

MULTILATERAL INTERACTIONS IMPROVE COOPERATION UNDER RANDOM FLUCTUATIONS*

Jin Li
Northwestern University

Michael Powell
Northwestern University

January 2017

Abstract

In an environment subject to random fluctuations, when does an increase in the breadth of activities in which individuals interact together help foster collaboration on each activity? We show that when players, on average, prefer to stick to a cooperative agreement rather than renegeing by taking their privately optimal action, then such an agreement can be approximated as equilibrium play in a sufficiently broad relationship. This is in contrast to existing results showing that a cooperative agreement can be sustained only if players prefer to adhere to it in every state of the world. We consider applications to favor exchange, multimarket contact, and relational contracts.

Keywords: repeated games, relational contracts, multimarket contact, favor exchange

JEL classifications: C73, L14

*Li: Kellogg School of Management, jin-li@kellogg.northwestern.edu. Powell: Kellogg School of Management, mike-powell@kellogg.northwestern.edu. We thank Makoto Hanazono, Rongzhu Ke, and Tadashi Sekiguchi for helpful comments.

1 Introduction

Long-term relationships allow self-interested individuals to achieve cooperation by putting the future at stake: individuals cooperate with each other, because the long-term gain from future interactions exceeds the short-term gains they forgo by pursuing privately optimal courses of action. However, random fluctuations in the environment can lead to extreme situations in which short-run gains are high, and these extreme situations can strain otherwise-healthy relationships and destroy gains from cooperation, even if they are rare.

In this paper, we study the extent to which increasing the *breadth* of interactions can help foster cooperation. Field studies have highlighted the importance of broad relationships: for example, in describing the interactions between farmers in Shasta County, California, Ellickson reports that, “Rural residents deal with one another on a large number of fronts, and most residents expect those interactions to continue far into the future... They interact on water supply, controlled burns, fence repair, social events, staffing the volunteer fire department, and so on.” (1994, p.55) Theoretical work pioneered by Bernheim and Whinston (1990) has shown that in broad relationships, parties may use the fact that cooperation is easy to sustain in some aspects of their relationship to help foster cooperation in more difficult aspects, but for the most part, this work has primarily focused on environments with no random fluctuations. We ask whether and to what extent increasing the breadth of relationships in environments with random fluctuations aids cooperation.

Specifically, we consider a repeated simultaneous-move game composed of M identical and independent *component games*. We refer to a strategy profile of the component stage game as a *component agreement*. For a single component game (i.e., $M = 1$), the condition for a component agreement to be sustained as an equilibrium—to correspond in each period to the on-path play of a subgame-perfect equilibrium—is that for each player, the future value of adhering to the component agreement exceeds his *maximal deviation gain* across all states of the world. Since these games are identical, when $M > 1$, the same condition has to be satisfied for the component agreement to be *perfectly replicated*, that is, to correspond to equilibrium play in each of the M component games.

If we relax the requirement of perfect replication slightly, however, the required condition changes substantially. Specifically, we say that a component agreement can be *almost-perfectly replicated* if for any $\varepsilon > 0$, for sufficiently large M , there is a subgame-perfect equilibrium in which equilibrium play coincides with the component agreement in each component game with probability greater than $1 - \varepsilon$. For a component agreement to be almost-perfectly replicated, the condition is that, for each player, his *mean deviation gain*—rather than his maximal deviation gain—is smaller than his future value of adhering to the component agreement.

There are therefore gains from multilateral cooperation, even if the component games are identical. To see why the mean component deviation gain matters for almost-perfect replication, note that if we were to perfectly replicate a component agreement, most of the time the *average deviation gain* is sufficiently close to the mean component deviation gain as the number of component games grows. We can therefore construct agreements—strategy profiles of the stage game—that “chop and replace” play in states in which players have large average deviation gains with stage-game Nash equilibrium play, and these agreements can be sustained as an equilibrium.

While this limit result provides a condition for almost-perfect replication, it does not prescribe an *optimal* agreement in a relationship of finite breadth. We therefore study optimal agreements for finite values of M in several important economic settings. Specifically, we consider applications to favor-exchange games (Mobius, 2001) and to Bertrand collusion games with demand shocks (Rotemberg and Saloner, 1986). Constructing optimal agreements in these settings involves modifying our “chop and replace” construction so as to be judicious with the “replace” aspect in a way that depends on the particular setting.

Our results also extend to a class of games with extensive-form stage games and private actions. In particular, we consider two classes of models of relational incentive contracts with noisy performance measures: one in which there is a single agent who exerts effort in M activities, and one in which there are M agents who each exert effort in a single activity. We show how our limit result can be applied to these settings, and we characterize optimal relational contracts for finite values of M .

Related Literature Bernheim and Whinston (1990) show that, in a deterministic environment, conditioning play in one component game on outcomes in other component games can aid cooperation only if these component games are not identical. In this paper, we identify key conditions under which, in environments characterized by random fluctuations, conditioning play in one component game on outcomes in others can help foster cooperation.

The most closely related work is a contemporaneous paper by Sekiguchi (2015), which studies how increasing the breadth of interaction affects the optimal degree of collusion in a multimarket contact setting with stochastic demand realizations. Sekiguchi derives a critical discount factor such that for higher discount factors, almost-perfect collusion can be sustained as the number of markets approaches infinity. Our paper differs in several ways. First, our main limit result (Theorem 1) is weaker in a probabilistic sense, though it is proven in a more general economic environment. Further, our limit result characterizes necessary and sufficient conditions for *any* component agreement to be almost-perfectly replicated, whereas Sekiguchi focuses on optimal symmetric subgame-perfect

equilibria. Our conditions for perfect collusion to be almost-perfectly replicated coincide with the critical discount factor that Sekiguchi identifies. Finally, in addition to our limit results, we analyze optimal equilibrium agreements away from the limit in several classes of games, and we also consider a class of repeated extensive-form games with private actions.

Our application to relational incentive contracts with multiple activities is related to Bond and Gomes (2009), which characterizes the optimal formal contract in an agency problem with multiple tasks and a fixed upper bound on total wage payments. In our application, the upper bound on total payments arises from the principal's limited commitment. Barron (2013) considers endogenous breadth of relational contracts in a supply-chain context and shows that suppliers underspecialize relative to the first-best. Fong and Li (Forthcoming) shows that intertemporal garbling of public information helps link incentive constraints in relational contracts so that shocks can be smoothed over time.

2 The Model

There are N players who play a repeated game in discrete time, which is indexed by $t = 1, 2, \dots$, and all players share a common discount factor $\delta \in (0, 1)$. The **stage game** is composed of M identical **component stage games**, indexed by $m \in \mathcal{M} = \{1, \dots, M\}$, which are played simultaneously. We refer to the repeated stage game as the **supergame**, the repeated component stage game as a **component supergame**, and we refer to M as the **breadth of interaction**. Throughout, variables with tildes will correspond to variables from component supergames, and variables without tildes will correspond to variables from the supergame.

In each period t , within each component stage game m , a state $\tilde{s}_{m,t} \in \mathcal{S} = \{1, \dots, S\}$ is drawn with probability $p_{\tilde{s}_{m,t}}$. Denote the vector of states realized in period t across all M component stage games as $s_t = (\tilde{s}_{1,t}, \dots, \tilde{s}_{M,t}) \in \mathcal{S}^M$. States are independent and identically distributed across periods and across component stage games: the probability that a state s_t is realized is $\Pr[s_t] = \prod_{m=1}^M p_{\tilde{s}_{m,t}}$. Within each component stage game m , the set of actions available to player i is given by the compact set \mathcal{A}_i , and we denote $\mathcal{A} = \prod_{i=1}^N \mathcal{A}_i$.

At the beginning of each period t , players commonly observe s_t and then simultaneously choose their actions, which are commonly observed. Denote player i 's actions in period t by the vector $a_{i,t} = (\tilde{a}_{i,1,t}, \dots, \tilde{a}_{i,M,t}) \in \mathcal{A}_i^M$. Player i 's payoff in component stage game m in period t is given by $\tilde{u}_i(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t})$, where $\tilde{a}_{-i,m,t} \in \mathcal{A}_{-i} \equiv \prod_{j \neq i} \mathcal{A}_j$. Payoffs are additive across component stage games, so player i 's stage-game payoff in period t is given by $u_i(a_{i,t}, a_{-i,t}, s_t) = \sum_{m=1}^M \tilde{u}_i(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t})$. Component-stage-game payoffs are bounded: there exists $K > 0$

such that $\tilde{u}_i(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t}) \in [-K, K]$ for all i, m , and $(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t})$.

Cooperative Agreements, Strategies, and Equilibrium For expositional purposes, we refer to a pure strategy profile of the component stage game, $\tilde{A} : \mathcal{S} \rightarrow \mathcal{A}$, as a **component agreement**. In each period t , within component game m , therefore, $\tilde{A}_i(\tilde{s}_{m,t}) \in \mathcal{A}_i$ denotes player i 's action under \tilde{A} in state $\tilde{s}_{m,t}$. A **Nash equilibrium of a component stage game** is a component agreement \tilde{A} such that $\tilde{u}_i(\tilde{A}_i(\tilde{s}_{m,t}), \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t}) \geq \tilde{u}_i(\tilde{a}_{i,m}, \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t})$ for all $\tilde{a}_{i,m} \in \mathcal{A}_i$ for all i and for all $\tilde{s}_{m,t} \in \mathcal{S}$.

We refer to a pure strategy profile of the stage game, $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$, as an **agreement**. Note that for each s_t , $A(s_t)$ is an $M \times N$ matrix specifying actions for each player in each component game. In each period t , denote, $A_i(s_t) \in \mathcal{A}_i^M$ as player i 's M -dimensional action vector under A in state s_t , denote $A_{(m)}(s_t) \in \mathcal{A}$ as the N -dimensional vector of players' actions under A in component game m in state s_t , and $A_{i,m}(s_t) \in \mathcal{A}_i$ as player i 's action in component game m in state s_t . A **Nash equilibrium of the stage game** is an agreement A such that $u_i(A_i(s_t), A_{-i}(s_t), s_t) \geq u_i(a_i, A_{-i}(s_t), s_t)$ for all $a_i \in \mathcal{A}_i^M$ for all i and for all $s_t \in \mathcal{S}^M$.

We next define strategies in a component supergame and in the supergame. Throughout the paper, we restrict attention to trigger strategies. In particular, a **trigger strategy of a component supergame** specifies a component agreement, $\tilde{A}(0) : \mathcal{S} \rightarrow \mathcal{A}$, and an N -tuple $(\tilde{A}(1), \dots, \tilde{A}(N))$, where each $\tilde{A}(i)$ is a Nash equilibrium of a component stage game. Play begins with $\tilde{A}(0)$, and if player i was the first to deviate from $\tilde{A}(0)$, play transitions to $\tilde{A}(i)$ for all future periods. A **trigger strategy of the supergame** specifies an agreement $A(0) : \mathcal{S}^M \rightarrow \mathcal{A}^M$, and an N -tuple $(A(1), \dots, A(N))$, where each $A(i)$ is a Nash equilibrium of the stage game. Play begins with $A(0)$, and if player i was the first to deviate from $A(0)$, play transitions to $A(i)$ for all future periods.

We will say that a component agreement \tilde{A}^* is an **equilibrium component agreement** if there exists a trigger strategy of the component game with $\tilde{A}(0) = \tilde{A}^*$ that is a subgame-perfect Nash equilibrium of the component supergame. An agreement A^* is an **equilibrium agreement** if there exists a trigger strategy of the supergame with $A(0) = A^*$ that is a subgame-perfect Nash equilibrium of the supergame.

Throughout, we assume that the component stage game has at least one pure-strategy Nash equilibrium. For each player i , denote by \tilde{v}_i his lowest expected payoff in any pure-strategy Nash equilibrium of a component stage game, and denote by \underline{v}_i his lowest expected payoff in any pure-strategy Nash equilibrium of the stage game. Note that $\underline{v}_i = M\tilde{v}_i$, which we prove in Lemma A1

in the appendix. At the end of Section 3, we discuss our restriction to trigger strategies, pure-strategy agreements, and pure-strategy Nash equilibrium punishments, and we discuss the role of our assumption that states in a given period are independent across component stage games.

3 General Limit Results

We study the degree to which an increase in the breadth of interaction helps sustain—in a way we will make precise below—a component agreement \tilde{A}^* as an equilibrium outcome. We first establish a necessary and sufficient condition for \tilde{A}^* to be an equilibrium component agreement. Recall that $\tilde{u}_i(\tilde{A}_i(\tilde{s}_{m,t}), \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t})$ is player i 's payoff in component stage game m in period t in state $\tilde{s}_{m,t}$ under component agreement \tilde{A} . Now, define $\tilde{v}_i(\tilde{A}) \equiv \sum_{\tilde{s}_{m,t}} p_{\tilde{s}_{m,t}} \tilde{u}_i(\tilde{A}_i(\tilde{s}_{m,t}), \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t})$ as player i 's expected component-stage-game payoff under \tilde{A} . We refer to the quantity

$$\tilde{d}_i(\tilde{s}_{m,t}; \tilde{A}) = \max_{\tilde{a}_{i,m} \in \mathcal{A}_i} \tilde{u}_i(\tilde{a}_{i,m}, \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t}) - \tilde{u}_i(\tilde{A}_i(\tilde{s}_{m,t}), \tilde{A}_{-i}(\tilde{s}_{m,t}), \tilde{s}_{m,t})$$

as player i 's **component deviation gain in state $\tilde{s}_{m,t}$ under \tilde{A}** and to the quantity $\tilde{d}_i^{\max}(\tilde{A}) = \max_{\tilde{s}_{m,t}} \{\tilde{d}_i(\tilde{s}_{m,t}; \tilde{A})\}$ as player i 's **maximal component deviation gain under \tilde{A}** .

