Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 15-008

"When and How the Punishment Must Fit the Crime"

by

George J. Mailath, Volker Nocke, and Lucy White

http://ssrn.com/abstract=2565868

When and How the Punishment Must Fit the Crime*

George J. Mailath† Volker Nocke‡ Lucy White§

February 15, 2015

Abstract

In repeated normal-form (simultaneous-move) games, *simple penal codes* (Abreu, 1986, 1988) permit an elegant characterization of the set of subgame-perfect outcomes. We show that the logic of simple penal codes fails in repeated *extensive*-form games. By means of examples, we identify two types of settings in which a subgame-perfect outcome may be supported *only* by a profile with the property that the continuation play after a deviation is tailored not only to the identity of the deviator, but also to the nature of the deviation.

Keywords: Simple Penal Code, Subgame Perfect Equilibrium, Repeated Extensive Game, Optimal Punishment.

***JEL* Classification Codes:** C70, C72, C73.

*Mailath thanks the National Science Foundation for research support (grants SES-0350969). Nocke is grateful for financial support by the European Research Council (grant no. 313623). This is a major revision of "When the Punishment Must Fit the Crime: Remarks on the Failure of Simple Penal Codes in Extensive-Form Games," originally circulated in 2004. We thank Larry Samuelson for helpful discussions.

†Department of Economics and PIER, University of Pennsylvania, and Research School of Economics, Australian National University.

‡University of Mannheim, CESifo, and CEPR.

§Harvard Business School and CEPR.

My object all sublime
I shall achieve in time
To let the punishment fit the crime,
The punishment fit the crime;
And make each prisoner repent
Unwillingly represent
A source of innocent merriment,
Of innocent merriment!

W. S. Gilbert (1885), *The Mikado*

1 Introduction

Many popular applications of game theory are naturally modeled as simultaneous-move stage games, such as the Cournot oligopoly model of collusion. But other interesting applications have a dynamic structure whose stage game interactions are more naturally represented by nontrivial extensive-form stage games. Examples include the interaction between government and the private sector in the time-inconsistency literature, between upstream and downstream firms in the vertical relations literature, between bidders in open-outcry auctions, between firms which first invest or choose standards and then compete, between proposer and responders in bargaining games, and between principal and agent in contracting games. Such applied models often feature a socially desirable (or 'cooperative') outcome which is inconsistent with equilibrium in a one-shot game because of opportunistic actions available to players, but which might be attainable in a repeated game. Economists are often interested in exploring how changes in institutions and/or repetition of the game might allow the more desirable outcomes to be supported as equilibria (see, e.g., the seminal work of Friedman (1971)). But with a few exceptions,[1] such analysis has been largely confined to repeated *normal-form* games, where all players are modeled as moving simultaneously in the stage game.

When players are sufficiently patient, one may ignore the detailed dynamic structure within the stage-game interaction and apply standard folk theorem arguments (Wen, 2002). The reason is that even a small difference in continuation values will dominate any short-term gain for patient players, so deviations are relatively easy to deter. However, applied theory is often concerned with the impact of a change in institutions or the environment on the set of equilibrium outcomes, and such an analysis is meaningful only for *impatient* players. In this paper we show that for impatient players, the dynamic structure of the stage game can be important and should not be neglected.

[1] Repeated extensive-form games have been analyzed in the relational contracting literature (Levin, 2003), in the policy games literature (Athey, Atkeson, and Kehoe, 2005), and in the vertical relations literature (Nocke and White, 2007).

The literature has so far directed surprisingly little attention towards the study of repeated extensive-form games with impatient players. For applications of repeated simultaneous-move games, the techniques developed by Abreu (1986, 1988) are central. Abreu (1988) shows that any pure-strategy subgame-perfect equilibrium outcome can be supported by a set of simple punishment strategies called *simple penal codes*: If a player, i say, deviates from the proposed equilibrium play in a given period, in the next period, players switch to player i's worst equilibrium play (called i's *optimal penal code*). A similar rule applies to any player who deviates from play during an optimal penal code. In other words, the continuation play after a deviation by a player is *independent* of the nature of the deviation, depending only on the identity of the deviator. This result vastly simplifies the task of finding the set of equilibria that can be supported in repeated simultaneous-move games: one needs only to characterize a worst equilibrium for each player, and from there one can proceed to fill in all other equilibria that can be supported by using these punishment strategies. Abreu's result has been correspondingly important for applications employing normal-form stage games.

But what about applications that involve extensive-form stage games? In this paper, we show that a similar simplification is not available for characterizing the set of subgame perfect equilibrium outcomes when the stage game has a nontrivial dynamic structure. We begin by discussing the appropriate restrictions that simple penal codes should satisfy in this case. We then present two settings in which simple penal codes can fail to support equilibria that are supportable with more complicated punishments for a given discount factor. The forces driving optimal punishment in these settings are intuitive and natural.

Consider a deviation by some player in a repeated simultaneous-move game. Since the stage game is a simultaneous-move game, the other players can respond only in the next period. Abreu's (1988) results rely on the observation that, in repeated simultaneous-move games, every subgame is strategically equivalent to the original repeated game.[2] Hence, the worst punishment is simply the worst subgame perfect equilibrium of the original game. Consider now a deviation in a repeated extensive-form game. In contrast to normal form games, the other players may be able to respond not only in the next period but also within the same period. Moreover, the deviation may lead to a subgame that is not strategically equivalent to the original game. Consequently, the appropriate notion of simple penal code for a repeated extensive-form game is not obvious. Nonetheless, for any history that ends at the end of a period, the associated subgame *is* strategically equivalent to the original repeated game. This suggests that any definition of a simple penal code should have the feature that after a player deviates from the candidate equilibrium, subsequent play *beginning in the next period* should be independent of the precise nature of the deviation.

Our two sets of examples show that optimal punishments in repeated extensive-form games need not have this property: even though it is feasible to play the same outcome path

[2]This observation is a direct implication of the property that every subgame's initial node is at the beginning of some period.

after two different deviations starting in the next period, sustaining a candidate equilibrium outcome may require that different outcome paths are played.

Say that an action for a player is *myopically suboptimal* if, given the specified behavior for the other players, that action is not optimal for that player when payoffs from future periods are ignored. The logic of simple penal codes fails in repeated extensive-form games because some equilibria require the use of within-period myopically-suboptimal "punishments" to ensure that deviations are not profitable. In contrast, in repeated normal-form games, the sequential rationality of "within-period" punishments is not an issue. In each of our examples, the within-period punishment is myopically suboptimal for the potential punisher(s), but is required to sustain the desired equilibrium.

