

Games with Imperfectly Observable Actions in Continuous Time.

Yuliy Sannikov

January 18, 2005

Abstract

This paper investigates a new class of 2-player games in continuous time, in which the players' observations of each other's actions are distorted by Brownian motions. These games are analogous to repeated games with imperfect monitoring in which the players take actions frequently. Using a differential equation we find the set $\mathcal{E}(r)$ of payoff pairs achievable by all public perfect equilibria of the continuous-time game, where r is the discount rate. The same differential equation allows us to find public perfect equilibria that achieve any value pair on the boundary of the set $\mathcal{E}(r)$. These public perfect equilibria are based on a pair of continuation values as a state variable, which moves along the boundary of $\mathcal{E}(r)$ during the course of the game. In order to give players incentives to take actions that are not static best responses, the pair of continuation values is stochastically driven by the players' observations of each other's actions along the boundary of the set $\mathcal{E}(r)$.¹

¹I would like to thank especially Bob Wilson, Andy Skrzypacz, Peter DeMarzo, Paul Milgrom, Dilip Abreu, Manuel Amador, Darrell Duffie, Drew Fudenberg, Mike Harrison, Eddie Lazear, George Mailath, Ennio Stacchetti, Ivan Werning, Ruth Williams, David Ahn, Anthony Chung, Willie Fuchs, Patricia Lasso, Deishin Lee, Day Manoli, Gustavo Manso, David Miller, William Minozzi, Dan Quint, Korok Ray, Alexei Tchistyi and all seminar participants at Stanford, Berkeley, Harvard, Princeton, Northwestern, NYU, MIT, the University of Chicago, Yale, the University of Minnesota, UCSD, Humboldt, Oxford, the Minnesota Workshop in Macroeconomic Theory, Rochester, the University of Pennsylvania and the University of Michigan for very valuable feedback on this paper. Also, I would like to thank the editor, Andrew Postlewaite, and two anonymous referees for very thoughtful comments.

1 Introduction.

This paper analyzes a new class of two-player games in continuous time that are related to repeated games with imperfect monitoring (i.e. imperfectly observable actions). In these continuous-time games players do not see each other's actions directly; they only see signals that are distorted by Brownian motions. We are interested in the set of payoff pairs that can be achieved in an equilibrium of the entire game. The benefit of modeling dynamic interactions as continuous-time games lies in the clarity with which the set of equilibrium payoffs can be characterized. The continuous-time approach also allows for a simple description of equilibrium strategies that achieve the extreme points of the set of equilibrium payoffs.

We study *public perfect equilibria* (PPE) and the set of payoff pairs that can be achieved by PPE in a game with imperfectly observable actions. This set of payoff pairs is denoted by $\mathcal{E}(r)$ for a discount rate r . A PPE is a pair of strategies that depend only on the commonly observable public outcomes such that each player's strategy is a best response after all public histories. The purpose of this paper is *not* to prove a Folk Theorem for this class of games, but to precisely characterize the set $\mathcal{E}(r)$ as well as public perfect equilibria.

It turns out that the boundary of the set $\mathcal{E}(r)$ can be found using an ordinary differential equation, which we call the *optimality equation*. The optimality equation also allows us to construct equilibria that achieve any payoff pair on the boundary of the set $\mathcal{E}(r)$. The dynamics of such equilibria are based on a pair of continuation values as a state variable, which moves along the boundary of the set $\mathcal{E}(r)$ during the course of the game. At any moment of time, a player's continuation value is his future expected payoff in the remaining game. The current continuation values determine the players' actions and the impact of observed signals on motion of continuation payoffs.

The optimality equation relates incentives, the equilibrium motion of continuation values, and the geometry of the set $\mathcal{E}(r)$. In equilibrium, a player's incentives stem from the influence of the signal about his actions on his future continuation values. The player's actions are optimal when they maximize his instantaneous payoff flow plus the expected rate of change of his continuation value. Because signals are stochastic, so is the motion of continuation values. The optimality equation, shown informally on the third panel of Figure 1, ties together four measures:

1. inefficiency, how much continuation values v fall behind the flow of payoffs $g(a)$
2. incentives, the sensitivity of continuation values towards public signals

3. the amount of noise in signals
4. the curvature of the set $\mathcal{E}(r)$.

We see that noise, curvature, and the necessity to provide incentives contribute positively to inefficiency. In equilibrium, as continuation values move on the boundary of the set $\mathcal{E}(r)$, the tangent line gives the ratio at which players can instantaneously transfer future equilibrium payoffs in order to create incentives. Because of the curvature of the set $\mathcal{E}(r)$ players cannot transfer utility between each other indefinitely at the same constant rate. Curvature, together with the magnitude of noise in the public signal, quantifies the informational inefficiency.² The greater the curvature, the more costly it is to provide incentives and the greater should be the difference between the continuation values and the flow of payoffs.

The optimality equation also assigns an equilibrium action pair a to each point v on the boundary of the set $\mathcal{E}(r)$. That action pair optimally resolves the trade-off between inefficiency and incentives to stretch the boundaries of the set $\mathcal{E}(r)$ as far out as possible.

This paper contributes to the theory of repeated games with imperfect monitoring, which has been developed by Abreu, Pearce, and Stacchetti (1990), hereafter APS, and Fudenberg, Levine, and Maskin (1994), hereafter FLM. Specifically, continuous-time games illustrate the pattern of equilibrium dynamics in such games and clearly outline the trade-offs involved in the choice of equilibrium actions. The contributions of APS, FLM, and continuous-time games are illustrated in three panels of Figure 1, in which the horizontal and vertical axes represent the players' payoffs.

APS investigate sequential equilibria of repeated games with imperfect monitoring. These games have a great multiplicity of equilibria. APS make the problem of finding equilibrium payoffs much more manageable. They show that any equilibrium payoff vector can be achieved by a recursive equilibrium, in which the players' continuation values are state variables. In equilibrium continuation values change location after every observation of the public signal. The arrows in the left panel of Figure 1 illustrate the potential jumps of continuation values after different signals. The challenge behind our understanding of discrete-time games is that it is difficult to see a pattern behind these jumps and the connection between the equilibrium dynamics and the shape of the set of equilibrium payoffs. Continuous-time games illuminate the connection between the equilibrium motion of

²As noise increases, the variance of continuation values necessary to provide incentives increases. Because of the curvature of $\mathcal{E}(r)$, that increases inefficiency.

continuation values, incentives, and the shape of the set $\mathcal{E}(r)$. In particular, the optimality equation leads naturally to a simpler computational procedure in a continuous-time setting.

FLM show that under appropriate conditions the Folk Theorem holds for repeated games with imperfect information: any smooth convex payoff set \mathcal{W} inside the set \mathcal{V}^* of all feasible and individually rational payoffs can be achieved in equilibrium as long as the players are sufficiently patient. The key insight behind FLM's proof of the Folk Theorem is to consider a specific pattern of the motion of continuation values. Specifically, any payoff pair v on the boundary of \mathcal{W} is achievable if the future continuation values (denoted by w , w' and w'' in the middle panel of Figure 1) are chosen on a tangent line. The continuous-time setting allows us to do more: for *any* discount rate r we can characterize the optimal equilibrium motion of continuation values. It turns out that this motion stays on the boundary of the set of equilibrium payoffs $\mathcal{E}(r)$, i.e. it is locally tangential.




Figure 1: Illustration of the methods of APS, FLM and this paper.

One may be surprised that the informational problem persists in our continuous-time setting. After all, if players can change their actions fast, why can they not instantaneously punish all deviations? A critical feature of our model is that while the players can adjust their actions as quickly as they want, the faster they react, the less information they observe. This feature is in a sharp contrast with the model adopted in FLM, where as the duration of a period is shrunk to 0, the amount of information that the players learn per period nevertheless remains the same. This issue was addressed by Abreu, Milgrom and Pearce (1991) in a model with Poisson arrival of signals.³

³Also, Kandori (1992) shows that the set of payoffs achieved in PPE increases in the accuracy of monitoring.

Brownian motion was first applied to the problem of dynamic incentive provision in Holmstrom and Milgrom (1987). Their paper is a good example that in some situations a continuous-time formulation allows us to better recognize patterns and prove clean results. The information flow in our paper is similar to Holmstrom and Milgrom (1987) in the sense that players learn about each other's actions from a continuous process with i.i.d. increments.⁴

Simon and Stinchcombe (1989) illustrate a lot of difficulties associated with the modeling of games in continuous time. For example, a simple description of a strategy in discrete time often has no equivalent in continuous time. These difficulties arise when the actions of one player instantaneously create information available to his opponent. This issue is not a problem in our framework. In our continuous-time games, information is defined exogenously in terms of all possible signals, and a strategy of a player simply defines a probability measure over all possible signals.

Recently, a number of authors have enriched the problem of optimal incentive provision in a dynamic setting using the mathematical tools of optimal control of diffusion processes. Sannikov (2004) and Williams (2004) both introduce a new method of analyzing the informational problem in a dynamic principal-agent relationship. In both models, the agent drives a stochastic state X with his choice of controls, but the agent's choice is not directly observable.⁵ Both papers analyze models with one-sided imperfect information, where only the agent takes hidden actions. This paper extends the continuous-time method to a two-sided setting, where both players take hidden actions.

This paper is organized as follows. Section 2 provides two motivating examples. Section 3 formally describes the class of continuous-time games analyzed in this paper. Section 4 describes several standard game-theoretic concepts in our setting: the stage game, the minmax payoff and the sets of unconstrained payoffs and Nash equilibrium payoffs. Section 5 identifies incentive compatibility conditions, discusses the concept of a continuation value, and describes PPE in terms of the stochastic motion of continuation values. Section 6 interprets this description of PPE as a stochastic control problem and characterizes the set $\mathcal{E}(r)$ as well as PPE that achieve its extreme points. Section 7 summarizes the main

⁴We do not allow statistically meaningful jumps in the players' observations, as in the Poisson model of Abreu, Milgrom and Pearce (1991). As a result, the noise has the form of a Brownian motion.

⁵Williams (2004) characterizes the optimal contract with a partial differential equation based on the following state variables: time, state X , the agent's value, and possibly other variables in an enriched formulation. In Sannikov (2004) the optimal contract can be derived using an ordinary differential equation based on a single state variable, the agent's continuation value.

results and provides an intuitive discussion of public perfect equilibria. Section 8 presents computational techniques and examples. Section 9 concludes the paper.

2 Examples.

In this section, we illustrate continuous-time games and the results of this paper with two examples. Each game is defined by the matrix of expected payoffs of a stage game and the noise structure. Players continuously take actions at each time $t \in [0, \infty)$, but they do not directly observe each other's past actions. Instead, players publicly observe random processes that carry some information about their past actions. This information is obscured by noise. Specifically, player i 's actions determines the drift of the publicly observed process X^i . Here are the two examples:




Figure 2: Matrix of Static Payoffs and Set $\mathcal{E}(r)$ in Partnership.