Recall that our definition of an equilibrium component agreement involves the use of trigger strategies. For the purposes of characterizing whether a particular agreement \tilde{A} is an equilibrium agreement, it is without loss of generality to assume that, following a deviation by player i , players repeatedly play the Nash equilibrium of the component stage game that yields payoff \tilde{v}_i for player i in each period. Therefore, a deviation by player i results in a **component continuation loss under \tilde{A}** given by $\delta \tilde{V}_i(\tilde{A}) \equiv \delta (\tilde{v}_i(\tilde{A}) - \tilde{v}_i) / (1 - \delta)$. For \tilde{A}^* to be an equilibrium component agreement, a necessary and sufficient condition is that the component continuation loss exceeds the maximal component deviation gain for each player i . That is,

$$\tilde{d}_i^{\max}(\tilde{A}^*) \leq \delta \tilde{V}_i(\tilde{A}^*) \text{ for all } i. \quad (1)$$

Next, we consider the effect of an increase in the breadth of interaction on sustaining a component agreement as part of an equilibrium agreement. Given an agreement $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$, we define player i 's **deviation gain in state s_t under A** as

$$d_i(s_t; A) = \max_{a_i \in \mathcal{A}_i^M} u_i(a_i, A_{-i}(s_t), s_t) - u_i(A_i(s_t), A_{-i}(s_t), s_t).$$

We also define $d_i^{\max}(A) = \max_{s_t} \{d_i(s_t; A)\}$ as player i 's **maximal deviation gain under A** and $d_i(s_t; A)/M$ to be player i 's **average deviation gain under A** . Similarly, a deviation by

player i results in a **continuation loss** given by $\delta V_i(A) = \delta(v_i(A) - \underline{v}_i) / (1 - \delta)$, where $v_i(A) = \sum_{m=1}^M \tilde{v}_i(A_{(m)})$. Given M , we next define an ε -**neighborhood of component agreement** \tilde{A} by

$$\mathcal{N}^M(\tilde{A}, \varepsilon) = \left\{ A : \Pr \left[A_{i,m}(s_t) = \tilde{A}_i(\tilde{s}_{m,t}) \text{ for all } i \text{ and } m \right] \geq 1 - \varepsilon \right\}.$$

That is, an ε -neighborhood of component agreement \tilde{A} is the set of all agreements A that prescribe the same play, component-stage-game by component-stage-game as \tilde{A} on a set of states that occurs with probability greater than $1 - \varepsilon$. As ε approaches zero, for any $A \in \mathcal{N}^M(\tilde{A}, \varepsilon)$, each players' payoffs in each component game converge to the payoffs they would receive under \tilde{A} . Note that if $A \in \mathcal{N}^M(\tilde{A}, 0)$, then A perfectly coincides with \tilde{A} in each component game, that is, for all s_t , $A_{i,m}(s_t) = \tilde{A}_i(\tilde{s}_{m,t})$ for all i and m . We will say that \tilde{A} can be M -**perfectly replicated** if there exists an equilibrium agreement $A \in \mathcal{N}^M(\tilde{A}, 0)$. By definition, \tilde{A} is an equilibrium component agreement if and only if it can be 1-perfectly replicated.

It is well-known that the conditions required for \tilde{A}^* to be an equilibrium component agreement are the same as the conditions required for any $A^* \in \mathcal{N}^M(\tilde{A}^*, 0)$ to be an equilibrium agreement for any M (Bernheim and Whinston, 1990; Mailath and Samuelson, 2006, p.162), because the component games are identical. Specifically, for any $A^* \in \mathcal{N}^M(\tilde{A}^*, 0)$, each player i 's continuation loss is given by $\delta V_i(A^*) = M\delta\tilde{V}_i(\tilde{A}^*)$, and his maximal deviation gain is $d_i^{\max}(A^*) = Md_i^{\max}(\tilde{A}^*)$, which occurs in the state s_t in which his component deviation gain in each component game is equal to his maximal component deviation gain. As a result, the necessary and sufficient condition for $A^* \in \mathcal{N}^M(\tilde{A}^*, 0)$ to be an equilibrium agreement is

$$Md_i^{\max}(\tilde{A}^*) \leq M\delta\tilde{V}_i(\tilde{A}^*) \text{ for all } i,$$

which is identical to (1). In particular, M -component punishments deter M -component deviations if and only if single-component punishments deter single-component deviations.

The previous result shows that increasing the breadth of interaction does not relax the conditions required for \tilde{A} to be M -perfectly replicated. We now show that if we relax the requirement of *perfect* replication, the required conditions change substantially. To do so, we say that a component agreement \tilde{A}^* can be **almost-perfectly replicated** if for any $\varepsilon > 0$, there exists $M(\varepsilon)$ such that for all $M \geq M(\varepsilon)$, there exists an equilibrium agreement $A^M \in \mathcal{N}^M(\tilde{A}^*, \varepsilon)$.

Our definition of almost-perfect replication describes the set of component agreements for which equilibrium play in a sufficiently broad relationship almost perfectly coincides with those component agreements in each component game. Define $\tilde{d}_i^{\text{mean}}(\tilde{A}) = \sum_{\tilde{s}_{m,t} \in \mathcal{S}} p_{\tilde{s}_{m,t}} \tilde{d}_i(\tilde{s}_{m,t}; \tilde{A})$ as player i 's **mean component deviation gain** under \tilde{A} . The result below shows that the condition for \tilde{A}^*

to be almost-perfectly replicated does not depend on the maximal component deviation gain but rather on the mean component deviation gain.

Theorem 1. *Consider a component agreement \tilde{A}^* . The following are true:*

1. *If $\tilde{d}_i^{mean}(\tilde{A}^*) < \delta\tilde{V}_i(\tilde{A}^*)$ for all i , then \tilde{A}^* can be almost-perfectly replicated.*
2. *If $\tilde{d}_i^{mean}(\tilde{A}^*) > \delta\tilde{V}_i(\tilde{A}^*)$ for some i , then \tilde{A}^* cannot be almost-perfectly replicated.*

The proof of Theorem 1 and all other results are in the appendix. The proof of Part 1 involves constructing a sequence of agreements that “chop off” the states in which players have large deviation gains and asymptotically approximates \tilde{A}^* in each of the component games. In particular, given \tilde{A}^* , we construct a sequence of agreements A^M so that for M sufficiently large, A^M is an equilibrium agreement, and A^M is in an ε -neighborhood of component agreement \tilde{A}^* . In this sequence, A^M coincides with \tilde{A}^* in each component game for all states except for those in which the average deviation gain under A^M exceeds the mean component deviation gain under \tilde{A}^* . In those states, A^M prescribes play that corresponds to a Nash equilibrium of the stage game. By the weak law of large numbers, as the number of component games grows, the probability of the set of states that are “chopped off” in this construction goes to zero, ensuring that A^M is in an ε -neighborhood of \tilde{A}^* and that the continuation loss approaches $M\delta\tilde{V}_i(\tilde{A}^*)$. It then follows that in a sufficiently broad relationship, $d_i^{max}(A^M) \leq \delta V_i(A^M)$, and therefore \tilde{A}^* can be almost-perfectly replicated.

Part 1 provides sufficient conditions for \tilde{A}^* to be almost-perfectly replicated, and part 2 shows that these conditions are almost necessary. In particular, \tilde{A}^* cannot be almost-perfectly replicated if the mean component deviation gain exceeds the component continuation loss for any player i . When $\tilde{d}_i^{mean}(\tilde{A}^*) > \delta\tilde{V}_i(\tilde{A}^*)$, the pigeonhole principle implies that it is impossible to allocate the deviation gains across states in a way that guarantees they are smaller than the continuation losses in every state, so no $A \in \mathcal{N}^M(\tilde{A}^*, \varepsilon)$ can be an equilibrium agreement. The component agreement \tilde{A}^* , therefore, cannot be almost-perfectly replicated.

Theorem 1 implies there are gains from linking play in the component games together, even if the component games are identical and independent. The reason for the gains from multilateral cooperation follows from the familiar logic of cross-subsidization of constraints. Even if the component games are identical, the component deviation gains in each component game depend on the realized state. Therefore, component games with small realized component deviation gains can be used to cross-subsidize those with large realized component deviation gains. Theorem 1 shows that in sufficiently broad relationships, the cross-subsidization can be made nearly perfect in the sense

that only the mean component deviation gain is relevant for almost-perfect replication. Theorem 1 also describes the limits of such gains, and the next section illustrates how these gains can be best realized for finite M in two widely studied applications.

The next proposition shows that when \tilde{A} satisfies $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$, whether it can be almost-perfectly replicated depends on the details of the underlying game. Recall that $\tilde{v}_i(\tilde{A})$ is player i 's expected component-stage-game payoff under \tilde{A} . We say that a component agreement \tilde{A} is **Pareto-optimal in the component stage game** if for any other component agreement \tilde{A}' , $\tilde{v}_i(\tilde{A}) > \tilde{v}_i(\tilde{A}')$ for some i implies $\tilde{v}_j(\tilde{A}) < \tilde{v}_j(\tilde{A}')$ for some j . If \tilde{A} can be almost-perfectly replicated, then we can define, for any $\varepsilon > 0$, the smallest integer $\hat{M}(\varepsilon)$ such that for all $M \geq \hat{M}(\varepsilon)$, there is an equilibrium agreement $A \in \mathcal{N}^M(\tilde{A}, \varepsilon)$.

Proposition 1. *Consider a component agreement \tilde{A} . The following are true:*

1. *If $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for all i , then \tilde{A} can be M -perfectly replicated for all M if and only if $\tilde{d}_i(\cdot; \tilde{A})$ is constant.*
2. *If $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for all i , $\tilde{d}_i(\cdot; \tilde{A})$ is not constant for some i , and \tilde{A} is Pareto-optimal in the component stage game, then \tilde{A} cannot be almost-perfectly replicated.*
3. *Suppose $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for any i , and \tilde{A} can be almost-perfectly replicated. then $\lim_{\varepsilon \rightarrow 0} \varepsilon \sqrt{\hat{M}(\varepsilon)} = \infty$.*

When players' deviation gains do not fluctuate, and $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for all i , \tilde{A} is an equilibrium component agreement. The component agreement \tilde{A} can therefore be M -perfectly replicated for all M as stated in Part 1 of the proposition. When players' deviation gains do fluctuate, Part 2 shows that \tilde{A} cannot be almost-perfectly replicated if \tilde{A} is Pareto-optimal in the stage game. If \tilde{A} is not Pareto-optimal in the stage game, there are games in which \tilde{A} can be almost-perfectly replicated. In these cases, however, the rate of convergence is limited, as Part 3 shows.

Discussion of Model Assumptions We now discuss how the results of Theorem 1 would continue to hold if we relaxed several of our models' assumptions.

First, it is not important that the component agreement \tilde{A} is a pure strategy of the component stage game. Rather, what is important is that any deviations from \tilde{A} are commonly observed. As long as this is the case, we can allow for mixed-strategy component agreements. For example,

suppose each player submits a probability distribution to a randomization device, and a deviation occurs—and is commonly observed—if he submits a different probability distribution.

Next, for consistency, we have restricted \tilde{v}_i to be the lowest pure-strategy stage-game Nash equilibrium payoff for player i . We can redefine \tilde{v}_i to be the lowest (possibly mixed) Nash equilibrium payoff for player i , and in our proposed trigger strategies, let $\tilde{A}(i)$ correspond to the (possibly mixed) Nash equilibrium in which player i receives \tilde{v}_i , and the statement of Theorem 1 would remain unchanged. Therefore, the assumption that a pure-strategy Nash equilibrium exists in the component stage game is not crucial for our results.

Third, we have restricted attention to trigger strategies. This restriction is without loss of generality if for each player i , \tilde{v}_i is his minmax payoff in the component stage game. This is the case in all the applications that follow.

Finally, when the component-stage-game states are independent from each other, we can apply the weak law of large numbers to show that if $\bar{A} \in \mathcal{N}^M(\tilde{A}, 0)$,

$$\lim_{M \rightarrow \infty} \Pr \left[d_i(s_t; \bar{A}) / M > \delta \tilde{V}_i(\tilde{A}) \right] = 0.$$

This is the critical condition for Part 1 of Theorem 1 to hold. This condition can also hold even if the component-stage-game states are not independent from each other—for example, if the component deviation gains are m -dependent (see Billingsley (1986)), the results of Part 1 of Theorem 1 would hold. The results of Part 2 of Theorem 1 hold for any correlation structure.

4 Optimal Agreements under Finite Breadth in Applications

While the limit result described in Theorem 1 provides a set of conditions for almost-perfect replication, it does not prescribe an *optimal* agreement when $M < \infty$. In this section, we construct optimal agreements for finite M in two sets of applications that are common in the literature on repeated games. The first is a favor-exchange setting in which multiple activities that require help arrive randomly each period (Mobius, 2001; Hauser and Hopenhayn, 2008). The second is Bertrand collusion under multimarket contact (Bernheim and Whinston, 1990) with demand fluctuations (Rotemberg and Saloner, 1986; Sekiguchi, 2015). The proof of Theorem 1 involves constructing a sequence of agreements that “chop and replace” play in the states in which players have large average deviation gains with stage-game Nash equilibrium play. Constructing optimal agreements in our two settings involves modifying this “chop and replace” construction, being judicious with the “replace” aspect in a way that depends on the particular setting.

We will say that an agreement $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$ is an **optimal agreement** if it is an equilibrium

agreement, and it maximizes the sum of players' payoffs over all equilibrium agreements. In the component stage game of each application, there is a Nash equilibrium in which each player receives his minmax payoff. It is therefore without loss of generality to focus on trigger strategies when characterizing optimal agreements.

4.1 Favor Exchange with Multiple Activities

Our first application is to multilateral favor-exchange games. In a favor-exchange game, in each period, each player may independently need a favor that the other player can grant. Granting a favor is costly to the player who grants it, but it benefits the other player if he needs it. Cooperation increases total surplus, but concerns about whether favors will be returned in the future limit a selfish player's incentives to grant favors today. A number of papers have studied the conditions under which favors are exchanged in equilibrium when players play a single favor-exchange game (Mobius, 2001; Hauser and Hopenhayn, 2008; Abdulkadiroglu and Bagwell, 2013).

We analyze a setting in which multiple favor-exchange games are played simultaneously. As an application of Theorem 1, we show that when full cooperation in the component game—in which each player grants a favor whenever the other player needs one—is not an equilibrium, there does not exist an equilibrium agreement in the M -component game in which all needed favors are granted. This occurs whenever the costs of granting a favor are larger than the future surplus generated by continued interaction. However, favors are not needed in every period, and as long as the *expected* costs of granting needed favors is smaller than the future surplus generated by continued interaction, then full cooperation in the component game can be almost-perfectly replicated: there exists a sequence of equilibrium agreements in which all needed favors are granted with probability approaching 1 as M grows.

Further, we characterize optimal agreements when M is finite. We first show that optimal agreements are symmetric threshold agreements: there is an integer $\bar{H}^*(M)$ such that each player grants the other player up to $\bar{H}^*(M)$ needed favors in each period. We show how to calculate $\bar{H}^*(M)$, and we calculate the limit of $\bar{H}^*(M)/M$ as $M \rightarrow \infty$.

