Continuously Repeated Games with Private Opportunities to Adjust

Asaf Plan*

UC Berkeley

April 30, 2010

Abstract

This paper considers a new class of dynamic, two-player games, where a stage game is continuously repeated but each player can only move at random times that she privately observes. A player's move is an adjustment of her action in the stage game, for example, a duopolist's change of price. Each move is perfectly observed by both players, but a foregone opportunity to move, like a choice to leave one's price unchanged, would not be directly observed by the other player. Some adjustments may be contrained in equilibrium by moral hazard, no matter how patient the players are. For example, a duopolist would not jump up to the monopoly price absent costly incentives. These incentives are provided by strategies that condition on the random waiting times between moves; punishing a player for moving slowly, lest she silently choose not to move. In contrast, if the players are patient enough to maintain the status quo, perhaps the monopoly price, then doing so does not require costly incentives. Deviation from the status quo would be perfectly observed, so punishment need not occur on the equilibrium path. Similarly, moves like jointly optimal price reductions do not require costly incentives. Again, the tempting deviation, to a larger price reduction, would be perfectly observed.

This paper provides a recursive framework for analyzing these games following Abreu, Pearce, and Stacchetti (1990) and the continuous time adaptation of Sannikov (2007). For a class of stage games with monotone public spillovers, like differentiated-product duopoly, I prove that optimal equilibria have three features corresponding to the discussion above: beginning at a "low" position, optimal, upward moves are impeded by moral hazard; beginning at a "high" position, optimal, downward moves are

*Email: asaf@econ.berkelϵy.edu. This paper has benefited from discussions with Matthew Rabin, Xin Guo, Kyna Fong, Robert Anderson and Steve Tadelis, and seminar participants at UC Berkeley.

unimpeded by moral hazard; beginning at an intermediate position, optimally maintaining the status quo is similarly unimpeded. Corresponding cooperative dynamics are suggested in the older, non-game-theoretic literature on tacit collusion.

1 Introduction

This paper considers a new class of dynamic, two-player games, where a stage game is continuously repeated but each player can only move at random times that she privately observes as they occur. As usual, imperfect observation yields issues of moral hazard, but the monitoring structure here departs from those considered previously. Call a player's action within the stage game her *position*, which is a perfectly observed state variable of the dynamic games considered here. Here, a player's action is an adjustment or deliberate non-adjustment of that position. The only action that is imperfectly observed is deliberate non-adjustment, which the other player cannot directly distinguish from a lack of opportunity. If non-adjustment is tempting relative to the equilibrium moves, then the constraints of moral hazard are binding. This paper studies the resulting dynamics of the players' positions in optimal equilibria.

A central application is a model of duopoly where each firm at each point is unsure of the length of time required before the other is able to adjust its price. Recall, the existing literature on repeated games with imperfect public monitoring instead provides models where each firm is unsure about the prices that its rival has set. Both types of uncertainty are suggested in the earlier, non-game-theoretic literature.[1] Recall, if prices are imperfectly observed, any price profile above the competitive level cannot be indefinitely maintained in equilibrium.[2] Even though the duopolists initially adhere to the proposed price, after random signals suggesting that one has deviated to a lower price, punishment like lower

[1] Chamberlin (1929), in one of the first papers to discuss tacit collusion, suggests that a firm may be unsure "not as to what his rival will do, but as to when he will do it" in response to a price cut. See section 6 for further discussion.

[2] This conclusion holds under the standard technical assumptions, including in particular that the support of the signal's probability distribution does not depend on the prices.

pricing by its rival must be carried out. Otherwise, the firms would prefer to deviate to a lower price.

Here, it is only intentional non-adjustment of one's price that would be imperfectly observed. Suppose the present price is optimal among those that can be maintained in equilibria, for example the monopoly price given enough patience. Each duopolist might be tempted to adjust to some lower price when able, but this action would be perfectly observed, so the incentive constraints of the moral hazard problem are not immediately binding. Suppose instead the present price is too high, perhaps following a one-time drop in the duopolists' production cost or a similar shock. The optimal course of adjustments is downward. Each firm might be tempted by an even larger downward adjustment, but again this would be perfectly observed. Lastly suppose the duopolists begin at low prices, like the competitive price level. The optimal course of adjustments is upward, and here non-adjustment may be tempting — I show that it will be at some point along the course. At some point, the moral hazard constraints will bind, and the course of upward adjustments will be constrained relative to the case where adjustment opportunities are publicly observed.

Such a strategic difficulty specifically regarding upward price adjustment is suggested in the older, non-game-theoretic literature on tacit collusion. Salop (1986) writes that a duopoly beginning at the competitive price may suffer "transitional difficulties of achieving the co-operative outcome," that is, the monopoly price. During the "transition period" where one firm but not the other has raised its price, of course the leader's profits are reduced and the follower's increased. Salop then writes, "the follower has every incentive to delay its price increase. Fear of further delays may convince [the leader] that it should return to the [competitive price] or should forgo the price increase to begin with." Galbraith (1936) suggests that collusive price increase is more difficult than decrease: "The problem of price changes under oligopoly is probably even greater when a price increase is involved. Here unanimity of action is essential and there is a positive premium to the individual who fails to conform to the increase."[3] Sweezy (1939) presents a model of kinked demand curves, where each firm expects its rivals to match price reductions but not price increases. That famous model has widely been attacked as being without foundation. In the model here, I find, perhaps in accordance with Sweezy's intuition, that the firm's rival might be tempted to leave its price unchanged rather than match an increase, but not a decrease. Out-of-equilibrium price increases would generally not be matched, while out-of-equilibrium price drops would generally be matched and worse.[4]

As Salop further remarks, "It may appear that the 'transitional' difficulties of achieving the co-operative outcome are only a one-time problem. However this view overlooks the dynamic elements of oligopoly interaction. As cost and demand parameters change over time, the joint profit-maximizing point changes as well. Thus, oligopolists face repeated transitional problems." One may interpret the model here as applying to the limiting case where there has just been a shock to such external parameters, but no further change is expected. A future paper might consider the more general case, which adds an external state variable. I anticipate that the results here correspond to richer but related results in the more general model. Galbraith writes, "It is fair to suppose that the day when a price increase is necessary is never far from mind when price decreases seem desirable." With private adjustment opportunities, a momentarilly beneficial price decrease would exacerbate the moral hazard issue on that day "when a price increase is necessary." This might lead to the often noted pattern of asymmetric price adjustment.[5]

[3]However, in the model I present, there is no strategic difficulty in decreasing prices toward the competitive level.

