Attainable payoffs in repeated games with interdependent private information

David A. Miller*
UCSD

This version: January 13, 2009†
First version: November 26, 2003

Abstract

This paper proves folk theorems for repeated games with private information, communication, and monetary transfers, in which signal spaces may be arbitrary, signals may be statistically interdependent, and payoffs for each player may depend on the signals of other players.

JEL Classifications: C72 (Noncooperative games), C73 (Stochastic & dynamic games), D82 (Asymmetric & private information)

*Email: econdm@umich.edu. This paper is based on the first chapter of my Ph.D. dissertation at Stanford University. I thank Susan Athey, Ilya Segal, and Steven Tadelis, my dissertation committee, for their valuable guidance. I am indebted to Satoru Takahashi for pointing out a serious error in a previous version of this paper. I have also benefited from discussions with Jonathan Levin, Navin Kartik, Mark Machina, Andrzej Skrzypacz, Joel Sobel, and Joel Watson, and from the comments of seminar participants at UCSD and the 2005 Southwest Economic Theory Conference. This research was supported in part by the Stanford Institute for Economic Policy Research.

†Minor edits in May 2015

1 Introduction

Consider a class of infinitely repeated games with private information and communication. Each stage game begins with each player observing his own private signal. (Signals are identically and independently distributed over time, but not necessarily across players.) The players then have the opportunity to communicate by sending each other simultaneous public messages. After communicating, each player selects an action. In such games, players face the tension of collectively wanting to share information in order to coordinate their activities, while each of them individually faces an incentive to misrepresent his own information to the others. This tension complicates the problem of coordination by requiring that the players give each other incentives to reveal their private information.

Fudenberg, Levine, and Maskin (1994, hereafter FLM) prove a folk theorem (Theorem 8.1) for such games: as players become more and more patient, they can approximate, as average payoffs in a *perfect public equilibrium (PPE)* in the repeated game, any stage game payoffs that are "feasible" and "individually rational." FLM focus mainly on games with imperfect public monitoring, in which actions, rather than information, are private. To apply their public monitoring results to games with private information, FLM invoke three limiting assumptions: (i) the signal space is finite, (ii) the private signals are statistically independent, (iii) each player's payoff does not depend on other players' signals ("independent payoffs").

In this paper I prove a folk theorem for games with monetary payments that bypasses these limitations. In such games, sufficiently patient players can support any incentive compatible and individually rational payoffs in the stage game as the average payoffs in a stationary PPE of the repeated game. Transferable utility simplifies the provision of incentives, allowing a direct proof.

In a previously circulated version of this paper I claimed to also have proven folk theorems for games without monetary payments, but a significant error in my proof was discovered by Satoru Takahashi.

In addition to the folk theorems, this paper extends the "mechanism design approach," initiated by Athey and Bagwell (2001) for the case of collusion with hidden costs, to apply to a general class of games. In the mechanism design approach, the problem of designing a PPE in the repeated private information game is converted to the problem of constructing a *recursive mechanism*, as if there were an institution that in each stage receives the players' messages and then instructs them how to act. The mechanism is recursive because in each period it promises them a certain level of utility in the continuation game, and when the next period begins this promised utility determines the mechanism to be used in the continuation game. The combination of monetary payments at the end of the period and changes in promised future utility provide the incentives for players to reveal their information truthfully. The problem of constructing a recursive mechanism, in turn, can be expressed (in part) as a static mechanism design problem, to which standard techniques apply. The mechanism design approach is compared with the FLM approach in Section 1.1 and developed formally in Section 3.

1.1 Comparison with the FLM approach

A comparison with FLM's approach to games with private information shows why the mechanism design approach yields a broader set of conclusions. In both approaches, the problem of constructing a PPE in the repeated game is successively converted into simpler problems, until the result is attained. At each step of simplification, there are conditions that must be imposed. The FLM approach, as described above, begins by converting a private information game to a public monitoring game. Specifically, FLM converted a private information game to a public monitoring game with a "product structure," where each player's action gives rise to a separate monitoring indicator, and these indicators are statistically independent. The monitoring indicator for each player is the realization of her message, but for these messages to be statistically independent there can be no interdependence among the private signals that the players observe. Hence FLM's first step imposes statistical independence on the signals.[1]

The second step of the FLM approach is to prove that in games with product structures any smooth set $\mathcal{W}$ of feasible and individually rational payoff profiles is "decomposable on tangent hyperplanes." That is, for each payoff profile on the boundary of $\mathcal{W}$, the hyperplane tangent to $\mathcal{W}$ at that profile separates $\mathcal{W}$ from some extreme payoff profile, and the outcome function that yields this extreme payoff profile can be supported with appropriate incentives by a set of continuation rewards that lie in the hyperplane. FLM were able to prove decomposability on tangent hyperplanes by assuming finite signal spaces and independent payoffs, and then applying their results for public monitoring games with product structures.

Finally, FLM proved that if $\mathcal{W}$ is decomposable on tangent hyperplanes and the players are sufficiently patient then any point in $\mathcal{W}$ can be supported as the average payoffs of a PPE. Since $\mathcal{W}$—as an arbitrary smooth set of attainable and individually rational payoff profiles—can include points arbitrarily close to any attainable and individually rational payoff profile, this last step yields the folk theorem.

The mechanism design approach follows a different path. As described above, the approach begins by shifting the focus from PPE to recursive mechanisms. This step is based on the dynamic programming approach developed by Abreu, Pearce, and Stacchetti (1990), FLM, and others, but with an additional simplification. Though in the game they can send any arbitrary messages and take their subsequent actions non-cooperatively, when participating in the recursive mechanism

[1] A more minor limitation is that for this step FLM need to assume that the players take no additional actions subsequent to sending their messages—that an outcome is automatically selected based on the realized messages, such as by an exogenous mechanism designer. I allow the players to take actions non-cooperatively after communicating. See Section 3 for details.

the players are assumed to reveal their information directly to each other and to obey the mechanism's instructions.

The second step of the mechanism design approach is to convert the problem of designing a recursive mechanism into a static mechanism design problem, by lumping together the monetary payments and changes in promised future utility into a "static transfer function." Static mechanisms should satisfy *interim incentive compatibility (IIC)*[2]—the requirement that each player be willing to reveal his information truthfully given his interim beliefs about other players.

