Essays in Revision Games

Dissertation Advisor: Drew Fudenberg Yuichiro Kamada

Essays in Revision Games

ABSTRACT

This dissertation consists of three essays related to revision games.

The first essay proposes and analyzes a new model that we call "revision games," which captures a situation where players in advance prepare their actions in a game. After the initial preparation, they have some opportunities to revise their actions, which arrive stochastically. Prepared actions are assumed to be mutually observable. We show that players can achieve a certain level of cooperation. The optimal behavior of players can be described by a simple differential equation.

The second essay studies a version of revision games in which revision opportunities are asynchronous across players. In 2-player "common interest" games where there exists a best action profile for all players, this best action profile is the only equilibrium outcome of the revision game. In "opposing interest" games which are 2 x 2 games with Pareto-unranked strict Nash equilibria, the equilibrium outcome of the revision game is generically unique and corresponds to one of the stage-game Nash equilibria. Which equilibrium prevails depends on the payoff structure and on the relative frequency of the arrivals of revision opportunities for each of the players.

The third essay studies a multi-agent search problem with a deadline: for instance, the situation that arises when a husband and a wife need to find an apartment by September 1. We provide an understanding of the factors that determine the positive search duration in reality. Specifically, we show that the expected search duration does not shrink to zero even in the limit as the search friction vanishes. Additionally, we find that the limit duration increases as more agents are involved, for two reasons: the *ascending acceptability effect* and the *preference heterogeneity effect*. The convergence speed is high, suggesting that the mere existence of *some* search friction is the main driving force of the positive duration in reality. Welfare implications and a number of discussions are provided. Results and proof techniques developed in the first two essays are useful in proving and understanding the results in the third essay.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

First, I am deeply grateful for the continual guidance and support of my thesis committee: Drew Fudenberg, Attila Ambrus, Alvin Roth, and Tomasz Strzalecki. Especially, Drew Fudenberg has advised me since the beginning of (actually, even before the start of) the PhD program at Harvard University. I also would like to thank Hirohisa Kishino and Akihiko Matsui for their generous support and encouragement during and after my transition from Agriculture to Economics. I cannot imagine writing a thesis in economics without Professors Fudenberg, Kishino, and Matsui.

I have been fortunate to work with many coauthors: Attila Ambrus, Eduardo Azevedo, Riccardo Calcagno, Drew Fudenberg, Daisuke Fujii, Ryota Iijima, Yuhta Ishii, Michihiro Kandori, Fuhito Kojima, Scott Kominers, Stefano Lovo, Nozomu Muto, Alex Peysakhovich, Neel Rao, Takuo Sugaya, Yuki Takagi, and Jun Wako. They hugely affected the way I view economics, and taught me how to write papers. Especially, the three essays in this dissertation are all coauthored. I owe my thanks to these and other coauthors.

I thank my family, Yaemi and Aiko Kamada, who supported my life in Boston, and my family back in Japan, Yoichiro, Keiichiro, Izumi, Chihiro, and Minato Kamada, for their love and encouragement. This dissertation is dedicated to my grandmother, Kazuko Koyano, who passed away when I was in the United States.

1. INTRODUCTION

This dissertation consists of three essays related to revision games.

The first essay proposes and analyzes a new model that we call "revision games," which captures a situation where players in advance prepare their actions in a game. After the initial preparation, they have some opportunities to revise their actions, which arrive stochastically. Prepared actions are assumed to be mutually observable. We show that players can achieve a certain level of cooperation. The optimal behavior of players can be described by a simple differential equation.

The second essay studies a version of revision games in which revision opportunities are asynchronous across players. In 2-player "common interest" games where there exists a best action profile for all players, this best action profile is the only equilibrium outcome of the revision game. In "opposing interest" games which are 2 x 2 games with Pareto-unranked strict Nash equilibria, the equilibrium outcome of the revision game is generically unique and corresponds to one of the stage-game Nash equilibria. Which equilibrium prevails depends on the payoff structure and on the relative frequency of the arrivals of revision opportunities for each of the players.

The third essay studies a multi-agent search problem with a deadline: for instance, the situation that arises when a husband and a wife need to find an apartment by September 1. We provide an understanding of the factors that determine the positive search duration in reality. Specifically, we show that the expected search duration does not shrink to zero even in the limit as the search friction vanishes. Additionally, we find that the limit duration increases as more agents are involved, for two reasons: the *ascending acceptability effect* and the *preference heterogeneity effect*. The convergence speed is high, suggesting that the mere existence of *some* search friction is the main driving force of the positive duration in reality. Welfare implications and a number of discussions are provided. Results and proof techniques developed in the first two essays are useful in proving and understanding the results in the third essay.

These essays enhance our understanding of dynamic incentives in finite horizon, and raises many interesting questions for future research. These open questions are noted at the conclusion section of each chapter.

2. REVISION GAMES[1]

2.1 *Introduction*

In social or economic problems, agents often prepare their actions in advance before they interact. Consider researchers who are competing to win research grants. "Actions" in this context correspond to research proposals to be submitted by a prespecified deadline. Researchers prepare their proposals in advance, and proposals are usually subject to some revisions before submission. Since they have other obligation, such as teaching and committee work, a revision can be made only when an opportunity to work on the proposal arrives. Researchers may also obtain some information about their rivals' proposals. Based on such information, researchers revise their proposals, and they submit what they have when the deadline comes.

In the present paper, we introduce a stylized model to capture such a situation, which we call a *revision game*. In a revision game, a *component game* is played only once, and players must in advance prepare their actions. They have some opportunities to revise their prepared actions, and the opportunities for revision arrive stochastically. Prepared actions are assumed to be mutually observable, and the

[1] This is a joint work with Michihiro Kandori (Faculty of Economics, University of Tokyo). We thank Drew Fudenberg, Barton Lipman, Stephen Morris, Satoru Takahashi, and the seminar participants at the 2009 Far East and South Asia Meetings of the Econometric Society, 2009 Evolution of Cooperation Conference at IIASA, Boston University, University College London, Harvard University, and Oxford University.

final action in the last revision opportunity is played in the component game. We show that, under some regulatory conditions, players can achieve a certain level of cooperation.

Let us contrast our model with the well-known fact that players can cooperate in a repeated game. If players expect a sufficient future reward, they can sustain costly cooperation. It must be the players' best interest to carry out the future reward, which is guaranteed by reward in the further future, and so forth till indefinitely. In this paper we argue that players can sometimes cooperate *even though the game is played only once*. Cooperation can be sustained by revision process of players' actions.

The basic mechanism to sustain cooperation in a revision game is similar to that in a repeated game, although the mechanism operates in somewhat disguised way. This is best seen when the revision process is stationary. Suppose players prepare their action in each period, and the prepared actions are played in the component game with a (small) constant probability. Once a component game is played, the game is over and there is no further interaction. In Section 2.2 we present a simple observation that such a model is actually isomorphic to an infinitely repeated game with a (high) discount factor.

The heart of the paper analyzes a more realistic case, where the component game is played at a *predetermined deadline*. Players obtain revision opportunities according to a Poisson process, and the finally-prepared actions are played at the deadline. In the class of component games that we focus on, we will show that an optimal symmetric trigger-strategy equilibrium exists and it is essentially unique. The equilibrium is characterized by a simple differential equation, which we apply to a variety of economic examples. In particular, the revision game of a Cournot duopoly game can achieve, in expectation, more than 96% of the full collusive payoff.

The key difficulty in sustaining cooperation comes with the fact that the preparation phase ends at a predetermined deadline: As time approaches the deadline, the probability of being rewarded in the future shrinks to zero. This means that the instantaneous cost of cooperation (the gain from deviation) must shrink to zero as well for incentive compatibility to be met at each moment of time.[2] We construct a trigger strategy equilibrium with such a property. On the equilibrium path of play, players prepare action $x(t)$ if they obtain a revision opportunity at time t; upon deviation players revert to the (unique) Nash action. $x(t)$ is a full collusive action when time t is sufficiently far away from the deadline, and it (continuously in t) approaches the Nash action towards the deadline. At the deadline, no more opportunity for reward is expected, so the only sustainable action profile is the static Nash action profile. For a 2-player good exchange game, we depict in Figure 2.1 the path $x(t)$ of the optimal equilibrium among all the trigger strategy equilibria, and a sample equilibrium path of play given $x(t)$.[3]

As the action approaches the Nash equilibrium, the instantaneous cost of cooperation shrinks to zero. However, it turns out that this is not enough to sustain cooperation. We further need that *the instantaneous cost shrinks sufficiently fast*. To see this point, note that as the action approaches the Nash action, the magnitude of the benefit from the opponent's future cooperation (conditional on there being an opportunity) shrinks to zero as well.[4] Since these benefits realize with a vanishing

---

[2] Under certain regularity assumptions.

[3] The formal analysis can be found in Section 2.5.1.

[4] Under a continuity assumption.

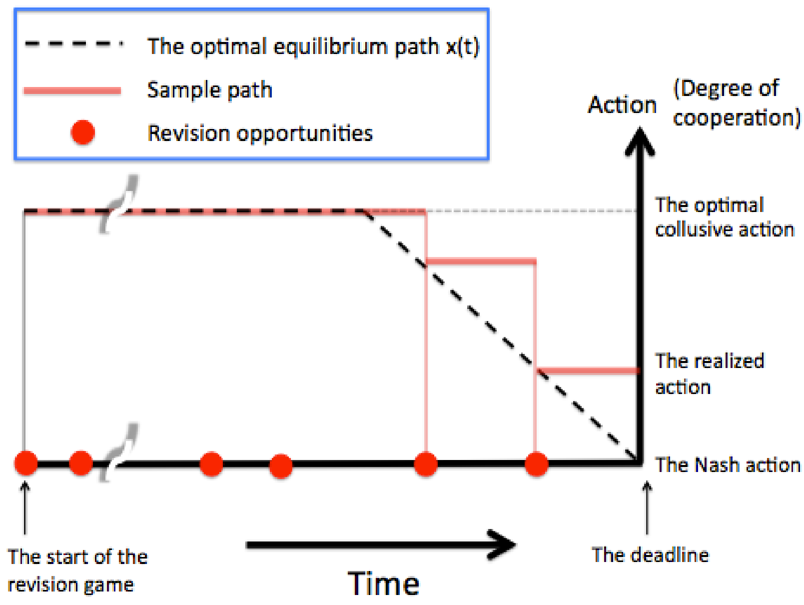

Figure 2.1: The optimal path and a sample path for a good exchange game.

probability, *the cost must be negligible relative to the benefit* when the action is close to Nash. We show that under an assumption on payoff structure that many economic applications satisfy,[5] the cost indeed shrinks fast enough.

2.1.1 *Related literature*

Although in revision games the component game is played only once, features of the model and the dynamic of the equilibrium that they imply are closely related to those in finitely repeated games. A striking fact that the repetition of defections is the only subgame perfect equilibrium in a repeated prisoner's dilemma is overcome by a variety of "twists" in the literature, such as multiple Nash payoffs.[6],

[5] This is expressed in Assumption A4 that we state in Section 2.4.

[6] Harrington (1987), Benoit and Krishna (1985, 1993).

bounded rationality[7], reputational effects[8], non-common knowledge of the timing of the deadline[9], social preferences[10], and so forth. Among others, the model of fu Chou and Geanakoplos (1988) is the most related to ours. They consider finite horizon repeated games in which a player can commit to a (contingent) action at the final period, and show that in "smooth games" a folk theorem obtains. The trigger-strategy equilibrium that we construct is reminiscent of theirs in that the action on the equilibrium path converges to the static Nash action, and the idea is related in the sense that in both models a small amount of cooperation at periods close to the deadline builds up a basis for a large cooperation in the entire game. The key difference, besides the fact that the component game is played only once in revision games, is that we do not use commitment to achieve cooperation—our players are fully rational. In our model rational players can cooperate even when the deadline is very close because there is no pre-determined "final period" at which players take actions with a positive probability.

Bernheim and Dasgupta (1995) consider infinite horizon repeated games in which the discount factor falls over time to approximate zero, and show that cooperation can be sustained if the speed at which the discount factor falls is sufficiently slow. They obtain a sufficient condition for the sustainability of cooperation but did not explore characterization of optimal equilibria. Although the mechanism to sustain cooperation in their model is similar to ours, in Section 2.6.2 we show some crucial differences and demonstrate that our model cannot be mapped into their

---

[7] Fudenberg and Levine (1983), Kalai and Neme (1992), Neyman (1985, 1998).

[8] Sobel (1985), Kreps, Milgrom, Roberts, and Wilson (1982), and Fudenberg and Maskin (1986).

[9] Neyman (1999).

[10] Ambrus and Pathak (2011).

model.

Pitchford and Snyder (2004) and Kamada and Rao (2009) consider situations in which two parties dynamically transfer a fixed amount of divisible goods that benefit the other party.[11] In the equilibria they construct, a failure to transfer the specified amount of goods at the specified date causes the opponent to stop transferring in the future. The remaining amount of the good in hand converges to zero as the transactions occur a number of times, so the relevant stake of the game gets smaller and smaller over time, reminiscent of our equilibrium form in which $x(t)$ approaches 0 as the deadline comes close. Also, as in our model, there cannot be a final transaction period, since if there can, then the parties do not have an incentive to make a transfer: the transactions need to occur indefinitely. The key difference is that in their models the transaction amount is specified in the way that the game becomes "isomorphic" from one period to the other in an appropriate sense (this is possible since the horizon is infinite), while in our equilibrium the balance of cost and reward for cooperation changes over time, as we have already discussed.

At a technical level, our model is related to that of Ambrus and Lu (2010a) who analyze a multilateral bargaining problem in a continuous-time finite-horizon setting where opportunities of proposal arrive via Poisson processes. If an agreement is reached at any time, the game ends then. If no offer is accepted until the deadline, players receive the payoff 0. They show that there is a unique Markov perfect equilibrium in which the first proposal is accepted, so the proposals that different players make converge to the same limit as the horizon length becomes large. Although their basic framework is similar to ours, there are two main differences. First, in their model the game can end before the deadline, if an agreement

[11] See related papers on gradualism, e.g. Admati and Perry (1991), Marx and Matthews (2000), and Compte and Jehiel (2003) and monotone games, e.g. Gale (1995, 2001).

is reached. Second, they focus on Markov perfect equilibrium, which in our model corresponds to the repetition of the component game Nash equilibrium.

The rest of the paper is organized as follows: The next section presents a class of revision games without a deadline to help the readers build up some intuition. The main model with a deadline is presented in Section 2.3. The results on a general setting with one-dimensional continuous strategies are given in Section 2.4.. Section 2.5 provides a number of applications. Section 2.6 discusses the robustness of our results to fine changes of the specification of the model, identifies the condition for the sustainability of cooperation, and compares our model to an infinite horizon model with a decreasing discount factor. Section 2.7 concludes.

2.2 An Example (Two Samurai): Stationary Revision Games

The purpose of this paper is to analyze a class of games where (i) a component game is played only once, (ii) players must prepare their actions in advance, (iii) prepared actions are observable, and (iv) the probability that the prepared actions are actually played is strictly positive but not one. We refer those games as *revision games*. In this section, we start with a simple case, where the problem is stationary in the sense that in each period $t = 0, 1, 2, ...$ there is a fixed, positive probability p with which the component game is played. We refer to this class of revision games as *stationary revision games*. This class will turn out to be isomorphic to a familiar class of games, and it helps to build some intuition on how revision games in general work. The point we make is a simple one, so we just present an example of stationary revision games.

Suppose that a rural village faces an attack of bandits. In each period $t = 0, 1, 2, ...$ the bandits attack the village with probability $p \in (0, 1)$ around midnight.

Table 2.1: The Samurai example.

	Defend	Hide
Defend	2, 2	-1, 3
Hide	3, -1	0, 0

They attack only once. The villagers hired two samurai, $i = 1, 2$, and they must prepare to defend the village (to show up at the village gate around midnight) or not (to hide away and watch the gate from a distance). Hence in each period they observe each other's prepared actions. The acts of preparation themselves (showing up and hiding away) have negligible effects on the samurai's payoffs. When the bandits attack, however, their prepared actions have huge impacts on their payoffs. Payoffs are summarized in Table 2.1. This is a Prisoner's Dilemma game. Now consider player i's expected payoff. We denote player i's payoff by $\pi_i(t)$, when the bandits' attack occurs at time t. We also assume that players have a common discount factor $\delta \in (0,1)$. Player i's expected payoff is

$$\begin{aligned} & p\pi_i(0) + \delta(1-p)p\pi_i(1) + \delta^2(1-p)^2 p\pi_i(2) + \cdots \\ = \ & p\sum_{t=0}^{\infty} \overline{\delta}^t \pi_i(t), \end{aligned}$$

where $\overline{\delta} := \delta(1-p)$. Hence, *stationary revision games are isomorphic to infinitely repeated games*, and cooperation can be sustained in a subgame perfect equilibrium if p is small. Even though the component game is played only once, when (i) a component game is played only once, (ii) players must prepare for their actions in advance, (iii) prepared actions are observable, and (iv) the probability that the prepared actions are actually played is strictly positive but not one (and the probability with which the game is played is sufficiently low as well as the discount factor is sufficiently high), then players manage to cooperate. The mechanism to sustain cooperation works, for example, as follows. As long as the samurai have been showing up at the gate, they continue to do so (to prepare to defend the village). If anyone hides away, however, they stop preparing to defend.

The next section deals with our main model, where there is a fixed deadline to prepare action in the component game. We will show that some cooperation can be sustained in such games (revision games with a deadline), and the basic mechanism to sustain cooperation is essentially the same as in this bandits story.

2.3 *Revision Games with a Deadline - The Main Model*

Consider a normal form game with players $i = 1, ..., N$. Player i's action and payoff are denoted by $a_i \in A_i$ and $\pi_i(a_1, ..., a_N)$, respectively. This game is played at time 0, but players have to prepare their actions in advance, and they also have some stochastic opportunities to revise their prepared actions. Hence, technically the game under consideration is a dynamic game with preparation and revisions of actions, where the normal-form game π is played at the end of the dynamic game (time 0). To distinguish the entire dynamic game and its component π, the former is referred to as a *revision game* and π is referred to as the *component game*.

Specifically, we consider two specifications. In both cases, time is continuous, $-t \in [-T, 0]$ with $T > 0$. At time $-T$, each player i chooses an action from A_i simultaneously. In time interval $(-T, 0]$, revision opportunities arrive stochastically, according to a process defined shortly. There is no cost of revision. At period 0, the payoffs $\pi(a') = (\pi_1(a'_1), \ldots, \pi_N(a'_N))$ materialize, where a'_i is i's finally-revised action.

1. *Synchronous revision game*: There is a Poisson process with arrival rate $\lambda > 0$ defined over the time interval $(-T, 0]$. At each arrival, each player i chooses an action from A_i simultaneously. We analyze this case in the present paper.

2. *Asynchronous revision game*: For each player i, there is a Poisson process with arrival rate $\lambda_i > 0$ defined over the time interval $(-T, 0]$. At each arrival, i chooses an action from A_i. We analyze this case in Kamada and Kandori (2011a).

We assume that players observe all the past events in the revision game, and analyze subgame perfect equilibria. In synchronous revision games, if the component game has a unique pure Nash equilibrium, one obvious subgame perfect equilibrium is the strategy profile in which players choose a static Nash action at time $-T$, and they do not revise their actions until time 0. In what follows, we show that, under some regulatory conditions, revision games have other subgame perfect equilibria, where players are better off than in the static Nash equilibrium.

2.4 Characterization of Optimal Trigger Strategy Equilibrium

In this section, we consider the case of synchronous moves. We restrict ourselves to two players with one-dimensional continuous action space. This case subsumes, for example, good exchange games, Cournot duopolies, Bertrand competitions, and so forth. These applications are discussed in Section 2.5. We assume two players, but this is just to simplify the exposition: Our results easily extend to the case of N players. The assumption of continuous actions is discussed in a great depth in Section 2.6.

Consider a general two-person symmetric component game with action $a_i \in A_i$

and payoff function π_i. Two players are denoted $i = 1, 2$, and a player's action space A_i is a convex subset (an interval) in $\mathbb{R}$: Examples include $A_i = [\underline{a}_i, \overline{a}_i]$ or $[0, \infty)$. Symmetry means $A_1 = A_2 =: A$ and $\pi_1(a, a') = \pi_2(a, a')$ for all $a, a' \in A$.[12] We assume that the component game has a unique symmetric pure Nash equilibrium (a^N, a^N), whose payoff is $\pi^N := \pi_i(a^N, a^N)$. Here we confine our attention to symmetric revision game equilibrium $x(t)$ that uses the "trigger strategy." The action path $x(t)$ means that, when a revision opportunity arrives at time $-t$, players are supposed to choose action $x(t)$, given that there has been no deviations in the past. If any player deviates and does not choose the prescribed action $x(t)$, then in the future players prepare the Nash equilibrium action of the component game a^N, whenever a revision opportunity arrives. This is what we mean by the trigger strategy in revision games. Below we identify the optimal symmetric equilibrium in the class of trigger strategy equilibria. By the "optimal equilibrium" in a given class of equilibria, we mean that the strategy profile achieves (ex ante) the highest payoffs in that class. Let the symmetric payoff function be

$$\pi(a) := \pi_1(a, a) = \pi_2(a, a),$$

and define the best symmetric action $a^* := \arg\max_{a \in A} \pi(a)$ and let $\pi^{*:} = \pi(a^*)$ denote the highest symmetric payoff.[13] We assume the following regularity conditions. Unless otherwise noted, *these assumptions are imposed only in this section.*

1. **A1:** A pure symmetric Nash equilibrium (a^N, a^N) exists, and it is different from the best symmetric action profile (a^*, a^*).

[12] When we write $\pi_i(x, x')$, x is player i's action and x' is player $-i$'s action.

[13] Assumption A2 that we state shortly ensures that all these pieces of notation are well-defined.

2. **A2:** The payoff function π_i for each $i = 1, 2$ is twice continuously differentiable.[14]

3. **A3:** There is a unique best reply $BR(a)$ for all $a \in A$.

4. **A4:** At the best reply, the first and second order conditions are satisfied: For each $i = 1, 2$,
$$\frac{\partial \pi_i(BR(a), a)}{\partial a_i} = 0, \quad \frac{\partial^2 \pi_i(BR(a), a)}{\partial^2 a_i} < 0.$$

5. **A5:** $\pi(a)$ is strictly increasing if $a < a^*$ and strictly decreasing if $a^* < a$.

6. **A6:** The gain from deviation
$$d(a) := \pi_i(BR(a), a) - \pi_i(a, a) \tag{2.1}$$
is strictly decreasing if $a < a^N$ and strictly increasing if $a^N < a$.

A *trigger strategy equilibrium* is characterized by its equilibrium path (revision plan) $x : [0, T] \to A$ (recall that $x(t)$ denotes the equilibrium action to be taken when a revision opportunity arrives at time $-t$). The expected payoff at the beginning of the game (i.e., at time $-T$) associated with x is defined by

$$V(x) := \pi(x(T))e^{-\lambda T} + \int_0^T \pi(x(t))\lambda e^{-\lambda t} dt. \tag{2.2}$$

We say that a path x is *feasible* if the expected payoff (2.2) is well-defined. Since (2.2) represents an expected payoff, the second term in (2.2) should be regarded as

---

[14] When A is not an open set, we assume that there exists an open interval $\tilde{A}$ such that $A \subset \tilde{A}$ and π_i can be extended to a function $\tilde{\pi}_i$ over $\tilde{A} \times \tilde{A}$ that is twice continuously differentiable, i.e. $\tilde{\pi}_i(a, a') = \pi_i(a, a')$ if $(a, a') \in A \times A$.

Lebesgue integral. Consequently the *set of feasible paths* is formally defined by

$$X := \{x : [0, T] \to A \mid \pi \circ x \text{ is measurable}\}.$$

Given a feasible path $x \in X$, the incentive constraint at time t is

$$\text{(IC(t))}: d(x(t))e^{-\lambda t} \leq \int_0^t \left(\pi(x(s)) - \pi^N\right) \lambda e^{-\lambda s} ds, \tag{2.3}$$

where $d(x(t))$ represents the gain from deviation (see (2.1)). The set of trigger strategy equilibrium paths is formally defined as

$$X^* := \{x \in X \mid \text{IC(t) holds for all } t \in [0, T]\}.$$

Thus, by optimal path we mean the path that achieves (ex ante) the highest payoff with in X^*.

First, we show formally that an optimal trigger strategy equilibrium path exists and it is differentiable (Proposition 1). Based on this, we then derive a differential equation to characterize the optimal path (Theorem 1).

Proposition 1. *There exists an optimal trigger strategy equilibrium $\overline{x}(t)$ ($V(\overline{x}) = \max_{x \in X^*} V(x)$) which is (i) continuous for all t, (ii) differentiable in t when $\overline{x}(t) \neq a^N, a^*$. Furthermore, $\overline{x}(t)$ satisfies the following binding incentive constraint if $\overline{x}(t) \neq a^*$:*

$$d(\overline{x}(t))e^{-\lambda t} = \int_0^t \left(\pi(\overline{x}(s)) - \pi^N\right) \lambda e^{-\lambda s} ds. \tag{2.4}$$

(Sketch of the proof): This proposition is proved by a series of propositions in Appendix A.1. First, we show that the optimal trigger strategy equilibrium exists and it is continuous in t. Then we use the continuity to show that it is differentiable.

The proofs rely on the following three elementary technical facts:

1. *For a collection of countably many measurable functions $\pi^n(t)$, $n = 1, 2, ...$, $\sup_n \pi^n$ is measurable.* We use this fact to construct a candidate optimal payoff $\overline{\pi}(t)$ that is measurable (as the supremum of a sequence $\pi(x^n(t))$, where x^n is a sequence in X^* whose payoffs approach $\sup_{x \in X^*} V(x)$: Proposition 33). Then, pretending that this is the optimal payoff, we construct the *candidate* optimal path $\overline{x}(t)$ by the binding "pseudo incentive constraint"

$$d(\overline{x}(t))e^{-\lambda t} = \int_0^t \left(\overline{\pi}(s) - \pi^N\right) \lambda e^{-\lambda s} ds. \tag{2.5}$$

   Note that we have yet to show that this implies the *true* incentive constraint.

2. *Lebesgue integral $\int_0^t f(s)ds$ is continuous in t for any measurable function f.* This fact shows that the right-hand side of the above equation (2.5), whose integrand is measurable by Step 1, is continuous in t, leading to the continuity of $\overline{x}(t)$ (Proposition 33). We also show that $\overline{\pi}(t) \leq \pi(\overline{x}(t))$ so that the pseudo incentive constraint (2.5) implies the true incentive constraint $d(\overline{x}(t))e^{-\lambda t} \leq \int_0^t \left(\pi(\overline{x}(s)) - \pi^N\right) \lambda e^{-\lambda s} ds$. We go on to show that this weak inequality is actually satisfied with equality (Proposition 34), so that we have the binding incentive constraint (2.4).

3. *Lebesgue integral $\int_0^t f(s)ds$ is differentiable in t if f is continuous.*[15] This fact shows that the right-hand side of the binding incentive constraint (2.4), whose integrand is continuous by Step 2, is differentiable in t, and this leads to the differentiability of $\overline{x}$ (Proposition 34). Q.E.D.

[15] When f is continuous, $\int_0^t f(s)ds$ is equal to Riemann integral and this is just the well-known fundamental theorem of calculus: $\frac{d}{dt} \int_0^t f(s)ds$ exists and equal to $f(t)$.

Next we show that there is *essentially a unique optimal path*. Note first that there are in fact multiple optimal paths which attain the same expected payoff. Let $\overline{x}(t)$ be the optimal trigger strategy equilibrium path identified by the previous proposition. Then,

$$x(t) := \begin{cases} a^N & \text{if } t \text{ is in a measure zero set} \\ \\ \overline{x}(t) & \text{otherwise} \end{cases} .$$

is also a trigger strategy equilibrium path that achieves the same expected payoff as $\overline{x}(t)$ does. However, it is easy to show that the following is true.

Proposition 2. *The optimal path is essentially unique: If $y(t)$is an optimal trigger strategy equilibrium path, then $y(t) = \overline{x}(t)$ almost everywhere, where $\overline{x}(t)$ is the optimal path that satisfies the binding incentive constraint (2.4).*

The proof is given in Appendix A.2. Hereafter, the continuous and differentiable optimal path $\overline{x}(t)$ identified in Proposition 1 is referred to as *the essentially unique optimal path*, or simply as *the optimal path*. Now we are ready to state our main result in this section: The optimal path is characterized by a differential equation.

Theorem 1. *The optimal path $\overline{x}(t)$ is the unique path with the following properties: (i) it is continuous in t and departs a^N at $t = 0$ (i.e., $\overline{x}(t) = a^N$ if and only if $t = 0$), (ii) for $t > 0$, it solves differential equation*

$$\frac{dx}{dt} = \frac{\lambda \left(d(x) + \pi(x) - \pi^N \right)}{d'(x)} =: f(x) \tag{2.6}$$

until $\overline{x}(t)$ hits the optimal action a^*, and (iii) if $\overline{x}(t)$ hits the optimal action a^*, it satays *there ($\bar{x}(t) = a^*$ implies $\bar{x}(t') = a^*$ if $t' > t$). Furthermore, if T is large enough, $\overline{x}(t)$ always hits the optimal action a^* at a fixed finite time,*

$$t(a^*) := \lim_{a \to a^N} \int_a^{a^*} \frac{1}{f(x)} dx \tag{2.7}$$

which is independent of T.

(Sketch of the proof): Technical details can be found in Appendices. Let us now confine our attention to the case $a^N < a^*$ (a symmetric proof applies to the case $a^* < a^N$). Lemma 2 in Appendix refapppaper1a implies that the essentially unique optimal path lies between the Nash and optimal actions ($\overline{x}(t) \in [a^N, a^*]$ for all t). By differentiating the binding incentive constraint (2.4), we obtain a differential equation (2.6) when $d'(x) \neq 0$. Under Assumption A6, we have $d'(a^N) = 0$ and $d'(x) > 0$ if $x \neq a^N$ (d' can be shown to exist (Lemma 3)). (Recall that $d'(a^N) = 0$ is the first order condition that the gain from deviation $d(x)$ is minimized at the Nash action $x = a^N$.) Thus we have obtained a differential equation on an open domain $(x, t) \in (a^N, a^* + \varepsilon) \times (-\infty, \infty)$, for some $\varepsilon > 0$.[16] Note well that the domain excludes the Nash action a^N, where $f(a^N)$ is not defined because $d'(a^N) = 0$.

The optimal path $\overline{x}(t)$ satisfies the following conditions:

- (i) it lies in $[a^N, a^*]$ for all t,
- (ii) it is continuous in t,
- (iii) it follows the differential equation when $x \in (a^N, a^*)$, and
- (iv) it starts with Nash action a^N at $t = 0$.

[16] If a^* is a boundary point of A, extend $f(x)$ to $(a^N, a^* + \varepsilon)$ by any continuously differentiable function and apply the same proof in what follows. This is possible under A2 and footnote 2.

It turns out that there are multiple paths which satisfy conditions (i)-(iv). For example, *trivial constant path* $x(t) \equiv a^N$ satisfies those conditions. In what follows, we identify all paths that satisfy conditions (i)-(iv) and find the optimal one among them.

The crucial step is to show that there is a *non-trivial* path to satisfy (i)-(iv). Is there any solution to $dx/dt = f(x)$ which departs from a^N and reach some action $a^0 \in (a^N, a^*]$ at some finite time? The answer is positive if and only if the following *finite time condition*

$$t(a^0) := \lim_{a \to a^N} \int_a^{a^0} \frac{1}{f(x)} dx < \infty \tag{2.8}$$

is satisfied. As we will show, $t(a^0)$ represents the time for a solution to the differential equation $\frac{dx}{dt} = f(x)$ to travel from a^N to a^0. The reason is the following. Under Assumptions A5 and A6, $f(x) = \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)} > 0$ when $x \in (a^N, a^*]$. Hence any solution $x(t)$ to $\frac{dx}{dt} = f(x)$ is strictly increasing in t. Therefore, $x(t)$ has inverse function $t(x)$, and its derivative is given by $\frac{dt}{dx} = \frac{1}{f(x)}$. This implies that $t(a^0) = \lim_{a \to a^N} \int_a^{a^0} \frac{dt}{dx} dx$ represents the time for a solution to the differential equation $\frac{dx}{dt} = f(x)$ to travel from a^N to a^0.

It is straightforward to check that this finite time condition (2.8) is satisfied for any $x^0 \in (a^N, a^*]$, under our assumptions (Lemma 6 in Appendix A.3). Given those observations, all paths that satisfy (i)-(iv) can be written as follows:

$$x_\tau(t) := \begin{cases} a^N & \text{if } t \in [0, \tau] \\ \\ x^*(t-\tau) & \text{if } t \in (\tau, \tau + t(a^*)) \\ \\ a^* & \text{if } t \in [\tau + t(a^*), \infty) \end{cases},$$

where $x^*(t)$ is the solution to $dx/dt = f(x)$ with boundary condition $x^*(t(a^*)) = a^*$. This path $x_\tau(t)$ departs from a^N at time τ, follows the differential equation, and then hits the optimal action a^* and stays there. (More precisely, we must consider the restriction of $x_\tau(t)$ on $[0, T]$.)

Those paths obviously satisfy (i)-(iv). Next we show the converse: any path satisfying (i)-(iv) is equal to $x_\tau(t)$ for some $\tau \in [0, \infty]$. This comes from the standard result in differential equation: $dx/dt = f(x)$ defined on an open domain $(x, t) \in (a^N, a^* + \varepsilon) \times (-\infty, \infty)$ has a unique solution given any boundary condition, if $f(x)$ is continuously differentiable. Under our assumptions, it is easy to check that $f(x)$ is indeed continuously differentiable on (a^N, a^*) (Lemma 4 in Appendix A.3). The uniqueness of the solution then implies that any path satisfying (i)-(iv) is equal to $x_\tau(t)$ for some $\tau \in [0, \infty]$.[17]

Among the paths $x_\tau(t)$, $\tau \in [0, \infty]$ the one that departs from a^N immediately (i.e., $x_0(t)$) obviously has the highest payoff. Therefore the optimal path is given by the restriction of $x_0(t)$ on $[0, T]$, which has the stated properties in Theorem 1. Q.E.D.

In the optimal trigger strategy equilibrium identified in the previous theorem, players act as follows. Recall that $\overline{x}(t)$ is the action to be taken at time $-t$. If the time horizon is long enough, (i.e., if $T \geq t(a^*)$), players start with the best action

---

[17] Formal proof goes as follows. The trivial path, which satisfy (i)-(iv), is equal to x_τ with $\tau = \infty$. Consider any non-trivial path $x^0(t)$ that satisfy (i)-(iv), where $x^0(t^0) =: a^0 \in (a^N, a^*)$ for some t^0. Define $t' := t^0 - \lim_{a \to a^N} \int_a^{a^0} \frac{1}{f(x)} dx$, so that $x^*(t - t')$ hits a^0 at time t^0. The uniqueness of the solution to the differential equation (for boundary condition $x(t^0) = a^0$) implies $x^0(t) = x^*(t - t')$. If $t' \geq 0$, we obtain the desired result $x^0(t) = x_\tau(t)$ for $\tau = t'$. If $t' < 0$, $x^0(0) = x^*(-t') > a^N$ and $x^0(0)$ cannot satisfy (iv). (We have $x^*(-t') > a^N$ because we are considering the case $a^N < a^*$, where the solution $x^*(t)$ is strictly increasing).

a^*, and even if a revision opportunity arrives, they do not revise their actions until time $-t(a^*)$ is reached. After that, if a revision opportunity arrives, they choose an action $\overline{x}(t)$, which is closer to the Nash action. The closer the timing of the revision opportunity is to the end of the game, the closer the revised action $\overline{x}(t)$ is to the Nash equilibrium. At the end of the game, the actions chosen at the last revision opportunity are implemented. Hence the best symmetric trigger strategy equilibrium induces a probability distribution of actions in between the best and Nash actions. The nature of this equilibrium distribution will be examined in the following propositions (Propositions 3 and 4).

One might expect that the outcome of the component game, and hence the payoffs, may depend on the arrival rate λ. The next proposition, which is actually nothing but a simple observation, shows that this is not the case. To state the proposition, we need to introduce the following notation. We denoted the first time to hit the optimal action by $t(a^*)$ in Theorem 1 (see (2.7)), but to explicitly show its dependence on arrival rate λ, let us now denote it by $t_\lambda(a^*)$.

Proposition 3. ***(Arrival Rate Invariance)*** *Under the best symmetric trigger strategy equilibrium, the probability distribution of action profile at period 0 is independent of the Poisson arrival rate λ, provided that the time horizon T is long enough. Specifically, Let $t_\lambda(a^*)$ be the (first) time to reach the optimal symmetric action, stated in Theorem 1. Then, as long as $t_\lambda(a^*) \leq T$, the probability distribution of the action profile at period 0 is independent of λ.*

Proof. Consider λ such that $t_\lambda(a^*) \leq T$ and call it Model 1. Rewrite this model by changing the time scale in such a way that one unit of time in Model 1 corresponds to λ units in the new model. Under the new time scale, the model is identical to the revision game with arrival rate 1 and time horizon λT. Call it Model 2. The best symmetric trigger strategy equilibrium path in Model 1 should be transformed into the best symmetric trigger strategy equilibrium path in Model 2. In particular, it must be the case that $t_1(a^*)$, the first time the optimal path hits a^* in Model 2, is equal to $\lambda t_\lambda(a^*)$, and this is smaller than the time horizon of Model 2 (λT). Hence in Model 2, there is no revision of action in the best symmetric trigger strategy equilibrium for $t \in [-\lambda T, -t_1(a^*)]$, and therefore the probability distribution of action profile at $t = 0$ is unchanged if the game starts at $-t_1(a^*)$ (instead of $-\lambda T$). Hence, the probability distribution of action profile at $t = 0$ under *any* arrival rate λ such that $t_\lambda(a^*) \leq T$ is equal to the probability distribution under arrival rate 1 and time horizon $t_1(a^*)$. □

Note that the fact that payoffs realize only at the deadline $t = 0$ played a crucial role in this proposition (otherwise, the expected payoffs would be affected by the arrival rate and the discount factor). Proposition 3 shows the following attractive feature of revision games: we can obtain a unique prediction that does not depend on the fine detail, namely the arrival rate λ of the revision opportunities. In particular, even if λ is sufficiently high (so that there are many chances to revise actions right before the component game), the expected outcome in the component game is the same as in the case of low λ.

The proof also shows how to calculate the cumulative distribution function of symmetric action a, denoted by $F(a)$. Again we consider the case with $a^N < a^*$ (a symmetric argument applies to the other case). Let $x_1(t)$ be the optimal trigger strategy equilibrium path under $\lambda = 1$, and denote the time for $x_1(t)$ to hit $a \in$

$[a^N, a^*]$ by $t_1(a)$. The latter is given by equation (2.8) for $\lambda = 1$:

$$t_1(a) := \lim_{a' \to a^N} \int_{a'}^{a} \frac{d'(x)}{d(x) + \pi(x) - \pi^N} dx. \tag{2.9}$$

For $a \in (a^N, a^*)$, $F(a) = \int_{\{t|x_1(t) \leq a\}} e^{-t} dt = \int_0^{t_1(a)} e^{-t} dt = 1 - e^{-t_1(a)}$. The first equality follows from the fact that the density of action $x_1(t) \leq a$ is the product of

- 1 (the density of revision at time t) and
- e^{-t} (the probability that the revised action at time t, $x(t)$, will never be revised again).

At a^*, the distribution function $F(a)$ jumps by $e^{-t_1(a^*)}$ and $F(a) = 1$ for $a \geq a^*$. This means that the optimal action a^* is played with probability $e^{-t_1(a^*)}$. This is the probability that no revision opportunity arises under $\lambda = 1$. Below we summarize our arguments.

Proposition 4. *Suppose that the time horizon is long enough so that the efficient action a^* is chosen at the beginning of the revision game, under the best symmetric trigger strategy equilibrium. When $a^N < a^*$, the cumulative distribution function of the symmetric action realized at $t = 0$ is given by*

$$F(a) = \begin{cases} 0 & \textit{if } a < a^N \\ 1 - e^{-t_1(a)} & \textit{if } a^N \leq a < a^* \\ 1 & \textit{if } a^* \leq a \end{cases},$$

where $t_1(a)$ is given by (2.9) and represents the time for the best symmetric trigger strategy action path to reach $a \in [a^N, x^*]$, when the arrival rate is $\lambda = 1$. When $a^* < a^N$, it is *given by*

$$F(a) = \begin{cases} 0 & \textit{if } a < a^* \\ e^{-t_1(a)} & \textit{if } a^* \leq a \leq a^N \\ 1 & \textit{if } a^N < a \end{cases} .$$

2.5 Applications

In this section, we use the general framework given in the previous section to analyze various games of interest. Specifically, we use the differential equation provided in Theorem 1 to analyze good exchange games (prisoner's dilemma), Cournot duopolies, Bertrand competition with product differentiation, and election campaign. Unless otherwise noted, the component games in these examples satisfy Assumptions A1-A6.

We will be considering two measures of the degree of cooperation. Let the expected payoff from the optimal trigger strategy equilibrium when T is sufficiently large be $\tilde{\pi}$. The two measures are:

$$R := \frac{\tilde{\pi}}{\pi(a^*)} \quad \text{and} \quad \tilde{R} := \frac{\tilde{\pi} - \pi^N}{\pi(a^*) - \pi^N}.$$

The first one simply takes the ratio of the equilibrium payoff to the fully collusive payoff. The second is a conservative one, which compares the improvement of the payoff relative to the Nash payoff (the static equilibrium payoff) with the maximum possible payoff improvement.

2.5.1 Good Exchange Game

For each player $i = 1, 2$, let the payoff function be $\pi_i(a_i, a_{-i}) = a_{-i}^k - c \cdot a_i^2$, where $c > 0$, $k \in (0, 2)$, and the action space is $a_i \in [0, \infty)$. This game represents the following situation. Two players $i = 1, 2$ exchange goods they produce. That is, player 1 produces one unit of good and gives it to player 2 (and *vice versa*). The quality of the good player i produces is equal to a_i^k, and i incurs a convex cost $c \cdot a_i^2$ to provide a good with quality a_i. Alternatively, one can interpret a_i^k as the quantity of goods i provides given effort level a_i and assume that $c \cdot a_i^2$ is the cost to exert the effort level a_i. Note that $a_i = 0$ is the dominant strategy, while the best symmetric action $a^* = \left(\frac{k}{2c}\right)^{\frac{1}{2-k}}$ is strictly positive. Hence this can be regarded as a version of the prisoner's dilemma game with continuous actions. Notice that the larger k is, the smaller is the gain from a very small amount of action (i.e., $a^k < a^{k'}$ if $k > k'$ and a is small).

The differential equation in Theorem 1 for this example is

$$\begin{aligned} \frac{dx}{dt} &= \frac{\lambda\left(d(x) + \pi(x) - \pi^N\right)}{d'(x)} \\ &= \frac{\lambda(cx^2 + (x^k - cx^2) + 0)}{(cx^2)'} = \frac{\lambda x^k}{2cx}. \end{aligned}$$

Note that, since 0 is a dominant action, the Nash payoff π^N is zero, and the best reply to any action is zero: $BR(a_{-i}) = 0$. The latter implies $d(x) = cx^2$. The above differential equation has a solution $x(t) = \left(\frac{2-k}{2c}\lambda t\right)^{\frac{1}{2-k}}$ which departs from 0 (the Nash action) at time $t = 0$. The time at which $x(t)$ reaches the best action, denoted $t(a^*)$, can be calculated by (2.7), but it is equivalently obtained by solving $a^* = x(t(a^*)) = \left(\frac{2-k}{2c}\lambda t(a^*)\right)^{\frac{1}{2-k}}$. We summarize our findings as follows.

Proposition 5. *In the good exchange game, the optimal trigger strategy equilibrium, $x(t)$,*

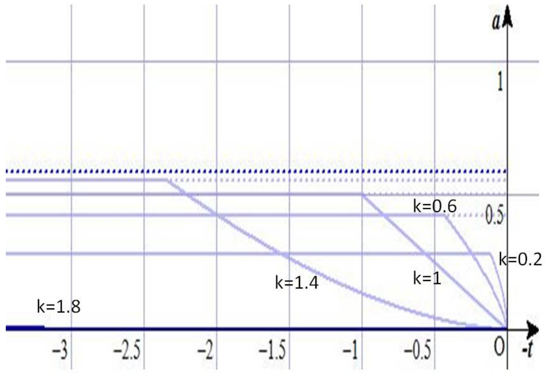

Figure 2.2: The optimal path for the good exchange game: $\lambda = 1$.

is characterized by

$$x(t) = \begin{cases} \left(\frac{2-k}{2c}\lambda t\right)^{\frac{1}{2-k}} & \text{if} \quad t < t(a^*) \\ a^* = \left(\frac{k}{2c}\right)^{\frac{1}{2-k}} & \text{if} \quad t(a^*) \leq t \end{cases},$$

where $t(a^*) = \frac{k}{\lambda(2-k)}$.

The path characterized in Proposition 5 is depicted in Figure 2.2.

In Figure 2.2, as k increases, the time that the path departs from the optimal action ($t(a^*)$) becomes larger, and the path approaches 0 more quickly. These observations suggest that it is more difficult to cooperate when parameter k is large. This is in line with our earlier observation that a larger k implies a smaller gain from cooperation around the Nash equilibrium $(0, 0)$ (hence it is more difficult to sustain cooperation).

Corollary 1. *In the good exchange game, R ($= \tilde{R}$) is decreasing in k. It approaches 1 as $k \searrow 0$, and approaches 0 as $k \nearrow 2$.*

The proof is straightforward calculation and therefore omitted. We can also explicitly calculate the expected payoff when the parameter k is just in the middle of $(0,2)$. That is, when $k = 1$, the expected payoff is $\frac{1}{2ec^2}$, which implies $R = \tilde{R} = \frac{2}{e} \cong 0.74$ (this is independent of the value of c). The revision game attains 74% of the fully cooperative payoff in this case.

Although there cannot be any cooperation in the Nash equilibrium of the component game, in revision games players can achieve around three fourths of the fully cooperative payoff. The degree of cooperation decreases as the gain from small cooperation decreases. Thus, a higher degree of overall cooperation is more difficult to achieve the less gain there is given a small amount of cooperation.

2.5.2 *Cournot Duopoly: Collusion Through Output Adjustment Achieves 97% of The Monopoly Profit*

In this subsection, we consider a Cournot duopoly game with a linear demand curve $P = a - b(q_i + q_j)$ (q_i denotes agent i's quantity) and constant (and identical) marginal cost c. Hence the (component game) payoff function for player i is $\pi_i = \left(a - b(q_i + q_j) - c\right) q_i$. We suppose $a > c > 0$ and $b > 0$. The differential equation is

$$\begin{aligned}\frac{dq}{dt} &= \frac{\lambda\left(d(q) + \pi(q) - \pi^N\right)}{d'(q)} \\ &= \frac{\lambda}{18}(q - 5\frac{a-c}{3b}).\end{aligned}$$

This comes from $d(q) = \frac{(a-c-3bq)^2}{4b}$, $\pi(q) = (a - c - 2bq)q$, and $\pi^N = \frac{(a-c)^2}{9b}$. The differential equation admits a simple solution $q(t) = \frac{a-c}{3b}(5 - 4e^{\frac{\lambda}{18}t})$ which departs from the Cournot Nash output $q^N = \frac{a-c}{3b}$ at $t = 0$, and this path hits the optimal

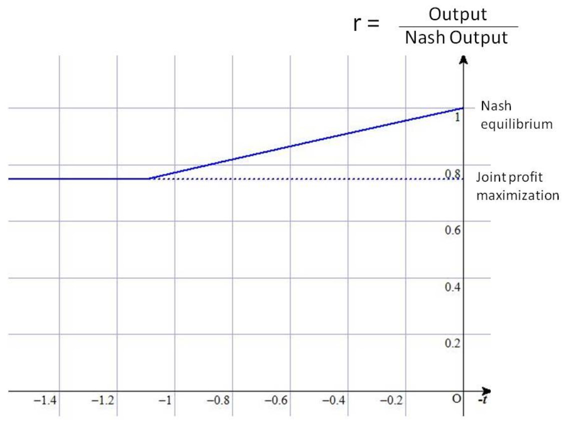

Figure 2.3: The optimal path for Cournot duopoly: $\lambda = 1$.

output $q^* = \frac{a-c}{4b}$ at $t(q^*) = \frac{18}{\lambda} \ln\left(\frac{17}{16}\right)$. Therefore, we have obtained the following.

Proposition 6. *In the Cournot duopoly game, the optimal trigger strategy equilibrium, $q(t)$, is characterized by*

$$q(t) = \begin{cases} \frac{a-c}{3b} \cdot r(t) & \text{if} \quad t < t(q^*) \\ q^* = \frac{a-c}{4b} & \text{if} \quad t(q^*) \leq t \end{cases},$$

where we let $t(q^*) = \frac{18}{\lambda} \ln\left(\frac{17}{16}\right)$ and $r(t) = 5 - 4e^{\frac{\lambda}{18}t}$.

Note that $r(t)$ is the ratio of the equilibrium quantity at time $-t$ to the static equilibrium quantity, $\frac{a-c}{3b}$.

When the firms collude, they produce less than the Nash quantity, and therefore the optimal trigger equilibrium path that we characterize is *deceasing* in t. That is, the ratio r starts from 1 (due to the initial condition), and decreases monotonically to $\frac{3}{4}$. The path of r with respect to t is depicted in Figure 2.3.