Noisy Partnership/Prisoners' Dilemma. Two players participate in a joint venture in continuous time. Each player has a set of actions $\mathcal{A}_1 = \mathcal{A}_2 = \{0, 1\}$, where 0 means “no effort” and 1 means “effort.” The action of player $i = 1, 2$ at time t is denoted by A_t^i . Players do not directly observe each other's actions, but instead observe random processes

$$X_t^1 = \int_0^t A_s^1 ds + Z_t^1, \quad X_t^2 = \int_0^t A_s^2 ds + Z_t^2,$$

where Z^1 and Z^2 are independent standard Brownian motions. The increments of the process X^i reflect how much the actions of player i contribute to the success of the joint venture. Players enjoy the success of their joint venture, but they dislike effort. The actual payoffs of players 1 and 2 are defined by

$$r \int_0^\infty e^{-rt} (2 dX_t^1 + 2 dX_t^2 - 3A_t^1 dt) \quad \text{and} \quad r \int_0^\infty e^{-rt} (2 dX_t^1 + 2 dX_t^2 - 3A_t^2 dt).$$

Note that the instantaneous payoff of player i depends on A_t^i , dX_t^1 and dX_t^2 . The expected payoffs can be written as

$$E \left[r \int_0^\infty e^{-rt} g_1(A_t^1, A_t^2) dt \right] \quad \text{and} \quad E \left[r \int_0^\infty e^{-rt} g_2(A_t^1, A_t^2) dt \right],$$

where

$$g_1(a_1, a_2) = 2a_2 - a_1 \quad \text{and} \quad g_2(a_1, a_2) = 2a_1 - a_2.$$

Static payoff functions g_1 and g_2 give the expectation of the rates at which the players receive their payoffs for any pair of actions. This payoff matrix of the stage game is shown in Figure 2.




Figure 3: Sample Paths of Continuation Values.

To give a taste of our results, Figure 2 also shows a computed set $\mathcal{E}(r)$ for $r = 0.2$.

Let us discuss an equilibrium that achieves payoff pair C on the boundary of the set $\mathcal{E}(r)$. During the equilibrium play, the pair of continuation values follows a diffusion process on the boundary of $\mathcal{E}(r)$, driven by the realizations of X . The pair of continuation values has a drift and a volatility. The tangential component of the drift is shown in Figure 2: it is directed away from points A , C and E , towards points B , D , and the origin. Players choose their effort levels depending on the current pair of continuation values as shown in Figure 2. Both players put effort on the thick portion of the boundary of $\mathcal{E}(r)$.

Figure 3 gives three sample paths of the players' continuation values in the equilibrium that achieves payoff pair C . The vertical axis represents the boundary of $\mathcal{E}(r)$, denoted by $\partial\mathcal{E}(r)$, with points A , B , C , D and E clearly marked.

In Figure 3, the drift of continuation values is directed away from the solid horizontal lines, towards the dotted lines. The dotted lines represent the boundaries, where one of the players switches from effort to no effort. Because of the drift pattern, players typically spend considerable amounts of time in “unequal” regimes, where one player puts effort and the other alternates between effort and no effort. These regimes are denoted by $(1/0,1)$ and $(1,1/0)$ in Figure 3. The realizations of X cause players to switch from one unequal regime to another, until they become absorbed in the static Nash equilibrium, in which players stop putting effort. We see from Figure 3 that the collapse into Nash equilibrium is fast because the drift towards the Nash equilibrium point becomes stronger as the continuation values approach that point.

Duopoly with Differentiated Products. Consider two firms, whose products are imperfect substitutes. The private actions of firm i are supply rates from set $\mathcal{A}^i = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$. The instantaneous prices of firms 1 and 2 are given by the increments of the processes

$$dP_t^1 = (25 - 2A_t^1 - A_t^2) dt + \text{noise} \quad \text{and} \quad dP_t^2 = (30 - 2A_t^2 - 2A_t^1) dt + \text{noise}.$$

Prices are publicly observable, and the noise structure is such that firms can isolate a signal about each firm's quantity from the prices

$$\begin{cases} dX_t^1 = \frac{1}{2} dP_t^2 - dP_t^1 + 10 dt = A_t^1 dt + dZ_t^1 \\ dX_t^2 = dP_t^1 - dP_t^2 + 5 dt = A_t^2 dt + dZ_t^2 \end{cases}$$

where Z^1 and Z^2 are independent standard Brownian motions. The payoffs of firms 1 and

2 are given by

$$r \int_0^\infty e^{-rt} A_t^1 dP_t^1 \quad \text{and} \quad r \int_0^\infty e^{-rt} A_t^2 dP_t^2.$$

The payoff functions can be identified as

$$g_1(a_1, a_2) = a_1(25 - 2a_1 - a_2) \quad \text{and} \quad g_2(a_1, a_2) = a_2(30 - 2a_2 - 2a_1).$$

This stage game has a unique Nash equilibrium (5, 5), but ideally firms could collude by producing (4, 4). The PPE of this game are described in Section 8.2.

3 The Setting.

Time $t \in [0, \infty)$ advances continuously. At every moment of time t , players $i = 1, 2$ choose actions A_t^i from finite sets \mathcal{A}^i . Players do not see each other's actions directly, but publicly observe the path of a d -dimensional public signal X_t that depends on the players' actions and noise. When players $i = 1, 2$ take actions $A^i = \{A_t^i, t \in [0, \infty)\}$, then the signal about their actions is

$$X_t = \int_0^t \mu(A_s) ds + dZ_t,$$

where $A_s = (A_s^1, A_s^2)$ and Z is a d -dimensional Brownian motion. A strategy of player i is *public* if his action at time t is determined by the public history, which consists of the signals $\{X_s, s \in [0, t]\}$ and possibly other public information introduced for public randomization. Formally, public information is captured by a filtration $\{\mathcal{F}_t\}$.

Each player receives a flow of payoff discounted at a common rate r , which depends on his action and the public signal, but not directly on his opponent's action. The instantaneous payoff of player i at time t is given by $c_i(A_t^i) dt + b_i(A_t^i) \cdot dX_t$ for some functions $c_i : \mathcal{A}^i \rightarrow \mathfrak{R}$ and $b_i : \mathcal{A}^i \rightarrow \mathfrak{R}^d$. Since X_t has drift $\mu(A_t)$, player i 's expected payoff flow at time t is

$$g_i(A_t) = c_i(A_t^i) + b_i(A_t^i) \cdot \mu(A_t).$$

Definition. A pair of *public* strategies $A = (A_t^1, A_t^2; 0 \leq t < \infty)$ is a *perfect public equilibrium* (PPE) if for $i = 1, 2$ A_i maximizes the expected discounted payoff of player i given the strategy A_j of his opponent after all public histories.

Formally, the expected discounted payoff (a.k.a. continuation value) of player i after a

public history at time t is

$$W_t^i(A) = E_t \left[r \int_t^\infty e^{-r(s-t)} g_i(A_s) ds \right], \quad (1)$$

where the expectation is conditioned on the public information at time t .

In the next sections we characterize the set $\mathcal{E}(r)$ of payoff pairs achievable by all PPE and pay special attention to the PPE that achieve extreme value pairs of the set $\mathcal{E}(r)$. We find that the equilibrium actions in those PPE are essentially unique and that those PPE do not use public randomization. The equilibrium dynamics are described in terms of the stochastic motion of continuation values on the boundary of $\mathcal{E}(r)$ driven by the public signals.

Games with Product Structure.

Some characteristics of our equilibria, especially the provision of incentives, are easier to understand within the context of a subclass of these games that has product structure. In this class the set of actions of each player is one-dimensional and $\mu(a_1, a_2) = (a_1, a_2)$, i.e. each player's action affects a separate signal, and the two signals are independent. The examples from the previous section are all from this class. After we derive our results in a general setting, we discuss incentives in equilibria of the games from this class in the end of Section 7.

4 Important Sets.

Let us review several concepts that are familiar from the theory of repeated games. A stage game G has the set of players $N = \{1, 2\}$, an action set of each player \mathcal{A}_i and payoff functions g_i .

$$G = \{N, (\mathcal{A}_i)_{i \in N}, (g_i)_{i \in N}\}.$$

Denote the set of all action profiles of the stage game G by $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$, and the set of pure strategy Nash equilibria, by $\mathcal{A}^N \subseteq \mathcal{A}$. Let \mathcal{N} be the convex hull of all pure strategy Nash equilibrium payoff pairs of game G , and \mathcal{V} , the convex hull of all feasible payoff pairs:

$$\mathcal{N} \equiv \text{co} \{(g_1(a), g_2(a)) \mid a \in \mathcal{A}^N\} \quad \mathcal{V} \equiv \text{co} \{(g_1(a), g_2(a)) \mid a \in \mathcal{A}\}.$$

The pure strategy minmax payoff of player i is

$$\underline{v}_i \equiv \min_{a_j} \max_{a_i} g_i(a_i, a_j) \quad (2)$$

Player i can guarantee himself his pure strategy minmax payoff for any strategy of the opponent. Define by

$$\mathcal{V}^* \equiv \{v \in \mathcal{V} \mid v_i \geq \underline{v}_i \text{ for } i = 1, 2\},$$

the subset of \mathcal{V} on which each player receives at least his minmax payoff.

Due to the possibility of public randomization, the set $\mathcal{E}(r)$ of payoff pairs achievable by all PPE is convex. As in repeated games in discrete time, we have

$$\mathcal{N} \subseteq \mathcal{E}(r) \subseteq \mathcal{V}^* \subseteq \mathcal{V},$$

as illustrated in Figure 4.




Figure 4: Sets \mathcal{N} , $\mathcal{E}(r)$, \mathcal{V}^* and \mathcal{V} for the Noisy Partnership with $r = 0.2$.

Remark. The reason why player i can guarantee himself a pure-strategy minmax payoff of \underline{v}_i is that the strategies of both players must be public, i.e. adapted to the public

filtration $\{\mathcal{F}_t\}$. We can imagine an alternative formulation, in which each player's strategy is adapted to his private filtration $\{\mathcal{F}_t^i\}$ and the public filtration satisfies $\mathcal{F}_t \subseteq \mathcal{F}_t^1 \cap \mathcal{F}_t^2$. Under this alternative formulation, there may be equilibria that are not public, and the achievable payoff set may be larger than $\mathcal{E}(r)$. Kandori and Obara (2003) show that the use of private strategies can improve equilibrium payoffs even in games with public monitoring. However, to defend PPE, it can be shown that if player i follows a public strategy then player j has a best response that is also a public strategy.