The Favor-Exchange Component Game There are $N = 2$ players. In each period t , each player i independently needs a favor with probability p .¹ We denote the state by $\tilde{s}_{m,t} = (\tilde{s}_{1,m,t}, \tilde{s}_{2,m,t})$, where $\tilde{s}_{i,m,t} \in \{0, 1\}$, $\tilde{s}_{i,m,t} = 1$ indicates that player i needs a favor in period t , and $\tilde{s}_{i,m,t} = 0$ indicates that he does not. There are therefore four states, $\mathcal{S} = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$, and state $\tilde{s}_{m,t} = (k, \ell)$ occurs with probability $p_{(k,\ell)} = p^{k+\ell} (1-p)^{2-(k+\ell)}$. Players simultaneously

¹When players' needs are correlated within a component game, the results in this section continue to hold.

choose whether to grant a favor to the other. Player i chooses $\tilde{a}_{i,m,t} \in \mathcal{A}_i = \{0, 1\}$ at cost $c\tilde{a}_{i,m,t}$, where $\tilde{a}_{i,m,t}$ indicates that player i grants a favor to player $-i$. If player $-i$ needed the favor, he receives benefits $b\tilde{a}_{i,m,t}$. We assume that $b > c$. Player i 's payoff in period t is therefore $\tilde{u}_i(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t}) = b\tilde{a}_{-i,m,t}\tilde{s}_{i,m,t} - c\tilde{a}_{i,m,t}$.

A component agreement is a function $\tilde{A} : \mathcal{S} \rightarrow \mathcal{A}$ specifying whether each player grants a favor as a function of which player(s) need a favor. Note that in the component stage game, there is a Nash equilibrium in which both players choose $\tilde{a}_{i,m,t} = 0$ for all $\tilde{s}_{m,t}$, and they earn 0, which is their minmax payoff. We will refer to the component agreement \tilde{A}^C in which all favors are granted whenever they are needed, $\tilde{A}_i^C(\tilde{s}_{m,t}) = \tilde{s}_{-i,m,t}$, as **component full cooperation**. Under component full cooperation, player i 's component deviation gain is $\tilde{d}_i(\tilde{s}_{m,t}; \tilde{A}^C) = c\tilde{s}_{-i,m,t}$. His maximal component deviation gain, therefore, is $\tilde{d}_i^{\max}(\tilde{A}^C) = c$, and his mean component deviation gain is $\tilde{d}_i^{\text{mean}}(\tilde{A}^C) = pc$. His expected per-period utility is $p(b - c)$, and therefore his component continuation loss is $\delta\tilde{V}_i(\tilde{A}^C) = \delta p(b - c) / (1 - \delta)$. Component full cooperation is an equilibrium component agreement if and only if $\tilde{d}_i^{\max}(\tilde{A}^C) \leq \delta\tilde{V}_i(\tilde{A}^C)$ or

$$\frac{1}{p} \frac{c}{b - c} \leq \frac{\delta}{1 - \delta}. \quad (2)$$

The Favor-Exchange Supergame Players engage in M simultaneous and independent favor-exchange component games. A state is a vector $s_t = (\tilde{s}_{1,t}, \dots, \tilde{s}_{M,t})$, where $\tilde{s}_{m,t} = (\tilde{s}_{1,m,t}, \tilde{s}_{2,m,t})$ specifies a pair of component-game needs for each component game. An agreement $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$ specifies which favors each player will grant as a function of the needs of both players across all component games. Denote by $f_i(s_t; A) = \sum_{m=1}^M A_i(s_t)$ the number of favors player i grants in state s_t under agreement A , and let $f_i^{\max}(A) = \max_{s \in \mathcal{S}^M} f_i(s; A)$ denote the maximal number of favors granted by player i . Player i 's maximal deviation gain occurs whenever he is asked to grant $f_i^{\max}(A)$ favors, and his maximal deviation gain is $d_i^{\max}(A) = cf_i^{\max}(A)$. We refer to the agreement A^C in which all needed favors are granted in each component game, $A_{i,m}^C(s_t) = \tilde{s}_{-i,m,t}$, as **full cooperation**. Of course, $A^C \in \mathcal{N}^M(\tilde{A}^C, 0)$.

Replicating Component Full Cooperation As an illustration of our limit results from Section 3, we first ask when component full cooperation can be M -perfectly replicated, and then we ask when component full cooperation can be almost-perfectly replicated. Component full cooperation can be M -perfectly replicated if full cooperation is an equilibrium agreement. Under full cooperation, player i 's per-period expected utility is $Mp(b - c)$, and therefore his continuation loss is $\delta V_i(A^C) = \delta Mp(b - c) / (1 - \delta)$. His maximal deviation gain occurs in states in which he is asked

to grant M favors, so full cooperation is an equilibrium agreement if and only if $Mc \leq \delta V_i(A^C)$, which coincides exactly with (2). Engaging in M simultaneous favor-exchange component games does not affect the necessary and sufficient conditions for full cooperation to be an equilibrium agreement.

Nevertheless, increasing the breadth of interaction can help improve cooperation. In particular, while the maximal component deviation gain under component full cooperation is c , the mean component deviation gain is only pc . Theorem 1 therefore shows that, as long as

$$\frac{c}{b-c} < \frac{\delta}{1-\delta},$$

component full cooperation can be almost-perfectly replicated. First-best payoffs can therefore be approximately attained as $M \rightarrow \infty$ when this condition is satisfied.

Optimal Agreements when $M < \infty$ Our next result further illustrates how multilateral interactions can improve cooperation in relationships of finite breadth. We do so by characterizing an optimal agreement, that is, an equilibrium agreement that maximizes players' joint surplus over all equilibrium agreements. To describe an optimal agreement, it is useful to introduce a couple pieces of notation and terminology.

Denote by $P(H, M)$ the CDF of a binomial distribution with parameters (p, M) , so that $P(H, M)$ is the probability that there are H or fewer successes in M trials when the success probability for each trial is given by p . In addition, let $h_i(s_t) = (\tilde{s}_{i,1,t}, \dots, \tilde{s}_{i,M,t})$ denote player i 's needs in state s_t , and refer to $H_i(s_t) = \sum_{m=1}^M \tilde{s}_{i,m,t}$ as player i 's **total needs** in state s_t . We will say that an agreement A is a **symmetric cooperation agreement** if the favors that player i grants are independent of his own needs and if the favors he grants as a function of player $-i$'s needs are the same as the favors player $-i$ grants as a function of player i 's needs. That is, if $A_i(s_t) = \hat{A}_i(h_{-i}(s_t))$ for some function $\hat{A}_i : \mathcal{S}_{-i}^M \rightarrow \mathcal{A}_i^M$, and $\hat{A}_1(h_2(s_t)) = \hat{A}_2(h_1(s_t))$ for all s_t . A symmetric cooperation agreement A is an **\bar{H} -threshold cooperation agreement** if, whenever $H_{-i}(s_t) \leq \bar{H}$, $A_{i,m}(s_t) = \tilde{s}_{-i,m,t}$, and whenever $H_i(s_t) > \bar{H}$, $A_{i,m}(s_t) = \tilde{s}_{-i,m,t}$ in exactly \bar{H} component games in which $\tilde{s}_{-i,m,t} = 1$. In an \bar{H} -threshold cooperation agreement, which we denote by $A^{\bar{H}}$, player i 's continuation loss is given by $\delta V_i(A^{\bar{H}}) = \delta V(\bar{H}, M)$, where

$$V(\bar{H}, M) \equiv \frac{1}{1-\delta} \left[Mp(b-c) - \sum_{m=\bar{H}}^M (1-P(\bar{H}, M))(b-c) \right].$$

These payoffs correspond to the payoffs players earn in each period if they grant up to \bar{H} needed favors to each other.

Proposition 2. When $\frac{\delta}{1-\delta} \in \left(\frac{c}{b-c}, \frac{1}{p} \frac{c}{b-c}\right)$ any optimal agreement is an $\bar{H}^*(M)$ -threshold cooperation agreement, where $\bar{H}^*(M)$ is the largest integer satisfying $\bar{H}^*(M) c \leq \delta V(\bar{H}^*(M), M)$.

Proposition 2 shows that an optimal agreement is symmetric, and it takes the form of a threshold cooperation agreement: each player grants up to $\bar{H}^*(M)$ needed favors for the other player. The surplus of the relationship is determined by the expected number of needed favors granted in each period. Given a total surplus of $V(\bar{H}, M)$, the cutoff $\bar{H}^*(M)$ is the maximal number of needed favors granted such that the renegeing temptation $\bar{H}^*(M) c \leq \delta V(\bar{H}, M)$. The optimal threshold maximizes the maximal number of needed favors that are granted, and therefore, it maximizes expected total surplus.

Proposition 3. When $\frac{\delta}{1-\delta} \in \left(\frac{c}{b-c}, \frac{1}{p} \frac{c}{b-c}\right)$, the optimal threshold $\bar{H}^*(M)$ is increasing in M , and for all n , $\bar{H}^*(nM) \geq n\bar{H}^*(M)$. As the breadth of the relationship goes to infinity, $V(\bar{H}^*(M), M)/M \rightarrow p(b-c)$ and

$$\bar{H}^*(M)/M \rightarrow \frac{\delta}{1-\delta} p \frac{b-c}{c} \in (p, 1)$$

Proposition 3 describes how the optimal threshold varies with the breadth of the relationship. It is clear that the optimal threshold is weakly increasing, since total surplus is weakly increasing in the breadth of the relationship for a given threshold \bar{H} . Moreover, Proposition 3 shows that the ratio of the optimal threshold relative to the total number of activities is in general increasing.² It is clear from Theorem 1 that as $M \rightarrow \infty$, the limit of this ratio must exceed p , since full component cooperation can be almost-perfectly replicated. Proposition 3 provides the exact value for the limit. Note that, even if $M \rightarrow \infty$, the discount factor still constrains the fraction of needed favors that are granted in an optimal agreement. Of course, as δ increases, this fraction increases.

4.2 Bertrand Collusion and Multimarket Contact

The second application is to Bertrand collusion between duopolists who interact in multiple markets. When demand fluctuates between high- and low-demand states, Rotemberg and Saloner (1986) show that optimal collusive agreements involve more aggressive price competition when demand is higher. In the same setting, Bernheim and Whinston (1990) demonstrate that contact in two markets with perfectly negatively correlated demand shocks helps sustain collusion in each relative to single-market collusion. Bernheim and Whinston (1987) consider more general correlation structures.

In this section, we consider the same model as in Bernheim and Whinston (1990), but we assume demand shocks are independent across M separate markets. As in our application to favor-exchange

²It is not true that $\bar{H}^*(M)/M$ is always increasing, because of integer problems: there will always exist M such that $\bar{H}^*(M+1) = \bar{H}^*(M)$.

games, we show that optimal agreements can be described by a threshold: firms charge monopoly prices whenever the number of high-demand markets does not exceed a cutoff $\bar{H}^*(M)$; otherwise, firms charge prices yielding total profits of $\pi^*(M)$ across all M markets. We explicitly compute the limit of $\bar{H}^*(M)/M$ as M goes to infinity.

The Bertrand Component Game There are $N = 2$ firms. Both firms produce identical products at zero marginal cost, and they simultaneously choose prices $\tilde{a}_{1,m,t}$ and $\tilde{a}_{2,m,t}$, and $\mathcal{A}_i = [0, \bar{a}]$ for some \bar{a} large. The **market price** is the lowest price chosen by the two firms, and we denote it by $\tilde{a}_{m,t} = \min\{\tilde{a}_{1,m,t}, \tilde{a}_{2,m,t}\}$. Market demand is given by $\tilde{q}(\tilde{a}_{m,t}, \tilde{s}_{m,t})$, where $\tilde{s}_{m,t} \in \mathcal{S} = \{h, l\}$ is a **market demand state** that is either high or low. In each period, market demand is high with probability p and low with probability $1 - p$. We assume that $\tilde{q}(\tilde{a}_{m,t}, l) \leq \tilde{q}(\tilde{a}_{m,t}, h)$ for all $\tilde{a}_{m,t}$. The total market profits of the two firms is therefore $\tilde{\pi}(\tilde{a}_{m,t}, \tilde{s}_{m,t}) = \tilde{a}_{m,t} \tilde{q}(\tilde{a}_{m,t}, \tilde{s}_{m,t})$, which we assume to be continuous in $\tilde{a}_{m,t}$. We also assume that there is a monopoly price, $\tilde{a}^*(\tilde{s})$ that uniquely maximizes total market profits in market demand state $\tilde{s} \in \{h, l\}$. We denote the associated monopoly profits as $\tilde{\pi}^*(\tilde{s}) \equiv \tilde{a}^*(\tilde{s}) q(\tilde{a}^*(\tilde{s}), \tilde{s})$. Notice that $\tilde{\pi}^*(h) \geq \tilde{\pi}^*(l)$.

We assume that the firm with a lower price captures the entire market, and when both firms choose the same price, they split the market demand equally. Firm i 's profits are therefore given by

$$\tilde{u}_i(\tilde{a}_{i,m,t}, \tilde{a}_{-i,m,t}, \tilde{s}_{m,t}) = \begin{cases} \tilde{a}_{i,m,t} \tilde{q}(\tilde{a}_{m,t}, \tilde{s}_{m,t}) & \tilde{a}_{i,m,t} < \tilde{a}_{-i,m,t} \\ 0 & \tilde{a}_{i,m,t} > \tilde{a}_{-i,m,t} \\ \tilde{a}_{i,m,t} \tilde{q}(\tilde{a}_{m,t}, \tilde{s}_{m,t}) / 2 & \tilde{a}_{i,m,t} = \tilde{a}_{-i,m,t} \end{cases}$$

A component agreement is a function $\tilde{A} : \mathcal{S} \rightarrow \mathcal{A}$ specifying each player's price as a function of the market demand state. Note that there is a unique Nash equilibrium of the component stage game in which both players choose $\tilde{a}_{i,m,t} = 0$ for all $\tilde{s}_{m,t}$, and they earn 0, which is their minmax payoff. We will refer to the component agreement \tilde{A}^C in which $\tilde{A}_i^C(\tilde{s}) = \tilde{a}^*(\tilde{s})$ for $\tilde{s} \in \{h, l\}$ as **component perfect collusion**. The next proposition describes necessary and sufficient conditions for component perfect collusion to be an equilibrium component agreement.