In the first set of examples, the need for different continuations arises because the interests of the deviator and the potential punisher are aligned, though imperfectly. Many important settings have this structure – for example, policy games, models of relational contracting, and models where players engage in repeated investment or production. In such settings, using continuation play to reward a player for carrying out myopically suboptimal within-period punishment of the earlier deviator may necessarily also reward the deviator himself. Consequently, the rewards for the punisher (and hence the outcome path following the deviation) may have to be fine-tuned to the particular deviation chosen by the deviator. Within-period punishment is valuable after some, but not all, deviations, depending on how effective and costly is within-period punishment after a particular deviation. We first analyze a simple example to illustrate the difficulty and then show how the same logic affects a repeated game of bilateral investment with hold-up in the spirit of Klein, Crawford, and Alchian (1978) and Grossman and Hart (1986).

In the second set of examples, we highlight a contrasting problem, which arises because players moving after a deviator are required to coordinate to inflict effective within-period punishment. In contrast to the first set, where complications arise when interests are aligned, here, difficulties occur when the interests of the players required to inflict punishment are in conflict. The simplest case has three players, and sustaining the desired equilibrium requires that the two later movers both inflict within-period punishment on the first mover in the event of a deviation. Think, for instance, of a situation where a deviation is profitable if and only if one other player 'accepts' it. Such a structure is common in many important settings. Examples include colluding upstream firms (Nocke and White, 2007), attempted expropriation by a sovereign power (where citizen groups can cooperate to successfully resist expropriation, as in Weingast (1995, 1997)), bargaining over a series of proposals where a majority vote is sufficient for acceptance (Baron and Ferejohn, 1989), or attempted entry when an entrant requires more than one customer to break even (as in Rasmusen, Ramseyer, and Wiley (1991) and Segal and Whinston (2000)). The conflict of interest between the two potentially enforcing players in these cases means that it is not possible to provide the maximum "reward" to both the punishers in any given equilibrium. Intuitively, the continuation equilibrium chosen (and hence the "reward" which each punisher receives) must depend on the degree of sacrifice which the punisher makes

in inflicting the myopically suboptimal punishment. In other words, the continuation play, starting in the period after the deviation by the first mover, is not simple, but rather depends on *which* deviation the first mover has chosen among those that are available to him. Which of the two later moving players receives more of the "carrot" in future play optimally depends on for which of them it was more costly to apply the "stick". Again, we first analyze a simple example and then turn to the classic "naked exclusion game" of Rasmusen, Ramseyer, and Wiley (1991) and Segal and Whinston (2000) to illustrate the operation of optimal punishments in that context.

Apart from Rubinstein and Wolinsky (1995), Sorin (1995), Wen (2002), and Mailath and Samuelson (2006), the repeated-game literature has focused on repeated normal-form games, ignoring the dynamic structure within the stage game.[3] Rubinstein and Wolinsky (1995) present some examples illustrating the difference between the set of subgame-perfect equilibrium payoffs of repeated extensive and normal-form games for patient players when the standard full dimensionality condition of Fudenberg and Maskin (1986) does not hold. Sorin (1995) discusses the implications of the different information that players have available across periods in repeated extensive, rather than normal, form games. Wen (2002) extends the arguments of Abreu (1988) to prove a folk theorem for repeated sequential-move games under a weaker condition than Fudenberg and Maskin's (1986) full dimensionality condition. Finally, Mailath and Samuelson (2006, Section 9.6) prove a folk theorem for repeated extensive form games via an extension of the tools of Abreu, Pearce, and Stacchetti (1990). None of these papers are concerned with penal codes or with characterizing the set of subgame-perfect equilibrium payoffs with impatient players.

2 The Punishment Should Fit The Crime

> "Is it her fault or mine?
> The tempter or the tempted - who sins the most?"
>
> William Shakespeare *Measure for Measure*, Act 2, Scene 2.

In this section, we highlight the value of tailoring the punishment to fit the deviation in games where the potential deviator and punisher have a commonality of interest. Thus our games have a strong coordination flavor. The most effective punishments can be complicated in such settings, because it is difficult to both punish the deviator and reward the punishing player for applying the costly punishment.

We first present a simple stylized example to illustrate the issues that arise. In the following subsection, we will present a more interesting application to a game of repeated

[3] Despite the name, the literature on *asynchronously repeated games* (for example, Lagunoff and Matsui (1997)) does not study repeated games. In each period only some players can choose (adjust) an action, with the other players' actions fixed from earlier periods. The canonical example is the asynchronous move prisoners' dilemma, where player 1 chooses in odd periods, and player 2 in even periods.

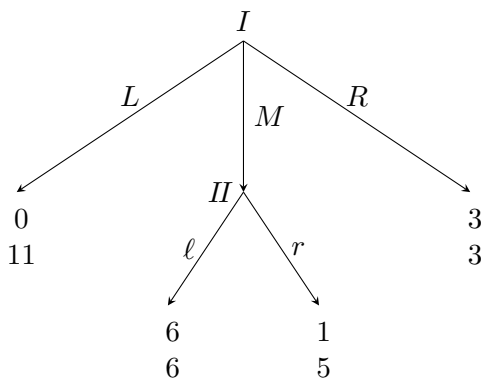

(a) The extensive-form game Γ_1.

	A	B	C
A	$6,6$	$0,0$	$0,0$
B	$0,0$	$4,4$	$0,0$
C	$0,0$	$0,0$	$2,2$

(b) The coordination game.

Figure 1: The two period game.

investment and hold-up.

2.1 A Simple Example

We begin with the extensive-form game Γ_1, presented in Figure 1a. In this stage game, the unique backward induction equilibrium is for player I to play M, and player II to play ℓ (after M). Player II of course, prefers that player I play L.[4]

Before we describe the repeated extensive-form game, we first analyze a simpler two-period game. In the first period, Γ_1 is played, while in the second period, the coordination game of Figure 1b is played. Each player's payoff is the sum of payoffs from the two periods. In the second period, the payoffs of the two players are, by construction, identical. It is impossible to punish (or reward) one player without simultaneously doing the same to the other.

We are interested in equilibria that support L as a choice by player I in the first period. The payoffs in this example have been chosen so that the variation in second-period payoffs alone is insufficient to deter player I from playing M when player II chooses ℓ, her myopic best reply: Player I's first-period incentive to deviate by playing M is then 6, while the largest punishment the second period can impose is the profile CC, with associated loss of payoff of 4.

[4]The game Γ_1 has an interpretation as an entry game between a potential entrant (player I) and an incumbent (player II), where the entrant can decide to stay out (play L), enter with product line M (play M) or enter with product line R (play R). (One can also interpret M as small-scale entry and R as large-scale entry.) Following entry with product M, the incumbent can choose to acquiesce (play ℓ) or fight (play r). We assume here for simplicity that fighting is possible only when product M is chosen, but one can also allow fighting after choice of R with, for example, payoffs (-1,1) without changing any of the conclusions reached in the analysis below.