5 Properties of Continuation Values in PPE.

In this section, we characterize public perfect equilibria in terms of the stochastic properties of the continuation values $W_t^1(A)$ and $W_t^2(A)$. Our analysis proceeds as follows. We start with the definition of a player's continuation value: it is his future expected equilibrium payoff after a given public history. As time passes and the history unfolds, the continuation values change: they will move in the set $\mathcal{E}(r)$. In a public perfect equilibrium this motion is determined by the public information: the signals X_t and public randomization. Proposition 1 represents the relationship between public information and the motion of continuation values formally, and shows that this motion must satisfy a *promise keeping* condition. This condition relates a player's current continuation value, his current payoff flow and the expected change of his continuation value.

Next, we find that incentives come from the relationship between the public signals and the motion of continuation values. A player may have incentives to take an action different from a static best response because (1) actions affect public signals, and (2) public signals affect future continuation values. Proposition 2 provides an instantaneous *incentive compatibility* condition. The analogue of this condition in a discrete-time game is that a deviation for one period is not profitable. Proposition 2 proves that a strategy of player i is optimal in response to the strategy of his opponent at all times if and only if the instantaneous incentive compatibility condition always holds. Together, Propositions 1 and 2 show that in any PPE, the motion of continuation values must satisfy two conditions: *promise keeping* and *incentive compatibility*. Propositions 2 and 3 imply the converse: if the motion of two random processes satisfies these conditions, then there is a PPE in which these random processes are continuation values. Theorem 1 summarizes these results of this section.

Recall that the continuation value of player i is defined as

$$W_t^i(A) = E_t \left[r \int_t^\infty e^{-r(s-t)} g_i(A_s) ds \right]. \quad (3)$$

First, we need a representation that identifies the drift of $W^i(A)$ and the sensitivity of $W^i(A)$ to X^i .

Proposition 1. (Representation and Promise Keeping). *Player i 's continuation value has drift $r(W_t^i(A) - g_i(A_t))$ and can be represented as*

$$W_t^i(A) = W_0^i(A) + r \int_0^t (W_s^i(A) - g_i(A_s)) ds + r \int_0^t \beta_s^i \cdot (dX_s - \mu(A_s)ds) + \check{\epsilon}_t^i, \quad (4)$$

where $\beta^i = (\beta^{i1} \dots \beta^{id})$ are chosen so that $\check{\epsilon}_t^i$ is a martingale orthogonal to X .

This proposition formalizes the fact that the path of player i 's continuation value is determined by public information: the signal X and a public randomization process $\check{\epsilon}$. The vector of coefficients β_t^i captures the exposure of player i 's continuation value to dX_t . The shorthand form of equation (4) is

$$dW_t^i(A) = r(W_t^i(A) - g_i(A_t)) dt + r\beta_t^i \cdot (dX_t - \mu(A_s)dt) + d\check{\epsilon}_t^i.$$

It can be interpreted as an instantaneous regression of $dW_t^i(A)$ onto a constant and dX_t , with a regression error $d\check{\epsilon}_t^i$.

Note that $r(g_i(A_t) - W_t^i(A))$ is the drift of player i 's continuation value. Figure 5 shows this intuitively from the fact that $W_t(A)$ is a weighted average of the current payoff flow $g(A_t)$ and expected continuation values an instant later. Thus, we call the condition that $W_t^i(A)$ has drift $r(g_i(A_t) - W_t^i(A))$ the *promise keeping* condition.




Figure 5: W_t is a weighted average of $g(A_t)$ and $E_t[W_{t+\epsilon}]$.

Proof. Define $\beta_s^{ik} = \frac{d\langle W^i(A), X^k \rangle_t}{dt}$, where $\langle W^i(A), X^k \rangle$ is called the cross-variation between $W^i(A)$ and X^k . Then $\check{\epsilon}^i$, defined by (4), is orthogonal to X . Also note that

$$V_t^i(A) = r \int_0^t e^{-rs} g_i(A_s) ds + e^{-rt} W_t^i(A)$$

is a martingale. Differentiating with respect to t , we get

$$\underbrace{dV_t^i(A)}_{\text{driftless}} = r e^{-rt} g_i(A_t) dt - r e^{-rt} W_t^i(A) dt + e^{-rt} dW_t^i(A),$$

which proves that $r(W_t^i(A) - g_i(A_t))$ is the drift of $W_t^i(A)$, so $\check{\epsilon}^i$ in representation (4) must be a martingale. \square

The process β^i represents the extent, to which player i 's value is driven by the public signal X . Therefore, β^i is responsible for player i 's incentives, as shown below:

Proposition 2. (Incentive Compatibility). *Strategy A^i of player i is optimal in response to strategy of A^j at all times if and only if the incentive compatibility condition*

$$\forall a'_i \in \mathcal{A}^i, \quad g_i(A_t) + \beta_t^i \cdot \mu(A_t) \geq g_i(a'_i, A_t^j) + \beta_t^i \cdot \mu(a'_i, A_t^j) \quad (5)$$

holds for all t . Therefore, a pair of strategies A is a PPE if and only if (5) holds for both players.

Let us interpret the incentive compatibility condition. Suppose that player i is contemplating a deviation to an alternative action a'_i at time t . This will change his instantaneous payoff flow by $g_i(a'_i, A_t^j) - g_i(A_t)$. At the same time, this deviation changes the drift of X_t by $\mu(a'_i, A_t^j) - \mu(A_t)$. Since β_t^i is the sensitivity of player i 's continuation value towards dX_t , this will change player i 's continuation value at rate $\beta_t^i \cdot (\mu(a'_i, A_t^j) - \mu(A_t))$. If the incentive compatibility condition holds, such an instantaneous deviation will cause player i a loss of

$$g_i(a'_i, A_t^j) - g_i(A_t) + \beta^i \cdot (\mu(a'_i, A_t^j) - \mu(A_t)) \leq 0. \quad (6)$$

Therefore, condition (5) states that an instantaneous deviation is not profitable.

In the proof of Proposition 2, we show that instantaneous incentive compatibility implies full incentive compatibility. Instantaneous losses from deviations integrate to a loss globally.

Proof. Let us find the payoff to player i from an alternative strategy \hat{A}^i . If player i follows strategy \hat{A}^i until time t and then switches back to A^i , he gets

$$V_t^i(\hat{A}^i, A^j) = r \int_0^t e^{-rs} g_i(\hat{A}_s^i, A_s^j) ds + e^{-rt} W_t^i(A). \quad (7)$$

We can use (4) to derive that

$$dV_t^i(\hat{A}^i, A^j) = r e^{-rt} \left(g_i(\hat{A}_t^i, A_t^j) dt - W_t^i(A) + (W_t^i(A) - g_i(A_t)) dt + \beta_t^i \cdot (dX_t - \mu(A_t) dt) + d\check{\epsilon}_t^i \right).$$

The expected payoff from following the strategy \hat{A}^i forever can be found as follows

$$\begin{aligned} W_0^i(\hat{A}^i, A^j) &= E[V_\infty(\hat{A}^i, A^j)] = E \left[\underbrace{V_0^i(\hat{A}^i, A^j)}_{W_0^i(A)} + \int_0^\infty dV_t^i(\hat{A}^i, A^j) \right] = \\ &= W_0^i(A) + E \left[r \int_0^\infty e^{-rt} \left(g_i(\hat{A}_t^i, A_t^j) - g_i(A_t) + \beta_t^i \cdot (\mu(\hat{A}_t^i, A_t^j) - \mu(A_t)) \right) dt \right], \end{aligned}$$

where we used the fact that X_t has drift $\mu(\hat{A}_t^i, A_t^j)$ and $\check{\epsilon}_t^i$ is a martingale.

If condition (5) holds for all t , then $W_0^i(\hat{A}^i, A^j) \leq W_0^i(A)$ and player i does not have a profitable deviation at time 0. By a similar argument, player i will not have a profitable deviation after any public history. Conversely if (5) fails, choose a strategy \hat{A}^i such that \hat{A}_t^i maximizes $g_i(a_i, A_t^j) + \beta_t^i \cdot \mu(a_i, A_t^j)$ for all t . Then, $W_0^i(\hat{A}^i, A^j) > W_0^i(A)$ and A^i is not an optimal response to the strategy A^j . \square

Definition. A $2 \times d$ matrix

$$B = \begin{bmatrix} \beta^1 \\ \beta^2 \end{bmatrix} = \begin{bmatrix} \beta^{11} \dots \beta^{1d} \\ \beta^{21} \dots \beta^{2d} \end{bmatrix}$$

enforces action pair $a \in \mathcal{A}$ if for $i = 1, 2$,

$$\forall a'_i \in \mathcal{A}^i, \quad g_i(a) + \beta^i \cdot \mu(a) \geq g_i(a'_i, a_j) + \beta^i \cdot \mu(a'_i, a_j) \quad (8)$$

To characterize the set $\mathcal{E}(r)$ in the next section we need to:

- (1) be able to show that some value pairs cannot be achieved in equilibrium.

(2) be able to construct PPE that achieve value pairs in our conjectured set $\mathcal{E}(r)$.

For (1), we can rely on the restrictions on the motion of continuation values from Propositions 1 and 2. These restrictions are *promise keeping*, which determines the drift of continuation values, and *incentive compatibility*, which determines their volatility. For (2), we prove a result converse to Propositions 1 and 2: if a pair of strategies $A = (A^1, A^2)$ together with a pair of random processes $W = (W^1, W^2)$ satisfy *promise keeping* and *incentive compatibility*, then W are continuation values and A is a PPE. This result allows us to construct PPE by defining the motion of a state variable, a pair promised values W , with an appropriate drift and volatility. Proposition 3 assures us that if the drift of W^i accounts for promise keeping then W^i coincides with player i 's true continuation value.

Proposition 3. (Promise Keeping). *Consider a pair of strategies A . Suppose W^i is a bounded random process that satisfies*

$$W_t^i = W_0^i + r \int_0^t (W_s^i - g_i(A_s)) ds + \tilde{W}_t^i \quad (9)$$

for some martingale \tilde{W}_t^i . Then W_t^i equals to player i 's true continuation value $W_t^i(A)$.