When considering games with monetary payments, I look for mechanisms that satisfy *ex post budget balance (EPBB)*—the requirement that the static transfers sum to zero across players after every realization of the private signals. Under a condition on the probability measure over private signals ("Condition C", after d'Aspremont, Crémer, and Gérard-Varet (2003), henceforth ACGVa), any IIC-implementable outcome function is also implementable with EPBB. When there are two players, Condition C is equivalent to statistical independence, but when there are three or more players Condition C is satisfied for generic distributions if the signal space is finite (d'Aspremont, Crémer, and Gérard-Varet (2004), henceforth ACGVb). For games with monetary payments, it is then possible to prove the folk theorem directly, by explicitly constructing a stationary equilibrium.

In sum, the mechanism design approach allows me to impose milder conditions than FLM: Condition C is generically satisfied when the signal space is finite and there are three or more players, while the notion that transfers should be uniformly bounded is implicit in any plausible application.

1.2 Other related literature

Following FLM, a number of papers in the literature investigate the possibility of efficient equilibria in repeated games with private information and communication. None of them, however, relax more than two of the three key limitations in FLM's folk theorem, and each is tied to the structure of a particular game.

Athey and Bagwell (2001) study a collusion game with discrete types, statistical independence, and independent payoffs. They use the mechanism design approach to construct optimal equilibria for impatient firms. Athey and Bagwell (2008) examine a similar model in which each firm's costs are correlated over time.

Aoyagi (2003, 2007) considers collusion in repeated auctions. Aoyagi (2003) allows continuous signal spaces, statistical interdependence, and interdependent payoffs, but the equilibrium he constructs is not asymptotically efficient. Aoyagi (2007) restricts the signal space to be finite, and constructs an equilibrium that is asymptotically efficient as the players become more patient. In work conducted simultaneously with the present paper, Martin and Vergote (2004) expand on Aoyagi's results by constructing an asymptotically efficient equilibrium with continuous signal

[2] Often called "Bayesian incentive compatibility."

spaces, but with statistical independence and independent payoffs. All of the papers mentioned in this paragraph use bid rotation schemes, an approach that does not extend to general games.[3]

In Miller (2007), I consider the same model as in the present paper, but add the restriction that the equilibrium should be robust to the possibility that players may observe arbitrary, payoff-irrelevant signals that are correlated with the payoff-relevant signals. I show that under this robustness restriction efficiency cannot be approximated by any PPE, and I characterize optimal robust equilibria. In Athey and Miller (2007) we consider the setting of repeated trade under this robustness restriction, with continuous signal spaces, but assuming statistical independence and independent payoffs.

2 The model

The dynamic game consists of infinite repetitions of a stage game. $\mathcal{N} = \{1, \ldots, n\}$ is the set of players. The players share a common discount factor $\delta \in (0, 1)$.

At the beginning of each stage, each player $i \in \mathcal{N}$ observes a private signal drawn from the set Θ_i, according to the common prior probability measure ϕ on $\Theta \equiv \Theta_1 \times \cdots \times \Theta_N$.[4] I write ϑ for the random variable distributed according to ϕ; θ is a typical realization. The players then send simultaneous public messages; each player i's message space is $\mathcal{M}_i \supset \Theta_i$,[5] with $\mathcal{M} \equiv \mathcal{M}_1 \times \cdots \times \mathcal{M}_N$. After the communication is concluded, each player i simultaneously chooses an action from $\mathcal{X}_i$, with $\mathcal{X} \equiv \mathcal{X}_1 \times \cdots \times \mathcal{X}_N$; all actions are publicly observed. Each player i receives a payoff from the actions and signals; this is given by $\pi_i : \Theta \times \mathcal{X} \to \mathbb{R}_+$, with $\pi \equiv \{\pi_1, \ldots, \pi_N\}$. Assume that π_i is integrable and uniformly bounded for all i, and that there exists a perfect Bayesian equilibrium in the stage game.

For part of this paper, I also allow the players to voluntarily "pay" utility to one another. I model this by assuming that they have utility that is quasilinear in "money," which is a continuous good in zero net supply. Monetary payments are bilateral and unidirectional: each player can pay any non-negative amount of money to each other player, and (for simplicity) such payments cannot be refused by the recipient. All such payments are publicly observed. Player i's utility in the stage game is equal to her payoff plus any monetary payments she receives minus any monetary payments she gives to other players.

[3]Several papers have investigated collusion in repeated auctions without communication, when the auctioneer releases only the identity of the winner. Skrzypacz and Hopenhayn (2004) show that in such games efficiency cannot be approximated by any PPE. Blume and Heidhues (2006) construct an equilibrium in private strategies that outperforms any PPE. It is not yet known whether there exist asymptotically efficient equilibria in their setting.

[4]The measure ϕ is associated with a product topology that I do not specify. I assume that all functions are measurable with respect to this topology; furthermore, I assume that the conditional probability measures $\{\{\phi_{-i|i}(\cdot|\theta_i)\}_{\theta_i \in \Theta_i}\}_{i=1}^N$ are well defined.

[5]More generally, since messages can be relabeled, what is actually required is that there exist a homeomorphism from Θ_i to a subspace of $\mathcal{M}_i$. Assuming $\Theta_i \subset \mathcal{M}_i$ simplifies the notation.

This setup allows an arbitrary measurable signal space Θ_i for each player and statistical interdependence among signals; it also allows each player's payoffs to depend on the entire vector of signals.

The equilibrium concept is *perfect public equilibrium* (PPE), after FLM, which is a refinement of perfect Bayesian equilibrium in which players play behavior strategies that in each period are conditioned only on the public history and their current private information. The public history comprises the messages, actions, and payments that have transpired, which are common knowledge. Since the mechanism design approach will simplify matters in the following section, I relegate the definitions of strategies and equilibria to Appendix A.

3 The mechanism design approach

This section develops a "mechanism design approach" to repeated games with private information. Abreu, Pearce, and Stacchetti (1986, 1990) showed that PPEs can be constructed recursively without loss of generality with respect to attainable utilities. The recursive construction yields a dynamic program in which the future path of play is summarized by a profile of average continuation utilities. I interpret this program as a recursive mechanism design problem.

This approach greatly simplifies the notation and analysis, but its validity must be proven. This is accomplished in Theorem 1, which shows that if the players are sufficiently patient then this approach yields (nearly) the same set of utilities that are attainable in PPE. Recursive mechanisms and their properties are defined below; Theorem 1 is stated at the end of the section.