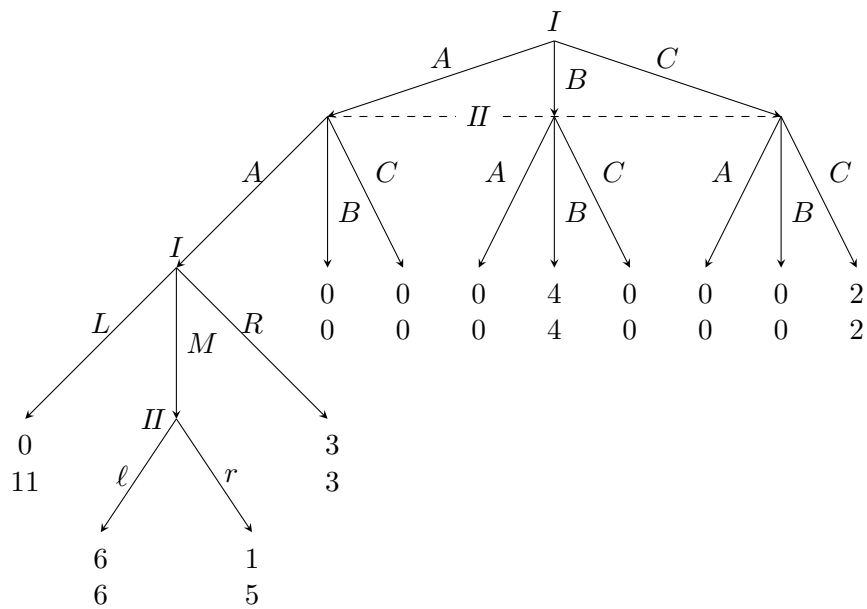

Figure 2: The game Γ_1^*.

On the other hand, if player II can be induced to play r, the benefit to player I of M is drastically reduced (from 6 to 1). The trick is in providing appropriate incentives to prevent II from myopically optimizing, i.e., in specifying a higher continuation payoff after Mr than after $M\ell$. At the same time, the continuation play after Mr must not ignore I's original deviation. This motivates the following specification of second period play: play AA after L, BB after Mr, and CC after $M\ell$ and r. It is straightforward to check that this specification supports Lr as equilibrium first period choices.

In the profile described in the previous paragraph, different continuation equilibria are specified after I's deviation to M and to R. We have already seen that the second period play BB after Mr is needed to make II's choice of r optimal. At the same time, a play of BB after R does not provide a sufficient disincentive for I.

The play of L in the first period cannot therefore be sustained by a "simple" penal code in which continuation play is independent of the particular deviation chosen by player I. As we have seen, the play of L is sustained by a more complex strategy profile which employs different punishments after different deviations.

Consider now the extensive-form game, Γ_1^*, presented in Figure 2, where the two players first simultaneously choose A, B or C, and Γ_1 follows a simultaneous choice of A. Since Γ_1 has a unique subgame-perfect equilibrium with payoffs $(6, 6)$, the pure strategy subgame perfect equilibrium payoffs of Γ_1^* coincide with the pure-strategy Nash equilibrium payoffs

	$A\ell$	Ar	B	C
AL	$1, 11$	$1, 11$	$0, 0$	$0, 0$
AM	$6, 6$	$1, 5$	$0, 0$	$0, 0$
AR	$3, 3$	$3, 3$	$0, 0$	$0, 0$
B	$0, 0$	$0, 0$	$4, 4$	$0, 0$
C	$0, 0$	$0, 0$	$0, 0$	$2, 2$

Figure 3: The (reduced) normal form of Γ_1^*.

of the game of Figure 1b.

The game Γ_1^* is played in two periods. The discussion above shows that there is a subgame-perfect equilibrium of the repeated game in which in the first period, both players choose A, and then player I chooses L. (The only issue we have not addressed is a deviation by I or II to B or C in the first period. Player II clearly cannot benefit from such a deviation. For player I, the period-1 payoff from this deviation is 0, the same as from Mr in Γ_1, and the second-period play is the same as well, and so the deviation is not profitable.)

It is also an implication of the above discussion that this outcome cannot be achieved in any equilibrium with the property that play in the second period is independent of the nature of I's deviation in the first period.

To conclude our discussion of this example, it is useful to compare our analysis with that of an analysis of the repeated normal form of Γ_1^*. The normal form is given in Figure 3. Treating the simultaneous-move normal form of Figure 3 as the stage game, the profile in which $(AL,\ Ar)$ is played in the first period and $(AR,\ A\ell)$ in the second, with *any* deviation by player 1 resulting in CC in the second period, is a subgame-perfect equilibrium of the repeated game. However, this profile is a subgame-perfect equilibrium only because the simultaneity of moves means that there is no subgame beginning with II's choice between ℓ and r, and so subgame perfection does not require that choice to be optimal.

2.2 Application: Bilateral Investment and Hold-Up

We now consider a simplified version of the classic hold-up model with bilateral investment in the spirit of Klein, Crawford, and Alchian (1978) and Grossman and Hart (1986). Here, the more standard Nash bargaining stage is replaced by a simple non-cooperative bargaining model where one player makes a take-it-or-leave-it offer to the other player. The stage game has three stages:

Stage 1 (Investment) Both players simultaneously decide whether or not to make a

relation-specific investment. The cost of the investment is $c > 0$. Unless both players have made the investment, the game ends, with player i's payoff being $-c$ if $i \in \{1, 2\}$ had invested and zero otherwise. If both players have made the investment, the game proceeds to the next stage.

Stage 2 (Offer) Player 1 makes a take-it-or-leave-it offer $x \geq 0$ to player 2.

Stage 3 (Acceptance) Player 2 decides whether to accept or reject the offer. If the offer is accepted, payoffs are $(B - x - c, x - c)$; otherwise payoffs are $(-c, -c)$, where B is the investment revenue.

The interesting case arises when investment is efficient: $B > 2c$. Despite the efficiency of investment, it is well known that ex post bargaining over terms leads to inefficiently low investment – here, the one-shot game has a unique subgame-perfect equilibrium with no investment by either player, resulting in payoffs $(0, 0)$.

2.2.1 The Infinitely Repeated Game

Now consider the infinite repetition of the stage game just described, with δ denoting the common discount factor. We characterize the best equilibrium for player 2, and show that this equilibrium depends on whether behavior is restricted to simple penal codes. In either case, this best-for-player-2 equilibrium will involve in each period either (i) no investment by either player (as in the static equilibrium), resulting in payoffs $(0, 0)$, or (ii) investment by both parties, with player 1 making a subsequent offer of $\hat{x}$, which is accepted by player 2, resulting in payoffs $(B - c - \hat{x}, \hat{x} - c)$. As each player can ensure himself a payoff of 0 by not investing, if an investment equilibrium exists, player 1's offer $\hat{x}$ must satisfy $B - c \geq \hat{x} \geq c$.

2.2.2 Optimal Punishments

We first analyze the features of optimal punishments without the restriction to simple penal codes. Since infinite repetition of the static equilibrium (i.e., the subgame-perfect equilibrium of the stage game), resulting in no investment, is trivially a subgame-perfect equilibrium of the repeated game, we turn to the incentive constraints for any equilibrium involving investment by both parties and an accepted offer of $\hat{x}$ in every period.