Proof. It follows that $W_t^i = W_t^i(A)$ if we prove that the processes

$$V_t^i = r \int_0^t e^{-rs} g_i(A_s) ds + e^{-rt} W_t^i \quad \text{and} \quad V_t^i(A) = r \int_0^t e^{-rs} g_i(A_s) ds + e^{-rt} W_t^i(A)$$

coincide. First note that both V_t^i and $V_t^i(A)$ are martingales. Indeed, using (9), $dV_t^i = e^{-rt} d\tilde{W}_t^i$, so V^i is a martingale. From (3), $V_t^i(A)$ is also a martingale because

$$V_t^i(A) = E_t \left[\int_0^\infty e^{-rs} g_i(A_s) ds \right].$$

These martingales converge as $t \rightarrow \infty$ because both $e^{-rt} W_t^i$ and $e^{-rt} W_t^i(A)$ converge to 0. We conclude that

$$V_t^i = E_t[V_\infty^i] = E_t[V_\infty^i(A)] = V_t^i(A) \quad \Rightarrow \quad W_t^i = W_t^i(A)$$

for all t . □

If continuation values under the strategy pair A satisfy the incentive compatibility

condition, then Proposition 2 already guarantees that A is a PPE. We summarize our characterization of PPE in the following theorem:

Theorem 1. Characterization of PPE. *In any PPE A , the pair of continuation values W is a process in \mathcal{V}^* that satisfies*

$$W_t = W_0 + r \int_0^t (W_s - g(A_s)) ds + r \int_0^t B_s \underbrace{(dX_s - \mu(A_s) ds)}_{dZ_s} + \check{\epsilon}_t, \quad (10)$$

where

- (i) $\check{\epsilon}$ is a 2-dimensional martingale orthogonal to X and
- (ii) B_t enforces action pair A_t for all t

Conversely, if W is a bounded 2-dimensional process that satisfies equation (10) for A , B and $\check{\epsilon}$ that satisfy properties (i) and (ii), then W is a pair of continuation values in public perfect equilibrium A .

Therefore, $\mathcal{E}(r)$ is the largest subset of \mathcal{V}^* such that a controlled process W , given by (10) can be kept in the set $\mathcal{E}(r)$ by controls A , B and $\check{\epsilon}$ that satisfy conditions (i) and (ii) starting from any point $W_0 \in \mathcal{E}(r)$.

The last part of the Theorem formulates the problem of finding the set $\mathcal{E}(r)$ as a problem from optimal stochastic control. We will use this result to characterize the set $\mathcal{E}(r)$ and PPE that achieve extreme points of the set $\mathcal{E}(r)$ in the next section.

6 PPE with Extreme Values: Derivation.

In this section, we use the characterization of PPE from Theorem 1 to describe the set $\mathcal{E}(r)$ by an ordinary differential equation and find PPE that achieve extreme points of $\mathcal{E}(r)$. First, we informally describe the results that are derived further in the section. Then, to highlight the game theoretic intuition and to keep technicalities in the background, we derive our main result under the assumption that $\mathcal{E}(r)$ is a compact set with a piecewise continuous curvature. In the Appendix we modify our argument to prove the result without the prior assumptions of compactness and differentiability.

6.1 Informal Discussion.

Let us review the properties of PPE from the previous section, and then introduce the main results of this section about the geometry of the set $\mathcal{E}(r)$ and PPE that achieve extreme value pairs of $\mathcal{E}(r)$. According to Theorem 1, $\mathcal{E}(r)$ is the largest subset of \mathcal{V}^* such that the process

$$W_t = W_0 + r \int_0^t (W_s - g(A_s)) ds + r \int_0^t B_s dZ_s + \check{\epsilon}_t \quad (11)$$

can be kept in the set $\mathcal{E}(r)$ by an appropriate choice of controls. We have the freedom to choose actions A , volatilities B that enforce those actions, and public randomization $\check{\epsilon}$. If the initial value pair W_0 is inside the set $\mathcal{E}(r)$, this freedom gives a lot of room for very many equilibria. However, if the initial value pair W_0 is an extreme point of the set $\mathcal{E}(r)$, the choice of controls is severely restricted because continuation values cannot escape from the set $\mathcal{E}(r)$. In fact, we will show that in an equilibrium that achieves an extreme value pair of the set $\mathcal{E}(r)$

- (a) future continuation values W_t must be extreme points of $\mathcal{E}(r)$
- (b) there is no public randomization, i.e. $\check{\epsilon} = 0$
- (c) the span of B_t is in the tangential direction to the set $\mathcal{E}(r)$ at point W_t at all times
- (d) the choice of A_t and B_t is generically unique at all times
- (e) if there are static Nash equilibrium payoff pairs on the boundary of $\mathcal{E}(r)$, players are eventually absorbed in a static Nash equilibrium with probability 1.

We should point out that the spirit of properties (a)-(e) is present in the existing literature on repeated games. However, in discrete time these properties hold only under special continuity assumptions or in approximation. In relation to (a) and (b) in discrete time, one can always choose extreme continuation values if there is public randomization. Without public randomization, APS show that future continuation values can be chosen to be extreme points of the equilibrium value set if the distribution of signals is non-atomic. Moreover, under certain analyticity conditions, future continuation values have to be extreme. The property (c) that B_t must have a tangential span is related to FLM's concept of enforceability of action pairs on tangent hyperplanes that is used to prove the Folk Theorem. Although this has not been demonstrated formally, one has a sense that having continuation values on tangent hyperplanes is required for a Folk Theorem. (For example,

the Folk Theorem fails for Green and Porter type of equilibria that have two regimes and do not involve continuation values on tangent lines.) Point (d) holds only under very strict continuity assumptions (e.g. the analyticity assumptions of APS that guarantee that continuation values *must* be extreme points). For point (e) if there is a unique way to support any extreme value pair, extreme Nash equilibrium payoff pairs must be absorbing states. However, in discrete time it should be possible for continuation values to never reach an absorbing state.⁶

Even though the spirit of properties (a)-(e) is present in discrete-time games, it is difficult to formalize them. However, they come out cleanly in our setting.

Besides proving (a)-(e) we also derive an ordinary differential equation for the boundary of set $\mathcal{E}(r)$, which we call the *optimality equation*. This equation connects the geometry of the boundary with the equilibrium actions and the stochastic motion of continuation values, and can be used for computation. To understand this equation, we must first provide an analogue of FLM's concept of enforceability on tangent hyperplanes in our setting:

Definition. A $2 \times d$ matrix B enforces action pair $a \in \mathcal{A}$ on tangent $\mathbf{T} = (t_1, t_2)$ if B is of the form

$$B = \mathbf{T}\phi = \begin{bmatrix} t_1\phi_1 \dots t_1\phi_d \\ t_2\phi_1 \dots t_2\phi_d \end{bmatrix}.$$

for some $\phi \in \mathfrak{R}^d$. Let $\mathcal{A}(\mathbf{T})$ be the set of action pairs that are enforceable on tangent \mathbf{T} . Of all vectors ϕ such that $\mathbf{T}\phi$ enforces $a \in \mathcal{A}(\mathbf{T})$, let $\phi(a, \mathbf{T})$ be the one of the smallest length.

Here is the optimality equation:

$$\kappa(w) = \max_{a \in \mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N} \frac{2(g(a) - w) \cdot \mathbf{N}(w)}{r|\phi(a, \mathbf{T}(w))|^2}, \quad (12)$$

where $\kappa(w)$, $\mathbf{N}(w)$ and $\mathbf{T}(w)$ are the curvature and unit normal and tangent vectors at point w on the boundary of the set $\mathcal{E}(r)$, and $|\phi|$ denotes the length of vector ϕ .

We derive these results in next subsection.

6.2 Derivation.

In this subsection we assume that the set $\mathcal{E}(r)$ is compact and has piecewise continuous curvature, and derive a characterization of this set and of PPE that achieve its extreme

⁶See Hauser and Hopenhayn (2004) for a continuous-time example with Poisson signal arrival, in which continuation values never reach static Nash equilibria.

value pairs. The compactness and differentiability of $\mathcal{E}(r)$ are justified in the Appendix. Let us briefly outline our derivation. We find that whenever a pair of continuation values W_t is an extreme point of $\mathcal{E}(r)$, the trajectories of continuation values must be tangential to the boundary of $\mathcal{E}(r)$. The normal component of the drift of W_t causes these trajectories to be locally bent (see Proposition 4.) The natural curvature of the trajectories depends on the drift and volatility of W_t , which are determined by the equilibrium actions. We show that at extreme value pairs W_t the equilibrium actions and volatilities that enforce them must satisfy the constraint that the natural curvature of continuation values cannot be smaller than the curvature of $\mathcal{E}(r)$: otherwise continuation values would escape from $\mathcal{E}(r)$. In optimal equilibria this constraint is binding, and thus the equilibrium actions to achieve extreme payoff pairs are generically unique. The optimality equation (12) connects the curvature of $\mathcal{E}(r)$ with the equilibrium choice of actions.

Let us go through the details of our argument. Recall that the motion of continuation values is described by

$$dW_t = r(W_t - g(A_t)) dt + rB_t dZ_t + d\check{\epsilon}_t \quad (13)$$

We can immediately make two observations about this motion for those times t (including $t = 0$) when W_{t-} is an extreme point of the set $\mathcal{E}(r)$.⁷ First, matrix B_t must have span in the tangential direction to the set $\mathcal{E}(r)$ at point W_{t-} . Indeed, a normal component of volatility would instantaneously throw future continuation values outside the set $\mathcal{E}(r)$ with positive probability. Second, there is no public randomization at moment t because an extreme value pair cannot be represented as a convex combination of other points in $\mathcal{E}(r)$, so $W_{t-} = W_t$.⁸ Since matrix B_t has a tangential span, we can represent it as $B_t = \mathbf{T}(W_t)\phi_t$, where $\mathbf{T}(W_t)$ is a unit tangent vector at point $W_t \in \text{ext } \mathcal{E}(r)$ and $\phi_t \in \mathfrak{R}^d$.

It turns out that when the span of $B_t \neq 0$ is focused along one line, the trajectories of continuation values become locally bent with a curvature that depends on the drift of W_t . This property, which is formalized in the next proposition, will help us write an equation that connects the geometry of the set $\mathcal{E}(r)$ with the stochastic motion of continuation values.

Proposition 4. *Suppose that W_t is on the curve \mathcal{C} , and that $B_t = \mathbf{T}\phi_t$ where the unit*

⁷ W_{t-} denotes a pair of continuation values at time t immediately before public randomization.

⁸One can object that we proved only that public randomization in the form of a jump is impossible. This is enough, because continuous public randomization can be done using an extra uninformative dimension of Z .

vector \mathbf{T} is tangent to \mathcal{C} at point W_t , and \mathbf{N} is an outward unit normal vector. Let $(\theta, f(\theta))$ be a parameterization of \mathcal{C} in normal and tangential coordinates and let

$$D_{t+\epsilon} = \mathbf{N} \cdot W_{t+\epsilon} - f(\mathbf{T} \cdot W_{t+\epsilon}).$$

Then D has volatility zero and drift

$$r(W_t - g(A_t)) \cdot \mathbf{N} + \frac{\kappa}{2} r^2 |\phi_t|^2,$$

at time t , where $\kappa = -f''(\mathbf{T} \cdot W_t)$ is the curvature of \mathcal{C} at point W_t .⁹




Figure 6: The definition of $D_{t+\epsilon}$.