[4]More precisely, in optimal equilibria, price increases greater than on the equilibrium path would not be matched or at least not matched quickly enough to make them unilaterally worthwhile.

[5]One can imagine a simple collusive arrangement like the following. If the current price is more than two times the current cost, lower your price to two times the cost, at the first opportunity. If the current price is less than the current cost, raise your price to the current cost. This scheme approximately has the feature that no costly incentives are required; the firms are never called to raise price above cost. (Perhaps the optimal scheme looks like this if the firms are very impatient because incentives would then be very costly.) If costs are secularly increasing, then on the equilibrium path of this scheme, the firms will often find themselves with prices near costs. Consequently, they will rarely be lowering their prices in response to cost reductions, while they will often be raising their prices in response to cost increases.

The issue of moral hazard surrounding adjustment is driven by uncertainty regarding whether or not temporary non-adjustment by one's rival is intentional. Section 2 formally presents the model, where such uncertainty is parsimoniously captured by the assumption that opportunities to move are random, following a Poisson process.[6] The model yields dynamic games with a non-standard imperfect public monitoring structure. These games allow a recursive formulation of equilibria adapting Sannikov's (2007) continuous-time methods, which in turn build on the earlier work by Abreu, Pearce, and Stacchetti (1990) in discrete time. Section 3 presents this formulation. The problem of determining optimal equilibria then corresponds to a stochastic optimal-control problem. Section 4 presents this general problem and two preliminary results: First, in extremal equilibria, continuation payoffs generally remain extreme. Second, when the moral hazard incentive constraints bind, continuation payoffs drift along the achievable boundary between adjustments, against the player with the greater incentive constraint. When these constraints do not bind, in extremal equilibria there is generally no drift between adjustments; that is, the path of adjustments does not condition on the realized time between adjustments. The games here do not permit a general characterization of the set of achievable payoffs like Sannikov's optimality equation.[7] Instead, section 5 considers a class of stage games with monotone public spillovers, like differentiated-product duopoly, and the three main results, which correspond to the three courses of duopoly price adjustment described above. In these stage games, there is a ranking of "higher" and "lower" actions. The set of Pareto optimal positions is above the static Nash point. I find that beginning at a low position, the moral hazard incentive constraints eventually bind on the course of optimal, upward adjustments. In contrast, beginning at

[6] This may be the simplest assumption that yields such uncertainty. A more realistic model where each player is uncertain particularly about the other player's opportunities to move would raise issues of adverse selection in addition to moral hazard. Further realism would be gained by endogenizing future opportunities, or at least their rate. If opportunities are in fact endogenous, then perhaps the model here can be viewed as a reduced form where the players have failed to communicate or coordinate their respective opportunities. Uncertainty about temporary non-adjustment seems reasonable in some situations of tacit collusion. This is discussed further in section 6.

[7] Sannikov's optimality equation derives from the application of Ito's formula in his continuous, diffusion setting. The noise here is Poisson; there are jumps. Additionally the games here are dynamic rather than repeated.

a high position, these constraints never directly bind on the course of optimal downward adjustments. Similarly, beginning at some intermediate position, it is optimal to maintain that position perpetually, and again the incentive constraints do not directly bind. Finally, section 6 discusses the assumptions that opportunities to adjust are random and private, and argues that the model here may be interpreted as a reduced form of one with informational and rationality constraints suggested in the literature. Such a more general model would still lead to the issues of moral hazard regarding adjustments that are analyzed in the relatively simpler model here.

2 The game

Consider a symmetric, two-player stage game.[8] Each player's action space is $\mathcal{A}^1 = \mathcal{A}^2 \subset \mathbb{R}$; let $\mathcal{A} = \mathcal{A}^1 \times \mathcal{A}^2$. (Superscripts will generally identify the players, rather than denoting exponents. The index $i \in \{1, 2\}$ will often denote the player in question, while $j \neq i$ denotes the other player.) The joint payoff function is $g : \mathcal{A} \to \mathbb{R}^2$, which is symmetric: $g_1(a_1, a_2) = g_2(a_2, a_1)$. The game satisfies the standard restrictions, from Mailath and Samuelson (2006): $\mathcal{A}^i$ is either finite, or a compact and convex Euclidean subspace. In the latter case, g is continuous and g_i is quasiconcave in a_i. This paper considers a corresponding dynamic game with privately Poisson-distributed adjustment opportunities, where this stage game is continuously repeated. At each moment $t \in [0, \infty)$, each player $i \in \{1, 2\}$ takes an action $A_t^i \in \mathcal{A}^i$, which I interpret as an adjustment target. Payoffs do not directly depend on this target, but on an associated position process, $P_t \in \mathcal{A}$, which evolves as follows:

$$dP_t = (A_t - P_t) \cdot dO_t, \tag{1}$$

where O^1 and O^2 are independent Poisson processes, each with intensity α. That is, P^i adjusts from its previous value to A^i at player i's opportunity times, where $dO_t^i = 1$. Flow payoffs are $g(P_t)$, because the position in the dynamic game corresponds to the action profile in the stage game. Note, the waiting time to a player's next adjustment opportunity is always distributed exponentially with constant intensity α.[9] Players do not directly observe each other's actions, A^j, or opportunities, dO^j, but perfectly observe the joint position process, P. While P^i is constant, player j is uncertain whether player i has no opportunity ($dO_t^i = 0$) or chooses not to adjust ($A_t^i = P_t^i$).