There are two possible ways in which deviations by player 1 to offers less generous than $\hat{x}$ can be punished. First, and standardly, deviations could be deterred by the threat of reversion to the static equilibrium in all future periods. Such reversion may well be sufficient to deter a small deviation by player 1 (to an offer x not much smaller than $\hat{x}$) because his short-run benefit from such a deviation is small. But such a punishment scheme may not deter a large deviation (to a small offer x) since the short-term benefit of having a very low offer accepted may be too large. To deter a large deviation, therefore, it may be necessary to use the second way of deterring a deviation, which is to have player 2 reject

the low offer (leaving player 1 with a payoff of $-c$). For player 2 to be willing to carry out such a within-period punishment, however, requires that he is subsequently rewarded by a sufficiently attractive continuation play. But for an impatient player 2 such a reward may not be large enough to induce rejection of an offer x not much smaller than $\hat{x}$. Thus there may be a need to fine-tune the continuation play to the size of the deviating offer.

These considerations suggest the following features for an optimal punishment scheme: There is a cutoff $\tilde{x}$ such that any deviation to an offer $x \in [\tilde{x}, \hat{x})$ is accepted by player 2 and the continuation play is reversion to the static equilibrium, while any offer $x \in [0, \tilde{x})$ is rejected by player 2 with the continuation play reverting to an offer of $\hat{x}$ in every future period (acceptance by player 2 is followed by reversion to the static equilibrium).

Player 1 will find it unprofitable to make a deviant offer $x \in [\tilde{x}, \hat{x})$ only if the smallest deviant offer $\tilde{x}$ (inducing the same continuation play) is unprofitable, i.e.,

$$\frac{B - c - \hat{x}}{1 - \delta} \geq (B - c - \tilde{x}),$$

which implies the following lower bound on $\tilde{x}$, the lowest deviant offer that player 2 can accept:

$$\tilde{x} \geq \frac{\hat{x} - \delta(B - c)}{1 - \delta}. \tag{1}$$

Player 2 will be willing to reject a deviant offer $x \in [0, \tilde{x})$ only if

$$-(1 - \delta)c + \delta(\hat{x} - c) \geq (1 - \delta)(\tilde{x} - c),$$

which implies an upper bound on $\tilde{x}$, the largest deviant offer that player 2 is willing to reject:

$$\frac{\delta(\hat{x} - c)}{1 - \delta} \geq \tilde{x}. \tag{2}$$

Combining inequalities (1) and (2), we obtain as a necessary condition, an upper bound on $\hat{x}$, the equilibrium offer in the best-for-player-2 equilibrium,

$$\hat{x} \leq \frac{\delta}{1 - \delta}(B - 2c). \tag{3}$$

In equilibrium, players' individual rationality constraints must also be satisfied, which requires that each player's payoff be nonnegative. So we also must have

$$\hat{x} \geq c \quad \text{and} \quad B - \hat{x} \geq c. \tag{4}$$

Thus, if $(\tilde{x}, \hat{x})$ describes the best-for-player-2 equilibrium, $\hat{x}$ is the largest offer satisfying (3) and (4). Since

$$\frac{\delta}{1 - \delta}(B - 2c) \geq c \iff \delta \geq \frac{c}{B - c} \equiv \delta',$$

if $\delta < \delta'$, there is no such $\hat{x}$.[5] If $\delta \geq \delta'$, then

$$\hat{x} = \min\left\{\frac{\delta}{1-\delta}(B-2c),\ B-c\right\}.$$

Let δ'' be the smallest value of δ at which $\hat{x} = B - c$, i.e.,

$$\delta'' \equiv \frac{B-c}{2B-3c}.$$

We are now in a position to describe the best equilibria for player 2 (the verification is straightforward):[6]

Proposition 1 *Suppose $\delta \in (\delta', \delta'')$. All best equilibria for player 2 have the following structure: Set*

$$\hat{x} = \frac{\delta}{1-\delta}(B-2c)$$

and fix an $\tilde{x} < \hat{x}$ satisfying

$$\tilde{x} \in \left[\frac{\hat{x} - \delta(B-c)}{1-\delta}, \frac{\delta(\hat{x}-c)}{1-\delta}\right].$$

The equilibria have two phases, "invest" and "don't invest", and begin in the "invest" phase. In the "invest" phase,

1. *on the path of play, both players invest and player 1 offers $\hat{x}$ to player 2,*
2. *player 2 accepts all offers $x \geq \tilde{x}$, and*
3. *player 2 rejects all offers $x < \tilde{x}$.*

Play stays in the "invest" phase as long as both players invest, the last offer satisfies $x \notin [\tilde{x}, \hat{x})$, which player 2 accepts if $x \geq \hat{x}$ and rejects if $x < \tilde{x}$. Otherwise, play switches to the "don't invest" phase, in which the static no-investment equilibrium is played in every period.

The equilibria described in Proposition 1 do not have a simple penal code structure, since deviations by player 1 to relatively generous offers $x \geq \tilde{x}$ lead to different continuation play than deviations to less generous offers $x < \tilde{x}$.

[5] Indeed, if $\delta < \delta'$, then there is no stationary equilibrium with investment, since even an offer $\hat{x} = \delta(B-c)$ (the largest offer consistent with $\tilde{x} \leq 0$) violates player 2's individual rationality.

[6] In addition to the stationary equilibria described in Proposition 1, there are also equilibria where the cutoff $\tilde{x}$ depends upon history. This history dependence off the equilibrium path does not affect equilibrium payoffs.

However, for $\delta \geq \delta''$, we can take $\tilde{x} = \hat{x} = B - c$, since for large δ,

$$B - c \leq \frac{\delta(\hat{x} - c)}{1 - \delta}.$$

(It is straightforward to verify that the relevant incentive constraints hold.) In this case, the simple penal code in which player 2 rejects all offers less than $B - c$ (expecting a continuation value of $(B - c)/(1 - \delta)$) supports the equilibrium.

2.2.3 Simple Penal Code for $\delta \in (\delta', \delta'')$

We have just shown that in this game, a simple penal code can be optimal as long as $\delta \geq \delta''$, and the only equilibrium involves no investment for $\delta \leq \delta'$. We now investigate the efficacy of simple penal codes for intermediate values of $\delta \in (\delta', \delta'')$. Let $x^\dagger$ be the largest offer that player 2 can receive in any equilibrium supported by a simple penal code. There are two possibilities for a simple penal code in this setting: either any deviating offer results in no investment in the future, or any deviating offer is rejected, in which case rejection leads to $x^\dagger$ in the future. (Failure by player 2 to reject such a deviating offer would in this case lead to the no investment equilibrium.)