To interpret the proposition, note that $D_{t+\epsilon}$ is related to the distance from $W_{t+\epsilon}$ to the curve \mathcal{C} , as illustrated in Figure 6. If $A_t \notin \mathcal{A}^N$ and

$$\kappa = \frac{2(W_t - g(A_t)) \cdot \mathbf{N}}{r|\phi_t|^2}$$

then $dD_t = 0$. In this case, we call κ a *curvature of enforcement* of action pair A_t at W_t on tangent \mathbf{T} .¹⁰

⁹Curvature is the rate at which the tangential angle changes with arc length.

¹⁰An action pair $a \in \mathcal{A}(\mathbf{T}) \setminus \mathcal{A}^N$ usually has many curvatures of enforcement at w on tangent \mathbf{T} . The

Proof. By projecting equation

$$dW_t = r(W_t - g(A_t)) dt + r\mathbf{T}\phi_t dZ_t. \quad (14)$$

onto the tangent axis we get

$$d(\mathbf{T} \cdot W_t) = r(W_t - g(A_t)) \cdot \mathbf{T} dt + r\phi_t dZ_t \quad (15)$$

Using Ito's Lemma,

$$df(\mathbf{T} \cdot W_t) = \underbrace{f'(\mathbf{T} \cdot W_t)}_0 d(\mathbf{T} \cdot W_t) + \underbrace{f''(\mathbf{T} \cdot W_t)}_{-\kappa} \frac{r^2|\phi_t|^2}{2} dt \quad (16)$$

By projecting (14) onto the normal axis we get

$$d(W_t \cdot \mathbf{N}) = r(W_t - g(A_t)) \cdot \mathbf{N} dt. \quad (17)$$

Combining (16) and (17) we get the desired result

$$dD_t = d(W_t \cdot \mathbf{N} - f(\mathbf{T} \cdot W_t)) = \left(r(W_t - g(A_t)) \cdot \mathbf{N} + \kappa \frac{r^2|\phi_t|^2}{2} \right) dt.$$

□

Using Proposition 4 we can characterize the curvature $\kappa(w)$ of the set $\mathcal{E}(r)$ at any point $w \in \text{ext } \mathcal{E}(r) \setminus \mathcal{N}$ by equation (18) below. We do it in two steps.

First, suppose $W_t(A) = w \in \text{ext } \mathcal{E}(r) \setminus \mathcal{N}$. We claim that action pair A_t cannot be a static Nash equilibrium and must have a curvature of enforcement $\kappa \geq \kappa(w)$. Indeed, if $A_t \in \mathcal{A}^N$ then the drift of continuation values at time t is directed outside $\mathcal{E}(r)$ as shown in the left panel of Figure 7. If $A_t \notin \mathcal{A}^N$ and $\kappa < \kappa(w)$ then by Proposition 4, the drift of the distance from W_t to the set $\mathcal{E}(r)$ would be positive, as shown in the right panel of Figure 7. In either case, continuation values instantaneously escape from $\mathcal{E}(r)$, which leads to a contradiction.

Second we claim that a curvature of enforcement cannot be strictly greater than $\kappa(w)$

smallest one is

$$\kappa = \frac{2(w - g(a)) \cdot \mathbf{N}}{r|\phi(a, \mathbf{T})|^2},$$

where \mathbf{N} is a unit vector orthogonal to \mathbf{T} .




Figure 7: Demonstrating that $\kappa \geq \kappa(w)$.

for any $a \in \mathcal{A}(\mathbf{T}) \setminus \mathcal{A}^N$. Otherwise, informally speaking, value pair w could be achieved by using action pair a at time 0 and continuation values on a curve with curvature $\kappa > \kappa(w)$, which lie inside the set $\mathcal{E}(r)$, as shown in Figure 8. Inspecting the figure, we realize that we must be able to achieve a value pair $w + \epsilon \mathbf{N}$ outside the set $\mathcal{E}(r)$ by using the same action pair and continuation values in the set $\mathcal{E}(r)$, which leads to a contradiction. Formally, this is shown in Proposition 6.

From these two steps, we conclude that the curvature of $\mathcal{E}(r)$ satisfies the following equation at all points $w \in \text{ext } \mathcal{E}(r) \setminus \mathcal{N}$:

$$\kappa(w) = \max_{a \in \mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N} \frac{2(g(a) - w) \cdot \mathbf{N}(w)}{r|\phi(a, \mathbf{T}(w))|^2}, \quad (18)$$

Different action pairs $a \in \mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N$ can have different enforcement curvatures at w on tangent $\mathbf{T}(w)$.¹¹ The action pair (generically unique) with the largest enforcement curvature is played at point $w \in \text{ext } \mathcal{E}(r) \setminus \mathcal{N}$.

Equation (18) does not completely characterize the set $\mathcal{E}(r)$ yet, because there may be many sets that satisfy (18) at all extreme points that are not in \mathcal{N} . We need to add some finishing touches to our derivation, i.e.

- (a) prove that all points $w \notin \mathcal{N}$ on the boundary of $\mathcal{E}(r)$ are extreme
- (b) show that $\mathcal{E}(r)$ is the largest set whose boundary satisfies (18) outside \mathcal{N} and
- (c) formalize the argument behind Figure 8.

¹¹Some action pairs may not be enforceable on tangent $\mathbf{T}(w)$.




Figure 8: Demonstrating that the curvature of enforcement cannot be greater than $\kappa(w)$.

Proposition 5 plays several roles to provide finishing touches: it implies (b) directly, and it is a building block to show (a) and (c).

Proposition 5. *Suppose that the curve \mathcal{C} satisfies equation (18). Furthermore, suppose that either \mathcal{C} is a closed curve, or has endpoints achievable by some PPE. Then $\mathcal{C} \subset \mathcal{E}(r)$.*

Proof. By Theorem 1, to achieve $W_0 \in \mathcal{C}$ in a PPE it is sufficient to construct a bounded process W_t that satisfies

$$W_t = W_0 + r \int_0^t (W_s - g(A_s)) ds + r \int_0^t B_s dZ_s, \quad B_t \text{ enforces } A_t \quad (19)$$

for all t . Denote by $a : \mathcal{C} \rightarrow \mathcal{A} \setminus \mathcal{A}^N$ the maximizer in (18). Let $W_t, t \leq \tau$ be a process that

- starts at W_0
- stays on \mathcal{C} until a stopping time τ when W_t hits an endpoint of \mathcal{C} (where $\tau = \infty$ if \mathcal{C} is a closed curve)
- has tangential drift $r(W_t - g(A_t)) \cdot \mathbf{T}(W_t)$ and volatility $r\mathbf{T}(W_t)\phi(A_t, \mathbf{T}(W_t))$, where $A_t = a(W_t)$.

Then by an application of Ito's lemma the normal component of the drift of W_t is

$$-\kappa(W_t) \frac{r^2 |\phi(A_t, \mathbf{T}(W_t))|^2}{2} = r(W_t - g(A_t)) \cdot \mathbf{N}(W_t)$$

until time τ . Let us extend process W beyond time τ by letting it follow the path of a PPE that achieves value W_τ . Then W becomes a bounded random process that satisfies equation (19) until time ∞ . By Theorem 1, we found a PPE that achieves W_0 . \square

Proposition 5 implies that if the boundary of any set \mathcal{W} satisfies equation (18) at all points outside \mathcal{N} , then $\mathcal{W} \subseteq \mathcal{E}(r)$. We can also use Proposition 5 to formalize the argument behind Figure 8.

Proposition 6. *Let $w \notin \mathcal{N}$ be a point on the boundary of $\mathcal{E}(r)$. Then there is no action pair $a \notin \mathcal{A}^N$, for which*

$$\kappa(w) < \frac{2(g(a) - w) \cdot \mathbf{N}(w)}{r|\phi(a, \mathbf{T}(w))|^2}. \quad (20)$$

Proof. Suppose not. Then the curve \mathcal{C} that solves equation (18) from initial conditions $(w, \mathbf{T}(w))$ has greater curvature at w than the boundary of $\mathcal{E}(r)$, as shown in Figure 8. This means that \mathcal{C} enters the interior of $\mathcal{E}(r)$ around point w . From the continuity of the solutions to (18) in initial conditions (see Appendix), the solution \mathcal{C}' from initial conditions $(w + \epsilon \mathbf{N}(w), \mathbf{T}(w))$ must enter the interior of $\mathcal{E}(r)$ as well. But then, by Proposition 5, $\mathcal{C}' \subseteq \mathcal{E}(r)$, a contradiction. \square

Corollary 1. *Any point $w \notin \mathcal{N}$ on the boundary of $\mathcal{E}(r)$ is extreme.*

Proof. If not, then $\kappa(w) = 0$. There is no action pair $a \notin \mathcal{A}^N$ with $(g(a) - w) \cdot \mathbf{N}(w) > 0$ that is enforceable on tangent $\mathbf{T}(w)$, because otherwise (20) holds. Point w must be contained in a segment on the boundary of $\mathcal{E}(r)$ with extreme endpoints w_L and w_H . Without loss of generality assume that $w_H \notin \mathcal{N}$. If there was an equilibrium that achieved value w_H , then continuation values would escape from the set $\mathcal{E}(r)$ instantaneously due to drift if $A_0 \in \mathcal{A}^N$ or tangential volatility if $A_0 \notin \mathcal{A}^N$, a contradiction. \square

This completes the derivation of our main result which is summarized at the next section. The next section also provides an intuitive discussion of the set $\mathcal{E}(r)$ and PPE that achieve extreme payoff pairs.

7 The Main Section: Summary and Discussion.

The following theorem characterizes the set $\mathcal{E}(r)$ and the public perfect equilibria (PPE) that achieve extreme value pairs of $\mathcal{E}(r)$.

Theorem 2. Characterization. $\mathcal{E}(r)$ is the largest closed subset of \mathcal{V}^* with curvature

$$\kappa(w) = \max_{a \in \mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N} \frac{2(g(a) - w) \cdot \mathbf{N}(w)}{r |\phi(a, \mathbf{T}(w))|^2}, \quad (21)$$

at all points $w \notin \mathcal{N}$ on the boundary of $\mathcal{E}(r)$, where $\mathbf{T}(w)$ and $\mathbf{N}(w)$ are unit tangent and outward normal vectors at w . We call (21) the optimality equation.¹²

PPE with extreme values. Denote by $a : \partial\mathcal{E} \setminus \mathcal{N} \rightarrow \mathcal{A} \setminus \mathcal{A}^N$ the maximizing action pairs in equation (21), where $\partial\mathcal{E}(r)$ denotes the boundary of $\mathcal{E}(r)$. Any value pair $W_0 \in \text{ext } \mathcal{E}(r)$ is achieved by a PPE with the following characteristics. The pair of continuation values under this PPE satisfies the SDE

$$W_t = W_0 + \int_0^t \underbrace{r(W_s - g(A_s))}_{\text{drift}} ds + \int_0^t \underbrace{r\mathbf{T}(W_s)\phi(A_s, \mathbf{T}(W_s))}_{B_s, \text{volatility}} (dX_s - A_s ds) \quad (22)$$

until time τ when W_t hits the set \mathcal{N} . For $t < \tau$, the players take action pairs $A_t = a(W_t)$. After time τ , the players follow a static Nash equilibrium with value W_τ . When $\partial\mathcal{E}(r) \cap \mathcal{N} = \emptyset$, then $\tau = \infty$. Otherwise, players become absorbed in a static Nash equilibrium with probability 1 in finite time.¹³

In the remainder of this section we discuss the implications of this result on various questions of interest: the equilibrium dynamics, the nature of inefficiency, the choice of equilibrium actions and the provision of incentives.