3.1 Recursive mechanisms

A *stage mechanism* consists of an outcome function x, a payment function t, and a continuation reward function w.

Definition 1. A *stage mechanism* is a triplet $\langle x, t, w\rangle : \Theta \to \mathcal{X} \times \mathbb{R}^N \times \mathbb{R}^N$ such that t and w are integrable and $\sum_i t_i(\theta) = 0$ for all $\theta \in \Theta$.

In addition to a collection of stage mechanisms, a recursive mechanism also specifies a set $\mathcal{V}$ of promised (average) utility state variables over which it is defined, as well as an initial condition v^0. Formally:

Definition 2. A *recursive mechanism* is a triplet $\langle \mathcal{V}, \{\langle x(\cdot; v), t(\cdot; v), w(\cdot; v)\rangle : v \in \mathcal{V}\}, v^0\rangle$, abbreviated as $\left\langle \mathcal{V}, \left\{\langle x, t, w\rangle(\cdot; v)\right\}, v^0\right\rangle$, such that:

(i) $\mathcal{V} \subset \mathbb{R}^N$,

(ii) $\{\langle x, t, w\rangle (\cdot; v)\}$ is a collection of stage mechanisms indexed by $v \in \mathcal{V}$,

(iii) $v^0 \in \mathcal{V}$.

Such a mechanism is called "recursive" because in each period the stage mechanism is selected based on the promised utility v carried over from the previous period, and the stage mechanism generates a continuation reward $w(\cdot; v)$ that becomes the promised utility for the next period. In contexts that do not invite confusion, I drop the $(\cdot; v)$ notation for clarity.

Given a signal profile $\theta \in \Theta$ and a vector of announcements $\hat{\theta} \in \Theta$ in a stage mechanism $\langle x, t, w \rangle$, each player's total utility is her stage game utility plus the present value of her continuation reward:

$$u_i(\theta, \hat{\theta}; \delta, \langle x, t, w \rangle) \equiv \pi_i(\theta, x(\hat{\theta})) + t_i(\hat{\theta}) + \frac{\delta}{1-\delta} w_i(\hat{\theta}). \tag{1}$$

The space of promised utilities $\mathcal{V}$ must be feasible: the utilities that are promised must actually be attainable in the game, and the mechanism must actually deliver what it promises.

Definition 3. With respect to $\delta \in (0,1)$, a recursive mechanism $\left\langle \mathcal{V}, \left\{ \langle x, t, w \rangle (\cdot; v) \right\}, v^0 \right\rangle$ is *feasible* if the following are satisfied:

(i) Attainability: $\mathcal{V} \subset \text{co}\left(\left\{ v \in \mathbb{R}^N : \sum_i v_i = \sum_i \mathbb{E}\left[\pi_i\left(\theta, x(\theta)\right) \right] \text{ for some } x : \Theta \to \mathcal{X} \right\} \right)$ (where $\text{co}(\cdot)$ is the convex hull operator);

(ii) Promise keeping: $v = (1-\delta)\, \mathbb{E}\left[u(\theta, \theta; \delta, \langle x, t, w \rangle (\cdot; v)) \right]$ for all $v \in \mathcal{V}$;

(iii) Coherence: $w\left(\theta; v\right) \in \mathcal{V}$ for all $\theta \in \Theta$ and all $v \in \mathcal{V}$.

3.2 The equivalent static mechanism

The following definition transforms the dynamic problem into an equivalent static problem.

Definition 4. A *static mechanism* is a pair $\langle x, y \rangle : \Theta \to \mathcal{X} \times \mathbb{R}^N$ such that y is integrable. Given $\delta \in (0,1)$ and a stage mechanism $\langle x, t, w \rangle$, the *equivalent static mechanism* is a static mechanism $\langle x, y \rangle$ such that $y(\theta) \equiv t(\theta) + \frac{\delta}{1-\delta} w(\theta)$.

Accordingly, y is termed the *equivalent static transfer function*. Note that by construction any equivalent static transfer function is integrable. Naturally, in a static mechanism players seek to maximize $\pi_i(\theta, x(\hat{\theta})) + y_i(\hat{\theta})$.

Definition 5. A static transfer function y is *ex post budget balanced (EPBB)* if $\sum_i y_i(\theta) = 0$ for all θ. A static mechanism $\langle x, y \rangle$ is EPBB if y is EPBB.

Definition 6. A static mechanism $\langle x, y \rangle$ is *uniformly bounded* if there exists a bound $B < \infty$ such that $\left| y_i(\theta) \right| < B$ for all $\theta \in \Theta$ and all $i \in \mathcal{N}$.

3.3 Incentive compatibility

Incentive compatibility requires that each player must prefer to make a truthful announcement rather than deviate to some other "on-schedule" announcement (that is, claiming to have observed some other signal $\hat{\theta}_i \neq \theta_i$, with $\hat{\theta}_i \in \Theta_i$.). Interim incentive compatibility (IIC) requires that each player, taking interim expectations over the other players' announcements, prefer to make a truthful announcement rather than deviate to some other on-schedule announcement. IIC can be thought of as incentive compatibility that applies when all announcements are made simultaneously.

Definition 7. With respect to a discount factor $\delta \in (0,1)$, a stage mechanism $\langle x,t,w\rangle$ is *interim incentive compatible (IIC)* if

$$\mathbb{E}\left[u_i(\vartheta,\vartheta;\delta,\langle x,t,w\rangle)\middle|\vartheta_i=\theta_i\right] \geq \mathbb{E}\left[u_i(\vartheta,(\hat{\theta}_i,\vartheta_{-i});\delta,\langle x,t,w\rangle)\middle|\vartheta_i=\theta_i\right] \tag{2}$$

for all $\hat{\theta}_i \in \Theta_i$, for ϕ_i-almost all $\theta_i \in \Theta_i$, and for all $i \in \mathcal{N}$. A recursive mechanism $\left\langle \mathcal{V}, \left\{\langle x,t,w\rangle(\cdot;v)\right\}, v^0\right\rangle$ is IIC if $\langle x,t,w\rangle(\cdot;v)$ is IIC for all $v \in \mathcal{V}$. A static mechanism $\langle x,y\rangle$ is IIC if

$$\mathbb{E}\left[\pi_i(\vartheta,x(\vartheta))+y_i(\vartheta)\middle|\vartheta_i=\theta_i\right] \geq \mathbb{E}\left[\pi_i(\vartheta,x(\hat{\theta}_i,\vartheta_{-i}))+y_i(\hat{\theta}_i,\vartheta_{-i})\middle|\vartheta_i=\theta_i\right] \tag{3}$$

for all $\hat{\theta}_i \in \Theta_i$, for ϕ_i-almost all $\theta_i \in \Theta_i$, and for all $i \in \mathcal{N}$.