First consider punishing deviations by infinite reversion to the static (no investment) equilibrium. Since this static equilibrium is the worst not only for player 1 but also for player 2, player 2 will accept any (strictly positive) deviant offer. Hence, player 1 will have no incentive to make a deviant offer $x < x^\dagger$ only if he has no incentive to make an arbitrarily small deviant offer, i.e.,

$$\frac{B - c - x^\dagger}{1 - \delta} \geq B - c,$$

which implies the following upper bound on $x^\dagger$:

$$x^\dagger \leq \delta(B - c). \tag{5}$$

It is straightforward to check that the profile in which both players invest, player 1 always makes the offer $\delta(B - c)$, which is accepted (as is any lower off-the-equilibrium path offer), and any deviation by player 1 to an inferior offer results in future no investment is an equilibrium for $\delta \in (\delta', \delta'')$.

Now consider the other possible simple penal code: inducing player 2 to reject all deviating offers, followed by a return to an investment equilibrium if and only if he does so. Player 2 will be prepared to reject any offer $x \in [0, x^\dagger)$ only if

$$-(1 - \delta)c + \delta(x^\dagger - c) \geq (1 - \delta)(x^\dagger - c),$$

which implies the upper bound on offers that player 2 will be willing to reject:

$$x^\dagger \leq \frac{\delta c}{2\delta - 1}. \tag{6}$$

Individual rationality for player 1 requires $B - c \geq x^{\dagger}$, which with (6) implies $\delta \geq \delta''$, contradicting our assumption on δ. Hence, if $\delta \in (\delta', \delta'')$, the best simple penal code involves the play of the no investment equilibrium in all periods following a deviant offer by player 1. The largest offer that player 2 can receive in any equilibrium supported by a simple penal code is $\delta(B - c)$, which is strictly smaller than $\hat{x}$.

2.2.4 Discussion

Whenever it improves on simple penal codes, the optimal punishment scheme has an interesting feature. If the offerer deviates from the expected price by shading just slightly, the responder will accept (shrug it off) but will no longer invest in the relationship (one can also think about the parties as walking away from the relationship and as receiving outside options of zero on the spot market). Whereas if the offerer deviates from the expected price by making a much lower offer, then instead the – perhaps insulting – offer is rejected, but in the expectation that next period the investment relationship will be re-established on better terms. The structure of the punishment might at first glance be considered counter-intuitive, since one might think that the relationship would be more likely to continue if shading of the price is only slight and not large. But if the relationship is to be continued, costs must be inflicted within the period by refusing to deal, otherwise shading will always occur; whereas if the relationship is to be sacrificed (and the pair are to return to the no-investment or spot market equilibrium), only small amounts of shading can be prevented. Moreover, the ability to inflict large costs by failing to make a mutually beneficial trade in a period of deviation (cutting off ones nose to spite ones face) can help sustain the relationship beyond what would otherwise be possible. So the optimal punishment scheme yields a prediction about the pattern of punishment strategies in relationships with hold-up: other things being equal, small deviations result in walking away whereas large deviations result in a costly standoff and then resumption of trade. By contrast, a simple penal code would involve the responder taking the same action (in particular, walk away) independently how low the offer he receives is.

3 The Reward Should Fit the Temptation

> "You oughtn't to yield to temptation."
> "Well somebody must, or the thing becomes absurd."
>
> Anthony Hope, *The Dolly Dialogues*.

We have just highlighted how the value of deviation-dependent punishments can arise from the *commonality* of interest between the punisher and the punished. We now examine how the value of deviation-dependent punishments arises for a very different reason – because there is a *conflict* of interest between two players who are not supposed to acquiesce

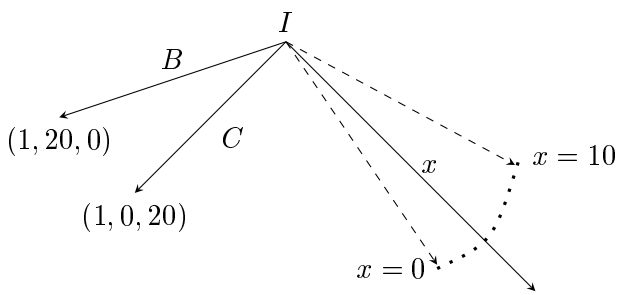

		III: Yes	III: No
II	Yes	$10,\ x,\ 10-x$	$10,\ x,\ 0$
	No	$10,\ 0,\ 10-x$	$0,\ 0,\ 0$

Figure 4: The extensive form game, Γ_2. The choice x for player I ranges over the nonnegative integers, $\{0, 1, \ldots, 10\}$.

to a deviation by a third player. As in the first set of examples, we begin with a simple stylized example to illustrate the main force at work, which is the need to choose the continuation play in order that the reward to each player must be adapted to the sacrifice that they made in inflicting within-period punishment. We then provide an application to a more complex game of greater applied interest – in this case, a repeated version of the "Naked Exclusion" game analyzed by Rasmusen, Ramseyer, and Wiley (1991) and Segal and Whinston (2000).

3.1 A Simple Example

The stage game for our second simple example is the extensive form Γ_2, presented in Figure 4. We interpret the choice of $x \in \{0, 1, \ldots, 10\}$ by player I as a bribe to player *II* (with $10-x$ the bribe to player *III*).[7] If player I chooses either one of the "cooperative" actions B or C, then the stage game ends. If player I chooses instead to offer a bribe of 10 to players *II* and *III* (with x representing the split), players *II* and *III* then simultaneously decide whether to accept or reject the bribe. The stage game then ends, and all actions become common knowledge. The game has many subgame-perfect equilibria, but they all share some common features: player I attempts to bribe the other players rather than behave cooperatively, and both players *II* and *III* accept any positive bribe offered. Moreover, the set of subgame-perfect equilibrium payoffs is given by

$$\{(10, x, 10-x) : x \in \{0, 1, \ldots, 10\}\}.$$

[7] The extensive form Γ_2 has a natural interpretation as a bargaining game where player I has a pie of 10 to be split between himself and two others, with decisions being taken by majority voting (to make this interpretation literal, simply replace the payoff of 1 for player I after actions A and B with a payoff of 0).

Note that the payoffs of players *II* and *III* are negatively related across the equilibria of the stage game.

We assume randomizations between B and C are observable to order to maintain the pure strategy nature of the analysis (equivalently, we could assume there is an observable "sharing" fraction y that I uses to allocate the surplus of 20 between *II* and *III*).

We are interested in the possibility of using these multiple equilibria to construct an equilibrium of the once-repeated game where player I behaves cooperatively in the first period by playing a (possibly degenerate) probability distribution over the pure cooperative actions B and C.

We begin by arguing that it is impossible to support this cooperative play by I in the first period using continuation play in the second period that is independent of the nature of I's bribe: If player I deviates by attempting the bribe x, then equilibrium requires that players *II* and *III* both reject player I's bribe. For *II* to reject, her continuation payoff from rejection must be at least x. If the continuation play is independent of I's bribe, this implies that *II*'s continuation payoff after rejection is 10 (otherwise *II* would accept a bribe of 10). At the same time, for *III* to reject, his payoff must be at least $10 - x$, again for all x. But this requires that *III*'s continuation payoff after rejection is also 10, which is impossible.

