Multi-Agent Search with Deadline*

Yuichiro Kamada† Nozomu Muto‡

January 16, 2015

Abstract

We study 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 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. The limit expected duration increases 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.

Keywords: Multi-agent search, finite horizon, duration, continuous time.

*We thank David Ahn, Attila Ambrus, Pol Antràs, Katie Baldiga, Alessandro Bonatti, Georgy Egorov, Jeff Ely, 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, Benny Moldovanu, Ichiro Obara, Akira Okada, Wojciech Olszewski, Daisuke Oyama, Debraj Ray, Al Roth, Yuval Salant, Larry Samuelson, Bruno Strulovici, Tomasz Strzalecki, Takuo Sugaya, Steven Tadelis, Satrou Takahashi, Kentaro Tomoeda, Takashi Ui, Takahiro Watanabe, Alex Wolitzky, Yuichi Yamamoto, Yosuke Yasuda, and seminar/conference participants at Brown University, Columbia University, Harvard University, Hitotsubashi University, Kyoto University, Toulouse School of Economics, Universitat Autònoma de Barcelona, University of California, Berkeley, University of California, Los Angeles, University of Pennsylvania, University of Tokyo, Washington University in St. Louis, Yale University, Yokohama National University, the 22nd Summer Festival on Game Theory (International Conference on Game Theory) at Stony Brook, SWET 2011 at Sapporo, GDRI Workshop Marseille, Game Theory Workshop 2012 at Hamamatsu, GAMES 2012 at Istanbul, 66th European Meeting of the Econometric Society at Málaga, and 14th SAET Conference at Tokyo for helpful comments. 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 supports from the Spanish Ministry of Science and Innovation through grant "Consolidated Group-C" ECO2008-04756 and FEDER, and from JSPS Grants-in-Aid for Scientific Research No. 26780116. The version with additional materials in the main text can be found at `http://ykamada.com/pdf/maswd.pdf`.

†Haas School of Business, University of California Berkeley, 545 Student Services Building, #1900 Berkeley, CA 94720-1900, e-mail: `y.cam.24@gmail.com`.

‡Department of Economics, Yokohama National University, 79-3 Tokiwadai, Hodogaya-ku, Yokohama 240-8501, Japan, e-mail: `nozomu.muto@gmail.com`.

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 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 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 a sellers' 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.

Search problems in finite horizon are abundant in real economic situations. In single-agent search problems, a worker may want to find a job before the term of the current contract ends at a known date; a student may want to find a job before graduating; a single may be searching an apartment before her moving. For multiple-agent cases, apart from the above apartment search problem, there can be a firm that needs to fill a position before the start of a new project whose recruiting committee consists of several key decision makers; a family-operated manufactory may want to receive orders before the current orders are completed in order to prevent the facilities from becoming idle.

Many questions arise regarding these situations: What are the incentives in the presence of finite horizon? How do they change when search is conducted with other agents (e.g., when the wife needs to search with the husband)? What are the implications of these changes on equilibrium behaviors?

To understand the answers to these questions, we consider an n-player search problem with a deadline. Time is continuous and "opportunities" arrive according to a Poisson process. Opportunities are i.i.d. realizations of payoff profiles. After viewing an opportunity, the players respond with "accept" or "reject." The search ends if and only if all players accept. If the search does not end by the deadline, players obtain an a priori specified payoff. Notice that the arrival rate of the 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 whenever there are two or more players, we analyze an (appropriately defined) trembling-hand equilibrium that we show is (essentially) unique. Our focus in this paper is on an analysis of search duration in trembling-hand equilibrium, which is one of the observable characteristics of equilibrium behaviors.[1]

Analyzing finite horizon problems, however, is difficult because of the non-stationarity, so we take an indirect approach in order to understand the search duration. First we analyze the limit expected search duration as the arrival rate goes to infinity—so there are many offers until the deadline is reached—, which is relatively easier to analyze. Then we argue that the limit expected duration is reasonably close to the expected durations given finite arrival rates. Specifically, our analysis consists of three steps. The first two steps are about the limit expected duration, and the third step is the "closeness" argument.

In the first step, we show that for any number of players and under minimal assumptions on the payoff distribution, the expected search duration does not shrink to zero even in the limit as the search friction 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 future opportunities to come offers a possibility of an incremental gain in payoffs, but an increased probability of reaching the deadline. Both the reward and the cost go to zero as the search friction vanishes; the optimal balance is difficult to quantify because agents need to make decisions before observing all future realizations of offers. For this reason, we employ an indirect proof that bounds the acceptance probability at each moment in time. The result is in contrast with what the existing models of multi-agent search with infinite horizon (Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013)) imply because the limit expected duration is zero due to discounting in these models.[3]

The second step investigates the agents' incentives when they face opponents, and explores the implication of these incentives on equilibrium behaviors. Specifically, our comparative-statics results show that in the limit, the expected duration increases with the number of agents involved in the search. The reason for this, which we call the "ascending acceptability effect," is 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. In addition, for a fixed number of agents, we demonstrate that the limit expected search duration increases as heterogeneity of preferences are magnified. We call this the "preference heterogeneity effect." Given this observation, we solve for the formula of the limit expected search duration as a function of the distribution of payoff profiles of n agents. The formula enables us to understand how the number of players and the payoff distribution affect the search duration.