I consider pure, public strategies. Such a strategy for player i is a stochastic process that

[8] The extension of this section to non-symmetric stage games is straightforward; to be done in a future draft.

[9] Higher intensity α means that opportunities come more frequently; the mean waiting time is $1/\alpha$. For simplicity rather than realism, the intensity is independent of the position vector, previous waiting time and the arrival of the other player's opportunities. A somewhat generalized opportunity process might allow α to vary with such circumstances.

maps each public history $H_t = \{P_s\}_{s=0}^t,\ t \in [0,\infty)$ to an adjustment target A_t^i. Given a strategy profile A and an initial position profile P_0, the vector of discounted-average payoffs up to time t is

$$U_t(P_0, A) = r\int_0^t e^{-rs} g(P_s)ds, \tag{2}$$

which is stochastic. U depends on the evolution of the position process, which depends on the given strategy profile and the opportunity process. At each time t, the players will maximize their respective components of the (expected) continuation payoffs:

$$W_t(P_t, A) = \mathrm{E}_t\left[r\int_t^\infty e^{-r(s-t)} g(P_s)ds\right]. \tag{3}$$

where P evolves as described in (1).

Let $\mathcal{V} = \mathrm{co}\{g(a) : a \in \mathcal{A}\}$, the convex hull of feasible payoffs in the stage game (note it is bounded). In the following sections we seek the mapping $\mathcal{E} : \mathcal{A} \to P(\mathcal{V})$, where $\mathcal{E}(p)$ is the set of equilibrium payoffs beginning at position p.

As described above, this is a dynamic game with a non-standard, imperfect public monitoring structure. It is equivalent to a game corresponding to the motivating story, where players only act at their opportunity times. In this case, A_t^i is the component of player i's strategy describing the adjustment action that she would take at time t if she had an opportunity ($dO_t = 1$). (In the game above, the target A_t^i is itself player i's action at time t, though it has effect only if $dO_t = 1$.) The game is dynamic rather than simply repeated, as with asynchronous adjustments, the position, P, is a state variable. If adjustment opportunities were instead publicly observable, this would be a particular asynchronously repeated game in continuous time, termed Poisson Revisions by Lagunoff and Matsui (1997).

Here the stage game itself is continuously repeated, without any change in parameters like the duopolists' production costs. A main interest is the equilibrium dynamics of the position profile given an exogenous starting value. One can view this starting value as determined by prior parameter values, in which case the analysis here corresponds to a situation where there has just been a shock to these parameters but no further change is expected. A more general analysis would incorporate these parameters as exogenously varying state variables.

3 The evolution of PPE continuation payoffs

This section follows Sannikov (2007) in formulating the players' continuation payoffs as stochastic processes. Given this formulation, the problem of determining optimal equilibria corresponds to a stochastic optimal control problem, which is pursued in the next section. Here strategies are pure public strategies unless otherwise stated.

The following proposition and proof closely follow Sannikov (2007, Proposition 1).

Proposition 1 (Representation & Promise Keeping). *A bounded stochastic process W^i is the continuation value $W^i(A)$ of player i under strategy profile A if and only if there exist finitely-valued processes J_t^{ki} such that for all $t \geq 0$,*

$$dW_t^i = r(W_t^i - g_i(P_t))dt + \sum_{k=1,2} (J_t^{ki}(dP_t^k + P_t^k)dO_t^k - \alpha J_t^{ki}(A_t^k)dt), \tag{4}$$

where P is determined by 1, and $J_t^k(P_t^k) = 0$.

The proposition states that continuation payoffs can be decomposed into a Poisson-Martingale determined by realized adjustments plus a drift term that compensates for the difference between expected discounted-average and current flow payoffs ("promise keeping"). The "if" direction of the proof relies on the Poisson-margingale representation theorem, while "only if" relies on martingale convergence.

Proof. The following process is a martingale:

$$V_t^i(A) = r\int_0^t e^{-rs} g_i(P_s)ds + e^{-rt}W_t^i(A) = \mathrm{E}_t\left[r\int_0^\infty e^{-rs} g_i(P_s)ds \middle| A\right].$$

By the Poisson-martingale representation theorem (see Hanson (2007, Theorem 12.11)), we get a representation

$$V_t^i(A) = V_0^i(A) + \int_0^t e^{-rs} \sum_k J_s^{ki}(d\tilde{O}_s^k - \alpha ds),$$

where J_t^k is finite and $\tilde{O}$ represents the public portion of the Poisson opportunity process; $d\tilde{O}_s^k = 1_{A_s^k \neq P_s^k} dO_s^k$. Combining the previous two expressions and differentiating with respect to t yields,

$$re^{-rt} g_i(P_t)dt - re^{-rt} W_t^i(A)dt + e^{-rt} dW_t^i(A) = e^{-rt} \sum_k J_t^{ki}(dO_t^k - \alpha dt),$$

where $J_t^{ki} = 0$ if $A_t^k = P_t^k$. Which yields the desired expression:

$$dW_t^i(A) = r\left(W_t^i(A) - g_i(P_t)\right) dt + \sum_k J_t^{ki}(dO_t^k - \alpha dt).$$

Regarding the converse,

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

is a Martingale under the strategies A. Further Martingales V_t^i and $V_t^i(A)$ converge as $e^{-rt}W_t^i$ and $e^{-rt}W_t^i(A)$ converge to 0. Because $V_t^i = \mathrm{E}_t[V_\infty^i] = \mathrm{E}_t[V_\infty^i(A)] = V_t^i(A)$, we have $W_t^i = W_t^i(A)$, as desired. □

At this point we depart somewhat from Sannikov's line of argument, as the monitoring structure here is qualitatively different.