Next, we consider the welfare implication of the revision game of the Cournot duopoly game. One can compute the equilibrium expected payoff, and it turns out that a surprisingly high degree of collusion can be achieved in this game. The next corollary implies that, when two firms gradually adjust their outputs before the market is open (and if they closely monitor each other's output adjustment processes), then they can achieve almost 97% of the fully collusive profit (this amounts to 72% of the gain relative to the Nash profit). We emphasize that those numbers are independent of the position and the slope of the demand curve (*a* and *b*) and the marginal cost *c* (and also independent of the arrival rate λ of revision opportunities, as Proposition 3 shows).

Corollary 2. *In the Cournot duopoly game,* $R \cong 0.968$ *and* $\tilde{R} \cong 0.714$*, independent of the values of parameters a, b, and c.*

The following story might fit the Cournot revision game. Two fishing boats depart from a harbor early in the morning, and they must return when the fish market at the harbor opens at 6:00 am. They would like to restrict their catch so as to increase the price at the fish market. They first catch a small amount of fish (the collusive quantity). They are operating side by side, closely monitoring each other's behavior. Fish schools occasionally visit them, by Poisson process.[18] The arrival rate is $\lambda = 0.1$ (and the time unit is a minute), so that a fish school comes every ten minutes on average. Since $t(q^*) = \frac{18}{\lambda} \ln\left(\frac{17}{16}\right) = 10.912$ minutes, they do not catch any additional fish until eleven minutes before the market opens. In the last eleven minutes, whenever a fish school visit them, they catch additional fish. If no fish school visits, they deliver the collusive amount to the market. If a fish school comes right before 6:00am, they catch Nash amount. If the last visit of a fish school is somewhat before, they catch a smaller amount. On average,

[18] Pun not intended.

they encounter only one revision opportunity in the last eleven minutes (because $\lambda \cdot t(q^*) \simeq 1$), and they can achieve 97% of fully collusive profit.

2.5.3 *Bertrand Competition with Product Differentiation*

In this subsection we consider a Bertrand competition with product differentiation. We would like to examine how the degree of product differentiation affects the possibility of collusion in the revision game. To this end, we employ the Hotelling model of spatial competition with price setting firms. This model has an advantage of incorporating the case with no product differentiation as a special case.

A continuum of buyers are distributed uniformly over $[0,1]$. Two firms $i = 1,2$ are located at positions 0 and 1, respectively. A buyer at x receives payoff $a - d|x-y| - p$ if she buys from the firm at y with price p, where $d \in [0, \frac{2}{3}a)$. Notice that d is the cost of transportation for the buyers, and it measures the degree of product differentiation. In particular, $d = 0$ corresponds to the case in which there is no product differentiation, and a high d means a high degree of product differentiation. If the buyer does not buy anything, the payoff is 0. No buyer would want to buy two or more products. Each firm's marginal cost is normalized at 0, and the firm's payoff is the average revenue from a buyer.[19]

For relatively high product differentiation, namely for $d \in (\frac{2}{7}a, \frac{2}{3}a)$, the differ-

[19] Firm i's payoff function has a kink for example when $p_i = p_{-i} - d$, so A2 is violated. However, A2 can be shown to be satisfied at relevant regions ((p,p) and $(BR(p),p)$ for all p weakly between the Nash price and the fully collusive price) when $d \in (\frac{2}{7}a, \frac{2}{3}a)$. When $d \in (0, \frac{2}{7}a)$, A2 is not satisfied at $(BR(p),p)$ for one p on the equilibrium path, but one can show that the optimal path is characterized by solving two differential equations, one for prices below such p and the other for prices above such p.

ential equation is

$$\frac{dp}{dt} = \lambda \frac{p+3d}{2}.$$

This comes from $\pi^N = \frac{d}{2}$, $\pi(p) = \frac{p}{2}$, and $d(p) = \frac{(p-d)^2}{8d}$.

When $d \in (0, \frac{2}{7}a]$, however, the degree of differentiation is so small that when the opponent sets a price close to the best collusive price, the best reply is to set a price just enough to take all the buyers, that is, $BR(p) = p - d$ and hence $d(p) = \frac{p}{2} - d$. π^N and $\pi(p)$ are the same as before. Using these formulas, the differential equation in this case can be written as

$$\frac{dp}{dt} = \lambda(2p - 3d).$$

Overall, we obtain the following:

Proposition 7. *In the Bertrand competition game, the optimal trigger strategy equilibrium path, $p(t)$, is characterized as follows:*

1. *If $d \in (\frac{2}{7}a, \frac{2}{3}a)$,*

$$p(t) = \begin{cases} d\left(4e^{\lambda\frac{t}{2}} - 3\right) & \text{if} \quad t < t(p^*) \\ p^* = a - \frac{d}{2} & \text{if} \quad t(p^*) \le t \end{cases},$$

where $t(p^*) = \frac{2}{\lambda}\ln\left(\frac{a}{4d} + \frac{5}{8}\right)$.

2. *If $d \in (0, \frac{2}{7}a]$,*

$$p(t) = \begin{cases} d\left(4e^{\lambda\frac{t}{2}} - 3\right) & \text{if} \quad t < t^1 \\ d\left(\frac{8}{27}e^{2\lambda t} + \frac{3}{2}\right) & \text{if} \quad t^1 \le t < t^2 \\ p^* = a - \frac{d}{2} & \text{if} \quad t^2 \le t \end{cases},$$

where $t^1 = \frac{2}{\lambda}\ln\left(\frac{3}{2}\right)$ *and* $t^2 = \frac{2}{\lambda}\ln\left(\frac{a}{d} - 2\right)$.

3. *If* $d = 0$, $p(t) = 0$ *for all* t.

The proposition claims that there is a cooperative path if and only if there is a product differentiation. This highlights the importance of Assumption A4. When $d = 0$, the first order condition does not hold at the static Nash equilibrium, so there cannot be a collusive path. The intuition is as follows: If there is no product differentiation, an infinitesimal price cut can increase the profit discontinuously almost to the double whenever the current price is strictly higher than the marginal cost (which is 0 in this example). This is because all buyers switch to the deviating firm. Hence, if the current price is not equal to the Nash equilibrium, the gain from deviation is not of a smaller order in magnitude than the gain from cooperation. This makes cooperation impossible. If there is a product differentiation, however, only a small fraction of buyers switch to the deviating firm, and this makes the cooperation sustainable.

Note also that p is increasing in d. Hence, the more differentiated the products are, the more collusion there is. This makes sense: When the degree of product differentiation is large, the instantaneous gain from deviation when firms set prices close to the Nash price is small relative to the loss from the punishment because firms need to decrease the price a lot to steal the opponent's share. The path characterized in Proposition 7 when $d > 0$ is depicted in Figure 2.4. In the figure, we fix $a = 10$ and draw the optimal paths for $d = 1$, $d = 2$, $d = 3$, and $d = 5$. As expected, the collusive path is close to the best collusive price when the degree of product differentiation is high.

The expected payoff under the optimal trigger strategy path can be computed as follows:

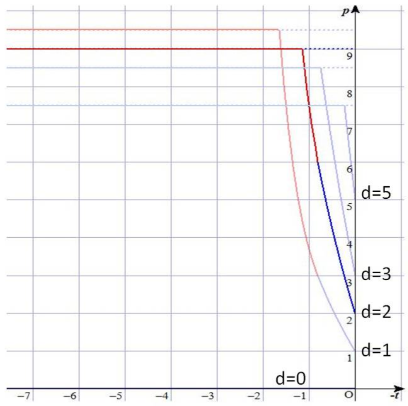

Figure 2.4: The optimal path for Bertrand competition: $\lambda = 1$, $a = 10$.

Corollary 3. *Let $h := \frac{d}{a}$ be the degree of product differentiation. Then, the level of collusion achieved in the revision game, measured by R and $\widetilde{R}$, is expressed as follows.*

1. *If $h \in (\frac{2}{7}, \frac{2}{3})$, $R = \frac{2h(10-7h)}{(2-h)(2+5h)}$ and $\tilde{R} = \frac{8h}{5h+2}$.*

2. *If $h \in (0, \frac{2}{7}]$, $R = \frac{2h(3-7h)}{(2-h)(1-2h)}$ and $\tilde{R} = \frac{2h(2-5h)}{(1-2h)(2-3h)}$.*

3. *If $h = 0$, $R = \tilde{R} = 0$.*

4. *Both R and $\tilde{R}$ are strictly increasing in h.*

The ratios stated in the corollary The ratio of expected payoff relative to the fully collusive payoff is calculated in Table 2.2 for several values of h. The table shows that in the revision game, firms can achieve quite a bit of cooperation to obtain high expected payoffs. For example, if $h = .5$ then, on average, a buyer's willingness to pay to the worse good is 71.4% of that of the preferred good. In such a circumstance, the table shows that 96% (89% even under the conservative

Table 2.2: Degrees of product differentiation and cooperation.

Degree of product differentiation ($r = \frac{d}{a}$)	0	.1	.2	.3	.5	$.6\dot{6}$
$\frac{\text{Expected payoff}}{\text{Collusive payoff}}$ (R)	0	.30	.59	.80	.96	1
$\frac{\text{Expected payoff - Nash payoff}}{\text{Collusive payoff - Nash payoff}}$ ($\tilde{R}$)	0	.22	.48	.69	.89	1

measure of the degree of cooperation) of payoffs can be achieved by the revision game.

2.5.4 *Election Campaign: Policy Platforms Gradually Converge*

In this subsection we consider a simple election model with policy-motivated candidates. The policy space is an interval $[0, 1]$. As in the standard model of policy-motivated candidates, the position of the median voter is unknown, but its distribution is known as the uniform distribution over the policy space, $[0, 1]$. There are two candidates, $i = 1, 2$, where candidate 1 chooses policy y_1 and candidate 2 chooses policy y_2.

Given a policy profile (y_1, y_2), let a random variable $w(y_1, y_2)$ represent the "winner" of the election. Let candidate i's realized payoff be

$$g_i(y_i, y_{-i}) = a \cdot \mathbb{I}_{\{i = w(y_1, y_2)\}} + b(|y_{w(y_i, y_{-i})} - \overline{y}_i|)$$

where $a \in (\frac{1}{2}, 1)$ is a positive constant representing the utility of winning *per se*, and $b(\cdot)$ is a "policy preference term," which depends on the distance between the winner's policy (the policy actually implemented) and candidate i's "bliss point," denoted $\overline{y}_i$. We assume that $\overline{y}_1 = 0$ and $\overline{y}_2 = 1$. That is, candidate 1 is "left wing" and candidate 2 is "right wing."

There are two key assumptions that we impose on this standard election model with policy-motivated candidates: *First*, we assume that the payoff function corresponding to the policy preference term is *convex*. As Kamada and Kojima (2012) discuss, such policy preferences are especially relevant for issues that contain religious content (e.g. same-sex marriage, abortion, gun control, and so forth), as in these policy issues it is natural to assume that a player's utility sharply decreases as the implemented policy departs from her bliss point.[20] Convex utility function implies that a profile $(0, 1)$ Pareto-dominates the Nash profile, so there is a potential room for cooperation in a revision game. For simplicity, we assume the following functional form: $b(z) = \max\{\frac{1}{2} - z, 0\}$.[21] *Second*, we assume that candidate 1 chooses policy y_1 from $[0, \frac{1}{2}]$ and candidate 2 chooses policy y_2 from $[\frac{1}{2}, 1]$. The motivation behind this assumption is that candidate 1 (resp. candidate 2) faces a reputational concern, so that she never wants to set a policy to the right (resp. left) of the middle (Remember that candidate 1 (resp. candidate 2) is "left wing" (resp. "right wing)). Without this assumption, the best response is always to set a policy as close as possible to the other candidate, and thus there is a huge gain by deviating from the profile close to the Nash equilibrium, which makes cooperation impossible in a revision game (by the violation of A4).

The payoff functions are not symmetric as they are, but by redefining actions by

$$x_1 = y_1 \text{ and } x_2 = 1 - y_2$$

[20] See Osborne (1995) for a criticism on the use of concave utility function for preferences over electoral policies.

[21] This functional form does not satisfy Assumption A2, but it is straightforward to check that A2 is satisfied over the relevant domain. The assumption that the candidate is exactly indifferent between two policies that are both further away from her bliss point is made only for the purpose of simplicity, and does not play any crucial role in our argument.

we can retain the symmetry. Notice that the probability of i's winning the election can be calculated as $\frac{1+x_i+x_{-i}}{2}$.

Now, let us explain what the revision game of this election game corresponds to. We interpret the revision phase as the time period for "election campaign." In the revision phase, candidates obtain opportunities to express their policy positions, for example at an open broadcast on radio or television. At each opportunity, candidates can choose their policy announcement that is possibly different from what they have said before (as is often the case). At "time 0" of the revision game, the election takes place, and candidates are committed to implementing her finally announced policy, given that she is elected.[22]

The differential equation for candidate 1's policy platform $y_1(t) = x(t)$ is

$$\frac{dx}{dt} = \lambda \frac{2x - 2a - 7}{4}.$$

This follows from $\pi(x) = \frac{1}{2}\left(a + \frac{1}{2} - x\right)$, $\pi^N = \frac{1}{2}$, and $d(x) = \frac{1}{8}(a - \frac{1}{2} - x)^2$. This has a solution $x(t) = \frac{7+2a-8\cdot e^{\frac{\lambda}{2}t}}{2}$ which departs from Nash action $x^N = \frac{2a-1}{2}$ at $t = 0$.

Proposition 8. *In the election campaign game, the optimal trigger strategy equilibrium, $(y_1(t), y_2(t))$, is characterized by*

$$y_1(t) = \begin{cases} \frac{7+2a-8\cdot e^{\frac{\lambda}{2}t}}{2} & \text{if} \quad t < t^* \\ 0 & \text{if} \quad t^* \leq t \end{cases},$$

where $t^* = \frac{2}{3\lambda} \ln\left(\frac{7}{2} + a\right)$ and $y_2(t) = 1 - y_1(t)$.

The above proposition shows that in the election campaign game, each candi-

[22] This "policy announcement game" is proposed and analyzed in Kamada and Sugaya (2011), in which they analyze the case where candidates cannot announce inconsistent policies while they have an option to announce an "ambiguous policy."

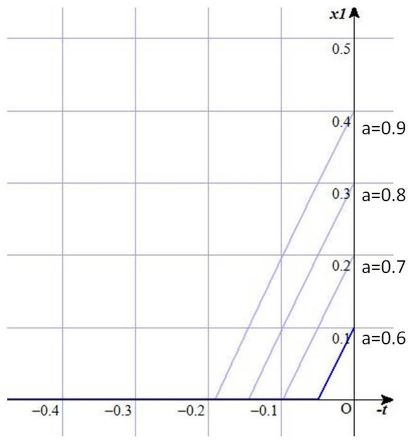

Figure 2.5: The optimal path for election campaign (the path of x_1): $\lambda = 1$.

date starts from announcing their most preferred policies until the time of election becomes close, and then begin catering to the middle in the end. Thus the model captures the well-observed phenomena of candidates changing their policy announcements, moving to the middle when the election is close. The path characterized in Proposition 8 is depicted in Figure 2.5.

We note that this result does not hold if the policy preference term b is concave, as usually assumed in the political science literature. If candidates' policy preferences are convex, they prefer a diverging policy profile $(0, 1)$ to a converging one $(\frac{1}{2}, \frac{1}{2})$. This is because, for example, candidate 1 does not care about the difference between policies $\frac{1}{2}$ and 1 while she perceives a huge difference between policies 0 and $\frac{1}{2}$. This is why there can be a nontrivial equilibrium path.

2.6 Discussion

2.6.1 Robustness of Cooperation

As should be clear at this point, the key to the sustainability of cooperation in revision games is the fact that as the deadline comes close, the gain from defection becomes arbitrarily smaller than the payoff from cooperation. This was made possible because we assume continuous action space and continuous time. First, note that if each player has a dominant action and time is discrete, then by backwards induction it is obvious that the only equilibrium is for each player to play the dominant action at any revision opportunity. Again, drop A1-A6 in this section.

Proposition 9. *Consider a component game with an action set A_i with a strictly dominant action a_i^N for each player i, and consider either of the following two cases:*

1. *A_i is finite.*

2. *There exists $\epsilon > 0$ such that all players are restricted to use strategies that, at any time $-t$, do not condition on what has happened in time $(-t+\epsilon, -t)$.*

Then, whether in synchronous or asynchronous revision games (with homogeneous or heterogeneous arrival rates), there exists a unique subgame perfect equilibrium. In this equilibrium, each player i plays action a_i^N conditional on any history.

For part 1, the proof for the result is straightforward. First observe that if A_i is finite then there exists $\epsilon > 0$ such that given any action of the opponent, a_i^N gives i a payoff at least ϵ greater than any other actions in A_i. This means that, if it is true that each player j prepares an action a_j^N whenever j gets a revision opportunity strictly after time $-t$, then by assumption i's payoff from preparing a_i^N is at least ϵ' greater than preparing any other action for some $\epsilon' > 0$. By continuity of payoffs with respect to probability, this means that there exists $\epsilon'' > 0$ such that i strictly prefers preparing a_i^N to any other action in the time interval $(-t - \epsilon'', -t]$, hence whenever i gets a revision opportunity in $(-t - \epsilon'', -t]$ she prepares a_i^N. This establishes the result.[23] The intuition is that if the time left until the deadline is very little, it is a dominant strategy for players to follow the dominant actions, irrespective of the opponents' strategies.[24] Notice that the above proof is invalid in our main model because there does not exist such $\epsilon > 0$ that we took above.

The restriction on the strategy stated in part 2 describe the situation where there exists a fixed positive "response time," so that any player cannot respond to a defection that has happened in a very close past. The proof is again straightforward. In time $(-\epsilon, 0]$, there is no reason for any player to play an action other than the strictly dominant action, as the preparation in that time interval does not affect the opponents' future behavior at all. Then, no action prepared in $(-2\epsilon, -\epsilon]$ affects the opponent's future behaviors, so players prepare a_i^N in this time interval as well. Going backwards, we get the desired result.

The proposition shows that *both* the continuity of action space *and* time are needed to obtain cooperation in revision games. In this particular sense, cooperation is not a robust result. However it is not clear why this is the robustness that we should consider with the first-order importance. On the other hand, recent experimental results show that economic agents have altruism motives. In our model, cooperation is retained by a very slight addition of such a behavioral

---

[23] Whether or not ϵ'' depends on t does not matter for the result (In our case, we can actually take ϵ'' independent of t). For a formal proof of this, see Lemma 1 in Calcagno, Kamada, Lovo, and Sugaya (2012).

[24] Calcagno and Lovo (2010) obtained a similar result when the component game is a two-player prisoner's dilemma.

element to the model. To illuminate this issue, we focus on the continuity of time and the situation where players are indifferent between very small cooperation and no cooperation, which is interpreted as an existence of incentives to "give away" a very small amount.[25] Consider a very simple example with the following payoff function:

$$\pi_i(a_i, a_j) = 2a_j - \max\{a_i - \epsilon, 0\}, \qquad \epsilon \geq 0, \quad a_i, a_j \geq 0.$$

This is a version of the good exchange game in Subsection 2.5.1, where the cost of cooperation takes a different form. We call this game as a *modified good exchange game*.

The cost term is constant at zero near action 0 if $\epsilon > 0$ but it increases linearly otherwise. Notice that when $\epsilon = 0$, there is only one Nash equilibrium in which each player i plays action 0. On the other hand, when $\epsilon > 0$, there are multiple equilibria. In particular, both $(0, 0)$ and (ϵ, ϵ) are Nash equilibria, where the former gives each player the payoff of 0 but the latter gives $2\epsilon > 0$. Also notice that this payoff function does not satisfy Assumption A4 when $\epsilon = 0$.[26]

Now we consider a discrete time version of synchronous revision game (this specification applies only in this subsection). Time is $-t = \ldots, -2, -1, 0$, and at each period, both players have a revision opportunity with probability $p > 0$. The component game is played at time 0 (assume that the revision opportunity may come also at time 0 (with probability p) before the game is played). We construct the optimal symmetric trigger strategy equilibrium path (analogously defined as was done so far) that converges to (ϵ, ϵ) as $t \to 0$ but triggers to $(0, 0)$ upon deviation.

---

[25] An analogous discussion for continuity of actions can be easily done.

[26] Also, this does not satisfy Assumption A2 when $\epsilon > 0$ (as there is a "kink" at $a = \epsilon$).

A straightforward calculation shows that the path is characterized by[27]:

$$x(t) = \epsilon \cdot \left(\frac{1+p}{1-p}\right)^t.$$

Notice that when $\epsilon = 0$, the path is a trivial one, i.e. $x(t) = 0$ for all t. However, if $\epsilon > 0$, there exists a cooperative path. The nonexistence of cooperative path when $\epsilon = 0$ is straightforward from backwards induction. The existence of cooperative path when $\epsilon > 0$ is that the cost of cooperation does not grow near the Nash action so that players can use the worse equilibria as a threat, and they can use this tiny threat as a foothold for long-run cooperation. This intuition is analogous to the logic of the sustainability of cooperation in Benoit and Krishna (1985), who consider a model of finitely repeated games, and show that an approximate "folk theorem" holds as the horizon becomes infinitely long when each player has multiple Nash payoffs.[28]

Now we turn to our setting with continuous time and smooth payoff function. Specifically, consider a payoff function from the exchange of goods game, $\pi_i(a_i, a_j) = 2a_j - a_i^2$ with $a_i, a_j \geq 0$. There is only one equilibrium at $(0, 0)$, so in the above discrete time setting, there is only one equilibrium in the revision game, by part 2 of Proposition 9. However, recall that there exists a cooperative path when $\epsilon > 0$ in the modified good exchange game, and the sustainability of the path hinges on the fact that the cost of cooperation does not grow near the Nash action so that players can use the worse equilibria as a threat, and they can use this tiny threat as a foothold for long-run cooperation. In the above payoff function, the

[27] Letting $p = \lambda\Delta\tau$, $\tau = (\Delta\tau) \cdot t$ and taking the limit as $\Delta\tau \to 0$, this converges to the optimal path of $x(\tau) = \epsilon \cdot e^{2\lambda\tau}$ in continuous time, which can also be obtained by a direct computation.

[28] Strictly speaking, Benoit and Krishna (1985) consider the case of flow payoffs (players receive payoffs each period) thus the two settings are slightly different from each other.

cost of cooperation, a_i^2, has approximately zero growth near the Nash equilibrium. In the discrete time setting we needed exactly zero growth, but with continuous time, since at no time except at time 0 players are sure that there exists no more revision opportunities, the "growth of approximate zero" (which corresponds to Assumption A4) works as a foothold for long-run cooperation. This is of course not a rigorous proof for why there exists a cooperative path in our model, but this is one of the key parts of the intuition behind our result.

Notice that what is important in the above argument is not the first order condition (A4) per se, but the fact that the gain from defection is smaller than the loss associated with it by a positive order. To make this point clear, consider the following example.

Consider $\pi_i(a_i, a_j) = \sqrt{a_j} - a_i$, with $a_i, a_j \in [0, \infty)$. Note the Nash equilibrium action $a_i = 0$ is a corner solution and the first order condition is not satisfied (the slope is -1). Nevertheless,

$$x(t) = \frac{\lambda^2}{4}t^2$$

constitutes a symmetric trigger strategy path because it satisfies the differential equation (2.6) in Theorem 1 with $x(0) = 0 = a^N$ (therefore $x(t)$ satisfies the binding incentive constraint $d(x(t))e^{-\lambda t} = \int_0^t \left(\pi(x(s)) - \pi^N\right) \lambda e^{-\lambda s} ds$).[29] This example shows that the first order condition at the Nash equilibrium is not necessary for a nontrivial path to be sustained. What is important in this example is the fact that *the gain from deviation, a, is one order smaller than the value of cooperation,* $\sqrt{a} - a$ *(which can be lost after a deviation) near the Nash action* of $a^N = 0$. In what follows we

[29] Note that (unlike in our model in Section 2.4) the differential equation $\frac{dx}{dt} = f(x) \equiv \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$ is well-defined at $x = a^N$, because $d'(a^N)$ does not vanish ($d' \equiv 1$). In particular, $f(a^N) = 0$. In this case, the differential equation $dx/dt = f$ with boundary condition $x(0) = a^N = 0$ has two solutions. One is $x(t) = \frac{\lambda^2}{4}t^2$, and the other is the constant path $x(0) \equiv a^N = 0$.

formulate this observation in a precise way.

First let us generalize the Assumptions A1-A6 imposed in Section 2.4. In particular we consider a general component game with symmetric action space A and payoff function π.

Proposition 10. *Suppose that there is a symmetric isolated Nash equilibrium* (a^N, a^N) *and that there exists* $\epsilon > 0$ *such that* $[a^N, a^N + \epsilon) \subseteq A \subseteq \mathbb{R}$ *for each player i. Suppose also that there exists* $\epsilon' > 0$*,* $r > s > 0$ *and* $k, k' > 0$ *such that for all* $a \in (a^N, a^N + \epsilon')$*,*

$$d(a) \leq k(a - a^N)^r \text{ and} \tag{2.10}$$

$$k'(a - a^N)^s \leq \pi(a) - \pi^N. \tag{2.11}$$

Then, in a synchronous revision game, there exists a subgame perfect equilibrium such that non-Nash profiles are prepared at all time $t > 0$ *on the path of play.*

Proof. Take an $\hat{\epsilon} > 0$ such that $[a^N, a^N + \hat{\epsilon}] \subseteq A$, conditions (2.10) and (2.11) hold for all $[a^N, a^N, a^N + \hat{\epsilon}]$ with constants k, k', r, and s, and $\frac{k'\lambda e^{-\lambda t}(r-s)}{s+r} t^{\frac{s+r}{r-s}} > kt^{\frac{2r}{r-s}}$. Such $\hat{\epsilon} > 0$ exists if the premise of the proposition holds. We are going to show that a trigger strategy path

$$x(t) = \begin{cases} t^{\frac{2}{r-s}} + a^N & \text{if } t < \hat{\epsilon}^{\frac{r-s}{2}} \\ \\ \hat{\epsilon} + a^N & \text{if } t \geq \hat{\epsilon}^{\frac{r-s}{2}} \end{cases} . \tag{2.12}$$

satisfies the incentive constraint

$$\int_0^t \left(\pi(x(\tau)) - \pi^N\right) \lambda e^{-\lambda\tau} d\tau \geq d(x(t)) e^{-\lambda t} \tag{2.13}$$

for all $t \in [0, T]$. To see this, first consider the case $t < \hat{\epsilon}$. We have

$$\begin{aligned}
\int_0^t \left(\pi(x(\tau)) - \pi^N\right) \lambda e^{-\lambda\tau} d\tau &\geq \int_0^t k' \left(x(\tau) - a^N\right)^s \lambda e^{-\lambda\tau} d\tau = \int_0^t k' \tau^{\frac{2s}{r-s}} \lambda e^{-\lambda\tau} d\tau \\
&\geq k' \lambda e^{-\lambda t} \frac{1}{\frac{2s}{r-s} + 1} t^{\frac{2s}{r-s}+1} = \frac{k' \lambda e^{-\lambda t}(r-s)}{s+r} t^{\frac{s+r}{r-s}}. \\
&= k t^{\frac{2r}{r-s}} = k(x(t) - a^N)^r \geq d(x(t)) e^{-\lambda t}.
\end{aligned}$$

Next, consider the case $t \geq \hat{\epsilon}$. We have

$$\begin{aligned}
&\int_0^t \left(\pi(x(\tau)) - \pi^N\right) \lambda e^{-\lambda\tau} d\tau \\
\geq\ & e^{-\lambda\hat{\epsilon}} \left(\pi(x(\hat{\epsilon})) - \pi^N\right) + \int_0^{\hat{\epsilon}} \left(\pi(x(\tau)) - \pi^N\right) \lambda e^{-\lambda\tau} d\tau \\
\geq\ & d(x(\hat{\epsilon})) e^{-\lambda\hat{\epsilon}} = d(x(t)) e^{-\lambda\hat{\epsilon}} \geq d(x(t)) e^{-\lambda t}.
\end{aligned}$$

Hence, the non-trivial path (2.12) satisfies the incentive constraint (2.13) for all $t \in [0, T]$. By definition, on the path of play of the subgame perfect equilibrium characterized by this path, non-Nash profiles are prepared for all $t > 0$. This completes the proof. □

The proposition says that a nontrivial path exists when *the gain from deviation $d(x)$ converges to zero faster than the value of cooperation $\pi(x) - \pi^N$ does, as $x \to a^N$.* If these conditions are met, we can construct a trigger strategy path. Note that those conditions are satisfied in our example (with $d(a) = a$, $a^N = \pi^N = 0$, and $\pi(a) - \pi^N = \sqrt{a} - a$).

A couple of remarks are in order:

- The intuition behind the above proposition can be expressed as follows. As the deadline comes closer and closer, the probability of punishment upon deviation converges to zero. Hence, to maintain the incentive to follow a non-trivial path, the instantaneous gain from deviation need to be infinitesimal relative to the future gain from cooperation, which roughly corresponds to the static loss from reverting to the Nash equilibrium.

- Remember that this condition fails in Bertrand competition without product differentiation and the aforementioned modified good exchange game with $\epsilon = 0$. Thus nonexistence of cooperative path in those examples are consistent with this proposition.

- The above proposition provides a sufficient condition for the existence of a non-trivial path. A necessary and sufficient condition is the finite time condition presented in Section 2.4 (see the discussion following condition (2.8)).

Now let us consider a partial converse of this result. Assume that the payoff for each player i has an upper bound $\bar{\pi}$.

Proposition 11. *Suppose that there exists a unique Nash equilibrium a^N and its payoff $\pi(a^N) = \pi^N$. Suppose that $\inf_{a \in A} \frac{d(a)}{\pi(a) - \pi^N} > 0$. Then, there exists a unique trigger strategy equilibrium. In this equilibrium, each player prepares a^N given any history.*

Proof. Let $\inf_{a \in A} \frac{d(a)}{\pi(a) - \pi^N} =: m > 0$. We will show that there exists $\epsilon > 0$ such that for any $t \in [0, T]$, if for all time strictly after $-t$ each player prepares a^N given any history then for all time in $(-t - \epsilon, -t]$, each player prepares a^N given any history in any subgame perfect equilibrium. This gives us the desired result.

So take some $t \in [0, T]$ and suppose that for all time strictly after $-t$ each player prepares a^N given any history. Suppose further that at time $-t - \epsilon$ with $\epsilon > 0$, an action profile a is played on the path of play. Then, by the incentive compatibility constraint, it is necessary that

$$d(a)e^{-\lambda(t+\epsilon)} \leq e^{-\lambda t}\int_0^{\epsilon}(\bar{\pi}-\pi^N)\lambda e^{-\lambda\tau}d\tau.$$

This implies

$$\begin{aligned} d(a)e^{-\lambda\epsilon} \leq \lambda(\bar{\pi}-\pi^N)\epsilon \quad &\iff \quad d(a) \leq \lambda(\bar{\pi}-\pi^N)\epsilon e^{\lambda\epsilon} \\ &\iff \quad \pi(a)-\pi^N \leq \frac{\lambda(\bar{\pi}-\pi^N)\epsilon e^{\lambda\epsilon}}{m} \end{aligned}$$

Hence, again by the incentive compatibility constraint it is necessary that

$$d(a)e^{-\lambda(t+\epsilon)} \leq e^{-\lambda t}\int_0^{\epsilon}\frac{\lambda(\bar{\pi}-\pi^N)\epsilon e^{\lambda\epsilon}}{m}\lambda e^{-\lambda\tau}d\tau.$$

This in turn implies that $d(a) \leq \frac{\lambda^2(\bar{\pi}-\pi^N)(\epsilon e^{\lambda\epsilon})^2}{m}$. Iterating, we have that

$$d(a) \leq \frac{\lambda^n}{m^{n-1}}(\bar{\pi}-\pi^N)(\epsilon e^{\lambda\epsilon})^n \quad \text{for all } n=1,2,\ldots.$$

Since the right hand side of this inequality goes to zero as n goes to infinity if $\epsilon < \frac{\lambda e^{\lambda\epsilon}}{m}$, $d(a)$ must be zero if $\epsilon < \frac{\lambda e^{\lambda\epsilon}}{m}$. But this means that a must be a Nash equilibrium a^N. Hence in time interval $(-t-\epsilon, -t]$, each player prepares a^N given any history. This completes the proof. □

The proof is based on the idea that the right hand side of the incentive compatibility condition is at most some constant times the time left to the deadline. That is, if the time left to the deadline is very short, the instantaneous gain from deviation must be very small relative to the payoff from cooperation (See the first remark after Proposition 10). If the ratio of the gain from deviation to the benefit of cooperation has a strictly positive lower bound then this is impossible when the remaining time is sufficiently small.

2.6.2 *Comparison with Infinite Repeated Games with Decreasing Discount Factors*

To compare a revision game with a repeated game, let us employ the standard way to measure time: a revision game is played over $[0, T]$ where 0 is the start of the problem and T is the end. The payoff in the revision game at time t is:

$$e^{-\lambda(T-t)}u(a_t) + \int_t^T e^{-\lambda(T-s)}u(a_s)\lambda ds = e^{-\lambda(T-t)}\left[u(a_t) + \int_t^T e^{\lambda(s-t)}u(a_s)\lambda ds\right].$$

Igroring the constant $e^{-\lambda(T-t)}$, we can regard that a player's objective function at time t (i.e., when a revision opportunity arrives at time t) is equal to

$$u(a_t) + \int_t^T e^{\lambda(s-t)}u(a_s)\lambda ds. \tag{2.14}$$

This highlights the similarity and difference between a revision game and repeated game with shrinking discount factor (Bernheim and Dasgupta, 1995). The objective function in their model at time t is given by

$$u(a_t) + \sum_{s=t+1}^{\infty} u(a_s) \prod_{\tau=t+1}^{s} \delta(\tau),$$

where the time dependent discount factor $\delta(\tau)$ shrinks over time ($\delta(\tau) \to 0$, as $\tau \to \infty$). One obvious (but minor) difference is that their model is in discrete time while ours is in continuous time. Our continuous time formulation enables us to characterize the optimal path by means of a simple differential equation. To compare their model with ours more closely, let us consider a continuous time version of their model, where the stage game is played according to Poisson arrival time. A continuous time version of their objective function would be

$$u(a_t) + \int_t^{\infty} e^{-\int_t^{\tau} \rho(\tau) d\tau} u(a_s) \lambda ds. \tag{2.15}$$

where instantaneous discount rate diverges ($\rho(\tau) \to \infty$, as $\tau \to \infty$). This is similar to our model in the sense that as time passes by (when t is large), the impact of future payoffs shrinks. However, note the crucial difference that the weight attached to future payoff $u(a_s)$ in our objective function (2.14), namely $e^{\lambda(s-t)}$, is *increasing* in s. That is, *a larger weight is attached to future payoff in a revision game.* This is an essential feature - as the deadline comes closer, the probability that the prepared action today is implemented becomes larger. One important implication of this fact is that full cooperation cannot be sustained in a revision game. There is always a positive probability that something very close to the Nash equilibrium (an action prepared near the deadline) is played. In contrast, in a repeated game with shrinking discount factor, payoffs in the distant future do not much affect the average payoff, and the full efficiency can be approximately achieved.

The fact that *a large*r weight is attached to future payoff in a revision game implies that there is no natural way to map our objective function to theirs. For example, one may "stretch" the time in our model to map our time domain $[0, T]$ to $[0, \infty)$ by some increasing function $t' = F(t)$, but such a transformation does not alter the property of our model that the weight attached to $u(a_s)$ is increasing in s.

2.7 *Concluding Remarks*

We analyzed a new class of games that we call "revision games," a situation where players in advance prepare their actions in a game. After the initial preparation, they have some opportunities to revise their actions, which arrive stochastically. Prepared actions are assumed to be mutually observable. We showed that players can achieve a certain level of cooperation in such a class of games. Specifically, in the class of component games that we focused on, we showed that an optimal symmetric trigger-strategy equilibrium exists and it is essentially unique. We characterized the equilibrium by a simple differential equation and applied it to analyze a variety of economic examples.

While we are circulating the earlier versions of the present paper, several follow-up papers have been written. Calcagno and Lovo (2010) and Kamada and Sugaya (2010) consider revision games with finite action space and assume that revision opportunities arrive independently across players (asynchronous revision). In contrast to the present paper, they show that the addition of revision phase sometimes narrows down the set of equilibria when the component game has multiple equilibria. They show that when the component game has a strictly Pareto-dominant Nash equilibrium, it is the only profile that can realize in a corresponding revision game when some regulatory conditions are met.[30] They also show that in battle of the sexes games one of the pure Nash equilibrium is played generically. Kamada and Sugaya (2011) introduce the first model of dynamic election campaigns into the literature on election by using a variant of the revision games framework. In their model, the revision phase corresponds to an election cam-

[30] Ishii and Kamada (2011) identify the condition under which this result is generalized to the case of a hybrid version of synchronous and asynchronous revisions. Romm (2011) examines the effect of reputation in a variant of revision games proposed by Kamada and Sugaya (2010).

paign phase where candidates announce their policies, and the component game corresponds to the standard Hotelling-Downs election game.[31] The rich dynamic structure of revision games enables them to endogenize the order of policy announcements, which are exogenously specified in the literature.[32]

We suggest several possible directions for future research. First, we investigate the case of asynchronous revision in a companion paper (Kamada and Kandori, 2011a) and show that cooperation is still possible in such a setting. Second, we used trigger strategy equilibrium to sustain cooperation, in which players revert to Nash actions upon deviation. Although this class of strategies is a natural one worth investigation, a severer punishment might be possible. In our continuation work, we consider severer punishment schemes than Nash reversion.

[31] In their model a policy announcement at each opportunity is restricted by previous announcements in a particular manner, while in our analysis in Section 2.5.4 no restriction is imposed.

[32] Other recent papers on variants of revision games include Ambrus and Burns (2010) and Kamada and Muto (2011b).

3. ASYNCHRONICITY AND COORDINATION IN COMMON AND OPPOSING INTEREST GAMES [2]

3.1 Introduction

It has been broadly argued that the addition of a pre-play phase to a game has a nontrivial effect on the outcome of the game. Cheap talk models à la Farrell (1987), Rabin (1994) and Aumann and Hart (2003) describe such a situation, where any action in the pre-play phase has no binding force. They show that the set of achievable outcomes *widens* with the addition of a pre-play phase. This paper analyzes an opposite situation, where any action in the pre-play phase has effects on the outcome of the game with strictly positive probability. During the pre-play phase, players prepare the actions that will be played at a predetermined deadline. The

---

[2] This is a joint work with Riccardo Calcagno (Department of Economics, Finance and Control, EMLYON Business School), Stefano Lovo (Finance Department, HEC, Paris and GREGHEC), and Takuo Sugaya (Department of Economics, Princeton University), and is the result of a merger between two independent projects: Calcagno and Lovo's "Preopening and Equilibrium Selection" and Kamada and Sugaya's " Asynchronous Revision Games with Deadline: Unique Equilibrium in Coordination Games." We thank Dilip Abreu, Gabriel Caroll, Sylvain Chassang, Drew Fudenberg, Michihiro Kandori, Fuhito Kojima, Barton Lipman, Thomas Mariotti, Sebastien Pouget, Stephen Morris, Assaf Romm, Satoru Takahashi, Tristan Tomala, Nicolas Vieille, and, particularly, Johannes Hörner and Yuhta Ishii for useful comments and suggestions on either project. We also thank seminar participant at the GDR Conference in Luminy 2009, SAET Conference in Ischia 2009, Toulouse School of Economics, Bocconi University, Research in Behavior in Games seminar at Harvard University, Student Lunch Seminar in Economic Theory at Princeton University, and at The 21st Stony Brook International Conference on Game Theory for helpful comments. We are grateful to three anonymous referees and the co-editor of *Theoretical Economics* for insightful comments and suggestions that substantially improved this paper. Stefano Lovo gratefully acknowledges financial support from the HEC Foundation and from the ANR Grant ANR-10-BLAN 0112.

action taken at the deadline solely determines players' payoffs. Prepared actions can be revised at stochastic (Poisson) opportunities in continuous time before the deadline. In this framework, Kamada and Kandori (2011b) show that the addition of pre-play phase can *widen* the set of achievable payoffs. This paper uncovers another role that the pre-play phase can play. We show that it can *narrow down* the set of achievable payoffs. The key assumptions that drive this difference will be discussed shortly.

We study this problem in two classes of games where coordination is an issue. The first is "common interest" games, in which there is an action profile that all players strictly prefer to all other profiles. For this class of games, we show that, in 2-player games, this "best profile" is the unique outcome of the revision game. The second class of games is the "opposing interest" games, which are two-player 2 x 2 games with two Pareto-unranked strict Nash equilibria. In this class of games, we show that generically there is a unique outcome of the revision game, which corresponds to one of the strict Nash equilibria. Which equilibrium prevails in the revision game depends on the payoff structure and the relative frequency of arrivals of revision opportunities for each player.

Besides the importance of assuming that revision opportunities are stochastic, there are other three key assumptions that are crucial to our results. The first one is *observability*. If a player is unable to observe what the other player has prepared, then the revision phase has no binding force, and so the outcome of revision games would be identical to that of static games. The second is *asynchronicity*. If revision opportunities are synchronous all the time, then any repetition of static Nash equilibria is subgame perfect. Hence, uniqueness would not hold if there are multiple static Nash equilibria. However, if opportunities are asynchronous, each player's action must be contingent on the opponent's current action (by observability). Thus a player can induce the opponent to play some particular action, by using as a threat the possibility that she may not be able to revise her own action before the deadline. The third key ingredient is *finite actions and strict incentives*. As we will argue, uniqueness is due to a backward induction argument. If there are only finitely many actions and the static game best replies to pure actions are strict, then each player has a single best reply near the deadline (by asynchronicity) in the revision game, and this constitutes the starting point of our backward induction argument.

These assumptions seem natural in many real-life and economic contexts where coordination is crucial. For example, such a situation arises in the daily practice of some financial markets, such as Nasdaq or Euronext for example, where half an hour before the opening of the market, participants are allowed to submit orders which can be withdrawn and changed until opening time. These orders and the resulting (virtual) equilibrium trading price are publicly posted during the whole "pre-opening" period. Only orders that are still posted at the opening time are binding and hence executed. In this framework, it is natural to assume asynchronicity and that traders do not always manage to withdraw old orders or submit new orders instantaneously because it takes a certain random time to fill in the new order faultlessly. Observability holds as the posted orders are displayed on the screen, and the number of payoff-relevant orders is practically finite.[3,4]

Another example is the case where two firms are contemplating the possibility

---

[3] Given this application, Calcagno and Lovo (2010) call the revision game a "preopening game".

[4] Biais, Cristophe, and Sebastien (2008) present an experiment simulating preopening in a financial market where the actual play is preceded by (only) one round of pre-play communication, which is either completely binding or completely non-binding. In both specifications players choose their actions simultaneously and there are multiple SPE equilibria. Consistently, they find both Pareto superior and Pareto inferior outcomes are observed in the experiment.

of investment at some fixed date and need to prepare for it (e.g. negotiating with banks to have enough liquidity, allocating agents working for the project and so forth). There are two actions (invest or not) and the investment is profitable only if the other firm invests as well. Revision opportunities are naturally asynchronous and firms may face several constraints such as administrative procedures or obligations to other projects. Since the firms cannot change their decisions every time they want, their opportunities would well be stochastic. Finally, if two firms are large, then it is natural to assume that they observe each other's preparation.

The rest of the paper is organized as follows. Section 3.2 reviews the related literature. Section 3.3 introduces the model. In Section 3.4, we present a simple but a useful lemma that allows us to implement a backward induction argument in continuous time. Section 3.5 considers 2-player common interest games and Section 3.6 studies 2-player opposing interest games. Section 3.7 discusses further results. Section 3.8 concludes. Some of the proofs are relegated to the Appendix.

3.2 Literature Review

Cheap talk. It is important to make a distinction between our model and cheap-talk models such as those in Farrell (1987), Rabin (1994) and Aumann and Hart (2003). In these models, players are allowed to be involved in preplay non-binding communication. Quite to the opposite, in our model, at each moment of time, the prepared action will become the final payoff-relevant action with a strictly positive probability. For this precise reason, the outcome can be affected by the addition of a revision phase in our model.

Equilibrium selection. It is instructive to compare our selected outcome with those in that literature. In many works on equilibrium selection, risk-dominant equilibria of John and Selten (1988) are selected in 2×2 games. In our model, however, a different answer is obtained: a strictly Pareto-dominant Nash equilibrium is played even when it is risk-dominated. Roughly speaking, since we assume perfect and complete information with non-anonymous players, there is only a very small "risk"of mis-coordination when the deadline is far. There are three lines of the literature in which risk-dominant equilibria are selected: models of global games, stochastic learning models with myopia and models of perfect foresight dynamics.[5,6] Since the model of perfect foresight dynamics seems closely related to ours, let us discuss it here.

Perfect foresight dynamics and repeated games. Perfect foresight dynamics, proposed by Matsui and Matsuyama (1995), are evolutionary models in which players are assumed to be patient and "foresighted" that is, they value the future payoffs and take best responses given (correct) beliefs about the future path of play.[7] There is a continuum of agents who are randomly and anonymously matched over infinite horizon according to a Poisson process. In this setup, they select the risk-dominant action profile in 2 x 2 games with two Pareto-ranked (static)

[5] The literature on global games was pioneered by Rubinstein (1989), and analyzed extensively in Carlsson and van Damme (1993), Morris and Shin (1998), and Sugaya and Takahashi (2011). They show that the lack of almost common knowledge due to incomplete information can select an equilibrium. The type of incomplete information they assume is absent in our model. Stochastic learning models with myopia are analyzed in Kandori, Mailath, and Rob (1993) and Young (1993). They consider a situation in which players interact repeatedly, and each player's action at each period is stochastically perturbed. The key difference between their assumptions and ours is that in their model players are myopic, while we assume that players take actions anticipating the opponents' future moves. In addition, the game is repeated infinitely in their models, while the game is played once and for all in our model.

[6] As an exception, Young (1998) shows that in the context of contracting, his evolutionary model does not necessarily lead to risk-dominant equilibrium (p-dominant equilibrium in Morris, Rob, and Shin (1995)). But he considers a large anonymous population of players and repeated interaction, so the context he focuses on is very different from the one of this paper.

[7] See also Oyama, Takahashi, and Hofbauer (2008).

Nash equilibria. The key difference is that they assume anonymous agents while we assume non-anonymous agents. For the "best action profile" to be selected in our model, it is important for each player to expect that if she prepares an action corresponding to the best profile, then that preparation can affect the other player's future preparation. This consideration is absent with anonymous players.

Common interest games and asynchronous moves. Farrell and Saloner (1985) and Lagunoff and Matsui (1997) are early works on the topic of obtaining the unique outcome in common interest games.[8] Dutta (2003) shows convergence to the unique outcome and Takahashi (2005) proves uniqueness of subgame perfect equilibria when players move asynchronously. One difference is that we assume the stochastic order of moves while they consider the fixed order. Also, we obtain a uniqueness result in a wider environment than in Lagunoff and Matsui (1997) due to the finite horizon.

War of attrition. The intuition behind the result for the opposing interest games is similar to the one for the "war of attrition". [9] Although the structure of the equilibria in war of attrition is similar to the equilibria in our model, the reasoning is different: in our model, players use the probability of not having future revision opportunities as a "commitment power" while the literature in the war of attrition assumes the existence of "commitment types" a priori.

Switching cost. Our model assumes there is no cost associated with revision.

---

[8] According to Dutta (1995), this result in Lagunoff and Matsui (1997) is due to the lack of full dimensionality of the feasible and individually rational payoff set. See also Lagunoff and Matsui (2001), Yoon (2001), and Wen (2002). Rubinstein and Wolinsky (1995) show that, even when the discount factor is arbitrarily close to one, the set of SPE payoff vectors of the repeated games resulting from the repetition of the extensive form game may not coincide with the one resulting from the normal form game, if the individually rational payoffs are different or full dimensionality is not satisfied.

[9] For example, among others, see Abreu and Gul (2000) and Abreu and Pearce (2007).

Several papers consider a finite-horizon model with switching cost and show that a unique outcome prevails in their respective games. Typically, the existence of switching cost results in different implications on the equilibrium behavior. See, for example, Lipman and Wang (2000) and Caruana and Einav (2008) for details.

Revision games. Kamada and Kandori (2011b) introduce the model of revision games. They show that, among other things, non-Nash "cooperative" action profiles can be played at the deadline when a certain set of regularity conditions is satisfied. Hence their focus is on expanding the set of equilibria when the static Nash equilibrium is inefficient relative to non-Nash profiles.[10] We ask a very different question in this paper: we consider games with multiple efficient static Nash equilibria and ask which of these equilibria is selected.[11] What derives this difference is that the action space is finite in our paper, whereas it is not in Kamada and Kandori (2011b). Kamada and Sugaya (2011) consider a revision game model with finite action set in the context of election campaign. The main difference is that they assume once a player changes her action, she cannot revise it further. Thus the characterization of the equilibrium is essentially different from the analysis in the present paper because in our model, when another opportunity arrives, a player can always change her preparation to the previously-prepared action.[12]

Further results. Finally, further results beyond what we have in this paper can be found in either or both of Calcagno and Lovo (2010) and Kamada and Sugaya (2010). We refer to these papers whenever appropriate.

---

[10] The possibility of cooperation in finite horizon in Kamada and Kandori (2011b) is closely related to that of finitely repeated games with multiple Nash equilibria (Benoit and Krishna, 1985).

[11] See also Ambrus and Lu (2010a) for a variant of revision games model of bargaining in which the game ends when an offer is accepted.

[12] van Damme and Hurkens (1996) analyze a related model of "timing games," in which players can choose the timing of their move out of two periods and they cannot switch back.