[1]Search duration has indeed attracted attention in the job search literature, e.g., Mortensen (1970) in a stationary environment, and Kiefer and Neumann (1979) and Heckman and Singer (1984) in non-stationary environments. See Kiefer (1988) for a survey.

[2]Indeed, in Appendix A.9, we offer examples in which our assumptions do not hold and the result fails.

[3]We discuss a comparison later in the Introduction as well as in Section 6 and Appendix A.1.

Number of agents		1	2	3	5	10	100
Limit expected duration	Cube	.333	.500	.600	.714	.833	.980
	Smooth Pareto frontier	.333	.571	.692	.806	.901	.990

Table 1: The limit expected search durations with horizon length 1.

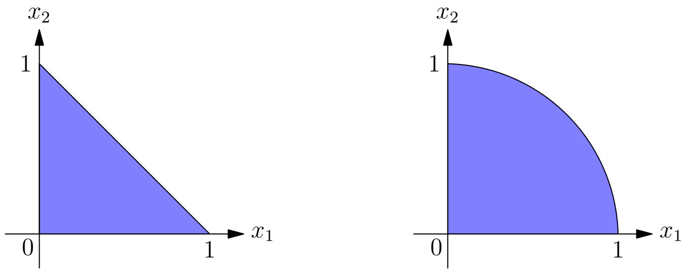

Figure 1: Examples of the domain of feasible payoff profiles

In the third step, we analytically show that the speed of convergence for the expected search duration is fast. Moreover, we use numerical examples to show that the limit expected duration of search is actually close to expected durations with small (but nonnegligible) 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 understanding the duration in reality.

Table 1 illustrates how the limit expected search durations change with respect to the number of agents n. When the feasible payoff set is an n-dimensional cube $[0, 1]^n$, the limit expected search duration increases in n. We formally show in Step 2 that this is general: When the distributions of payoffs are independent across players, the limit expected search duration increases in the number of agents. In a fully general setting, however, we do not assume independence and the duration formula depends on the distribution in a complicated manner. In a special case where the feasible payoff set has a smooth Pareto frontier as in the two-player examples in Figure 1 and the distribution has a strictly positive and continuous density over its support, the limit expected search duration depends only on and increases in n. As the table shows, the difference from the case with $n = 1$ under distributions with smooth Pareto frontiers is larger than the difference under independent distributions, and our formula of the limit expected duration clarifies why it is so. In short, when the Pareto frontier is smooth the preferences are more heterogeneous around the limit expected payoff profile, which makes the limit expected duration longer. Figure 2 depicts the cumulative distribution functions of the durations given various arrival rates when the feasible payoff set is as depicted in the left panel of Figure 1. The figure shows that the limit cumulative distribution function is non-degenerate, and the convergence speed with respect to the arrival rate is fairly fast. For example, the arrival rate of 10 (and the horizon length of 1) corresponds to the

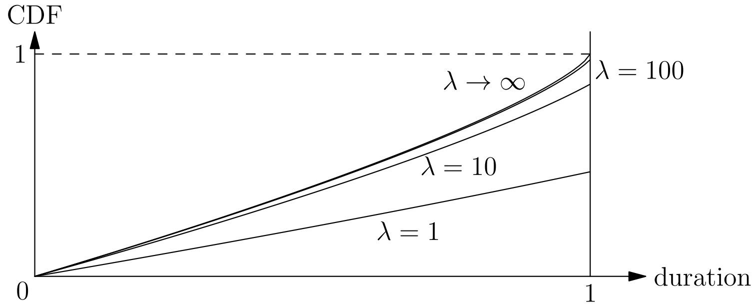

Figure 2: The cumulative distribution functions of the durations for the case with two players where the payoff profile is uniformly distributed on the set shown in the left panel of Figure 1. The parameter λ is the Poisson arrival rate, and the horizon length is 1.

case where there are ten weeks to search an apartment, and the information of a new apartment comes only once in 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 little effect on the duration. A straightforward calculation shows that the expected search duration in this case is 0.608, which is only about 6.5% higher than the limit.