The previous result does not require that A is an equilibrium. It is if and only if at each point, for each player, the proposed action maximizes the jump in her payoff given an adjustment opportunity; that is, $A_t^k = \arg\max_{\tilde{a} \in \mathcal{A}^k} J_t^{kk}(\tilde{a})$. (The one-shot deviation principle holds here by the usual argument.) In the game considered here, we distinguish between two types of deviations. A player may choose to not adjust, $\tilde{A}_t^k = P_t^k$, in which case $J_t^{kk}(\tilde{A}_t^k) = 0$, because this deviation is not revealed in the public history. Secondly, a player may choose to make an alternative adjustment, although such a deviation is instantly, publicly revealed. Alternative adjustments may be disincentivized by then minmaxing the player, as in the usual dynamic games of perfect information.

Lemma 2 (Equilibrium restrictions on J)**.** *Player i's strategy A^i is optimal in response to some $\tilde{A}^j$ if and only if there exists a process W satisfying the conditions of Proposition 1 subject to the following restrictions on J on the equilibrium path, for all $t \geq 0$, and $i = 1, 2$,*

$$J_t^{ii}(A_t^i) \geq 0 \tag{IC}$$

$$W_t^i + J_t^{ii}(A_t^i) \geq \underline{v}(P_t^j) = \max_{\tilde{a}^i} \min\{\tilde{w}^i : \tilde{w} \in \mathcal{E}(\tilde{a}^i, P_t^j)\}. \tag{IR}$$

Proof. Take $\tilde{A}^j$ to coincide with A^j on the equilibrium path but to inflict the minmax punishment $\underline{v}$ after any deviation. □

The IC condition implies for each player that realizing her equilibrium adjustment is better than non-adjustment. The IR conditions implies for each player that realizing her equilibrium adjustment is better than any alternative adjustment followed by the worst continuation equilibrium for her at the new position.

Combining the previous two results, we get a characterization of equilibrium payoffs.

Theorem 3 (Characterization of PPE)**.** *In any equilibrium A, the pair of continuation values is a process in $\mathcal{V}$ that satisfies*

$$dW_t = r(W_t - g(P_t))dt + \sum_{i=1,2} J_t^i(dO_t^k - \alpha dt), \tag{5}$$

where J is finite-valued and satisfies (IC) and (IR), and P is determined by 1.

Conversely, if W satisfies these conditions, it corresponds to some equilibrium having the same outcome as A.

Definition 4. A mapping $\mathcal{W}(p) \subset \mathbb{R}^2$ is self-generating if and only if, for all $p \in \mathcal{A}$, for any point $w \in \mathcal{W}(p)$, for P starting at p there exists a process W that starts at w, stays in $\mathcal{W}(P)$, and satisfies the conditions of the previous Theorem.

Corollary 5. *$\mathcal{E}$ is self-generating and its graph contains those of all other bounded, self-generating mappings.*

In general the set of equilibrium payoffs under mixed strategies may be larger than that under pure strategies, butFudenberg and Levine (1994, Theorem 5.2) show that this is not so in repeated games with a product monitoring structure, where the public signals do not jointly condition on both players' actions. The monitoring structure here has this feature. I conjecture that their result can be adapted to this class of dynamic games. (Intuitively, here neither player's private history, her forgone opportunities and unrealized adjustments, reveals anything about the other's; so it cannot be used to generate further incentives.)

Conjecture 6. *Here the set of pure-strategy payoffs coincides with the set of mixed-strategy payoffs.*

As the noise structure here is (intrinsically) Poisson rather than Brownian, I am not able to give a characterization of $\mathcal{E}$ like Sannikov's optimality equation.[10] Instead, in the next two sections, I analyze how and when the IC constraint binds. First, two remarks.

Remark 7. At each position P, the set of equilibrium payoffs, $\mathcal{E}(P)$, is contained in the set of payoffs achieveable in equilibrium of the public opportunities benchmark game (which requires only IC not IR). In turn, this is contained in the set of feasible payoffs (which requires neither IC nor IR), which is contained in the convex hull of the set of stage-game payoffs.

[10]The optimal equation is an ordinary differential equation for the boundary of the set of achievable payoffs in the class of imperfect-monitoring games that Sannikov considers. It is derived from Ito's lemma, exploiting the Brownian noise structure of those games. Kalesnik (2005) parallels Sannikov's approach for continuously repeated games of imperfect monitoring but with a Poisson noise structure; he is not able to provide a crisp characterization of achievable payoffs like the optimality equation.

Even when publicly observed, the fact that adjustment opportunities are asynchronous affects the set of equilibrium payoffs. (See Yoon, 2001; Lagunoff and Matsui, 1997; Wen, 2002) I am interested in the additional effect of opportunities being private. We see already that (IC) will only shrink the set of equilibrium payoffs. We will see that this shrinking is not uniform; (IC) does not eliminate certain extremal payoffs. It is easy to see to that maintaining the status quo satisfies IC:

Remark 8. Consider the outcome where the present position is perpetually maintained, $A = P$. The IC constraint is trivially satisfied on this path, because $J^i = 0$.[11]

Beginning at some positions, like the monopoly price given enough patience, maintaining the status quo is optimal, and IC does not directly bind on this optimal outcome. Consider the restriction on cooperation in standard supergames with fixed discounting. Some action profiles are not achievable in equilibrium as the payoff from a single period of deviation is too tempting. This can be cast as a restriction on maintaining cooperative positions. Similarly consider the restriction on cooperation imposed by standard imperfect monitoring in discretely or continuously repeated games. Some action profiles may no longer be achievable in optimal equilibrium as the necessary incentives are too costly to be worthwhile. Further, beginning at a cooperative position it will be necessary to leave it in equilibrium following bad signals, in order to provide incentives. These too seem like restrictions on maintaining cooperative positions. The restriction on cooperation due to IC is qualitatively different. IC does not directly restrict the maintenance of a cooperative position, but may restrict the achievement of such a position from a different starting position. The next section shows that IC may limit both the speed and course of cooperative adjustment. Section 5 shows that this restriction on adjustment is asymmetric in a subclass of games with a monotone positive externality; beginning at a lower position, IC will restrict optimal upward adjustment, but beginning at a higher position, downward adjustment is unrestricted.