Proposition 4. *Component perfect collusion is an equilibrium component agreement if and only if*

$$\tilde{\pi}^*(h) \leq \frac{\delta}{1 - \delta} (p \tilde{\pi}^*(h) + (1 - p) \tilde{\pi}^*(l))$$

Proposition 4 shows that, in this game, deviation gains and continuation losses can be pooled across players, and perfect collusion can be sustained if the total market profits when market demand is high are not too large. As in Rotemberg and Saloner (1986), this is the state in which

firms' deviation gains are highest, so as long as the sum of those deviation gains, $\tilde{\pi}^*(h)$ are not too high, perfect collusion can be sustained.

Multimarket Contact Firms engage in simultaneous Bertrand competition in M separate and independent markets. A demand state is a vector $s_t = (\tilde{s}_{1,t}, \dots, \tilde{s}_{M,t}) \in \mathcal{S}^M$ of market demand states. An agreement $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$ specifies the price level that each firm sets in each market when the demand state is s_t . Denote by $u_i(s_t; A)$ firm i 's associated per-period profits summed up over all M markets in state s_t under agreement A . We refer to the agreement A^C in which $A_{i,m}^C(s_t) = \tilde{a}^*(\tilde{s}_{m,t})$ as **perfect collusion**.

Replicating Component Perfect Collusion When there are M markets, for perfect collusion to be an equilibrium agreement, it has to be the case that for each state s_t , the market price in component game m is equal to the monopoly price associated with state $\tilde{s}_{m,t}$ for all m . As in the favor-exchange application, increasing the number of markets does not affect the conditions required for perfect collusion to be an equilibrium agreement. Under perfect collusion, the maximal deviation gain occurs when the market demand state is high in all markets, and this implies that both the maximal deviation gains and continuation losses are linear in M . This reasoning is similar to the argument given in Bernheim and Whinston (1990), which shows that multimarket contact does not affect the conditions required for sustaining collusion when markets are identical, and there are no demand fluctuations.

While the conditions required for perfect collusion to be an equilibrium agreement do not depend on M , multimarket contact nevertheless expands firms' ability to collude. Under component perfect collusion, the sum of players' mean component deviation gains is $p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)$. It follows from Theorem 1 that as long as $\delta/(1-\delta) > 1$ or $\delta > 1/2$, component perfect collusion can be almost-perfectly replicated.

Optimal Collusive Agreements when $M < \infty$ Our next set of results describe optimal agreements, which maximize firms' joint profits among the set of all equilibrium agreements. To describe an optimal agreement, as in the favor-exchange application, let $P(H, M)$ be the CDF of a binomial distribution with parameters (p, M) , and with a slight abuse of notation, let $p(H, M)$ denote the associated probability mass function. In addition, define $H(s_t)$ to be the total number of markets in which the market demand is high: $H(s_t) = \sum_{m=1}^M 1_{\{\tilde{s}_{m,t}=h\}}$, where 1 is an indicator function.

An agreement A is a **symmetric collusive agreement** if $A_{1,m}(s_t) = A_{2,m}(s_t)$ for all s_t . A

symmetric collusion agreement is an \bar{H} -**threshold collusive agreement** if, whenever $H(s_t) \leq \bar{H}$, $A_{i,m}(s_t) = \tilde{a}^*(\tilde{s}_{m,t})$, and whenever $H(s_t) > \bar{H}$, $A_{i,m}(s_t) = \tilde{a}^*(h)$ in exactly \bar{H} markets for which $\tilde{s}_{m,t} = h$. In an \bar{H} -threshold collusive agreement $A^{\bar{H}}$, firm i 's continuation loss is given by $\delta V_i(A^{\bar{H}}) = \delta V(\bar{H}, M)$, where

$$V(\bar{H}, M) \equiv \frac{\sum_{m=0}^{\bar{H}} (m\tilde{\pi}^*(h) + (\bar{H} - m)\tilde{\pi}^*(l)) p(\bar{H}, M)}{1 - \delta - \delta(1 - P(\bar{H}, M))}.$$

The next proposition characterizes optimal agreements in this setting.

Proposition 5. When $\frac{\delta}{1-\delta} \in \left(1, \frac{\tilde{\pi}^*(h)}{p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)}\right)$, any optimal agreement is an $\bar{H}^*(M)$ -threshold collusive agreement, where $\bar{H}^*(M)$ is the largest number \bar{H} satisfying $\bar{H}\tilde{\pi}^*(h) + (M - \bar{H})\tilde{\pi}^*(l) \leq \delta V(\bar{H}, M)$. Moreover, for all s_t satisfying $H(s_t) > \bar{H}^*(M)$,

$$\sum_{m=1}^M \left[\tilde{u}_1 \left(A_{(m)}^{\bar{H}^*(M)}(s_t), \tilde{s}_{m,t} \right) + \tilde{u}_2 \left(A_{(m)}^{\bar{H}^*(M)}(s_t), \tilde{s}_{m,t} \right) \right] = \delta V(\bar{H}^*(M), M).$$

The results of Proposition 5 mirror those of Proposition 2. An optimal agreement is a threshold agreement: the market price is equal to the monopoly price in each market as long as the total number of high-demand markets is less than $\bar{H}^*(M)$. When the number of high-demand markets exceeds $\bar{H}^*(M)$, the total profits across all markets is exactly equal to $\delta V(\bar{H}^*(M), M)$.

Proposition 6. When $\frac{\delta}{1-\delta} \in \left(1, \frac{\tilde{\pi}^*(h)}{p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)}\right)$, the optimal threshold $\bar{H}^*(M)$ is increasing in M , and for all n , $\bar{H}^*(nM) \geq n\bar{H}^*(M)$. As the number of markets goes to infinity, $V(\bar{H}^*(M), M)/M \rightarrow p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)$ and

$$\bar{H}^*(M)/M \rightarrow \frac{\delta(p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)) - (1-\delta)\tilde{\pi}^*(l)}{(1-\delta)(\tilde{\pi}^*(h) - \tilde{\pi}^*(l))} \in (p, 1).$$

Proposition 6 describes how the optimal threshold varies with the number of markets. As in the favor-exchange model, the optimal threshold is weakly increasing, and the limit is strictly increasing in the discount factor.

5 Imperfect Observability with Applications to Relational Contracts

We now show that the limit results in Theorem 1 also extend to a class of extensive-form stage games with imperfectly observable actions. In Section 5.1, we establish a limit result that corresponds to the results in Theorem 1. Sections 5.2 and 5.3 examine optimal agreements in two models of relational incentive contracts: a model with multiple activities and a single agent and a model with

multiple agents, each of whom performs a single activity. Section 5.2 involves a single principal and a single agent who chooses an unobserved effort level in each of M independent tasks, a setting similar to the settings in Laux (2001) and Bond and Gomes (2009). The analysis in Section 5.3 involves a single principal and M agents who each choose a single unobserved effort level, as in the model of Levin (2002).

5.1 Replication in Games with Imperfect Observability

This section describes limit results for a class of extensive-form games with imperfect observability. The class of games includes the application in Section 5.2. At the end of this section, we comment on how the limit results could similarly be extended to the class of games in Section 5.3.

There are N players. In each component stage game, each player i chooses a **private action** $\tilde{a}_{i,m,t} \in \mathcal{A}_i$ that is unobserved by others and a **public action** $\tilde{b}_{i,m,t} \in \mathcal{B}_i$ that is commonly observed. Player i 's private action determines the distribution over his state, $\tilde{s}_{i,m,t} \in \mathcal{S}_i$, which is distributed according to distribution function $G(\cdot | \tilde{a}_{i,m,t})$, and each player's state is commonly observed. Denote $\mathcal{A} = \prod_{i=1}^N \mathcal{A}_i$, $\mathcal{B} = \prod_{i=1}^N \mathcal{B}_i$, and $\mathcal{S} = \prod_{i=1}^N \mathcal{S}_i$ with generic element $\tilde{s}_{m,t}$. The timing of the component stage game is: (1) players simultaneously choose private actions, (2) players' states are commonly observed, (3) players simultaneously choose public actions. Payoffs in each component stage game are additively separable across the two stages and can be written as $\tilde{u}_i(\tilde{b}_{i,m,t}, \tilde{b}_{-i,m,t}, \tilde{s}_{m,t}) - \tilde{c}_i(\tilde{a}_{i,m,t})$. For player i , the private actions of the other players only affect his payoff inasmuch as they determine the distribution over the state $\tilde{s}_{m,t}$.

A component agreement specifies a set of private actions and a set of state-contingent public actions for each player within a component game, that is, it is a pair $\tilde{\alpha} = (\tilde{A}, \tilde{B})$, where $\tilde{A} \in \mathcal{A}$ and $\tilde{B} : \mathcal{S} \rightarrow \mathcal{B}$. We assume that there is a unique SPNE of the component stage game, and it gives each player his minmax payoff, which we normalize to 0. We will refer to this SPNE as a **punishment component agreement**. This assumption is satisfied by our applications but is not essential for our limit results. A **public trigger strategy of a component supergame** specifies a component agreement $\tilde{\alpha}(0)$ and a punishment component agreement $\tilde{\alpha}(1)$ such that play begins with $\tilde{\alpha}(0)$, and if any publicly observable deviations occur, play transitions to $\tilde{\alpha}(1)$ for all future periods. Note that any deviation in public actions is publicly observable, and if $\tilde{\alpha}(0)$ specifies private actions such that $G(\cdot | \tilde{a}_{i,m,t})$ has full support, then deviations in public actions are the only publicly observable deviations. We say that a component agreement $\tilde{\alpha}^*$ is an **equilibrium component agreement** if there exists a public trigger strategy of the component supergame with $\tilde{\alpha}(0) = \tilde{\alpha}^*$ that is a Perfect Public Equilibrium (PPE) of the component supergame.

An agreement, then, is a pair $\alpha = (A, B)$, where $A \in \mathcal{A}^M$ and $B : \mathcal{S}^M \rightarrow \mathcal{B}^M$. A **punishment agreement**, a **public trigger strategy of the supergame**, and a **public equilibrium agreement** are defined similarly. Accordingly, we can define an ε -neighborhood of component agreement $\tilde{\alpha}$ and therefore what it means for $\tilde{\alpha}$ to be almost-perfectly replicated.

Unlike the game with perfect observability, the players might deviate on the private actions. For an agreement $\tilde{\alpha}$, let $\tilde{d}_i^{priv}(\tilde{\alpha})$ denote the maximal net gain player i can achieve by deviating to another private action. Denote $\tilde{d}_i^{pub}(\tilde{s}_{m,t}; \tilde{\alpha})$ as the maximal gain player i can achieve by deviating to another public action in state $\tilde{s}_{m,t}$, and $\tilde{d}_i^{pub,mean}(\tilde{\alpha})$ as the expected value of $\tilde{d}_i^{pub}(\tilde{s}_{m,t}; \tilde{\alpha})$. As in our main model, we denote $\delta\tilde{V}_i(\tilde{\alpha})$ to be player i 's continuation loss following a deviation in his public action.

In this setting, results analogous to those in Theorem 1 hold. In particular, consider a component agreement $\tilde{\alpha}^*$ with $\tilde{d}_i^{priv}(\tilde{\alpha}^*) < 0$. Then, (1) if $\tilde{d}_i^{pub,mean}(\tilde{\alpha}^*) < \delta\tilde{V}_i(\tilde{\alpha}^*)$ for all i , $\tilde{\alpha}^*$ can be almost-perfectly replicated and (2) if $\tilde{d}_i^{pub,mean}(\tilde{\alpha}^*) > \delta\tilde{V}_i(\tilde{\alpha}^*)$ for some i , then $\tilde{\alpha}^*$ cannot be almost-perfectly replicated.

These results follow from the same logic as in the proof of Theorem 1. For Part 1, we can construct a sequence of agreements that specify the same private actions as in $\tilde{\alpha}$ but that “chop off” the states in which players have large public deviation gains and replaces the specified public actions by, say, static Nash equilibrium play in the second stage of the game. Again, by the law of large numbers, the probability of the set of states that are “chopped off” goes to zero, and therefore $\tilde{\alpha}$ can be almost-perfectly replicated. The only additional consideration is that under this sequence of agreements, players might also want to deviate in terms of their private actions. This possibility is ruled out by our assumption that players have a strict incentive not to deviate in private actions under $\tilde{\alpha}$. For Part 2, when $\tilde{d}_i^{pub,mean}(\tilde{\alpha}^*) > \delta\tilde{V}_i(\tilde{\alpha}^*)$, for any agreement that approximates $\tilde{\alpha}^*$, it is again impossible to allocate public deviation gains across states in a way that guarantees they are smaller than the continuation losses in every state.

In the component game in the model of Section 5.3, there is a principal and agent m . Agent m chooses a private action $\tilde{a}_{m,t} \in \mathcal{A}$, which determines the distribution over the state $\tilde{s}_{m,t} \in \mathcal{S}$. The state is observed, and then the principal chooses a public action $\tilde{b}_m \in \mathcal{B}$. A component agreement, then, is a private action by agent m , and a state-contingent public action by the principal. In this model, the number of component games is equal to the number of agents, and so the question is whether a component agreement with a single agent can be almost-perfectly replicated as the number of agents grows large. In this setting, the results of Theorem 1 can again be extended in a similar fashion. In particular, the condition for a component agreement to be almost-perfectly

replicated is again that the mean public deviation gain is smaller than the continuation loss for each player.

5.2 Relational Incentive Contracts with Multiple Activities

In this section, we consider optimal relational contracts between a risk-neutral principal and a risk-neutral agent who chooses unobserved binary effort in M activities. This application is similar to the models of Laux (2001) and Bond and Gomes (2009), which analyze optimal formal contracts in this setting when there are exogenous bounds on payments. The key complication in this setting is that local incentive constraints are not sufficient: in a given agreement, even if the agent does not want to deviate by choosing a different effort in a single activity, he may want to deviate and choose a different effort level in many activities simultaneously. As a result, non-local constraints may—and in fact do—bind.

Relational Incentive Contracts Component Game There are $N = 2$ players. Player 1 is a risk-neutral principal, and player 2 is a risk-neutral agent. At the beginning of each period, the agent chooses an effort level $\tilde{a}_{2,m,t} \in \mathcal{A}_2 = \{0, 1\}$ at cost $c\tilde{a}_{2,m,t}$. The agent's effort determines the distribution over output $\tilde{s}_{m,t} \in \mathcal{S} = \{0, 1\}$, which accrues to the principal. The probability of high output is $\Pr[\tilde{s}_{m,t} = 1 | \tilde{a}_{2,m,t}] = p\tilde{a}_{2,m,t}$, where $p < 1$. The principal then chooses whether to pay the agent a bonus $\tilde{b}_{1,m,t} \in \mathcal{B} = [0, \bar{b}]$ for \bar{b} large. The principal's payoffs are therefore $\tilde{s}_{m,t} - \tilde{b}_{1,m,t}$, and the agent's payoffs are $\tilde{b}_{1,m,t} - c\tilde{a}_{2,m,t}$.