Note that a stage mechanism is IIC if and only if its equivalent static mechanism is IIC.

3.4 Individual rationality

To support the recursive mechanism as an equilibrium, if a player is ever observed to deviate (an "off-schedule" deviation), then he should be punished by switching to a PPE that gives him a low payoff starting in the next period, rather than $w_i(\theta)$. If the punishment outweighs his gain from deviating, then it is "individually rational" for him not to choose an observable deviation.

Definition 8. With respect to a discount factor $\delta \in (0,1)$, a *punishment payoff profile* is a payoff profile $p \in \mathbb{R}^n$ such that for each player i there exists a PPE that yields a payoff profile $p^{(i)} \in \mathbb{R}^n$ with $p_i^{(i)} = p_i$.

By Abreu, Pearce, and Stacchetti (1990) Theorem 6, if p is a punishment payoff profile with respect to $\hat{\delta} < 1$, then p is also a punishment payoff profile with respect to any $\delta \in [\hat{\delta}, 1]$.[6]

Definition 9. With respect to a discount factor $\delta \in (0,1)$ and a punishment payoff profile $p \in \mathbb{R}^n$, a stage mechanism $\langle x,t,w\rangle$ is *ex post individually rational* (EPIR) if the following are satisfied for all $i \in \mathcal{N}$:

[6]Note that Abreu, Pearce, and Stacchetti (1990) Theorem 6 is stated in terms of total payoffs. Translating it to average payoffs yields the claim.

(i) for ϕ_i-almost all $\theta_i \in \Theta_i$,

$$\mathbb{E}\left[u_i\left(\vartheta, \vartheta; \delta, \langle x, t, w\rangle\right)\middle|\vartheta_i = \theta_i\right] \geq \mathbb{E}\left[\sup_{\chi \in \mathcal{X}^*(\vartheta; i)} \pi_i(\vartheta, \chi)\middle|\vartheta_i = \theta_i\right] + \frac{\delta}{1-\delta} p_i, \tag{4}$$

where $\mathcal{X}^*(\theta; i)$ is the set of action profiles that can arise in perfect Bayesian equilibrium in the remainder of the stage game given θ, after player i makes an off-schedule announcement;[7]

(ii) for almost all $\theta \in \Theta$ and all $\hat{\theta}_i \in \Theta_i$,

$$u_i\left(\theta, (\hat{\theta}_i, \theta_{-i}); \delta, \langle x, t, w\rangle\right) \geq \sup_{\chi_i \neq x_i(\hat{\theta}_i, \theta_{-i})} \pi_i\left(\theta, \left(\chi_i, x_{-i}(\hat{\theta}_i, \theta_{-i})\right)\right) + \frac{\delta}{1-\delta} p_i; \tag{5}$$

(iii) for almost all $\theta \in \Theta$,

$$t_i(\theta) + \delta \mathbb{E}\left[u_i\left(\vartheta, \vartheta; \delta, \langle x, t, w\rangle\right)\right] \geq \frac{\delta}{1-\delta} p_i. \tag{6}$$

A recursive mechanism $\left\langle \mathcal{V}, \left\{\langle x, t, w\rangle(\cdot; v)\right\}, v^0\right\rangle$ is EPIR if $\langle x, t, w\rangle(\cdot; v)$ is EPIR for all $v \in \mathcal{V}$.

Eq. 4 ensures that player i's expected utility, given his own signal θ_i, is at least what he could get by choosing an off-schedule announcement, under strategies specifying that after he deviates the other players would not pay him any money in that period. Eq. 5 ensures that, regardless of whether his announcement in the communication phase was truthful, player i's utility starting in the action phase from playing as if in equilibrium is at least what he could get by choosing a deviant action, also under strategies specifying that after he deviates the other players would not pay him any money in that period. Eq. 6 ensures that once actions have been taken, player i prefers to make his specified payment, rather than pay zero.[8]

EPIR is helpful for justifying the mechanism design approach, but for the purpose of proving the folk theorem it is stronger than necessary. Since payoffs are uniformly bounded, if the punishment payoffs are strictly less than the equilibrium payoffs, then there will always be some discount factor such that the stage mechanism satisfies EPIR. This motivates a simpler notion of individual rationality.

Definition 10. With respect to a payoff profile $p \in \mathbb{R}^n$, a payoff profile $v \in \mathbb{R}^n$ is *strictly individually rational* (SIR) if $\min_{i \in \mathcal{N}}\{v_i - p_i\} > 0$. A set of payoff profiles $\mathcal{V} \subset \mathbb{R}^n$ is SIR if $\inf_{v \in \mathcal{V}} \min_{i \in \mathcal{N}}\{v_i - p_i\} > 0$.

[7] For simplicity, we can take $\mathcal{X}^*(\theta; i)$ to be the set of equilibria given θ when the other players have passive beliefs after observing player i's off-schedule announcement.

[8] Without loss of generality, strategies can be constrained to specify that player i is supposed to pay money to any other player(s) only if $t_i(\theta) < 0$.

When utility is transferable, there is the possibility that even if a mechanism violates the SIR constraint for player i, the other players' SIR constraints may be slack. Then a money payment to player i from players with slack constraints can yield a new mechanism that satisfies SIR and implements the same outcome function x. An outcome function is thus "potentially IR" if there exists an SIR mechanism that implements it.

Definition 11. With respect to a payoff profile $p \in \mathbb{R}^n$, a static mechanism $\langle x, y \rangle$ is *potentially individually rational* (PIR) if $\sum_i \mathbb{E}\left[\pi_i(\vartheta, x(\vartheta))\right] > \sum_i p_i$ for all $i \in \mathcal{N}$. An outcome function x is PIR there exists a static transfer function$y$ such that $\langle x, y \rangle$ is PIR.