[11] However IC may bind out of equilibrium on the punishment path, thus increasing the minimum payoff necessary to satisfy IR. I will generally consider cases where r/α is small enough that IR is satisfied, even given the potentially reduced punishments satisfying IC. Punishment equivalent to Nash reversion should satisfy IC, but is not so easy to define here with asynchronous adjustments; see Dutta (1995).

4 PPE with extreme values

Given a starting position $p \in \mathcal{A}$, consider the equilibrium continuation payoffs with the largest weighted sum in the direction N,

$$w(N;p) = \arg\max_{w \in \mathcal{E}(p)} w \cdot N, \ |N| = 1. \tag{6}$$

This section characterizes the instantaneous values of the "controls," $a^i, \hat{w}^i = w + j^i(a^i)$ and $\dot{w}$, given the maintained assumption that $\partial\mathcal{E}(p)$ is differentiable at this point $w(N;p)$. (Here I drop the subscript $_t$ and take these lower case variables to denote time-t values of the corresponding processes, for some generic t. I also normalize $r + 2\alpha = 1$.) From Theorem 3 we have the following characterization of payoffs in terms of the instantaneous controls, and equilibrium constraints on those controls:

Corollary 9. *The extremal payoff in direction N satisfies the following Hamilton-Jacobi-Bellman equation,* [12]

$$w(N;p) = \max_{a^i, \hat{w}^i, \dot{w}} rg(p) + \alpha \sum_i \hat{w}^i + \dot{w},$$

where $\hat{w}^i = w + j^i$ and the following conditions are satisfied.

$$N \cdot \dot{w} \leq 0, \tag{Fd}$$

$$\hat{w}^i \in \mathcal{E}(a^i, p^j), \tag{Fj}$$

$$\hat{w}^{ii} \geq w^i, \tag{IC^i}$$

$$\hat{w}^{ii} \geq \underline{v}(p^j). \tag{IR^i}$$

Condition (Fj) states that the continuation payoffs after adjustment are themselves equi-

[12] That is, $w = g(p) + \frac{\mathrm{E}[dW]}{r\,dt}$. Recall dW is described in Theorem 3.

librium payoffs at the new position. Here, condition (Fd) implies that the drift between adjustments does not take us out of the present set of equilibrium payoffs. We saw the IC and IR constraints before, which imply that neither player prefers to hide an adjustment opportunity or to make an out-of-equilibrium adjustment, respectively.

The results in this section are based on analysis of the Lagrangian corresponding to this constrained maximization,

$$\mathcal{L}(a^i, \hat{w}^i, \dot{w}; N, p) = \underbrace{\left(rg(p) + \alpha \sum_i \hat{w}^i + \dot{w} \right) \cdot N}_{\mathcal{H}} - \rho \underbrace{(N \cdot \dot{w})}_{\leq 0 \text{ (Fd)}}$$
$$+ \sum_i \left(\lambda^i \underbrace{\left(\hat{w}^{ii} - rg_i(p) - \alpha \hat{w}^{ji} - \dot{w}^i \right)}_{\geq 0 \text{ (IC}^i\text{)}} + \mu^i \underbrace{\left(\hat{w}^{ii} - \underline{v}(p^j) \right)}_{\geq 0 \text{ (IR}^i\text{)}} - \nu^i \underbrace{\left(d(\hat{w}^i, \mathcal{E}(a^i, p^j)) \right)}_{\leq 0 \text{ (Fj}^i\text{)}} \right), \tag{7}$$

where $d(w, \mathcal{E}(p))$ is defined as follows: if $w \notin \mathcal{E}(p)$, it is the (positive) distance between the point and the set, while if $w \in \mathcal{E}(p)$, it is minus the distance between w and the complement of $\mathcal{E}(p)$, so negative. Notice that the constraints are linearly independent, so the conditions of the Karush-Kuhn-Tucker theorem are met, yielding the following results:

Stationarity:	$\nabla \mathcal{L} = 0$ (over the controls: $a^i, \hat{w}^i, \dot{w}$)
Primal Feasibility:	All the constraints (Fd,IC,IR,Fj) are satisfied
Dual Feasibility:	All the multipliers, $\lambda^i, \mu^i, \nu^i, \rho$ are ≥ 0
Complementary Slackness:	Each constraint holds with equality or the associated multiplier is zero

Lemma 10 (Payoffs after adjustment are extremal.)**.** *Generically, the payoffs after adjustment are on the new boundary:* $\hat{w}^i \in \partial \mathcal{E}(a^i, p^j)$ *if* $N^i \neq -1$.
Further, the adjustment targets maximize this boundary locally: $\hat{w}^i \in \partial \left(\bigcup_{\tilde{a}^i \in \mathcal{A}} \mathcal{E}(\tilde{a}^i, p^j) \right)$ *if* $N^i \neq -1$.

Proof. Case $N^i > 0$:

$$0 = \frac{\partial \mathcal{L}}{\partial \hat{w}^{ii}} = \alpha N^i + \lambda^i + \mu^i - \nu^i d^i(\hat{w}^i, \mathcal{E}(a^i, p^j)) \quad \text{(Stationarity)}$$

$$\Rightarrow \nu^i > 0 \text{ (and } d^i(...) > 0) \quad \text{(Dual feasibility for } \lambda \text{and } \mu)$$

$$\Rightarrow d(\hat{w}^i, \mathcal{E}(a^i, p^j)) = 0 \quad \text{(Complementary slackness for (Fj) with } \nu)$$

(This also shows that $\hat{N}^{ii} > 0$.)