On the other hand, it is easily verified that the following profile is a subgame-perfect equilibrium: In the first period, player I plays cooperatively; if he were to deviate all bribes would be rejected by both players *II* and *III*. In the second period, after cooperative play in the first period, player I bribes at some level x in the second period (any level works). If player I had deviated in the first period, and his bribe x was rejected by both *II* and *III*, player I offers the *same* bribe in the second period, which is accepted by both *II* and *III*. (If I offers another bribe in the second period, both players *II* and *III* accept.) If only one player accepts a bribe, then in the next period, player I offers the bribe that leaves the accepting player with 0. It is irrelevant whether that player accepts the bribe in the second period (the other player of course accepts). Finally, if both players accept the first-period bribe, an arbitrary continuation equilibrium is played (since both players accepting is a simultaneous deviation by *II* and *III*, these payoffs are irrelevant for the purposes of checking for subgame perfection).

We now show that this logic extends if the stage game is infinitely repeated and payoffs are discounted. Consider supporting the play of a (constant or time-dependent) probability distribution over the pure cooperative actions B and C in every period of the infinitely repeated game. In case player I deviates by offering bribes $(x, 10-x)$, the best simple penal code prescribes that the bribes be rejected by both players *II* and *III*, which is then followed by the infinite reversion to the pure cooperative actions B and C, each with probability 1/2; this maximally punishes the deviant player I and optimally rewards players *II* and *III*, given that the continuation play cannot depend on x by definition of a simple penal code. (As in the once-repeated game, accepting a bribe would leave that deviant player with a payoff of 0 in all future periods.) Cooperative play can be sustained by this best

simple penal code if and only if $\delta \geq 1/2$.[8]

The optimal punishment, by contrast, provides a larger reward to the player who was tempted more: in case player I deviates from cooperative play to offer bribes $(x, 10 - x)$, players *II* and *III* reject the bribes, which is then followed by the perpetual play of cooperative actions B and C with probabilities y_B on B and $y_C = 1 - y_B$ on C. This is effective if player *II* does not find it profitable to accept the bribe,

$$x(1-\delta) \leq \delta 20 y_B \iff y_B \geq \frac{x(1-\delta)}{20\delta},$$

and player *III* does not find it profitable to accept the bribe,

$$(10-x)(1-\delta) \leq \delta 20(1-y_B) \iff y_B \leq 1 - \frac{(10-x)(1-\delta)}{20\delta}.$$

These two inequalities are consistent as long as $\delta \geq 1/3$, in which case choosing the midpoint,

$$y_B = \frac{1}{2} + \left(\frac{1-\delta}{\delta}\right)\left(\frac{2x-10}{40}\right),$$

provides sufficient incentive to support the cooperative outcome.

This example shares some features with the three-person alternating-offer bargaining game of Shaked (described in Sutton (1986, p. 721)), though Shaked's example is not a repeated game: it ends once players have agreed on a division of the pie. In Shaked's game, when players are sufficiently patient, any division of the pie between the three bargainers is consistent with subgame perfection. Those equilibria, like here, specify different continuation equilibria *as a function of the identity of the player who is supposed to reject* (in his game, one veto is enough). But our example illustrates that the continuation equilibria may need to be finely tuned *to the original deviation*. In Shaked's example, by contrast, it is always sufficient to promise the entire pie to the rejecter, while both players must be rewarded in our game, since both must reject.

3.2 Application: Naked Exclusion

In this section, we analyse a repeated version of the classic "naked exclusion" model of Rasmusen, Ramseyer, and Wiley (1991) and Segal and Whinston (2000). The game has three long-lived players, an incumbent monopolist (I) and two buyers (B_1 and B_2). Each period, the same incumbent faces a challenge from a different short-lived entrant who is more efficient than the incumbent. For simplicity, we do not model the potential entrants as

[8] The relevant incentive constraint is that neither player *II* nor player *III* prefer to accept a bribe of 10,

$$10(1-\delta) \leq 10\delta \iff \delta \geq \frac{1}{2}.$$

players. In each period, the incumbent can choose whether or not to offer exclusive dealing contracts to the two buyers. Exclusive dealing contracts last one period and the incumbent can make a transfer payment to the buyers to compensate them for signing an exclusive dealing contract. The incumbent's exclusive dealing offers are publicly observable, and buyers make their acceptance/rejection decisions simultaneously. If at least one buyer signs an exclusive dealing contract for the current period, no entry occurs, and the incumbent subsequently charges the monopoly price to each buyer, earning a profit π^m on each. If no buyer has signed an exclusive dealing contract for the current period, the entrant enters the market, resulting in zero profit for the incumbent and an increase in the rents for the buyers equal to S per buyer.

The stage game proceeds as follows:

Stage 1 I chooses between offering no exclusive dealing contracts, N, and a tuple $(x_1, x_2) \in [0, \infty)^2$, where $x_i \geq 0$ is the offered transfer payment to B_i in return for signing an exclusive dealing contract. If I chooses N, then the period ends, and payoffs for I and the two buyers are $(0, S, S)$; otherwise the game proceeds to Stage 2. (Here, S denotes the increase in buyer surplus due to entry.)

Stage 2 Facing public offers (x_1, x_2), B_1 and B_2 simultaneously choose whether to accept the offer ($a_i = 1$) or not ($a_i = 0$). If at least one buyer accepts the offer, payoffs are $(2\pi^m - a_1x_1 - a_2x_2, a_1x_1, a_2x_2)$, where π^m is the monopoly profit that the incumbent can extract from each buyer; if both buyers reject the offer, payoffs are $(0, S, S)$.

We assume that payoffs satisfy $2\pi^m > S > \pi^m > 0$. The first inequality ensures that the incumbent's monopoly profit is sufficiently large that it is worthwhile for him to offer a large enough "bribe" to one buyer to make it a dominant strategy for that buyer to accept the exclusive dealing offer. The second inequality implies that entry is efficient. While the stage game has multiple subgame-perfect equilibria, there is no entry in any of these equilibria and aggregate payoffs are $2\pi^m$. The best equilibrium for buyer B_1 (buyer B_2) involves I offering $(S, 0)$ (resp., $(0, S)$), making it a dominant strategy for that buyer to accept the offer and resulting in incumbent profit of $2\pi^m - S$. These are the worst equilibria for I. The best equilibrium for I is the one in which I offers $(0, 0)$ and both buyers accept the offer, resulting in incumbent profit of $2\pi^m$.

3.2.1 The Infinitely Repeated Game

We now investigate the effects of the infinite repetition of this game. We are interested in determining the conditions under which it is possible to sustain the play of N (no exclusive dealing, and thus entry) in each period as an equilibrium of the game. In this equilibrium, the incumbent's payoff is zero, and each buyer receives S in every period. The common discount factor is denoted $\delta \in (0, 1)$.