¹²In our model, we normalized each component of the signal X to be independent of the others and have volatility 1. Alternatively, if the players observed signals

$$dX_t = \mu(A_t^1, A_t^2) dt + \Sigma dZ_t,$$

where the volatility matrix Σ has full rank, then after appropriate rescaling the optimality equation would be

$$\kappa(w) = \max_{a \in \mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N} \frac{2(g(a) - w) \cdot \mathbf{N}(w)}{r |\phi(a, \mathbf{T}(w))\Sigma|^2},$$

where $\phi(a, \mathbf{T})$ is defined the same way as before.

¹³There is a great multiplicity of equilibria that achieve non-extreme values. In those equilibria players do not need to become absorbed in a static Nash equilibrium.

Let us describe dynamics in a PPE that achieves an extreme value pair W_0 . As soon as the game begins, the players' continuation values W_t start moving along the boundary of the set $\mathcal{E}(r)$.¹⁴ This motion is a diffusion process defined by equation (22). Point W_t plays the role of a single state variable in this equilibrium. As a state variable, W_t determines the actions which the players take in a given instant, and the law by which W_t itself evolves based on the observations of signal X . If there are Nash equilibrium payoff pairs on the boundary of $\mathcal{E}(r)$, then a pair of continuation values must eventually hit one of them with probability 1. When that happens, the players become absorbed in a static Nash equilibrium forever. Of course, if all static Nash equilibrium payoff pairs are inside the set $\mathcal{E}(r)$, then players never become absorbed in a Nash equilibrium, and the motion of continuation values never stops.

At times $t < \tau$ before the players become absorbed in a static Nash equilibrium (if ever), they choose action pairs A_t and receive the payoff flow $g(A_t) \notin \mathcal{E}(r)$. The pair of continuation values W_t has drift directed away from point $g(A_t)$ inside the set $\mathcal{E}(r)$. This drift accounts for promise keeping: the current continuation value W_t is always a weighted average of the current payoff flow $g(A_t)$ and the expected continuation value a moment later $E_t[W_{t+\epsilon}]$, as shown in Figure 9.




Figure 9: The drift and volatility of continuation values.

It may seem surprising that the drift of continuation values is directed inside $\mathcal{E}(r)$ even though continuation values stay on the boundary. We can reconcile these two facts as follows: because continuation values diffuse along the boundary due to tangential volatility and because the boundary has curvature, the expectation of future continuation values must be inside the set $\mathcal{E}(r)$.

¹⁴Typically, as in all our examples, the pair of continuation values will diffuse along the entire boundary of $\mathcal{E}(r)$, not just its Pareto efficient portion.

The equilibrium actions pairs A_t come from the optimality equation (21). The objective of this equation is to describe the largest set of payoff pairs achievable in equilibrium. The choice among action pairs involves a trade-off between the extremity of payoffs and the incentives required to enforce them. The extremity of a payoff pair is measured by the payoff gain in the direction of the normal vector (see the numerator of (21)). The incentives are measured by the instantaneous tangential variance of continuation values (see the denominator of (21)). An optimal action pair achieves the maximum in (21). This action pair can be enforced by using continuation values on the boundary of $\mathcal{E}(r)$. If we tried to enforce a suboptimal action pair, the required drift and tangential volatility of W_t would take future continuation values outside $\mathcal{E}(r)$.

Let us discuss the provision of incentives. Before time τ , actions A_t are not static Nash equilibria, so players must have incentives to take actions that are not static best responses. These incentives arise because actions affect the drift of the public signals, which in turn affect continuation values. The volatility matrix in equation (22) is the sensitivity of continuation values to the signal X . From Section 5 we know that player i has incentives not to deviate from action pair A_t if his action maximizes the sum of his instantaneous payoff and the expected change of his continuation value, i.e.

$$g_i(A_t) + \beta_t^i \cdot \mu(A_t) = \max_{a_i} g_i(a_i, A_t^j) + \beta_t^i \cdot \mu(a_i, A_t^j), \quad (23)$$

where β_t^i is row i of the volatility matrix at time t . In an equilibrium that achieves an extreme payoff pair, the volatility matrix must be of the form $\mathbf{T}(w)\phi_t$ to have a tangential span. Generally, there could be many ways to enforce A_t on a tangent line, but only the smallest tangential variance must be used in an equilibrium, for which W_0 is extreme. In our general setting, not all action pairs can be enforced on all tangent lines so the Folk Theorem may fail.

7.1 Incentives in Games with Product Structure.

The provision of incentives is especially clear in a special subclass of games with product structure that was outlined in the end of Section 3. For that class of games $\mathcal{A}_1, \mathcal{A}_2 \subset \mathfrak{R}$, the public signal is two-dimensional and has drift $\mu(a_1, a_2) = (a_1, a_2)$. Therefore, there is a separate signal that is indicative of each player's actions. The examples in Section 2 come from this class of games.

For this class of games the volatility matrix is 2×2 and condition (8) reduces to

$$g_i(A_t) + \beta_t^{ii} A_t^i = \max_{a_i} g_i(a_i, A_t^j) + \beta_t^{ii} a_i, \quad (24)$$

where β_t^{ii} for $i = 1, 2$ are the diagonal entries of the volatility matrix $B_t = \begin{bmatrix} \beta_t^{11} & \beta_t^{22} \frac{t_1}{t_2} \\ \beta_t^{11} \frac{t_2}{t_1} & \beta_t^{22} \end{bmatrix}$. Here (t_1, t_2) denotes a unit tangent vector and the off-diagonal entries are determined uniquely from the requirement that B_t has a tangential span. To enforce action pair A_t with minimal volatility, we must choose β_t^{ii} of the smallest absolute value for (24) to hold.

Definition of γ . Consider all values of β^{ii} for which a_i maximizes $g_i(a_i, a_j) + \beta^{ii} a_i$ given a_j . Of these values, define $\gamma_i(a_i, a_j)$ to be the smallest in terms of absolute value. Note that $\gamma_i(a)$ is defined for all action pairs a if and only if g_i is concave in a_i .




Figure 10: How to find γ .

Rates γ_i can be computed very easily from g_i , as illustrated in Figure 10. If we fix a_j and plot $g_i(\cdot, a_j)$, then $-\gamma_i(a_i, a_j)$ equals the slope of function $g_i(\cdot, a_j)$, between point a_i and the nearest profitable deviation.

In terms of γ 's, the volatility matrix must be of the form

$$B_t = \begin{bmatrix} \gamma_1(A_t) & \gamma_2(A_t) \frac{t_1}{t_2} \\ \gamma_1(A_t) \frac{t_2}{t_1} & \gamma_2(A_t) \end{bmatrix}.$$

We can make several useful observations about incentive provision for the games of this paper in general, and for this specific class with product structure.

- Generally, because the local motion of continuation values is restricted to a tangent line, the necessity to provide incentives to one player affects the continuation value of another player. In this smaller class of games, player j 's continuation value has sensitivity $\gamma_i(A_t)t_j/t_i$ towards the signal X_i , which reflects player i 's action exclusively.
- If a game has product structure, then incentives provided to different players do not interfere; there is a separate signal and a separate column of the matrix B_t that is responsible for the actions of each player. In general this is not true because the same signal can be affected by both players.

Let us comment on the enforceability of action pairs on tangent lines and the Folk Theorem in this class of games. If g_i is concave in a_i , then $\gamma_i(a)$ is well defined for all $a \in \mathcal{A}$ and $i = 1, 2$. Then all action pairs can be enforced on all tangent lines, except for horizontal or vertical lines (because the off-diagonal entries of B_t may blow up.) On horizontal lines we can enforce any pair of actions in which player 2 chooses a static best response, including one that maximizes player 2's payoff and one that minmaxes player 2. A similar statement holds for vertical lines and player 1. For that reason, the Folk Theorem holds.¹⁵ If the function g_i is concave in a_i for $i = 1, 2$, the Folk Theorem may fail.

8 Computation.

In this section we discuss the computation of the set $\mathcal{E}(r)$, and present the outcomes of computation in our two examples: noisy partnership and a duopoly with differentiated products.

The optimality equation can be written in a form suitable for computation, for example, as follows

$$\begin{cases} dw(\theta)/d\theta = \mathbf{T}(\theta)/\kappa(\theta), \\ \kappa(\theta) = \max_{\mathcal{A}(\mathbf{T}(w)) \setminus \mathcal{A}^N} \frac{2(g(a) - w(\theta)) \cdot \mathbf{N}(\theta)}{r|\phi(a, \mathbf{T}(\theta))|^2}, \end{cases} \quad (25)$$

where $\mathbf{T}(\theta) = (-\sin \theta, \cos \theta)$, $\mathbf{N}(\theta) = (\cos \theta, \sin \theta)$, and angle θ parameterizes the curve. Generally, $|\phi(a, \mathbf{T})|^2$ can be found by solving a quadratic program

$$|\phi(a, (t_1, t_2))|^2 = \min_{\phi} |\phi|^2$$

¹⁵From the optimality equation we can see immediately that the Folk Theorem holds when $\gamma_i(a)$ is defined for all i and a . Indeed, as r decreases to 0, the numerator $2(g(a) - w) \cdot \mathbf{N}$ in the optimality equation also decreases to 0, making the set $\mathcal{E}(r)$ expand towards the boundaries of \mathcal{V}^* .

$$\text{s.t.} \quad \forall a \in \mathcal{A} \forall i = 1, 2 \forall a'_i \in \mathcal{A}^i, \quad g_i(a) + t_i \phi \cdot \mu(a) \geq g_i(a'_i, a_j) + t_i \phi \cdot \mu(a'_i, a_j).$$

For the subclass of games with product structure,

$$|\phi(a, (t_1, t_2))|^2 = \gamma_1(a)^2/t_1^2 + \gamma_2(a)^2/t_2^2.$$

We can solve equation (25) numerically starting from any initial conditions $(w, \theta) \in \mathcal{V}^* \times [0, 2\pi)$. We present computation in our examples in an increasing order of difficulty.