Difficulties

The two key features in our model, *deadline* and *multiple agents*, give rise to new theoretical challenges. 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 second steps) which is relatively easier to characterize, and then argue that the limit case approximates the cases with finite arrival rates reasonably well (the thrid 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 accept, the problem essentially boils down to a single-player search problem. This argument misses an important key feature of our model. It is indeed true in equilibrium 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" are different precisely due to the nonstationarity—hence these two difficulties derived from the two key features of our model interact and produce an additional difficulty. It will become clear in our analysis that it is this interaction that is crucial to our argument in the three steps.

These conceptual challenges invite technical difficulties that we need to handle. On one hand we want a model to be tractable enough to characterize the search duration, so we employ a continuous-time framework.[4] On the other hand, when time is continuous

[4]Sakaguchi (1978) analyzed a continuous-time model with finite horizon as ours and considered a

the recursive dependence of the actions on future strategies that we discussed above continues indefinitely. This implies that there is a possibility for an infinite sequence of punishments, which induces the potential for multiplicity of equilibria that would prevent us from conducting unambiguous comparative statics. To overcome this difficulty, we use trembling-hand equilibrium as a refinement concept.[5] The idea is that the small probability of a tremble by the opponents pins down the behavior when the offer is either highly desirable or highly undesirable. We show that when the horizon is finite so that the continuation payoff at the deadline is uniquely determined, even a small probability of a tremble is enough to actually pin down the behavior at almost all the payoff realizations at any moment of time. Uniqueness can be proven under subgame perfect equilibrium if the order of moves at each opportunity is sequential. However, due to the possibility for an infinite sequence of punishments, this would need a proof similar to the one for the uniqueness of trembling-hand equilibrium payoffs in our model. One merit of employing trembling-hand equilibrium is that the timing of responses at a given opportunity (e.g., simultaneous vs. sequential) does not affect the set of equilibrium outcomes, while it does under subgame perfect equilibrium.

Even under the continuous-time setting, it is difficult to directly analyze the duration so we employ the indirect approach in which we first consider the limit as the arrival rate goes to infinity. The characterization of the limit behavior depends on the probability distribution of offers around the efficiency frontier of the feasible payoff set, and it is easier to handle when we impose the technical assumptions employed in the literature of multi-agent search (Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013)), namely the compact and convex feasible payoff set with strictly positive and continuous density.[6] However, we would like to be as agnostic as possible about reasonability of different assumptions in part because the simplifying assumptions in the literature rule out several cases that some other strings of the literature on search have paid a particular attention to (e.g. log-normal or exponential distributions).[7] For this reason, we do not use the proof technique used in the literature of multi-agent search that makes extensive use of particular distributional assumptions; rather we employ a new proof technique that is free from almost all distributional assumptions. This enables us to better understand the incentives faced by agents in our model; for example, it turns out that the result that the limit expected duration is strictly positive is orthogonal to the simplifying assumptions. The proof of positive limit expected duration identifies a lower bound of the cutoff for each moment of time, and then use that to bound the acceptance probability. This lower

special class of distributions of payoffs. He focused on Markov cutoff strategies, and did not prove uniqueness even in that class. He did not analyze equilibrium durations either.

[5]The uniqueness of trembling-hand equilibrium is straightforward in a discrete-time model with finite horizon since the usual backward induction works.

[6]Cho and Matsui (2013, Section 4.4) generalize some of their results to non-convex cases.

[7]Restricting attention to strictly positive densities over a compact feasible payoff set rules out the distributions that are continuous over $\mathbb{R}^n_+$.

bound is computed by specifying a continuation strategy that is not necessarily a best response. We find that the specified continuation strategy is necessarily nonstationary for the bound to be tight enough. The duration formula can be obtained under general distributions as well, and it gives us insights on how different aspects of distributions affect the duration.[8] For example, we can discuss which term in the duration formula is derived from what aspect of the payoff distribution.

A Brief Summary of the Relation to the Literature

The present paper is related in various ways to a number of lines of the literature. Due to its volume we opted to provide a more comprehensive review in Appendix B. Here, we delve into the detail of the comparison to the literature on multi-agent search with infinite horizon because, in the subsequent sections, we repeatedly compare our results to the ones in that literature.

Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013) consider multi-agent search models with infinite horizon in which a unanimous agreement is required to accept an alternative, and show that the equilibrium outcome approaches the Nash bargaining solution irrespective of the distribution of offers, as the frequency of offers goes to infinity. In the discussion section (Section 6), we consider a model in which payoffs realize as soon as an agreement is reached and show a similar result (convergence to the Nash bargaining solution) even with the presence of a deadline. The difference between these results and our main result is that the payoffs are discounted in the former. With discounting, the effect of the deadline (if any) vanishes as the deadline becomes far away, so the continuation payoff profile converges to a point that does not depend on the payoff distribution around low payoffs, which are not accepted anyway in the relevant future. In our main model with payoffs realizing at the deadline (or equivalently, without discounting), however, the effect of the deadline does not vanish, and the equilibrium dynamics close to the deadline, which are affected by the payoff distribution at low payoffs, critically determines the continuation payoffs at times far from the deadline. This is because the continuation payoff is increasing in the remaining time (as the more offers there are the better off players are), so if the continuation payoff profile at some point in time is close to the Pareto frontier, then further back in time it stays close to that point.