A **component relational contract** is a component agreement $\tilde{\alpha} = (\tilde{A}, \tilde{B})$, where $\tilde{A} \in \mathcal{A}$ and $\tilde{B} : \mathcal{S} \rightarrow \mathcal{B}$ specifying an effort level $\tilde{A}_2 = \tilde{e}$, and a bonus payment as a function of the realization of output. Without loss of generality, we can restrict attention to relational contracts that pay $\tilde{B}_1(0) = 0$ and $\tilde{B}_1(1) = b$ for some $b \geq 0$. In the component stage game, there is a subgame-perfect equilibrium in which the agent chooses $\tilde{a}_{2,m,t} = 0$, the principal chooses $\tilde{b}_{1,m,t} = 0$, and both players earn their minmax payoff of 0. We will refer to a component relational contract $\tilde{\alpha}$ specifying $\tilde{e} = 1$ as an **component effort-inducing relational contract**.

Under an effort-inducing component relational contract, the agent will choose $e = 1$ if and only if $pb \geq c$. The principal is willing to pay the bonus b rather than 0 when $\tilde{s}_{m,t} = 1$ if $b \leq \delta(p - pb) / (1 - \delta)$, where the right-hand side is the discounted expected payoff for the principal under an effort-inducing component relational contract. Such a bonus level b exists, and therefore an effort-inducing component relational contract is an equilibrium component agreement if and

only if

$$\frac{c}{p} \leq \frac{\delta}{1-\delta} (p-c).$$

If an effort-inducing component relational contract $\tilde{\alpha}$ specifies b , then the principal's deviation gain $\tilde{d}_1^{pub}(1; \tilde{\alpha}) = b$ and $\tilde{d}_1^{pub}(0; \tilde{\alpha}) = 0$, so her maximal deviation gain is $\tilde{d}_1^{pub, \max}(\tilde{\alpha}) = b$ and her mean deviation gain is $\tilde{d}_1^{pub, mean}(\tilde{\alpha}) = pb$. The agent's maximal net gain is $\tilde{d}_2^{priv}(\tilde{\alpha}) = c - pb$.

Multi-Activity Relational Contract In each period, production consists of M independent activities. A **relational contract** is an agreement $\alpha = (A, B)$, where $A \in \mathcal{A}$ and $B : \mathcal{S}^M \rightarrow \mathcal{B}^M$ specifying an effort level $A_{2,m} = \tilde{e}_m$ in each activity and a bonus level $B_m(s_t)$, which depends on the vector of realized outputs $s_t = (\tilde{s}_{1,t}, \dots, \tilde{s}_{m,t})$. An **effort-inducing relational contract** specifies $\tilde{e}_m = 1$ for all m . Among the set of all effort-inducing relational contracts that are equilibrium agreements, we will say that an **optimal relational contract** is the one that minimizes the principal's maximal deviation gain. In other words, it is an equilibrium effort-inducing relational contract for the largest range of discount factors. For the purposes of characterizing optimal relational contracts, we can without loss of generality focus on agreements that pay a bonus that depends only on the number of activities in which output was high: $b_{\hat{m}} \equiv \sum_{m=1}^M B_m(s_t)$ where $\sum_{m=1}^M \tilde{s}_{m,t} = \hat{m}$. We refer to $b_{\hat{m}}$ as a **bonus scheme**.

Replicating Component Effort-Inducing Relational Contracts When there are M activities, a component effort-inducing relational contract can be M -perfectly replicated if and only if $pb \geq c$ and $b \leq \delta(p - pb) / (1 - \delta)$. As in the applications in the previous section, increasing the number of activities does not affect the conditions for M -perfect replication. Next, we characterize the condition required for almost-perfect replication.

Proposition 7. *If $pb > c$ and $pb < \delta(p - pb) / (1 - \delta)$, then a component effort-inducing relational contract can be almost-perfectly replicated.*

Combining the conditions in Proposition 7, we see that as long as $c < \delta(p - c) / (1 - \delta)$, there exists an effort-inducing agreement that can be almost-perfectly replicated. This implies that increasing the breadth of interaction helps sustain effort.

Optimal Relational Contracts As described above, an effort-inducing component relational contract is an equilibrium agreement and can therefore be M -perfectly replicated if and only if $c/p \leq \delta(p - c) / (1 - \delta)$. If, however, $c/p > \delta(p - c) / (1 - \delta) > c$, then an effort-inducing component relational contract is M -perfectly replicable, but it is almost-perfectly replicable. The bonus scheme

constructed in the proof of Proposition 7 in general is suboptimal, that is, it does not minimize the principal's maximal deviation gain. The next proposition describes the bonus scheme that does so when $M < \infty$.

Proposition 8. If $\frac{c}{p} > \frac{\delta}{1-\delta}(p-c) > c$, then an optimal effort-inducing relational contract has the following bonus scheme:

$$b_m = \begin{cases} 0 & \sum_{m=1}^M \tilde{s}_{m,t} < m^*(M) \\ \gamma & \sum_{m=1}^M \tilde{s}_{m,t} = m^*(M) \\ \beta & \sum_{m=1}^M \tilde{s}_{m,t} > m^*(M) \end{cases}$$

for some integer $m^*(M)$, where $0 \leq \gamma \leq \beta$.

The intuition for the proof is as follows. The problem of finding the bonus scheme that implements effort in each activity with the smallest maximal bonus is relatively complicated. This is because there are M potential deviations by the agent, since he can choose to exert effort in L activities for any $0 \leq L < M$. It suffices, however, to ignore all but two of the agent's incentive-compatibility constraints: the local constraint that ensures he does not prefer to choose effort in $L = M - 1$ activities and the global constraint that ensures he does not prefer to choose effort in $L = 0$ activities. Showing that the solution to this relaxed problem is also a solution to the full problem is non-trivial and is similar to the analysis of Bond and Gomes (2009).

If the global constraint did not need to be satisfied, the principal would optimally choose a threshold bonus scheme in which $b_m = 0$ if output is sufficiently low, and $b_m = \beta$ if output is sufficiently high. The threshold \hat{m} would be chosen such that the bonus β that implements effort given this threshold rule is minimized, as this would be the scheme with the smallest maximal bonus that is capable of satisfying the agent's local constraint.

However, this bonus scheme does not satisfy the agent's global constraint, as it also minimizes the agent's interim rents. In order to satisfy the agent's global constraint, the optimal contract chooses a lower threshold $m^* < \hat{m}$ such that if m^* were larger, the global constraint would be slack. At this value of m^* , interim rents are then adjusted downward at the cutoff value by decreasing the amount paid at the cutoff value to $\gamma < \beta$ until the global constraint is exactly binding.

That the form of the optimal bonus scheme in this setting is not linear in realized output implies that the optimal relational contract pools together performance in the M independent tasks. Slack is transferred from states of the world in which the linear contract calls for large total bonus payments to states of the world in which it does not. Since the agent is motivated by his expected bonus, his incentives remain in place.

5.3 Relational Incentive Contracts with Multiple Agents

Our final application is to a model of relational contracts with multiple agents and imperfect public monitoring of output but perfect public monitoring of bonus payments, similar to the model of Levin (2002). In the optimal relational contract in this setting, the principal pays out a fixed bonus pool as long as one or more agents produces high output. The bonus in a given period is shared equally by all agents who produce high output in that period. We study how the number of agents affects the conditions required for an effort-inducing relational contract to be an equilibrium agreement.

Single-Agent Component Game There are two players, player 0 and, for reasons that will become clear, player m . Player 0 is a risk-neutral principal, and player m is a risk-neutral agent. The component game is essentially the same as in the previous subsection, except that high output is not necessarily a perfect signal of effort. That is, following an effort choice $\tilde{a}_{m,t} \in \{0, 1\}$, by the agent, the probability that output is high is p when $\tilde{a}_{m,t} = 1$ and $r < p$ when $\tilde{a}_{m,t} = 0$, so that $\Pr[\tilde{s}_{m,t} = 1 | \tilde{a}_{m,t}] = p\tilde{a}_{m,t} + r(1 - \tilde{a}_{m,t})$. A component relational contract $\tilde{\alpha}$ specifies a wage payment \tilde{w}_m made to the agent, an effort choice \tilde{e}_m by the agent, and a bonus payment $\tilde{B}_0(\tilde{s}_{m,t})$ made to the agent depending on the realization of output. As above, we can restrict attention, without loss of generality, to relational contracts that pay $\tilde{B}_0(0) = 0$ and $\tilde{B}_0(1) = b$ for some $b \geq 0$. Throughout, we will assume that $p - c > r$, so that high effort increases total surplus, and we will refer to a component relational $\tilde{\alpha}$ specifying $\tilde{e}_m = 1$ as an component effort-inducing relational contract. The base wage can be chosen so that the principal captures all the surplus, so we can write the principal's expected surplus as $\tilde{v}_0 = p - c$.

For an effort-inducing component relational contract to be an equilibrium component agreement, the bonus level b must be high enough that the agent prefers to choose $\tilde{e}_m = 1$ rather than $\tilde{e}_m = 0$. That is, we must have $b \geq c/(p - r)$. To ensure the principal is willing to pay the bonus, the bonus must be smaller than her expected future surplus: $b \leq \delta\tilde{v}_0/(1 - \delta)$. It is well-known that a necessary and sufficient condition for an effort-inducing component relational contract to be an equilibrium component agreement is given by adding these two inequalities together so that $c/(p - r) \leq \delta\tilde{v}_0/(1 - \delta)$, or equivalently

$$\frac{\delta}{1 - \delta} \geq \frac{c/(p - r)}{\tilde{v}_0}.$$

Note that the right-hand side is the ratio between the principal's maximal deviation gain and the surplus of the relationship.

Multilateral Relational Contracts Now, suppose the principal (player 0) interacts with M agents (players $1, \dots, M$) simultaneously. Again, assume that for each agent m , his output in period t is given by $\Pr[\tilde{s}_{m,t} = 1 | \tilde{a}_{m,t}] = p\tilde{a}_{m,t} + r(1 - \tilde{a}_{m,t})$, and his cost of effort is c . Agents' outputs are independent. Levin (2002) shows that in this setting, an optimal relational contract specifies a bonus pool B_M to be shared by agents who produce high output. No bonuses are paid out when all outputs are low. In addition, the principal sets the wage payments w_m to extract all the surplus.

An effort-inducing relational contract is one that specifies $\tilde{e}_m = 1$ for all m , and we now describe the conditions required for an effort-inducing relational contract to be an equilibrium agreement. To motivate each agent m to choose high effort, his expected bonus for high output must exceed $c/(p - r)$, or, for all m ,

$$\sum_{k=1}^{M-1} \frac{B_M}{k+1} \Pr \left[\sum_{\hat{m} \neq m} \tilde{s}_{\hat{m},t} = k \right] \geq \frac{c}{p-r}.$$

The left-hand side is the expected bonus of agent m conditional on having a high output. To ensure the principal is willing to pay the bonus, the condition is

$$B_M \leq \frac{\delta}{1-\delta} M \tilde{v}_0,$$

where the right-hand side is the principal's future surplus. As above, the necessary and sufficient conditions for an effort-inducing relational contract to be an equilibrium agreement is that the sum of these conditions is satisfied:

$$\frac{\delta}{1-\delta} M \tilde{v}_0 \geq \frac{c/(p-r)}{\sum_{k=1}^{M-1} \frac{1}{k+1} \Pr \left[\sum_{\hat{m} \neq m} \tilde{s}_{\hat{m},t} = k \right]}. \quad (3)$$

The next proposition simplifies this expression and makes clear how this condition varies with the number of agents, M .

Proposition 9. *An effort-inducing relational contract is an equilibrium agreement if and only if*

$$\frac{\delta}{1-\delta} \geq \frac{pc/(p-r)}{\left(1 - (1-p)^M\right) \tilde{v}_0}.$$

The condition in Proposition 9 is easier to satisfy for larger M : there are increasing returns to scale in sustaining effort in multilateral relational contracts. The reason why effort is easier to sustain in multilateral relational contracts, rather than in bilateral relational contracts is again because it allows for cross-subsidization of incentive constraints. When there are M agents, slack

can be transferred from the output realizations in which the principal's deviation gain is small to output realizations in which her deviation gain is large while maintaining each agent's incentives to exert effort. This proposition shows that this cross-subsidization can be made more effective when the number of agents grows. In particular, when $M \rightarrow \infty$, the condition required for an effort-inducing relational contract to be an equilibrium agreement becomes

$$\frac{\delta}{1 - \delta} > \frac{pc/(p - r)}{\tilde{v}_0},$$

where notice that $pc/(p - r)$ is the average bonus paid in an effort-inducing bilateral relational contract. In larger firms, therefore, cooperation is limited not by the magnitude of the maximal bonus that needs to be paid out to a particular agent, but rather by the magnitude of the mean bonus paid out to all of the agents.

6 Conclusion

This paper explores the degree of efficiency gains that can be achieved by increasing the number of activities that players engage in an infinitely repeated game. We study a repeated game that is composed of M identical and independent component games. For a component agreement to be an equilibrium component agreement, the maximal deviation gain for each player must be smaller than the future surplus he can earn by adhering to the component agreement. As M tends toward infinity, the component agreement can be almost-perfectly replicated as long as the mean deviation gain is smaller for each player than his future surplus. This is a considerably weaker condition, especially in very volatile environments.

We apply this result to three applications: favor exchange, multi-market contact, and multilateral relational contracts. In all three applications, we characterize the optimal agreement in games of finite breadth. A common feature among all applications is that the optimal agreement takes a threshold form in the sense that a player will cooperate as much as possible as long as the deviation gain does not exceed an endogenous threshold.

Our results can be taken in several directions. For example, in a model with endogenous governance structures (such as firm boundaries), the number of activities the firm engages in affects the optimal governance structure. Baker, Gibbons, and Murphy (2011) takes the view that governance structures should be designed to support cooperation by minimizing the maximal deviation gain. Our analysis suggests that as the number of activities increases, the relevant deviation gain is the mean, rather than the maximal, deviation gain. In other words, optimal governance structures can change as businesses expand the breadth of their activities.