3.3 The Model

We consider a two-player normal-form finite game $\left((X_i)_{i=1,2}, (u_i)_{i=1,2}\right)$ (referred to as the "*component game*" in the following) where X_i is the finite set of player i's actions with $|X_i| \geq 2$, $X = X_1 \times X_2$ is the set of action profiles, and $u_i : X \to \mathbb{R}$ is player i's utility function. Let $u = (u_1, u_2)$. We use a female (resp. male) pronoun for player 1 (resp. player 2).

Before players actually take actions, they need to "prepare" their actions. We model this situation similarly to Kamada and Kandori (2011b): time is continuous, $t \in [-T, 0]$, and the component game is played once and for all at time 0. The game, "*revision game*" henceforth, proceeds as follows. First, at time $-T$, the initial action profile is exogenously given.[13] In the time interval $(-T, 0]$, each player independently obtains opportunities to revise their prepared action according to two random Poisson processes $\mathbf{p}_1$ and $\mathbf{p}_2$ with arrival rates λ_1 and λ_2 respectively, where $\lambda_i > 0$, $i = 1, 2$. As Poisson processes $\mathbf{p}_1$ and $\mathbf{p}_2$ are independent, the probability that the two players revise their actions simultaneously is nil. In other words, only asynchronous revision opportunities arise.[14] At $t = 0$, the action profile that has been prepared most recently is actually taken and each player receives the payoff that corresponds to the payoff specification of the component game.

In order to define the strategy space of the revision game, consider the game has reached time t. We assume here that each player i at any time t has perfect information about all past events including whether i has a revision opportunity

[13] As we will see, the uniqueness results in Sections 5 and 6 become even sharper if players simultaneously choose actions at time $-T$.

[14] We refer to Section 3.7 for the discussion of the role played by this assumption. See Calcagno and Lovo (2010) and Ishii and Kamada (2011) for more general processes underlying the arrival of revision opportunities.

at t but excluding whether the opponent gets an opportunity at t.[15] Formally, let $t_{i,k} \leq t$ be the time when player i has received the k-th revision opportunity until t and $x_{i,k}$ be the action prepared by player i at $t_{i,k}$.[16] A *history* for player i at t is

$$h_i(t) = X \times \left\{ \left(\{t_{i,k}\}_{k,\, t_{i,k} \leq t} \ , \ \{x_{i,k}\}_{k,\, t_{i,k} < t} \right), \left(\{t_{j,k}\}_{k,\, t_{j,k} < t} \ , \ \{x_{j,k}\}_{k,\, t_{j,k} < t} \right), \ t \right\},$$

where $j \neq i$, and an element in X denotes the exogenous choice of the action profile at time $-T$. Let $H_i(t)$ denote the set of all possible histories for player i at t.[17] A strategy for player i is a mapping $\sigma_i : \cup_{t=-T}^{0} H_i(t) \rightarrow \{\emptyset\} \cup \Delta(X_i)$ where $\sigma_i(h_i(t)) \subseteq \Delta(X_i)$ if there exists k such that $t_{i,k} = t$ (i.e. at t, player i receives a revision opportunity) and $\sigma_i(h_i(t)) = \emptyset$ otherwise (i.e. at t, player i does not receive a revision opportunity). For any given history $h_i(t)$, let $x_i(t) := x_{i,k^*,} \in X_i$ with $k^* := \arg\max_k \{t_{i,k} < t\}$ be player i's prepared action resulting from his last revision opportunity (strictly) before t. We shall denote $x(t) := \{x_i(t)\}_{i=1,2}$ the *last prepared action profile* before time t (time t "PAP" henceforth). Note that $x_i(t)$ will be player i's payoff-relevant action in $t = 0$ in the event where i receives no further revision opportunities from time t included, until time 0.

A strategy profile σ^* forms a subgame perfect equilibrium (SPE) of the revision

[15] This assumption is expressed in the definition of history $h_i(t)$, where we use a strict inequality for a profile of $t_{j,k}$'s. Since simultaneous revision opportunities occur with zero probability, our result does not depend on this assumption.

[16] Notice that $-T = t_{i,1} < \cdots < t_{i,k} \leq t$, that is, we count the revision opportunities from the first one $k = 1$ after the beginning of the revision game.

[17] Note that $H_i(-T)$ is defined to be X.

game if for all t, $h_i(t)$, i and $j \neq i$,[18]

$$\sigma_i^* \in \arg\max_{\sigma_i} E\left[u_i(x(0))|h_i(t), \sigma_i, \sigma_j^*\right].$$

Our main results will concern the case when T is large. However, we note that the model with arrival rate (λ_1, λ_2) and horizon length T is essentially equivalent to the model with arrival rates $(a\lambda_1, a\lambda_2)$ and horizon length $\frac{T}{a}$, for any positive constant a.[19] Hence all our results obtained for T large enough and fixed revision frequencies (λ_1, λ_2) can be obtained by keeping fixed the horizon T and having revisions frequent enough.

To avoid ambiguity, in the rest of the paper, we will use terminology "*revision equilibrium*" for a SPE of the whole revision game and "*Nash equilibrium*" for a (strict) Nash equilibrium of the component game.

3.4 Backward Induction in Continuous Time

The proofs of our main results will rely on the idea of backward induction. The standard backward induction argument starts from proving a statement for "the last period" and then given the statement is true there, it proves the statement for the "second-last period" and so forth. However, this argument is not immediately applicable to our continuous-time setting, as there is no obvious definition of "second-last period". In this section we present a lemma that allows us to implement a backward-induction-type argument in continuous time. The proof is in the

[18] Strictly speaking, $x(0)$ is the last action profile prepared before time 0, thus in this formulation players do not maximize the expected payoff prepared exactly at 0. However, since the probability that any player obtains a revision opportunity exactly at time 0 is nil, this issue does not affect the solution of the maximization problem.

[19] See the "arrival rate invariance" result discussed in Kamada and Kandori (2011b).

Appendix.

Lemma 1. *Suppose that for any t, there exists $\epsilon > 0$ such that statement $A_{t'}$ is true for all $t' \in (t-\epsilon, t]$ if statement $A_{t''}$ is true for any $t'' > t$.*[20] *Then, for any t, statement A_t is true.*

It is noteworthy that the ϵ in the statement of the lemma can depend on t. Hence, in particular, the lemma goes through even though the required ϵ shrinks to zero as t approaches some finite constant, and then jumps discontinuously there.[21]

3.5 Common Interest Games

In this section, we consider a component game with an action profile that strictly Pareto-dominates all other action profiles. Formally, we say that an action profile x^* is **strictly Pareto-dominant** if $u_i(x^*) > u_i(x)$ for all i and all $x \in X$ with $x \neq x^*$. We say that a game is a **common interest game** if it has a strictly Pareto-dominant action profile. Notice that if x^* is strictly Pareto-dominant, then it is a Nash equilibrium.

For example, games in Table 3.1 are common interest games where (U, L) in each game is strictly Pareto-dominant, but those in Table 3.2 are not.

The first main result of this paper is the following:

[20] Note that if $t = 0$ then it is vacuously true that statement $A_{t''}$ is true for any $t'' > t$. Thus the premise of the lemma holds if and only if there exists $\epsilon > 0$ such that $A_{t'}$ is true for all $t' \in (-\epsilon, 0]$.

[21] A version of the lemma that switches the order of quantifiers (so that ϵ cannot depend on t) appears in Chao (1919).

Table 3.1: Common interest games.

	L	R
U	$2,2$	$-10,1$
D	$1,-10$	$1,1$

	L	C	R
U	$2,2$	$1,0$	$-4,3$
M	$1,1.5$	$1,1$	$-3,1$
D	$1,0$	$0,1$	$0,0$

Table 3.2: Non-common interest games.

	L	R
U	$2,2$	$3,0$
D	$0,3$	$1,1$

	L	R
U	$2,2$	$0,0$
D	$2,0$	$1,1$

Theorem 2. *Consider a common interest component game and let x^*be the strictly Pareto-dominant action profile. Then for any $\epsilon > 0$, there exists T' such that for all $T > T'$, in all revision equilibria, $x(0) = x^*$ with probability higher than $1 - \epsilon$.*

3.5.1 *Intuition*

The proof consists of two steps. First, we show that x^* is absorbing in the revision game: since the action space is finite, the difference between $u_i(x^*)$ and i's second best payoff is strictly positive. Therefore, when the PAP is x^*, no player wants to prepare another action and to create a possibility that she cannot have further revision opportunities and will be forced to take a second best or even worse action profile.

Second, given the first step, each player i "knows" that if the opponent $-i$ has a revision opportunity while player i prepares x_i^*, then the opponent will prepare x_{-i}^*. Hence, the lower bound of the equilibrium payoff for each player is given by always preparing x_i^* whenver she receives the revision opportunity. If T is sufficiently large, then this strategy gives her a payoff very close to $u_i(x^*)$, which means x^* should be taken with high probability at the deadline in any revision equilibrium.

3.5.2 *Proof of Theorem 1*

Now we offer the formal proof. Steps 1 and 2 in the formal proof correspond to those in the intuitive explanation above.

Step 1:

Let $m := \min_{i,x \neq x^*}(u_i(x^*) - u_i(x))$ be the minimum payoff difference between the best payoff and the second best payoff. Since X is finite and x^* is strictly Pareto-dominant, the minimum is well defined and $m > 0$.

Fix $t \leq 0$ arbitrarily. Suppose that for all time after time $t < 0$, each player i has a strict incentive to prepare x_i^* if the opponent $-i$ prepared action is x_{-i}^*.[22] Suppose also that player i obtains a revision opportunity at time $t - \epsilon$ and $-i$ prepared action is x_{-i}^*. Then, the payoff from preparing x_i^* is at least

$$u_i(x^*) - (1 - e^{-(\lambda_i + \lambda_{-i})\epsilon})M \qquad (3.1)$$

where $M := \max_{i,x \neq x^*}(u_i(x^*) - u_i(x)) < \infty$, because with probability at least $e^{-(\lambda_i+\lambda_{-i})\epsilon}$, no further revision opportunities arrive between $t - \epsilon$ and t and the PAP at time t is x^*. In such a case, action x^* will be played at the deadline by assumption. On the other hand, the payoff from preparing an action $x_i \neq x_i^*$ is at

[22] For $t = 0$, this is vacuously true.

most

$$u_i(x^*) - e^{-\lambda_i(-t+\epsilon)}m, \tag{3.2}$$

because with probability $e^{-\lambda_i(-t+\epsilon)}$, player i never has a revision opportunity again and in such a case, the action profile at the deadline cannot be x^*.

Notice that expression (3.1) is strictly greater than expression (3.2) for $\epsilon = 0$. Also by continuity of (3.1) and (3.2) with respect to ϵ, there exists $\epsilon' > 0$ such that for all $\epsilon \in (0, \epsilon')$, expression (3.1) is strictly greater than expression (3.2).[23] Hence, by Lemma 1, we have that for any $t < 0$, each player i has a strict incentive to prepare x_i^* if the opponent $-i$ prepares x_{-i}^*.

Step 2: Since in any subgame perfect equilibrium, players can guarantee at least the payoff that can be obtained by always playing the action x_i^*, it suffices to show that the payoff of always preparing x_i^* converges to the strictly Pareto-dominant payoff as T goes to infinity. This will imply that the probability of action being x^* at the deadline converges to 1, as desired.

By Step 1, the action profile x^* is the absorbing state: each player has a strict incentive to prepare x_i^* if the opponent $-i$ prepares x_{-i}^*. In 2-player games, since player i is the unique opponent of player $-i$, player $-i$ prepares x_{-i}^* if player i prepares x_i^*. Therefore, the payoff of always preparing x_i^* guarantees a payoff which converges to $u_i(x^*)$. □

3.5.3 *Remarks*

Four remarks are in order at this stage.

First, if players choose their actions at $-T$, then we can pin down the behavior

[23] Note that here we again use the assumption that the action space is finite, so that the maximum payoff difference is bounded.

of players on the equilibrium path. In fact, in the appendix, we show that in a common interest game defined as above, players prepare the strictly Pareto-dominant profile x^* at all times $t \in [-T, 0]$ on the (unique) path of play in any revision equilibrium.

Second, notice that if there exist two strict Pareto-ranked Nash equilibria in a 2×2 component game, then the game is a common interest game. Hence in such a case, the Pareto-superior Nash equilibrium is the outcome of the revision game.[24]

Third, the outcome of the revision game is the strictly Pareto-dominant profile even if it is risk-dominated by another Nash equilibrium. For example, in the left payoff matrix in Table 3.1, the action profile (U, L) is risk-dominated while it is the outcome of the revision game. The key is that, whenever a player prepares x_i^* (the action that corresponds to the Pareto-dominant profile), the opponent will move to the Pareto-dominant profile whenever she can revise and they stay at this profile until the deadline (Step 1 of the proof in the previous subsection). Therefore, if the remaining time is sufficiently long, then the "risk of mis-coordination" by preparing x_i^* can be arbitrarily small (Step 2).[25]

Fourth, notice that we allow for the component game to be different from a pure coordination game (i.e. a game in which two players have identical payoff functions). This result is in a stark difference from Lagunoff and Matsui (1997), whose result only applies to pure coordination games (see Yoon, 2001). This difference comes from the different assumptions on the horizon: since their models have an

---

[24] Note that Kamada and Kandori (2011b) prove that if each player has a strictly dominant action when the action space is finite, then it is played in asynchronous revision games.

[25] With more than two players, if all the players are preparing actions different from the Pareto dominant one, no player can create a situation where it is enough for only one player to change her preparation in order to go to the Pareto dominant action profile. Hence, the same proof does not work. See Kamada and Sugaya (2010) for the details.

Table 3.3: Opposing interest games.

	L	R
U	$3,3$	$0,1$
D	$0,5$	$1,6$

	L	C
U	$2,1$	$0,0$
M	$0,0$	$1,2$

Table 3.4: The general opposing interest game.

	L	R
U	$u_1(U,L), u_2(U,L)$	$u_1(U,R), u_2(U,R)$
D	$u_1(U,D), u_2(U,D)$	$u_1(D,R), u_2(D,R)$

infinite horizon, there can be an infinite sequence of punishments. On the other hand, in our model, there is a deadline so the incentives near the deadline can be perfectly pinned down as x^* is strictly Pareto-dominant. Hence, we can implement backward induction starting from the deadline.

3.6 Opposing Interest Games

In the previous section, we analyzed games in which there is the "best" action profile for both players. Now we turn to the class of games in which different players have different "best" action profiles. Examples of games that we consider in this section are given in Table 3.3.

Generally, we consider 2-player component games with the payoff matrix as in Table 3.4 with two strict Nash equilibria (U, L) and (D, R) such that

$$u_1(U,L) > u_1(D,R) \quad \text{and} \quad u_2(U,L) < u_2(D,R). \tag{3.3}$$

The first inequality implies that player 1 strictly prefers (U, L) to (D, R) among pure Nash equilibria while the second implies that player 2's preference is opposite. Note that, since (U, L) and (D, R) are strict Nash equilibria of this component game, these conditions imply (U, L) (resp. (D, R)) gives player 1 (resp. player 2) a strictly better payoff than any other action profile.

Let

$$t_1^* = -\frac{1}{\lambda_1 + \lambda_2} \ln \left(\frac{\lambda_1}{\lambda_2} \frac{u_1(D,R) - u_1(U,R)}{u_1(U,L) - u_1(D,R)} + \frac{u_1(U,L) - u_1(U,R)}{u_1(U,L) - u_1(D,R)} \right), \tag{3.4}$$

and

$$t_2^* = -\frac{1}{\lambda_1 + \lambda_2} \ln \left(\frac{\lambda_2}{\lambda_1} \frac{u_2(U,L) - u_2(U,R)}{u_2(D,R) - u_2(U,L)} + \frac{u_2(D,R) - u_2(U,R)}{u_2(D,R) - u_2(U,L)} \right). \tag{3.5}$$

Theorem 3. *Suppose that a component game of a revision game satisfies condition (3.3). If $t_1^* \neq t_2^*$, then there exists a unique revision equilibrium for all T. As T converges to infinity,*

1. *if $t_1^* > t_2^*$, then the equilibrium payoffs converge to $u_i(U, L)$.*

2. *if $t_1^* < t_2^*$, then the equilibrium payoffs converge to $u_i(D, R)$.*

Notice that $t_1^* = t_2^*$ happens only in a knife-edge set of parameter. In this non-generic case, the revision game has multiple equilibria.[26]

Theorem 2 proves that for almost all parameter values, there is a unique revision equilibrium payoff and the outcome at the deadline will form one of the underlying game Nash equilibria with probability that converges to 1 as T increases.

[26] See Kamada and Sugaya (2010) for a characterization of the set of revision equilibrium payoffs for the case $t_1^* = t_2^*$.

Which Nash equilibrium is prepared depends on a joint condition on the payoff function (u) and the ratio of arrival rates ($\frac{\lambda_1}{\lambda_2}$), as t_1^* and t_2^* depend on these parameters. In Table 3.3, if $\lambda_1 = \lambda_2$, then $t_1^* > t_2^*$ in the left game and $t_1^* = t_2^*$ in the right game. Hence if $\lambda_1 = \lambda_2$, then (U, L) is the (limit) outcome in the left game, while the theorem does not cover the case in the right game. However, if $\lambda_1 < \lambda_2$ (resp. $\lambda_1 > \lambda_2$) then the theorem implies that in the right game, the (limit) outcome is (U, L) (resp. (D, R)).

In the proof, we actually completely pin down the behavior at any time t in this unique revision equilibrium. In particular, players prepare the action corresponding to the limit payoff profile for sufficiently long time on the path of play. This implies that if they were to choose actions simultaneously at $-T$, then they choose these actions and they never revise them on the path of play.

In Subsection 3.6.1, we provide an interpretation of this result. Subsection 3.6.2 provides the proof, and Subsection 3.6.3 fully describes the equilibrium dynamics, including off-path plays.

3.6.1 *Interpretation of Theorem 3*

The first step of the proof of Theorem 2 shows that when t is close to zero, each player strictly prefers to prepare the component game best response to the last prepared action of his opponent. Hence, in the game of Figure 3, when getting closer to time zero, players will move away from PAP (U, R), to reach either (U, L) or (D, R) and then stay there until the deadline.[27] If t becomes increasingly far from 0, player i's expected continuation payoff from PAP (U, R) gets closer to a convex combination of $u_i(U, L)$ and $u_i(D, R)$ since the probability that no players revise

[27] Note that the incentive is strict at the deadline $t = 0$.

their actions between t and 0 gets smaller. Hence, there is a finite time t^* such that, when the PAP is (U, R), one player, whom we will call the *strong player*, becomes indifferent at time t^* between (a) preparing the component game best response to his/her opponent's prepared action and (b) preparing the action necessary to form his/her preferred component game Nash equilibrium. Strictly before t^*, the strong player strictly prefers choice (b) in all PAPs. As the proof in the next subsection clarifies, $t^* = \min\{t_1^*, t_2^*\}$ is the time such that the strong player is indifferent between these two actions. The condition $t_1^* > t_2^*$ implies the player 1 is the strong player. In other words, player 1 can stick to non-Nash profiles longer than player 2 to induce player 2 to coordinate on her own preferred Nash equilibrium. The condition $t_1^* > t_2^*$ thus means that the strength with which player 1 can stick to a non-Nash profile is greater than that of player 2.

To see how this "strength" is affected by the parameters of the model, we consider two special cases. First, suppose that $\lambda_1 = \lambda_2$. In this case, $t_1^* > t_2^*$ is equivalent to

$$\frac{u_2(D,R) - u_2(U,R)}{u_2(D,R) - u_2(U,L)} > \frac{u_1(U,L) - u_1(U,R)}{u_1(U,L) - u_1(D,R)}.$$

The formula compares how strongly each player likes (U, R) relative to the other two Nash equilibria. If player 1 likes it more, then she suffers less from miscoordination at (U, R), so as a consequence it is more likely that the inequality is satisfied. If player 1 likes (D, R) less, then she expects less from moving away from (U, R) to (D, R), and so the inequality is more likely to be satisfied if we decrease $u_1(D, R)$.[28]

Second, consider the case with symmetric payoff functions: $u_1(U, L) = u_2(D, R)$, $u_1(D, R) = u_2(U, L)$, and $u_1(U, R) = u_2(U, R)$. In this case, $t_1^* > t_2^*$ is

[28] Note that a risk-dominated Nash equilibrium in the component game may be the (limit) outcome

equivalent to $\lambda_1 < \lambda_2$. This means that $\lambda_1 < \lambda_2$ (resp. $\lambda_1 > \lambda_2$) implies that (U, L) (resp. (D, R)) is the outcome of the revision game. More generally, $\frac{|t_1^*|}{|t_2^*|}$ is increasing in $\frac{\lambda_1}{\lambda_2}$: if player 1's relative frequency of the arrival of revision opportunities compared to player 2's frequency decreases, then player 1's commitment power becomes stronger, so (U, L) is more likely to be selected.

These results are reminiscent of the findings in the bargaining literature. Player i's bargaining power increases in the disagreement payoff $u_i(U, R)$, decreases with the steepness of preference over the two "agreement outcomes" ($|u_1(U, L) - u_1(D, R)|$ for player 1 and $|u_2(D, R) - u_2(U, L)|$ for player 2) and increases in the ability to commit $1/\lambda_i$ on a proposal.

3.6.2 *Proof of Theorem 3*

In this subsection, we provide the proof of the convergence of the equilibrium payoff in Theorem 3.[29] The proof consists of the following three steps.

Step 1:

First, for each player i, we define t_i^* as the infimum of time t such that given that both players prepare the component game best responses against the opponent's action at any $t' > t$, i strictly prefers to prepare a component game best response to any other action. Since the incentive to take a static best response in the component

---

of the revision equilibrium: Consider the payoff matrix

	L	R
U	$2+\epsilon, 1$	$0, 0$
D	$2\epsilon, 0$	$1, 2$

with $\epsilon > 0$. (U, L) is risk-dominated by (D, R), while it is the (limit) outcome of the revision equilibrium when $\lambda_1 = \lambda_2$.

[29] The intuition behind this proof idea is analogous to the one provided in Kamada and Sugaya (2010)'s "three-state example." We thank an anonymous referee for suggesting the way to extend it.

game is strict at the deadline, this is true for t close enough to 0. By this definition and continuity of expected payoffs (with respect to probabilities and so to time), player i must be indifferent between the two actions at t_i^* given that (i) the PAP at t_i^* is (U, R) and that (ii) both players prepare best responses against the opponent's action at time $t > t_i^*$. Then, from a straightforward calculation contained in the Appendix, we show that for each $i = 1, 2$, t_i^* defined in this way coincides with t_i^* defined in (3.4) and (3.5).

Step 2:

Suppose w.l.o.g. that $t_1^* > t_2^*$ and fix $t \in (-\infty, 0]$. Suppose that the following statements are true for any $t' > t$:

1. $t_1^* \le t'$ or player 1 strictly prefers preparing U at t' whatever the opponent's current prepared action is;

2. $t_1^* \le t'$ or player 2 strictly prefers preparing L at t' when player 1's current prepared action is U.

These two statements are trivially true for t' close enough to 0. We will show that there exists $\epsilon > 0$ such that these two statements are true also for all $t' \in (t - \epsilon, t]$, which proves that the statements are true for any t, by Lemma 1.

Step 2-1: First, consider player 1's incentive when she obtains an opportunity at time $t < t_1^*$ (In the other case (i.e. $t \ge t_1^*$), the conclusion trivially holds). Suppose first that player 2 is currently preparing L, or he has a chance to revise strictly after time t but strictly before time t_1^*. If player 1 prepares action U, then statements (1) and (2) and Step 1 imply that the action profile at the deadline is (U, L), which gives player 1 the largest possible payoff that she can obtain in this revision game. On the other hand, if she prepares D, then there is a positive probability that she will obtain no other chances to revise. In such a case, the action profile at the deadline is not (U, L). Hence, player 1 receives a payoff strictly less than the best possible payoff $u_1(U, L)$.

Suppose next that the current action of player 2 is R, and he will not have any chance to revise strictly after time t but strictly before time t_1^*. In this case player 1's expected payoff is the same as the continuation payoff when player 2's prepared action is R at time t_1^*.[30] Hence, player 1 must be indifferent between U and D at t_1^* by Step 1.

Overall, player 1 is strictly better off by preparing U at time t. Hence statement (1) is true at time t.

Step 2-2: Now consider player 2's incentive when he obtains an opportunity at time $t < t_1^*$ (Again, the other case is trivial). Suppose that player 1's current action is U (Note that statement (2) concerns only such a case). If player 2 prepares L, then statements (1) and (2) and Step 1 imply that both players never change their actions in the future. Hence, the action profile at the deadline is (U, L), which leads to the payoff of $u_2(U, L)$. On the other hand, suppose that he prepares R. Player 2 prepares L if he obtains a revision opportunity strictly after time t but strictly before time t_1^*, which results in the payoff of $u_2(U, L)$. If he does not obtain any revision chance within that interval, then the expected payoff is the same as the continuation payoff given action profile (U, R) at time t_1^*. The latter is strictly less than $u_2(U, L)$, since, by the assumption that $t_2^* < t_1^*$, player 2 has a strict incentive to prepare L given that player 1 is preparing U at all $t > t_1^*$.

Overall, player 2 is strictly better off by preparing L when player 1 prepares U at time t. Hence statement (2) is true at time t.

[30] Note that the probability of player 2 getting a revision opportunity at t_1^* is zero.

Step 2-3: By continuity (of expected payoffs with respect to time), Steps 2-1 and 2-2 imply that there exists $\epsilon > 0$ such that for all $t' \in (t - \epsilon, t]$, both statements (1) and (2) hold. Thus by Lemma 1, we have the desired result.

Step 3:

Statement (1) in Step 2 shows that at any $t < t_1^*$, player 1 prepares U. Hence for any finite $t' < t_1^*$, the probability that player 1 prepared action is U at t' converges to 1 as T increases. If player 1 prepared action is U at t', then between t' and t^*, by statement (2), player 2 must prepare L and by statement (1) player 1 keep preparing U. Hence the probability that the PAP at t_1^* is (U, L) can be made arbitrarily close to 1 by setting T large enough. Considering that the probability of revision at time t_1^* is zero, Step 1 implies that, if the PAP at t_1^* is (U, L), then players keep preparing (U, L) until the deadline. □

3.6.3 *Equilibrium Dynamics*

The proof in the previous subsection characterizes the strong player's equilibrium strategy fully but the weak player's equilibrium strategy only after the strong player prepares the action corresponding to the strong player's preferred Nash equilibrium.[31] Here we provide the full characterization of the equilibrium dynamics, which will imply that the equilibrium strategy is unique. The proof of the result stated in this subsection is provided in Calcagno and Lovo (2010) and Kamada and Sugaya (2010).

The equilibrium dynamics are summarized in Figure 3.1 for the case $t_1^* > t_2^*$. The dynamics consists of three phases. In each phase, the arrow that originates

[31] If players choose their actions simultaneously at $-T$, then it is common knowledge that the strong player prepares the action corresponding to the strong player's preferred Nash equilibrium at $-T$. Hence, the proof is sufficient to fully characterize the path of play in the revision equilibrium.

from a PAP x represents what players will prepare if they are given an opportunity to revise during that phase when the PAP is x. More specifically, an arrow from PAP (x_i, x_{-i}) to PAP (x'_i, x_{-i}) means that if player i is given an opportunity to revise at that phase when the PAP is $x \in \{(x_i, x_{-i}), (x'_i, x_{-i})\}$, then player i would prepare x'_i. If a player does not switch her action, then there is no arrow corresponding to that strategy. Hence, in particular, if there are no arrows originating from x, then that means that no player would change actions if given a revision opportunity.

When the deadline is close, each player prepares a component game best response to the PAP (each player "equilibrates"). This phase is $(t_1^*, 0]$ shown in the far-right panel of Figure 3.1 where t_1^* is given in Step 1 of the proof of Theorem 3. Since t_1^* is the time at which player 1 is indifferent between U and D, given that player 2 is preparing R, in the next phase the direction of the arrow that connects (U, R) and (D, R) is flipped. This is shown in the middle panel.

The proof shows that the directions of arrows in this figure stay unchanged for all t further back from t_1^*, except the one that connects (D, L) and (D, R). Direct calculation in Calcagno and Lovo (2010) and Kamada and Sugaya (2010) show that the direction of the arrow does change at some finite t^{**} and then stays unchanged for all further t's back.

In summary, for large T, the dynamics start from the phase where both players try to go to the (U, L) profile irrespective of the current PAP. When the deadline comes closer, there comes the second phase where player 2 would choose R given that 1 chooses D. Finally, when the deadline is close, both players prepare the component-game best-reply to the PAP.

The above observation implies that if players choose their actions at $-T$, then they will immediately select (U, L) and on the equilibrium path, no player will

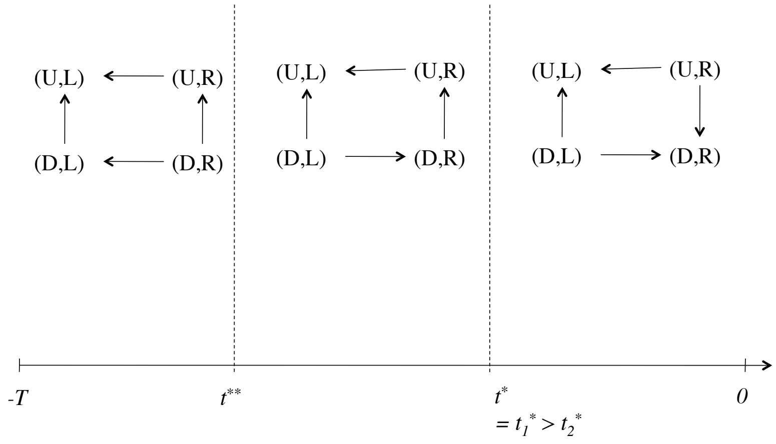

Figure 3.1: On and off equilibrium dynamics.

change actions.

3.7 Homogeneity and Asynchronicity

In the main section, we assumed that the Poisson processes are homogeneous across time (the arrival rate λ_i is time-independent) and perfectly asynchronous. In this section, we discuss the role of these assumptions.

First, consider the following non-homogeneous Poisson processes: the arrival rate for player i at time s, $\lambda_i(s)$, is a measurable function of s and so the expected number of revision opportunities between t and t' is

$$L_i(t,t') = \int_t^{t'} \lambda_i(s)ds.$$

We maintain the assumption that the Poisson processes are perfectly asynchronous.

Since the proofs of Theorems 2 and 3 do not use the fact that $\lambda_i(s)$ is constant over time, Theorems 2 and 3 hold for non-homogeneous Poisson processes. The only difference is in the expression of t_i^*, which is given in the Appendix.

Second, consider the effect of different degrees of asynchronicity. For this purpose, in addition to the two independent processes specified in Section 3.3, consider another independent Poisson process $\mathbf{p}_{12}$ with arrival rate $\lambda_{12} > 0$, at which both players revise simultaneously. For simplicity, we assume the Poisson process is homogeneous. At the time of decision corresponding to each revision opportunity, player i does not know whether such an opportunity is driven by the process $\mathbf{p}_i$ or by $\mathbf{p}_{12}$. If $\lambda_1 = \lambda_2 = 0$ and $\lambda_{12} > 0$, then all revision opportunities are synchronous and it is straightforward to show that any repetition of a Nash equilibrium is an equilibrium of the revision game. The following result shows that it is not enough to have only a slightest degree of asynchronicity to rule out the multiplicity of equilibria.

Theorem 4. *Suppose that $x \in X$ is a strict Nash equilibrium of the component game. Then, there exist strictly positive $\lambda_1, \lambda_2, \lambda_{12}$ such that it is a SPE for each player i to always prepare action x_i all the time.*

The proof and a detailed discussion can be found in Calcagno and Lovo (2010). Note that λ_1 and λ_2 in the theorem are required to be strictly positive. This means that the only slight degree of asynchronicity is not enough to eliminate multiple equilibria. This raises the question of how much of asynchronicity is needed to obtain equilibrium uniqueness in a revision equilibrium. Ishii and Kamada (2011) characterize the parameter regions such that multiplicity persists in common interest games; in particular, their results imply that the complete asynchronicity assumed in the present paper is not a knife-edge case.

3.8 Concluding Remarks

We analyzed revision games where the component game is a coordination game. Two players prepare their actions before they play a normal-form coordination game at a predetermined date where in the preparation stage, players obtain opportunities to revise their actions according to independent Poisson processes, and the finally-revised action is played at the deadline. In common interest games, the strictly Pareto-dominant profile is the only outcome of the revision game. In opposing interest games, generically there is a unique outcome of the revision game, which corresponds to one of the strict Nash equilibria. Which equilibrium prevails in the revision game depends on the payoff structure and the relative frequency with which revision opportunities arrive at each player.

Let us conclude this paper by coming back to the three key assumptions discussed in the Introduction and suggesting possible directions of future research. First, we assumed perfect observability of the opponent's play. In general, one could think of a model in which a player may only imperfectly observe the opponent's revision. It is an open question how unobservability affects the outcome of revision games.

Second, we assumed perfect asynchronicity. To understand the exact effect of asynchronicity, it is desirable to characterize the outcome in a model with both synchronous and asynchronous revision opportunities. Calcagno and Lovo (2010) formulate such processes. Section 3.7 of this paper and Ishii and Kamada (2011) analyze such a model and partially characterize the condition such that uniqueness obtains. A thorough investigation would be desirable to better understand the exact role of asynchronicity.

Third, we assume finite actions and strict incentives. Our proof hinges on these assumptions and we do not know whether these assumptions are necessary in all cases.[32] Although these assumptions hold in many applications and many coordination games discussed in the literature, it is of theoretical interest to investigate how it affects the outcome of revision games.

Finally, we focused on 2-player component games while the model can be easily extended to more-than-two-player games. At this point, we do not know which results of the 2-player set-up are robust. For example, Kamada and Sugaya (2010) provide a sufficient condition for an n-player asynchronous revision game to have a unique revision equilibrium and they give an example in which uniqueness does not hold when the condition is violated. A general characterization of revision equilibria is an important topic for future research.

[32] Note that Kamada and Kandori (2011b) show that with a continuous action space, it is possible that the set of equilibrium payoffs widens.

4. MULTI-AGENT SEARCH WITH DEADLINE[1]

4.1 Introduction

This paper studies a search problem with two features that arise in many real-life situations: The decision to stop searching is made by *multiple individuals*, and there is a *predetermined deadline* by which a decision has to be made. Our primary goal is to provide an understanding of the factors that determine the positive search duration in reality.

To fix ideas, imagine a couple who must find an apartment in a new city by September 1, as the contract with their current landlord terminates at the end of

[1] This is a joint work with Nozomu Muto (Departament d'Economia i d'Història Econòmica, Universitat Autònoma de Barcelona, and MOVE). We thank David Ahn, Attila Ambrus, Pol Antràs, Katie Baldiga, Alessandro Bonatti, Georgy Egorov, Drew Fudenberg, Chiaki Hara, Johannes Hörner, Chong Huang, Haruo Imai, Yuhta Ishii, Atsushi Kajii, Fuhito Kojima, David Laibson, Bart Lipman, Mihai Manea, Jordi Massó, Akihiko Matsui, Sho Miyamoto, Akira Okada, Wojciech Olszewski, Daisuke Oyama, Debraj Ray, Al Roth, Yuval Salant, Larry Samuelson, Bruno Strulovici, Tomasz Strzalecki, Takuo Sugaya, Satrou Takahashi, Kentaro Tomoeda, Takahiro Watanabe, Alex Wolitzky, Yuichi Yamamoto, Yosuke Yasuda, and seminar/conference participants at Universitat Autònoma de Barcelona, Brown University, Columbia University, Harvard University, Hitotsubashi University, Kyoto University, University of Tokyo, Yale University, the 22nd Summer Festival on Game Theory (International Conference on Game Theory) at Stony Brook, SWET 2011 at Hokkaido University, and GDRI Workshop Marseille for helpful comments. In addition, Morgan McClellon read through the previous version of this paper and gave us very detailed comments, which significantly improved the presentation of the paper. A portion of this research was conducted while Kamada was visiting Institut d'Anàlisi Econòmia at Universitat Autònoma de Barcelona; he thanks the university and especially Joan Esteban for the hospitality during the stay. Kamada thanks his advisors, Attila Ambrus, Al Roth, Tomasz Strzalecki, and especially Drew Fudenberg for extensive discussions and comments as well as continual guidance and support. Muto gratefully acknowledges support from the Spanish Ministry of Science and Innovation through grant "Consolidated Group-C" ECO2008-04756 and FEDER.

August. Since they are not familiar with the city, they ask a broker to identify new apartments as they become available. The availability of new apartments depends on many factors; there is no guarantee that a new apartment will become available every day. Whenever the broker finds an apartment, the husband and wife both express whether they are willing to rent it or not. If they cannot agree, they forfeit the offered apartment—since the market is seller's market, there is no option to "hold" an offer while searching for a better one. Although the couple agree on the need to rent some apartment, their preferences over specific apartments are not necessarily aligned. The search ends once an agreement is made; if the couple cannot agree on an apartment by September 1, they will be homeless.

To analyze these situations, we consider an n-player search problem with a deadline. Time is continuous and "opporunties" arrive according to a Poisson process. Opportunties are i.i.d. realizations of payoffs for each player. After viewing an opportunity, the players respond with "yes" or "no." The search ends if and only if all players say "yes." If the search does not end by the deadline, players obtain an a priori specified fixed payoff. Notice that the arrival rate of Poisson process captures "friction" inherent in the search process: larger arrival rates correspond to smaller friction. Since there is a trivial subgame perfect equilibrium in which all players always reject, we analyze an (appropriately defined) trembling-hand equilibrium of this game.

Our analysis consists of three steps. *In the first step*, we show that for any number of players and under very weak distributional assumptions, the expected duration of search does not shrink to zero even in the limit as the friction of search vanishes. Hence the mere existence of some search friction has a nonvanishing impact on the search duration. This result is intuitive but by no means obvious.[2] The incentives are complicated. Waiting for a future opportunity to arrive offers an incremental gain in payoffs, but an increased possibility of reaching the deadline. Both the rewards and the costs go to zero as the search friction vanishes; the optimal balance is difficult to quantify because agents need to make decision of before observing all future realizations of offers. For this reason, we employ an indirect proof that bounds the acceptance probability at each moment.

In the second step, we show that in the limit, expected duration increases with the number of agents involved in the search. This happens for two reasons, which we call the "ascending acceptability effect" and the "preference heterogeneity effect." Roughly put, the ascending acceptability effect refers to the fact that a player faces a larger incentive to wait if there are more opponents, as in equilibrium the opponents become increasingly willing to accept offers as time goes on. The preference heterogeneity effect refers to the fact that such these future opportunities include increasingly favorable offers for the player due to heterogeneity of preferences.

In the third step, we show that the speed of convergence for the expected search duration is fast. Moreover, we use numerical examples to show that as the friction disappears, the limit duration of search is actually close to durations with non-negligible search friction. This provides evidence that our limit analysis contains economically-meaningful content, and the mere existence of some friction is actually the main driving force of the positive duration in reality—so the effects that we identify in the first and second steps are the keys to understand the positive duration in reality.

In Figures 4.1 and 4.2, we depict how the duration can be decomposed into

[2] Indeed, we offer examples where our assumptions do not hold and the result fails.

the effects mentioned above when there are two players, the offer distribution is uniform over a feasible payoff set that has all nonnegative payoff profiles with the sum of coordinates being no more than 1 (Figure 4.3), the arrival rate is 10, and the horizon length is 1. This corresponds to the case where there are ten weeks to search an apartment, and the information of a new apartment comes once every week on average—quite a high friction. Even in this case, it is clear in the figure that the finiteness of arrival rates has a very little effect on the duration, while other effects are significant. The limit expected duration is directly computed from a key variable r determined by the details of the model (X, μ). The larger the r is, the longer the duration is. The increase in r from the one-player model to the two-player model is accounted for by the ascending acceptability effect and the preference heterogeneity effect. In this example, the former effect is larger than the latter.[3]

The two key features in our model, *deadline* and *multiple agents*, give rise to new theoretical challenges. In particular, these two things *interact* with each other. First, the existence of a deadline implies that the problem is *nonstationary*: the problems faced by the agents at different moments of time are different. Nonstationarity often results in intractability, but we partially overcome this by taking an indirect approach: we first analyze the limit expected duration (the first and the second steps) which is easier to characterize, and then argue that the limit case approximates the case with finite arrival rates reasonably well (the third step). Second, one may argue that since each player's decision at any given opportunity is essentially conditional on the situation where all other agents say "yes," the problem essentially boils down to a single-player search problem. This argument misses an

[3] At the end of the main section (Section 4.4), we will be explicit about how we conducted this decomposition.

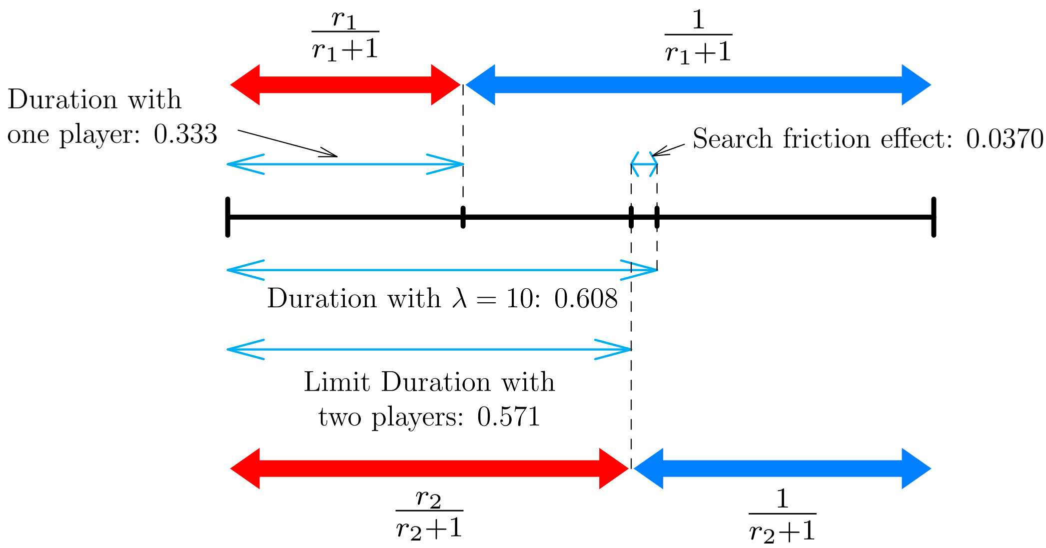

Figure 4.1: Decomposition of search durations: The case with uniform distribution over the space depicted in Figure 4.3 and the horizon length of 1. The one player duration is computed by assuming uniform distribution over the unit interval. r_1 and r_2 are illustrated in Figure 4.2.

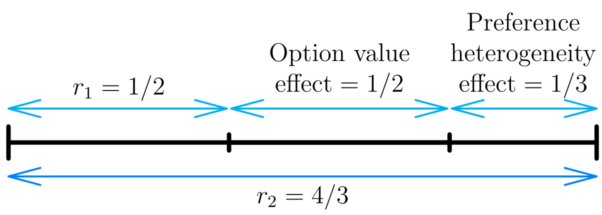

Figure 4.2: Decomposition of $r_2 - r_1$.

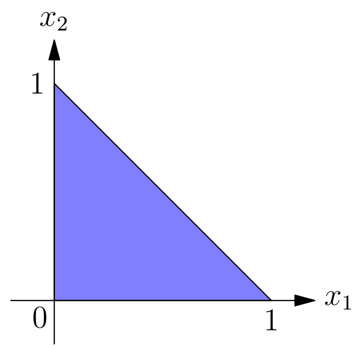

Figure 4.3: An example of the domain of feasible payoff profiles.

important key property of our model. It is indeed true that at each given opportunity the decisions by the opponents do not affect a player's decision. However, the player's expectation about the opponents' future decisions affects her decision today, and such *future* decisions by opponents are in turn affected by other agents' decisions even *further in the future*. The two "futures" discussed in the previous sentence are different precisely due to the nonstationarity—hence the two features interact. It will become clear in our analysis that it is this interaction that is crucial to our argument in the three steps.

Besides the results on duration, we provide a number of additional results. Most prominently, we study welfare implications our model. In order to isolate the effects of multiple agents and a finite horizon as cleanly as possible, the departure from the standard model is kept minimal. This enables us to modify our model in a wide variety of directions and also to conduct comparative statics. To give some examples, we study the case when payoffs realize upon agreement (corresponding to the situation where the couple can rent an apartment as soon as they sign a contract); the robustness of our results to different arrival processes; the case with the presence of fixed time costs; offer distributions varying over time; changes in bargaining power over time; the optimal choice of horizon length (in a market-design context); the case of majority rule rather than unanimity; the possibility of negotiation. All these and many other things can be and will be discussed in our framework.

4.1.1 Literature Review

Finite vs. infinite horizon with multiple agents.

First, although there is a large body of literature on search problems with a single agent and an infinite horizon, there are only few papers that diverge from these two assumptions.[4] Some recent papers in game theory discuss infinite-horizon search models in which a group of decision-makers determine when to stop. Wil-

[4] See Rogerson, Shimer, and Wright (2005) for a survey.

son (2001), Compte and Jehiel (2010), and Cho and Matsui (2011) consider search models in which a unanimous agreement is required to accept an alternative, and show that the equilibrium outcome is close to the Nash bargaining solution when players are patient. Despite the absence of a deadline, these convergence results to the Nash bargaining solution have a similar flavor to our result in Section 4.6 where payoffs realize as soon as an agreement is reached. In Section 4.7.3, we will discuss a common logic behind these convergence results. Compte and Jehiel (2010) also analyze general majority rules to discuss the power of each individual to affect outcomes of search, and the size of the set of limit equilibrium outcomes. Albrecht, Anderson, and Vroman (2010) consider general majority rules, and show that the decision-makers are less picky than the agent in the corresponding single-person search model, and the expected duration of search is shorter if they are sufficiently patient. Alpern and Gal (2009), and Alpern, Gal, and Solan (2010) analyze a search model in which a realized object is chosen when one of two decision-makers accepts it, unless one of them casts a veto which can be exercised only a finite number of times in the entire search process. Moldovanu and Shi (2010) analyze an infinite-horizon two-agent search problem with interdependent preferences with respect to private signals of the payoffs realized in every period. They also show that the duration becomes longer if the number of decision-makers increases from one to two while retaining the information structure.[5] Bergemann and Välimäki (2011) provide an efficient dynamic mechanism with a presence of monetary transfer in an n-agent model with private signals of agents' private values. Importantly, in all of these papers the search duration converges to zero as the frequency of offer

[5] Moldovanu and Shi (2010) show that agents are pickier when there is a larger conflict in preferences, whereas if the signals are public, they are less picky and the duration is shorter with a larger conflict.

arrivals tends to infinity. The key distinction is that discounting is assumed and payoffs realize upon agreement in these papers, while in our model payoffs realize at the deadline, the assumption that fits to our motivating example of apartment search.

Multiple vs. single agent search with deadline.

A single-agent search problem with deadline is explored in much detail in the operations research literature on the so-called “secretary problem.” There is an important difference between this literature and our model. In secretary problems, there are n potential candidates (secretaries) who arrive each date, and the decision maker makes acceptance decisions. The key difference from our analysis is that in secretary problems the decision maker does not have cardinal preferences but ordinal preferences, and attempts to maximize the probability that the best candidate is chosen. Since the number of candidates is finite, this is technically a search problem with finite horizon. The optimal policy as the number of candidates grows to infinity is to disregard all candidates for some time before choosing, so this model also has a positive limit search duration. The reason for positive duration is, however, different from ours. In secretary problems, the decision maker must gather information about available alternatives to make sure what she chooses is reasonably well-ranked. In our setting with cardinal preferences and known distribution of payoffs, there is no information gathering. Rather, what underlies the positive duration is the tradeoff between the potential gain from waiting for a better allocation in the future and the loss from reaching the deadline. This tradeoff is not an issue in secretary problems as the decision maker benefits only from the best candidate. See Ferguson (1989) for an extensive survey of the literature.

Single-agent search with infinite horizon.

The so-called "search theory" literature has focused mainly on a single-agent search problem with infinite horizon and extended such a model to the context of large population. Seminal papers by McCall (1970) and Mortensen (1970) explore models in which a single agent faces an iid draw of payoffs over an infinite horizon. These models are extended in many directions.[6] A common feature in these papers is that the model has some form of "waiting costs" either as a discounting or as a search cost, irrespective of the length of the horizon (finite or infinite). This assumption would be a reasonable one in their context as their main application was job search, where the overall horizon length (in finite horizon models) is several decades, and one period corresponds to a year or a month. On the other hand, our interest is in the case where the horizon length is rather short, as in the apartment search example we provided in the introduction. This naturally gives rise to the assumption that payoffs realize at the deadline—which would not have made sense in the job search application. Because of this difference, the limit search duration as the friction goes away in models of this line of the literature is zero, so they could not implement the exercise that we do in this paper. Later work extended the model to a large population model in which the search friction is given endogenously through a "matching function." Again, in a nutshell, these analyses are more or less extensions of the single-agent search model with infinite horizon, and thus there has been no question on the "limit duration" as the friction vanishes.

[6] An extensive survey of the literature can be found in Lippman and McCall (1976).

Multi-agent search with finite horizon.

A few papers consider multi-person decision problems (See Ferguson (2005) and Abdelaziz and Krichen (2007) for surveys), but none has looked at the search duration. Sakaguchi (1973) was the first to study a multi-agent search model with finite horizon. Sakaguchi (1978) proposed a two-agent continuous-time finite-horizon stopping game in which opportunities arrive according to a Poisson process as in our model. He derived the same ordinary differential equations (ODE) as ours and provided several characterizations,[7] and then computed equilibrium strategies in several specific examples.[8] However, no analysis on duration appeared in his papers. Note that obtaining the ODE constitutes only a preliminary part of our contribution; our focus is on the search duration implied by this equation.