Case $N^j < 0$: By a similar argument applied to $\partial \mathcal{L}/\partial \hat{w}^{ij}$ we have $d(\hat{w}^i, \mathcal{E}(a^i, p^j)) = 0$ and $\nu^i > 0$, $d^j(\hat{w}^i, \mathcal{E}(a^i, p^j)) < 0$ (so also have $\hat{N}^{ij} < 0$).

Case $N^i < 0$ and $N^j > 0$: Here $\lambda^i = 0 = \lambda^j, \rho = 1$. Again considering $\partial \mathcal{L}/\partial \hat{w}^{ij}$, $\nu^i > 0$, $d(\hat{w}^i, ...) = 0$, $d^j(...) > 0$ (so also have $\hat{N}^{ij} > 0$).

Case $N^i = 0, N^j = 1$: Here $\lambda^i = 0$ and $\lambda^j = (1 - \rho)$. Again considering $\partial \mathcal{L}/\partial \hat{w}^{ij}$, $\rho = \nu^i d^j(\hat{w}^i ...)$ so $d(\hat{w}^i, ...) = 0$, $d^j(...) > 0$ (so also have $\hat{N}^{ij} > 0$).

In each of these cases considering $\nu^i > 0$ in the FOC $0 = \partial \mathcal{L}/\partial a^i$ implies that $\frac{\partial}{\partial a^i} d(\hat{w}^i, \mathcal{E}(a^i, p^j)) = 0$. Combining this with the fact that the distance is zero implies that $\hat{w}^i$ is on the envelope of the boundaries as claimed.

As noted, these results don't hold in the last, non-generic case, $N^i = -1, N^j = 0$: Here $\lambda^j = 0 = \nu^i$, $\lambda^i = \rho - 1$. □

Lemma 11 (Payoffs stay extremal between adjustments.)**.** *Any drift is along the boundary:* $\dot{w} \cdot N = 0$.

Proof. Must have $\dot{w} \cdot N \leq 0$ (primal feasibility). Suppose $\dot{w} \cdot N < 0$ then $\rho = 0$ (comp. slackness), so $\lambda^i = N^i$ (from $\partial \mathcal{L}/\partial \dot{w}^i = 0$), so $\nu^i = 0$ (from $\partial \mathcal{L}/\partial \hat{w}^{ij} = 0$), so $\mu^i = -(1+\alpha) N^i$ (from $\partial \mathcal{L}/\partial \hat{w}^{ii} = 0$), so must have that $N^i = 0$, but similarly must have that $N^j = 0$ — contradiction. □

The previous two results imply that, in equilibria with extreme payoffs, continuation payoffs remain extreme; payoffs do not move into the interior of the achievable set.

Lemma 12 (Strategies condition on waiting times iff there is moral hazard.)**.** *If the IC constraints are slack, then there is no drift:* $\dot{w} = 0$*. If the IC constraints are binding, there is drift unless the corresponding Lagrange multipliers satisfy* $\lambda \cdot T = 0$*.*

Proof. Given $\dot{w} \cdot N = 0$, we can rewrite the Lagrangian with (Fd) holding with equality; separately, we can rewrite IC in terms of w:

$$\mathcal{L} = \underbrace{\left(rg(p) + \alpha \sum_i \hat{w}^i\right) \cdot N + f'(w \cdot T) \overbrace{\left(w - rg(p) - \alpha \sum_i \hat{w}^i\right)}^{=\dot{w}} \cdot T}_{\mathcal{H}}$$
$$+ \sum_i \left(\lambda^i \underbrace{\left(\hat{w}^{ii} - w^i\right)}_{\geq 0 \text{ (IC}^i)} + \mu^i \underbrace{\left(\hat{w}^{ii} - \underline{v}(p^j)\right)}_{\geq 0 \text{ (IR}^i)} - \nu^i \underbrace{\left(d(\hat{w}^i, \mathcal{E}(a^i, p^j))\right)}_{\leq 0 \text{ (Fj}^i)}\right)$$

Now consider the "envelope condition":

$$f'(w \cdot T) = \frac{\partial \mathcal{L}}{\partial w} \cdot T = f'(w \cdot T) + f''(w \cdot T)\dot{w} \cdot T - \lambda \cdot T \quad \Rightarrow f''(w \cdot T)\dot{w} \cdot T = \lambda \cdot T$$

so $\lambda = 0$ implies $\dot{w} \cdot T = 0$. Together with $\dot{w} \cdot N = 0$, this yields $\dot{w} = 0$, as desired. □

This implies that targets drift between adjustments only when IC is binding. If opportunities are public, targets do not drift. That is to say, the course of adjustments does not condition on the realized waiting times between adjustments.

These results are of limited practical interest, but they serve as the foundation for the analysis in the next section of a specific class of stage games.

5 Moral hazard binds moving "up" but not "down"

The previous section presents general features of extremal equilibria for generic stage games. This section presents more specific results for a class of stage games with monotone public spillovers. This class includes differentiated-product price competition. For an introduction to these results, first consider the following stage game.

Example (PD with a third option).

	C	c	d
C	0,0	-2,3	-4,4
c	3,-2	1,1	-1,2
d	4,-4	2,-1	0,0

The bottom-right 2x2 portion is the standard prisoner's dilemma. Here payoffs are additively separable across the two players' actions, $g_i(a) = 2a_j - a_i^2,\ j \neq i$ where the three actions have the following numerical values, $d = 0,\ c = 1,\ C = 2$. Each player's payoffs are monotone increasing in the other's action. However, the joint-payoff maximizing action profile is the middle one, c, c. Unlike c, the incremental cost to one of playing C is greater than the benefit to the other. The third option, C, involves too much self-sacrifice, it is "higher" than the optimal profile.