Another potential application is to organization formation. Our results suggest that organizations might arise precisely because cooperation is easier to sustain with a larger number of players. Our results further suggest that the types of individuals who should be included in an organization depends on the scale of the organization. Smaller-scale organizations should pay more attention to maximal deviation gains, and therefore should hire individuals whose performance is predictable. Larger-scale organizations should focus on mean deviation gains, and they can therefore tolerate individuals whose performance is more variable.

Appendix

The appendix contains five sections. We first establish some preliminary results in section A. The proof of our limit results are in Section B, and the proofs for the propositions related to our applications are in sections C, D, and E.

6.1 A. Preliminary Results

Lemma A1. $\underline{v}_i = M\tilde{v}_i$.

Proof of Lemma A1. It is clear that $\underline{v}_i \leq M\tilde{v}_i$, since any Nash equilibrium of the component stage game is M -perfectly replicable for any M . Next, since payoffs are additively separable across component games, if A is a Nash equilibrium of the stage game, then $A_{(m)}$ is a Nash equilibrium of the component stage game for each m . Otherwise, there exists a component game m' in which a profitable deviation is available for some player i . For this player, it is therefore a profitable deviation for him to choose an action vector that uses the profitable deviation in m' and uses the action specified in $A_{(m)}$ for all $m \neq m'$. Since $A_{(m)}$ is a Nash equilibrium of the component stage game, it follows that each player i receives at least \tilde{v}_i in that component game. Therefore, $\underline{v}_i \geq M\tilde{v}_i$. ■

Lemma A2. Let X and Y be two real-valued random variables. For any $k \in \mathbb{R}$, $\Pr[X > k] \geq \Pr[Y > k] - \Pr[X \neq Y]$.

Proof of Lemma A2. Define the events $A = \{X > k\}$, $B = \{Y > k\}$, and $C = \{X = Y\}$. Since $A^C \cap C^C = \{X = Y > k\}$, $A^C \cap C^C \subset B^C$, so $\Pr[A^C \cap C^C] \leq \Pr[B^C]$, which is equivalent to $\Pr[A \cup C] \geq \Pr[B]$. Finally, since $\Pr[A] + \Pr[C] \geq \Pr[A \cup C]$, the result follows. ■

B. Limit Result

Theorem 1. Consider a component agreement \tilde{A}^* . The following are true:

1. If $\tilde{d}_i^{mean}(\tilde{A}^*) < \delta\tilde{V}_i(\tilde{A}^*)$ for all i , then \tilde{A}^* can be almost-perfectly replicated.
2. If $\tilde{d}_i^{mean}(\tilde{A}^*) > \delta\tilde{V}_i(\tilde{A}^*)$ for some i , then \tilde{A}^* cannot be almost-perfectly replicated.

Proof of Theorem 1. We prove Part 1 first. Suppose $\tilde{d}_i^{mean}(\tilde{A}^*) < \delta\tilde{V}_i(\tilde{A}^*)$ for all i . Now, for any $\varepsilon_1 > 0$, there exists $\varepsilon \in (0, \varepsilon_1)$ such that

$$(1 + \varepsilon)\tilde{d}_i^{mean}(\tilde{A}^*) < \delta\left(\tilde{V}_i(\tilde{A}^*) - \varepsilon K/(1 - \delta)\right) \text{ for all } i,$$

where recall that B is the upper bound of the component-stage-game payoff. Define the event \mathcal{E} as the set of states for which the sample average deviation gains for i are below $\tilde{d}_i^{mean}(\tilde{A}^*)(1 + \varepsilon)$ for all i :

$$\mathcal{E} = \left\{ s_t : \frac{1}{M} \sum_{m=1}^M \tilde{d}_i(\tilde{s}_{m,t}; \tilde{A}^*) < \tilde{d}_i^{mean}(\tilde{A}^*)(1 + \varepsilon) \text{ for all } i \right\}.$$

By the weak law of large numbers, there exists $M(\varepsilon)$ such that for all $M \geq M(\varepsilon)$, $\Pr[\mathcal{E}] > 1 - \varepsilon$.

Now, for all $M \geq M(\varepsilon)$, consider the following agreement A^M :

$$A_{i,m}^M(s_t) = \begin{cases} \tilde{A}_i(\tilde{s}_{m,t}) & \text{if } s_t \in \mathcal{E} \\ \hat{a}_{i,m} & \text{if } s_t \notin \mathcal{E}, \end{cases}$$

where for some stage-game Nash equilibrium \hat{A} , $\hat{A}_{i,m}(s_t) = \hat{a}_{i,m}$. Consider the trigger strategy with $A^M(0) = A$ and each $A^M(i)$ being the Nash equilibrium of the stage game yielding \underline{v}_i for player i . Let u_i^M be the expected per-period payoff of player i under this trigger strategy and $\delta V_i(A^M) = \delta(v_i(A^M) - \underline{v}_i)/(1 - \delta)$ the corresponding continuation loss. Notice that $V_i^M \geq M(\tilde{V}_i(\tilde{A}^*) - \varepsilon K/(1 - \delta))$.

Now, for any $s_t \notin \mathcal{E}$, it is clear that no player has an incentive to deviate from $A^M(0)$. For $s_t \in \mathcal{E}$, player i 's gain from deviating is smaller than $M\tilde{d}_i^{mean}(\tilde{A}^*)(1 + \varepsilon)$ by the definition of \mathcal{E} . Player i 's continuation loss is given by $\delta V_i(A^M)$. Since

$$M\tilde{d}_i^{mean}(\tilde{A}^*)(1 + \varepsilon) \leq \delta M\left(\tilde{V}_i(\tilde{A}^*) - \frac{\varepsilon K}{1 - \delta}\right) \leq \delta V_i(A^M),$$

player i will not deviate. The agreement A^M is therefore an equilibrium agreement. In addition, $\Pr[A_{(m)}^M(s_t) = \tilde{A}(\tilde{s}_{m,t}) \text{ for all } m] \geq 1 - \varepsilon \geq 1 - \varepsilon_1$ by construction. This proves Part 1.

We next turn to Part 2. Without loss of generality, we may assume that $\tilde{d}_1^{mean}(\tilde{A}^*) > \delta\tilde{V}_1(\tilde{A}^*)$.

Suppose to the contrary that \tilde{A}^* can be almost-perfectly replicated. For any $\varepsilon > 0$, there exists an equilibrium agreement A^M and a set $\mathcal{F} \subset \mathcal{S}^M$ such that (1) $\Pr[\mathcal{F}] > 1 - \varepsilon$ and (2) for all $s_t \in \mathcal{F}$, $A_{1,m}^M(s_t) = \tilde{A}_1(\tilde{s}_{m,t})$ for all m .

Let u_1^M be player 1's per-period expected payoff associated with A^M and $\delta V_1(A^M)$ his continuation loss. Since component-stage-game payoffs are bounded from above by K , it follows that $V_1(A^M) \leq M(\tilde{V}_1(\tilde{A}^*) + \varepsilon K)$. Now, for A^M to be an equilibrium agreement, a necessary condition is that player 1 cannot benefit from deviating when $s_t \in \mathcal{F}$:

$$d_1(s_t; A^M) \leq \delta V_1(A^M) \leq \delta M(\tilde{V}_1(\tilde{A}^*) + \varepsilon K) \text{ for all } s_t \in \mathcal{F}.$$

It follows that

$$\sum_{s_t} \Pr[s_t] d_1(s_t; A^M) \leq \delta M(\tilde{V}_1(\tilde{A}^*) + \varepsilon K) \text{ for all } s_t \in \mathcal{F}.$$

Now, notice that

$$\begin{aligned} M\tilde{d}_1^{mean}(\tilde{A}^*) &= \sum_{s_t} \Pr[s_t] d_1(s_t; A^M) \\ &\leq \delta M(\tilde{V}_1(\tilde{A}^*) + \varepsilon K) \Pr[\mathcal{F}] + MK(1 - \Pr[\mathcal{F}]) \\ &= \delta M\tilde{V}_1(\tilde{A}^*) + \delta MK(1 - \varepsilon)\varepsilon + MK\varepsilon. \end{aligned}$$

Since this is true for all $\varepsilon > 0$, it follows that $\tilde{d}_1^{mean}(\tilde{A}^*) \leq \delta\tilde{V}_1(\tilde{A}^*)$, contradicting the original claim. This proves Part 2. ■

Proposition 1. *Consider a component agreement \tilde{A} . The following are true:*

1. *If $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for all i , then \tilde{A} can be M -perfectly replicated for all M if and only if $\tilde{d}_i(\cdot; \tilde{A})$ is constant.*
2. *If $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for all i , $\tilde{d}_i(\cdot; \tilde{A})$ is not constant for some i , and \tilde{A} is Pareto-optimal in the component stage game, then \tilde{A} cannot be almost-perfectly replicated.*
3. *If $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$ for any i , and \tilde{A} can be almost-perfectly replicated, then $\lim_{\varepsilon \rightarrow 0} \varepsilon \sqrt{\hat{M}(\varepsilon)} = \infty$.*

Proof of Proposition 1. For Part 1, let $\bar{A} \in \mathcal{N}^M(\tilde{A}, 0)$. Since $\tilde{d}_i(\cdot; \tilde{A})$ is constant, $d_i(s_t; \bar{A}) = M\tilde{d}_i^{mean}(\tilde{A}) = \delta M\tilde{V}_i(\tilde{A}) = \delta V_i(\bar{A})$ for all $s_t \in \mathcal{S}^M$. The agreement \bar{A} is therefore an equilibrium agreement for any M , so \tilde{A} can be M -perfectly replicated for all M . Moreover, if $\tilde{d}_i(\cdot; \tilde{A})$ is not constant, then there exists some state s_t and some i for which $d_i(s_t; \bar{A}) > M\tilde{d}_i^{mean}(\tilde{A}) = \delta V_i(\bar{A})$, which implies that \bar{A} is not an equilibrium agreement.

For Parts 2 and 3, let $\bar{A} \in \mathcal{N}^M(\tilde{A}, 0)$. By Lemma A2, for any agreement A and for any $k > 0$,

$$\Pr[d_i(s_t; A) > k] \geq \Pr[d_i(s_t; \bar{A}) > k] - \Pr[d_i(s_t; A) \neq d_i(s_t; \bar{A})].$$

For any $A \in \mathcal{N}^M(\tilde{A}, \varepsilon)$, if the left-hand side of the inequality is strictly positive for $k = \delta V_i(A)$, then A is not an equilibrium agreement. For any $A \in \mathcal{N}^M(\tilde{A}, \varepsilon)$, we necessarily have that $V_i(A) \leq V_i(\bar{A}) + \varepsilon MK$ and $\Pr[d_i(s_t; A) \neq d_i(s_t; \bar{A})] \leq \varepsilon$. A necessary condition for the above inequality to hold with $k = \delta V_i(A)$ is therefore

$$\Pr[d_i(s_t; A) > \delta V_i(A)] \geq \Pr[d_i(s_t; \bar{A})/M - \delta\tilde{V}_i(\tilde{A}) > \delta K\varepsilon] - \varepsilon. \quad (4)$$

For Part 2, note that for all i , $d_i(s_t; \bar{A})/M$ converges to $\tilde{d}_i^{mean}(\tilde{A}) = \delta\tilde{V}_i(\tilde{A})$. By the central limit theorem, the right-hand side of (4) converges to $1 - \Phi(\delta K\varepsilon) - \varepsilon$ as $M \rightarrow \infty$, where Φ is the cdf for a normal random variable with mean zero and variance $\sigma^2 = \text{Var}(\tilde{d}_i(\tilde{s}_{m,t}, \tilde{A}))$. Taking ε to zero, the right-hand side approaches 1/2, which implies that \tilde{A} cannot be almost-perfectly replicated.

For Part 3, if \tilde{A} can be almost-perfectly replicated, then as $\varepsilon \rightarrow 0$ and $M \rightarrow \infty$, the left-hand side of (4) has to go to zero. It is therefore necessary that the right-hand side also goes to zero as

$\varepsilon \rightarrow 0$, so that we must have

$$\Pr \left[\frac{d_i(s_t; \bar{A}) - \delta \hat{M}(\varepsilon) \tilde{V}_i(\bar{A})}{\sqrt{\hat{M}(\varepsilon)}} > \delta K \sqrt{\hat{M}(\varepsilon)} \varepsilon \right] \rightarrow 0.$$

This implies that $\lim_{\varepsilon \rightarrow 0} \varepsilon \sqrt{\hat{M}(\varepsilon)} = +\infty$. ■

C. Favor Exchange

Proposition 2. *When $\frac{\delta}{1-\delta} \in \left(\frac{c}{b-c}, \frac{1}{p} \frac{c}{b-c}\right)$ any optimal agreement is an $\bar{H}^*(M)$ -threshold cooperation agreement, where $\bar{H}^*(M)$ is the largest integer satisfying $\bar{H}^*(M)c \leq \delta V(\bar{H}^*(M), M)$.*

Proof of Proposition 2. To prove Proposition two, we first prove two preliminary results. First, we show that an optimal agreement must give both players the same expected payoff. Second, we show that if there is an optimal agreement, there is a symmetric optimal agreement.

Step 1. Any optimal equilibrium must give both players the same expected payoff. Take an arbitrary agreement A . Recall that $f_i(s_t; A) = \sum_{m=1}^M A_i(s_t)$ is the number of favors player i grants in state s_t under agreement A . Define the expected number of favors player i grants under A as $F_i(A)$.

Now define $\mathcal{E}_i(F_i)$ as the set of stage-game strategy profiles in which player i grants F_i favors in expectation. Each agreement A with $A_i \in \mathcal{E}_i(F_i)$ is associated with a maximal reneging temptation for player i . Let $d_i(F_i)$ be the smallest maximal reneging temptation among all such agreements, and let $\hat{f}_i(F_i) = d_i(F_i)/c$. Define $\mathcal{E}_i^*(F_i) \subset \mathcal{E}_i(F_i)$ as the set of stage-game strategy profiles in which player i grants F_i favors in expectation, and so his maximal deviation gain is $d_i(F)$. The game is symmetric, so $\hat{f}_1(F) = \hat{f}_2(F) \equiv \hat{f}(F)$ for all F , and it is clear that $\hat{f}(F)$ is weakly increasing in F .