Ferguson (2005)'s main interest is in existence and uniqueness of the stationary cutoff subgame perfect equilibrium with discrete time, general voting rules, varying distributions over time, and presence of fixed costs of search.[9,10] The sufficient condition for uniqueness that he obtains is different from ours.[11]

---

[7] Specifically, he showed that (a) the cutoffs are nondecreasing and concave in the time variable, and (b) in the independent environment, players are less picky than in the single player case.

[8] Examples he examined are (1) the Bernoulli distribution on a binary domain, (2) $h(x,y) = f(x)g(y)(1+\gamma(1-2F(x))(1-2G(x)))$ for f, g being arbitrary density functions, and γ being a parameter that measures correlation, (3) an exponential distribution, and (4) a direct product of exponential and uniform distributions. Apart from case (1) in which the limit search duration is trivially zero, our results imply that all cases have positive limit durations.

[9] He mentions the idea of trembling-hand equilibrium only verbally, and does not provide a formal definition. Instead, there is an assumption that agreement probability is always positive.

[10] He also analyzes an exponential case and does a comparative statics in terms of individual search costs.

[11] The condition states that the distribution of offers is independent across agents and the distance to the conditional expectation above value v_i is decreasing in v_i for all player i.

Multi-agent search vs. bargaining.

The multi-agent search problems are similar to bargaining problems in that both predict what outcome in a prespecified domain is chosen as a consequence of strategic interaction between agents. However, as discussed by Compte and Jehiel (2004, 2010), the search models are different from bargaining models in that in the former, players just make an acceptance decision on what is exogenously provided to them, while in the latter, players have full control over what to propose. Our model is a search model, and thus in our model players are "passively" assess exogenous opportunities. This assumption captures the feature of situations that we would like to analyze. For example many potential tenants do not design their houses for themselves, but they simply wait for a broker to pass them information regarding new apartments. The distinction between these "passive" and "active" players is important when we consider the difference between our work and the standard bargaining literature.[12]

Another important issue in relation with the bargaining literature is the distinction between positive search duration and so-called "bargaining delay." Bargaining delay is particularly important because it is often associated with inefficiency caused by discounting. In our model payoffs realize at the deadline (so in essence agents do not discount the future), so the positive-duration result does not necessary imply inefficiency. Actually, we prove that generically the expected payoff profile cannot be Pareto inefficient in the limit as the search friction vanishes. We do not view this as necessarily detrimental to our contribution, as our primary aim

[12] Cho and Matsui (2011) present another view: A drawn payoff profile in the search process can be considered as an outcome of a (unique) equilibrium in a bargaining game which is not explicitly described in the model and does not depend on the future equilibrium strategy profile. According to this interpretation, every player is "active" although the "activeness" is embedded in the model.

is to provide a deeper understanding of the positive duration in reality.[13]

Multi-agent search vs. bargaining with finite horizon.

Ambrus and Lu (2010a), Gomes, Hart, and Mas-Colell (1999) and Imai and Salonen (2011) consider a bargaining model with finite horizon, in which players obtain an opportunity to propose a share distribution of the surplus at asynchronous timings, having full control over proposals, and analyze the equilibrium payoffs.[14] The important distinction from our search model is that without any further assumptions (such as private information) that can be resolved over time or an "option to wait" as assumed in Ma and Manove (1993), the first player who obtains the opportunity makes an offer that all players would accept in equilibrium. This is in line with the intuition of Rubinstein (1982)'s canonical model of alternating-offer bargaining, and implies that as the timing of proposals becomes frequent the duration until the agreement can become arbitrarily small.[15] In our model, however, there is a trade-off as the search friction decreases between more arrivals today and more arrivals in the future. Our main objective of this paper is to discuss the effects driven (at least in part) by this trade-off, while bargaining models do not have such a trade-off (thus a question on duration is trivial).

A part of results by Gomes, Hart, and Mas-Colell (1999) and Imai and Salonen (2011) shows that in some cases the limit equilibrium is the Nash bargaining solution. Although these results about equilibrium payoff profiles is reminiscent of our

[13] In our framework, we can also ask a normative question: In Section 4.7.9, we examine a market designer's problem to tune parameters of the model (the horizon length and the distribution of offers) when the search friction is finite so the Pareto-efficiency result does not have bite.

[14] See Ambrus and Lu (2010b) for an application of their model to legislative processes.

[15] A finite horizon version of Rubinstein (1982)'s model with Poisson opportunities is a special case of Ambrus and Lu (2010a)'s model, so the limit duration is zero in such a model.

result in Section 4.6, the results are different in the conditions that determine the limit profiles.[16]

Revision games.

Broadly, this paper is part of a rapidly growing literature on "revision games," which explores implications of adding a revision phase before a predetermined deadline at which actions are implemented and players receive payoffs. The first papers on revision games by Kamada and Kandori (2011b,a) show the possibility of cooperation in such a setting,[17] and Calcagno and Lovo (2010), Kamada and Sugaya (2010), and Ishii and Kamada (2011) examine the effect of asynchronous timings of revisions on the equilibrium outcome in revision games. Kamada and Sugaya (2011) apply the revision games setting to election campaigns. Romm (2011) analyzes the implication of introducing a "reputational type" in a variant of a revision game introduced by Kamada and Sugaya (2010). General insights from these works are that when the action space is finite (as in our case) the set of equilibria is typically small and the solution can be obtained by (appropriately implemented) backwards induction, and that a differential equation is useful when characterizing the equilibrium. In our paper we follow and extend these methods to characterize equilibria and apply the framework to the context of search situations that often arise in reality. Some examples we provide in this paper are reminiscent of those provided in Kamada and Sugaya (2010).

---

[16] See Remark 2 in the previous version of this paper (Kamada and Muto (2011b)) for a more comprehensive comparison between our work and these papers. There we argue that under different conditions the limit equilibrium payoff profile is the Nash bargaining solution in each model when the discount rate and the frequency of opportunities converge simultaneously.

[17] See Ambrus and Burns (2010) for a related work on an analysis of eBay-like auctions.

The paper is organized as follows. Section 4.2 provides a model. In Section 4.3 we provide some preliminary results. In particular, we show that trembling-hand equilibria take the form of cutoff strategies, by which we mean each player at each moment of time has a "cutoff" of payoffs below which they reject offers and otherwise accept. Section 4.4 is the main section of the paper. Subsections 4.4.1, 4.4.2, and 4.4.3 correspond to Steps 1, 2, and 3 of our argument, respectively. Section 4.5 provides a welfare analysis of our main model. Section 4.6 considers the case in which payoffs realize upon agreement and there is a discounting—the case analogous to analyses in the previous work. In Section 4.7, we provide a number of discussions. Among others, we show that even if agents can negotiate and transfer utilities after each realization of payoffs, our basic result of positive duration is still valid. Section 4.8 concludes. Proofs are given in the Appendix unless otherwise noted.

4.2 Model

The Basic Setup

There are n players searching for an indivisible object. Let $N = \{1, \dots, n\}$ be the set of players. A typical player is denoted by i, and the other players are denoted by $-i$. The players search within a finite time interval $[-T, 0]$ with $T > 0$, on which opportunities of agreement arrive according to the Poisson process with arrival rate $\lambda > 0$. At each opportunity, nature draws an indivisible object which is characterized by a payoff profile $x = (x_1, \dots, x_n)$ following an identical and independent probability measure μ defined on the Borel sets of $\mathbb{R}^n$. A payoff profile $x \in \mathbb{R}^n$ is often referred to as an allocation. After allocation x is realized, each player simultaneously responds by either accepting or rejecting x without a lapse of time. Let $B = \{\text{accept}, \text{reject}\}$ be the set of responses in this search process. If all players accept, the search ends, and at time 0 the players receive the corresponding payoff profile x. If at least one of the players rejects the offer, then they continue to search. If players reach no agreement before the deadline at time 0, they obtain the disagreement payoff profile normalized at $x^d = (0, \dots, 0) \in \mathbb{R}^n$.[18]

Support and Pareto Efficiency

Let $X = \{x \in \mathbb{R}^n \mid \mu(Y) > 0 \text{ for all open } Y \ni x\}$ be the support of μ. Note that $X \subseteq \mathbb{R}^n$ is a closed subset on which μ is full support. Without loss of generality, we assume that $X \subseteq \mathbb{R}^n_+$.[19] An allocation $x = (x_1, \dots, x_n) \in X$ is *Pareto efficient* in X if there is no allocation $y = (y_1, \dots, y_n) \in X$ such that $y_i \geq x_i$ for all $i \in N$ and $y_j > x_j$ for some $j \in N$. An allocation $x \in X$ is *weakly Pareto efficient* in X if there is no allocation $y \in X$ such that $y_i > x_i$ for all $i \in N$. The set of all (weakly) Pareto efficient allocations in X is called the (*weak*, resp.) *Pareto frontier* of X. We sometimes consider weak Pareto efficiency also on $\hat{X} = \{v \in \mathbb{R}^n_+ \mid x \geq v \text{ for some } x \in X\}$ which is the nonnegative region of the comprehensive extension of X.

Assumptions

We make the following weak assumptions throughout the paper.

Assumption 1. (a) The expectation $\int_X x_i \, d\mu$ is finite for all $i \in N$.

---

[18] This is without loss of generality as long as payoffs realize at the deadline. When the payoffs realize upon agreement as in Section 4.6, this does change some of the analysis (the initial condition of the differential equation (4.10) changes), but we restrict ourselves to $x^d = (0, \dots, 0)$ as the change is minor.

[19] This is without loss of generality as long as there is a positive probability in $\mathbb{R}^n_+$ since the strategic environment is identical to the case where the arrival rate is adjusted to $\mu(\mathbb{R}^n_+)\lambda$ because players prefer to reject any negative payoffs.

(b) If $n \geq 2$, for all $i \in N$, i's marginal distribution of μ has a density function that is locally bounded.[20]

If condition (a) is violated, a player always wants to wait for better payoffs before the deadline, so a best response does not exist. Condition (b) rules out a distribution which has infinitely large density at some point, while it still allows for a distribution under which there is a positive probability that an allocation falls on degenerate subsets such as a line segment which is not horizontal or vertical. In Section 4.7.5, we will provide an example that demonstrates the need for Condition (b) for our main results to hold.

Histories and Strategies

Let us define strategies in this game. A history at $-t \in [-T, 0]$ where players observed k (≥ 0) offers in $[-T, -t)$ consists of

1. a sequence of times $(t^1, \dots, t^k)$ when there were Poisson arrivals before $-t$, where $-T \leq -t^1 < -t^2 < \cdots < -t^k < -t$,
2. allocations $x^1, \dots, x^k$ drawn at opportunities $t^1, \dots, t^k$, respectively,
3. acceptance/rejection decision profiles $(b^1, \dots, b^k)$, where each decision profile b^l ($l = 1, \dots, k$) is contained in $B^n \setminus \{(\text{accept}, \dots, \text{accept})\}$,
4. allocation $x \in X \cup \{\varnothing\}$ at $-t$ ($x = \varnothing$ if no Poisson opportunity arrives at $-t$).

[20] A function $g(y)$ defined on $\mathbb{R}$ is locally bounded if for all y, there exists $C > 0$ and $\varepsilon > 0$ such that $|g(y')| \leq C$ for all $y' \in (y - \varepsilon, y + \varepsilon)$.

We denote a history at time $-t$ by $\big((t^1, x^1, b^1), \ldots, (t^k, x^k, b^k), (t, x)\big)$. Let $\tilde{\mathcal{H}}_t^k$ be the set of all such histories at time $-t$, $\tilde{\mathcal{H}}_t = \bigcup_{k=0,1,2,\ldots} \tilde{\mathcal{H}}_t^k$ and $\tilde{\mathcal{H}} = \bigcup_{-t \in [-T,0]} \tilde{\mathcal{H}}_t$.[21] Let

$$\mathcal{H}_t^k = \big\{ \big((t^1, x^1, b^1), \ldots, (t^k, x^k, b^k), (t, x)\big) \in \tilde{\mathcal{H}}_t^k \,|\, x \neq \varnothing \big\}$$

be the history at time $-t$ when players have an opportunity and there have been k opportunities in the past. Let $\mathcal{H}_t = \bigcup_{k=0,1,2,\ldots} \mathcal{H}_t^k$ and $\mathcal{H} = \bigcup_{-t \in [-T,0]} \mathcal{H}_t$. A (behavioral) *strategy* σ_i of player i is a function from $\mathcal{H}$ to the set of probability distributions over the set of responses B. Let Σ_i be the set of all strategies of i, and $\Sigma = \times_{i \in N} \Sigma_i$. For $\sigma \in \Sigma$, let $u_i(\sigma)$ be the expected payoff of player i when players play σ.[22]

Equilibrium Notions

A strategy profile $\sigma \in \Sigma$ is a *Nash equilibrium* if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$ for all $\sigma_i' \in \Sigma_i$ and all $i \in N$. Let $u_i(\sigma \,|\, h)$ be the expected continuation payoff of player i given that a history $h \in \tilde{\mathcal{H}}$ is realized and strategies taken after h is given by σ. A strategy profile $\sigma \in \Sigma$ is a *subgame perfect equilibrium* if $u_i(\sigma_i, \sigma_{-i} \,|\, h) \geq u_i(\sigma_i', \sigma_{-i} \,|\, h)$ for all $\sigma_i' \in \Sigma_i$, $h \in \mathcal{H}$, and all $i \in N$. A strategy $\sigma_i \in \Sigma_i$ of player i is a *Markov strategy* if for history $h \in \mathcal{H}_t$ at $-t$, $\sigma_i(h)$ depends only on the time $-t$ and the drawn allocation x. A strategy profile $\sigma \in \Sigma$ is a *Markov perfect equilibrium* if σ is a subgame

[21] Precisely speaking, there are histories in which infinitely many opportunities arrive. We ignore these possibilities since such histories happen with probability zero.

[22] The function $u(\sigma)$ is well-defined for the following reason: $\mathcal{H}^k := \bigcup_t \mathcal{H}_t^k$ is seen as a subset of $\mathbb{R}^{(2n+1)k+(n+1)}$, and thus endowed with a Borel sigma-algebra. We assume that $\mathcal{H} = \bigcup_k \mathcal{H}^k$ is endowed with a sigma-algebra induced by these sigma-algebras on $\mathcal{H}^k$, and a strategy must be measurable with respect to this sigma-algebra. The measurability ensures that a strategy profile generates a probability measure on the set of terminal nodes. See Stinchcombe (1992) for a detailed treatment of strategies in general continous-time games.

perfect equilibrium and σ_i is a Markov strategy for all $i \in N$. We will later show that players play a Markov perfect equilibrium (except at histories in a zero-measure set) if they follow a trembling-hand equilibrium defined below. For $\varepsilon \in (0, 1/2)$, let Σ^ε be the set of strategy profiles which prescribe probability at least ε for both responses in $\{\text{accept}, \text{reject}\}$ after all histories in $\mathcal{H}$. A strategy profile $\sigma \in \Sigma$ is a *trembling-hand equilibrium* if there exists a sequence $(\varepsilon^m)_{m=1,2,...}$ and a sequence of strategy profiles $(\sigma^m)_{m=1,2,...}$ such that $\varepsilon^m > 0$ for all m, $\lim_{m\to\infty} \varepsilon^k = 0$, $\sigma^m \in \Sigma^{\varepsilon^m}$, σ^m is a Nash equilibrium in the game with a restricted set of strategies Σ^{ε^m} (ε^m-constrained game) for all m, and $\lim_{m\to\infty} \sigma^m(h) = \sigma(h)$ for all $h \in \mathcal{H}$ according to the pointwise convergence in histories.[23]

4.3 Preliminary Results

In this section, we present preliminary results which will become useful in the subsequent sections. We will show that there exists an essentially unique trembling-hand equilibrium, in which every player plays a "cutoff strategy." We will derive an ordinary differential equation that characterizes the cutoff profile in the equilibrium. In addition, we will observe a basic invariance: The change in equilibrium continuation payoffs when raising the arrival rate is the same as that when stretching the duration from the deadline with the same ratio. Finally, by examining the differential equation, the limit equilibrium payoff as $\lambda \to \infty$ is shown to be weakly efficient.

The next proposition shows that trembling-hand equilibrium is essentially

[23] This equilibrium concept is an analog of extensive-form trembling-hand equilibrium, as opposed to its normal-form counterpart. Although our extensive-form game involves uncountably many nodes and hence the standard definitions of trembling-hand equilibria are not directly applicable, it is for this reason that we call this notion a trembling-hand equilibrium.

unique and Markov.

Proposition 12. *Suppose that σ and σ' are two trembling-hand equilibria. Then $u_i(\sigma \mid h) = u_i(\sigma' \mid h')$ for almost all histories $h, h' \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ and all $i \in N$.*

That is, regardless of the history, any two trembling-hand equilibria give rise to the same continuation payoff at time $-t$. Three remarks are in order: First, we ruled out histories in $\mathcal{H}_t$, because different realization of payoffs that players accept clearly give rise to different continuation payoffs on the equilibrium path. Second, since agents move simultaneously, there exist subgame perfect equilibria in which all players reject any allocations.[24] We introduced the trembling-hand equilibrium to rule out such trivial equilibria. In an ε-constrained game, a player will optimally accept a favorable allocation for herself, expecting the others to accept it with a small probability. Third, and relatedly, there exist sequential equilibria in which every player has a strict incentive at almost all histories.[25] Our trembling-hand equilibrium rules out such equilibria.

A Markov strategy σ_i of player $i \in N$ is a *cutoff strategy* with cutoff $v_i(t) \geq 0$ if player i who is to respond at time $-t$ accepts allocation $x \in X$ whenever $x_i \geq v_i(t)$, and rejects it otherwise. For a profile $v = (v_1, \ldots, v_n) \in \mathbb{R}^n_+$, we define a set of allocations by $A(v) = \{x \in X \mid x_i \geq v_i \text{ for all } i \in N\}$. When players play cutoff strategies with cutoff profile $(v_1(t), \ldots, v_n(t))$, we sometimes call $A(v(t))$ an "acceptance set" as they agree with an allocation x at time $-t$ if and only if $x \in A(v(t))$. We often denote this acceptance set by $A(t)$ with a slight abuse of

[24] If players respond sequentially, we can show that any subgame perfect equilibrium consists of cutoff strategies. Therefore our results are essentially independent of the timing of responses of players.

[25] An example similar to the one in Cho and Matsui (2011, Proposition 4.4) can be used to show this result.

notation when the cutoff profile in consideration is not ambiguous. Suppose that all players play Markov strategies σ, and there is no Poisson arrival at time $-t \in [-T, 0]$. Then player i has an expected payoff $u_i(\sigma \,|\, h)$ at $-t$ independent of history $h \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ played before time $-t$. We denote the continuation payoff at time $-t$ by $v_i(t, \sigma) = u_i(\sigma \,|\, h)$. For simplicity of notations, we hereafter omit to write a cutoff strategy profile σ explicitly, and denote by $v_i(t)$ the continuation payoff of player i at time $-t$.

The following proposition shows that there exists a trembling-hand equilibrium that consists of cutoff strategies, and characterizes the path of cutoffs.

Proposition 13. *There exists a trembling-hand equilibrium that consists of (Markov) cutoff strategies. Moreover, an equilibrium continuation payoff profile $v(t) = (v_1(t), \ldots, v_n(t))$ at time $-t \in [-T, 0]$ is given by a solution of the following ordinary differential equation (ODE)*

$$v'(t) = \lambda \int_{A(t)} \big(x - v(t)\big) \, d\mu \tag{4.1}$$

with an initial condition $v(0) = (0, \ldots, 0)$.

This proposition is shown by the following argument. An equilibrium continuation payoff $v_i(t)$ supported by a cutoff strategy profile is given by the following recursive expression: For $i \in N$,

$$\begin{aligned} v_i(t) &= \int_0^t \left(\int_{X \setminus A(\tau)} v_i(\tau) \, d\mu + \int_{A(\tau)} x_i \, d\mu \right) \lambda e^{-\lambda(t-\tau)} \, d\tau \\ &= \int_0^t \left(v_i(\tau) + \int_{A(\tau)} \big(x_i - v_i(\tau)\big) \, d\mu \right) \lambda e^{-\lambda(t-\tau)} \, d\tau. \end{aligned} \tag{4.2}$$

After time $-t$, players receive the first Poisson opportunity at time $-\tau$ with probability density $\lambda e^{-\lambda(t-\tau)}$. If player i finds that the drawn payoff x_i to her is no worse than her continuation payoff $v_i(\tau)$, i optimally accepts this allocation x, otherwise, i rejects it. If all players accept x, i.e., $x \in A(\tau)$, then they reach an agreement with x. If some player rejects x, then search continues with continuation payoff profile $v(\tau)$. This discussion shows that a cutoff strategy profile with cutoffs $v(t)$ characterized by equation (4.2) is a Markov perfect equilibrium.

Bellman equation (4.2) implies that $v_i(t)$ is differentiable in t. Multiplying both sides of (4.2) by $e^{\lambda t}$ and differentiating both sides yield the ordinary differential equation (4.1) of continuation payoff profile $v(t)$ defined in $\hat{X}$.

Now, a standard argument of ordinary differential equations shows that ODE (4.1) has a solution whenever Assumption 1 holds.[26] The above argument only shows that the cutoff strategy profile with a cutoff profile given by this solution of ODE (4.1) is a Markov perfect equilibrium. In the Appendix, we will show that it is in fact a trembling-hand equilibrium.

By Proposition 12, the solution of ODE (4.1) is unique. Therefore the game has essentially a unique trembling-hand equilibrium for any given X and μ satisfying Assumption 1. Let us denote the unique solution of (4.1) by $v^*(t;\lambda)$, the continuation payoff profile in the trembling-hand equilibrium. We simply denote this by $v^*(t)$ as long as there is no room for confusion. The probability that all players accept a realized allocation at time $-t$ on the equilibrium path conditional on the event that an opportunity arrives at $-t$ is referred to as the "acceptance probability" at time $-t$.

Let us make a couple of observations about ODE (4.1). Figure 4.4 describes an illustration of a typical path and the velocity vector that appear in this ODE for

[26] This is because Assumption 1 (b) ensures continuity in v of the right hand side of (4.1). See Coddington and Levinson (1955, Chapter 1) for a general discussion about ODE.

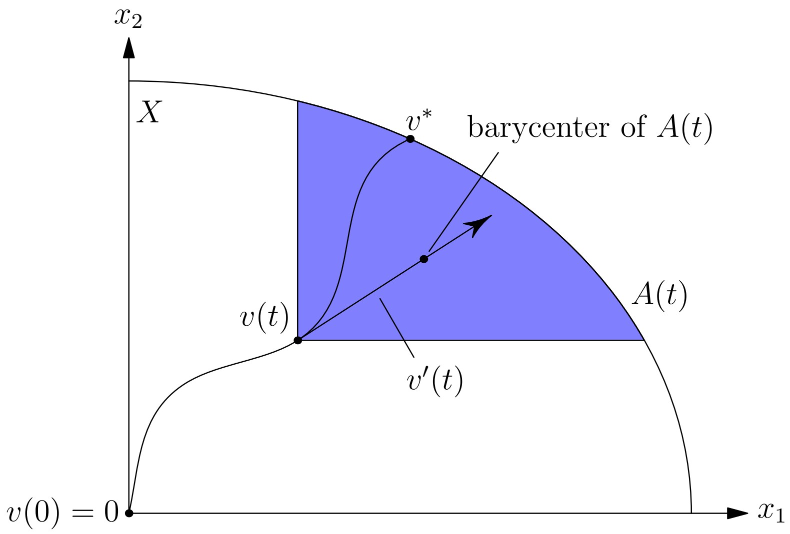

Figure 4.4: The path and the velocity vector of ODE (4.1).

$n = 2$. The shaded area shows the acceptance set $A(t)$, whose barycenter with respect to the probability measure μ is $\int_{A(t)} x d\mu / \mu(A(t))$. Therefore the velocity vector $v^{*\prime}(t)$ is parallel to the vector from $v^*(t)$ to the barycenter of $A(t)$, which represents the gain upon agreement relative to $v^*(t)$. The absolute value of $v^{*\prime}(t)$ is proportional to the weight $\mu(A(t))$. Note that ODE (4.1) immediately implies $v_i^{*\prime}(t) \geq 0$ for all t and $i \in N$, and $v_i^{*\prime}(t) = 0$ if and only if $\mu(A(t)) = 0$. For each $i \in N$, the continuation payoff $v_i^*(t)$ grows as t increases, and eventually either converges to a limit payoff v_i^*, or diverges to infinity.

Since the right hand side of ODE (4.1) is linear in λ, we have $v^*(t; \alpha\lambda) = v^*(\alpha t; \lambda)$ for all $\alpha > 0$ and all t such that $-t, -\alpha t \in [-T, 0]$. By considering the limit as $\alpha \to \infty$, we have the following proposition:

Proposition 14. *The two limits of $v^*(T;\lambda)$ coincide, i.e.,* $\lim_{\lambda\to\infty} v^*(T;\lambda) = \lim_{T\to\infty} v^*(T;\lambda)$*, if one of them exists.*

We henceforth denote this limit by v^*. In the next section, we sometimes deal with these two limits interchangeably. Note that the equality implies $\lim_{\lambda\to\infty} v^*$ $(T;\lambda)$ does not depend on $T > 0$.

Finally, we argue weak Pareto efficiency of the limit allocation. Suppose that $v^* = \lim_{\lambda\to\infty} v^*(T) = \lim_{T\to\infty} v^*(T)$ exists but is not weakly Pareto efficient. Then there exists $x \in X$ that strictly Pareto dominates v^*. Since x belongs to the support of μ, $\mu(Y) > 0$ for any open set $Y \subseteq \mathbb{R}^n_+$ that includes x. For Y sufficiently small, $A(v^*)$ contains Y, and thus $\mu(A(v^*)) > 0$. This implies that the right hand side of ODE (4.1) is positive, contradicting the fact that $v^* = \lim_{\lambda\to\infty} v^*(t) = \lim_{t\to\infty} v^*(t)$. Hence we obtain the following proposition:

Proposition 15. *Let $t > 0$ fixed, and suppose that the solution $v^*(t;\lambda)$ of equation* (4.1) *converges to $v^* \in \hat{X}$ as $\lambda \to \infty$. Then v^* is weakly Pareto efficient.*[27]

We will have further discussions about efficiency in Section 4.5, in which we will show that the limit allocation v^* is Pareto efficient for almost all distributions μ satisfying mild assumptions, and Pareto efficient for all convex X.

4.4 Duration of Search

In this section, we will discuss the duration of search in our model. Our argument consists of three steps: In Section 4.4.1 we will show that even under the quite weak conditions in Assumption 1, search takes a positive time even in the limit as the friction vanishes. In Section 4.4.2, we argue that the limit duration becomes longer as the number of involved agents gets larger. This extra duration is caused

[27] Note that this does not necessarily imply weak Pareto efficiency in the convex hull when X is nonconvex. That is, the convex hull can contain allocations that Pareto-dominate the limit expected payoff profile, while such allocations cannot be achieved under a trembling-hand equilibrium. See footnote 38 for a further discussion on this.

by two effects called the "ascending acceptability effect" and the "preference heterogeneity effect." We will provide a method to decompose the extra duration by these two effects. In Section 4.4.3, we demonstrate that the limit duration is "close" to the durations for finite arrival rates. This provides evidence that our limit analysis contains economically-meaningful content, and the mere existence of some friction is actually the main driving force of positive duration in reality—so the effects that we identify in Steps 1 and 2 are the keys to understand the positive duration in reality.

First, let us explain how we compute the expected duration. Given arrival rate $\lambda > 0$, by the differential equation (4.1) we can compute the equilibrium path of the cutoff profile $v^*(t;\lambda)$. Given $v^*(t;\lambda)$, one can compute the acceptance probability $p(t;\lambda)$ at each time $-t$ as follows:

$$p(t;\lambda) = \mu(A(v^*(t;\lambda))).$$

Let $P(t;\lambda)$ be the probability that there is no agreement until time $-t$:

$$P(t;\lambda) = e^{-\int_t^T \lambda p(s;\lambda)\,ds}. \tag{4.3}$$

Notice that $\frac{dP(t;\lambda)}{dt} = \lambda p(t;\lambda)P(t;\lambda)$. We often omit λ and simply denote $p(t)$ and $P(t)$ when there is no room for confusion. As an example, Figure 4.5 graphs $p(t)$ for $\lambda = 1, 10, 100$ when $n = 2$, $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$, μ is the uniform distribution on X, and $T = 1$. Figure 4.6 shows a graph of $P(t)$ in the same environment.

Let $D(\lambda)$ be the expected duration in the equilibrium for given λ when $T = 1$. Since we have $v^*(T;\lambda) = v^*(1;\lambda T)$ as discussed in Section 4.3, the search duration

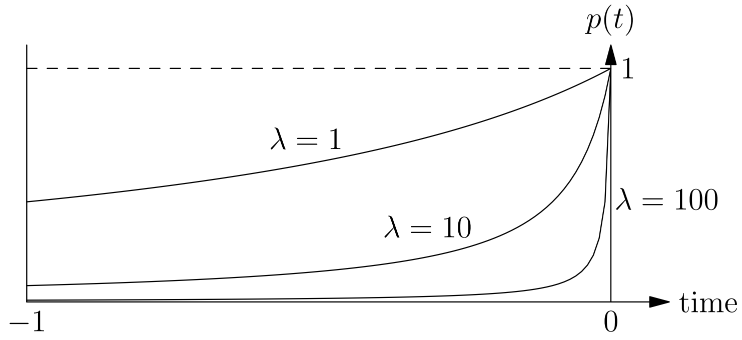

Figure 4.5: A numerical example of $p(t)$ for the case when $n = 2$, $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$, μ is the uniform distribution on X, and $T = 1$.

is proportional to T, and thus the expected duration for general T is written as $D(\lambda)T$. We use these $p(t;\lambda)$ and $P(t;\lambda)$ to solve for $D(\lambda)$, using integration by parts:

$$\begin{aligned} D(\lambda)T = T \cdot & \underbrace{P(0;\lambda)}_{\substack{\text{The probability of no agreement} \\ \text{until time } 0}} \\ + \int_0^T & \underbrace{(T-t)}_{\substack{\text{The duration when the search} \\ \text{ends at time } -t}} \cdot \underbrace{P(t;\lambda)}_{\substack{\text{The probability that} \\ \text{the search does not end until } -t}} \cdot \underbrace{\lambda p(t;\lambda)}_{\substack{\text{The probability density of} \\ \text{agreement at time } -t}} \, dt \\ &= T \cdot P(0;\lambda) + [(T-t)P(t;\lambda)]_0^T + \int_0^T P(t;\lambda)\, dt \\ &= \int_0^T P(t;\lambda)\, dt. \end{aligned} \tag{4.4}$$

This final expression has a direct interpretation: For each time $-t$, $P(t)$ is the probability that the duration is greater than $T - t$. Since $P(t) > P(t')$ for $t > t'$, $P(t)$ is integrated for the length of $T - t$ (from T to t). Thus the expression measures the expected duration.

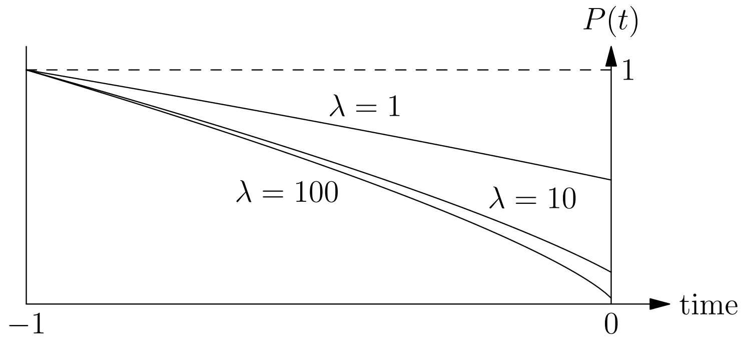

Figure 4.6: A numerical example of $P(t)$ for the case when $n = 2$, $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$, μ is the uniform distribution on X, and $T = 1$.

Finally, define

$$D(\infty) = \lim_{\lambda \to \infty} D(\lambda)$$

whenever the limit is well-defined.

In Steps 1 and 2, we will analyze $D(\infty)$. Then in Step 3 we will demonstrate that $D(\lambda)$ converges to $D(\infty)$ reasonably fast.

4.4.1 Step 1: Positive Duration

The first step of our argument shows the following: *For any number of players n and any probability distribution over feasible allocations μ satisfying fairly weak assumptions, the limit expected search duration as the search friction vanishes is strictly positive.*

We first show the result for the case with $n = 1$ (Theorem 5) and detail the intuition. Then, using this result, we generalize to an arbitrary number of players (Theorem 6).

Single Agent

Roughly, there are two effects of having a higher arrival rate. On one hand, for any (small) given time interval, there will be an increasing number of opportunities, thus it becomes easier to get a lucky draw. On the other hand, since there will be more and more opportunities in the future as well, the player becomes pickier. Our result shows that these two effects balance each other out. The incentives are complicated. Waiting for a future opportunity to arrive offers an incremental gain in payoffs, but an increased possibility of reaching the deadline. Both the rewards and the costs go to zero as the search friction vanishes; the optimal balance is difficult to quantify because agents need to make decision of before observing all future options.

To explain the detailed intuition for our result, let us specialize to the case of $X = [0, 1]$ and μ being the uniform distribution. We first show that if the acceptance probability at each time $-t$ is $O(\frac{1}{\lambda t})$ then the limit duration is strictly positive.[28] Then we show that the acceptance probability must be indeed $O(\frac{1}{\lambda t})$.

Suppose the acceptance probability at each time $-t$ is $O(\frac{1}{\lambda t})$. Then, the probability that the agreement does not take place by time $-\frac{T}{2}$ is at least

$$e^{-\lambda C \frac{1}{\lambda T/2}} = e^{-2\frac{C}{T}}$$

for some constant $C > 0$, and this is strictly positive. This means that the limit expected duration is at least a strict convex combination of 0 and $\frac{T}{2}$, and therefore is strictly positive.

Now we explain why we expect such a small acceptance probability. Fix time

[28] For functions $g(y)$ and $h(y)$, we say that $g(y) = O(h(y))$, if there exist $C > 0$ and $\bar{y}$ such that $|g(y)| \leq C \cdot |h(y)|$ for all $y \geq \bar{y}$.

$-t$. Note that the cutoff at $-t$ must be equated with the continuation payoff at $-t$ by optimality at $-t$, and the continuation payoff must be at least as good as the expected payoff by playing some arbitrarily specified strategy from time $-t$ on by optimality in the future. Also, the cutoff at $-t$ uniquely determines the acceptance probability at $-t$. That is, by specifying a future strategy, we can obtain a lower bound of continuation payoff which must be equal to the cutoff, and this gives us an upper bound of the acceptance probability:

$$\begin{aligned}
&\text{The acceptance probability} \\
&= 1 - \text{ the cutoff of the optimal strategy} \\
&= 1 - \text{ the continuation payoff from the optimal future strategy} \\
&\leq 1 - \text{ the continuation payoff from an arbitrarily specified future strategy.}
\end{aligned}$$

To see what type of future strategy will generate an interesting bound, first consider specifying a constant cutoff from $-t$ on. Suppose that at any time $-s$ after $-t$ the cutoff is $1 - O(\frac{1}{\lambda t})$. Then, a lower bound of the probability that there will be no acceptance in the future can be calculated as

$$e^{-\lambda C \frac{1}{\lambda t}} = e^{-\frac{C}{t}}$$

for some constant $C > 0$, and this is strictly greater than 0 irrespective of λ. This means that even in the limit as $\lambda \to \infty$, the probability of no agreement at time 0 does not shrink to zero. But then, the continuation payoff from this strategy must be at most a strict convex combination of a number at most 1 (the best possible payoff) and 0 irrespective of λ, which means that the acceptance probability is at least a positive number independent of λ. Hence $p(t)$ cannot be $O(\frac{1}{\lambda t})$.

Next, consider a future strategy such that at any time $-s$ after $-t$ the cutoff is such that the player accepts with a higher order than $\frac{1}{\lambda t}$ (thus she accepts with a higher probability; e.g., $\frac{1}{\sqrt{\lambda t}}$). Then the probability of acceptance in the future indeed tends to 1 as $\lambda \to \infty$, but the payoff conditional on acceptance is smaller than the best payoff (i.e., 1) by the amount of the order higher than $\frac{1}{\lambda t}$. Hence the cutoff at $-t$ must be smaller than the best payoff by such an amount, which means that the acceptance probability at $-t$ is of the order higher than $\frac{1}{\lambda t}$.

The analysis of the above two scenarios reveals the tradeoff faced by the player: Setting a high cutoff gives her a high payoff conditional on acceptance, but reduces the acceptance probability. On the other hand, setting a low cutoff results in a low payoff conditional on acceptance but raises the acceptance probability. This suggests thata good strategy must specify a high cutoff for a sufficiently long time to ensure a high payoff conditional on acceptance, and lower cutoffs towards the end to ensure a high enough acceptance probability. Specifically, consider the cutoff $1 - \frac{2}{\lambda s+2}$ for each time $-s$ after time $-t$. This plan has a feature that for any finite time $s > 0$, the acceptance probability is

$$\frac{2}{\lambda s + 2} = \frac{\lambda t + 2}{\lambda s + 2} \cdot \frac{2}{\lambda t + 2} = O\Big(\frac{1}{\lambda t}\Big),$$

thus for any positive future time, the player's payoff conditional on acceptance is smaller than the best payoff by the amount $O(\frac{1}{\lambda t})$. Yet this gives us the limit acceptance probability of 1, as the probability for no acceptance can be calculated as:

$$e^{-\int_0^t \lambda \frac{2}{\lambda s+2} ds} = e^{-[2\ln(\lambda s+2)]_0^t} = \Big(\frac{2}{\lambda t + 2}\Big)^2 \to 0 \text{ as } \lambda \to \infty.$$

A rough intuition for why this can achieve the limit acceptance probability of 1 is that, for each time subinterval $[-\frac{T}{2^{k-1}}, -\frac{T}{2^k}]$ for $k = 1, 2, \ldots$, this strategy makes an acceptance with probability at least

$$1 - e^{-\lambda \cdot \frac{2}{\lambda(T/2^{k-1})+2} \cdot \frac{T}{2^k}} = 1 - e^{-\frac{\lambda T}{\lambda T + 2^k}} \to 1 - \frac{1}{e} > 0 \quad \text{as } \lambda \to \infty.$$

Since $1 - \frac{1}{e}$ is a positive constant independent of k, the acceptance probability increases with an exponential speed as k increases. We can indeed check that the future cutoff scheme $\frac{2}{\lambda s+2}$ gives the player a continuation payoff of $\frac{2}{\lambda t+2} = O(\frac{1}{\lambda t})$ at time $-t$.

Overall, we have shown that when $X = [0, 1]$ and the distribution μ is uniform, the limit expected duration is strictly positive. This argument is generalized to the cases of general distributions satisfying Assumption 1 and the following assumption. Let $F(x)$ be the cumulative distribution function of μ.

Assumption 2. There exists concave function φ such that $1 - \varphi(x)$ is of the same order as $1 - F(x)$ in $\{x \in \mathbb{R} \mid F(x) < 1\}$.[29]

To see what this assumption implies, consider two separate cases—bounded X and unbounded X. If X is bounded, besides pathological cases where F is non-differentiable at infinitely many points, the assumption amounts to say that the slope of the cdf F cannot diverge to infinity at the maximum payoff. If X is unbounded, a simple sufficient condition to guarantee that the assumption holds is to require there exists $\tilde{x}$ such that F is concave on $(\tilde{x}, \infty)$, or equivalently, there exists a nonincreasing density function f on $(\tilde{x}, \infty)$. Concavity of φ lets us invoke the Jensen's inequality to bound the cumulative acceptance probability.

[29] For functions $g(y)$ and $h(y)$, we say that $g(y)$ is *of the same order as* $h(y)$ in $Y \subseteq \mathbb{R}$ if there exist $c, C > 0$ and $\bar{y} < \sup(Y)$ such that $c|h(y)| \leq |g(y)| \leq C|h(y)|$ for all $y \geq \bar{y}$.

Recall that $D(\lambda)T$ is the expected duration in the equilibrium for given arrival rate λ. Then we obtain the following:

Theorem 5. *Suppose $n = 1$. Under Assumptions 1 and 2,* $\liminf_{\lambda\to\infty} D(\lambda) > 0$.

In the Appendix we also provide a proof that the conclusion of this result holds when X is bounded with another assumption: for $\bar{x} \in X$, $\ln\big(\mu\big(\{x \in X \,|\, |\bar{x} - x| \leq \varepsilon\}\big)\big)$ is of the same order as $\ln \varepsilon$ when $\varepsilon > 0$ is small. A sufficient condition for this is that there exists $\alpha > 0$ such that $\mu\big(\{x \in X \,|\, |\bar{x} - x| \leq \varepsilon\}\big)$ is of the same order as ε^{α} when $\varepsilon > 0$ is small.

Multiple Agents

Now we extend our argument to the case of $n \geq 2$. The basic argument is the same as in the case of $n = 1$: We fix some strategies for players other than i, and consider bounding i's continuation payoff. However, it is not the case that we can implement this proof for any given strategies by the opponents. To see this point, consider the case of 2 players with $X = \{x \in \mathbb{R}^2_+ \,|\, x_1 = x_2 \leq 1\}$ and the uniform distribution. Suppose that we are given player 2's strategy to set the cutoff $v_2 = 0$ for the time interval $\left[-t, -\left(t - \frac{1}{\sqrt{\lambda t}}\right)\right]$, and then the cutoff $v_2 = 1$ for the rest of the time. Then, player 1's upper bound of acceptance probability cannot be given by $O(\frac{1}{\lambda t})$ because, to ensure the acceptance of a positive payoff, player 1 must accept within the time interval $\left[-t, -\left(t - \frac{1}{\sqrt{\lambda t}}\right)\right]$, and to do so she must set a low enough cutoff.[30]

What is missing in the above strategy of player 2 is the feature that a player's cutoff must be decreasing over time. In the above strategy, the cutoff starts from 0

[30] There also exist strategies for player 2 that are independent of λ and still give rise to a low cutoff for player 1, such as $v_2(t) = e^{-(T-t)}$.

and then jumps up to 1. We use the decreasingness to show our result.

To see how the decreasingness helps, fix t and consider player $-i$'s equilibrium cutoffs at time $-t$, and suppose for the moment that they will keep using these cutoffs in the future as well. Then, by the result in the case of $n = 1$, we know that the acceptance probability at $-t$ by playing optimally in the future against such strategies is $O(\frac{1}{\lambda t})$ as long as Assumption 2 is met for any cutoff profiles of the other players (sufficient conditions for this to hold are analogous to what we discussed after introducing Assumption 2). Let $p(s)$ for $s < t$ be the acceptance probability given by i's optimal strategy against $-i$'s fixed strategies. Now, consider the actual equilibrium cutoff strategy for $-i$ and consider a new future strategy for player i, which is to accept at each time $-s$ with probability $p(s)$. Notice that, since each opponent's cutoff is decreasing, the expected payoff conditional on acceptance at each time $-s$ must be greater than the case with fixed cutoffs for $-i$, while at each moment the acceptance probability is identical to that case. This means that i's continuation payoff at $-t$ must be higher than in the original case, which implies that the acceptance probability at $-t$ must be $O(\frac{1}{\lambda t})$.

Hence, we obtained the following: Recall that $D(\lambda)T$ is the expected duration in the equilibrium for given arrival rate λ. Let $F_i^{v_{-i}}$ be the marginal cumulative distribution function of player i's payoff conditional on cutoff profiles v_{-i} of the other players with $\mu(A(0, v_{-i})) > 0$.

Assumption 2′. *There is $i \in N$ such that for all v^{-i} with $\mu(A(0, v_{-i})) > 0$, there exists a concave function φ such that $1 - \varphi(x_i)$ is of the same order as $1 - F_i^{v_{-i}}(x_i)$ in $\{x_i \in \mathbb{R} \mid F_i^{v_{-i}}(x_i) < 1\}$.*[31]

Theorem 6. *Suppose $n \geq 2$. Under Assumptions 1 and 2′,* $\liminf_{\lambda \to \infty} D(\lambda) > 0$.

[31] This assumption reduces to Assumption 2 when $n = 1$.

4.4.2 Step 2: Effects of the Larger Number of Agents

The second step of our argument concerns the effect of having a larger number of players.

In what follows we demonstrate that *there are two reasons that we expect longer durations when there are more players.* The effects that underlie these reasons are called *ascending acceptability effect* and *preference heterogeneity effect*. We explain these effects in turn.

Ascending Acceptability Effect

In Section 4.4.1 we demonstrated that the decreasingness of the opponents' cutoffs can be used to reduce the acceptance probability (through the rise of continuation payoffs). The ascending acceptability effect is also based on the fact that the opponents' cutoffs are decreasing.

To isolate such an effect, let us consider the case when we add players whose preferences are independent of those of the existing players. Specifically, let two problems $(X, \mu), (Y, \gamma)$ satisfy Assumption 1 where $X \subseteq \mathbb{R}^n_+$, $Y \subseteq \mathbb{R}^m_+$, $\mu \in \Delta(X)$, $\gamma \in \Delta(Y)$, and $n, m \geq 1$. Consider three models: (i) n player model (X, μ), (ii) m player model (Y, γ), and (iii) $n + m$ player model $(X \times Y, \mu \times \gamma)$.

Theorem 7. *Suppose that the limit expected durations exist for models (i) and (ii) with $T = 1$, denoted by D_X and D_Y, respectively. Then the limit expected duration D_{XY} also exists in model (iii) with $T = 1$, and satisfies*

$$D_{XY} = 1 - \frac{(1 - D_X)(1 - D_Y)}{1 - D_X D_Y} \tag{4.5}$$

if $D_X D_Y < 1$, and $D_{XY} = 1$ if $D_X D_Y = 1$.

The reasoning of this proposition will be given in Section 4.4.2. Theorem 7 implies an immediate corollary:

Corollary 4. *Under the assumption in Theorem 7,* $D_{XY} > \max\{D_X, D_Y\}$ *if* $D_X, D_Y \in (0, 1)$.

In Section 4.4.2, we provide the explicit formula for the probability distribution of the expected duration. The formula in particular implies that the distributions of the durations in models (i) and (ii) are first-order stochastically dominated by that of model (iii), which implies Corollay 4.

There is a simple reasoning behind Corollay 4. Note first that the locus of the path in model (iii) projected on X is identical to the one in model (i) because, by (4.1), the direction of the vector is determined by the position of the barycenter in the acceptance set. Notice further that if we exogenously specify the strategies of additional m players to be the ones that accept any payoff profiles, then the time path of the cutoffs for the original n players should remain unchanged. In equilibrium, however, these m players' cutoffs are decreasing, so there will be more chances for desirable draws to be accepted ("ascending acceptability"). This is why we expect a longer duration with more players. Another way to put this is that the increase in the acceptance probability caused by additional m players corresponds to an increase in arrival rates over time. This means that a larger fraction of opportunities comes at the late stage of the game, so we expect a longer duration.

To understand the formula (4.5) in Theorem 7, manipulate it to get:

$$\frac{\text{The expected remaining time in model (iii)}}{\text{The expected remaining time in model (i)}} = \frac{1 - D_{XY}}{1 - D_X} = \frac{1 - D_Y}{1 - D_X D_Y} < 1. \quad (4.6)$$

Notice that $1 - D_X$ denotes the expected remaining time until the deadline at the time of agreement given X, and the same interpretation is valid for $1 - D_{XY}$. Thus, the left hand side of equation (4.6) is the ratio of remaining time with $X \times Y$ compared to that of X. This ratio is strictly smaller than 1 whenever $D_X < 1$, by the expression in the right hand side, and increases as D_Y grows. This is intuitive: Higher D_Y implies a slower speed for the continuation payoffs of the additional players to move. Thus players in X have more incentives to wait than in the case with a lower D_Y.

Preference Heterogeneity Effect

Theorem 7 considers the case where preferences of players in model (i) and those of players in model (ii) are independent. In many relevant cases, players' preferences are not independent. Specifically, they are often heterogeneous. We now analyze how heterogeneity in preferences, captured by the change in X and μ, affects the duration. In this subsection, we first provide a general duration formula to understand how preference heterogeneity affects the duration. Then we use this formula in specific examples to analyze the effect of preference heterogeneity.

