e^s, m)$ is a contract. Then $U_P^s - \frac{\overline{m}}{1+\delta_P\rho} = O_P$ or $U_A^s - \frac{\rho\overline{m}}{1+\delta_A\rho} = O_A$. Otherwise, $\overline{m}$ could be increased and the resulting action sequence would still be a contract, contradicting the maximality of $\overline{m}$. Similarly, let $\underline{m} < 0$ be the smallest m such that $C(m^s, e^s, m)$ is a contract. Then $U_A^s + \frac{m}{1+\delta_A\rho} = O_A$ or $U_P^s + \frac{\rho m}{1+\delta_P\rho} = O_P$.

The locus of contract payoffs $\mathcal{W} := \{(U_A^s + \frac{m}{1+\delta_A\rho}, U_P^s - \frac{m}{1+\delta_P\rho})\}_{m\in[\underline{m},\overline{m}]}$ is linear and has slope $\frac{-\delta_A}{\theta\delta_A+(1-\theta)\delta_P}$. Thus $\mathcal{W} \subset \mathcal{S} \subset V$. Clearly, if $\theta = 1$ $\mathcal{W} = \mathcal{S} = V$ and we're done.

So consider the case when $\theta \in [\frac{1}{1+\delta_A}, 1)$. Suppose there exists a point $p \in \mathcal{S}$ not in $\mathcal{W}$. WLOG, p is to the right of $\mathcal{W}$. This implies that $V_L \in \mathcal{W}$. Now, if it is not the case that every continuation payoff of every contract with payoff in $\mathcal{S} - \mathcal{W}$ is in $\mathcal{W}$, then there is a stationary payoff in $\mathcal{S}$. This follows from Kakutani's fixed point theorem and is a contradiction.

So pick a contract with payoff p and continuation payoff $p_1 \in \mathcal{W}$. Let (m', e') denote this contract's initial action. Since $p_1 \in \mathcal{W}$, it is the continuation payoff of an action sequence of the form $C(m^s, e^s, m)$ with payoff $(U_A^s + \frac{m}{1+\delta_A\rho}, U_P^s - \frac{m}{1+\delta_P\rho})$. Using arguments similar to before, one can show that $e' = e^s$. Now if $m' = \frac{m}{-\rho}$ then $p \in \mathcal{W}$. Contradiction. If $m' \neq \frac{m}{-\rho}$, then the slope between p and $(U_A^s + \frac{m}{1+\delta_A\rho}, U_P^s - \frac{m}{1+\delta_P\rho})$ is -1. Contradiction since $\theta \neq 1$.

Thus $\mathcal{W} = \mathcal{S}$. Now, using an argument similar to before, one can show that each $\mathcal{S}$-payoff is uniquely supported by a contract of the form $C(m^s, e^s, m)$.

Finally, let p be a payoff of V not in $\mathcal{W}$. WLOG, p is to the right of $\mathcal{W}$. Then $V_L \in \mathcal{W}$. Lemma 6 then implies that any contract with payoff p must have continuation payoff V_L. This simultaneously determines the initial action and all subsequent actions and the contract is completely characterized.

Case 2: $\theta < \frac{1}{1+\delta_A}$
Again, introduce the set $\mathcal{S}$. Like before, any $\mathcal{S}$-payoff is uniquely supported by an action sequence of the form $C(m^s, e^s, m)$. However, since $\theta < \frac{1}{1+\delta_A}$, any non-stationary $\theta < \frac{1}{1+\delta_A}$ will eventually break the participation constraint and is not feasible. Therefore, $\mathcal{S} = (U_A^s, U_P^s)$.

Suppose $\theta = \frac{\delta_P}{\delta_P+\delta_A}$. Consider the contract that supports the payoff V_R. Denote its time t continuation payoff by $(U_{R,A,t}, U_{R,P,t})$. Suppose $(U_{R,A,1}, U_{R,P,1}) \neq V_L$, then Lemma 6 implies that $(U_{R,A,2}, U_{R,P,2}) = V_R$ and the V_R contract is an oscillation contract. In this case, any payoff of V between $(U_{R,A,1}, U_{R,P,1})$ and V_R is supported by an oscillation contract. Moreover, any payoff of V to the left of $(U_{R,A,1}, U_{R,P,1})$ is supported by a contract whose continuation contract is the V_R-contract.

Now suppose $(U_{R,A,1}, U_{R,P,1}) = V_L$. If the continuation contract of the V_L-contract is the V_R-contract then the V_L-contract is an oscillation contract. And every non-stationary Pareto-optimal contract is an oscillation contract. If the continuation contract of the V_L-contract is not the V_R-contract then using the same argument as before, the V_L-contract is still an oscillation contract and contracts that support V are completely characterized.

Now suppose $\theta \in (\frac{\delta_P}{\delta_P+\delta_A}, \frac{1}{1+\delta_A})$. Using the same argument as before, one can show that either the V_L- or the V_R-contract is an oscillation contract. WLOG, the V_R-contract is an oscillation contract. Again, just like before, any payoff of V to the left of $(U_{R,A,1}, U_{R,P,1})$ is supported by a contract whose continuation contract is the V_R-contract. Moreover, Lemma 6 implies that any payoff between $(U_{R,A,1}, U_{R,P,1})$ and V_R is supported by a contract that explosively oscillates until it eventually becomes the V_R-contract.

Lastly, suppose $\theta < \frac{\delta_P}{\delta_P+\delta_A}$. Then Lemma 6 implies that every non-stationary Pareto-optimal contract exhibits damped oscillation and converges to (but never reaches) the stationary Pareto-optimal contract. □

References

ACEMOGLU, D., M. GOLOSOV, AND A. TSYVINSKI (2008): "Political economy of mechanisms," *Econometrica*, 76(3), 619–641.

AGUIAR, M., M. AMADOR, AND G. GOPINATH (2009): "Investment cycles and sovereign debt overhang," *Review of Economic Studies*, 76(1), 1–31.

LEHRER, E., AND A. PAUZNER (1999): "Repeated games with differential time preferences," *Econometrica*, 67(2), 393–412.

LEVIN, J. (2003): "Relational incentive contracts," *American Economic Review*, 93(3), 835–857.

OPP, M. M. (2012): "Expropriation risk and technology," *Journal of Financial Economics*, 103(1), 113 – 129.

RAY, D. (2002): "The time structure of self-enforcing agreements," *Econometrica*, 70(2), 547–582.

THOMAS, J., AND T. WORRALL (1988): "Self-enforcing wage contracts," *Review of Economic Studies*, 55(4), 541–54.

——— (1994): "Foreign direct investment and the risk of expropriation," *Review of Economic Studies*, 61(1), 81–108.