Now, take an optimal agreement A^* that gives player i expected per-period payoff of v_i^* . It follows that

$$\begin{aligned} v_1^* &= F_2^* b - F_1^* c \\ v_2^* &= F_1^* b - F_2^* c. \end{aligned}$$

Let d_i^{\max} be the maximal deviation gain for player i , and let f_i^{\max} be the associated maximal number of favors player i is asked to grant. Since A^* is an equilibrium agreement, we must have that for each i ,

$$c f_i^{\max}(A^*) \leq \frac{\delta}{1-\delta} v_i^*.$$

Now, suppose to the contrary that $v_1^* > v_2^*$. It follows that $F_2^* > F_1^*$. For small enough feasible $\varepsilon > 0$, choose an agreement A with $A_1 \in \mathcal{E}_1^*(F_1^* + \varepsilon)$ and $A_2 \in \mathcal{E}_2^*(F_2^*)$. This agreement gives per-period payoffs of $v_1^* - c\varepsilon$ for player 1 and $v_2^* + b\varepsilon$ for player 2. In addition, player 1's maximal deviation gain is $\hat{f}_1(F_1^* + \varepsilon)c$ and player 2's is $\hat{f}_2(F_2^*)c$. For small enough feasible ε , we have

$$\hat{f}_1(F_1^* + \varepsilon)c \leq \hat{f}_2(F_2^*)c \leq f_2^{\max}c \leq \frac{\delta}{1-\delta} v_2^* < \frac{\delta}{1-\delta} (v_1^* - c\varepsilon),$$

where the first inequality follows because \hat{f}_i is weakly increasing, and $F_1^* < F_2^*$; the second inequality follows because $A_2 \in \mathcal{E}_2^*(F_2^*)$; the third inequality follows, because the initial agreement is an equilibrium agreement; the final inequality follows, because $v_1^* > v_2^*$. This implies that player 1 will not deviate. In addition, for player 2,

$$\hat{f}_2(F_2^*)c \leq f_2^{\max}c < \frac{\delta}{1-\delta}(v_2^* + b\varepsilon),$$

so player 2 will not deviate either. Therefore, A is an equilibrium agreement. Moreover, its total surplus is larger than the original agreement's, since more needed favors are granted in expectation. This contradicts the claim that the original agreement was optimal.

Step 2. If there is an optimal agreement, there is a symmetric agreement that is optimal. For each s_t , we can define the **mirror image state** s'_t in which $\tilde{s}'_{i,m,t} = \tilde{s}_{-i,m,t}$. Take an optimal agreement A^* . By Step 1, the associated payoffs satisfy $v_1^* = v_2^* = v^*$, and therefore, the associated expected number of favors players grants are given by $F_1^* = F_2^* = F^*$. Let d_i^{\max} be the associated maximal deviation gain for player i . Since A^* is an equilibrium agreement, it must be the case that for $i = 1, 2$, $d_i^{\max} \leq \delta v^* / (1 - \delta)$.

Now, construct a new agreement \hat{A} such that for all s_t , $\hat{A}_1(s_t) = A_1^*(s_t)$ and $\hat{A}_2(s_t) = A_1^*(s'_t)$. Under agreement \hat{A} , players' payoffs are $\hat{v}_1 = \hat{v}_2 = v^*$, and both players' maximal deviation gains are d_1^{\max} , so \hat{A} is an equilibrium agreement.

Step 3. Proof of Proposition 2. Take an optimal symmetric agreement A^* that gives each player an expected per-period payoff of v^* . Let f_i^{\max} be the maximal number of favors player i is expected to carry out under A^* . Since A^* is an equilibrium agreement, it follows that $cf_i^{\max} \leq \delta v^* / (1 - \delta)$.

Suppose that under A^* , there is a state s'_t in which $\sum_{m=1}^M A_{1,m}^*(s'_t) \leq f_1^{\max} - 1$ and some \hat{m} for which $A_{1,\hat{m}}^*(s'_t) = 0$ and $\tilde{s}_{\hat{m},t} = 1$. Now, consider a new agreement \hat{A} which coincides with $A_{1,m}^*(s'_t)$ for all $m \neq \hat{m}$ and $\hat{A}_{1,\hat{m}}(s'_t) = 1$. Adjust \hat{A}_2 accordingly to preserve symmetry. This new agreement does not increase either player's maximal deviation gain, but it increases the expected number of favors granted in each period and therefore increases both players' payoffs. This contradicts the claim that A^* was an optimal agreement and proves that any optimal symmetric agreement is an \bar{H} -threshold cooperation agreement.

Finally, since total expected surplus is increasing in \bar{H} , it is clear that an optimal agreement specifies the largest \bar{H} such that players' maximal deviation gains are smaller than their continuation losses. Notice that $V(\bar{H}, M)$ is each player's expected surplus under an \bar{H} -threshold cooperation agreement. In particular, notice that when $\bar{H} = M$, $V(M, M) = Mp(b - c)$ is the first-best payoff. Moreover,

$$V(\bar{H}, M) - V(\bar{H} - 1, M) = (1 - P(\bar{H}, M))(b - c),$$

and this gives the expression for $V(\bar{H}, M)$. ■

Proposition 3. When $\frac{\delta}{1-\delta} \in \left(\frac{c}{b-c}, \frac{1}{p} \frac{c}{b-c}\right)$, the optimal threshold $\bar{H}^*(M)$ is increasing in M , and for all n , $\bar{H}^*(nM) \geq n\bar{H}^*(M)$. As the breadth of the relationship goes to infinity, $V(\bar{H}^*(M), M) / M \rightarrow p(b - c)$ and

$$\bar{H}^*(M) / M \rightarrow \frac{\delta}{1-\delta} p \frac{b-c}{c} \in (p, 1)$$

Proof of Proposition 3. We first argue that $\bar{H}^*(M)$ is weakly increasing. As M increases to

$M + 1$, one can construct an agreement with the same threshold $\bar{H}^*(M)$. Under this agreement, players' maximal deviation gains remain unchanged. In addition, players' expected payoffs are higher since in the $M + 1$ -component game, the expected number of favors needed in the first M component games is the same as in the M -component game, and players may receive favors from the $M + 1$ -st component game. Therefore, this agreement is an equilibrium agreement. Since total payoffs are increasing in the cutoff \bar{H} , it follows that $\bar{H}^*(M + 1) \geq \bar{H}^*(M)$.

To see that $\bar{H}^*(nM) \geq n\bar{H}^*(M)$, first let the expected payoffs for each player in the M -component game be V . Now notice that when there are nM component games, they can be partitioned into n subsets with M components each. In each subset, we can repeat the optimal agreement from the M -component game, that is, using the $\bar{H}^*(M)$ -threshold cooperation agreement. Under the new agreement, the expected payoffs for each player are nV , and the maximal deviation gains are $n\bar{H}^*(M)$. It is then clear that this agreement is an equilibrium agreement. Now consider the $n\bar{H}^*(M)$ -threshold cooperation agreement. This agreement generates a payoff larger than nV , since more needed favors are being granted. It follows that optimal agreement in the nM -component game has a threshold $\bar{H}^*(nM) \geq n\bar{H}^*(M)$.

Next, the result that $\lim_{M \rightarrow \infty} V(\bar{H}^*(M), M)/M = p(b - c)$ follows directly from Theorem 1 and the optimality of the $\bar{H}^*(M)$ -threshold agreement. Finally, notice that,

$$\bar{H}^*(M)c \leq \frac{\delta}{1 - \delta} V(\bar{H}^*(M), M) \leq (\bar{H}^*(M) + 1)c,$$

where the second inequality follows from the optimality of $\bar{H}^*(M)$ and that given the restrictions on δ , first-best payoffs cannot be achieved.

If we divide the above inequalities by Mc and take limits as $M \rightarrow \infty$, then by the squeeze theorem,

$$\lim_{M \rightarrow \infty} \frac{\bar{H}^*(M)}{M} = \lim_{M \rightarrow \infty} \frac{\delta}{1 - \delta} \frac{V(\bar{H}^*(M), M)}{Mc} = \frac{\delta}{1 - \delta} \frac{p(b - c)}{c}.$$

Finally, notice that $\frac{\delta}{1 - \delta} \frac{p(b - c)}{c} < 1$ because first-best payoffs cannot be achieved. And $1 < \frac{\delta}{1 - \delta} \frac{b - c}{c}$ follows because component full cooperation can be almost-perfect replicated. ■

D. Multimarket Contact

Proposition 4. *Component perfect collusion is an equilibrium component agreement if and only if*

$$\tilde{\pi}^*(h) \leq \frac{\delta}{1 - \delta} (p\tilde{\pi}^*(h) + (1 - p)\tilde{\pi}^*(l))$$

Proof of Proposition 4. For sufficiency, note that if the condition is satisfied, it is easy to check that component perfect collusion is an equilibrium component agreement. To prove necessity, take any component agreement \tilde{A} in which $\tilde{a}_{m,t}(\tilde{s}) = \tilde{a}^*(\tilde{s})$ for $\tilde{s} \in \{h, l\}$. Let $\tilde{u}_i(\tilde{s}; \tilde{A})$ be firm i 's profits in state \tilde{s} in this agreement. Component perfect collusion implies that $\tilde{u}_1(\tilde{s}; \tilde{A}) + \tilde{u}_2(\tilde{s}; \tilde{A}) = \tilde{\pi}^*(\tilde{s})$. In state h , firm i can guarantee itself a payoff of $\tilde{\pi}^*(h) - \varepsilon$ for any $\varepsilon > 0$ by setting a price slightly smaller than $\tilde{\alpha}^*(h)$. It follows that

$$\tilde{\pi}^*(h) - \varepsilon \leq \tilde{u}_i(h; \tilde{A}) + \frac{\delta}{1 - \delta} (p\tilde{u}_i(h; \tilde{A}) + (1 - p)\tilde{u}_i(l; \tilde{A}))$$

for each firm. Summing over i and using $\tilde{u}_1(\tilde{s}; \tilde{A}) + \tilde{u}_2(\tilde{s}; \tilde{A}) = \tilde{\pi}^*(\tilde{s})$, we obtain that

$$\tilde{\pi}^*(h) - \varepsilon \leq \frac{\delta}{1-\delta} (p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)).$$

Taking ε to zero, we obtain the desired result. ■

Proposition 5. When $\frac{\delta}{1-\delta} \in \left(1, \frac{\tilde{\pi}^*(h)}{p\tilde{\pi}^*(h)+(1-p)\tilde{\pi}^*(l)}\right)$, any optimal agreement is an $\bar{H}^*(M)$ -threshold collusive agreement, where $\bar{H}^*(M)$ is the largest integer satisfying $\bar{H}^*(M)\tilde{\pi}^*(h) + (M - \bar{H}^*(M))\tilde{\pi}^*(l) \leq \delta V(\bar{H}^*(M), M)$. Moreover, for all s_t satisfying $H(s_t) > \bar{H}^*(M)$,

$$\sum_{m=1}^M \left[\tilde{u}_1 \left(A_{(m)}^{\bar{H}^*(M)}(s_t), \tilde{s}_{m,t} \right) + \tilde{u}_2 \left(A_{(m)}^{\bar{H}^*(M)}(s_t), \tilde{s}_{m,t} \right) \right] = \delta V(\bar{H}^*(M), M).$$

Proof of Proposition 5. We first show that for any equilibrium agreement A , there is a symmetric collusive agreement that yields the same total profits. Given any agreement $A : \mathcal{S}^M \rightarrow \mathcal{A}^M$, define the associated total profits in state s_t as $\pi(s_t; A) = \sum_{m=1}^M \tilde{a}_{m,t}(s_t) \tilde{q}(\tilde{a}_{m,t}(s_t), \tilde{s}_{m,t})$ and the associated expected per-period profits as $\pi(A) = \sum_{s_t} \Pr[s_t] \pi(s_t; A)$. By the same type of argument as in the proof of Proposition 4, a necessary condition for A to be an equilibrium agreement is that

$$\max_{s_t} \pi(s_t; A) \leq \frac{\delta}{1-\delta} \pi(A).$$

If A satisfies this condition, construct agreement \hat{A} such that for each i , $\hat{A}_{i,m}(s_t) = \tilde{a}_{m,t}(s_t)$. Under this agreement, $u_i(s_t; \hat{A}) = \pi(s_t; A)/2$, so $d_i^{\max}(\hat{A}) = \max_{s_t} \pi(s_t; A)/2$ and $v(A) = \pi(A)/2$. \hat{A} is therefore an equilibrium agreement, and it is a symmetric collusive agreement.

Next, take an optimal agreement A^* , and let $\pi(A^*)$ denote the total per-period expected profits under A^* . For each s_t , if $\pi(s_t; A^*) < \delta\pi(A^*)/(1-\delta)$, then an argument similar to that in Proposition 2 implies that both firms set the monopoly price in each market, that is $A_{i,m}^*(s_t) = \tilde{a}^*(\tilde{s}_{m,t})$. This implies there exists a largest integer, \bar{H}^* , which satisfies $\bar{H}^*\tilde{\pi}^*(h) + (M - \bar{H}^*)\tilde{\pi}^*(l) \leq \delta\pi(A^*)/(1-\delta)$. Moreover, whenever $H(s_t) \leq \bar{H}^*$, both firms set the monopoly price in each market. Whenever $H(s_t) > \bar{H}^*$, then market profits are exactly equal to $\pi(A^*)$. It must therefore be the case that $\pi(A^*)$ solves the following equation:

$$\pi = \sum_{m=0}^{\bar{H}^*} (m\tilde{\pi}^*(h) + (M-m)\tilde{\pi}^*(l)) p(m, M) + \frac{\delta}{1-\delta} (1 - P(\bar{H}^*, M)) \pi.$$

This establishes Proposition 5.

Proposition 6. When $\frac{\delta}{1-\delta} \in \left(1, \frac{\tilde{\pi}^*(h)}{p\tilde{\pi}^*(h)+(1-p)\tilde{\pi}^*(l)}\right)$, the optimal threshold $\bar{H}^*(M)$ is increasing in M , and for all n , $\bar{H}^*(nM) \geq n\bar{H}^*(M)$. As the number of markets goes to infinity, $V(\bar{H}^*(M), M)/M \rightarrow p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)$ and

$$\bar{H}^*(M)/M \rightarrow \frac{\delta(p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)) - (1-\delta)\tilde{\pi}^*(l)}{(1-\delta)(\tilde{\pi}^*(h) - \tilde{\pi}^*(l))} \in (p, 1).$$

Proof of Proposition 6. The proof of monotonicity of $\bar{H}^*(M)$ and the proof that $\bar{H}^*(nM) \geq n\bar{H}^*(M)$ follow from identical arguments as in Proposition 4 and are omitted here. The result that $\lim_{M \rightarrow \infty} V(\bar{H}^*(M), M)/M = p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)$ follows directly from Theorem 1 and the optimality of the $\bar{H}^*(M)$ -threshold collusive agreement.