Let us define values $\underline{r}, \overline{r}$ as follows:

$$\underline{r} = \liminf_{t \to \infty} \sum_{i \in N} \underline{d}_i(v^*(t)) \cdot b_i(v^*(t)), \quad \overline{r} = \limsup_{t \to \infty} \sum_{i \in N} \overline{d}_i(v^*(t)) \cdot b_i(v^*(t)) \qquad (4.7)$$

where

$$
\begin{aligned}
b_i(v) &= g_i(A(v)) - v_i, \quad b(v) = (b_1(v), \ldots, b_n(v)), \\
\underline{d}_i(v) &= \frac{1}{\mu(A(v))} \cdot \\
&\quad \liminf_{\varepsilon \to 0} \frac{\mu(A((v_j + \varepsilon b_j(v))_{j<i}, (v_j)_{j \geq i})) - \mu(A((v_j + \varepsilon b_j(v))_{j \leq i}, (v_j)_{j>i})))}{\varepsilon |b(v)|}, \\
\overline{d}_i(v) &= \frac{1}{\mu(A(v))} \cdot \\
&\quad \limsup_{\varepsilon \to 0} \frac{\mu(A((v_j + \varepsilon b_j(v))_{j<i}, (v_j)_{j \geq i})) - \mu(A((v_j + \varepsilon b_j(v))_{j \leq i}, (v_j)_{j>i})))}{\varepsilon |b(v)|},
\end{aligned}
$$

and $g(Y) = (g_1(Y), \ldots, g_n(Y))$ denotes a barycenter of the set $Y \subseteq \mathbb{R}^n$ with respect to μ. Recall that $P(t;\lambda)$ is the probability of no agreement until time $-t$, and $D(\lambda)$ is the limit expected duration when $T = 0$. Now we can show that $P(t;\infty) = \lim_{\lambda \to \infty} P(t;\lambda)$ can be written in the following way:

Theorem 8. *Under Assumption 1, for all $-t \in [-T, 0]$*

$$
\left(\frac{t}{T}\right)^{1/\underline{r}} \leq \liminf_{\lambda \to \infty} P(t;\lambda) \quad \text{and} \quad \limsup_{\lambda \to \infty} P(t;\lambda) \leq \left(\frac{t}{T}\right)^{1/\overline{r}}, \text{ and}
$$

$$
\frac{1}{1 + \underline{r}^{-1}} \leq \liminf_{\lambda \to \infty} D(\lambda) \quad \text{and} \quad \limsup_{\lambda \to \infty} D(\lambda) \leq \frac{1}{1 + \overline{r}^{-1}}.
$$

Thus, if $\underline{r} = \overline{r} =: r$, then for all $-t \in [-T, 0]$

$$
P(t;\infty) = \left(\frac{t}{T}\right)^{1/r} \quad \text{and} \quad D(\infty) = \frac{1}{1 + r^{-1}}.
$$

Proof Sketch. Let us provide a proof when $r = \underline{r} = \overline{r}$, and μ has a density function. A formal proof in the general case is given in the Appendix. To show the result, we first prove

$$\lim_{\lambda\to\infty} p(t)\cdot\lambda t = \frac{1}{r}$$

where $p(t) = \mu(A(v^*(t)))$. By the ODE (4.1), $v^{*\prime}(t) = \lambda(g(A(v^*(t))) - v^*(t))\cdot p(t)$. Differentiating $p(t) = \mu(A(v^*(t)))$,

$$\begin{aligned}
p'(t) &= \sum_{i\in N} \frac{\partial \mu(A(v))}{\partial v_i}\Big|_{v=v^*(t)} v_i^{*\prime}(t)\\
&= \sum_{i\in N} \frac{\partial \mu(A(v))}{\partial v_i}\Big|_{v=v^*(t)} \lambda\cdot(g_i(A(v^*(t))) - v_i^*(t))\cdot p(t)\\
&= -\sum_{i\in N} d_i(v^*(t))p(t)\cdot\lambda b_i(v^*(t))\cdot p(t).
\end{aligned}$$

Therefore

$$\frac{p'(t)}{\lambda p(t)^2} = -\sum_{i\in N} d_i(v^*(t))b_i(v^*(t)).$$

This implies that r is the limit of $-p'(t)/\lambda p(t)^2$ as $t\to\infty$. If the limit exists, for any $\varepsilon > 0$ there exists $\bar{t}$ such that $t \ge \bar{t}$ implies

$$r-\varepsilon \le -\frac{p'(t)}{\lambda p(t)^2} \le r+\varepsilon. \tag{4.8}$$

This means that $p(t)$ is approximated by the solution of ODE $p'(t) = -r\lambda p(t)^2$ with an initial condition at $t = \bar{t}$. Solving this equation, for large t,

$$p(t) \approx \frac{1}{r\lambda(t-\bar{t}) + p(\bar{t})^{-1}}.$$

Hence we get $\lim_{\lambda\to\infty} p(t)\cdot\lambda t = \frac{1}{r}$. We can compute the approximated probabil-

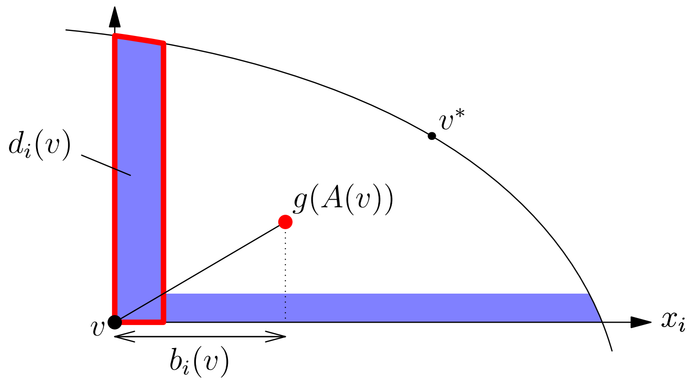

Figure 4.7: Density term and barycenter term.

ity of disagreement P by following formula (4.3), showing that $P(t;\infty) = \left(\frac{t}{T}\right)^{\frac{1}{r}}$.[32] Applying formula (4.4), one can easily obtain $D(\infty) = \frac{1}{1+r^{-1}}$. $\square$

Moreover, Theorem 8 immediately implies that if $\underline{r} = \overline{r} = r$, then $P(t;\infty)$ is increasing in t, and (a) for $r < 1$, $P(t;\infty)$ is concave and $\lim_{t\to 0} P'(t;\infty) = \infty$, (b) for $r = 1$, $P(t;\infty)$ is linear and $P'(t;\infty) = \frac{1}{T}$, and (c) for $r > 1$, $P(t;\infty)$ is convex and $\lim_{t\to 0} P'(t;\infty) = 0$.

The graphical intuition for the formula in Theorem 8 is depicted in Figure 4.7. The first term $d_i(v^*(t))$, which we call the *density term*, is i's marginal density at her continuation payoff conditional on the distribution restricted to the acceptance set. The second term $b_i(v^*(t))$, which we call the *barycenter term*, measures the distance between the barycenter of the acceptance set and the cutoff. Remember that the speed at which the cutoff moves towards the limit point is determined by this distance, by equation (4.1). Hence the formula for r in equation (4.7) measures the speed at which the acceptance set shrinks. This is consistent with the fact that

[32] The computation is given in Lemma 7 in the Appendix.

the duration formula in equation (4.7) is increasing in $r > 0$ because if the acceptance probability shrinks quickly, then players reject with high probability for a long time, resulting in a long duration.

Theorem 8 also explains the reasoning behind the formula in Theorem 7. Let r_X, r_Y, and r_{XY} be associated with models (i), (ii), and (iii), respectively. r_X and r_Y are well-defined as D_X and D_Y exist. Then, by definition, r_{XY} must be equal to $r_X + r_Y$. Hence the limit expected duration in model (iii) exists and it is $\frac{1}{1+r_{XY}^{-1}} = \frac{1}{1+(r_X+r_Y)^{-1}}$. Rearranging terms, we get the formula in Theorem 7.

Now we use the formula given in Theorem 8 to analyze specific classes of games to understand the preference heterogeneity effect.

Example 1 (Bounded X, smooth boundary, and continuous and strictly positive density).

Here we impose assumptions employed often in the literature on multi-agent search (Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2011)).

Assumption 3. (a) X is convex and compact subset of $\mathbb{R}^n_+$, and has a smooth Pareto frontier.

(b) The probability measure μ is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^n$, and admits a probability density function f that is continuous and bounded away from zero, i.e., $\inf_{x \in X} f(x) > 0$.

In this case, when the cutoffs for players are high enough, the acceptance set can be approximated with an n-dimensional pyramid by the assumption of smooth boundary, and the distribution over this acceptance can be approximated with a uniform distribution due to the assumption of continuous and strictly positive density. This allows us to explicitly compute the limit expected duration, as follows:

Table 4.1: Limit expected duration of search as opportunities arrive more and more frequently.

Number of agents	1	2	3	5	10	100
Limit expected duration	.333	.571	.692	.806	.901	.990

Proposition 16. *Under Assumptions 1 and 3,* $\lim_{\lambda\to\infty} D(\lambda) = \dfrac{n^2}{n^2+n+1}$.

Corollary 5. *Under Assumptions 1 and 3,* $\lim_{\lambda\to\infty} D(\lambda)$ *is increasing in n.*

In the proof in the Appendix, we show this result under much more general assumptions (Assumptions 1, 4, and 7).

The solution of the expected duration provided in Proposition 16 implies that, if only two players are involved in search, the expected duration is $\frac{4}{7}T$, and it monotonically increases to approach T as n gets larger. Table 4.1 shows the limit expected duration for several values of n when $T = 1$.

If Assumption 3 holds, then the limits in $\overline{d}_i$ and $\underline{d}_i$ coincide. Let us denote $d_i(v) = \overline{d}_i(v) = \underline{d}_i(v)$. When X is bounded, as assumed in Assumption 3, $d_i(v)$ grows to infinity as v comes close to the Pareto frontier, while $b_i(v)$ decreases to zero. For this reason we normalize these terms as follows: For $Y \subseteq X$, let $s(Y) = V(Y)^{\frac{1}{n}}$ be the "size" of Y in X. Let us define the normalized terms as

$$\begin{aligned}\tilde{d}_i(v) &= d_i(v)s(A(v))\\ &= \frac{s(A(v))}{\mu(A(v))}\cdot\\ &\quad \lim_{\varepsilon\to 0}\frac{\mu(A((v_j+\varepsilon b_j(v))_{j<i},(v_j)_{j\geq i}))-\mu(A((v_j+\varepsilon b_j(v))_{j\leq i},(v_j)_{j>i})))}{\varepsilon|b(v)|},\\ \tilde{b}_i(v) &= \frac{b_i(v)}{s(A(v))} = \frac{g_i(A(v))-v_i}{s(A(v))}.\end{aligned}$$

Notice that if $V(A(v)) = 1$, then $\tilde{d}_i(v) = d_i(v)$ and $\tilde{b}_i(v) = b_i(v)$.

When the acceptance set is approximated by an n-dimensional pyramid whose area is normalized to 1, the density term is $\left(n^{\frac{n-1}{n}}\right) \times n$ where "$\times n$" accounts for the fact that there are n agents, and the barycenter term is $n^{\frac{1}{n}}/(n+1)$.

Under Assumption 3, the ascending acceptability effect can be seen in Figure 4.7 by noting that the area that corresponds to the density term has two segments (n segments in the case of n players), each corresponding to each player. Thus, adding a player results in an extra region of payoffs that will be accepted in the future. The probability density in the extra region increases not only because the number of segments increases, but also because the length of each segment increases. This happens precisely because players' preferences are heterogeneous so the density of the marginal distribution of a player's payoff increases as her payoff decreases. This means that the "extra region" that a player's opponents accept in the future contains relatively more favorable allocations for the player when there are more opponents. Although the barycenter term deceases due to this preference heterogeneity as well, the overall effect is positive. We call this effect the preference heterogeneity effect. Note well that the effect of preference heterogeneity is to (only) magnify the ascending acceptability effect by lengthening the length of each segment.

The argument so far suggests that, under Assumption 3, we can make the following decomposition:

$$\text{The limit duration with } n \text{ agents} = \frac{1}{1 + \left(n \cdot \frac{n}{n+1}\right)^{-1}}$$

$$= \frac{1}{1 + \Bigg(\underbrace{n \cdot n^{\frac{n-1}{n}}}_{\text{density term}} \cdot \underbrace{\frac{n^{\frac{1}{n}}}{n+1}}_{\text{barycenter term}} \Bigg)^{-1}}.$$

□

We next use the formula (4.7) to conduct several comparative statics with respect to preference heterogeneity, holding fixed the number of players.

Example 2 (Change in the shape of X around the limit payoff profile).

First, consider the two-player case, under Assumptions 1 and 3 but dropping the assumption that the Pareto frontier is smooth (assumed in Assumption 3 (b)). In this environment, generally X has a kink at the limit expected payoff, so that the acceptance set when t is large can be approximated by a quadrilateral similar to $co\{(0,0)(1,0)(0,1)(q,q)\}$ after rescaling each axis. In this case the limit duration is computed as

$$\text{The limit duration} = \frac{1}{1 + \left(2 \cdot \frac{2q+1}{6q}\right)^{-1}} = \frac{1}{1 + \Bigg(\underbrace{2 \cdot \frac{1}{\sqrt{q}}}_{\text{density term}} \cdot \underbrace{\frac{2q+1}{6\sqrt{q}}}_{\text{barycenter term}} \Bigg)^{-1}}. \quad (4.9)$$

Notice that the term corresponding to the preference heterogeneity effect, $\frac{2q+1}{6q}$, is decreasing in q. This is consistent with how the shape of acceptance set changes with respect to q. As q grows, the kink of the boundary at the limit payoff becomes sharper, so the preferences among the players become less heterogeneous. This means that the “extra region” that a player’s opponent accepts in the future does not contain relatively favorable allocations for the player. As a result, the limit duration can be calculated as $D(\infty) = \frac{2q+1}{5q+1}$, and this is decreasing in q. □

Example 3 (**Change in the shape of X under Assumptions 1 and 3**).

Next, consider n-player symmetric X and μ. Consider a transformation of this problem in the following sense: let X^q and $\bar{X}^a$ defined by

$$X^q = \{x \in X \mid \max_{i \in N} x_i - \min_{j \in N} x_j \leq q\} \quad \text{and} \quad \bar{X}^a = \{y^a(x) \mid x \in X\}$$

where $y^a(x) = ax + (1-a)x^e$, $a \in (0,1]$, with $x^e = \left(\frac{x_1+\cdots+x_n}{n}, \ldots, \frac{x_1+\cdots+x_n}{n}\right)$. Define μ^q by $\mu^q(C) = \frac{1}{\mu(X^q)} \cdot \mu(C \cap X^q)$ and $\bar{\mu}^a$ by $\bar{\mu}^a\left(\{y^a(x) \mid x \in C\}\right) = \mu(C)$ for any $C \subseteq X$.

Both μ^q and $\bar{\mu}^a$ shrink the distribution to the middle: μ^q takes out the offers that give agents "too asymmetric" payoffs, while $\bar{\mu}^a$ moves each point by the amount proportional to the original distance to the equi-payoff line. See Figure 4.8 for a graphical description in the case of two players. Proposition 16 shows that as long as Assumptions 1 and 3 are met, expected duration is unaffected by the specificality of distribution μ. That is,

Proposition 17. *If (X, μ) satisfies Assumptions 1 and 3, then the limit expected durations with (X^q, μ^q) and $(\bar{X}^a, \bar{\mu}^a)$ are the same as in the case with (X, μ) for any $a \in [0, 1)$.*

The intuition is simple. In both cases, the distribution is still uniform around the limit point and the boundary is smooth even under $\bar{\mu}^a$, so exactly the same calculation as in the case with μ suggests that the limit duration is $\frac{n^2}{n^2+n+1}$. In this case, however, durations with finite arrival rates are affected by the change in preferences. Table 4.2 shows that the duration becomes shorter as we make the preferences less heterogeneous, in the case $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$. □

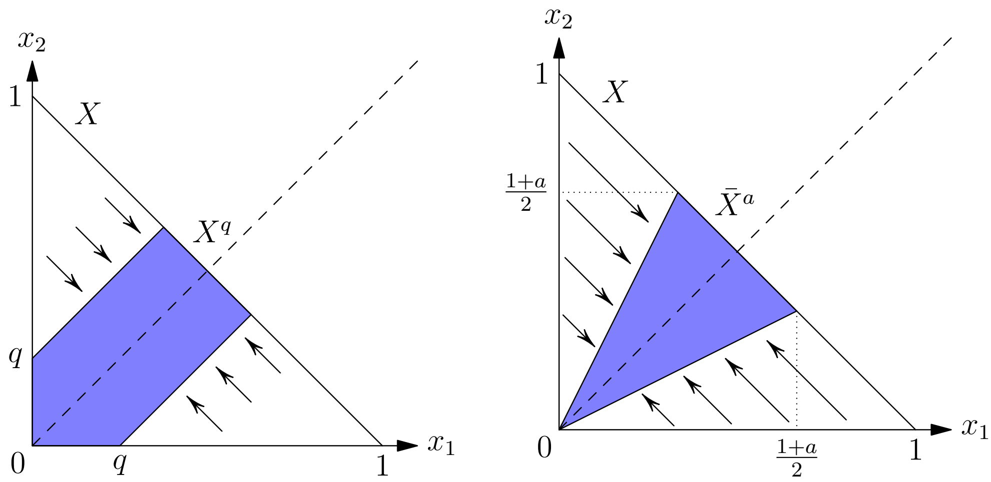

Figure 4.8: Transfer of allocations in the negotiation for μ^q (left) and $\bar{\mu}^a$ (right).

Example 4 (Change in the distribution over unbounded X).

Consider 2-player symmetric $X = \mathbb{R}^2_+$ and μ which is associated with a density function f_σ parameterized by σ as follows:

$$f_\sigma(x_1, x_2) \propto \begin{cases} e^{-(x_1+x_2)} \cdot \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_1-x_2)^2}{2\sigma^2}} & \text{if } (x_1, x_2) \in \mathbb{R}^2_+. \\ 0 & \text{otherwise,} \end{cases}$$

That is, we consider an exponential distribution in the direction of 45 degree line, and a normal distribution with variance σ^2 in the direction of 135 degree line. The parameter σ measures the heterogeneity of preferences. Notice that the limit distribution as $\sigma \to 0$ is the (degenerate) exponential distribution over 45 degree line, and the limit distribution as $\sigma \to \infty$ is the product measure in which each player's marginal distribution is an exponential distribution with parameter $\sqrt{2}$. We can solve for the limit duration in these two cases analytically using the duration formula in Theorem 8. In the former case the problem is isomorphic to that of one-

Table 4.2: Preference heterogeneity effect under Assumptions 1 and 3. q and a measure heterogeneity of preferences.

μ^q	λ				
	10	20	30	100	∞
$q=1$	0.608	0.591	0.585	0.576	0.571
$q=0.8$	0.607	0.590	0.584	0.575	0.571
$q=0.6$	0.600	0.586	0.581	0.575	0.571
$q=0.4$	0.579	0.574	0.573	0.572	0.571
$q=0.2$	0.515	0.534	0.544	0.562	0.571
$q=0$	0.398	0.366	0.355	0.340	0.333

$\bar{\mu}^a$	λ				
	10	20	30	100	∞
$a=1$	0.608	0.591	0.585	0.576	0.571
$a=0.8$	0.625	0.601	0.591	0.578	0.571
$a=0.6$	0.604	0.588	0.583	0.575	0.571
$a=0.4$	0.567	0.566	0.567	0.570	0.571
$a=0.2$	0.489	0.512	0.528	0.557	0.571
$a=0$	0.398	0.366	0.355	0.340	0.333

player case and the limit expected duration as $\lambda \to \infty$ is $\frac{1}{2}$, and in the latter case it is $\frac{2}{3}$. We use the duration formula to numerically compute the limit duration in the intermediate values of σ, and the result is given in the graph of Figure 4.9. For any $\sigma > 0$, the density term is a constant $\frac{1}{\sqrt{2}}$. It is the barycenter term that varies with σ. Specifically, the barycenter term rises with σ from $\frac{1}{\sqrt{2}}$ (when $\lambda \to 0$) to $\sqrt{2}$ (when $\lambda \to \infty$). This is because, the more heterogeneous the preferences are, the more realizations of payoffs are scattered outside of the acceptance set.[33] Since it

[33] The effect that the total probability on the acceptance set decreases for a given value of v does not matter, as we take the limit as $\lambda \to \infty$.

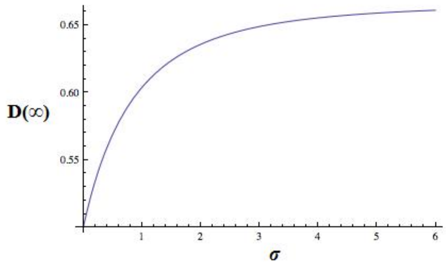

Figure 4.9: Preference heterogeneity (measured by σ) and the limit search duration.

is more difficult for a realization to stay in the acceptance set if the sum of payoffs are smaller, heterogeneity implies that, conditional on acceptance, payoffs are high on average. Thus, if preferences are more heterogeneous (σ is larger), the opponent's gain relative to the continuation payoff conditional on acceptance is higher. This means that the loss from a unit time passing is larger, so the opponent will decrease the cutoff faster. This makes the incentive to wait larger, implying a longer expected duration. □

4.4.3 *Step 3: Finite Arrival Rate*

Our results on the expected duration so far suggest that there are reasons to expect a positive duration of search even in the limit as the friction of search vanishes. To evaluate the significance of these reasons, we now consider cases with finite arrival rates. We will show that *the expected duration converges to the limit duration very fast*, provides evidence that our limit analysis contains economically-meaningful content—so the effects that we identify in the previous discussion are the keys to understand the positive duration in reality.

First, we show that the convergence speed of the duration is high. Recall that

$D(\lambda)$ and $D(\infty)$ are the expected durations under arrival rate λ and the limit expected duration for $T = 1$, respectively. Theorem 8 ensures the existence of $D(\infty)$ if $\underline{r} = \bar{r}$.

Theorem 9. *Under Assumption 1, if $\underline{r} = \bar{r}$, then $|D(\lambda) - D(\infty)| = O(\frac{1}{\lambda})$.*

This is a fast rate of convergence; for example, when payoffs realize upon agreement and there is a positive discount rate (with a finite or infinite horizon), $|D(\lambda) - D(\infty)|$ is of the same order as $\frac{1}{\lambda^{\frac{1}{n+1}}}$ under Assumptions 1 and 3.

We further support our claim numerically through a number of examples. We find that the limit duration of Proposition 16 is not far away from those with finite λ in many cases. The differential equation (4.1) does not have a closed-form solution in general, and even if it does, $D(\lambda)$ may not have a closed-form solution as it involves further integration. For this reason, we solve the differential equation and integration numerically to obtain the values of $D(\lambda)$ for specific values of λ. We considered the following distributions standard in the literature with $T = 1$. Note that, in the apartment search example, if the couple has ten weeks before the deadline and a broker provides information of an apartment once per week on average (very *in*frequent case), the situation corresponds to $\lambda = 10$.

Case 1: μ is the uniform distribution over $X = \{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i \leq 1\}$ for $n = 1, 2, 3$ and $\lambda = 10, 20, 30, 100$.

Case 2: μ is the uniform distribution over $X = \{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i^2 \leq 1\}$ for $n = 1, 2$ and $\lambda = 10, 20, 30, 100, 1000$.

Case 3: μ is the uniform distribution over $X = \{x \in \mathbb{R}^n_+ \mid \max_{i \in N} x_i \leq 1\}$ for $n = 1, 2, 3$ and $\lambda = 10, 20, 30, 100$.

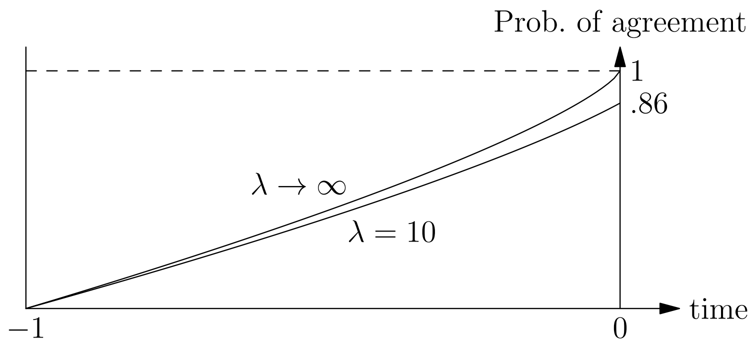

Figure 4.10: A numerical example of the cumulative probability of agreement.

Case 4: μ is the product measure over $X = \mathbb{R}^n_+$ where each marginal corresponds to an exponential distribution with parameter $a_i > 0$ for $n = 1, 2, 3, 10$ and $\lambda = 10, 20, 30, 100$.

Case 5: μ is the product measure over $X = \mathbb{R}^n_+$ where each marginal corresponds to a log-normal distribution with mean 0 and standard deviation $\sigma = \frac{1}{4}, 1, 4$ for $n = 1$ and $\lambda = 10, 20, 30, 100$.

Figure 4.10 shows a graph of the cumulative probability of agreement for $\lambda = 10$ (i.e., $1 - P(t; 10)$) and for $\lambda \to \infty$ (i.e., $1 - \lim_{\lambda \to \infty} P(t; \lambda)$) of Case 1 with $n = 2$. Also, Figure 4.11 shows the probability density function of the duration of search in such a case (i.e., $P(t; \lambda) \cdot p(t; \lambda)$).

In Table 4.3, we provide the computed values for selected choices of parameter values and cases. We provide the complete description of all the computed values in the Appendix.[34]

According to our calculation, $D(\lambda)$ is within 10% difference from $D(\infty)$ except

[34] Some values are computed analytically: The results for Case 1 for $n = 1, 2$, Case 2 for $n = 1$, Case 4 for $n = 1, 2, 3, 10$ are analytical.

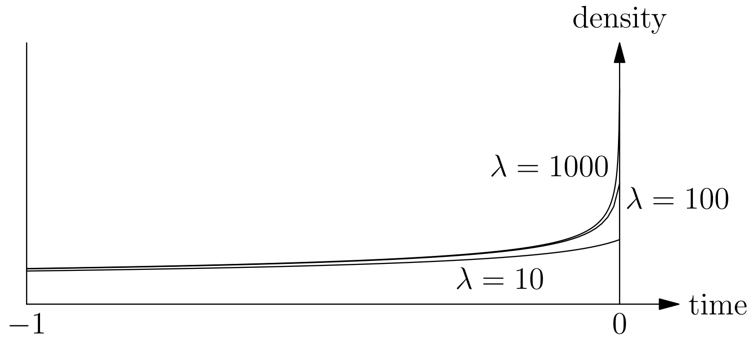

Figure 4.11: The probability density of the duration of search.

for a single case where the difference is 19.4%, which happens in Case 1 with $n = 1$. Generally the percentage falls as the number of agents becomes larger and the arrival rate goes up.[35] For example if we add another player in Case 1, the difference falls down dramatically to 6.5%, and if we increase the arrival rate to 20 (fixing the number of players at $n = 1$), the difference becomes 9.9%. In all other cases the difference is much smaller and often less than 5%.[36] Notice that we predict "overshooting" of the expected duration in Case 2 than in 1. This is because when the continuation value is far away from the boundary, the shape of the acceptance set is close to a square with which we expect a shorter duration, and gradually the shape approaches a triangle (precisely, the density effect *would* be smaller than the case of a triangle *if* the limit shape of the acceptance set *were* the same as that of X). This suggests that convexity of the set of available allocations, which is often assumed in the literature, facilitates a faster convergence. When X is unbounded,

---

[35] The monotonicity with respect to arrival rates can be analytically proven for Case 1 with $n = 2$. However, the monotonicity fails in general. To see this, consider the case in which $D(\infty) = 1$. By optimality it must be the case that $D(\lambda) < 1$ for any finite λ, so in this case the duration cannot be decreasing in λ. Note also that in Case 2, after the "overshooting" the duration comes back to the limit duration, thus $D(\lambda)$ is nonmonotonic.

[36] We are planning to extend the analysis to more cases beyond the setting provided here.

Table 4.3: Expected durations for finite arrival rates.

Case 1

		λ				
		10	20	30	100	∞
$n = 1$	Expected duration	0.398	0.366	0.355	0.340	0.333
	Percentage (%)	19.4	9.92	6.64	2.00	0
$n = 2$	Expected duration	0.608	0.591	0.585	0.576	0.571
	Percentage (%)	6.48	3.44	2.35	0.731	0

Case 2

		λ					
		10	20	30	100	1000	∞
$n = 1$	Expected duration	0.398	0.366	0.355	0.340	0.334	0.333
	Percentage (%)	19.4	9.92	6.64	2.00	0.200	0
$n = 2$	Expected duration	0.582	0.568	0.565	0.562	0.567	0.571
	Percentage (%)	1.90	-0.541	-1.21	-1.61	-0.798	0

Case 4

		λ				
		10	20	30	100	∞
$n = 1$	Expected duration	0.545	0.524	0.516	0.505	0.5
	Percentage (%)	9.09	4.76	3.23	0.990	0
$n = 2$	Expected duration	0.693	0.681	0.676	0.670	0.667
	Percentage (%)	3.91	2.11	1.45	0.465	0

the computed difference was much smaller (Cases 4 and 5).

The discussion so far enables us to perform the decomposition mentioned in the Introduction. The expected duration in the 2-player model with the uniform distribution over $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$ and $\lambda = 10$ is 0.608. The limit expected duration as $\lambda \to \infty$ in this case is $\frac{4}{7}$, so the difference is 0.037. This limit duration $\frac{4}{7}$ is calculated from the number r that we denote by $r_2 := \frac{4}{3}$. When there is only one player and the distribution is uniform over $[0, 1]$, the limit duration is $\frac{1}{3}$, and the number r is $r_1 := \frac{1}{3}$. The difference between r_2 and r_1—the difference caused by adding one more player—is determined by two effects, the ascending acceptability effect and the preference heterogeneity effect. To calculate the ascending acceptability effect, we compute r that we would obtain if this additional agent's distribution over feasible payoffs is independent of the original player's, and the distribution corresponds to the uniform distribution over $[0, 1]$. The duration and r in this case are $\frac{1}{2}$ and $r_{\text{op}} := 1$, respectively, and the difference in terms of r is given by $r_{\text{op}} - r_1 = 1 - \frac{1}{2} = \frac{1}{2}$. Now the preference heterogeneity effect is the change in r caused by the change in distribution from this product measure to X. This is given by $r_2 - r_{\text{op}} = \frac{4}{3} - 1 = \frac{1}{3}$. In general, fixing an n-player model (X, μ) and an $(n + m)$-player model (Y, γ), we can solve for the ascending acceptability effect by computing the difference between the r in the model (X, μ) and the r in the model $(X \times [0, 1]^m, \mu \times (U[0, 1])^m)$. Then the preference heterogeneity effect can be computed by solving for the difference in the latter r and the r in the model (Y, γ).[37] This decomposition is well-defined in the sense that the ascending acceptability effect (resp. preference heterogeneity effect) of changing the models from an n-player model (X, μ) to an $(n + m)$-player model (Y, γ) is identical to the sum of ascending acceptability effect (resp. preference heterogeneity effect) of changing the models from an n-player model (X, μ) to an $(n + l)$-player model (Z, δ) and the ascending acceptability effect (resp. preference heterogeneity effect) of changing models from an $(n + l)$-player model (Z, δ) to an $(n + m)$-player model (Y, γ) where $l < m$, since

[37] The uniform distribution over $[0, 1]$ can be replaced with any distribution with a positive continuous density over a compact interval without changing the computation.

r is additive.

4.5 *Welfare Implications*

In Section 4.3, we showed that the limit expected payoff must be weakly Pareto efficient if the limit exists. In this section we seek further welfare implications. Let us impose the following assumption to rule out uninteresting cases:

Assumption 4. (a) X is a compact subset of $\mathbb{R}^n$.

(b) X coincides with the closure of its interior (with respect to the standard topology of $\mathbb{R}^n$).

(c) The probability measure μ is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^n$, and admits a probability density function f.

(d) The probability density f is bounded above and away from zero, i.e., $\sup_{x\in X} f(x) < \infty$ and $\inf_{x\in X} f(x) > 0$.

Condition (a) in Assumption 4 is a standard assumption when we consider welfare implications. Note that we do not assume convexity here. Condition (b) rules out irregularities involving lower dimensional subsets. For example, if X has an isolated point this condition is violated, because the interior of X does not contain any isolated points. Condition (c) implies that $\mu(Y) = 0$ for any $Y \subseteq X$ that has (n-dimensional) Lebesgue measure zero. Condition (d) is a condition that makes our analysis tractable.

In general, v^* is not necessarily (strictly) Pareto efficient in X even if v^* exists. There is an example of a distribution μ satisfying Assumptions 1 and 4 in which $v^*(t)$ converges to an allocation that is not Pareto efficient.

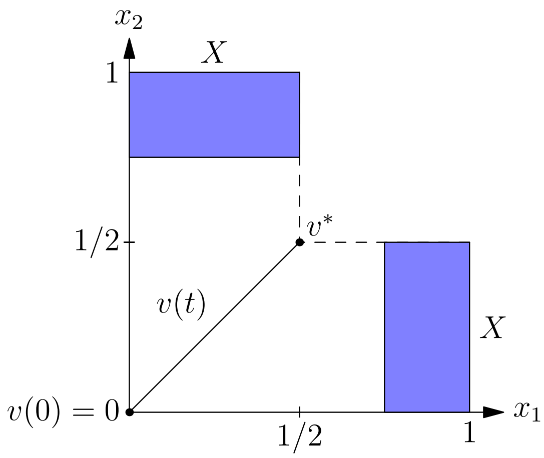

Figure 4.12: A path that converges to a weakly Pareto efficient allocation.

Example 5. Let $n = 2$, $X = \big([0, 1/2] \times [3/4, 1]\big) \cup \big([3/4, 1] \times [0, 1/2]\big)$, and f be the uniform density function on X, which is shown in Figure 4.12. By the symmetry with respect to the 45 degree line, we must have $v_1^*(t) = v_2^*(t)$ for all t. Therefore $v^* = (1/2, 1/2)$, which is not Pareto efficient in X.[38] □

Note that v^* is weakly Pareto efficient, and that X is a non-convex set in this example. In fact, we can show that v^* is strictly Pareto efficient if X is convex. Furthermore, even if X is not convex, we can show v^* is "generically" Pareto efficient, that is, v^* is Pareto efficient in X for any generic f that satisfies Assumptions 1 and 4.

[38] There are (non-trembling-hand) subgame perfect equilibria in which players obtain a more efficient payoff profile than $(1/2, 1/2)$. For example, consider a strategy profile in which players agree with allocations close to $(1, 1/2)$ or $(1/2, 1)$, and if one of the players rejects such allocations, both players reject all allocations after the deviation. This is a subgame perfect equilibrium and gives players expected payoffs close to $(3/4, 3/4)$ in the limit. Similar constructions show that any allocations in the convex hull of general nonconvex X can be an expected payoff profile supported by a subgame perfect equilibrium. However, we rule out such subgame perfect equilibria in a view that rejecting anything after deviation is not a credible threat if a player expects the others to accept with a small probability.

Formally, let $\mathcal{F}$ be the set of density functions that satisfy Assumptions 1 and 4. We consider a topology on $\mathcal{F}$ defined by the following distance in $\mathcal{F}$: For $f, \tilde{f} \in \mathcal{F}$,

$$|f - \tilde{f}| = \sup_{x \in X} |f(x) - \tilde{f}(x)|.$$

Proposition 18. *Under Assumptions 1 and 4, the set $\{f \in \mathcal{F} \mid v^* \text{ is Pareto efficient in } X\}$ is open and dense in $\mathcal{F}$.*

This proposition shows that v^* is efficient only for generic f. However, if X is convex, then v^* is efficient for all f.

Proposition 19. *Suppose that X is a convex set. Under Assumptions 1 and 4, v^* is Pareto efficient in X.*

Pareto efficiency implies that players reach an agreement almost surely if t is very large. To see this, let $\pi(t)$ be the probability that players reach an agreement in equilibrium before the deadline given that no agreement has been reached until time $-t$. Then the expected continuation payoffs $v^*(t)$ must fall in the set $\{\pi(t)v \mid v \in co(\hat{X})\}$ where $co(\hat{X})$ is the convex hull of $\hat{X}$. This implies $v^*(t)/\pi(t) \in co(\hat{X})$. We have $v_i^*(t) > 0$ for all $t > 0$ and $i \in N$ since $v_i^*(t)$ is nondecreasing and $v_i^{*\prime}(0) > 0$ by equation (4.1). Since there is a positive probability that no opportunity arrives before the deadline, $\pi(t)$ is smaller than one. Therefore $v^*(t)/\pi(t) \in co(\hat{X})$ Pareto dominates $v^*(t)$. Since $co(\hat{X}) \cap A(v^*(t))$ converges to a singleton as $v^*(t)$ goes to v^* if v^* is Pareto efficient, this implies $\lim_{t\to\infty} \pi(t) = 1$. That is, we have the following proposition:

Proposition 20. *Suppose that v^* is Pareto efficient. Then the probability of agreement before the deadline converges to one as $\lambda \to \infty$.*

We note that this proposition fails if v^* is only weakly Pareto efficient. In Example 5, it is evident that players reach no agreement before the deadline with a positive probability, since the limit allocation is $(1/2, 1/2)$ while players should find a good allocations close to $(1, 1/2)$ or $(1/2, 1)$ in the limit as $T \to \infty$.

In Propositions 18 and 19, we showed that $v^*(t)$ almost always converges to the Pareto frontier of X. Now, we consider an inverse problem: For any Pareto efficient allocation w in X which is not at the edge of the Pareto frontier,[39] we show that one can find a density f that satisfies Assumptions 1 and 4 such that the limit of the solution $v^*(t)$ of equation (4.1) is w.

Proposition 21. *Suppose that $X \subseteq \mathbb{R}^n_+$ satisfies Assumption 4 (a), (b). Suppose that $w \in \mathbb{R}^n_{++}$ is a Pareto efficient allocation in X, and is not located at the edge of the Pareto frontier of X. Then there exists a probability measure μ with support X such that Assumptions 1 and 4 hold, and $\lim_{\lambda \to \infty} v^*(t) = w$ for all $t \in (0, T]$.*

In the proof, we construct a probability density function f to have a large weight near $w \in X$, and show that the limit continuation payoff profile is w if there is a sufficiently large weight near w. Note that this claim is not so obvious as it seems. Indeed, we will see in Section 4.6 that the limit is independent of density f if there is a positive discount rate $\rho > 0$, as long as Assumptions 1 and 4 hold.

4.6 The Payoffs Realizing upon Agreement

In this section, we consider the case where the payoffs realize as soon as an agreement is reached, as opposed to the assumption in the previous sections that the payoffs realize only at the deadline. We suppose that if a payoff profile $x =$

[39] We formally define this property in the proof given in the Appendix.

$(x_1, \ldots, x_n)$ is accepted by all players at time $-t \in [-T, 0]$ then player i obtains a payoff $x_i e^{-\rho(T-t)}$ where $\rho \geq 0$ is a discount rate. If no agreement has been reached until time 0, each player obtains the payoff 0.[40] First, we note that if $\rho = 0$, exactly the same analyses as in the previous sections apply. This is because with $\rho = 0$, player i's payoff when an agreement occurs at time $-t$ is $x_i e^{-\rho(T-t)} = x_i e^{-0\cdot(T-t)} = x_i$, which is independent of t. In this section, we focus on the case where $\rho > 0$. Under Assumption 1, an easy computation shows that the differential equation (4.1) is modified in the following way:

$$v'(t) = -\rho v(t) + \lambda \int_{A(t)} (x - v(t))\, d\mu \tag{4.10}$$

with an initial condition $v(0) = (0, \ldots, 0) \in \mathbb{R}^n$.

Suppose Assumptions 1 and 4 hold. Let $v^*(t; \rho, \lambda)$ be the (unique) solution of ODE (4.10).[41] If λ is large, the right hand side of equation (4.10) is approximated by the right hand side of equation (4.1) when the value of the integral is not too small. Therefore, $v^*(t; \rho, \lambda)$ is close to the solution of equation (4.1) in the case of $\rho = 0$, for λ large relative to ρ. This resemblance of trajectories holds until $v^*(t; \rho, \lambda)$ approaches the boundary of $\hat{X}$. In particular, we can show that the path of $v^*(t; \rho, \lambda)$ approaches $v^*(t; 0, \infty) = \lim_{\lambda \to \infty} v^*(t; \rho, \lambda)$ arbitrarily closely as $\lambda \to \infty$, where $v^*(t; 0, \infty)$ is the limit of the solution of equation (4.1).

[40] This is with loss of generality but setting a nonzero threat point payoff leads only to minor modifications of the statements of our results.

[41] Essential uniqueness of trembling-hand equilibrium is obtained by a proof analogous to that for Proposition 12. The unique solution of equation (4.10) gives the cutoff profile that characterizes a trembling-hand equilibrium.

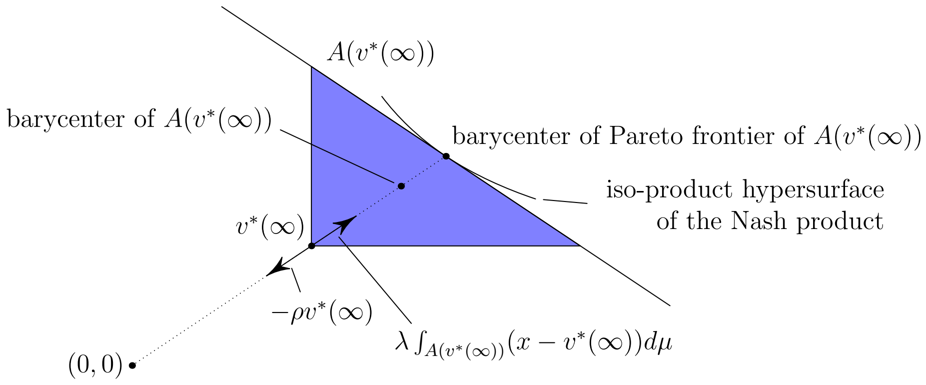

Figure 4.13: Vectors when $t \to \infty$.

Proposition 22. *For all $\varepsilon > 0$, there exists $\bar{\lambda} > 0$ such that for all $\lambda \geq \bar{\lambda}$,*

$$|v^*(t; 0, \infty) - v^*(t; \rho, \lambda)| \leq \varepsilon \quad \textit{for some } t.$$

Remark 1. Before analyzing $v^* = \lim_{\lambda \to \infty} v^*(t; \rho, \lambda)$, let us consider another limit $v^*(\infty) = \lim_{t \to \infty} v^*(t; \rho, \lambda)$. Since the right hand side of equation (4.10) is not proportional to λ, these two limits do not coincide for positive $\rho > 0$. If the limit $v^*(\infty)$ exists, this must satisfy

$$\rho v^*(\infty) = \lambda \int_{A(v^*(\infty))} \big(x - v^*(\infty)\big)\, d\mu. \tag{4.11}$$

For $\rho > 0$, equality (4.11) shows $\mu(A(v^*(\infty))) > 0$, which implies that $v^*(\infty)$ is Pareto inefficient in X. This will contrast with efficiency of $v^* = \lim_{\lambda \to \infty} v^*(t; \rho, \lambda)$ that we will show in Proposition 23.

Equality (4.11) also implies that $v^*(\infty)$ is parallel to the vector from $v^*(\infty)$ to the barycenter of $A(v^*(\infty))$, as shown in Figure 4.13 in the two-dimensional case. □

To avoid complications, we impose the following assumption in addition to Assumptions 1 and 4:

Assumption 5. (a) The weak Pareto frontier of $\hat{X}$ is smooth.

(b) Every component of the normal vector at any Pareto efficient allocation in X is strictly positive.

(c) There exists $\varepsilon > 0$ such that X contains a set $\{x \in \mathbb{R}^n_+ \mid w \geq x$, and $|w - x| \leq \varepsilon$ for some weakly Pareto efficient $w \in X\}$.

(d) The density function f is continuous.

Now suppose that λ is very large. Then $\mu(A(v^*(\infty)))$ must be very small, which means that $v^*(\infty)$ is very close to the Pareto frontier of X, where $v^*(\infty)$ is defined as in Remark 1. By Assumptions 1 and 4, the density f is approximately uniform in $A(v^*(\infty))$ if $A(v^*(\infty))$ is a set with a very small area. To obtain an intuition, suppose that $A(v^*(\infty))$ is a small n-dimensional pyramid. The vector in the right hand side of equality (4.11) is parallel to the vector from $v^*(\infty)$ to the barycenter of $A(v^*(\infty))$. We use this property to show that the boundary of $A(v^*(\infty))$ at its barycenter is tangent to the hypersurface defined by $\prod_{i \in N} x_i = a$ for some constant a. We refer to such a Pareto efficient allocation as a *Nash point*, and the set of all Nash points as the *Nash set* of $(\hat{X}, 0)$ (Maschler, Owen, and Peleg (1988), Herrero (1989)). The Nash set contains all local maximizers and all local minimizers of the Nash product. If X is convex, there exists a unique Nash point, called the Nash bargaining solution.

The above observation leads to the next proposition.

Proposition 23. *Suppose that Assumptions 1, 4, and 5 hold, and that any Nash point is isolated in X. Then the limit* $v^* = \lim_{\lambda \to \infty} v^*(t; \rho, \lambda)$ *exists and belongs to the Nash* *set of the problem* $(\hat{X}, 0)$ *for all* $t > 0$*. If* X *is convex, this limit coincides with the Nash bargaining solution of* $(X, 0)$*.*

Therefore, the trajectory of $v^*(t)$ for very large λ starts at $v^*(t) = 0$, approaches $v^*(t; 0, \infty)$, and moves along the Pareto frontier until reaching a point close to a Nash point.

Finally we consider the duration of search in the equilibrium. In contrast to Theorems 5 and 6 in the case when payoffs realize at the deadline (or $\rho = 0$), we show that an agreement is reached almost immediately if λ is very large.

Proposition 24. *Suppose that Assumptions 1, 4, and 5 hold. If* $\rho > 0$*, then* $D(\infty) = 0$*.*

4.7 Discussions

4.7.1 Non-Poisson Arrival Processes

In the main sections we considered Poisson processes to make the presentation of the results easier. The Poisson process assumes that the probability of opportunity arrival is zero at any moment, so in particular the probability of receiving one more opportunity shrinks continuously to zero as the deadline approaches. However in some circumstances it would be more realistic to assume there is a well-defined "final period" that can be reached with a positive probability. In this section we generalize our model to encompass such cases and show that our results are unaffected.

Specifically, consider dividing the time horizon of length T into small subintervals each with length Δt (so there are $\frac{T}{\Delta t}$ periods in total). At the end of each subinterval, players obtain an opportunity with probability $\pi(\Delta t)$. Notice that the Poisson process corresponds to the case when $\pi(\Delta t) = \lambda \Delta t$ for some $\lambda > 0$

and we let $\Delta t \to 0$. Here we allow for general π function, such as $\pi(\Delta t) = a$ or $\pi(\Delta t) = a\sqrt{\Delta t}$ for some constant $a > 0$. Let $v_i(n)$ be the continuation payoff at time $n\Delta t$. Then,

$$\begin{aligned} v_i\Big(\frac{t}{\Delta t}+1\Big) &= \big(1-\pi(\Delta t)\big)v_i\Big(\frac{t}{\Delta t}\Big) + \pi(\Delta t)\left(\int_{X\setminus A(v(\frac{t}{\Delta t}))} v_i\Big(\frac{t}{\Delta t}\Big)\,d\mu + \int_{A(v(\frac{t}{\Delta t}))} x_i\,d\mu\right) \\ &= v_i\Big(\frac{t}{\Delta t}\Big) + \pi(\Delta t)\int_{A(v(\frac{t}{\Delta t}))}\Big(x_i - v_i\Big(\frac{t}{\Delta t}\Big)\Big)\,d\mu. \end{aligned}$$

Hence,

$$v_i\Big(\frac{t}{\Delta t}+1\Big) - v_i\Big(\frac{t}{\Delta t}\Big) = \pi(\Delta t)\int_{A(v(\frac{t}{\Delta t}))}\Big(x_i - v_i\Big(\frac{t}{\Delta t}\Big)\Big)\,d\mu. \tag{4.12}$$

Notice that if we set $\pi(\Delta t) = \lambda\Delta t$ and take the limit as $\Delta t \to 0$, the left hand side divided by Δt converges to $v_i'(t)$ in the Poisson model and the right hand side divided by Δt converges to $\lambda \int_{A(v(t))}(x_i - v_i(t))d\mu$, consistent with equation (4.1).

Proposition 25. *If* $\lim_{\Delta t\to 0} \frac{\pi(\Delta t)}{\Delta t} = \infty$*, under Assumptions 1 and 3, the limit expected duration is* $\frac{n^2}{n^2+n+1}T$.

Note that this result is consistent with Proposition 16 where we consider the Poisson process and take a limit of $\lambda \to \infty$. Thus our limit result is robust to the move structure.

4.7.2 *Relative Importance of Discounting and Search Friction*

In the main sections, we have shown that if $\rho = 0$, the limit expected duration as $\lambda \to \infty$ is positive under certain assumptions, and the limit equilibrium payoff profiles are efficient but depend on the distribution μ. In Section 4.6, in contrast, the limit duration is zero, and the limit payoffs are the Nash bargaining solution if $\rho > 0$ is fixed. In this section, we show that the limit duration and the limit equilibrium payoffs as $\lambda \to \infty$ and $\rho \to 0$ simultaneously depend on the limit of $\lambda\rho^n$.

Proposition 26. *Suppose that Assumptions 1, 3, and 5 hold. The limit expected duration $D(\infty)$ and the limit allocation $v^* = \lim_{\lambda\to\infty,\rho\to 0} v^*(t;\rho,\lambda)$ satisfy the following claims: (i) If $\lambda\rho^n \to 0$, then $D(\infty) > 0$, and $v^* = \lim_{\lambda\to\infty} v^*(t;0,\lambda)$, which is the limit analyzed in Sections 4.4 and 4.5. (ii) If $\lambda\rho^n \to \infty$, then $D(\infty) = 0$, and $v^* = \lim_{\lambda\to\infty} v^*(t;\rho,\lambda)$ for $\rho > 0$, which is the limit shown in Section 4.6.*

An insight behind this result is as follows: The limit of the expected payoffs depend on whether the first term in ODE (4.10) is negligible or not when compared to the second term. Let $z(t;\rho,\lambda)$ be the Hausdorff distance from $v^*(t;\rho,\lambda)$ to the Pareto frontier of X. If ρ is very small and λ is not very large, Proposition 22 shows that $v^*(t;\rho,\lambda)$ is close to $\lim_{\lambda\to\infty} v^*(t;0,\lambda)$ which is on the Pareto frontier. Then we can apply an analogous argument to the one provided in the discussion in the proof sketch of Theorem 8 to show that $z(t;\rho,\lambda)$ is approximately proportional to $\lambda^{-1/n}$. Since $\mu(A(t))$ approximates $z(t;\rho,\lambda)^n$ (times some constant), and the length of the vector from $v^*(t;\rho,\lambda)$ to the barycenter of $A(t)$ is linear in $z(t;\rho,\lambda)$, the second term is of order $\lambda \cdot \lambda^{-1/n} \cdot \lambda^{-1} = \lambda^{-1/n}$. Therefore if $\lambda\rho^n \to 0$ the first term, which approximates ρv^*, is negligible because ρ vanishes more rapidly than $\lambda^{-1/n}$. Thus the limits in this case are the same as in Sections 4.4 and 4.5. If $\lambda\rho^n \to \infty$, the first term is significant because ρ does not vanish rapidly compared to $\lambda^{-1/n}$. This corresponds to Section 4.6. An analogous argument can be made for the limit of durations.