8.1 Noisy Partnership.

From symmetry considerations, the boundary of the set $\mathcal{E}(r)$ must contain a point on the 45-degree line with an outward unit normal $\mathbf{N} = (\cos(45^\circ), \sin(45^\circ))$. Also, point $(0, 0)$ will be on the boundary as well. For all points w on the line segment between the origin and point $(1, 1)$, consider the curve $\mathcal{C}(w)$ that solves the optimality equation from initial conditions (w, \mathbf{N}) . To compute the set $\mathcal{E}(r)$, we search along the 45-degree line and find point w , removed furthest from the origin, such that the curve $\mathcal{C}(w)$ reaches the origin. First, we do a grid search to identify an interval where the desired point w is located. After that, we do a binary search within the interval to compute w exactly. Figure 11 illustrates the computational procedure for $r = 0.2$.




Figure 11: Computation of $\mathcal{E}(r)$ in Noisy Partnership.

From the grid search on Figure 11a, we know that there are two symmetric closed curves

which satisfy the optimality equation everywhere except in the origin: one in the interval $(0.2, 0.3)$, and one in the interval $(0.8, 0.9)$. We are interested in the latter curve, because it is larger. That curve can be found by means of a binary search in the interval $(0.8, 0.9)$.

The computed boundary of $\mathcal{E}(r)$, along with recommended action pairs at every point, is shown in Figure 11b.

8.2 Duopoly with Differentiated Products.




Figure 12: Computation: an asymmetric game.

Computation for duopoly with differentiated products is more difficult because it is an asymmetric game and because the static Nash equilibrium payoff is in the interior of \mathcal{V}^* . For this game the computational procedure is illustrated in Figure 12. We start at an arbitrary point w^1 on the boundary of the set \mathcal{V}^* , and compute the solutions of the optimality equation from initial conditions (w^1, θ) for $\theta \geq \pi$. We raise θ continuously, until the corresponding solution \mathcal{C}^1 (for some angle $\hat{\theta}^1$) hits point w^1 after making a loop, as shown in Figure 12a. We claim that the resulting solution must enclose the set $\mathcal{E}(r)$. If not, as we vary θ continuously between π to $\hat{\theta}^1$, some solution would have to be tangent to $\mathcal{E}(r)$. However, this is impossible, because then the solution would have to coincide with $\mathcal{E}(r)$ (from the uniqueness of solutions given the initial conditions at the point of tangency).

Next, take point w^2 on the curve \mathcal{C}^1 with an outward unit normal $(1, 0)$. Again, we compute the solutions of the optimality equation from initial conditions (w^2, θ) for $\theta \geq 0$. We raise θ continuously, until the corresponding solution \mathcal{C}^2 (for some angle $\hat{\theta}^2$) hits point w^2 after making a loop, as shown in Figure 12b. Then the curve \mathcal{C}^2 must enclose the set

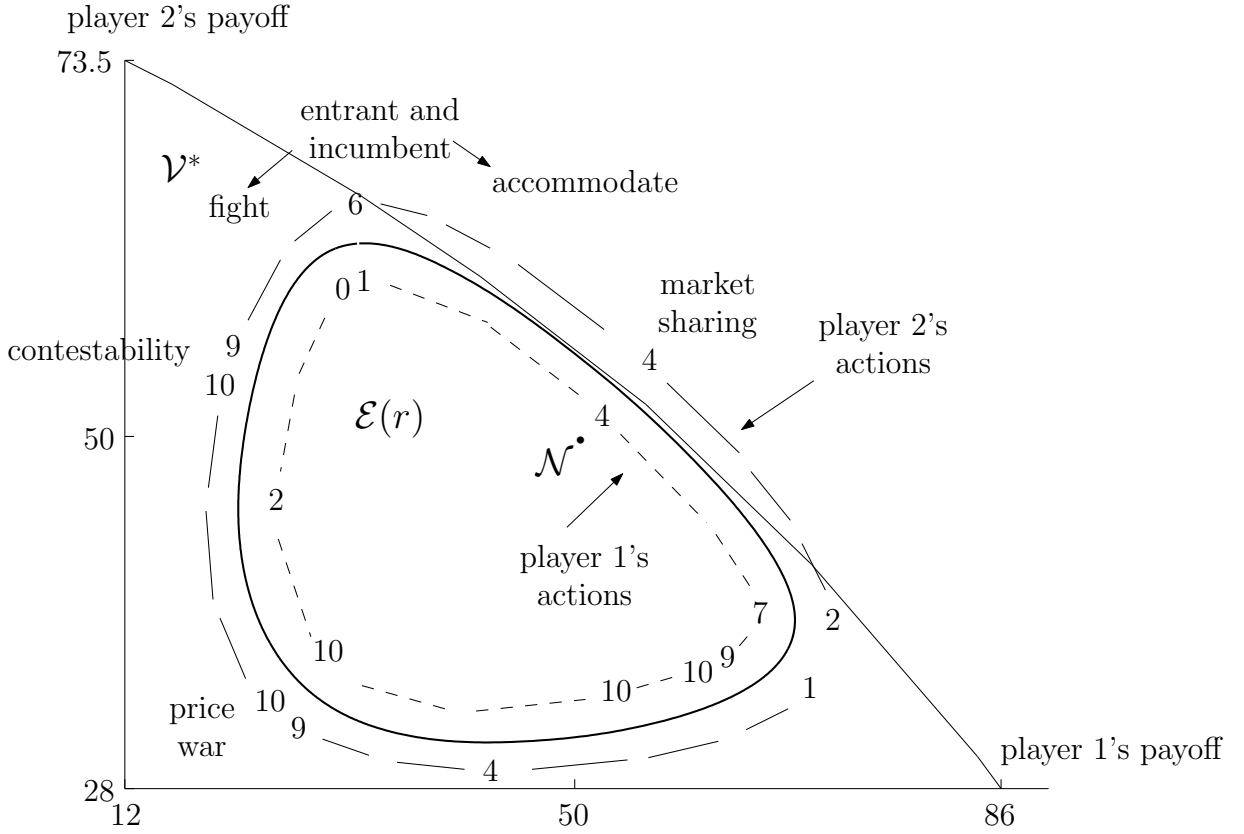


Figure 13: Set $\mathcal{E}(r)$ in Duopoly.

$\mathcal{E}(r)$ inside. By continuing this procedure iteratively, we will converge to the set $\mathcal{E}(r)$.

Figure 13 illustrates the outcome of computation for discount rate $r = 1.5$. The boundary of $\mathcal{E}(r)$ is divided into many segments on which players keep their actions constant. Figure 13 illustrates the general pattern of actions, as well as an interpretation of each portion on the boundary of the set $\mathcal{E}(r)$.¹⁶ For comparison, recall that a static Nash equilibrium is $(5, 5)$. Along the Pareto frontier of $\mathcal{E}(r)$, players collude by producing less than their static best responses. We call this regime “market sharing.” In this regime, when a player’s continuation value increases, his market share also increases. Therefore, players are rewarded for underproducing by an increased future market share. On top of the set $\mathcal{E}(r)$, player 2 receives the maximal payoff that he possibly could in a PPE. At that point, player 1 produces very little, while player 2 produces close to his monopoly quantity. While player 2 chooses a static best response, player 1 needs strong incentives to “stay

¹⁶I am thankful to William Fuchs for helping me find these interpretations.

out.” To reward player 1 for “staying out,” player 2 accommodates, and to punish player 1 for cheating, player 2 fights. We call this regime “entrant and incumbent.” On the left side of $\mathcal{E}(r)$ player 1 is acting passively by producing a static best response, while player 2 is overproducing aggressively. At this point, player 2 is rewarded for overproducing by being able to drive player 1 out of the market. We call this regime “contestability.” At the bottom left portion of $\mathcal{E}(r)$, players are fighting a “price war” by overproducing. They have incentives to do so because the player that looks more aggressive will come out as a winner of the price war. The winner gets his reward by becoming a monopolist for some period of time.

9 Conclusion.

This paper introduces a new class of games in continuous time, in which the players’ observations of each other’s actions are distorted by Brownian motion. In these games, the set of value pairs which are achievable in public perfect equilibria has a clean characterization. The form of public perfect equilibria that achieve values on the boundary of the set $\mathcal{E}(r)$ and the way by which the players organize the provision of incentives are intuitive. We saw examples of various economic interactions that can be modeled as continuous-time games. Besides our examples of a partnership and a duopoly, our model can be applied to principal-agent problems, risk-sharing models, etc. One is hopeful that the simplicity of characterizations in continuous-time models will allow deeper analysis of applications to various dynamic incentive problems with imperfect information.

Let us discuss several questions for development of future theory. First, it is necessary to illustrate the connection between discrete-time repeated games and continuous-time games and to understand how continuous-time games can be used to approximate repeated games in discrete time. Second, it is beneficial to extend the continuous-time approach to games with private information. DeMarzo and Sannikov (2004) show how to attack the issue of private information in a setting with one-sided imperfect information. Third, one has to extend the continuous-time approach to settings where more than one state variable is required. Finally, it would be interesting to explore other computational procedures to find the set $\mathcal{E}(r)$.

Appendix: Technical Details for Section 6.

Section 6 assumes that the set $\mathcal{E}(r)$ is compact and has a piecewise continuous curvature in order to simplify the argument leading to our characterization of $\mathcal{E}(r)$. Here we explain how to alter the argument to avoid these prior assumptions. We show that the set $\mathcal{E}(r)$ is closed and that its boundary satisfies the optimality equation at all points, except for those in \mathcal{N} . To prove this, consider an arbitrary point $w \in \partial\mathcal{E}(r) \setminus \mathcal{N}$. We will show that a tangent solution to the optimality equation through point w coincides with the boundary. Therefore the boundary of $\mathcal{E}(r)$ has piecewise continuous curvature given by the optimality equation. By Proposition 5 the set $\mathcal{E}(r)$ contains its boundary, so $\mathcal{E}(r)$ must be closed. The following Theorem summarizes the results required to carry the argument of Section 6 without extra assumptions.

Theorem 3. Tangent curves. *There is a unique tangent vector $\mathbf{T}(w)$ at any point $w \in \partial\mathcal{E}(r) \setminus \mathcal{N}$. Also, the curve \mathcal{C} that solves equation (21) from initial conditions $(w, \mathbf{T}(w))$ coincides with the boundary of $\mathcal{E}(r)$ in a neighborhood of w .*

Proof. The proof goes in two steps. First, we show that the curve \mathcal{C} cannot go outside the boundary of the set $\mathcal{E}(r)$ in a neighborhood of w . Otherwise, by altering initial conditions slightly, we would be able to find a curve \mathcal{C}' that solves equation (21) and cuts through the boundary of the set $\mathcal{E}(r)$ as shown in Figure 14. Lemma 1 shows that this leads to a contradiction. From the first step we also conclude that the tangent vector is unique at any point $w \in \partial\mathcal{E}(r) \setminus \mathcal{N}$. If the tangent vector was not unique then a tangent solution would go outside the set. Second, we show that the curve \mathcal{C} does not enter the interior $\mathcal{E}(r)$. Otherwise, we would be able to construct PPE that achieves a value pair outside $\mathcal{E}(r)$, as shown in Lemma 2. Throughout the analysis, we use the fact that we can adjust solutions to the optimality equations continuously, which follows from Lemma 3 in the end of this proof.