3.5 Justifying the mechanism design approach

The following theorem justifies the mechanism design approach to repeated games with private information. The proof, in Appendix B, is conceptually simple. When the players select their recursive mechanism (this selection is not modeled), they are collectively acting as their own mechanism designer. Their recursive mechanism specifies the equilibrium path they should follow; incentive compatibility discourages on-schedule deviations while the trigger punishment discourages off-schedule deviations. Hence a feasible, IIC, and EPIR recursive mechanism describes an equilibrium path and implies the existence of a PPE that supports that path. In the other direction, the equilibrium path of any SIR PPE can be described by a history-dependent dynamic mechanism. Similarly to the logic of Abreu, Pearce, and Stacchetti (1986, 1990), it is without loss of generality with respect to attainable payoffs to substitute a recursive mechanism for a dynamic mechanism.

Theorem 1. *Given a payoff profile $p \in \mathbb{R}^n$,*

(i) for any $\delta \in (0,1)$, if p is a punishment payoff profile and a recursive mechanism $\left\langle \mathcal{V}, \left\{ \langle x, t, w \rangle(\cdot; v) \right\}, v^0 \right\rangle$ is feasible, IIC, and EPIR, then there exists a PPE that yields the same announcements, actions, and net payments along the equilibrium path;

(ii) if $\mathcal{V} \subset \mathbb{R}^N$ is the set of average utility profiles yielded at the beginning of any period along the equilibrium path of some PPE, $\mathcal{V}$ is SIR with respect to p, and there exists $\hat{\delta} \in (0,1)$ with respect to which p is a punishment payoff profile, then there exists $\underline{\delta} \in [\hat{\delta}, 1)$ such that for any $\delta \in (\underline{\delta}, 1)$ and any $v \in \mathcal{V}$ there exists a recursive mechanism with initial promised utility v that is feasible, IIC, and EPIR.

4 Budget balanced static mechanisms

Leave aside for a moment the repeated game structure, and consider the problem of designing static mechanisms that satisfy both IIC and EPBB. It is widely known that in many situations efficient IIC static mechanisms can be implemented with EPBB if IR constraints are ignored (see, e.g., d'Aspremont and Gérard-Varet 1979; Johnson, Pratt, and Zeckhauser 1990; and Fudenberg, Levine, and Maskin 1995). Most generally, ACGVb showed that, in games with finite type spaces, any IIC-implementable outcome function can be implemented with EPBB if and only if their "Condition C" is satisfied.

Definition 12. A game satisfies *Condition C* if, for $\lambda = (1, \dots, 1)$, any function $R : \Theta \to \mathbb{R}$ can be expressed as $R(\theta) = \sum_i \lambda_i r_i(\theta)$, where $r : \Theta \to \mathbb{R}^N$ satisfies

$$\mathbb{E}\big[r_i(\vartheta)\big|\vartheta_i = \theta_i\big] \geq \mathbb{E}\big[r_i(\hat{\theta}_i, \vartheta_{-i})\big|\vartheta_i = \theta_i\big] \tag{7}$$

for ϕ_i-almost all $\theta_i \in \Theta_i$, for all $\hat{\theta}_i$ in the support of ϕ_i, and for all $i \in \mathcal{N}$.

Lemma 1 (ACGVb)**.** *If Condition C is satisfied and the static mechanism $\langle x, y\rangle$ is IIC and uniformly bounded, then there exists a static mechanism $\langle x, y^\lambda\rangle$ that is IIC, uniformly bounded, and satisfies EPBB.*

The only implicit requirement for this conclusion is that the conditional expectations of the transfer function are well-defined almost everywhere, which is assured by the assumption that $\langle x, y\rangle$ is IIC.

Remark 1 (Statistical independence). When $\theta_1, \dots, \theta_N$ are statistically independent according to ϕ, the conclusion of Lemma 1 is always satisfied. To see this, given an IIC and uniformly bounded mechanism $\langle x, y\rangle$, let

$$\hat{y}_i(\theta) = \mathbb{E}\big[y_i(\vartheta)\big|\vartheta_i = \theta_i\big] - \frac{1}{(N-1)} \sum_{j \neq i} \mathbb{E}\big[y_j(\vartheta)\big|\vartheta_j = \theta_j\big]. \tag{8}$$

Then $\langle x, \hat{y}\rangle$ is uniformly bounded and satisfies $\sum_i y_i^\lambda(\theta) = 0$ for all $\theta \in \Theta$ and for all $i \in \mathcal{N}$. It is also IIC, since the second term does not depend on θ_i while the first term satisfies

$$\begin{aligned} &\mathbb{E}\big[\pi_i(\vartheta, x(\hat{\theta}_i, \vartheta_{-i})) + y_i(\hat{\theta}_i, \vartheta_{-i})\big|\vartheta_i = \theta_i\big] \\ &\qquad = \mathbb{E}\Big[\pi_i(\vartheta, x(\hat{\theta}_i, \vartheta_{-i})) + \mathbb{E}\big[y_i(\hat{\theta}_i, \vartheta_{-i})\big|\vartheta_i = \hat{\theta}_i\big]\Big|\vartheta_i = \theta_i\Big] \end{aligned} \tag{9}$$

for all $\theta_i \in \Theta_i$, for all $\hat{\theta}_i \in \Theta_i$, and for all $i \in \mathcal{N}$.

Furthermore, for $N \geq 3$ ACGVb prove, in a model with finite Θ, that Condition C is satisfied for an open and dense set of probability distributions. Unfortunately their method of proof rests crucially on finiteness, so it is an open question whether Condition C is satisfied generically for infinite Θ.

5 Folk theorem for games with transfers

I now return to the repeated game context to prove the folk theorem with monetary payments, using grim trigger punishments for off-schedule deviations. The theorem shows that the problem of finding a PPE with an individually rational payoff profile can essentially be solved by finding a static mechanism that is IIC and uniformly bounded, and whose outcome function provides the desired sum of payoffs.

Theorem 2 (Folk theorem with money). *For any payoff profile $p \in \mathbb{R}^n$, suppose that (i) the static mechanism $\langle x, y\rangle$ is IIC and uniformly bounded; (ii) $\langle x, y\rangle$ is PIR; (iii) either $\langle x, y\rangle$ is EPBB or Condition C is satisfied; and (iv) there exists $\hat{\delta} \in (0,1)$ such that p is a punishment payoff profile. Then for any payoff profile $v^0 \in \mathbb{R}^n$ that is SIR and satisfies $\sum_i v_i^0 = \sum_i \mathbb{E}\big[\pi_i(\vartheta, x(\vartheta))\big]$, there exists $\underline{\delta} \in (\hat{\delta}, 1)$ such that, for all $\delta > \underline{\delta}$, there exists a stationary PPE that yields a payoff profile of v^0.*

Proof. The mechanism I construct is stationary on the equilibrium path; i.e., it uses the same stage mechanism in every period, or $\langle \bar{x}, \bar{t}, \bar{w}\rangle(\cdot; v) = \langle x, t, w\rangle$ for all v.