4.7.3 Infinite-Horizon and Static Games

Although we consider a finite-horizon model, our convergence result in Proposition 23 is suggestive of that in infinite-horizon models such as Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2011), all of whom consider the limit of stationary equilibrium outcomes as the discount factor goes to one in discrete-time infinite-horizon models. This is because the threatening power of disagreement at the deadline is quite weak if the horizon is very far away, and thus the infinite-horizon model is similar to a finite-horizon model with $T \to \infty$ if $\rho > 0$. In fact, we can show that the iterated limit as $T \to \infty$ and then $\rho \to 0$ is the Nash bargaining solution in our model if X is convex. By Proposition 26, $\lim_{\rho \to 0} v^*(T; \alpha\rho, \alpha\lambda)$ with $\alpha = \rho^{-a}$ is the Nash bargaining solution for all $a > n/(n+1)$. As $a \to \infty$, we see that the iterated limit $\lim_{\rho \to 0} \lim_{\alpha \to \infty} v^*(T; \alpha\rho, \alpha\lambda)$ is also the Nash bargaining solution. Since enlarging T is equivalent to raising both λ and ρ in the same ratio by the form of ODE (4.10), the iterated limit as $T \to \infty$ first and then $\rho \to 0$ must be the Nash bargaining solution. For the same reason, the expected duration in the limit as the discount factor goes to one in the infinite-horizon model is zero, being analogous to our Proposition 24 in which we send λ to ∞ while $\rho > 0$ is fixed. Therefore we obtained the following proposition:

Proposition 27. *In the infinite-horizon search model, the expected duration in a stationary equilibrium converges to zero as the discount factor goes to one.*

Propositions 12 and 13 imply that the limit continuation payoff of a player is essentially equal to the cutoff, which is expressed by a single variable. In this sense, there is some connection between our model and a static game considered by Nash (1953) himself, who provided a characterization of the Nash bargaining solution by introducing a static demand game with perturbation described as follows.[42] Suppose that X is convex. The basic demand game is a one-shot strategic-form game in which each player i calls a demand $x_i \in \mathbb{R}_+$. Players obtain $x = (x_1, \dots, x_n)$ if $x \in X$, or 0 otherwise. In the perturbed demand game, players fail to obtain $x \in X$ with a positive probability if x is close to the Pareto frontier. Under certain conditions, he showed that the Nash equilibrium of the perturbed demand game converges to the Nash bargaining solution as the perturbation vanishes.

Let us compare the perturbed Nash demand game with our multi-agent search model with a positive discount rate. Let $p(x) = \mu\big(A(x)\big)$ be the probability that players come across an allocation which Pareto dominates or equals $x \in X$ at an opportunity. If T is very large and t is close to T, players at time $-t$ choose almost the same cutoff profile, say x, contained in the interior of X. The average duration that players wait for an allocation falling into $A(x)$ is almost $1/\lambda p(x)$. During this time interval, payoffs are discounted at rate ρ. Since x_i must be equal to her continuation payoff in an equilibrium, i would lose nearly $(1 - e^{-\rho/\lambda p(x)})y_i$ on average by insisting on cutoff x_i where y is the expected allocation conditional on $y \in A(x)$. Note that this loss vanishes as $\rho \to 0$ for every x in the interior of X. Let probability $P(y)$ satisfy $P(y) = e^{-\rho/\lambda p(x)}$. Player i loses the same expected payoff when $y \in X$ is demanded in the perturbed demand game where the probability of successful agreement is $P(y)$.

The key tradeoff in this game, the attraction to larger demands or the fear of failure of agreement, is parallel to that in the multi-agent search, to be pickier or to avoid loss from discounting.

---

[42] We here follow a slightly modified game considered by Osborne and Rubinstein (1990, Section 4.3). Despite the difference, the model conveys the same insight as the original.

4.7.4 Time Costs

In the model of the main sections, whether or not players discount the future does not affect the outcome of the game, as payoffs are received at the deadline. However, there may still be a time cost associated with search. In this subsection we analyze a model with time costs, and show numerically that the search durations with "reasonable parameter values" are close to the limit duration with zero time cost that we solved for in the main sections.

Consider a model in which each player incurs a flow cost $c > 0$ until the search ends. In this model, it is straightforward to see that the differential equation (4.1) is modified in the following way:

$$v_i'(t) = -c + \lambda \int_{A(t)} \big(x_i - v_i(t)\big)\, d\mu \tag{4.13}$$

for each $i \in N$, with an initial condition $v(0) = (0, \ldots, 0) \in \mathbb{R}^n$.

The analysis of this differential equation is similar to the one in Section 4.6, with an exception that under Assumptions 1, 3, and 5, the limit expected payoff profile as $\lambda \to \infty$ for a fixed cost $c > 0$ is now a point that maximizes the sum of the payoffs, denoted v^S. Let $v^*(t; c, \lambda)$ be the expected payoff at time $-t$ when parameters c and λ are given. A proof similar to the one for Proposition 26 shows the following:

Proposition 28. *Suppose that Assumptions 1, 3, and 5 hold. The limit expected duration* $D(\infty)$ *and the limit allocation* $v^* = \lim_{\lambda\to\infty, c\to 0} v^*(t; c, \lambda)$ *satisfy the following claims: (i) If* $\lambda c^n \to 0$*, then* $D(\infty) > 0$*, and* $v^* = \lim_{\lambda\to\infty} v^*(t; 0, \lambda)$*, which is the limit analyzed in Sections 4.4 and 4.5. (ii) If* $\lambda c^n \to \infty$*, then* $D(\infty) = 0$*, and* $v^* = v^S$ *for* $c > 0$*.*

The proposition suggests that for a high arrival rate λ, the expected duration

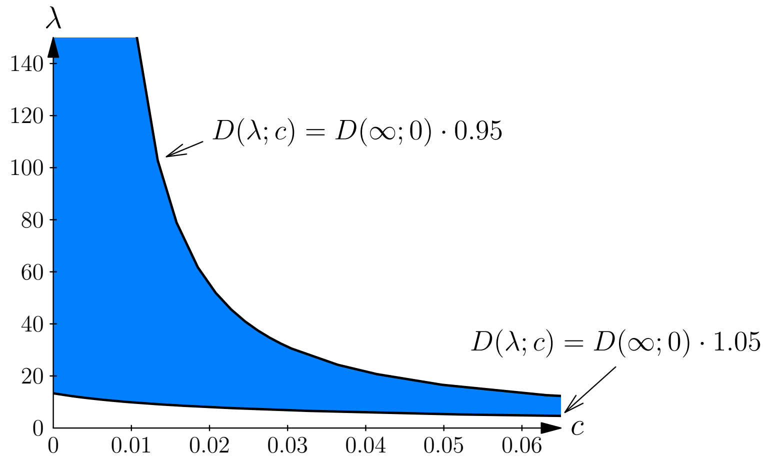

Figure 4.14: Time costs and arrival rates. The shaded region describes the set of pairs (c, λ) with which the expected search duration is within 5% difference from the limit duration.

does not change so much when we increase the cost from zero to a small but positive number. Combined with our argument in Step 3, this suggests that whenever the cost is sufficiently small, our limit arguments in Steps 1 and 2 are economically meaningful. Now we numerically show that the degree to which the cost should be small is not too extreme. Specifically, we consider the case when $n = 2$ and μ is a uniform distribution over $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$, and solve for the range of pairs of costs and arrival rates such that the expected search duration is within 5% difference from the limit duration. As shown in Figure 4.14, such a range contains a wide variety of pairs of parameter values (note that the expected limit payoff is 0.5 in this game, so the cost of 0.05 corresponds to the setting with a fairly high cost). When $n = 2$ and μ is an independent distribution such that each player's marginal is an exponential distribution, whenever the cost c is less than 10% of the expected payoff given with $c = 0$ and $\lambda = 100$, we find that λ for which the expected duration is of 95% of the limit duration is more than 100, and that of 105% is less than

10.[43] These results suggest that the limit argument that we conducted in Steps 1 and 2 of the main sections is economically reasonable.[44]

4.7.5 *Counterexamples of Positive Duration of Search*

In Theorems 5 and 6, we showed that the limit expected duration of search is positive if certain assumptions hold. In this section, we present examples of distributions under which assumptions are not satisfied and the expected duration may not be positive as $\lambda \to \infty$.

First, note that it is straightforward to see that if μ assigns a point mass to a point that Pareto-dominates all other points in the support of μ, the limit expected duration is zero. Less obvious is the situation where μ allows for point masses while no point Pareto-dominates all the other points. Even in this case, Kamada and Sugaya (2010)'s "three-state example" shows that trembling-hand equilibrium may not be unique, and the duration can be zero in a trembling-hand equilibrium.[45] Our Assumption 1 (b) requires a more stringent condition that the marginal distribution must have a locally bounded density function, and thus does not have a point mass. Here we present an example in which μ does not have a point mass while its marginal has a point mass, and there are multiple equilibria and some of them have zero limit duration.

Example 6. Consider $X = co\{(2,1),(1,1)\} \cup co\{(1,2),(1,1)\}$ and let μ be a uniform distribution over this X. First, consider a strategy profile in which agents accept

[43] In this case, we can show that the expected duration is positive even in the limit as $\lambda \to \infty$, exhibiting a stark contrast to Proposition 24. See Kamada and Muto (2011a) for details.

[44] We are planning to extend this argument to more cases beyond the setting provided here.

[45] Consider μ that assigns equal probabilities to $(2,1)$ and $(1,2)$.

a payoff strictly above 1 until time $-t^*$ and accept all offers after $-t^*$, where t^* satisfies the indifference condition at $-t^*$:

$$1 = \frac{1 - e^{-\lambda t^*}}{2}(1 + 1.5) + e^{-\lambda t^*} \cdot 0,$$

or $t^* = \frac{1}{\lambda}\ln(5)$. Since given this strategy profile the continuation payoff for both players is 1 if $-t \leq -t^*$ and it is strictly less than 1 otherwise, this indeed constitutes a trembling-hand equilibrium.

However, there exist other equilibria. For example, consider a strategy profile which is exactly the same as the above one except that both agents accept the first offer regardless of its realization. Since the continuation payoff at the time of the first arrival is 1 for both players as we have argued, this also constitutes an trembling-hand equilibrium. Thus there are multiple trembling-hand equilibria. Also the limit expected duration under the second equilibrium is trivially zero, suggesting the need for Assumption 1 (b) for Theorem 6 to hold.

The key to multiplicity and zero duration is the fact that payoff profiles at which players are indifferent arrive with positive probability due to the atom on marginals. Assumption 1 rules out such a situation. □

Next, we show that even if μ has no point mass, the limit expected duration may be zero when μ does not satisfy Assumption 2 (nor Assumption 6 in Appendix C.2.4).

Example 7. For $n = 1$, let F be a cumulative distribution function defined by $F(x) = 1 + \dfrac{1}{\ln(1-x)}$ for $x \in [1 - e^{-1}, 1)$, and $F(1) = 1$. The density is $f(x) =$

$\frac{1}{(1-x)(\ln(1-x))^2}$. Recalling formula (4.7), the density term is

$$d(v) = \frac{f(v)}{1-F(v)} = -\frac{1}{(1-v)\ln(1-v)},$$

and the barycenter term $b_1(v)$ is clearly smaller than $1-v$. Since $\lim_{\lambda\to\infty} v^*(t) = 1$,

$$\begin{aligned} r &= \lim_{v\to 1} d_1(v) b_1(v) \\ &\leq \lim_{v\to 1} \frac{-1}{\ln(1-v)} = 0. \end{aligned}$$

By Theorem 8, the limit duration is zero.

In this example, it is easy to show that for all $\alpha > 0$, there exists $\varepsilon > 0$ such that $1-(1-x)^\alpha \geq F(x)$ for all $x \in [1-\varepsilon, 1]$. Distribution F is very close to a discrete distribution, in that $F(x)$ converges to 1 as $x \to 1$ at a speed slower than any polynomial functions. In such a case, the above computation shows that the limit duration can be zero, which is the same as the case with discrete distributions.

□

4.7.6 *The Effect of a Slight Change in the Distribution*

The limit result in Proposition 16 depends crucially on the assumption of smooth boundary and continuous positive density. Although this is the assumption that is often invoked in the literature, it is desirable to know how robust this result is. To this end, consider $X = \mathbb{R}^n_+$ and a distribution over X, μ, which may or may not be full-support. Introduce a notion of distance between two distributions, $d(\mu, \gamma) = \sup_{A\subseteq X} |\mu(A) - \gamma(A)|$.

A standard argument on ordinary differential equations shows the following:

Proposition 29. *For any λ, the limit duration is continuous in distribution almost everywhere.*

That is, for any finite arrival rate, the limit duration is not substantially affected by a slight change in distribution. Combined with the result that our limit result approximates the situation with a finite but high arrival rate, this suggests that our limit duration is relevant even for the distributions that are not very different from a distribution that satisfies our assumptions (Assumptions 1 and 4).

4.7.7 Time Varying Distributions

In the main model we considered the case in which the distribution μ is time-independent. This benchmark analysis is useful in understanding the basic incentive problems that agents face, but in some situations it might be more realistic that the distribution changes over time. In this section, we examine whether the positive duration result in Theorem 5 (the case with a single agent) is robust to this independence assumption. An analogous argument can be made for the multiple-agent case. Let F_t be the cumulative distribution function of the payoff at time $-t$.

First, consider the case in which the distribution becomes better over time in the sense of first order stochastic dominance. In this case, it is easy to see that the expected duration is still positive and it becomes longer at least in certain cases: For each t, consider the cutoff at each time $-s \in (-t, 0]$ that equates the acceptance probability with the one that the agent would get at $-s$ if the distribution in the future were fixed at F_t. This gives a higher continuation payoff at $-t$ as the distribution becomes better over time. Thus the cutoff at $-t$ must be greater than the continuation payoff at $-t$ that the agent would obtain by fixing the distribution at F_t ever after. This means that at any $-t$, the acceptance probability is smaller than the one obtained by fixing the distribution at F_t ever after. Hence the acceptance probability at $-t$ is $O(\frac{1}{\lambda t})$, so we have a positive duration. If $(F_t)_{t\in[0,T]}$ is such that the acceptance probability at time $-s$, $p(s)$, when the payoff is drawn by the fixed F_t independently over time does not depend on t, then the above argument also implies that the duration becomes longer.

Now consider the case when the distribution may become worse off. First, if the support of the distribution becomes worse off, then there is no guarantee of positive duration. For example, if the upper bound of the support decreases exponentially then the analysis of the duration becomes equivalent to that for the case with discounting, in which Proposition 24 has already shown that the limit expected duration is zero.

If the support does not change, then the positive duration result holds quite generally: In the proof of Theorem 5 provided in the Appendix, we did not use the fact that F does not depend on t. The following modification of Assumption 2 guarantees the positive duration.

Assumption 2$''$. *There exists a concave function φ such that $1-\varphi(x)$ is of the same order as $1-F_t(x)$ in $\{x \in \mathbb{R} \mid F_t(x) < 1\}$ for all t.*

Notice that we require the existence of $\bar{x}$ and φ that are applicable to all F_t.

Proposition 30. *Suppose $n = 1$. Under Assumptions 1 and 2$''$, $\liminf_{\lambda\to\infty} D(\lambda) > 0$.*

4.7.8 Dynamics of the Bargaining Powers

Consider the case where $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$ and a density f such that $f(x) > f(x')$ if $x_2 - x_1 > x'_2 - x'_1$. Suppose that the payoff realizes upon agreement as in Section 4.6, and the discount rate $\rho > 0$ is very small. In this case, the

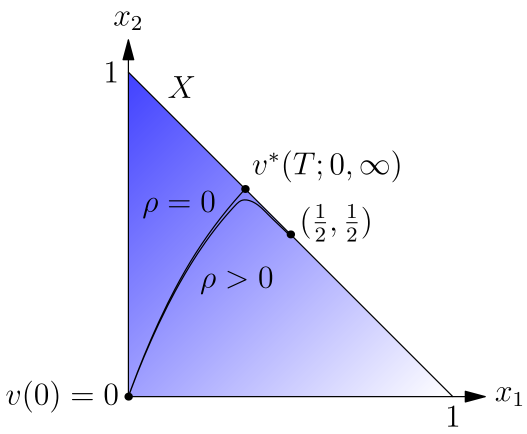

Figure 4.15: Paths of continuation payoffs. The probability density is low near $(1,0)$, and high near $(0,1)$.

limit of the solution of ODE (4.1) with $\rho = 0$, denoted $v^*(T;0,\infty)$, locates at the boundary of X by Proposition 19, and it is to the north-west of $(\frac{1}{2},\frac{1}{2})$, which is the Nash bargaining solution and is the limit of the solution of ODE (4.10). Hence, by Proposition 22, the continuation payoff when the players receive payoffs upon the agreement starts at a point close to $(\frac{1}{2},\frac{1}{2})$, and goes up along the boundary of X and reaches a point close to $v^*(T;0,\infty)$, and then goes down to $(0,0)$. On this path of play, player 1's expected payoff is monotonically decreasing over time. On the other hand, player 2's expected payoff changes non-monotonically. Specifically, it rises up until it reaches close to $v_2^*(T;0,\infty)$, and then decreases over time. Figure 4.15 illustrates this path.

Underlying this non-monotonicity is the change in the bargaining powers between the players. When the deadline is far away, there will be a lot of opportunities left until the deadline, so it is unlikely that players will accept allocations that are far from the Pareto efficient allocations, so the probability distribution over such allocations matters less. Since X is convex and symmetric, two players expect roughly the same payoffs. However, as the time passes, the deadline comes closer, so players expect more possibility that Pareto-inefficient allocations will be accepted. Since player 2 expects more realizations favorable to her than player 1 does, player 2's expected payoff rises while player 1's goes down. Finally, as the deadline comes even closer, player 2 starts fearing the possibility of reaching no agreement, so she becomes less pickier and the cutoff goes down accordingly.

4.7.9 Market Designer's Problem

In this section we consider problems faced by a market designer who has a control over some parameters of the model.

First, consider the case when the payoffs realize at the deadline, and the designer can tune the horizon length T. In this case there is no point in making the horizon shorter, as the continuation payoff $v(t)$ is increasing in t.

Second, still in the case with payoffs realizing at the deadline, suppose that the designer can instead affect the probability distribution over potential payoff profiles, by "holding off" some offers. Formally, given μ, let the designer choose a distribution μ' such that $\mu'(C) \leq \mu(C)$ for all $C \subseteq X$.[46] In this case the designer faces a tradeoff: On one hand, tuning the distribution can affect the path of continuation payoffs and the ex ante expected payoff at time $-T$ (an analogous argument to Proposition 21). On the other hand, however, changing the distribution will decrease the expected number of offer arrivals in the finite horizon, so $v(T)$ is lower than the case when the distribution is instead given by μ'' such that $\mu''(C) = \frac{\mu'(C)}{\mu(X')}$ for all $C \subseteq X$. The explicit form of an optimal design would depend on the specificalities of the problem at hand and the objective function of the designer, but basi-

[46] Note that μ' may not be a probability measure because it might be the case that $\mu'(X) < 1$.

cally if the horizon length T is high then reducing probabilities would not lead to too much loss.

Next, consider the case with payoffs realizing upon agreement. In this case there can be a benefit from reducing T. As in the case with payoffs realizing at the deadline, lower T means that the expected payoff at time $-T$ is less close to the Pareto boundary. However, if the solution when payoffs realizing at the deadline, $v^*(T;0,\infty)$, is socially desirable than the Nash bargaining solution, then by reducing T appropriately the expected payoff profile will come closer to $v^*(T;0,\infty)$ (provided that the expected payoffs are in between these two payoffs before shortening T; remember that by Proposition 22 the expected payoffs for intermediate values of time t is close to $v^*(T;0,\infty)$).

On the other hand, tuning the distribution has a smaller effect than the case with payoffs realizing at the deadline, as we know that the payoffs eventually converge to the Nash bargaining solution. However, since $v^*(T;0,\infty)$ depends on the distribution, Proposition 22 implies that the direction from which the payoff converges varies as the designer varies the distribution.

4.7.10 Majority Rule

In the main sections we considered the case when players use *unanimous rule* for their decision making. This is a reasonable assumption in many applications such as the apartment search, but there are certain other applications in which *majority rule* fits the reality better. This section is devoted to the analysis of such a case.

Precisely speaking, by majority rule we mean the decision rule such that the object of search is accepted if and only if $k < n$ players say "accept" upon its arrival.

First of all, it is straightforward to check that Propositions 12 and 13 (the trembling-hand equilibrium is essentially unique and players use cutoff strategies) carry over to this case. If X is convex, and satisfies Assumption 4, the limit expected payoff cannot be weakly Pareto efficient. To see this, suppose that the limit payoff is weakly Pareto efficient. Then, as shown in Figure 4.16, there is a region with a positive measure such that the acceptance takes place. However the barycenter of these regions is in the interior of X by convexity, and hence the limit payoff profile must be an interior point as well. This contradicts the assumption that the limit payoff profile is weakly Pareto efficient. Now, let the true limit point be $\tilde{v}$. Since any payoffs that strictly Pareto dominate $\tilde{v}$ must be accepted by all players, and the measure of this region is strictly positive, the limit duration must be zero.[47] We summarize our finding as follows:

Proposition 31. *Under the majority rule with $k < n$, if X is convex and satisfies Assumption 4, then the limit expected payoff profile is not weakly efficient, and the limit expected duration of search is zero.*

We note that the conclusion of this proposition still holds even if the smoothness assumption in Assumption 4 is replaced by the following assumption, which essentially says that the Pareto frontier is downward-sloping with respect to other players' payoffs: For all $i \in N$, $\overline{x}_i(x_{-i}) = \sup\{x_i \mid (x_i, x_{-i}) \in X\}$ is decreasing in x_j for all $j \neq i$.

4.7.11 Negotiation

Our model assumes that players cannot transfer utility after agreeing on an allocation. We believe our model keeps the deviation from the standard single-

[47] This discussion is parallel to Compte and Jehiel (2010, Proposition 7) who consider majority rules in a discrete-time infinite-horizon search model.

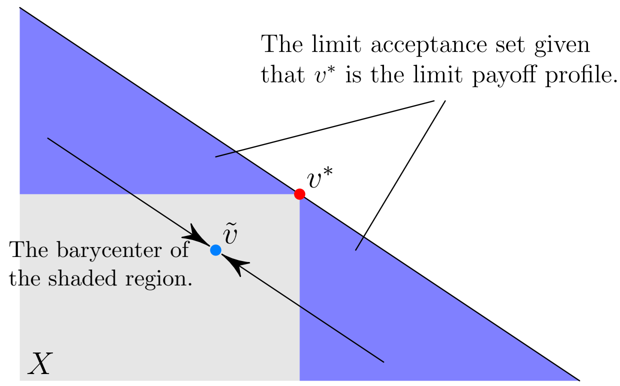

Figure 4.16: Equilibrium continuation payoffs under a majority rule: If the point v^* were the limit payoff profile, then all payoffs in the shaded region are in the acceptance set, and the barycenter of the shade region, $\tilde{v}$, is in the interior of X as X is convex.

agent infinite-horizon search model minimal so that the analysis purifies the effect of modifying the number of agents and the length of the horizon. Also, our primary interest is in the case where such negotiation is impossible or the case where the stake of the object is very high so even if players could negotiate, the impact on the outcome is negligible. However, in some cases negotiation may not be negligible. Here we discuss such cases. We will show that the duration continuously changes with respect to the degree of impact of negotiation, hence our results are robust with respect to the introduction of negotiation. Our extension also lets us obtain intuitive comparative statics results.

Suppose that players can negotiate after they observe a payoff profile $x \in X$ at each opportunity at time $-t$. Players can shift their payoff profile by making a transfer, and may agree with the resulting allocation. We assume that the allocation they agree with is the Nash bargaining solution where a disagreement point is the continuation payoff profile at the time $-t$ in the equilibrium defined for this modified game.[48] When making a transfer, we suppose that a linear cost is incurred: If player i gives player j a transfer z, j obtains only az for $a \in [0,1)$. This cost may be interpreted as a misspecification of resource allocation among agents, or a proportional tax assessed on the monetary transfer. Note that a measures the degree of impact of negotiation. Our model in the main sections corresponds to the case of $a = 0$.

To simplify our argument we restrict attention to a specific model with two players.[49] Specifically, we consider the case with costly transferable utility: Suppose that $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$ with the uniform distribution μ on X. For each arrival of payoff profile x, players can negotiate among the set of feasible allocations defined by

$$S(x) = \left\{ x' \in X \;\middle|\; \begin{array}{ll} a(x_1 - x_1') \geq x_2' - x_2 & \text{if } x_1 \geq x_1', \\ x_1' - x_1 \geq a(x_2 - x_2') & \text{if } x_1 < x_1' \end{array} \right\}.$$

We suppose that each player says either "accept" or "reject" to the Nash bargaining solution obtained from the feasible payoff set $S(x)$ and the disagreement point given by the continuation payoff profile $v(t)$.

By looking at geometric properties of the Nash bargaining solution, we can compute the limit expected duration in this environment.

Proposition 32. *Under Assumptions 1, 4, and 7, the limit expected duration in the game with negotiation is* $D(\infty) = \dfrac{4 + 4a^2}{7 + 6a + 7a^2}$ *for* $a \in [0,1)$.

---

[48] This use of Nash bargaining solution is not critical to our result. Similar implications are obtained from other bargaining solutions such as the one given by take-it-or-leave-it offers by a randomly selected player.

[49] We expect that nothing substantial would change even if we extended the argument to the cases of three or more players.

Since $D(\infty)$ is decreasing in a, the limit expected duration becomes shorter in the presence of negotiation. This is intuitive, as negotiation essentially precludes extreme heterogeneity in the offer realization, thus the agreement can be reached soon. Notice also that the proposition claims that the duration must be strictly positive even with negotiation, and $D(\infty)$ converges to $4/7$ as $a \to 0$, which is the same duration as we claimed in Proposition 16. That is, our main result is robust to the introduction of negotiation.

Note that the proposition does not apply to the case in which utilities are perfectly transferable, i.e., $a = 1$. However, this is due to the fact that the analysis above is in a knife-edge case because the Pareto frontier consists of a straight line: If X is strictly convex, then even if $a = 1$, the acceptance set shrinks with a faster speed than the case that we analyzed above, and the resulting duration is longer in such case.[50]

4.8 Conclusion

This paper analyzed a modification of the standard search problem by introducing multiplicity of players and a finite horizon. Together, these extensions significantly complicate the usual analysis. Our main results identified the reasons behind the widely-observed phenomenon that such searches often take a long (or at least some) time. We first showed that the search duration in the limit as the search friction vanishes is still positive, hence the mere existence of some search friction has a nonvanishing impact on the search duration. This limit duration is increasing in the number of players as a result of two effects: the ascending ac-

[50] This suggests that with negotiation, preference heterogeneity may help shortening the search duration. We plan to explore this issue in our continuation work.

ceptability effect and the preference heterogeneity effect. In short, the ascending acceptability effect states that a player has an extra incentive to wait as the opponents accept more offers in the future, and the preference heterogeneity effect states that such "extra offers" include increasingly favorable offers for a player due to heterogeneity of preferences. Then we showed that the convergence speed of the duration as the friction vanishes is high, and numerically demonstrated that durations with positive frictions are reasonably close to the limit duration in our examples. This provides evidence that our limit analysis contains economically-meaningful content, and the mere existence of some friction is actually the main driving force of the positive duration in reality—so the effects that we identify in Steps 1 and 2 are the keys to understand the positive duration in reality.

We also conducted a welfare analysis, and showed that the limit expected payoff is generically Pareto efficient, and depends on the distribution of offers. Lastly, we provided a wealth of discussions to examine the robustness of our main conclusions and to analyze a variety of alternative specifications of the model.

Our paper raises many interesting questions for future research. First, it would be interesting to consider the case where agents can search for another offer even after they agree on an offer (i.e., search with recall). In this case the search duration in equilibrium must be always T, but the duration until the first agreement is not obvious, because players' preferences are heterogeneous: Player 1 may not want to agree on the offer that gives player 2 a high payoff, expecting 2's future reluctance to accept further offers. In our continuation work Kamada and Muto (2011c) we analyze this case and find that under some assumptions, the duration until the first acceptance is positive even in the limit as the friction vanishes. In that paper we also find that players may no longer use cutoff strategies, and as a result the shape of the acceptance set is quite complicated.

Second, it would be interesting to consider a large market model where at each period a fixed number of agents from a large population are matched and some payoff profile is realized. If all agents agree on the profile, they leave the market. There are at least two possible specifications for such a model. First, we can consider the situation where an overlapping generation model with agents facing different deadlines, and there is a constant inflow of agents. In our ongoing research, we solve for a steady state equilibrium strategy and characterize the expected search duration of each agent in the population under certain regularity assumptions. On the other hand, if all agents share the same deadline, the arrival rate must decrease or the set of feasible payoffs must shrink over time to reflect the change in the measure of agents who remain in the market, and it is not obvious whether the positive duration results carry over. Our result on time-varying distributions in Section 4.7.7 may be useful in such an analysis.

Finally, in order to isolate the effects of multiple agents and a finite horizon as cleanly as possible, we attempted to minimize the departure from the standard model. Inevitably, this entailed ruling out some properties that would be relevant in particular applications. For example, in some cases there may be uncertainty (perhaps resolving over time) about the distribution over outcomes or the opponents' preferences. We conjecture such uncertainty would increase the duration. Another example would be the possibility of agents using effort to increase the arrival rate or perhaps sacrificing a monetary cost to postpone the deadline. Again this would increase the search duration, as players could make these decisions conditional on the time left to the deadline. These extensions of our model are left for future work.

APPENDIX

A. APPENDIX TO CHAPTER 2

A.1 Proof of Proposition 1

We provide the proof of Proposition 1 (the existence and differentiability of the optimal path). First, we present a simple but useful lemma. Recall that we are assuming that the optimal action a^* is different form the Nash action a^N (A1). Then consider

$$A^* := \begin{cases} [a^N, a^*] \text{ if } a^N < a^* \\ \\ [a^*, a^N] \text{ if } a^* < a^N \end{cases}$$

The next lemma shows that we can restrict our attention to the trigger strategy equilibria whose action always lies in A^*.

Lemma 2. *For any trigger strategy equilibrium $x \in X^*$, there is a trigger strategy equilibrium $\widehat{x} \in X^*$ such that $\forall t\ \ \widehat{x}(t) \in A^*$ and $\pi(\widehat{x}(t)) \geq \pi(x(t))$.*

Proof. We show this for the case of $A^* = [a^N, a^*]$. By assumptions A2 and A5, the graph of π is continuous and "single peaked", and therefore if $\pi(a^N) < \pi(x(t))$ and $x(t) \notin A^*$ then there must be $\widehat{x}(t) \in A^*$ such that $\pi(\widehat{x}(t)) = \pi(x(t))$ and $\widehat{x}(t) < x(t)$ (see Figure A.1).

Replace such $x(t)$ by $\widehat{x}(t) \in A^*$ defined above. If $\pi(a^N) \geq \pi(x(t))$, replace $x(t)$ by $\widehat{x}(t) \equiv a^N$. If $x(t) \in A^*$, let $\widehat{x}(t) = x(t)$. Note that $\pi(\widehat{x}(t)) = \max\left\{\pi(x(t)), \pi^N\right\}$ and this is measurable (so that $\widehat{x}$ is feasible). Lastly, we show that $\widehat{x}$ satisfies the

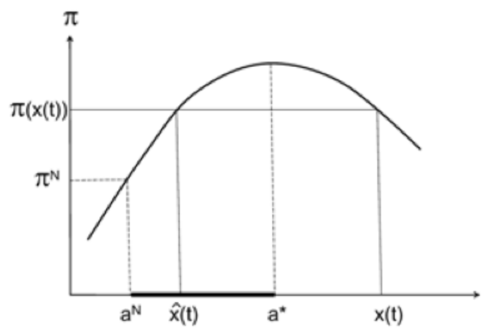

Figure A.1: The graph of $\pi(\cdot)$.

incentive constraint (2.3). Since $\pi(\widehat{x}(t)) \geq \pi(x(t))$, the right hand side of (2.3) is weakly larger under $\widehat{x}$ for all t. Hence we only need to show $d(\widehat{x}(t)) \leq d(x(t))$ for all t. This is trivially true when $\widehat{x}(t) = a^N$. Otherwise, we have $a^N < \widehat{x}(t) \leq x(t)$. Since $d(a)$ is increasing for $a > a^N$ (by A6), we have $d(\widehat{x}(t)) \leq d(x(t))$. □

This Lemma shows that the optimal trigger strategy (if any) can be found in the set X^{**} of trigger strategy equilibria whose range is $A^*(= [a^N, a^*]$ or $[a^*, a^N])$:

$$X^{**} := \{x \in X^* | \forall t \; x(t) \in A^*\}.$$

Proposition 33. *There is an optimal trigger strategy equilibrium $\overline{x}(t)$ (i.e., $\overline{x} \in X^*$ and $V(\overline{x}) = \max_{x \in X^*} V(x)$, where V denotes the expected payoff associated with x) which is continuous in t.*

Proof. We show that there is a trigger strategy equilibrium in X^{**} that attains max $_{x \in X^{**}} V(x)$ (by Lemma 2, it is the true optimal in X^*). We consider the case $a^N < a^*$, so that $x(t) \in A^* = [a^N, a^*]$.

Since $V(x)$ is bounded above by $\pi(a^*) = \max_a \pi(a)$, $\sup_{x \in X^{**}} V(x)$ is a finite number. Hence, by Lemma 2, we can find a sequence x^n, $n = 1, 2, ...$ in X^{**} such that $\lim_{n\to\infty} V(x^n) = \sup_{x \in X^{**}} V(x)$.

Note that $\{\pi(x^n(\cdot))\}_{n=1,2,...}$ is a collection of *countably* many measurable functions. This implies that $\overline{\pi}(t) := \sup_n \pi(x^n(t))(< \infty)$ is also measurable. Now let us define $\overline{x}(t)$ to be the solution to

$$\textbf{Problem P(t)}: \max_{x(t) \in [a^N, a^*]} \pi(x(t))$$

$$\text{s.t. } d(x(t))e^{-\lambda t} \leq \int_0^t \left(\overline{\pi}(s) - \pi^N\right) \lambda e^{-\lambda s} ds. \tag{A.1}$$

Note that the right hand side of the constraint (A.1) is well-defined, because $\overline{\pi}(\cdot)$ is measurable. Also note that the right hand side is nonnegative by $\overline{\pi}(s) \geq \pi^N$.[1]

Under Assumptions A5 and A6, both $\pi(a)$ and $d(a)$ are increasing on $[a^N, a^*]$. Hence the solution $\overline{x}(t)$ to Problem P(t) is either a^* or the action in $[a^N, a^*)$ with the binding constraint (A.1) by continuity of d (which follows from A2). Let us write down the solution in the following way. Note first that, by Assumptions A2 and A6, d is continuous and strictly increasing on $[a^N, a^*]$, and therefore its continuous inverse d^{-1} exists (if we regard d as a function from $[a^N, a^*]$ to $d([a^N, a^*]) = [0, d(a^*)]$). Then the optimal solution $\overline{x}(t)$ can be expressed as

$$\overline{x}(t) = \begin{cases} a^* & \text{if } d(a^*) < h(t) \\ d^{-1}(h(t)) & \text{otherwise} \end{cases}, \tag{A.2}$$

[1] By A5, $x^n(t) \in [a^N, a^*]$ implies $\pi(x^n(t)) \geq \pi^N$. Hence $\overline{\pi}(t) = \sup_n \pi(x^n(t)) \geq \pi^N$.

where

$$h(t) := e^{\lambda t} \int_0^t \left(\overline{\pi}(s) - \pi^N \right) \lambda e^{-\lambda s} ds.$$

A crucial step in the proof is to note that $h(t)$ is continuous in t for any measurable function $\overline{\pi}(\cdot)$.[2] Since d^{-1} is continuous, $\overline{x}(t)$ is continuous whenever $x(t) \in [a^N, a^*)$. Moreover, since $h(t)$ is increasing in t, (A.2) means that $x(t) = a^*$ implies $x(t') = a^*$ for all $t' > t$. Hence $\bar{x}$ is continuous for all t.

Lastly, we show that $\overline{x}$ is a trigger strategy equilibrium. The continuity of $\overline{x}$ and π implies that $\pi(\overline{x}(\cdot))$ is a measurable function. Therefore, $\overline{x}$ is feasible. We show that $\overline{x}$ also satisfies the (trigger strategy) incentive constraint IC(t) for all t. Recall that x^n is a trigger strategy equilibrium for all $n = 1, 2, ...$ Then we have

$$\begin{aligned} d(x^n(t))e^{-\lambda t} &\leq \int_0^t \left(\pi(x^n(s)) - \pi^N \right) \lambda e^{-\lambda s} ds \quad (x^n \text{ is an equilibrium}) \\ &\leq \int_0^t \left(\overline{\pi}(s) - \pi^N \right) \lambda e^{-\lambda s} ds. \qquad (\text{by definition of } \overline{\pi}) \end{aligned}$$

This means that $x^n(t)$ satisfies the constraint of Problem P(t). Since $\overline{x}(t)$ is the solution to Problem P(t), we have

$$\forall n \; \forall t \quad \pi(\overline{x}(t)) \geq \pi(x^n(t)) \tag{A.3}$$

and therefore

$$\forall t \;\; \pi(\overline{x}(t)) \geq \overline{\pi}(t) = \sup_n \pi(x^n(t)). \tag{A.4}$$

[2] Note to ourselves (may be omitted): The standard result in measure theory shows that, for any measurable function $f(t)$, the Lebesgue integral $\int_0^t f(s)ds$ is absolutely continuous in t, so it is continuous in t. (See, for example, S. Ito Thm 19.2).

Hence, for all t, $\overline{x}(t)$ satisfies the incentive constraint IC(t):

$$\begin{aligned} d(\overline{x}(t))e^{-\lambda t} &\leq \int_0^t \left(\overline{\pi}(s) - \pi^N\right) \lambda e^{-\lambda s} ds \;\; (\overline{x}(t) \text{ satisfies (A.1)}) \\ &\leq \int_0^t \left(\pi(\overline{x}(t)) - \pi^N\right) \lambda e^{-\lambda s} ds. \end{aligned}$$

Thus we have shown that $\overline{x}$ is a trigger strategy equilibrium ($\overline{x} \in X^*$), and $V(\overline{x}) \geq V(x^n)$ for all n (by (A.3)). By definition $\lim_{n\to\infty} V(x^n) = \sup_{x\in X^{**}} V(x)$, and the above inequality implies $V(\overline{x}) \geq \sup_{x\in X^{**}} V(x)$. Since $\overline{x} \in X^{**}$, we must have $V(\overline{x}) = \sup_{x\in X^{**}} V(x) = \max_{x\in X^{**}} V(x) (= \max_{x\in X^{*}} V(x)$ by Lemma 2). Hence we have established that there is an optimal and continuous trigger strategy equilibrium $\overline{x}$. □

Next, we show that $\overline{x}(t)$ satisfies binding incentive constraint and is differentiable. The continuity of $\overline{x}$ plays a crucial role in the proof.

Proposition 34. *The optimal trigger strategy equilibrium $\overline{x}(t)$ satisfies the binding incentive constraint*

$$d(\overline{x}(t))e^{-\lambda t} = \int_0^t \left(\pi(\overline{x}(t)) - \pi^N\right) \lambda e^{-\lambda s} ds.$$

if $\overline{x}(t) \neq a^*$, and $\overline{x}(t)$ is differentiable when $\overline{x}(t) \neq a^*, a^N$.

Proof. The proof of Proposition 33 shows that, if $\overline{x}(t) \neq a^*$, then

$$d(\overline{x}(t))e^{-\lambda t} = \int_0^t \left(\overline{\pi}(s) - \pi^N\right) \lambda e^{-\lambda s} ds$$

$$\leq \int_0^t \left(\pi(\overline{x}(t)) - \pi^N\right) \lambda e^{-\lambda s} ds. \tag{A.5}$$

We now show that the weak inequality above is actually an equality (and therefore we have the binding incentive constraint). If the above inequality were strict for some t, by (A.4), we would have

$$e^{-\lambda T}\overline{\pi}(T) + \int_0^T \overline{\pi}(s)\lambda e^{-\lambda s} ds < e^{-\lambda T}\pi(\overline{x}(T)) + \int_0^T \pi(\overline{x}(s))\lambda e^{-\lambda s} ds = V(\overline{x}).$$

Since $\overline{\pi}(s) := \sup_n \pi(x^n(t))$, the left hand side is more than or equal to $V(x^n)$ for all n. Since $\lim_{n\to\infty} V(x^n) = \sup_{x\in X^{**}} V(x)$, the above inequality implies $\sup_{x\in X^{**}} V(x) < V(\overline{x})$. This contradicts $\overline{x} \in X^{**}$. Hence (A.5) should be satisfied with an equality (i.e., $\overline{x}$ satisfies the binding incentive constraint), if $\overline{x}(t) \neq a^*$.

Next we show the differentiability. We continue to consider the case $a^N < a^*$, so that $\overline{x}(t) \in A^* = [a^N, a^*]$. By Assumptions A2 and A6, d is continuous and strictly increasing on $[a^N, a^*]$ and therefore its inverse d^{-1} exists. Hence, if $\overline{x}(t) \neq a^*$, the binding incentive constraint implies

$$\overline{x}(t) = d^{-1}\left(e^{\lambda t}\int_0^t \left(\pi(\overline{x}(s)) - \pi^N\right)\lambda e^{-\lambda s} ds\right).$$

The continuity of $\overline{x}$ implies that $\left(\pi(\overline{x}(s)) - \pi^N\right)\lambda e^{-\lambda s}$ is continuous, and the fundamental theorem of calculus shows that $\int_0^t \left(\pi(\overline{x}(s)) - \pi^N\right)\lambda e^{-\lambda s} ds$ is differentiable with respect to t (with the derivative $\left(\pi(\overline{x}(t)) - \pi^N\right)\lambda e^{-\lambda t}$). Hence the argument of d^{-1} is differentiable with respect to t, and therefore $\overline{x}(t)$ is differentiable whenever $(d^{-1})'$ exists. Note that $(d^{-1})' = 1/d'(\overline{x}(t))$ indeed exists if $\overline{x}(t) \neq a^N$, because d' exists (Lemma 3 in Appendix C) and $d'(\overline{x}(t)) > 0$ (Assumption A6). $\square$

A.2 Proof of Proposition 2

We provide the proof of Proposition 2 (essential uniqueness of the optimal path):

Proof. Suppose $H := \{t|\pi(y(t)) > \pi(\overline{x}(t))\}$ has a positive measure. Then, define

$$z(t) := \begin{cases} y(t) & \text{if } t \in H \\ \\ \overline{x}(t) & \text{otherwise} \end{cases}.$$

This has a measurable payoff $\pi(z(t)) = \max\{\pi(y(t)), \pi(\overline{x}(t))\}$ and achieves strictly higher expected payoff than $\overline{x}(t)$. Furthermore, z satisfies the incentive constraints

$$\forall t \;\; d(z(t))e^{-\lambda t} \leq \int_0^t \left(\pi(z(s)) - \pi^N\right) \lambda e^{-\lambda s} ds.$$

This follows from the incentive constraints for $\overline{x}$ and y, together with $\pi(z(t)) = \max\{\pi(y(t)), \pi(\overline{x}(t))\}$. Hence, z is a trigger strategy equilibrium path, which achieves a higher payoff than $\overline{x}(t)$ does. This contradicts the optimality of $\overline{x}(t)$, and therefore H must have measure zero. Hence $\pi(y(t)) \leq \pi(\overline{x}(t))$ almost everywhere. If $\{t|\pi(y(t)) < \pi(\overline{x}(t))\}$ has a positive measure, y attains a strictly smaller payoff than $\overline{x}(t)$ does, which contradicts our premise that y is optimal. Therefore we conclude that $\pi(y(t)) = \pi(\overline{x}(t))$ almost everywhere.

Finally we show that $y(t) = \overline{x}(t)$ almost everywhere. Note that π is not monotone and therefore $\pi(y(t)) = \pi(\overline{x}(t))$ may not imply $y(t) = \overline{x}(t)$. Since any trigger strategy equilibrium must play a^N at $t = 0$, suppose $\pi(y(t)) = \pi(\overline{x}(t))$ but $y(t) \neq \overline{x}(t)$, for $t > 0$. This means $y(t) \neq a^*$, so suppose $y(t) \neq a^*$. This will lead to a contradiction.

Consider the case of $a^N \leq a^*$. We must have $a^N < \overline{x}(t) < a^* < y(t)$ (see the graph of π (Figure A.1)). Since the incentive constraint is binding when $a^N \leq \overline{x}(t) < a^*$,

$$d(\overline{x}(t))e^{-\lambda t} = \int_0^t \left(\pi(\overline{x}(s)) - \pi^N\right) \lambda e^{-\lambda s} ds.$$

This implies that y does not satisfy the incentive constraint, because (i) Assumption A6 and $a^N < \overline{x}(t) < y(t)$ imply $d(\overline{x}(t)) < d(y(t))$, and (ii) $\pi(y(s)) = \pi(\overline{x}(s))$ almost everywhere. This is a contradiction, and therefore $y(t) = \overline{x}(t)$ almost everywhere. □

A.3 Auxiliary Lemmas for Theorem 1

We provide auxiliary lemmas to prove Theorem 1. First, we show that d' and d'' exist and are continuous under our assumptions.

Lemma 3. *Under A2-A4, both $d'(x)$ and $d''(x)$ exist and are continuous. In particular,*

$$d'(x) = \frac{\partial \pi_1(BR(x), x)}{\partial x_2} - \frac{\partial \pi_1(x, x)}{\partial x_1} - \frac{\partial \pi_1(x, x)}{\partial x_2}, \tag{A.6}$$

$$d''(x) = -\left(\frac{\partial^2 \pi_1(BR(x), x)}{\partial x_1 \partial x_2}\right)^2 / \frac{\partial^2 \pi_1(BR(x), x)}{\partial^2 x_1} + \frac{\partial^2 \pi_1(BR(x), x)}{\partial^2 x_2}$$
$$- \frac{\partial^2 \pi_1(x, x)}{\partial^2 x_1} - 2\frac{\partial^2 \pi_1(x, x)}{\partial x_1 \partial x_2} - \frac{\partial^2 \pi_1(x, x)}{\partial^2 x_2}, \text{ and} \tag{A.7}$$

$$d''(a^N) = \frac{-\left(\frac{\partial^2 \pi_1(a^N, a^N)}{\partial^2 x_1} + \frac{\partial^2 \pi_1(a^N, a^N)}{\partial x_1 \partial x_2}\right)^2}{\frac{\partial^2 \pi_1(a^N, a^N)}{\partial^2 x_1}} \tag{A.8}$$

Proof. We first examine the properties of $BR(x)$. To this end, we apply the implicit function theorem to the first order condition $\frac{\partial \pi_1(BR(x),x)}{\partial x_1} = 0$ (A4). The assumptions of implicit function theorem are satisfied:

- $\frac{\partial^2 \pi_1(BR(x),x)}{\partial^2 x_1} \neq 0$ (by A4) and
- $\frac{\partial \pi_1(x_1,x_2)}{\partial x_1}$ is continuously differentiable (A2).

Hence $BR(x)$ is a continuously differentiable function (and therefore also continuous), with

$$BR'(x) = -\frac{\partial^2 \pi_1(BR(x),x)}{\partial x_1 \partial x_2} / \frac{\partial^2 \pi_1(BR(x),x)}{\partial^2 x_1},$$

and it is finite. Given this, differentiating $d(x) := \pi_1(BR(x),x) - \pi_1(x,x)$ and using the first order condition $\frac{\partial \pi_1(BR(x),x)}{\partial x_1} = 0$ (A4), we obtain (A.6). Differentiating this once again and using the above formula for $BR'(x)$, we obtain (A.7). By the twice continuous differentiability of π_1 (A2), $\frac{\partial^2 \pi_1(BR(x),x)}{\partial^2 x_1} \neq 0$ (by A4), and the continuity of $BR(x)$, both d' and d'' are continuous. Lastly, (A.8) is obtained from (A.7), by noting that $BR(x) = x$ when x is equal to the Nash action a^N. □

Next we show $f(x) := \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$, which defies the differential equation $dx/dt = f$, is continuously differentiable.

Lemma 4. *Function* $f(x) := \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$ *is continuously differentiable for* $x \neq a^N$.

Proof. Note that $d'(x) \neq 0$ if $x \neq a^N$ (A6). Then, $f' = \frac{\lambda((d'+\pi')d' - \left(d+\pi-\pi^N\right)d'')}{(d')^2}$ is a continuous function, by Lemma 3. □

We now examine the behavior of $dx/dt = f(x)$ when x is close to a^N. In particular, we evaluate $f^N := \lim_{x \to a^N} \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$.

Lemma 5.

$$f^N \quad : \quad = \lim_{x\to a^N} \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$$
$$= \quad -\lambda\ \pi'(a^N) \Big/ \frac{\left(\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}+\frac{\partial^2\pi_1(a^N,a^N)}{\partial x_1\partial x_2}\right)^2}{\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}}.$$

Under Assumptions A1-A6, f^N is always non-zero, and $f^N = \infty$ or $-\infty$ if and only if $\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}+\frac{\partial^2\pi_1(a^N,a^N)}{\partial x_1\partial x_2}=0$.