Consider the corresponding dynamic game with Poisson-adjustment opportunities. Beginning at any position, the symmetric-optimal outcome requires that each player takes her first opportunity to adjust to c and then stays there. With public opportunities, beginning anywhere this outcome is achievable in equilibrium given that r/α is small enough. With private opportunities, beginning at a position in the upper-left four squares, this outcome is also achievable given the identical threshhold on r/α. Here the equilibrium adjustments are downward, and each player independently wants to make them. The players would individually like to adjust further downward but can be dissuaded by trigger strategies that threaten d, d following an out-of-equilibrium adjustment. These incentives are identical with public and private opportunities, and they are costless as the threat is never carried out in equilibrium. In contrast, consider the private-opportunities game beginning at one of the five squares on the right and lower edges. Here the proposed outcome is not achievable in equilibrium. Some player is called on to adjust upward to c from d, which she prefers not to do absent some incentive. As her opportunities are private, incentives must take the form of reward or punishment conditioning on how long it takes her to adjust. These incentives are costly as the players are not able to efficiently transfer payoffs between themselves and the reward/punishment must be carried out with some probability in equilibrium, as opportunities are stochastic. With private opportunities, the players may be able to adjust upward to c, c, but not both certainly and as quickly as feasible.

The results of this section follows this example. I consider a class of games with monotone externalities, and some other regularity conditions. Each player's flow playoffs are increasing in the other's action. Here private opportunities directly limit upward adjustments but not downward adjustments nor non-adjustment. I believe that these are the first game-theoretic foundations for such a restriction on cooperative dynamics.

I make the following assumption on the stage game payoffs, $g : \mathcal{A} \to \mathbb{R}^2$. (1) g is symmetric across the two players, (2) g_i is increasing in a_j, $j \neq i$, (3) g is twice differentiable, (4) g is concave, (5) the actions are strategic complements, that is the cross partial is non-negative.Consider the frontier of Pareto optimal action profiles, $\mathcal{P}^O$. The N-optimal profile is, $P^O(N) = \arg\max_{p \in \mathcal{P}^O} g(p) \cdot N$.

If the players are not patient enough to go higher, then maintaining the present position is optimal, and the moral hazard constraints do not directly bind:

Proposition 13 (Status quo)**.** *Suppose the initial position is on the Pareto frontier in direction N, $P^O(N)$. If IR is satisfied maintaining this point, then so is IC, and so the N-optimal is achievable in equilibrium, where the players maintain this position indefinitely.*

Proof. Given the assumption that $g \cdot N$ is quasi-concave, staying indefinitely at any point on $\mathcal{P}^O$ is not Pareto dominated by any other mixture over positions, feasible or otherwise. Thus it is N-optimal to stay at the original position indefinitely, $A = P^O(N)$. Here IC is trivially satisfied with equality, $W^i = W = g(P^O(N))$. (Note IC is saturated but not binding.) If IR is not satisfied, it may be that IC impacts the minmax value. □

If the players find themselves at a higher position, then the moral hazard constraints do not bind on the optimal course of downward adjustments.

Proposition 14 (Downward adjustment–Symmetric case)**.** *Suppose the initial position P is strictly above the Pareto frontier and symmetric. If IR does not bind in maintaining $P^O(N^s)$, $N_1^s = N_2^s$, then the symmetric-optimal feasible outcome is achievable. On the equilibrium path, each player's adjustment target depends only on the other's position. The players adjust downward quickly in the sense that they do at least as well if they adjusted to the eventual position as quickly as is feasible.*

Proof. I will show: The adjustments are monotone downward. While each player may not like the other's downward adjustment, each is better off after every pair of adjustments, beginning with their own, then they would be if they stopped adjusting.

Suppose player 2 is the first to adjust. Then player 1's kth adjustment, p_1^k satisfies the FOC,

$$0 = \frac{d}{dp_1}\left(g\left(p_1, p_2^k\right) + \frac{\alpha}{r+\alpha} g\left(p_1, p_2^{k+1}\right)\right) \cdot N,$$

by the envelope theorem; and similary for player 2. The adjustments are monotone downward: Starting at the symmetric position above the Pareto frontier, $p_2^1 < p_2^0$. Then, as the actions are strategic complements, $p_1^1 \leq p_2^1 < p_1^0$, and so on. Write $p_1^k = a(p_2^k)$

I want to show that each player is always better off following at least one more equilibrium adjustment. That is for player one,

$$0 \leq g_1(p^n, p^{n+1}) - g_1(p^{n-2}, p^{n-1}) = \int_{p^{n-2}}^{p^n} \frac{d}{dp_1} g_1(p_1, a(p_1)) dp_1$$

So it suffices that

$$0 \geq \frac{d}{dp_1} g_1(p_1, a(p_1)) = \frac{\partial}{\partial p_1} g_1(p_1, a(p_1)) + \frac{\partial}{\partial p_2} g_1(p_1, a(p_1)) * a'(p_1)$$

The last two claims follow from the results of the previous section for IC slack.

□

If the players finds themseleves at a lower position, then the moral hazard constraints eventually must bind on the upward course of optimal adjustments:

Proposition 15 (Upward adjustment)**.** *If the initial position p is weakly above the competitive level and strictly below the Pareto frontier, then no optimal feasible outcome is achievable. Upward adjustments are eventually constrained by IC.*

Proof. Consider the N-optimal path beginning with player i. I want to show that some player prefers to forgo some adjustment. If there are only a finite number of adjustments, and they are monotone upward, then this is true at least for the last adjustment. Suppose instead there are a countable number of adjustments.