We first consider sustaining the perpetual play of N by a simple penal code. Any pair of deviant offers (x_1', x_2') will be accepted by both buyers if $x_1' + x_2' > 2\pi^m$, and rejected by both buyers otherwise. In either case, this is followed by the play of N in all future periods. Note that this simple penal code maximally punishes the deviating I, and maximally rewards the two buyers, given that the rewards cannot be made dependent on the deviant offers. If one of the buyers were to deviate and accept the exclusive dealing offer when $x_1' + x_2' \leq 2\pi^m$, then that deviant buyer would be maximally punished by the play of an equilibrium that gives him a zero payoff in all future periods. (For instance, by the perpetual play of the static equilibrium that gives S to the other buyer.)

Facing deviant offers (x_1', x_2') such that $x_1' + x_2' \leq 2\pi^m$,[9] buyer B_i is willing to reject his offer x_i' if and only if

$$x_i' \leq \frac{S}{1-\delta}.$$

As this incentive constraint has to hold for any $x_i' \leq 2\pi^m$, provided $x_i' + x_{-i}' \leq 2\pi^m$, the perpetual play of N can be sustained by a simple penal code if and only if

$$\delta \geq \frac{2\pi^m - S}{2\pi^m} \equiv \hat{\delta}^{SPC}.$$

We now show that, using an optimal punishment, the perpetual play of N can be sustained for even lower discount factors, namely if and only if

$$\delta \geq \frac{2\pi^m - S}{4\pi^m - S} \equiv \hat{\delta}^{OPC} < \hat{\delta}^{SPC}.$$

The optimal punishment differs from the simple penal code above only in the event in which I's deviant offers (x_1', x_2') are such that $\max(x_1', x_2') \in (S, 2\pi^m)$ and $x_1' + x_2' \leq 2\pi^m$. Such deviant offers are rejected by both buyers, and are followed, from the next period onward, by play of the equilibrium which maximizes the payoff of buyer B_j, $j \equiv \arg\max(x_1', x_2')$, i.e., by providing maximal reward for the buyer who received the larger deviant offer. As we show below, in this continuation equilibrium, buyer B_j gets a per-period payoff of $2\pi^m$, whereas buyer B_{-j} and the incumbent I both get zero. So, I is maximally punished.

Turning to the buyers' incentive constraints following such deviant offers, note first that buyer B_{-j} receives S this period (and zero in all future periods) from rejecting his offer but only $x_{-j}' \leq 2\pi^m - x_j' < S$ in this period (and zero in all future periods) from accepting. So his incentive constraint is satisfied for all discount factors. Buyer B_j is willing to reject his offer if and only if

$$x_j' \leq S + \frac{\delta}{1-\delta} 2\pi^m.$$

[9] If instead the deviant offers sum to more than $2\pi^m$, equilibrium prescribes that both buyers accept the offers. As each buyer has a myopic incentive to do so, provided the other buyer does as well, the incentive constraints are trivially satisfied in that case.

As this inequality has to hold for all x'_j not exceeding $2\pi^m$, B_j's incentive constraint is satisfied if and only if $\delta \geq \hat{\delta}^{OPC}$.

It remains to show that, for any $\delta \geq \hat{\delta}^{OPC}$, there is an equilibrium that gives buyer B_j a per-period payoff of $2\pi^m$.[10] Along the equilibrium path, I offers $x_j = 2\pi^m$ to B_j and $x_{-j} = 0$ to B_{-j}, both buyers accept, and no entry takes place. (Note that neither buyer has a myopic incentive to deviate, so equilibrium may as well prescribe play of the same equilibrium in all future periods following a deviation by a buyer.) Consider now a deviation by I to offers (x'_1, x'_2). If $x'_1 + x'_2 > 2\pi^m$, both offers are accepted, and the same equilibrium (giving I a zero payoff) is played forever after. (Accepting the offer is myopically optimal for each buyer, provided the other buyer does so as well, so no dynamic incentives need to be provided.) If instead $x'_1 + x'_2 \leq 2\pi^m$, both offers are rejected (so that entry occurs, yielding a payoff of S to each buyer in the current period); equilibrium play from the next period onward is exactly the same as following the same vector of deviant offers in the optimal (non-simple) punishment scheme for sustaining the perpetual play of N described above. As our analysis above shows, both buyers have an incentive to reject I's deviant offers if $\delta \geq \hat{\delta}^{OPC}$. Hence, I does not have a profitable deviation.

3.2.2 Discussion

In this extended example, 'no exclusive dealing' can be supported as an equilibrium outcome in a repeated version of the "naked exclusion" game. On the equilibrium path, the incumbent is not supposed to offer exclusive dealing contracts to his retailers – but if he does, both of the retailers must reject these offers in order to inflict within-period punishment on the incumbent and reduce his temptation to deviate. Inducing such rejection does not require dynamic incentives if the incumbent offers a payment for exclusive dealing of less than S, since (if rejection by the other retailer is expected), rejection is myopically optimal. But inducing rejection of an exclusive dealing offer with a payment of more than S does require dynamic incentives since such rejection would decrease the retailer's current profit. Retailers must therefore be 'rewarded' in the continuation game for rejecting such a tempting exclusive dealing contract.

In contrast to the first set of examples, in this setting, there is no trade-off between rewarding the punisher and punishing the deviator, since equilibria which reward the punishers (exclusive dealing or equilibria with entry) can be constructed which yield zero profits to the incumbent. There is, however, a trade-off between rewarding the two retailers who must both reject the deviant offers. The most that both retailers can receive at the same time is S per retailer (when entry occurs in every period), so the best simple penal code involves the play of this equilibrium after any deviation by the incumbent. But there exist equilibria with exclusive dealing that can provide more than S to one retailer (and less than S to the other one, and zero to the incumbent). Thus, if the incumbent makes a

[10] As both I and B_{-j} can ensure themselves a payoff of at least zero, there does not exist an equilibrium in which B_j earns more than $2\pi^m$ per period.

deviation involving a bribe of more than S to one retailer (and less than S to the other, otherwise the deviation is unprofitable), the optimal punishment scheme provides a larger 'carrot' to the retailer most tempted. In this setting, as in Section 2, the continuation play again optimally depends on the particular deviation chosen, but this time because the reward provided must fit the temptation resisted (or the sacrifice made) when rejecting the deviation.

The principle that the reward provided should be tailored to the sacrifice is natural. This is particularly so in the dynamic variant of the repeated naked exclusion game set out above, where, rather than entrants being short-lived, the entrant replaces the incumbent in the period following rejection of all exclusive dealing contracts, and hence becomes the new incumbent, offering exclusive dealing contracts against a new entrant, the following period.[11] In this dynamic game, if an incumbent deviates by offering an asymmetric set of exclusive dealing contracts (more than S to one retailer, less than S to another), then it is natural that the entrant, when he becomes the incumbent next period, plays a continuation equilibrium which provides larger rewards to the retailer which made a greater sacrifice in enabling the entrant to replace the old incumbent.