Proof. By de l'Hopital rule,

$$\lim_{x\to a^N} \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)} = \frac{\lambda\pi'(a^N)}{d''(a^N)}$$

where we used $d'(a^N)=0$ (A6). Then the expression of the lemma directly follows from (A.8) in Lemma 3. The numerator is non-zero, because $\pi'=0$ only at the optimal action a^* (A5). By the second order condition at the Nash equilibrium (A4), $\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}<0$. Hence, $f^N\neq 0$ in general, and $f^N=\infty$ or $-\infty$ if and only if $\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}+\frac{\partial^2\pi_1(a^N,a^N)}{\partial x_1\partial x_2}=0$. ☐

Remark 2. *The condition for the finiteness of f^N, $\frac{\partial^2\pi_1(a^N,a^N)}{\partial^2 x_1}+\frac{\partial^2\pi_1(a^N,a^N)}{\partial x_1\partial x_2}\neq 0$ is equivalent to $BR'(a^N)\neq 1$. This follows from the implicit function theorem $BR' = -\frac{\partial^2\pi_1}{\partial x_1\partial x_2}/\frac{\partial^2\pi_1}{\partial^2 x_1}$.*

Finally we show that the finite time condition (2.8) is satisfied under our assumptions. Recall that we are looking at the case where $a^N<a^*$.

Lemma 6. *For any $x^0\in(a^N,a^*]$, $t(x^0) := \lim_{a\to a^N}\int_a^{x^0}\frac{1}{f(x)}dx<\infty$.*

Proof. Recall
$$\frac{1}{f} = \frac{d'(x)}{\lambda\left(d(x)+\pi(x)-\pi^N\right)}$$

and it is finite when $x \in (a^N, a^*)$ because the numerator is finite by Lemma 3 and the denominator is nonzero by A5 and A6. Note that $1/f(x)$ is not defined for $x = a^N$ (both the numerator and denominator of the right hand side is zero at $x = a^N$). By Lemma 5, we have

$$\lim_{a \to a^N} \frac{1}{f(a)} = \frac{1}{f^N} = \frac{-\left(\frac{\partial^2 \pi_1(a^N, a^N)}{\partial^2 x_1} + \frac{\partial^2 \pi_1(a^N, a^N)}{\partial x_1 \partial x_2}\right)^2}{\lambda \pi'(a^N) \frac{\partial^2 \pi_1(a^N, a^N)}{\partial^2 x_1}},$$

By Assumption A5 (and $a^N < a^*$), we have $\pi'(a^N) > 0$, and also $\frac{\partial^2 \pi_1(a^N, a^N)}{\partial^2 x_1} < 0$ by A4. Also the numerator is finite by A2. Therefore $\lim_{a \to a^N} \frac{1}{f(a)}$ is a finite number. Hence, $t(x^0) := \lim_{a \to a^N} \int_a^{x^0} \frac{1}{f(x)} dx$ is a finite number. □

B. APPENDIX TO CHAPTER 3

B.1 Proof of Lemma 1

Proof. Suppose that the premise of the lemma holds. Let t^* be the supremum of t such that A_t is false. If $t^* = -\infty$, then we are done. Suppose that $t^* > -\infty$. Then it must be the case that for any $\epsilon > 0$, there exists $t' \in (t^* - \epsilon, t^*]$ such that $A_{t'}$ is false. But by the definition of t^*, there exists $\epsilon' > 0$ such that statement $A_{t'}$ is true for all $t' \in (t^* - \epsilon', t^*]$ because the premise of the lemma is true. Contradiction. $\square$

B.2 A Sharper Result for the Case when Players Move at $-T$

Proposition 35. *Suppose that players choose their action at $-T$ and consider a component game of a revision game with a strictly Pareto-dominant action profile x^*. Then there exists T' such that for all $T > T'$, in all SPE, $x(0) = x^*$ with probability* 1.

Proof. Suppose without loss of generality that $\lambda_1 \leq \lambda_2$. Consider first the case of $\lambda_1 < \lambda_2$. Fix an SPE strategy profile where player 1 prepares $x_1 \neq x_1^*$ in $T < 0$. In this case, player 1's expected payoff at time $-T$ is at most

$$u_1(x^*) - e^{-\lambda_1 T} m,$$

as with probability $e^{-\lambda_1 T}$, player 1 has no further revision opportunities. On the other hand, one possible deviation is to play x_i^* for all $[-T, 0]$, and in that case the expected payoff is

$$u_1(x^*) - e^{-\lambda_2 T} M,$$

since by Step 1 in the proof of Theorem 1, it follows that player 2 will switch to x_2^* as soon as he has a chance to revise, and afterwards the PAP never changes.[1] However the assumption that $\lambda_1 < \lambda_2$ implies that for sufficiently large T, the latter value becomes strictly greater than the former, implying that in any SPE, player 1 must prepare x_1^* when T is sufficiently large. Given this, player 2 has a strict incentive to prepare x_2^* at $-T$ as that would give him the highest possible expected payoff in equilibrium, while preparing some other action results in a strictly lower payoff because there is a strictly positive probability that he has no chance to revise the action in the future.

Suppose $\lambda_1 = \lambda_2 \equiv \lambda$. Fix a revision equilibrium, and let $V_i^t(x)$ be player i's value from the revision equilibrium when the PAP is x at time t. Let $v_1^t(x_2) = \max_{x_1 \neq x_1^*} V_1^t(x_1, x_2)$, that is, player 1's maximum value at t when player 2 prepares x_2 conditional on that player 1 does not prepare x_1^*. On the other hand, for any t, the lower bound of always taking x_1^* at t is

$$u_1(x^*) - e^{\lambda_2 t} M \tag{B.1}$$

since player 1 can stick to x_1^* in the continuation game after t. It suffices to show that for any $x_2 \neq x_2^*$, there exists $\bar{t}$ such that player 1 strictly prefers taking x_1^* at $\bar{t}$ when PAP by player 2 is x_2:

---

[1] Here we use the fact that there are only two players. If there are two or more opponents, Step 1 cannot be used to conclude that *all* the opponents will switch to actions prescribed by x^*.

$$v_1^{\bar{t}}(x_2) < u_1(x^*) - e^{\lambda_2 \bar{t}} M. \tag{B.2}$$

To see this, suppose player 2 prepares x_2 at $t < \bar{t}$. Player 1 wants to prepare x_1^* from t and $\bar{t}$. Conditional on that player 2 can move once by $\bar{t}$, player 1 gets the highest payoff $u_1(x^*)$. Conditional on that player 2 cannot move by $\bar{t}$, PAP at $\bar{t}$ is x_2 and player 1's value is highest at x_1^*. Hence, x_1^* becomes a strictly dominant strategy eventually and player 1 takes x_1^* for sure at the beginning of the game. Given that, player 2's unique best response is to take x_2^*.

Take any $x_2 \neq x_2^*$ and S with $e^{\lambda S} < 1$. Since both players need to move afterwards to go to x^* from (x_1, x_2),

$$v_1^S(x_2) \leq u_1(x^*) - m\left(e^{\lambda_1 S} + e^{\lambda_2 S} - e^{(\lambda_1+\lambda_2)S}\right). \tag{B.3}$$

Below, we will show that, at nS, one of the following two is correct: (i) we have (B.2) with $\bar{t} = nS$ or (ii)

$$v_1^{nS}(x_2) \leq u_1(x^*) - \left(n + 1 - ne^{\lambda S}\right) me^{n\lambda S}. \tag{B.4}$$

For $n = 1$, (B.4) is (B.3) with $n = 1$. Hence, (ii) is true for $n = 1$. Suppose (i) or (ii) holds for $n = k$. If (i) is the case, we are done. Otherwise, at $(k+1)\, S$, the upper bound of not taking x_1^* is determined as follows: if player 2 cannot move by $-kS$, PAP by player 2 at $-kS$ is x_2. Then, since (i) is not the case, the maximum of player 1's payoff is given by $v_1^{-kS}(x_2)$. Therefore,

$$
\begin{aligned}
& u_1(x^*) - \underbrace{e^{\lambda S}}_{\text{Pr of 2 not entering by } -kS} \times \underbrace{\left(k+1-ke^{\lambda S}\right) me^{k\lambda S}}_{\substack{\text{from above,}\\ \text{this is the least loss compared to } u_1(x^*).}} \\
& - \underbrace{\left(1-e^{\lambda S}\right)}_{\text{Pr of 2 entering by } -kS} \times \underbrace{e^{\lambda S}}_{\text{Pr of 1 not entering by } -kS} \times \underbrace{me^{\lambda kS}}_{\text{the least loss for } x \neq x^*} \\
= \; & u_1(x^*) - \left(k+2-(k+1)\,e^{\lambda S}\right) me^{\lambda(k+1)S},
\end{aligned}
$$

which is (B.4) with $n = k+1$ as desired. Since (ii) implies (B.2) for sufficiently large n, this completes the proof. □

B.3 *Proof of Theorem 3: Derivation of* t_i^*

We provide a derivation of t_1^*. The value of t_2^* can be solved for in a symmetric manner. By the definition of t_1^*, assuming that both players play best responses to the PAP at any time strictly after t_1^*, the payoff from playing a best response against R at t_1^* and playing otherwise must be equal. Thus, it must be the case that

$$
\begin{aligned}
u_1(D,R) = & \underbrace{e^{(\lambda_1+\lambda_2)t_1^*}u_1(U,R)}_{\text{nobody moves until 0}} + \underbrace{\tfrac{\lambda_1}{\lambda_1+\lambda_2}(1-e^{(\lambda_1+\lambda_2)t_1^*})u_1(D,R)}_{\text{player 1 moves first}} \\
& + \underbrace{\tfrac{\lambda_2}{\lambda_1+\lambda_2}(1-e^{(\lambda_1+\lambda_2)t_1^*})u_1(U,L)}_{\text{player 2 moves first}}.
\end{aligned}
$$

Solving this equation with respect to t_1^*, we obtain the desired expression. □

B.4 *Derivation of t_i^* for Non-Homogeneous Poisson Processes*

For any non-empty interior time interval $[t, t'] \subseteq [-T, 0]$, let $Q_i(t, t')$ be the probability that player i has at least one revision opportunity in this interval that arrives strictly before all revision opportunities of the other player in the same interval. Let $Q_0(t, t')$ denote the probability that no player receives a revision opportunity in $[t, t']$. Namely,

$$\begin{aligned} Q_i(t,t') &= 1 - \frac{L_i(t,t')}{L_1(t,t') + L_2(t,t')} \exp\left(-\left(L_1(t,t') + L_2(t,t')\right)\right) \text{ for all } i \\ &\quad \text{and} \\ Q_0(t,t') &= \exp\left(-\left(L_1(t,t') + L_2(t,t')\right)\right). \end{aligned}$$

Given this, define t_i^{**} as the unique solution for

$$\begin{aligned} u_1(D,R) &= \underbrace{Q_0(t_1^{**},0)u_1(U,R)}_{\text{nobody moves until 0}} + \underbrace{Q_1(t_1^{**},0)u_1(D,R)}_{\text{player 1 moves first}} + \underbrace{Q_2(t_1^{**},0)u_1(U,L)}_{\text{player 2 moves first}}, \\ u_2(U,L) &= \underbrace{Q_0(t_2^{**},0)u_2(U,R)}_{\text{nobody moves until 0}} + \underbrace{Q_1(t_2^{**},0)u_2(D,R)}_{\text{player 1 moves first}} + \underbrace{Q_2(t_2^{**},0)u_2(U,L)}_{\text{player 2 moves first}}. \end{aligned}$$

Under non-homogeneous Poisson processes, Theorem 2 holds by replacing t_i^* in its statement with t_i^{**} defined above.

C. APPENDIX TO CHAPTER 4

C.1 Numerical Results for Finite Arrival Rates

C.1.1 Uniform Distribution over Multi-Dimensional Triangle (Case 1)

Consider the distribution given by the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i \leq 1\}$. The result is summarized in C.1.

C.1.2 Uniform Distribution over a Sphere (Case 2)

Consider the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i^2 \leq 1\}$. We get the result shown in C.2. Note that the limit duration for $n = 1$ is the same as in the case of uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i \leq 1\}$.

C.1.3 Uniform Distribution over a Cube (Case 3)

Consider the distribution given by the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \max_{i \in N} x_i \leq 1\}$. The result is summarized in C.3.

C.1.4 Exponential Distribution (Case 4)

Consider the exponential distribution with parameter a_i for each player i. The result is summarized in Table C.4.

Table C.1: The distribution given by the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i \leq 1\}$ (Case 1).

		λ				
		10	20	30	100	∞
$n = 1$	Expected duration	0.398	0.366	0.355	0.340	0.333
	Percentage (%)	19.4	9.92	6.64	2.00	0
$n = 2$	Expected duration	0.608	0.591	0.585	0.576	0.57143
	Percentage (%)	6.48	3.44	2.35	0.731	0
$n = 3$	Expected duration	0.734	0.716	0.709	0.698	0.692
	Percentage (%)	5.97	3.35	2.35	0.780	0

Table C.2: The uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i^2 \leq 1\}$ (Case 2).

		λ					
		10	20	30	100	1000	∞
$n = 1$	Expected duration	0.398	0.366	0.355	0.340	0.334	0.333
	Percentage (%)	19.4	9.92	6.64	2.00	0.200	0
$n = 2$	Expected duration	0.582	0.568	0.565	0.562	0.567	0.571
	Percentage (%)	1.90	−0.541	−1.21	−1.61	−0.798	0

Table C.3: The distribution given by the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \max_{i \in N} x_i \leq 1\}$ (Case 3).

		λ				
		10	20	30	100	∞
$n = 1$	Expected duration	0.398	0.366	0.355	0.340	0.333
	Percentage (%)	19.4	9.92	6.64	2.00	0
$n = 2$	Expected duration	0.545	0.524	0.516	0.505	0.5
	Percentage (%)	9.09	4.76	3.23	0.990	0
$n = 3$	Expected duration	0.634	0.618	0.612	0.604	0.6
	Percentage (%)	5.62	3.00	2.05	0.643	0

Table C.4: The exponential distribution with parameter a_i for each player i (Case 4).

		λ				
		10	20	30	100	∞
$n = 1$	Expected duration	0.545	0.524	0.516	0.505	0.5
	Percentage (%)	9.09	4.76	3.23	0.990	0
$n = 2$	Expected duration	0.693	0.681	0.676	0.670	0.667
	Percentage (%)	3.91	2.11	1.45	0.465	0
$n = 3$	Expected duration	0.767	0.759	0.756	0.752	0.75
	Percentage (%)	2.27	1.24	0.864	0.284	0
$n = 10$	Expected duration	0.912	0.911	0.910	0.910	0.909
	Percentage (%)	0.370	0.206	0.145	0.0499	0

Table C.5: The log-normal distribution with $\mu = 0$ (Case 5).

		λ			
		10	20	30	100
	$\sigma = \frac{1}{4}$	0.449	0.462	0.469	0.484
$n = 1$	$\sigma = 1$	0.612	0.595	0.588	0.575
	$\sigma = 4$	0.961	0.952	0.946	0.926

C.1.5 Log-Normal Distribution (Case 5)

Consider the log-normal distribution with the following pdf:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}}.$$

Assume $\mu = 0$. The expected durations can be calculated and summarized in Table C.5.

C.2 Proofs of the Results

C.2.1 Computation of the Limit Durations

We here prove a lemma that computes the limit cumulative disagreement probability and the limit expected duration when the agreement probability $p(t)$ at time $-t$ is of the same order as $\frac{1}{\lambda t}$.

Lemma 7. *The following three statements hold:*

(i) If for all $\varepsilon > 0$, there exist $C > 0$ and $\bar{\lambda}$ such that $p(t) \leq \frac{C}{\lambda t}$ for all $t \geq \varepsilon$ and all $\lambda \geq \bar{\lambda}$, then $\liminf_{\lambda \to \infty} P(t; \lambda) \geq \left(\frac{t}{T}\right)^C$ for all $t \geq 0$, and $\liminf_{\lambda \to \infty} D(\lambda) \geq \frac{1}{1+C}$.

(ii) If for all $\varepsilon > 0$, there exist $c > 0$ and $\bar{\lambda}$ such that $p(t) \geq \frac{c}{\lambda t}$ for all $t \geq \varepsilon$ and all $\lambda \geq \bar{\lambda}$, then $\limsup_{\lambda \to \infty} P(t; \lambda) \leq \left(\frac{t}{T}\right)^c$ for all $t \geq 0$, and $\limsup_{\lambda \to \infty} D(\lambda) \leq \frac{1}{1+c}$.

(iii) If $\lim_{\lambda \to \infty} p(t)\lambda t = a > 0$ for all $t > 0$, then $P(t; \infty) = \left(\frac{t}{T}\right)^a$ for all $t \geq 0$, and $D(\infty) = \frac{1}{1+a}$.

Proof. First we prove (i). Let us fix $0 < \varepsilon < T$. By formula (4.3), for all $\lambda \geq \bar{\lambda}$ and all $t \geq \varepsilon$,

$$\begin{aligned} e^{-\int_t^T (C/s)ds} &\leq P(t; \lambda) \\ \left(\frac{t}{T}\right)^C &\leq P(t; \lambda). \end{aligned}$$

Since the above inequality is satisfied for all $\varepsilon > 0$ and sufficiently large λ, we have $\liminf_{\lambda \to \infty} P(t; \lambda) \geq \left(\frac{t}{T}\right)^C$ for all $t \geq 0$. By formula (4.4), $D(\lambda)T = \int_0^T P(t)dt$ is bounded as follows:

$$\begin{aligned} \int_\varepsilon^T \left(\frac{t}{T}\right)^C dt &\leq D(\lambda)T \\ \frac{T^{1+C} - \varepsilon^{1+C}}{(1+C)T^C} &\leq D(\lambda)T. \end{aligned}$$

Since the above inequality is satisfied for all $\varepsilon > 0$ and sufficiently large λ, we have $\liminf_{\lambda \to \infty} D(\lambda) \geq \frac{T^{1+C}}{(1+C)T^C \cdot T} = \frac{1}{1+C}$.

Next, a parallel argument shows (ii). Finally, (i) and (ii) together imply (iii). □

C.2.2 *Proof of Proposition 12*

Suppose that there exists at least one trembling-hand equilibrium. We show that the continuation payoff of player i at time $-t$ is unique for almost all histories in any trembling-hand equilibrium.

By Assumption 1 (a), the set of player i's expected payoffs given by any play of the game within $[-T, 0]$ is bounded by a value $\overline{x}_i$ for each $i \in N$. By Assumption 1 (b), we can find a Lipschitz constant L_i for $i \in N$ such that $\mu(\{x \in X \mid x_i \in [x_i', x_i'']\}) \le L_i |x_i' - x_i''|$ for all x_i', x_i'' in the above domain of payoffs. Let $L = \max_i L_i$.

Let $S_i(\sigma, t) \subseteq \mathbb{R}$ be the support of the continuation payoffs $u_i(\sigma \mid h)$ of player i after histories $h \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ realized at time $-t$ given a strategy profile σ. For $\varepsilon \in (0, \frac{1}{2})$, let $\overline{v}_i^\varepsilon(t)$ and $\underline{v}_i^\varepsilon(t)$ be the supremum and the infimum of

$$\bigcup_{\sigma\text{: Nash equilibrium in } \Sigma^\varepsilon} S_i(\sigma, t).$$

(Note that Assumption 1 (a) ensures boundedness of the support for finite t.) Let $w_i^\varepsilon(t) = \overline{v}_i^\varepsilon(t) - \underline{v}_i^\varepsilon(t)$, and $\bar{w}^\varepsilon(t) = \max_{i \in N} w_i^\varepsilon(t)$. We will show that $\bar{w}^\varepsilon(t) = 0$ for all $\varepsilon > 0$ for any time $-t \in [-T, 0]$. Note that $\bar{w}^\varepsilon(0) = 0$ for all ε.

Let us consider the ε-constrained game. If player i accepts an allocation $x \in X$ at time $-t$, she will obtain x_i with probability at least ε^{n-1}. Accepting x is a dominant action of player i if the following inequality holds:

$$\varepsilon^{n-1} x_i + (1 - \varepsilon^{n-1}) \underline{v}_i^\varepsilon(t) > \overline{v}_i^\varepsilon(t).$$

Rearranging this, we have

$$x_i > \overline{v}_i^\varepsilon(t) + \frac{1 - \varepsilon^{n-1}}{\varepsilon^{n-1}} w_i^\varepsilon(t).$$

Let $\tilde{v}_i^\varepsilon(t) = \overline{v}_i^\varepsilon(t) + \frac{1-\varepsilon^{n-1}}{\varepsilon^{n-1}} w_i^\varepsilon(t)$, the right hand side of the above inequality. Then $\tilde{v}_i^\varepsilon(t) - \underline{v}_i^\varepsilon(t) = \frac{1}{\varepsilon^{n-1}} w_i^\varepsilon(t)$.

Let $X_i^1(t) = \{x \in X \mid x_i > \tilde{v}_i^\varepsilon(t)\}$, $X_i^m(t) = \{x \in X \mid \underline{v}_i^\varepsilon(t) \le x_i \le \tilde{v}_i^\varepsilon(t)\}$, and

$X_i^0(t) = \{x \in X \mid x_i < \underline{v}_i^\varepsilon(t)\}$. Then $\mu(X_i^m) \leq \frac{L}{\varepsilon^{n-1}} w_i^\varepsilon(t)$. Any player i accepts $x \in X_i^1(t)$ and rejects $x \in X_i^0(t)$ with probability $1-\varepsilon$ after almost all histories at time $-t$. Note that $X = \big(\bigcup_{j \in N} X_j^m(t)\big) \cup \big(\bigcup_{(s_1,\ldots,s_n)\in\{0,1\}^n} \bigcap_{j\in N} X_j^{s_j}(t)\big)$ (where $X_j^m(t)$'s have a nonempty intersection). Then

$$\begin{aligned}\overline{v}_i^\varepsilon(t) \leq \int_0^t \Big(&\sum_{j\in N} \int_{X_j^m(\tau)} \overline{x}_i d\mu \\ &+ \sum_{(s_1,\ldots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big((1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)} x_i \\ &\quad + (1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}) \overline{v}_i^\varepsilon(\tau)\big) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau,\end{aligned}$$

and

$$\begin{aligned}\underline{v}_i^\varepsilon(t) \geq \int_0^t \Big(&\sum_{(s_1,\ldots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big((1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)} x_i \\ &\quad + (1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}) \underline{v}_i^\varepsilon(\tau)\big) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau.\end{aligned}$$

Therefore $w_i^\varepsilon(t) = \overline{v}_i(t) - \underline{v}_i(t)$ is bounded as follows:

$$\begin{aligned}
w_i^\varepsilon(t) \leq & \int_0^t \Big(\sum_{j\in N} \int_{X_j^m(\tau)} \overline{x}_i d\mu \\
& + \sum_{(s_1,\ldots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) w_i^\varepsilon(\tau) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau \\
\leq & \int_0^t \Big(\sum_{j\in N} \overline{x}_i \frac{L}{\varepsilon^{n-1}} w_j^\varepsilon(\tau) \\
& + \sum_{(s_1,\ldots,s_n)\in\{0,1\}^n} \int_X \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) w_i^\varepsilon(\tau) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau \\
\leq & \int_0^t \Big(\sum_{j\in N} \max_{k\in N}\{\overline{x}_k\} \frac{L}{\varepsilon^{n-1}} \\
& + \sum_{(s_1,\ldots,s_n)\in\{0,1\}^n} \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) \Big) \overline{w}^\varepsilon(\tau) \lambda e^{-\lambda(t-\tau)} d\tau.
\end{aligned}$$

Since the above inequality holds for all $i \in N$, there exists a constant $M > 0$ such that the following inequality holds:

$$\overline{w}^\varepsilon(t) \leq \int_0^t M \overline{w}^\varepsilon(\tau) e^{-\lambda(t-\tau)} d\tau.$$

Let $W^\varepsilon(t) = \int_0^t \overline{w}^\varepsilon(\tau) e^{\lambda\tau} d\tau$. Then

$$\begin{aligned}
W^{\varepsilon\prime}(t) &= \overline{w}^\varepsilon(t) e^{\lambda t} \\
&\leq M W^\varepsilon(t).
\end{aligned}$$

Therefore we have $\frac{d}{dt}\big(W^\varepsilon(t)e^{-Mt}\big) = \big(W^{\varepsilon\prime}(t) - MW^\varepsilon(t)\big)e^{-Mt} \leq 0$. Since $W^\varepsilon(0) = 0$ by the definition of $W^\varepsilon(t)$, $W^\varepsilon(t)e^{-Mt} \leq 0$ for all $t \geq 0$. This implies that $\overline{w}^\varepsilon(t) \leq MW^\varepsilon(t)e^{-\lambda t} \leq 0$ for all $t \geq 0$. Hence, $\overline{w}^\varepsilon(t) = 0$ for all $t \geq 0$ and all $\varepsilon \in (0, \frac{1}{2})$. Any trembling-hand equilibria yield the same continuation payoffs after almost all histories at time $-t \in [-T, 0]$.

C.2.3 *Proof of Proposition 13*

We show that a solution $v^*(t)$ of ODE (4.1) characterizes a trembling-hand equilibrium. For $s_i \in \{+, -\}$, and $v_i \in [0, \infty)$ let

$$I_i^{s_i}(v_i) = \begin{cases} [0, v_i] & \text{if } s_i = +, \\ [v_i, \infty) & \text{if } s_i = -, \end{cases}$$

and $p^+ = 1 - \varepsilon$, $p^- = \varepsilon$. For $\varepsilon > 0$, let us write down a Bellman equation similar to (4.2) with respect to a continuation payoff profile $v^\varepsilon(t)$ in the ε-constrained game:

$$\begin{aligned} &v_i^\varepsilon(t) = \\ &\int_0^t \Big(\sum_{s \in \{+,-\}^n} \int_{(I_1^{s_1}(v_i^\varepsilon(\tau)) \times \cdots \times I_n^{s_n}(v_i^\varepsilon(\tau))) \cap X} \big(p^{s_1} \dots p^{s_n} \cdot x_i + (1 - p^{s_1} \dots p^{s_n}) v_i^\varepsilon(\tau) \big) d\mu \Big) \\ &\qquad \cdot \lambda e^{-(\lambda+\rho)(t-\tau)} d\tau \end{aligned}$$

This implies that

$$\begin{aligned} v_i^{\varepsilon\prime}(t) = -(\lambda + \rho) v_i^\varepsilon(t) \quad &+ \lambda \sum_{s \in \{+,-\}^n} \\ &\int_{(I_1^{s_1}(v_i^\varepsilon(t)) \times \cdots \times I_n^{s_n}(v_i^\varepsilon(t))) \cap X} \big(p^{s_1} \dots p^{s_n} \cdot x_i + (1 - p^{s_1} \dots p^{s_n}) v_i^\varepsilon(t) \big) d\mu. \end{aligned}$$

This ODE has a unique solution because the right hand side is Lipschitz continuous in v_i^ε. Let $v^\varepsilon(t)$ be this solution, which is a cutoff profile of a Nash equilibrium in the ε-constrained game by construction. Since $A(t) = (I_1^+(v_i^\varepsilon(\tau)) \times \cdots \times I_n^+(v_i^\varepsilon(\tau))) \cap X$, and $p^+ \to 1$, $p^- \to 0$ as $\varepsilon \to 0$, ODE (4.1) is obtained by letting $\varepsilon \to 0$. Therefore

$v^{\varepsilon}(t)$ converges to $v^*(t)$ as $\varepsilon \to 0$ because the above ODE is continuous in ε.[1] Hence the cutoff strategy profile with cutoffs $v^*(t)$ is a trembling-hand equilibrium.

C.2.4 Proof of Theorem 5

By Assumption 2, there exists a nondecreasing and concave function φ and $\kappa \geq 1$ such that

$$1 - \varphi(x) \leq 1 - F(x) \leq \kappa\big(1 - \varphi(x)\big)$$

for all $x \geq 0$. Let us consider a cutoff strategy with the following cutoff $w(t)$:

$$w(t) = F^{-1}\Big(1 - \frac{2}{\lambda t + 2}\Big),$$

namely, the strategy with acceptance probability $\frac{2}{\lambda t+2}$ at time $-t$. By Assumption 2, we have $w(t) \geq \varphi^{-1}\big(1 - \frac{2}{\lambda t+2}\big)$. Let $P(t)$ be the probability that the search stops before time $-t$ when $w(t)$ is played. Then

$$\begin{aligned} P(t) &= 1 - \exp\Big(-\int_t^T \frac{2}{\lambda\tau + 2} \cdot \lambda d\tau\Big) \\ &= 1 - \Big(\frac{\lambda t + 2}{\lambda T + 2}\Big)^2. \end{aligned}$$

The expected continuation payoff obtained from this strategy is larger than

$$\int_t^0 w(\tau) dP(\tau) \geq \int_0^t \varphi^{-1}\Big(1 - \frac{2}{\lambda\tau + 2}\Big) d\big(1 - P(\tau)\big).$$

Let $W(t)$ be the payoff on the right hand side, and $Q(t)$ be the probability that the search stops before time $-t$ when the player plays a cutoff strategy with cutoff

[1] See, e.g., Coddington and Levinson (1955, Theorem 7.4 in Chapter 1).

$W(t)$.

$$
\begin{aligned}
Q(t) &= 1 - \exp\left(- \int_t^T \big(1 - F(W(\tau))\big)\lambda d\tau \right) \\
&\leq 1 - \exp\left(- \int_t^T \kappa\big(1 - \varphi(W(\tau))\big)\lambda d\tau \right).
\end{aligned}
$$

By concavity of φ, $\varphi(W(\tau))$ is bounded as follows:

$$
\begin{aligned}
\varphi(W(t)) &= \varphi\left(\int_0^t \varphi^{-1}\Big(1 - \frac{2}{\lambda\tau + 2}\Big) d\big(1 - P(\tau)\big) \right) \\
&\geq \int_0^t \varphi\left(\varphi^{-1}\Big(1 - \frac{2}{\lambda\tau + 2}\Big) \right) d\big(1 - P(\tau)\big) \\
&= \int_0^t \Big(1 - \frac{2}{\lambda\tau + 2}\Big) d\Big(\Big(\frac{\lambda\tau + 2}{\lambda T + 2}\Big)^2\Big) \\
&= 1 - \frac{4}{\lambda T + 2} + \frac{4}{(\lambda T + 2)^2} \\
&\geq 1 - \frac{4}{\lambda T + 2}.
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
Q(t) &\leq 1 - \exp\left(- \int_t^T \kappa\Big(\frac{4}{\lambda T + 2}\Big)\lambda d\tau \right) \\
&= 1 - \exp\left(-4\kappa \ln\Big(\frac{\lambda T + 2}{\lambda t + 2}\Big) \right) \\
&= 1 - \Big(\frac{\lambda t + 2}{\lambda T + 2}\Big)^{4\kappa},
\end{aligned}
$$

which is strictly lower than 1 for all $\lambda > 0$ and all $-t \in (-T, 0]$. Since $W(t)$ is the continuation payoff calculated from a strategy that is not necessarily optimal, an optimal strategy gives the player continuation payoffs larger than or equal to $W(t)$. Therefore an optimal strategy must possess a cutoff higher than or equal to $W(t)$. Hence, for all $-t \in (-T, 0]$, the search stops with probability strictly lower than 1 before time $-t$. This proves Theorem 5.

Next, we show a proposition under an independent assumption of Assumption 2 when the support is bounded.

Assumption 6. *If X is bounded, then for $\bar{x} = \sup X$, there exists $\alpha \geq \beta > 0$ such that $\varepsilon^{\alpha} \leq \mu\big([\bar{x} - \varepsilon, \bar{x}]\big) \leq \varepsilon^{\beta}$ for all $\varepsilon \in (0, 1)$.*

Proposition 36. *Suppose that X is bounded, and Assumptions 1 and 6 hold. Then $\liminf_{\lambda \to \infty} D(\lambda) > 0$.*

Proof. By ODE (4.1), the equilibrium continuation payoff $v^*(t)$ is the solution of

$$v'(t) = \lambda \int_{v(t)}^{\bar{x}} \big(x - v(t)\big) d\mu(x).$$

Let $z(t) = \bar{x} - v^*(t)$. Since $z(t) \to 0$ as $t \to \infty$, there exists $\bar{t}$ such that $z(t) < 1$ for all $t \geq \bar{t}$. For $t \geq \bar{t}$, $z(t)$ satisfies

$$\begin{aligned} z'(t) &= -\lambda \int_{\bar{x} - z(t)}^{\bar{x}} \big(\bar{x} - x + z(t)\big) d\mu(x) \\ &\leq -\lambda \int_{\bar{x} - \frac{z(t)}{2}}^{\bar{x}} \frac{z(t)}{2} d\mu(x) \\ &\leq -\lambda \cdot \frac{z(t)}{2} \cdot \Big(\frac{z(t)}{2}\Big)^{\alpha}. \end{aligned}$$

Solving this, $z(t) \leq \big(2^{-(1+\alpha)} \lambda (t - \bar{t}) + z(\bar{t})^{-\alpha}\big)^{-\frac{1}{\alpha}}$. Therefore,

$$\begin{aligned} p(t) &= \mu\big([\bar{x} - z(t), \bar{x}]\big) \\ &\leq z(t)^{\beta} \\ &\leq \big(2^{-(1+\alpha)} \lambda (t - \bar{t}) + z(\bar{t})^{-\alpha}\big)^{-\frac{\beta}{\alpha}}. \end{aligned}$$

By formulas (4.3) and (4.4), if $p(t)$ is of the order of $\frac{1}{\lambda t}$ or less, then by Lemma 7 we have that $\liminf_{\lambda\to\infty} D(\lambda) > 0$. This is the desired result. □

C.2.5 Proof of Theorem 8

First, we show that Assumption 1 (b) implies that $\mu(A(v))$ is continuous in $v \in \mathbb{R}^n$.

By ODE (4.1), $v_i'(t) = \lambda b_i(v(t)) \cdot p(t)$ for each $i \in N$. Since $\mu(A(v))$ is continuous in v,

$$\begin{aligned}\liminf_{\Delta t\to 0} \frac{p(t) - p(t+\Delta t)}{\Delta t} &\leq \sum_{i\in N} \underline{d}_i(v(t))p(t) \cdot v_i'(t) \\ &= \sum_{i\in N} \underline{d}_i(v(t))p(t) \cdot \lambda b_i(v(t))p(t).\end{aligned}$$

By the definition of $\underline{r}$, for all $\varepsilon > 0$, there exists $\bar{t}$ such that for all $t \geq \bar{t}$,

$$\frac{\liminf_{\Delta t\to 0} \frac{p(t)-p(t+\Delta t)}{\Delta t}}{\lambda p(t)^2} \geq \underline{r} - \varepsilon. \tag{C.1}$$

Integrating the both sides and letting $\lambda \to \infty$, we have

$$\limsup_{\lambda\to\infty} p(t) \cdot \lambda t \leq \underline{r}^{-1}. \tag{C.2}$$

An analogous argument shows that

$$\liminf_{\lambda\to\infty} p(t) \cdot \lambda t \geq \overline{r}^{-1}. \tag{C.3}$$

By Lemma 7, we obtain

$$\left(\frac{t}{T}\right)^{1/\underline{r}} \leq \liminf_{\lambda\to\infty} P(t;\lambda) \quad \text{and} \quad \limsup_{\lambda\to\infty} P(t;\lambda) \leq \left(\frac{t}{T}\right)^{1/\bar{r}}, \text{ and}$$
$$\frac{1}{1+\underline{r}^{-1}} \leq \liminf_{\lambda\to\infty} D(\lambda) \quad \text{and} \quad \limsup_{\lambda\to\infty} D(\lambda) \leq \frac{1}{1+\bar{r}^{-1}}.$$

C.2.6 Proof of Theorem 9

Let $r = \bar{r} = \underline{r}$. Then we can follow the discussion in the proof sketch of Theorem 8, and obtain inequality (4.8). Since this inequality holds for any λ, for all $\varepsilon > 0$, there is a large $\bar{t}$ such that for all $t \geq \bar{t}$

$$-(r+\varepsilon)p(t;1)^2 \leq p'(t;1) \leq -(r-\varepsilon)p(t;1)^2.$$

For any small $\eta > 0$, let $\bar{\lambda} = \bar{t}/\eta$. Since $p(t;\lambda) = p(t/\lambda;1)$ for all t and λ, we have

$$-(r+\varepsilon)p(t;\lambda)^2 \leq p'(t;\lambda) \leq -(r-\varepsilon)p(t;\lambda)^2$$

for all $t \geq \eta$, and all $\lambda \geq \bar{\lambda}$. Solving this with an initial condition at η, for $t \geq \eta$ and $\lambda \geq \bar{\lambda}$,

$$\frac{1}{(r-\varepsilon)\lambda(t-\eta)+p(\eta)^{-1}} \leq p(t) \leq \frac{1}{(r+\varepsilon)\lambda(t-\eta)+p(\eta)^{-1}}.$$

By formula (4.3), for $t \geq \eta$ and $\lambda \geq \bar{\lambda}$, we have

$$e^{-\int_t^T \frac{1}{(r+\varepsilon)(s-\eta)+p(\eta)^{-1}/\lambda} ds} \leq P(t) \leq e^{-\int_t^T \frac{1}{(r-\varepsilon)(s-\eta)+p(\eta)^{-1}/\lambda} ds}$$
$$\left(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\right)^{(r-\varepsilon)^{-1}} \leq P(t) \leq \left(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\right)^{(r+\varepsilon)^{-1}}.$$

By formula (4.4), we have

$$
\begin{aligned}
&\int_{\eta}^{T} \Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r-\varepsilon)^{-1}} dt \\
&\quad \leq D(\lambda)T \leq \eta + \int_{\eta}^{T} \Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r+\varepsilon)^{-1}} dt \\
&\frac{1}{1+(r-\varepsilon)^{-1}}\Big(T-\eta+\frac{1}{r\lambda p(\eta)}-\frac{1}{r\lambda p(\eta)(T-\eta)+1}\Big) \\
&\quad \leq D(\lambda)T \leq \frac{1}{1+(r+\varepsilon)^{-1}}\Big(T-\eta+\frac{1}{r\lambda p(\eta)}-\frac{1}{r\lambda p(\eta)(T-\eta)+1}\Big).
\end{aligned}
$$

Since the above inequalities are satisfied for all $\varepsilon > 0$ and $\eta > 0$ in the limit as $\lambda \to \infty$, and $D(\infty) = \frac{1}{1+r^{-1}}$, $|D(\lambda) - D(\infty)| = O(\frac{1}{\lambda})$.

C.2.7 *Proof of Proposition 16*

We prove Proposition 16 in an environment more general than Assumption 3.

Assumption 7. *(a) The limit $v^* = \lim_{\lambda\to\infty} v^*(t)$ is Pareto efficient in X.*

(b) The Pareto frontier of X is smooth in a neighborhood of v^*.

(c) For the unit normal vector $\alpha \in \mathbb{R}^n_+$ at v^*, $\alpha_i > 0$ for all $i \in N$.[2]

(d) For all $\eta > 0$, there exists $\varepsilon > 0$ such that $\{x \in \mathbb{R}^n_+ \mid |v^* - x| \leq \varepsilon, \alpha \cdot (x - v^*) \leq -\eta\}$ is contained in X, where "$\cdot$" denotes the inner product in $\mathbb{R}^n$.

(e) μ has a continuous density function.

Proposition 16′. *Under Assumptions 1, 4, and 7, $\lim_{\lambda\to\infty} D(\lambda) = \dfrac{n^2}{n^2+n+1}$.*

[2] We can show basically the same results without this assumption. We avoid complications derived from the indeterminacy of a normal vector on the boundary of X.

Proof. Let $f_H(t) = \max_{x \in A(t)} f(x)$, and $f_L(t) = \min_{x \in A(t)} f(x)$. Since f is continuous, both $f_H(t)$ and $f_L(t)$ are continuous and converge to $f(v^*)$ as $t \to \infty$. For $\varepsilon > 0$, there is $\bar{t}$ such that $|v^* - v^*(t)| \leq \varepsilon$ for all $t \geq \bar{t}$. For $\eta > 0$, let

$$\underline{A}(t) = \{x \in \mathbb{R}^n_+ \mid x \geq v^*(t), \alpha \cdot (x - v^*) \leq -\eta\},$$
$$\overline{A}(t) = \{x \in \mathbb{R}^n_+ \mid x \geq v^*(t), \alpha \cdot (x - v^*) \leq \eta\}.$$

The volume of $\underline{A}(t)$ (with respect to the Lebesgue measure on $\mathbb{R}^n$) is

$$V(\underline{A}(t)) = \frac{1}{n} \prod_{j \in N} \Big(\frac{\alpha \cdot (v^* - v^*(t)) - \eta}{\alpha_j} \Big), \tag{C.4}$$

and the volume of $\overline{A}(t)$ is

$$V(\overline{A}(t)) = \frac{1}{n} \prod_{j \in N} \Big(\frac{\alpha \cdot (v^* - v^*(t)) + \eta}{\alpha_j} \Big). \tag{C.5}$$

Suppose that $\varepsilon > 0$ is small and $\bar{t}$ is large. Then by Assumption 7, $\underline{A}(t) \subset A(t) \subset \overline{A}(t)$ holds for all $\eta > 0$ and all $t \geq \bar{t}$. The rest of the proof consists of two steps.

Step 1: We show that for any two players $i, j \in N$, $\lim_{t \to \infty} v_j^{*\prime}(t) / v_i^{*\prime}(t) = \alpha_i / \alpha_j$. The ith coordinate of the right hand side of equation (4.1) is bounded as

$$\begin{aligned} &f_L(\bar{t}) \int_{\underline{A}(t)} \big(x_i - v_i^*(t)\big) dx \\ &\quad \leq \int_{A(t)} \big(x_i - v_i^*(t)\big) f(x) dx \leq f_H(\bar{t}) \int_{\overline{A}(t)} \big(x_i - v_i^*(t)\big) dx. \end{aligned}$$

Therefore,

$$\frac{\lambda f_L(\bar{t})V(\underline{A}(t))}{n+1}\Big(\frac{\alpha\cdot(v^*-v^*(t))-\eta}{\alpha_i}\Big)$$
$$\leq v_i^{*\prime}(t)\leq\frac{\lambda f_H(\bar{t})V(\overline{A}(t))}{n+1}\Big(\frac{\alpha\cdot(v^*-v^*(t))+\eta}{\alpha_i}\Big)$$

for all $t\geq\bar{t}$ and $i\in N$. By substituting (C.4) and (C.5),

$$\frac{\lambda f_L(\bar{t})}{n(n+1)}\Big(\frac{\alpha\cdot(v^*-v^*(t))-\eta}{\alpha_i}\Big)\prod_{j\in N}\Big(\frac{\alpha\cdot(v^*-v^*(t))-\eta}{\alpha_j}\Big)$$
$$\leq v_i^{*\prime}(t)\leq\frac{\lambda f_H(\bar{t})}{n(n+1)}\Big(\frac{\alpha\cdot(v^*-v^*(t))+\eta}{\alpha_i}\Big)\prod_{j\in N}\Big(\frac{\alpha\cdot(v^*-v^*(t))+\eta}{\alpha_j}\Big) \tag{C.6}$$

for all $t\geq\bar{t}$ and $i\in N$. By letting $\eta\to 0$, $\varepsilon\to 0$, and $t\to\infty$, we have $\lim_{t\to\infty} v_j^{*\prime}(t)/v_i^{*\prime}(t)=\alpha_i/\alpha_j$.

Step 2: By Step 1, for i and small $\delta>0$, there exists $\tilde{t}\geq\bar{t}$ such that

$$(1-\delta)\frac{\alpha_i}{\alpha_j}\leq\frac{v_j^*-v_j^*(t)}{v_i^*-v_i^*(t)}\leq(1+\delta)\frac{\alpha_i}{\alpha_j}$$

for all $t\geq\tilde{t}$ and $j\in N$. Therefore,

$$n(1-\delta)(v_i^*-v_i^*(t))\leq\frac{\alpha\cdot(v^*-v^*(t))}{\alpha_i}\leq n(1+\delta)(v_i^*-v_i^*(t)).$$

By inequality (C.6), we have

$$\frac{\lambda f_L(\bar{t})}{n(n+1)}\Big(n(1-\delta)(v_i^* - v_i^*(t)) - \frac{\eta}{\alpha_i}\Big)\prod_{j\in N}\Big(n(1-\delta)\frac{\alpha_i}{\alpha_j}(v_i^* - v_i^*(t)) - \frac{\eta}{\alpha_j}\Big)$$
$$\leq v_i^{*\prime}(t) \leq$$
$$\frac{\lambda f_H(\bar{t})}{n(n+1)}\Big(n(1+\delta)(v_i^* - v_i^*(t)) + \frac{\eta}{\alpha_i}\Big)\prod_{j\in N}\Big(n(1+\delta)\frac{\alpha_i}{\alpha_j}(v_i^* - v_i^*(t)) + \frac{\eta}{\alpha_j}\Big)$$

for all $t \geq \tilde{t}$ and $j \in N$. Therefore,

$$\Big(n(1-\delta)(v_i^* - v_i^*(t)) - \frac{\eta}{\alpha_i}\Big)' \leq$$
$$-\frac{\lambda f_L(\bar{t})(1-\delta)}{n+1}\Big(\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)\Big(n(1-\delta)(v_i^* - v_i^*(t)) - \frac{\eta}{\alpha_i}\Big)^{n+1}, \text{ and}$$
$$\Big(n(1+\delta)(v_i^* - v_i^*(t)) + \frac{\eta}{\alpha_i}\Big)' \geq$$
$$-\frac{\lambda f_H(\bar{t})(1+\delta)}{n+1}\Big(\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)\Big(n(1+\delta)(v_i^* - v_i^*(t)) + \frac{\eta}{\alpha_i}\Big)^{n+1}$$

for all $t \geq \tilde{t}$ and $j \in N$. By solving differential equations given by the above inequalities with equality, we have

$$n(1-\delta)(v_i^* - v_i^*(t)) - \frac{\eta}{\alpha_i} \leq \Big(C_L + \frac{\lambda f_L(\bar{t})(1-\delta)nt}{n+1}\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)^{-\frac{1}{n}}$$
$$n(1+\delta)(v_i^* - v_i^*(t)) + \frac{\eta}{\alpha_i} \geq \Big(C_H + \frac{\lambda f_H(\bar{t})(1+\delta)nt}{n+1}\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)^{-\frac{1}{n}}$$

where C_L, C_H are constants determined by the initial condition at $t = \tilde{t}$. Deforming the above inequalities,

$$\frac{1}{n(1+\delta)}\Big(\frac{C_H}{\lambda t}+\frac{f_H(\bar{t})(1+\delta)n}{n+1}\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)^{-\frac{1}{n}}-\frac{(\lambda t)^{\frac{1}{n}}\eta}{n(1+\delta)\alpha_i}$$
$$\leq (v_i^*-v_i^*(t))(\lambda t)^{\frac{1}{n}}\leq \frac{1}{n(1-\delta)}\Big(\frac{C_L}{\lambda t}+\frac{f_L(\bar{t})(1-\delta)n}{n+1}\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)^{-\frac{1}{n}}+\frac{(\lambda t)^{\frac{1}{n}}\eta}{n(1-\delta)\alpha_i}.$$

As $\eta\to 0$, $\bar{t}\to\infty$, $\delta\to 0$, and $t\to\infty$, we have

$$\lim_{t\to\infty}(v_i^*-v_i^*(t))(\lambda t)^{\frac{1}{n}}=\frac{1}{n}\Big(\frac{f(v^*)n}{n+1}\prod_{j\neq i}\frac{\alpha_i}{\alpha_j}\Big)^{-\frac{1}{n}}$$
$$\lim_{t\to\infty}\alpha_i(v_i^*-v_i^*(t))(\lambda t)^{\frac{1}{n}}=\Big(\frac{n+1}{f(v^*)n^{n+1}}\prod_{j\in N}\alpha_j\Big)^{\frac{1}{n}},$$

which is a positive constant.

Step 3: By the definition of $\underline{A}(t),\overline{A}(t)$,

$$f_L(t)V(\underline{A}(t))\leq p(t)\leq f_H(t)V(\overline{A}(t)).$$

Inequalities (C.4), (C.5) implies that by letting $\eta\to 0$ we have

$$\lim_{t\to\infty}p(t)\cdot\lambda t=\lim_{t\to\infty}\frac{f(v^*)}{n}\prod_{j\in N}\Big(\frac{\sum_{i\in N}\alpha_i(v_i^*-v_i^*(t))(\lambda t)^{\frac{1}{n}}}{\alpha_j}\Big).$$

By the result of Step 2, this limit exists and computed as follows:

$$\lim_{t\to\infty}p(t)\cdot\lambda t=\frac{f(v^*)}{n}\prod_{j\in N}\left(\frac{n}{\alpha_j}\Big(\frac{n+1}{f(v^*)n^{n+1}}\prod_{k\in N}\alpha_k\Big)^{\frac{1}{n}}\right)$$
$$=\frac{n+1}{n^2}.$$

By Lemma 7, the limit expected duration is

$$D(\infty) = \frac{1}{1+\frac{n+1}{n^2}} = \frac{n^2}{n^2+n+1}.$$

□

C.2.8 Proof of Proposition 18

Let $v^*(t; f)$ be the solution of ODE (4.1) for density $f \in \mathcal{F}$, and $v^*(f) = \lim_{\lambda\to\infty} v^*(t; f) = \lim_{t\to\infty} v^*(t; f)$.