The N-optimal path converges to $P^O(N)$, and as it does so, $\nabla_{p_i} g \to k_i T$, where $T \perp N$ and k_i is a constant. For example, for $N_1 = N_2$ the continuuing adjustments get closer and closer to straight transfers between the two players. However, as there is delay between adjustments, each player requires that her transfer is returned multiplied by $\frac{r+\alpha}{\alpha}$ beginning at the other player's next adjustment. As $\nabla_{p_1} g - k \nabla_{p_2} g \to 0$ this is impossible. □

6 Discussion and Conclusion

This paper relies on two main assumptions: First, opportunities to adjust are randomly distributed according to a Poisson point process. Second, they are privately observed; each player directly observes only her own opportunities. Lagunoff and Matsui (1997) consider this first assumption, which they call "Poisson revisions." They study the dynamic game that arises specifically from a coordination stage game, for which they prove an anti-folk result. Calvo (1983) applies this assumption in a setting of individual decision making: price-setting by perfectly competitive firms. He studies the macroeconomic implications. For Lagunoff and Matsui, the opportunities are publicly observed, while for Calvo, it does not matter whether they are public or private. Hauser and Hopenhayn (2008) consider an assumption similar to the second one here: two players have random, private opportunities to provide a favor to the other. These are not opportunities to adjust one's position within a stage game and there is no intrinsic state variable in the game they consider, but similar issues arise as here of moral hazard regarding the choice whether or not to take one's opportunities.

While not novel, there is a way in which the first assumption, that one's opportunities to move are random, is unrealistic, for example in the setting of price adjustment. Before turning to this issue, I want to discuss the standard models of price adjustment. In continuous time models, firms may adjust their prices instantly and incessantly. One might protest on both counts, but such models have analytical appeal and may serve as a fair approximation when adjustment speed is not a strategic issue. For example, Scherer (1980) reports that the big tobacco companies typically matched each other's price changes within the day. In other industries, price responses seem to take an economically significant amount of time, for example, weeks for the makers of breakfast cereals. Even for airlines, which typically respond within a few days, it seems that the leading firm may bear a significant cost during the short period where it is priced above its rivals. The standard supergame model, where the stage game is discretely repeated, implies that players cannot instantly adjust. However, in the context of general price setting, the supergame model gets something wrong that was right in the continuous time model: While firms cannot adjust at every time, it seems that they ought to be able to adjust at any time. It is not clear what a period represents in dynamic price setting. Consider instead a pair of habitual criminals who are repeatedly arrested and must decide whether or not to confess. The prisoner's dilemma stage game may be literally repeated. Similarly in a market where firms are constrained to adjust prices only on January 1 of each year, the Bertrand pricing stage game may be literally repeated — but markets with such a definitive restriction on price adjustment are the exception. Instead, each "period" seems intended to capture a restriction on how quickly the players may adjust. In this case, one might think that opportunities to adjust ought to be asynchronous, random and privately observed, leading to the type of uncertainty about the rival's temporary non-adjustment that drives the model here.

The model here is approximately the simplest in which the issue of moral hazard regarding adjustment arises. One might view it as a reduced form of a more realistic model where opportunities are endogenous subject to restrictions of bounded rationality and costly information acquisition. I am not aware of satisfying and tractable models of these restrictions. Rotemberg (1987) suggests that there may be "somewhat random delay" between a firm's price adjustments due to costs of determining the optimal price coupled with some random, private observation of associated information. Many authors have suggested that a significant portion of "menu costs" are associated with such decision costs, rather than costs of executing a decided price change. I imagine that a realistic model of such price setting would result in a range of interesting behavior, including the moral hazard issue that arises in the reduced model here.[13] The existing model is still difficult to analyze in a strategic setting like oligopoly.

As discussed in the introduction, the issue of moral hazard surrounding adjustments, and the resulting strategic difficulty in attaining a collusive price from a lower price, seems to have been informally suggested in the older literature on tacit collusion. This paper aims

[13] Taken as a reduced form, this model can be interpreted as assuming that the firms have failed to collude on the times at which they will adjust. This seems reasonable in some instances of tacit collusion.

to provide the first game theoretic foundation and formal analysis. Note the restrictions on equilibrium collusion here are distinct from those in the standard supergame model and models of imperfect monitoring of prices following Green and Porter (1984). I say that in those models the restrictions are on “maintaining” a collusive price, while here the restriction is on attaining it.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383 – 398.

Chamberlin, E. H. (1929): "Duopoly: Value Where Sellers are Few," *The Quarterly Journal of Economics*, 44, 63–100.

Dutta, P. K. (1995): "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, 1 – 32.

Fudenberg, D. and D. K. Levine (1994): "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62, 103 – 135.

Green, E. J. and R. H. Porter (1984): "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87–100.

Hanson, F. (2007): *Applied Stochastic Processes and Control for Jump-Diffusions*, Philadelphia: Society for Industrial and Applied Mathematics.

Hauser, C. and H. Hopenhayn (2008): "Trading Favors: Optimal Exchange and Forgiveness," Working paper.

Kalesnik, V. (2005): "Continuous Time Partnerships with Discrete Events," Job market paper.

Lagunoff, R. and A. Matsui (1997): "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 65, 1467–1477.

Mailath, G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford Oxfordshire: Oxford University Press.

Øksendal, B. and A. Sulem (2007): *Applied Stochastic Control of Jump Diffusions, 2nd ed.*, Berlin: Springer.

Rotemberg, J. (1987): "The New Keynesian Microfoundations," in *NBER Macroeconomics Annual 1987, Volume 2*.

Salop, S. C. (1986): "Practices that (Credibly) Facilitate Oligopoly Coordination," in *New Developments in the Analysis of Market Structure*, ed. by J. Stiglitz and G. Mathewson, chap. 9.

Sannikov, Y. (2007): "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica*, 75, 1285–1329.

Scherer, F. (1980): *Industrial Market Structure and Economic Performance*.

Wen, Q. (2002): "Repeated Games with Asynchronous Moves," .

Yoon, K. (2001): "A Folk Theorem for Asynchronously Repeated Games," *Econometrica*, 69, 191–200.