The same forces are also at work in Nocke and White (2007). That paper studies collusion in an intermediate goods industry where several upstream firms compete to sell inputs to downstream retailers. In the stage game, which is infinitely repeated, the upstream firms first simultaneously make contract offers to downstream firms; then, the downstream firms simultaneously decide which contract(s) to accept; and, finally, the downstream firms compete in the retail market. An upstream firm can profitably deviate from the collusive equilibrium only if his deviant contract offer is accepted by at least one downstream retailer. Whilst Nocke and White concentrate on the simpler case where collusion is sustained by "Nash reversion", it can be shown that the optimal punishment scheme induces downstream retailers to help sustain upstream collusion by rejecting certain deviant offers (in particular, those which are very profitable for the deviating firm). In order that downstream firms do indeed reject such offers, the tempted downstream retailers need to be "rewarded" by the play of an equilibrium in the continuation game that is favorable to them in that case, with the likelihood of this reward optimally chosen to be increasing in the relative size of the temptation.

Weingast (1995, 1997) has also analyzed a game with a structure similar to the naked exclusion game set out above. In Weingast's game, an incumbent sovereign can expropriate one or both of his subject groups. The subject groups can successfully resist expropriation only if they both do so, however (just as exclusivity is successfully resisted only if both retailers reject it), and resistance is costly. As in the naked exclusion game, it is desirable to support an equilibrium with no expropriation by the sovereign, but this is impossible in

[11] Since we focus on equilibria in which the incumbent makes zero profit in all future periods, independent of whether or not he deviates, our analysis of the optimal punishment extends to this dynamic version of the game. In the dynamic version, it is the entrant who rewards the buyer rejecting the exclusive dealing contract.

the one-shot game because the sovereign can play a 'divide and conquer' strategy whereby he expropriates only one of the two groups, so that the other has no myopic incentive to resist. This is formally very similar to the incumbent making asymmetric exclusive dealing offers in the game analyzed above. Weingast discusses informally the potential benefits of repetition as a way to coordinate resistance and provide rewards for groups which join the resistance against their myopic self interest; these benefits have been confirmed experimentally by Cason and Mui (2013). Our analysis suggests that coordinated resistance will be most successful when the division of the spoils from successful resistance (provided in the continuation game) takes into account the relative sacrifices made in resistance. This requires that players anticipate that a new sovereign put in place following successful coordinated resistance will play an asymmetric equilibrium, where those who made the largest sacrifices in resisting the ancien régime are most rewarded when it is overthrown, a suggestion which seems not implausible. These examples show that the supposed "complexity" of optimal penal codes in extensive form games need not be seen as a disadvantage and a reason to focus instead on Nash-reversion in such games, but can instead be a source of richness and additional predictions for a model.

4 Conclusion

A major concern in any long run interaction is the provision of incentives to discipline the behavior of agents. In simple interactions, such as the repeated prisoners dilemma, an opportunistic deviation is immediately profitable (since the other players cannot immediately react) and can only be deterred by appropriate specifications of continuation play (or future punishments). In many applications, interactions are intrinsically dynamic and opportunistic deviations are only profitable if they are validated by the complicit behavior of at least one other player. Moreover, this complicit behavior is often myopically optimal (since, for example, it may involve accepting a "bribe").

When agents are impatient, deterring deviations may therefore require preventing the complicit behavior, which requires specifying deviation-dependent continuations. In this paper, we have indicated two likely causes of the need for deviation-dependent continuations: 1) Rewarding a potentially complicit player for not being complicit also rewards the original deviator; and 2) there are multiple potentially complicit players and there is a trade-off in rewarding players for not being complicit.

We have illustrated these causes by providing applications to a repeated bilateral investment game with hold-up in the spirit of Klein, Crawford, and Alchian (1978) and Grossman and Hart (1986), and to a repeated naked exclusion game modeled on Rasmusen, Ramseyer, and Wiley (1991) and Segal and Whinston (2000). But we expect that the same phenomena will arise in many other interesting applications. The first cause of the need for deviation-dependent continuations – the commonality of interest between potential punisher and punishee – arises in many environments, such as the interaction between parent

and child; between legislators of the same political party; between a monetary authority and the public; and between principal and agent in relational contracting. The second cause – the trade-offs between rewarding different punishing parties – arises in settings such as multi-player bargaining or lobbying; collusion in vertically related markets; and a defendant dealing with multiple claimants. As we hope our examples show, examining the particular structure of optimal punishment in an applied settings can yield new insights as to ways to make existing institutions work better.

Our analysis also indicates the importance of carefully modeling the within-period interactions. In this respect, the game theoretic literatures focus on folk theorems with patient players is unfortunate and potentially seriously misleading. In particular, such a focus allows researchers to use the normal form of the within-period interaction, which effectively ignores the within-period dynamic structure.

References

Abreu, D. (1986): "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39(1), 191–225.

——— (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56(2), 383–396.

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 73(5), 1431–1475.

Baron, D., and J. Ferejohn (1989): "Bargaining in Legislatures," *American Political Science Review*, 83(4), 1181–1206.

Cason, T. N., and V.-L. Mui (2013): "Coordinating Resistance through Communication and Repeated Interaction," *Economic Journal*, 124(566), F226–F256.

Friedman, J. W. (1971): "A Non-Cooperative Equilibrium for Supergames," *Review of Economic Studies*, 38(1), 1–12.

Fudenberg, D., and E. Maskin (1986): "The Folk Theorem in Repeated Games with Discounting or Incomplete Information," *Econometrica*, 54(3), 533–554.

Grossman, S. J., and O. D. Hart (1986): "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94(4), 691–719.

Klein, B., R. G. Crawford, and A. Alchian (1978): "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *Journal of Law and Economics*, 21(2), 297–326.

Lagunoff, R., and A. Matsui (1997): "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 6(5), 1467–1477.

Levin, J. (2003): "Relational Incentive Contracts," *American Economic Review*, 93(3), 835–857.

Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York, NY.

Nocke, V., and L. White (2007): "Do vertical mergers facilitate upstream collusion?," *American Economic Review*, 97(4), 1321–1339.

Rasmusen, E. B., J. M. Ramseyer, and J. S. Wiley, Jr. (1991): "Naked Exclusion," *American Economic Review*, 81(5), 1137–1145.

Rubinstein, A., and A. Wolinsky (1995): "Remarks on Infinitely Repeated Extensive-Form Games," *Games and Economic Behavior*, 9(1), 110–115.

Segal, I. R., and M. D. Whinston (2000): "Naked Exclusion: Comment," *American Economic Review*, 90(1), 296–309.

Sorin, S. (1995): "A Note on Repeated Extensive Games," *Games and Economic Behavior*, 9(1), 116–123.

Sutton, J. (1986): "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53(5), 709–724.

Weingast, B. R. (1995): "The Economic Role of Political Institutions: Market-Preserving Federalism and Economic Development," *Journal of Law, Economics, and Organization*, 11(1), 1–31.

——— (1997): "The Political Foundations of Democracy and the Rule of Law," *American Political Science Review*, 91(2), 245–262.

Wen, Q. (2002): "A Folk Theorem for Repeated Sequential Games," *Review of Economic Studies*, 69(2), 493–512.