Figure 14: Constructing a curve that cuts through $\mathcal{E}(r)$.

Suppose that the curve \mathcal{C} has parts outside the set $\mathcal{E}(r)$. Then, by adjusting initial conditions slightly, we can draw a curve \mathcal{C}' that also solves the optimality equation and cuts through a small portion of the boundary of $\mathcal{E}(r)$, as shown in Figure 14. The left panel shows that when the set $\mathcal{E}(r)$ has a kink at w , we can find \mathcal{C}' by moving initial conditions

inside the set. The right panel shows that when the set $\mathcal{E}(r)$ has a unique tangent at w , we can draw \mathcal{C}' from the same point w but with a rotated angle.

Lemma 1 shows that it is impossible to have a curve that cuts through a small portion¹⁷ of the boundary of $\mathcal{E}(r)$ near a point $w \in \partial\mathcal{E} \setminus \mathcal{N}$.

Lemma 1. *A solution \mathcal{C}' of the optimality equation with endpoints $v_L, v_H \in \partial\mathcal{E}(r)$ cannot pass through the interior of $\mathcal{E}(r)$, as shown on Figure 15 if*

- (i) *there is a unit vector \mathbf{N}' such that for any $x > 0$, $v^L + x\mathbf{N}' \notin \mathcal{E}(r)$ and $v^H + x\mathbf{N}' \notin \mathcal{E}(r)$.*
- (ii) *for all $w \in \mathcal{C}$ with an outward unit normal \mathbf{N} , we have $\max_{v \in \mathcal{N}} v \cdot \mathbf{N} < w \cdot \mathbf{N}$*




Figure 15: Impossible Curve.

Proof. Suppose such curve \mathcal{C}' existed. Then there must be a PPE that achieves any point W_0 strictly between the curve \mathcal{C}' and the boundary of the set $\mathcal{E}(r)$. Denote by W_t the continuation values in this PPE. We will show that such PPE is impossible.

Let \mathbf{T}' be a unit vector orthogonal to \mathbf{N}' . Let us introduce a rectangular coordinate system based on these vectors. Let $(\theta, f(\theta))$ be a parameterization of \mathcal{C}' between points v_L and v_H in these coordinates. Let

$$D_t = \mathbf{N}' \cdot W_t - f(\mathbf{T}' \cdot W_t)$$

until a stopping time τ when W_t hits the curve \mathcal{C}' . Observe that $D_0 > 0$.

¹⁷If $w \in \partial\mathcal{E} \setminus \mathcal{N}$ and the portion of the boundary where \mathcal{C} cuts through is sufficiently small, then conditions (i) and (ii) of the Lemma hold.

We will show that until time τ , D_t has either positive volatility or positive drift. Also, at moments of discontinuous public randomization, D_t increases in expectation. These statements (together with appropriate technical conditions) imply that D_t can become arbitrarily large, so continuation values would escape from the set $\mathcal{E}(r)$. This leads to a contradiction.

To finish the proof of the Lemma, suppose that the volatility of D_t is zero and let us show that it must have positive drift. Then the volatility of W_t is parallel to the curve \mathcal{C}' at point v given by coordinates $(\mathbf{T}' \cdot W_t, f(\mathbf{T}' \cdot W_t))$. Denote by \mathbf{T} and \mathbf{N} the unit tangent and normal vectors at that point on the curve \mathcal{C}' respectively. Then, ignoring public randomization (which only makes D_t increase in expectation),

$$dW_t = r(W_t - g(A_t)) dt + \mathbf{T}\phi_t dZ_t,$$

where ϕ_t enforces action pair A_t on tangent \mathbf{T} . By Proposition 4, the trajectories of continuation values have natural curvature

$$\kappa = \frac{2(W_t - g(A_t)) \cdot \mathbf{N}}{r|\phi_t|^2} > \frac{2(v - g(A_t)) \cdot \mathbf{N}}{r|\phi_t|^2} \geq \max_{a \in \mathcal{A}(\mathbf{T}) \setminus \mathcal{A}^N} \frac{2(v - g(a)) \cdot \mathbf{N}}{r|\phi(a, \mathbf{T})|^2} = -f''(\mathbf{T}' \cdot W_t).$$

This implies that the drift of D_t is positive. □

Next, we need to prove that a solution to the optimality equation that is tangent to the boundary of $\mathcal{E}(r)$ at an arbitrary point $w \in \partial\mathcal{E}(r) \setminus \mathcal{N}$ does not enter the interior of $\mathcal{E}(r)$, denoted by $\mathcal{E}(r)^\circ$.

Lemma 2. Tangent curves do not enter $\mathcal{E}(r)$. *Consider point $w \in \partial\mathcal{E}(r) \setminus \mathcal{N}$ with an outward unit normal vector \mathbf{N} . Then the curve \mathcal{C} , which solves equation (21) from initial conditions (w, \mathbf{N}) , lies completely outside or on the boundary of the set $\mathcal{E}(r)$. It does not enter the interior of $\mathcal{E}(r)$.*

Proof. Suppose there is $v \in \mathcal{C} \cap \mathcal{E}(r)^\circ$, as shown in Figure 16. We will show how to construct a curve \mathcal{C}' with two endpoints $v_L, v_R \in \mathcal{E}(r)$ and a point $W_0 \notin \mathcal{E}(r)$ between them.

Take a neighborhood N_δ around point v in the interior of $\mathcal{E}(r)$. Without loss of generality, assume that point v is found by moving in the clockwise direction from point w along the curve \mathcal{C} , as shown in Figure 6. Let us choose a normal vector \mathbf{N}' by rotating \mathbf{N} in the counterclockwise direction. Consider the curve \mathcal{C}' that solves the optimality equation from initial conditions (w, \mathbf{N}') . From the continuity of solutions of the optimality equation in initial conditions, if \mathbf{N}' is sufficiently close to \mathbf{N} , then the curve \mathcal{C}' will enter the neighborhood N_δ of v . Because \mathbf{N}' is rotated counterclockwise relative to \mathbf{N} , the curve \mathcal{C}' will pass above the line P_w tangent to $\mathcal{E}(r)$ at w , before it enters neighborhood N_δ . Because there a unique tangent line P_w at point w , as argued earlier, the curve \mathcal{C}' will enter the interior of $\mathcal{E}(r)$ in the counterclockwise direction from w . Therefore, we can choose $W_0 \notin \mathcal{E}(r)$ that is between points v_L and $v_R \in \mathcal{E}(r)^\circ$ on the curve \mathcal{C}' , as shown in Figure 6. By Proposition 5, there is a PPE that achieves the value pair W_0 , so $W_0 \in \mathcal{E}(r)$, a contradiction. We conclude that the curve \mathcal{C} cannot enter the interior of the set $\mathcal{E}(r)$. □




Figure 16: Proof of Lemma 2.

We also need to prove one last fact that we used several times during the proof and in Section 6.

Lemma 3. (Solutions of the Optimality Equation). *Solutions of the optimality equation (21) exist locally and are continuous in initial conditions whenever the right hand side of (21) is positive.*

Proof. Consider initial conditions (w_0, \mathbf{T}_0) . Existence and continuity in w_0 follows because the right hand side of (21) is Lipschitz continuous in w . This is true because

$$H_a(w, \mathbf{T}) = \frac{2(g(a) - w) \cdot \mathbf{N}}{r|\phi(a, \mathbf{T})|^2}$$

is continuously differentiable in w for all a that are enforceable on tangent \mathbf{T} , and $\kappa = \max_{a \in \mathcal{A}(\mathbf{T}) \setminus \mathcal{A}^N} H_a(w, \mathbf{T})$. $H(w, \mathbf{T})$ may be discontinuous in \mathbf{T} when some action becomes unenforceable as we change the tangential angle. Nevertheless, continuity in \mathbf{T}_0 follows from the continuity in w_0 when the right hand side of (21) is strictly positive. \square

This concludes the proof of the Theorem. \square

Bibliography.

- [1] Abreu, D., Milgrom, P., and Pearce, D. (1991) “Information and Timing in Repeated Partnerships.” *Econometrica* Vol. 59, pp. 1713-1733
- [2] Abreu, D., Pearce, D., and Stacchetti, E. (1986) “Optimal Cartel Equilibria with Imperfect Monitoring.” *Journal of Economic Theory* Vol. 39, pp. 251-269
- [3] Abreu, D., Pearce, D., and Stacchetti, E. (1990) “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica* Vol. 58, pp. 1041-1063
- [4] Athey, S. and Bagwell, K. (2001) : “Optimal Collusion with Private Information.” *RAND Journal of Economics* Vol. 32, No. 3, pp. 428-465
- [5] DeMarzo, P. and Sannikov, Y. (2004) : “A Continuous-Time Agency Model of Optimal Contracting and Capital Structure,” working paper
- [6] Fudenberg, D., Levine, D., and Maskin, E. (1994) “The Folk Theorem with Imperfect Public Information.” *Econometrica* Vol. 62, pp. 997-1039
- [7] Harrison, J.M. (1985) *Brownian Motion and Stochastic Flow Systems*. J. Wiley & Sons, New York
- [8] Holmstrom, B. and Milgrom, P (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives.” *Econometrica* Vol 55, pp. 303-328
- [9] Hauser, C. and H. Hopenhayn (2004): “Trading Favors: Optimal Exchange and Forgiveness,” working paper
- [10] Kandori, M. (1992): “The Use of Information in Repeated Games with Imperfect Monitoring,” *Econometrica* 62, 997-1040
- [11] Kandori, M., and I. Obara (2003): “Efficiency in Repeated Games Revisited: The Role of Private Strategies,” working paper
- [12] Karatzas, I. and Shreve, S.: *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York, 1991
- [13] Matsushima, H. (1989): “Efficiency in Repeated Games with Imperfect Monitoring,” *Journal of Economic Theory*, 48, 428-442
- [14] Phelan, C. and Townsend, R. (1991) : “Computing Multi-Period, Information-Constrained Optima.” *Review of Economic Studies* Vol. 58, pp. 853-881
- [15] Protter, P.: *Stochastic Integration and Differential Equations*. Springer-Verlag, New York, 1990

- [16] Sannikov, Y. (2004) “A Continuous-Time Version of the Principal-Agent Problem,” working paper
- [17] Simon, L. and Stinchcombe, M. (1989) “Extensive Form Games in Continuous Time: Pure Strategies.” *Econometrica* Vol. 57, pp. 1171-1214
- [18] Williams, N. (2004) “On Dynamic Principal-Agent Problems in Continuous Time,” working paper