First we show that the set is open, i.e., for all $f \in \mathcal{F}$ with $v^*(f)$ Pareto efficient, $\varepsilon > 0$, and a sequence $f_k \in \mathcal{F}$ $(k = 1, 2, \ldots)$ with $|f_k - f| \to 0$ $(k \to \infty)$, there exist $\delta > 0$ and $\bar{k}$ such that

$$|v^*(f_k) - v^*(f)| \leq \varepsilon$$

for all $k \geq \bar{k}$.

Since $\lim_{t\to\infty} v^*(t; f) = v^*(f)$, for all $\delta > 0$ there exists $\bar{t} > 0$ such that $|v^*(f) - v^*(t; f)| \leq \delta$ for all $t \geq \bar{t}$. By Pareto efficiency of $v^*(f)$, let $\delta > 0$ be sufficiently small so that $A\big(v^*(\bar{t}; f) - (\delta, \delta, \ldots, \delta)\big)$ is contained in the ε-ball centered at $v^*(f)$. Since the right hand side of ODE (4.1) is continuous in v by Assumption 4, the unique solution of (4.1) is continuous with respect to parameters in (4.1). Therefore, for a finite time interval $[0, T]$ including $\bar{t}$, there exists $\bar{k}$ such that $|v^*(t; f_k) - v^*(t; f)| \leq \delta$ for all $t \in [0, T]$ and all $k \geq \bar{k}$. This implies that $v^*(t; f_k) \in A\big(v^*(\bar{t}; f) - (\delta, \delta, \ldots, \delta)\big)$, thereby $v^*(f_k) \in A\big(v^*(\bar{t}; f) - (\delta, \delta, \ldots, \delta)\big)$. Hence we have $|v^*(f_k) - v^*(f)| \leq \varepsilon$.

Second we show that the set is dense, i.e., for all $f \in \mathcal{F}$ with $v^*(f)$ not strictly Pareto efficient in X and all $\varepsilon > 0$, there exists $\tilde{f} \in \mathcal{F}$ such that $|f - \tilde{f}| \leq \varepsilon$ and

$v^*(\tilde{f})$ is Pareto efficient. Since $v^*(f)$ is only weakly Pareto efficient in $\hat{X}$, there exists Pareto efficient $y \in X$ which Pareto dominates $v^*(f)$. Let $I = \{i \in N \mid y_i = v_i^*(f)\}$ and $J = N \setminus I$. Since y is Pareto efficient, there is $\delta > 0$ such that if $x \in X$ is weakly Pareto efficient, satisfies $|y - x| \leq \delta$, and $y_i = x_i$ for some $i \in N$, then there is no $\tilde{x} \in X$ such that $\tilde{x}_i > y_i$ and $|y - \tilde{x}| \leq \delta$.

By Assumption 4, for any small $\delta/2 > \eta > 0$, there is a small ball contained in X centered at $\tilde{y}$ with $|y - \tilde{y}| \leq \eta$. Let g be a continuous density function whose support is the above small ball, takes zero on the boundary of the ball, and the expectation of g is exactly $\tilde{y}$. Let $\tilde{f} = (1 - \frac{\varepsilon}{|f|+|g|})f + \frac{\varepsilon}{|f|+|g|}g \in \mathcal{F}$. Since f and g are bounded from above, $|f - \tilde{f}| \leq \varepsilon$.

Since $v^*(f)$ is weakly Pareto efficient, if $v^*(f) \in A(v)$, then $A(v) \subseteq \bigcup_{i \in N} \big([v_i, v_i^*(f)] \times \prod_{j \neq i}[0, \overline{x}_j]\big)$. If $|v^*(f) - v| \leq \xi$ where $\xi > 0$ is very small,

$$\begin{aligned}\int_{A(v)} (x_i - v_i) f(x) dx &\leq f_H \sum_{j \in N} (v_j^*(f) - v_j) \prod_{k \in N} \overline{x}_k \\ &\leq \xi n f_H \prod_{k \in N} \overline{x}_k\end{aligned}$$

If $v^*(f) \in A(v)$, $\min_{j \in N}(y_j - v_j) \geq 2\eta$ and $|v^*(f) - v| \leq \xi$, we have

$$\begin{aligned}\int_{A(v)} (x_i - v_i)\tilde{f}(x)dx - \int_{A(v)} (x_i - v_i) f(x) dx &= \int_{A(v)} (x_i - v_i)(\tilde{f}(x) - f(x))dx \\ &= \frac{\varepsilon}{|f| + |g|} \int_{A(v)} (x_i - v_i)(g(x) - f(x))dx \\ &\geq \frac{\varepsilon}{|f| + |g|}\Big((\tilde{y}_i - v_i) - \Big(\xi n f_H \prod_{k \in N} \overline{x}_k\Big)\Big).\end{aligned}$$

If $j \in J$ and $|v^*(f) - v| \leq \xi$ where $\xi > 0$ is very small, then

$$\int_{A(v)} (x_j - v_j)\tilde{f}(x)dx - \int_{A(v)} (x_j - v_j) f(x) dx \geq \frac{\varepsilon}{2(|f| + |g|)}(\tilde{y}_j - v_j^*(f)).$$

Let $w(t) = v^*(t;\tilde{f}) - v^*(t;f)$. Since ODE (4.1) is continuous in the parameters, for all $\zeta > 0$, there exists $\varepsilon > 0$ such that $|w(t)| \leq \zeta$ for all $t \in [0, T]$. Suppose that T and t are very large so that $|v^*(f) - v^*(t;f)| \leq \xi$. For $j \in J$, $w_j'(t)$ is estimated as follows:

$$\begin{aligned}
w_j'(t) &= \lambda \int_{A(v^*(t;\tilde{f}))} (x_j - v_j^*(t;\tilde{f}))\tilde{f}(x)dx - \lambda \int_{A(v^*(t;f))} (x_j - v_j^*(t;f))f(x)dx \\
&= \lambda \int_{A(v^*(t;\tilde{f}))} (x_j - v_j^*(t;\tilde{f}))\tilde{f}(x)dx - \lambda \int_{A(v_j^*(t;\tilde{f})))} (x_j - v_j^*(t;\tilde{f})))f(x)dx \\
&\quad + \lambda \int_{A(v^*(t;\tilde{f}))} (x_j - v_j^*(t;\tilde{f}))f(x)dx - \lambda \int_{A(v^*(t;f))} (x_j - v_j^*(t;f))f(x)dx \\
&\geq \frac{\lambda\varepsilon}{2(|f| + |g|)}(\tilde{y}_j - v_j^*(f)) - \lambda \int_{A(v^*(t;f)) \cap A(v^*(t;\tilde{f}))} w_j(t)f(x)dx \\
&\quad - \lambda \int_{A(v^*(t;f)) \setminus (A(v^*(t;f)) \cap A(v^*(t;\tilde{f})))} (x_j - v_j^*(t;f))f(x)dx \\
&\geq \frac{\lambda\varepsilon}{2(|f| + |g|)}(\tilde{y}_j - v_j^*(f) - \zeta) - \lambda\zeta\xi \sum_{k \in N} \prod_{l \neq k} \overline{x}_l - \lambda\xi n f_H \prod_{k \in N} \overline{x}_k.
\end{aligned}$$

Therefore when $\xi > 0$ is sufficiently small, $w_j'(t)$ is bounded away from zero:

$$w_j'(t) \geq \frac{\lambda\varepsilon}{4(|f| + |g|)}(\tilde{y}_j - v_j^*(f) - \zeta).$$

This implies that for small $\varepsilon > 0$ and large t, $v_j^*(t;\tilde{f}) > v^*(f)$ for all $j \in J$. Then the similar method to Step 3 in the proof of Proposition 19 shows that $v^*(t;\tilde{f})$ converges to a Pareto efficient allocation in X.

C.2.9 *Proof of Proposition 19*

Let $f_L = \inf_{x \in X} f(x) > 0$, $f_H = \sup_{x \in X} f(x)$, and $\overline{x}_i = \max\{x_i \mid x \in X\}$ for $i \in N$. Assumption 4 ensures existence of these values. Let $A = \{x \in X \mid x \geq v^*\}$, $I = \{i \in N \mid x_i = v_i^* \text{ for all } x \in A\} \subseteq N$, and $J = N \setminus I$. Suppose that there exists

$x \in X$ which Pareto dominates v^*, thereby $J \neq \emptyset$.

Step 1: We show that I is nonempty. If there is no such player, there exist $y(1), \ldots, y(n)$ such that $y(j) \in A$ and $y_j(j) > v_j^*$ for all $j \in N$. This implies that $y = \frac{1}{n} \sum_{j \in N} y(j)$ strictly Pareto dominates v^*. Since X is convex, y also belongs to A. This contradicts the weak Pareto efficiency of v^* shown in Proposition 15.

Step 2: Next we show that if v^* is not Pareto efficient in X, and $i \in I$, then $x_i \leq v_i^*$ for all $x \in X$.

Let i be the player in I. Suppose that there exists $y \in X$ with $y_i > v_i^*$. Since X is convex, $\alpha y + (1-\alpha)x \in X$ for all $0 \leq \alpha \leq 1$ and $x \in X$. Since we assumed that there exists $x \in X$ which Pareto dominates v^*, $x_j > v_j^*$ for $j \in J$. Then there exists $\alpha > 0$ such that $\alpha y + (1-\alpha)x \geq v^*$, and $\alpha y_j + (1-\alpha)x_j > v_j^*$ for some j. By Step 1, we must have $x_i = v_i^*$. Therefore, $\alpha y_i + (1-\alpha)x_i > v_i^*$, which contradicts the fact that $i \in I$.

Step 3: Finally we show that $v^*(t)$ converges to a Pareto efficient allocation in X as $t \to \infty$.

By convexity of X, one can find $y_j, \bar{y}_j$ ($j \in J$) such that $v_j^* < y_j < \bar{y}_j$, and $\prod_{i \in I}[v_i^* - \varepsilon, v_i^*] \times \prod_{j \in J}[y_j, \bar{y}_j]$ is contained in X for small $\varepsilon > 0$. Let $\varepsilon \in (0, 1/2)$ be sufficiently small such that $\varepsilon \leq \dfrac{2 f_L \prod_{j \in J}(\bar{y}_j - y_j)}{f_H \prod_{j \in J} \overline{x}_j}$. Since $v^*(t)$ converges to v^* as $t \to \infty$, there exists $\bar{t}$ such that $\max_{i \in N}\{v_i^* - v_i^*(t)\} \leq \varepsilon$ whenever $t \geq \bar{t}$. Let $Y(t) = \prod_{i \in I}[v_i^*(t), v_i^*] \times \prod_{j \in J}[y_j, \bar{y}_j] \subseteq A(t)$.

We have $A(t) \subseteq \prod_{i \in I}[v_i^*(t), v_i^*] \times \prod_{j \in J}[0, \overline{x}_j]$ since there is no $x \in A(t)$ with

$x_i > v_i^*$. By equation (4.1), for $i \in I$,

$$\begin{aligned}
v_i^{*\prime}(t) &= \lambda \int_{A(t)} \big(x_i - v_i^*(t)\big) d\mu \\
&\leq \lambda \int_{\prod_{i' \in I}[v_{i'}^*(t), v_{i'}^*]} \big(x_i - v_i^*(t)\big) \int_{\prod_{j \in J}[0, \overline{x}_j]} f_H \prod_{j \in J} dv_j \prod_{i' \in I} dv_{i'} \\
&\leq \frac{1}{2} \lambda f_H (v_i^* - v_i^*(\bar{t})) \prod_{i' \in I} (v_{i'}^* - v_{i'}^*(t)) \prod_{j \in J} \overline{x}_j
\end{aligned}$$

for all $t \geq \bar{t}$. On the other hand, for $j \in J$,

$$\begin{aligned}
v_j^{*\prime}(t) &= \lambda \int_{A(t)} \big(x_j - v_j^*(t)\big) d\mu \\
&\geq \lambda \int_{Y(t)} (y_j - v_j^*) d\mu \\
&= \lambda (y_j - v_j^*) \mu(Y(t)) \\
&\geq \lambda f_L (y_j - v_j^*) \prod_{i \in I} (v_i^* - v_i^*(t)) \prod_{j' \in J} (\bar{y}_{j'} - y_{j'}).
\end{aligned}$$

Then for $i \in I$ and $j \in J$,

$$\begin{aligned}
\frac{v_i^{*\prime}(t)}{v_j^{*\prime}(t)} \cdot \frac{v_j^* - v_j^*(\bar{t})}{v_i^* - v_i^*(\bar{t})} &\leq \frac{f_H (v_i^* - v_i^*(\bar{t}))(v_j^* - v_j^*(\bar{t})) \prod_{j \in J} \overline{x}_j}{2 f_L \prod_{j' \in J} (\bar{y}_{j'} - y_{j'})} \\
&\leq \frac{(v_i^* - v_i^*(\bar{t}))(v_j^* - v_j^*(\bar{t}))}{\varepsilon} \\
&\leq \varepsilon \leq \frac{1}{2}
\end{aligned}$$

for all $t \geq \bar{t}$. Therefore,

$$\frac{v_i^{*\prime}(\bar{t})}{v_j^{*\prime}(\bar{t})} \leq \frac{v_i^* - v_i^*(\bar{t})}{2\big(v_j^* - v_j^*(\bar{t})\big)}$$

holds for all $t \geq \bar{t}$. This inequality implies

$$v_i^*(t) - v_i^*(\bar{t}) \leq \frac{v_i^* - v_i^*(\bar{t})}{2\big(v_j^* - v_j^*(\bar{t})\big)}\big(v_j^*(t) - v_j^*(\bar{t})\big)$$

for all $t \geq \bar{t}$. By letting $t \to \infty$ in the above inequality, we have $0 < v_i^* - v_i^*(\bar{t}) \leq \big(v_i^* - v_i^*(\bar{t})\big)/2$, a contradiction. Hence v^* is strictly Pareto efficient in X.

C.2.10 *Proof of Proposition 21*

First, we define the notion of the edge of the Pareto frontier. Suppose that w is Pareto efficient in X, and $w_i > 0$ for all $i \in X$. Let us denote an $(n-1)$-dimensional subspace orthogonal to w by $D = \{z \in \mathbb{R}^n \mid w \cdot z = 0\}$. For $\xi > 0$, let D_ξ be an $(n-1)$-dimensional disk defined as

$$D_\xi = \{z \in D \mid |z| \leq \xi\},$$

and let S_ξ be its boundary. We say that a Pareto efficient allocation w in X is *not* located at the edge of the Pareto frontier of X if there is $\xi > 0$ such that for all vector $z \in D_\xi$ there is a scalar $\alpha > 0$ such that $\alpha(w+z)$ is Pareto efficient in X. We denote this Pareto efficient allocation by $w_z \in X$.

Let $B_\varepsilon(y) = \{x \in X \mid |w - x| \leq \varepsilon\}$ for $y \in X$ and $\varepsilon > 0$. We denote the volume of $B_\varepsilon(y)$ by $V_\varepsilon(y)$, and the volume of the n-dimensional ball with radius ε by V_ε. Note that $\min_{y \in X} V_\varepsilon(y) > 0$ by Assumption 4. Let g be a continuous density function on an n-dimensional ball centered at $0 \in \mathbb{R}^n$ with radius ε, assumed to take zero on the boundary of the ball. Let $\tilde{f}$ be the uniform density function on X. For a Pareto efficient allocation y, we define a probability density function f_y on X

by

$$f_y(x) = \eta \tilde{f}(x) + (1-\eta) g(y-x) \frac{V_\varepsilon}{V_\varepsilon(y)}$$

where $\eta > 0$ is small. Note that $f_y(x)$ is uniformly bounded above and away from zero in x and y.

For $z \in D_\xi$, let $\tilde{\varphi}(z)$ be the limit of the solution of ODE (4.1) with density f_{w_z}, and define a function φ from D_ξ to D by $\varphi(z) = \tilde{\varphi}(z) + \delta w \in D$ for some $\delta \in \mathbb{R}$. By the form of ODE (4.1), the solution of (4.1) with density f_{w_z} is continuously deformed if z changes continuously. Since w is not at the edge of the Pareto frontier, $\tilde{\varphi}(z)$ is also Pareto efficient in X and comes close to w if ξ, ε, and η are small. Therefore $\varphi(z)$ is a continuous function. The rest of the proof consists of two steps.

Step 1: We show that for any $\xi > 0$, there exist $\varepsilon > 0$ and $\eta > 0$ such that $|\varphi(z) - z| \leq \xi$ for all $z \in D_\xi$. If a density function has a positive value only in $B_\varepsilon(y)$ for some y in the Pareto frontier of X, then the barycenter of $A(t)$ is always contained in $B_\varepsilon(y)$. In such a case, the limit allocation with density f_y belongs to $B_\varepsilon(y)$. As $\eta \to 0$, f_y approaches the above situation. Therefore, for sufficiently small $\eta > 0$, the distance between the limit allocation and y is smaller than 2ε. For $y = w_z$ and letting ε very small, we have $|\varphi(z) - z| \leq \xi$. Since D_ξ is compact, such we can take such small $\varepsilon > 0$ and $\eta \to 0$ uniformly.

Step 2: We show that there is $z \in D_\xi$ such that $\varphi(z) = 0$. Let $\psi(z) = z - \varphi(z)$. By Step 1, $\psi(z)$ belongs to D_ξ for all $z \in D_\xi$. By Brouwer's fixed point theorem, there exists $z \in D_\xi$ such that $\psi(z) = z$. Therefore there exists $z \in D_\xi$ such that $\varphi(z) = 0$.

Hence for $z \in D_\xi$ such that $\varphi(z) = 0$, the limit allocation with density f_{w_z} coincides with w.

C.2.11 *Proof of Proposition 22*

Let $v^0(t;\lambda)$ be the solution of (4.1) for $\rho = 0$. Fix any $t \in [0, T]$. Recall that $v^0(t;\alpha\lambda) = v^0(\alpha t;\lambda)$ for all $\alpha > 0$. Since we defined as $\lim_{\lambda\to\infty} v^0(t;\lambda) = v^*(t;0,\infty)$, there exists $\bar{\lambda}^1 > 0$ such that

$$\begin{aligned}\left|v^*(t;0,\infty) - v^0(t;\lambda)\right| &= \left|v^*(t;0,\infty) - v^0(\lambda t;1)\right| \\ &\leq \varepsilon/2 \end{aligned} \tag{C.7}$$

for all $\lambda \geq \bar{\lambda}^1$.

Since the right hand side of ODE (4.10) is continuous in ρ, λ, and uniformly Lipschitz continuous in v, the unique solution $v^*(t;\rho,\lambda)$ is continuous in ρ, λ for all $t \in [0, T]$. Recall that $v^*(t;\rho,\alpha\lambda) = v^*(\alpha t;\rho/\alpha,\lambda)$ for all $\alpha > 0$. Therefore by continuity in ρ, there exists $\bar{\lambda}^2 > 0$ such that

$$\begin{aligned}\left|v^*(t;\rho,\lambda) - v^0(t;\lambda)\right| &= \left|v^*(\lambda t;\rho/\lambda,1) - v^0(\lambda t;1)\right| \\ &\leq \varepsilon/2 \end{aligned} \tag{C.8}$$

for all $\lambda \geq \bar{\lambda}^2$. By adding (C.7) and (C.8), we obtain the desired inequality for $\bar{\lambda} = \max\{\bar{\lambda}^1, \bar{\lambda}^2\}$.

C.2.12 *Proof of Proposition 23*

Let $v(t)$ be the solution of ODE (4.10). The proof consists of five steps.

Step 1: We show that for any $t > 0$, $\mu(A(t)) \to 0$ as $\lambda \to \infty$. If not, there exist a positive value $\varepsilon > 0$ and an increasing sequence $(\bar{\lambda}_k)_{k=1,2,...}$ such that $\mu(A(t)) \geq \varepsilon$ for all $\bar{\lambda}_k$. Since X is compact and f is bounded from above, there exists $\eta > 0$ such that $\mu\big(A(v(t)+(\eta,\ldots,\eta))\big) \geq \varepsilon/2$. In fact, since

$$\begin{aligned}\mu\big(A(v(t)) \setminus A(v(t)+(\eta,\ldots,\eta))\big) &\leq \sum_{i\in N} \mu\Big([v_i(t), v_i(t)+\eta] \times \prod_{j\neq i}[0,\overline{x}_j]\Big) \\ &\leq f_H \sum_{i\in N} \eta \prod_{j\neq i} \overline{x}_j,\end{aligned}$$

we have $\mu\big(A(v(t)+(\eta,\ldots,\eta))\big) \geq \varepsilon/2$ for $\eta = \dfrac{\varepsilon}{2f_H \sum_{i\in N}\prod_{j\neq i}\overline{x}_j}$. For this η, the integral in ODE (4.10) is estimated as

$$\begin{aligned}\int_{A(t)} \big(x_i - v_i(t)\big) d\mu &\geq \int_{A(v(t)+(\eta,\ldots,\eta))} \big(x_i - v_i(t)\big) d\mu \\ &\geq \int_{A(v(t)+(\eta,\ldots,\eta))} \eta d\mu \\ &\geq \eta\varepsilon/2.\end{aligned}$$

By ODE (4.10),

$$v_i'(t) \geq -\rho\overline{x}_i + \bar{\lambda}_k \eta\varepsilon/2,$$

which obviously grows infinitely as $\bar{\lambda}_k$ becomes large. This contradicts compactness of X.

Step 2: We compute the direction of $\int_{A(t)}\big(x_i - v_i(t)\big)d\mu$ in the limit as $\lambda \to \infty$. By Step 1, the boundary of X contains all accumulation points of $\{v_i(t) \mid \lambda > 0\}$ for fixed $t > 0$. Fix an accumulation point $v^*(t)$. There exists an increasing sequence $(\lambda_k)_{k=1,2,\ldots}$ with $v^*(t) = \lim_{k\to\infty} v(t)$. By Assumption 5, there exists a unit normal vector of X at $v^*(t)$, which we denote by $\alpha \in \mathbb{R}_{++}$.

Step 1 implies that $v(t)$ is very close to the boundary of X when λ_k is very large. By smoothness of the boundary of X, $A(t)$ looks like a polyhedron defined by convex hull of $\{v(t), v(t) + (z_1(t), 0, \ldots, 0), v(t) + (0, z_2(t), 0, \ldots, 0), \ldots, v(t) + (0, \ldots, 0, z_n(t))\}$ where $z_i(t)$'s are positive length of edges such that the last n vertices are on the boundary of X. This vector $z(t)$ is parallel to $(1/\alpha_1, \ldots, 1/\alpha_n)$. Let $r(t)$ be the ratio between the length of $z(t)$ and $(1/\alpha_1, \ldots, 1/\alpha_n)$, i.e., $r(t) = z_1(t)\alpha_1 = \cdots = z_n(t)\alpha_n$.

Since density f is bounded from above and away from zero, distribution μ looks almost uniform on $A(t)$ if λ_k is large. Then the integral $\int_{A(t)} (x_i - v_i(t)) d\mu$ is almost parallel to the vector from $v(t)$ to the barycenter of the polyhedron, namely, $z(t)/(n+1)$. Therefore $\int_{A(t)} (x_i - v_i(t)) d\mu$ is approximately parallel to $(1/\alpha_1, \ldots, 1/\alpha_n)$ when λ_k is large.

Step 3: We show that $\sum_{i \in N} \alpha_i v_i'(t) \geq 0$ for large λ. Let $(\lambda_k)_{k=1,2,\ldots}$ be the sequence defined in Step 2. For large λ_k, $A(t)$ again looks like a polyhedron with the uniform distribution. By Step 2, the ODE near $v_i(t)$ is written as

$$v_i'(t) = -\rho v_i(t) + \lambda_k \frac{z_i(t)}{n+1} \cdot \mu(A(t)). \tag{C.9}$$

Note that $v_i(t)$ is close to $v_i^*(t)$ and $\mu(A(t))$ is order n of the length of $z(t)$. By replacing the above equation by $r(t)$, ODE (C.9) approximates

$$r'(t) = \rho a - \lambda_k b r(t)^{n+1} \tag{C.10}$$

for some constants $a, b > 0$. Since $r(t)$ is large when t is small, the above ODE shows that $r(t)$ is decreasing in t. Therefore $\mu(A(t))$ is also decreasing in t. For large λ_k, this implies that

$$\alpha \cdot v'(t) = \sum_{i \in N} \alpha_i v_i'(t) \geq 0.$$

Step 4: We show that the Nash product is nondecreasing if λ is large. By ODE (C.9), we have

$$\alpha_i v_i'(t) = -\rho\alpha_i v_i(t) + \beta \tag{C.11}$$

where $\beta = \lambda_k \mu(A(t))/(n+1)$ independent of i. Let us assume without loss of generality that $\alpha_1 v_1'(t) \geq \cdots \geq \alpha_n v_n'(t)$. Then we must have $1/\alpha_1 v_1(t) \geq \cdots \geq 1/\alpha_n v_n(t)$.

Let $L(t) = \sum_{i\in N} \ln v_i(t)$ be a logarithm of the Nash product. Then $L'(t) = \sum_{i\in N} v_i'(t)/v_i(t)$. By Chebyshev's sum inequality,

$$\begin{aligned} L'(t) &= \sum_{i\in N} \frac{v_i'(t)}{v_i(t)} \\ &\geq \frac{1}{n}\Big(\sum_{i\in N} \alpha_i v_i'(t)\Big)\Big(\sum_{i\in N} \frac{1}{\alpha_i v_i(t)}\Big) \geq 0. \end{aligned}$$

Hence, $L(t)$ is nondecreasing if λ_k is large. Moreover, equality holds if and only if $\alpha_1 v_1'(t) = \cdots = \alpha_n v_n'(t)$ or $\alpha_1 v_1(t) = \cdots = \alpha_n v_n(t)$.

Step 5: We show that $v(t)$ converges to a point in the Nash set as $\lambda \to \infty$. Step 4 shows that $L'(t)$ converges to zero as $\lambda \to \infty$. Then $\alpha_1 v_1'(t) = \cdots = \alpha_n v_n'(t)$ or $\alpha_1 v_1(t) = \cdots = \alpha_n v_n(t)$ in the limit of $\lambda \to \infty$. The former case implies $v_i'(t) = 0$ for all $i \in N$ by Step 3. Then ODE (C.11) shows that the latter case holds. Therefore the latter case always holds in the limit of $\lambda \to \infty$. This implies that the boundary of X at $v^*(t)$ is tangent to the hypersurface defined by "Nash product $= \prod_{i\in N} v_i^*(t)$." Hence any accumulation point $v^*(t)$ belongs to the Nash set.

Since we assumed that the Nash set consists of isolated points, $v^*(t)$ is isolated. If $v(t)$ does not converge to $v^*(t)$, there is $\delta > 0$ such that for any $\bar{\lambda}$ there exists $v(t)$ with $\lambda \geq \bar{\lambda}$. Let $\delta > 0$ be small such that there is no point in the Nash set in

$\{x \in X \mid |v^*(t) - x| \leq \delta\}$. Since $v(t)$ is continuous with respect to λ, for any $\bar{\lambda}$, there exists $\lambda > \bar{\lambda}$ such that $\delta/2 \leq |v^*(t) - v(t)| \leq \delta$. Since $\{x \in X \mid \delta/2 \leq |v^*(t) - x| \leq \delta\}$ is compact, $v(t)$ must have an accumulation point in this set. This contradicts the fact that any accumulation point is contained in the Nash set. Furthermore, $v^*(t)$ does not depend on t since $v^*(t)$ is continuous in t.

C.2.13 *Proof of Proposition 24*

(Sketch of proof): The ODE (C.10) is approximated by a linear ODE, which has a solution converging to v^* with an exponential speed. Therefore for large λ, $r(t)$ is approximated by $r(t) = \left(\frac{\rho a}{\lambda b}\right)^{\frac{1}{n+1}}$. Since $\mu(A(t))$ is proportional to $r(t)^n$, $\mu(A(t)) = c\lambda^{-\frac{n}{n+1}}$ for a constant $c > 0$. the probability that players reach an agreement before time $-(T-s)$ is

$$1 - e^{-\int_{T-s}^{T} \mu(A(t))\lambda dt} = 1 - e^{-sc\lambda^{\frac{1}{n+1}}},$$

which converges to one as $\lambda \to \infty$.

C.2.14 *Proof of Proposition 25*

By equation (4.12), $v_i(\frac{t}{\Delta t})$ is a nondecreasing sequence. Since X is bounded and convex, $v_i(\frac{t}{\Delta t})$ converges to a Pareto efficient allocations as $\Delta t \to 0$. Let $v^*(\frac{t}{\Delta t})$ be the solution of equation (4.12), and $v^* = \lim_{\Delta t \to 0} v^*(\frac{t}{\Delta t})$ for $t > 0$.

The proof proceeds basically on the same route as that in Proposition 16. Let $f_h(\frac{t}{\Delta t}), f_L(\frac{t}{\Delta t})$, be defined as in the proof of Proposition 16. Then a parallel argument to Step 1 shows an inequality analogous to (C.6): For large $\bar{t}$,

$$\frac{\pi(\Delta t) f_L(\frac{\bar{t}}{\Delta t})}{n(n+1)} \Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta t})) - \eta}{\alpha_i}\Big) \prod_{j \in N} \Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta t})) - \eta}{\alpha_j}\Big)$$
$$\leq v_i^*\Big(\frac{t}{\Delta t} + 1\Big) - v_i^*\Big(\frac{t}{\Delta t}\Big) \leq$$
$$\frac{\pi(\Delta t) f_H(\frac{\bar{t}}{\Delta t})}{n(n+1)} \Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta t})) + \eta}{\alpha_i}\Big) \prod_{j \in N} \Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta t})) + \eta}{\alpha_j}\Big),$$

where notations are the same as in the proof of Proposition 16. Therefore we have $\lim_{\Delta t \to 0} \frac{v_j^*(\frac{t}{\Delta t} + 1) - v_j^*(\frac{t}{\Delta t})}{v_i^*(\frac{t}{\Delta t} + 1) - v_i^*(\frac{t}{\Delta t})} = \frac{\alpha_i}{\alpha_j}$. Similar computations as in Step 2 show an approximation for small Δt

$$\Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t} + 1\Big)\Big) - \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t}\Big)\Big) \approx$$
$$- \frac{\pi(\Delta t) f_m(\frac{\bar{t}}{\Delta t})}{n(n+1)} \Big(\prod_{j \neq i} \frac{\alpha_i}{\alpha_j}\Big) \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t}\Big)\Big)^{n+1},$$

where $f_m(\frac{t}{\Delta t})$ is the average density in $A(\frac{t}{\Delta t})$. Then we can show that

$$\lim_{\Delta t \to 0} \alpha_i \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t}\Big)\Big) \cdot \Big(\frac{\pi(\Delta t) t}{\Delta t}\Big)^{\frac{1}{n}} = \Big(\frac{n+1}{f(v^*) n^{n+1}} \prod_{j \in N} \alpha_j\Big)^{\frac{1}{n}}.$$

Here we used the fact that Δt is very small when compared to $\pi(\Delta t)$ if Δt is small, to ignore the constant derived from an initial condition.

A similar computation to Step 3 shows that

$$\lim_{\Delta t \to 0} p(t) \cdot \Big(\frac{\pi(\Delta t) t}{\Delta t}\Big) = \frac{n+1}{n^2},$$

and thus the limit expected duration is $D(\infty) = \frac{n^2}{n^2+n+1}$.

C.2.15 *Proof of Proposition 26*

(Sketch of proof): The approximated ODE (C.10) for large t in the proof of Proposition 23 is rearranged as follows:

$$\lambda^{\frac{1}{n}} r'(t) = \lambda^{\frac{1}{n}} \rho a - b \cdot \left(\lambda^{\frac{1}{n}} r(t)\right)^{n+1}$$

If $\lambda^{\frac{1}{n}} \rho \to 0$, this ODE is approximated as

$$\lambda^{\frac{1}{n}} r'(t) \approx -b\left(\lambda^{\frac{1}{n}} r(t)\right)^{n+1}$$

which yields $r(t) = O\left(\frac{1}{(\lambda t)^{\frac{1}{n}}}\right)$. This is the same case with $\rho = 0$. On the other hand, If $\lambda^{\frac{1}{n}} \rho \to \infty$, the ODE is approximated as

$$\begin{aligned}
\lambda^{\frac{1}{n}} r'(t) &= \left((\lambda^{\frac{1}{n}} \rho a)^{\frac{1}{n+1}} - b^{\frac{1}{n+1}} (\lambda^{\frac{1}{n}} r(t))\right) \\
&\quad \cdot \left((\lambda^{\frac{1}{n}} \rho a)^{\frac{n}{n+1}} + (\lambda^{\frac{1}{n}} \rho a)^{\frac{n-1}{n+1}} \cdot b^{\frac{1}{n+1}} (\lambda^{\frac{1}{n}} r(t)) + \cdots + b^{\frac{n}{n+1}} (\lambda^{\frac{1}{n}} r(t))^n\right) \\
&\approx (\lambda^{\frac{1}{n}} \rho a) - (\lambda^{\frac{1}{n}} \rho a)^{\frac{n}{n+1}} \cdot b^{\frac{1}{n+1}} (\lambda^{\frac{1}{n}} r(t))
\end{aligned}$$

which implies $r(t) = \left(\frac{\rho a}{\lambda b}\right)^{\frac{1}{n+1}} - O(e^{-t})$. This corresponds to the case with $\rho > 0$.

C.2.16 *Proof of Proposition 32*

By symmetry, $v_1^*(t) = v_2^*(t)$ and $v^* = (1/2, 1/2)$. Let $z(t) = v_i^* - v_i^*(t)$. Suppose that t is large and $z(t)$ is small, so that $z(t) \leq \frac{1-a}{2(1+a)}$. It is straightforward to see that an agreement is reached after negotiation with a costly transfer if and only if realized allocation $x \in X$ is in the triangle $T_1 \cup T_2 \cup T_3$ shown in Figure C.1, where the slopes of the line segments are $-a, a, 1/a, -1/a$, respectively, from southeast to northwest.

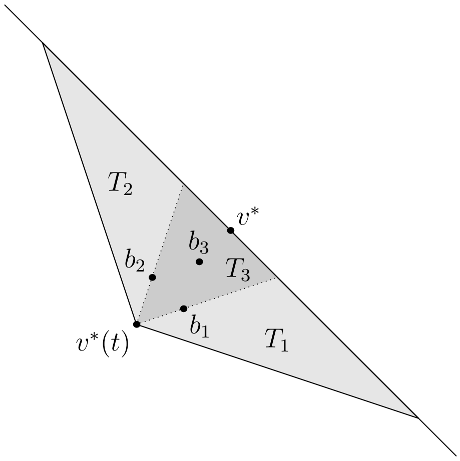

Figure C.1: The set of realized allocations that the players accept

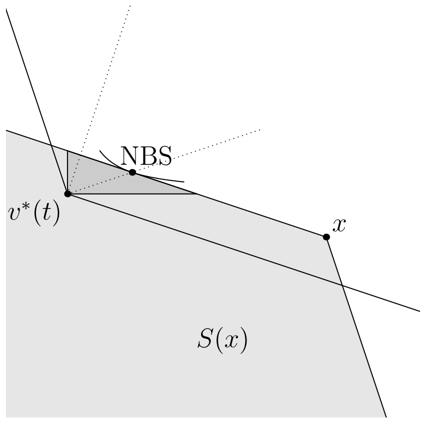

Figure C.2: The set of feasible allocations when $x \in T_1$ is realized

Suppose that realized allocation x belongs to the triangle T_1 in Figure C.1. Then the set $S(x)$ of feasible allocations is described in Figure C.2. Since the disagreement point is at $v^*(t)$, the Nash bargaining solution (NBS) is located on the borderline between T_1 and T_3. Therefore the ex post distribution of payoff profiles on agreement has a mass on the line segment between T_1 and T_3, and the barycenter b_1 of the mass is the intersection point between the line segment and the line drawn through the barycenter of T_1 with slope $-a$. The symmetric argument applies to the case of $x \in T_2$, and the barycenter b_2 of the mass on the borderline between T_2 and T_3 is computed correspondingly.

If x belongs to T_3, the Nash bargaining solution is x itself. The the barycenter b_3 of the set of ex post payoff profiles conditional on the realized allocation x being contained in T_3 is exactly the barycenter of T_3. A computation shows that

$$\begin{aligned}
b_1 &= v(t) + \Big(\frac{2}{3(1+a)}, \frac{2a}{3(1+a)}\Big) z(t), \quad b_2 = v(t) + \Big(\frac{2a}{3(1+a)}, \frac{2}{3(1+a)}\Big) z(t), \\
b_3 &= v(t) + \Big(\frac{2}{3}, \frac{2}{3}\Big) z(t), \\
\mu(T_1) = \mu(T_2) &= \frac{8a}{1-a^2} z(t)^2, \quad \mu(T_3) = \frac{2(1-a)}{1+a} z(t)^2.
\end{aligned}$$

Therefore the barycenter of the entire set of ex post payoff profiles is computed as a convex combination of b_1, b_2, b_3. By ODE (4.1),

$$\begin{aligned}
z'(t) &= -v_1'(t) \\
&= -\lambda\big((b_1 - v(t))\mu(T_1) + (b_2 - v(t))\mu(T_2) + (b_3 - v(t))\mu(T_3)\big) \\
&= -\lambda \cdot \frac{8(1+a^2)}{3(1-a^2)} z(t)^3.
\end{aligned}$$

Since $p(t) = \mu(T_1) + \mu(T_2) + \mu(T_3) = \dfrac{4(1+a)}{1-a} z(t)^2$,

$$\begin{aligned}
p'(t) &= \frac{8(1+a)}{1-a} z(t) z'(t) \\
&= \lambda \cdot \frac{4(1+a^2)}{3(1+a)^2} p(t)^2.
\end{aligned}$$

Therefore the constant r defined in Section 4.4.2 is $\frac{4(1+a^2)}{3(1+a)^2}$. By Theorem 8, the limit duration is

$$\begin{aligned}
D(\infty) &= \frac{1}{1 + \left(\frac{4(1+a^2)}{3(1+a)^2}\right)^{-1}} \\
&= \frac{4+4a^2}{7+6a+7a^2}.
\end{aligned}$$

BIBLIOGRAPHY

ABDELAZIZ, F., AND S. KRICHEN (2007): "Optimal stopping problems by two or more decision makers: a survey," *Computational Management Science*, 4(2), 89–111.

ABREU, D., AND F. GUL (2000): "Bargaining and Reputation," *Econometrica*, 68(1), 85–117.

ABREU, D., AND D. PEARCE (2007): "Bargaining, Reputation and Equilibrium Selection in Repeated Games with Contracts," *Econometrica*, 75(3), 653–710.

ADMATI, A. R., AND M. PERRY (1991): "Joint Project without Commitment," *Review of Economic Studies*, 58(2), 259–276.

ALBRECHT, J., A. ANDERSON, AND S. VROMAN (2010): "Search by committee," *Journal of Economic Theory*, 145(4), 1386–1407.

ALPERN, S., AND S. GAL (2009): "Analysis and design of selection committees: a game theoretic secretary problem," *International Journal of Game Theory*, 38(3), 377–394.

ALPERN, S., S. GAL, AND E. SOLAN (2010): "A sequential selection game with vetoes," *Games and Economic Behavior*, 38(1), 1–14.

AMBRUS, A., AND J. BURNS (2010): "Gradual bidding in ebay-like auctions," mimeo.

AMBRUS, A., AND S.-E. LU (2010a): "A continuous-time model of multilateral bargaining," mimeo.

——— (2010b): "Legislative bargaining with long finite horizons," mimeo.

AMBRUS, A., AND P. A. PATHAK (2011): "Cooperation over finite horizons: a theory and experiments," *Journal of Public Economics*, 95(7-8), 500–512.

AUMANN, R., AND S. HART (2003): "Long Cheap Talk," *Econometrica*, 71(6), 1619–1660.

BENOIT, J.-P., AND V. KRISHNA (1985): "Finitely Repeated Games," *Econometrica*, 53(4), 905–922.

——— (1993): "Renegotiation in Finitely Repeated Games," *Econometrica*, 61(2), 303–323.

BERGEMANN, D., AND J. VÄLIMÄKI (2011): "Efficient Search by Committee," mimeo.

BERNHEIM, B. D., AND A. DASGUPTA (1995): "Repeated Games with Asymptotically Finite Horizons," *Journal of Economic Theory*, 67(1), 129–152.

BIAIS, B., B. CRISTOPHE, AND P. SEBASTIEN (2008): "Equilibrium discovery and preopening mechanisms in an experimental market," mimeo.

CALCAGNO, R., Y. KAMADA, S. LOVO, AND T. SUGAYA (2012): "Asynchronicity and Coordination in Common and Opposing Interest Games," mimeo.

CALCAGNO, R., AND S. LOVO (2010): "Preopening and Equilibrium Selection," mimeo.

CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61(5), 989–1018.

CARUANA, G., AND L. EINAV (2008): "A Theory of Endogenous Commitment," *Review of Economic Studies*, 75(1), 99–116.

CHAO, Y. R. (1919): "A Note on "Continuous Mathematical Induction"," *Bulletin of the American Mathematical Society*, 26(1), 17–18.

CHO, I.-K., AND A. MATSUI (2011): "A Search Theoretic Foundation of Nash Bargaining Solution," mimeo.

CODDINGTON, E. A., AND N. LEVINSON (1955): *Theory of Ordinary Differential Equations*. McGraw-Hill.

COMPTE, O., AND P. JEHIEL (2003): "Voluntary contributions to a joint project with asymmetric agents," *Journal of Economic Theory*, 112(2), 334–342.

——— (2004): "Bargaining over Randomly Generated Offers: A new perspective on multi-party bargaining," mimeo.

——— (2010): "Bargaining and Majority Rules: A Collective Search Perspective," *Journal of Political Economy*, 118(2), 189–221.

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

——— (2003): "Coordination Need Not Be A Problem," mimeo.

FARRELL, J. (1987): "Cheap talk coordination and entry," *RAND Journal of Economics*, 18(1), 34–39.

FARRELL, J., AND G. SALONER (1985): "Standardization, Compatibility, and Innovation," *RAND Journal of Economics*, 16(1), 70–83.

FERGUSON, T. S. (1989): "Who solved the secretary problem?," *Statistical Science*, 4(3), 282–296.

——— (2005): "Selection by Committee," in *Advances in Dynamic Games*, ed. by A. S. Nowak, and K. Szajowski, vol. 7 of *Annals of the International Society of Dynamic Games*, pp. 203–209. Birkhäuser Boston.

FU CHOU, C., AND J. GEANAKOPLOS (1988): "The Power of Commitment," Cowles Foundation Discussion Paper NO. 885.

FUDENBERG, D., AND D. K. LEVINE (1983): "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31(2), 251–268.

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

GALE, D. (1995): "Dynamic Coordination Games," *Economic Theory*, 5(1), 1–18.

——— (2001): "Monotone Games with Positive Spillovers," *Games and Economic Behaviors*, 27(2), 295–320.

GOMES, A., S. HART, AND A. MAS-COLELL (1999): "Finite Horizon Bargaining and the Consistent Field," *Games and Economic Behavior*, 27(2), 204–228.

HARRINGTON, J. E. J. (1987): "Collusion in Multiproduct Oligopoly Games under a Finite Horizon," *International Economic Review*, 28(1), 1–14.

HERRERO, M. J. (1989): "The Nash Program: Non-Convex Bargaining Problems," *Journal of Economic Theory*, 49(2), 266–277.

IMAI, H., AND H. SALONEN (2011): "A characterization of a limit solution for finite horizon bargaining problems," *International Journal of Game Theory*, forthcoming.

ISHII, Y., AND Y. KAMADA (2011): "The Effect of Correlated Inertia on Coordination," mimeo.

JOHN, C. H., AND R. SELTEN (1988): *A general Theory of Equilibrium Selection in Games*. MIT press.

KALAI, E., AND A. NEME (1992): "The Strenght of a Little Perfection," *International Journal of Game Theory*, 20(4), 335–255.

KAMADA, Y., AND M. KANDORI (2011a): "Asynchronous Revision Games," mimeo.

——— (2011b): "Revision Games," mimeo.

KAMADA, Y., AND F. KOJIMA (2012): "Voter preferences, Polarization, and Electoral Policies," mimeo.

KAMADA, Y., AND N. MUTO (2011a): "Delay in Multi-Agent Search with Time Costs," mimeo.

——— (2011b): "Multi-Agent Search with Deadline," mimeo.

——— (2011c): "Multi-Agent Search with Recall," mimeo.

KAMADA, Y., AND N. RAO (2009): "Sequential Implementation of Unenforceable Contracts with Stochastic Transaction Costs," mimeo.

KAMADA, Y., AND T. SUGAYA (2010): "Asynchronous Revision Games with Deadline: Unique Equilibrium in Coordination Games," mimeo.

——— (2011): "Valence Candidates and Ambiguous Platforms in Policy Announcement Games," mimeo.

KANDORI, M., G. J. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61(1), 29–56.

KREPS, D. M., P. MILGROM, J. ROBERTS, AND R. WILSON (1982): "Rational cooperation in the finitely repeated prisoners' dilemma," *Journal of Economic Theory*, 27(2), 245–252.

LAGUNOFF, R., AND A. MATSUI (1997): "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 65(6), 1467–1477.

——— (2001): "Are "Anti-Folk Theorems" in repeated games nongeneric?," *Review of Economic Design*, 6(3-4), 397–412.

LIPMAN, B. L., AND R. WANG (2000): "Switching Costs in Frequently Repeated Games," *Journal of Economic Theory*, 93(2), 149–190.

LIPPMAN, S. A., AND J. J. MCCALL (1976): "The Economics of Job Search: A Survey," *Economic Inquiry*, 14(2), 155–189.

MA, C.-T. A., AND M. MANOVE (1993): "Bargaining with Deadlines and Imperfect Player Control," *Econometrica*, 61(6), 1313–1339.

MARX, L. M., AND S. A. MATTHEWS (2000): "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, 67(2), 327–358.

MASCHLER, M., G. OWEN, AND B. PELEG (1988): "Paths leading to the Nash set," in *The Shapley value: Essays in Honor of Lloyd S. Shapley*, ed. by A. E. Roth, pp. 321–330. Cambridge University Press.

MATSUI, A., AND K. MATSUYAMA (1995): "An Approach to Equilibrium Selection," *Journal of Economic Theory*, 65(2), 415–434.

MCCALL, J. J. (1970): "Economics of Information and Job Search," *The Quarterly Journal of Economics*, 84(1), 113–126.

MOLDOVANU, B., AND X. SHI (2010): "Search Committees," mimeo.

MORRIS, S., R. ROB, AND H. S. SHIN (1995): "p-Dominance and Belief Potential," *Econometrica*, 63(1), 145–167.

MORRIS, S., AND H. S. SHIN (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88(3), 587–597.

MORTENSEN, D. T. (1970): "Job Search, the Duration of Unemployment, and the Phillips Curve," *The American Economic Review*, 60(5), 847–862.

NASH, J. F. (1953): "Two-Person Cooperative Games," *Econometrica*, 21(1), 128–140.

NEYMAN, A. (1985): "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma," *Economics Letters*, 19(3), 227–229.

——— (1998): "Finitely Repeated Games with Finite Automata," *Mathematics of Operations Research*, 23(3), 513–552.

——— (1999): "Cooperation in Repeated Games When the Number of Stages Is Not Commonly Known," *Econometrica*, 67(1), 45–64.

OSBORNE, M. J. (1995): "Spatial Models of Political Competition Under Plurality Rule: A Survey of Some Explanations of The Number of Candidates and The Positions They Take," *Canadiann Journal of Economics*, 2, 261–301.

OSBORNE, M. J., AND A. RUBINSTEIN (1990): *Bargaining and Markets*. Academic Press.

OYAMA, D., S. TAKAHASHI, AND J. HOFBAUER (2008): "Monotone Methods for Equilibrium Selection under Perfect Foresight Dynamics," *Theoretical Economics*, 3(2), 155–192.

PITCHFORD, R., AND C. M. SNYDER (2004): "A Solution to the Hold-up Problem Involving Gradual Investment," *Journal Economic Theory*, 114(2), 88–103.

RABIN, M. (1994): "A model of pre-game communication," *Journal Economic Theory*, 63(2), 370–391.

ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43(4), 959–988.

ROMM, A. (2011): "Reputation in Revision Games," mimeo.

RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97–109.

——— (1989): "The Electronic Mail Game: A Game with Almost Common Knowledge," *American Economic Review*, 79(3), 389–391.

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

SAKAGUCHI, M. (1973): "Optimal Stopping in Sampling from a Bivariate Distribution," *Journal of the Operations Research Society of Japan*, 16(3), 186–200.

——— (1978): "When to Stop: Randomly Appearing Bivariate Target Values," *Journal of the Operations Research Society of Japan*, 21(1), 45–58.

SOBEL, J. (1985): "A Theory of Credibility," *Review of Economic Studies*, 52(4), 557–573.

STINCHCOMBE, M. B. (1992): "Maximal Strategy Sets for Continuous-Time Game Theory," *Journal of Economic Theory*, 56(2), 235–265.

SUGAYA, T., AND S. TAKAHASHI (2011): "Coordination Failure in Repeated Games with Private Monitoring," mimeo.

TAKAHASHI, S. (2005): "Infinite Horizon Common Interest Games with Perfect Information," *Games and Economic Behavior*, 53(2), 231–247.

VAN DAMME, E., AND S. HURKENS (1996): "Commitment Robust Equilibria and Endogenous Timing," *Games and Economic Behavior*, 15(2), 290–311.

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

WILSON, C. A. (2001): "Mediation and the Nash bargaining solution," *Review of Economic Design*, 6(3), 353–370.

YOON, K. (2001): "A Folk Theorem for Asynchronous Repeated Games," *Econometrica*, 69(1), 191–200.

YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61(1), 57–84.

——— (1998): "Conventional Contracts," *Review of Economic Studies*, 65(4), 773–792.