Suppose that Condition C is satisfied, and the mechanism $\langle x, y\rangle$ is IIC, PIR, and uniformly bounded. By Lemma 1, $\langle x, y\rangle$ is also EPBB without loss of generality. Suppose that v^0 satisfies SIR and $\sum_i v_i^0 = \sum_i \mathbb{E}\big[\pi_i(\vartheta, x(\vartheta))\big]$. Let $\hat{v} \equiv \mathbb{E}\big[\pi(x(\vartheta)) + y(\vartheta)\big]$.

I begin by constructing the stage mechanism $\langle x, t, w\rangle$ and establishing that it satisfies IIC. For all θ, let $t_i(\theta) = y_i(\theta) + v_i^0 - \hat{v}_i$ and $w_i(\theta) = v_i^0$. As required, the monetary payments are uniformly bounded and sum to zero for all θ, because $\langle x, y\rangle$ is EPBB and uniformly bounded, and $\sum_i (v_i^0 - \hat{v}_i) = 0$. Observe that $\langle x, t, w\rangle$ satisfies IIC, since its equivalent static mechanism differs from $\langle x, y\rangle$ by only a lump sum.

Next I show that the recursive mechanism is feasible. Since $w(\theta) = v^0$ for all θ, the recursive mechanism is $\big\langle \{v^0\}, \langle \bar{x}, \bar{t}, \bar{w}\rangle(\cdot; v), v^0 \big\rangle$. Observe that the attainability and coherence conditions for feasibility are satisfied by construction. The promise keeping condition for feasibility is also satisfied:

$$(1-\delta)\,\mathbb{E}\big[u(\vartheta, \vartheta; \delta, \langle x, t, w\rangle)\big] = (1-\delta)\,\mathbb{E}\Big[\hat{v} + v^0 - \hat{v} + \frac{\delta}{1-\delta} v^0\Big] = v^0. \tag{10}$$

For EPIR, sufficient conditions are that for all i, θ, and $\hat{\theta}_i$,

$$v_i^0 \geq p_i + \frac{1-\delta}{\delta}\mathbb{E}\bigg[\sup_{\chi \in \mathcal{X}^*(\vartheta; i)} \pi_i(\vartheta, \chi) - \pi_i\big(\vartheta, x(\vartheta)\big) - t_i(\vartheta)\bigg|\vartheta_i = \theta_i\bigg], \tag{11}$$

$$v_i^0 \geq p_i + \frac{1-\delta}{\delta}\left(\sup_{\chi_i \neq x_i(\hat{\theta}_i, \theta_{-i})} \pi_i\big(\theta, \big(\chi_i, x_{-i}(\hat{\theta}_i, \theta_{-i})\big)\big) - \pi_i\big(\theta, x(\hat{\theta}_i, \theta_{-i})\big) - t_i(\hat{\theta}_i, \theta_{-i})\right), \tag{12}$$

$$v_i^0 \geq p_i + \frac{1-\delta}{\delta} t_i(\theta). \tag{13}$$

These are satisfied for sufficiently high $\delta < 1$ because $v_i^0 > p_i$ and $\langle x, y \rangle$ is uniformly bounded. Let $\underline{\delta}$ be the minimal δ such that Eqs. 11–13 are satisfied for all i.

Finally, by Theorem 1, for any $\delta \in (\underline{\delta}, 1)$ there exists a PPE that yields the same messages, actions, and net payments along the equilibrium path, and hence yields a payoff profile of v^0. ■

6 Conclusion

This paper proves a folk theorems for games with private information, communication, and monetary transfers. Based on the work of Abreu, Pearce, and Stacchetti (1986, 1990), FLM, Athey and Bagwell (2001), and others, I adapt static mechanism design definitions and techniques to the analysis of repeated game equilibria. This mechanism design approach not only simplifies the description of equilibria, but also leads to a simpler and more broadly applicable extension of FLM, Theorem 8.1. Rather than adapt results from games with hidden action, where the statistical identifiability of actions is of serious concern, it is simpler to note that, under Condition C, IIC-implementability implies enforceability with respect to the relevant. This direct approach allows the folk theorem to be extended to games with interdependent signals, interdependent payoffs, and arbitrary signal spaces.

The notions of IIC and EPIR that I employ are "almost sure" in the sense of Balder (1996). It is not necessary to impose stronger conditions under the assumption (heretofore maintained implicitly) that there exists some perfect Bayesian equilibrium in the continuation game after any event that some $\theta \in \Theta$ is realized on which IIC or IR fails for some player. This is consistent with the construction of a perfect Bayesian equilibrium in the repeated game because failures of EPIR or IIC occur with zero probability, and so affect neither interim incentives nor the value of the recursive mechanism.

Appendix A Equilibrium

This section provides the notation necessary to describe strategies and equilibria. A public outcome in period τ is $h^\tau \in \mathcal{M} \times \mathcal{X} \times (\mathbb{R}_+^{N-1})^N$, while a private outcome for player i is $h_i^\tau \equiv (\theta_i, h^\tau)$. A pure stage strategy for player i is a triplet $s_i = \langle \hat{m}_i, \hat{x}_i, \hat{t}_i \rangle$ that contains a reporting rule $\hat{m}_i : \Theta_i \to \mathcal{M}_i$, an action rule $\hat{x}_i : \Theta_i \times \mathcal{M} \to \mathcal{X}_i$, and a payment rule $\hat{t}_i : \Theta_i \times \mathcal{M} \times \mathcal{X} \to \mathbb{R}_+^{N-1}$, with $\hat{t}_{i,j}(\theta_i, \mu, \chi) \in \mathbb{R}_+$ indicating the amount that player i pays to player j. A pure stage strategy profile is a vector $s \equiv (s_1, \ldots, s_N)$, or, equivalently, $s \equiv \langle \hat{m}, \hat{x}, \hat{t} \rangle$. The public history at the end of period τ is $H^\tau \equiv (h^1, \ldots, h^\tau)$, and the private history for player i is $H_i^\tau \equiv (h_i^1, \ldots, h_i^\tau)$; H^0 and H_i^0 are null histories. A "stage-behavior" strategy for player i is a function σ_i that maps player i's private history (of any length) to a probability distribution over pure stage strategies.[9] For convenience, I allow players to choose their stage strategies using an arbitrary public randomization