Finally, from the definition of $\bar{H}^*(M)$, we have

$$\begin{aligned} & \frac{\bar{H}^*(M)\tilde{\pi}^*(h) + (M - \bar{H}^*(M))\tilde{\pi}^*(l)}{M} \\ & \leq \frac{\delta}{1-\delta} \frac{V(\bar{H}^*(M), M)}{M} \\ & < \frac{(\bar{H}^*(M) + 1)\tilde{\pi}^*(h) + (M - \bar{H}^*(M) - 1)\tilde{\pi}^*(l)}{M}. \end{aligned}$$

By the squeeze theorem and the observation that $\lim_{M \rightarrow \infty} V(\bar{H}^*(M), M) = p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)$,

$$\lim_{M \rightarrow \infty} \frac{\bar{H}^*\tilde{\pi}^*(h) + (M - \bar{H}^*)\tilde{\pi}^*(l)}{M} = \frac{\delta}{1-\delta} (p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)).$$

It then follows that

$$\lim_{M \rightarrow \infty} \frac{\bar{H}^*(M)}{M} = \frac{\delta(p\tilde{\pi}^*(h) + (1-p)\tilde{\pi}^*(l)) - (1-\delta)\tilde{\pi}^*(l)}{(1-\delta)(\tilde{\pi}^*(h) - \tilde{\pi}^*(l))},$$

which establishes the result. ■

E. Relational Incentive Contracts

Proposition 7. *If $pb > c$ and $pb < \delta(p - pb)/(1 - \delta)$, then a component effort-inducing relational contract can be almost-perfectly replicated.*

Proof of Proposition 7. Given $pb > c$, there exists an $\varepsilon > 0$ such that $(1 - 2\varepsilon)pb > c$. Now, take m^* to be the smallest integer satisfying $m^* \geq p(1 + \varepsilon)M$. Now, consider an effort-inducing agreement with bonus scheme $b_m = \min\{m^*b, \hat{m}(s_t)b\}$, where $\hat{m}(s_t) = \sum_{m=1}^M \tilde{s}_{m,t}$ is the number of activities in which output was equal to 1. For M sufficiently large, the probability that $\hat{m}(s_t) \leq m^*$ is greater than $1 - \varepsilon$.

We first show that the agent does not want to deviate. In principle, the agent's maximal net gain may occur when he chooses not to exert effort on multiple activities. In other words, we need to check that if the agent has exerted effort in $L < M$ activities, he always benefits from exerting effort in $L + 1$ activities. To see this, note that the incremental cost of doing so is c , and the incremental benefit is the increase in his expected bonus, which is

$$\begin{aligned} & \Pr[\hat{m}(s_t) \leq m^* - 1 | e_m = 1 \text{ in } L \text{ activities}]pb \\ & \geq \Pr[\hat{m}(s_t) \leq m^* - 1 | e_m = 1 \text{ in } M \text{ activities}]pb \\ & \geq (1 - \varepsilon)pb, \end{aligned}$$

where the first inequality follows from MLRP, and the second inequality follows from the condition above. Since $(1 - 2\varepsilon)pb > c$, the agent's net deviation gain is at most $c - (1 - \varepsilon)pb < 0$.

We next show that the principal does not want to deviate. The necessary and sufficient condition is that

$$m^*b \leq M \frac{\delta}{1-\delta} (p - pb).$$

Note that this condition is essentially (taking into account the integer constraint)

$$p(1 + \varepsilon)Mb \leq M \frac{\delta}{1-\delta} (p - pb).$$

For ε small enough, this condition is satisfied whenever $pb < \delta(p - pb)(1 - \delta)$. An effort-inducing agreement can therefore be almost-perfectly replicated. ■

Proposition 8. If $\frac{\varepsilon}{p} > \frac{\delta}{1-\delta} (p - c) > c$, then an optimal effort-inducing relational contract has the following bonus scheme:

$$b_m = \begin{cases} 0 & \sum_{m=1}^M \tilde{s}_{m,t} < m^*(M) \\ \gamma & \sum_{m=1}^M \tilde{s}_{m,t} = m^*(M) \\ \beta & \sum_{m=1}^M \tilde{s}_{m,t} > m^*(M) \end{cases}$$

for some integer $m^*(M)$, where $0 \leq \gamma \leq \beta$.

Proof of Proposition 8. The proof of this proposition is broken into several steps.

Step 1. Because the salary component of the relational contract can be used to extract all the surplus from the agent at the beginning of each period, it suffices to find a bonus scheme $(b_m)_{m=0}^M$ that satisfies the agent's incentive-compatibility constraints and has the smallest maximal payment. This dual problem can be written as a linear program. Let $p(m, K)$ denote the probability that m outputs are high realized given the agent exerts effort in L tasks. We then have:

$$\min_{b_0, b_1, \dots, b_M, b} b \tag{5}$$

subject to $(IC - K)$

$$\sum_{m=0}^M [p(m, M) - p(m, L)] b_m - (M - L)c \geq 0 \text{ for } L = 0, 1, \dots, M - 1$$

and

$$0 \leq b_m \leq b \text{ for } m = 0, 1, \dots, M.$$

Step 2. In this step, we define a relaxed linear programming problem, which has a straightforward solution, and we will show that the solution to this program is also a solution to the full problem. The relaxed problem is the same as (5), except we ignore all the incentive-compatibility constraints, except for $(IC - M - 1)$ and $(IC - 0)$:

$$\min_{b_0, b_1, \dots, b_M, b} b \tag{6}$$

subject to the local incentive constraint (*LIC*)

$$\sum_{m=0}^M [p(m, M) - p(m, M-1)] b_m - c \geq 0,$$

the global incentive constraint (*GIC*)

$$\sum_{m=0}^M p(m, M) b_m - b_0 - Mc \geq 0,$$

and

$$0 \leq b_m \leq b \text{ for } m = 0, 1, \dots, M.$$

Step 3. The solution to (6) is given by

$$b_m = \begin{cases} 0 & \sum_{m=1}^M \tilde{s}_{m,t} < m^*(M) \\ \gamma & \sum_{m=1}^M \tilde{s}_{m,t} = m^*(M) \\ \beta & \sum_{m=1}^M \tilde{s}_{m,t} > m^*(M), \end{cases}$$

where $0 \leq \gamma \leq \beta$. To see this, we write the Lagrangian for (6):

$$\begin{aligned} \mathcal{L} = & -b + \lambda_1 \left[\sum_{m=0}^M (p(m, M) - p(m, M-1)) b_m - c \right] + \lambda_2 \left[\sum_{m=0}^M p(m, M) b_m - b_0 - Mc \right] \\ & + \sum_{m=0}^M [\mu_{m+} (b - b_m) + \mu_{m-} b_m], \end{aligned}$$

which yields, for $m \geq 1$, the following first-order conditions:

$$\lambda_1 (p(m, M) - p(m, M-1)) + \lambda_2 p(m, M) + (\mu_{m-} - \mu_{m+}) = 0.$$

There are several cases. First, suppose that $(\lambda_1 + \lambda_2) p(m, M) - \lambda_1 p(m, M-1) > 0$. Notice that this is the case if and only if

$$\frac{p(m, M)}{p(m, M-1)} > \frac{\lambda_2}{\lambda_1 + \lambda_2} \equiv \lambda.$$

Notice that $p(m, M) = \binom{M}{m} p^m (1-p)^{M-m}$ and $p(m, M-1) = \binom{M-1}{m} p^m (1-p)^{M-1-m}$, implying that the likelihood ratio

$$\frac{p(m, M)}{p(m, M-1)} = \frac{M}{M-m} (1-p).$$

$m > Mp$

The likelihood ratio is increasing in m , and this implies that $\frac{p(m,M)}{p(m,M-1)} > \lambda$ if and only if

$$m > \frac{p + \lambda - 1}{\lambda} M \equiv m^*(M).$$

This implies that $m > m^*(M)$, which implies that $\mu_{m+} > 0$ and thus $b_m^* = b$.

Second, suppose $(\lambda_1 + \lambda_2)p(m, M) - \lambda_1 p(m, M - 1) < 0$. This implies that $1 \leq m < m^*(M) = \frac{\lambda + p - 1}{\lambda} M$, which in turn implies that $b_m^* = 0$ and thus $\mu_{m-} > 0$.

Third, if $(\lambda_1 + \lambda_2)p(m, M) - \lambda_1 p(m, M - 1) = 0$, then $\mu_{m+} = \mu_{m-} = 0$, which implies that $m = m^*(M)$ and $b_m^* = \gamma \in [0, \beta]$.

Finally, when $m = 0$, the FOC is given by

$$\lambda_1 (p(0, M) - p(0, M - 1)) + \lambda_2 p(0, M) - \lambda_2 + (\mu_{0-} - \mu_{0+}) = 0.$$

Note that $\lambda_1 (p(0, M) - p(0, M - 1)) + \lambda_2 p(0, M) - \lambda_2 \leq (\lambda_1 + \lambda_2)p(0, M) - \lambda_1 p(0, M - 1) < 0$, where the last inequality follows from the argument above. This implies that $\mu_{0-} > 0$, and therefore $b_0^* = 0$.

Step 4. We now show that b_m^* solves (5). First, derive the marginal net benefit of exerting effort on the L th task for the conjectured solution:

$$\Delta(L) = \sum_{m=0}^M [p(m, L) - p(m, L - 1)] b_m^* - c.$$

We want to show that only $L = 0$ or $L = M$ is ever optimal. To do so, we examine how $\Delta(L)$ changes with L . Note that

$$\Delta(L) = p(m^* - 1, L - 1)p\gamma + p(m^*, L - 1)p(\beta - \gamma) - c.$$

The marginal benefit of effort on the L th task is γ times the probability that with only $L - 1$ tasks, the agent would have had $m^* - 1$ successes (receiving a bonus of 0) but instead has m^* successes (receiving a bonus of γ) plus $(\beta - \gamma)$ times the probability that the agent would have had m^* successes (receiving γ) with only $L - 1$ tasks but instead has $m^* + 1$ successes (receiving β).

It can be shown that $\Delta(L)$ for $L \geq m^*$ is either weakly decreasing or it initially increases and then decreases. We do not replicate the proof here, because it is similar to the result of Bond and Gomes (2009). Finally, note that $\Delta(M) \geq 0$, because the (*LIC*) constraint is satisfied. Given this structure of marginal benefits, the agent's problem either has one peak (at $L^* = M$) or two peaks (one at $L^* = 0$ and another at $L^* = M$), so the agent's optimal choice of effort is either $L^* = 0$ or $L^* = M$. Thus, the solution to the relaxed problem is in fact the solution to the full problem. ■

Proposition 9. *An effort-inducing relational contract is an equilibrium agreement if and only if*

$$\frac{\delta}{1 - \delta} \geq \frac{pc / (p - r)}{\left(1 - (1 - p)^M\right) \tilde{v}_0}.$$

Proof of Proposition 9. Note that

$$\begin{aligned}
\sum_{k=0}^{M-1} \frac{1}{k+1} \Pr \left[\sum_{\hat{m} \neq m} \tilde{s}_{\hat{m},t} = k \right] &= \sum_{k=0}^{M-1} \frac{1}{k+1} \binom{M-1}{k} p^k (1-p)^{M-1-k} \\
&= \frac{1}{Mp} \sum_{k=0}^{m-1} \binom{M}{k+1} p^{k+1} (1-p)^{M-1-k} \\
&= \frac{1 - (1-p)^M}{Mp}.
\end{aligned}$$

Substituting this into (3), we obtain the result. ■

References

- [1] ABDULKADIROGLU, ATILA AND KYLE BAGWELL. 2013. Trust, Reciprocity, and Favors in Cooperative Relationships. *American Economic Journal: Microeconomics*, 5(2): 213-259.
- [2] BAKER, GEORGE, ROBERT GIBBONS, AND KEVIN MURPHY. 2011. Relational Adaptation. *Working Paper*.
- [3] BARRON, DANIEL. 2013. Putting the Relationship First: Relational Contracts and Market Structure. *Working Paper*.
- [4] BERNHEIM, DOUGLAS AND MICHAEL WHINSTON. 1987. Multimarket Contact and Collusive Behavior. *Working Paper*.
- [5] BERNHEIM, DOUGLAS AND MICHAEL WHINSTON. 1990. Multimarket Contact and Collusive Behavior. *RAND Journal of Economics*, 21(1): 1-26.
- [6] BILLINGSLEY, PATRICK. 1986. *Probability and Measure*, 2nd Edition. Wiley Press, New York.
- [7] BOND, PHILIP AND ARMANDO GOMES. 2009. Multitask Principal-Agent Problems: Optimal Contracts, Fragility, and Effort Misallocation. *Journal of Economic Theory*, 114(1): 175-211.
- [8] ELLICKSON, ROBERT. 1994. *Order without Law—How Neighbors Settle Disputes*. Cambridge, Massachusetts: Harvard University Press.
- [9] FONG, YUK-FAI AND JIN LI. Forthcoming. Information Revelation in Relational Contracts. *Review of Economic Studies*.
- [10] HAUSER, CHRISTINE AND HUGO HOPENHAYN. 2008. Trading Favors: Optimal Exchange and Forgiveness. *Working Paper*.
- [11] LAUX, CHRISTIAN. 2001. Limited-Liability and Incentive Contracting with Multiple Projects. *RAND Journal of Economics*, 32(3): 514-526.
- [12] LEVIN, JONATHAN. 2002. Multilateral Contracting and the Employment Relationship. *Quarterly Journal of Economics*, 117(3): 1075-1103.
- [13] LEVIN, JONATHAN. 2003. Relational Incentive Contracts. *American Economic Review*, 93(3): 835-857.
- [14] MAILATH, GEORGE AND LARRY SAMUELSON. 2006. *Repeated Games and Reputations*. Oxford: Oxford University Press.
- [15] MOBIUS, MARKUS. 2001. Trading Favors. *Working Paper*.
- [16] ROTEMBERG, JULIO AND GARTH SALONER. 1986. A Supergame-Theoretic Model of Price Wars during Booms. *American Economic Review*, 76(3): 390-407.
- [17] SEKIGUCHI, TADASHI. 2015. Multimarket Contact under Demand Fluctuations. *International Journal of Game Theory*, 44: 1033-1048.