[9]Strictly speaking, a behavior strategy would specify a probability distribution over reports as a function of private type, a probability distribution over actions as a function of private type and message vector, and so on. Stage-behavior strategies are both more convenient and more general.

device, so that (indulging in some abuse of notation) $\sigma(\{H_i^{\tau-1}\}_{i=1}^N)$ need not be statistically independent. Given a stage-behavior strategy profile σ and a set of private histories $\{H_i^{\tau-1}\}_{i=1}^N$, the ex post stage game payoff for player i in period τ is

$$\begin{aligned}\hat{\pi}_i(\theta_i;\sigma(\{H_i^{\tau-1}\}_{i=1}^N)) = \pi(\theta,\hat{x}(\theta,\hat{m}(\theta))) &+ \sum_{j\neq i}\hat{t}_{j,i}\big(\theta_j,\hat{m}(\theta),\hat{x}(\theta,\hat{m}(\theta))\big)\\ &- \sum_{j\neq i}\hat{t}_{i,j}\big(\theta_j,\hat{m}(\theta),\hat{x}(\theta,\hat{m}(\theta))\big),\end{aligned} \tag{14}$$

where the understanding is that $\langle\hat{m},\hat{x},\hat{t}\rangle$ is the outcome of the randomization specified by the stage strategy $\sigma(\{H_i^{\tau-1}\}_{i=1}^N)$. A public strategy for player i is a stage-behavior strategy such that $\sigma_i(H_i^{\tau-1}) = \sigma_i(\tilde{H}_i^{\tau-1})$ whenever $H^{\tau-1} = \tilde{H}^{\tau-1}$; i.e., a strategy in which player i ignores her private history. When σ_i is a public strategy, I write $\sigma_i(H^{\tau-1})$ for simplicity. In the repeated game, given a profile of stage-behavior strategies σ, the value to player i of the private history H_i^τ is

$$\hat{v}_i(H_i^\tau;\sigma) = (1-\delta)\,\mathbb{E}\left[\sum_{\tilde{\tau}=\tau+1}^{\infty}\delta^{\tilde{\tau}-1}\hat{\pi}_i(\vartheta_i^{\tilde{\tau}};\sigma(\{H_i^{\tilde{\tau}-1}\}_{i=1}^N))\,\middle|\,H_i^\tau\right], \tag{15}$$

where $\vartheta_i^{\tilde{\tau}}$ is player i's private signal in period $\tilde{\tau}$.

Definition 13. A *perfect public equilibrium*, or PPE, is a public strategy profile σ such that $\hat{v}_i(H_i^\tau;\sigma) \geq \hat{v}_i(H_i^\tau;(\sigma_i',\sigma_{-i}))$ for all strategies σ_i', for all private histories H_i^τ, and for all i.

Appendix B Justifying the mechanism design approach

Proof of Theorem 1 (page 10). To prove part (i), suppose that given δ, $G = \big\langle \mathcal{V}, \big\{\langle x,t,w\rangle(\cdot;\nu)\big\}, \nu^0\big\rangle$ is a feasible, IIC, and EPIR recursive mechanism with respect to a punishment payoff profile p. Construct a strategy profile σ as follows.

First, consider any public history H^τ that contains an off-schedule deviation with respect to G. Since p is a punishment payoff profile, there exists a PPE s^* in the stage game that yields an expected average payoff of p; select this equilibrium.

Next, consider any public history H^τ that contains no off-schedule deviations with respect to G. Let $\{\sigma_i(\cdot;H^\tau)\}_{i=1}^N$ specify a pure public strategy profile such that, for all $i\in\mathcal{N}$ and all $j = 2,\ldots,N-1$,

$$\hat{m}_i(\theta_i;H^\tau) = \theta_i, \tag{16}$$

$$\hat{x}_i(\theta_i,\theta;H^\tau) = x_i(\theta;\nu(H^\tau;\sigma)), \tag{17}$$

$$\hat{t}_{1,2}\big(\theta_1,\theta,\hat{x}_i(\theta_i,\theta;H^\tau);H^\tau\big) = -t_1(\theta;\nu(H^\tau;\sigma)), \tag{18}$$

$$\hat{t}_{j,j+1}\big(\theta_j,\theta,\hat{x}_i(\theta_i,\theta;H^\tau);H^\tau\big) = -t_j(\theta;\nu(H^\tau;\sigma)) + \hat{t}_{j-1,j}\big(\theta_{j-1},\theta,\hat{x}_i(\theta_i,\theta;H^\tau);H^\tau\big), \tag{19}$$

with $\hat{t}_{i,j}\big(\theta_i,\theta,\hat{x}_i(\theta_i,\theta;H^\tau);H^\tau\big) = 0$ for any $(i,j)\in\mathcal{N}\times\mathcal{N}$ not specified above.

In addition, for any consistent beliefs following the realization of any message profile $\mu\notin\Theta$, choose $\hat{t}(\theta,\mu,\hat{x}(\theta,\mu;H^\tau);H^\tau) = (0,\ldots,0)$ and choose $\hat{x}(\theta,\mu;H_i^\tau)$ so as to constitute a perfect Bayesian equilibrium in the remainder of the stage game (assume that such an equilibrium exists); and for any consistent beliefs following the realization of any $\chi \neq \hat{x}(\theta, \mu; H^\tau)$, choose $\hat{t}(\theta, \mu, \chi; H^\tau) = (0, \ldots, 0)$ so as to constitute a perfect Bayesian equilibrium in the remainder of the stage game.

This strategy profile yields an equilibrium path equivalent to that of G by construction, and it yields a PPE in every subgame after off-schedule deviations by construction as well. Since G is EPIR with respect to p, no off-schedule deviation is profitable. On-schedule deviations are also unprofitable, because G is IIC. Every subgame after an on-schedule deviation is followed by a subgame that is equivalent to an equilibrium path subgame. Since the proposed strategies do not depend on private histories, they form a PPE in the repeated game.

For part (ii), define a *dynamic mechanism* $\Gamma = \{\langle x, t\rangle(\cdot; H^\tau)\}$ as a collection of pairs $\langle x, t\rangle$ indexed by H^τ. Let $v(H^\tau)$ be the vector of average utilities provided by Γ in the continuation game after history H^τ. Given $\delta \in (0,1)$, I say that Γ is EPIR and IIC with respect to a punishment payoff profile p if the stage mechanism $\left\langle x(\cdot; H^\tau), t(\cdot; H^\tau), \mathbb{E}\left[v(H^{\tau+1})\middle|\cdot; H^\tau\right]\right\rangle$ is EPIR and IIC with respect to p for all equilibrium path public histories H^τ.

Suppose that $\sigma = \{\langle \hat{m}, \hat{x}, \hat{t}\rangle(\cdot; H^\tau)\}$ is a pure strategy PPE, and $\mathcal{V}$ is the set of average utility profiles yielded at the beginning of any period along its equilibrium path. Suppose that $\mathcal{V}$ is SIR with respect to p, and there exists $\hat{\delta} \in (0,1)$ with respect to which p is a punishment payoff profile.

Fix $0 < \varepsilon < \inf_{v \in \mathcal{V}} \min_{i \in \mathcal{N}}\{v_i - p_i\}$; such ε exists because $\mathcal{V}$ is SIR with respect to p. I begin by showing that there exists an IIC dynamic mechanism Γ that yields the same announcements, actions, and net payments as σ on the equilibrium path. Consider any public history H^τ on the equilibrium path. For all $i \in \mathcal{N}$, let

$$x_i(\theta; H^\tau) = \hat{x}_i(\theta_i, \hat{m}(\theta; H^\tau); H^\tau), \tag{20}$$

$$t_i(\theta; H^\tau) = \sum_{j \neq i} \hat{t}_{j,i}\big(\theta_j, \hat{m}(\theta; H^\tau), x(\theta; H^\tau); H^\tau\big) - \sum_{j \neq i} \hat{t}_{i,j}\big(\theta_i, \hat{m}(\theta; H^\tau), x(\theta; H^\tau); H^\tau\big). \tag{21}$$

Although Γ is equivalent to σ in terms of equilibrium path actions and net payments, the players may have additional information at the time they choose actions and payments in any period, since true signals are announced publicly. Nevertheless, since both π and $\hat{t}$ are uniformly bounded by some $D < \infty$,[10] if $\delta > \frac{D}{D+\varepsilon}$ then no off-schedule deviation is profitable, so Γ is EPIR with respect to p. Since the same actions, net payments, and continuation rewards result from truthful announcements as under σ, no on-schedule deviation is profitable, so Γ is IIC. This argument extends in a natural way to mixed strategy PPEs, using randomized dynamic mechanisms.

Next I show that there exists a recursive mechanism $\langle \mathcal{V}, \{\langle x, t, w\rangle(\cdot; v)\}, v^0\rangle$ that is feasible, IIC, and EPIR with respect to p, such that $v^0 = \hat{v}(H^0)$. Note that for each equilibrium path public history H^τ, the portion of Γ that applies following H^τ is itself an IIC and EPIR dynamic mechanism. Thus the promised utility $\hat{v}(H^\tau)$ is supported by a range of continuation rewards generated by dynamic mechanisms; let $w(\cdot; H^\tau)$ be the mapping from announcements to continuation rewards. For each $\tilde{v} \in \mathcal{V}$, select from among the histories $\{H^\tau : \hat{v}(H^\tau) = \tilde{v}\}$ an arbitrary history, and call it $H^*(\tilde{v})$. Then construct a recursive mechanism G by setting $\langle x, t, w\rangle(\cdot; v) = \left\langle x(\cdot; H^*(v)), t(\cdot; H^*(v)), w(\cdot; H^*(v))\right\rangle$ for all $v \in \mathcal{V}$ and setting $v^0 = \hat{v}(H^0)$. G is IIC and EPIR with respect to p because, for each $v \in \mathcal{V}$, it specifies an action function, payment function, and

[10] If $\hat{t}$ were not uniformly bounded, then it would be a profitable deviation for some player to refuse to make his specified payment.

continuation reward function that offer incentives that are identical to the incentives offered by some dynamic mechanism drawn from an equilibrium path subgame of Γ. Finally, G is feasible by construction. ■

References

Dilip Abreu, David Pearce, and Ennio Stacchetti. Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory*, 39(1):251–69, June 1986.

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, September 1990.

Masaki Aoyagi. Bid rotation and collusion in repeated auctions. *Journal of Economic Theory*, 112(1):79–105, September 2003.

Masaki Aoyagi. Efficient collusion in repeated auctions with communication. *Journal of Economic Theory*, 134(1):61–92, May 2007.

Susan Athey and Kyle Bagwell. Optimal collusion with private information. *RAND Journal of Economics*, 32(3):428–465, Autumn 2001.

Susan Athey and Kyle Bagwell. Collusion with persistent cost shocks. *Econometrica*, 76(3):493–540, April 2008.

Susan Athey and David A. Miller. Efficiency in repeated trade with hidden valuations. *Theoretical Economics*, 2(3):299–354, September 2007.

Erik J. Balder. On the existence of optimal contract mechanisms for incomplete information principal-agent models. *Journal of Economic Theory*, 68(1):133–148, January 1996.

Andreas Blume and Paul Heidhues. Private monitoring in auctions. *Journal of Economic Theory*, 131(1): 179–211, November 2006.

Claude d'Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11(1):24–45, February 1979.

Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. Correlation, independence, and Bayesian incentives. *Social Choice and Welfare*, 21(2):281–310, October 2003.

Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. Balanced Bayesian mechanisms. *Journal of Economic Theory*, 115(2):385–396, April 2004.

Drew Fudenberg, David K. Levine, and Eric S. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, September 1994.

Drew Fudenberg, David K. Levine, and Eric S. Maskin. Balanced-budget mechanisms with incomplete information. Working paper, May 1995. URL http://levine.sscnet.ucla.edu/papers/balance7.pdf.

Scott Johnson, John W. Pratt, and Richard J. Zeckhauser. Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica*, 58(4):873–900, July 1990.

Alberto Martin and Wouter Vergote. Efficient bid rotation with a continuum of types. Working paper, October 2004.

David A. Miller. The dynamic cost of ex post incentive compatibility in repeated games of private information. Working paper, July 2007. URL http://dss.ucsd.edu/~d9miller/papers/d9miller_EPPPE.pdf.

Andrzej Skrzypacz and Hugo A. Hopenhayn. Tacit collusion in repeated auctions. *Journal of Economic Theory*, 114(1):153–169, January 2004.