Compte and Jehiel (2010) and Albrecht et al. (2010) consider general majority rules. Albrecht et al. (2010) analyze an infinite-horizon model with independent payoff distributions across agents. Among their results, the most related to ours is the one on the unanimity case, in which they show that the expected search duration increases in the number of agents. The logic is that the cutoff decreases in the number of agents while the expected gain conditional on future acceptance does not change so much due to their

[8]See Section 4.2.2 and Appendix A.15 for the detail.

distributional assumption, and hence the equilibrium condition implies that the expected wait time until the acceptance has to increase. This is because the cutoff is equal to the continuation payoff, which in turn is roughly equal to the discounted value of the sum of the cutoff payoff and the expected gain. We show the same comparative statics with respect to the number of agents. However, as we explain in Section 4.2.1, our logic relies on nonstationarity of cutoffs which is not present in their analysis where stationary equilibrium is assumed.

Under the settings of all the above papers,[9] the expected search duration shrinks to zero in the limit as the frequency of offers goes to infinity, while it often converges to a positive duration under our setting. The characterization of the expected search duration can be neatly done by analyzing the limit expected duration because of this positiveness.

Structure of the Paper

The paper is organized as follows. Section 2 provides the model. Section 3 provides 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 is the main section of the paper. Subsections 4.1, 4.2, and 4.3 correspond to Steps 1, 2, and 3 of our argument, respectively. Section 5 provides a welfare analysis of our main model. There, we discuss Pareto efficiency of the limit payoff profile, and examine how the payoff distribution affects the limit payoff profile.

In order to isolate the effects of multiple agents and a finite horizon as clearly as possible, the departure from the standard model is kept minimal. This enables us to extend and/or modify our model in a wide variety of directions. In Section 6, we provide discussions on such extensions/modifications. Section 7 concludes. The Appendix contains materials that we could not cover in the main text: additional discussion topics, numerical examples, and a comprehensive literature review. Proofs are given in Appendix D unless otherwise noted.

2 Model

The Basic Setup

There are n players searching for an indivisible object. Let $N = \{1, \ldots, n\}$ 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

[9] The "Finite vs. infinite horizon with multiple agents" section in Appendix B discusses additional papers: Alpern and Gal (2009), Alpern et al. (2010), Moldovanu and Shi (2013), Bergemann and Välimäki (2011), and Herings and Predtetchinski (2014).

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 $\mathbb{R}^n$. We call a payoff profile $x \in \mathbb{R}^n$ an allocation or an offer. After allocation x is realized, each player observes x and simultaneously responds by either accepting or rejecting x without a lapse of time.[10] 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 $x^d = (0, \dots, 0) \in \mathbb{R}^n$.[11]

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, by definition, $X \subseteq \mathbb{R}^n$ is a closed subset on which μ has a full support. Without loss of generality, we assume that $X \subseteq \mathbb{R}^n_+$.[12] 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 Pareto efficient allocations and that of all weakly Pareto efficient allocations in X are called the *Pareto frontier* and the *weak Pareto frontier* of X, respectively. We sometimes consider weak Pareto efficiency also in $\hat{X} = \{v \in \mathbb{R}^n_+ \mid x \geq v \text{ for some } x \in X\}$ which is the nonnegative region of the comprehensive extension of X.

Assumptions

We make the following mild assumptions throughout the paper.

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

(b) If $n \geq 2$, for all $i \in N$, the marginal distribution of μ on i's payoffs has a density function that is locally bounded.[13]

Condition (a) is necessary and sufficient for existence of a best response. If it 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

[10] The assumption that each player i observes not only x_i but also x_{-i} does not affect our analysis. That is, even in the model in which each player i only observes x_i, all of our results hold.

[11] This is without loss of generality as long as payoffs realize at the deadline because we can consider an equivalent game in which the origin is shifted to the disagreement point. When the payoffs realize upon agreement as in Appendix A.1, such a shifting cannot induce an equivalent game. However, the change would be minor, as we will explain in footnote 63 in Appendix A.1.

[12] 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 it will turn out that no player accepts any strictly negative payoffs according to the equilibrium concept we will employ.

[13] A function $g(y) : \mathbb{R} \to \mathbb{R}$ is locally bounded if for all $y \in \mathbb{R}$, there exist $C > 0$ and $\varepsilon > 0$ such that $|g(y')| \leq C$ for all $y' \in (y - \varepsilon, y + \varepsilon)$. This property reduces to the (global) boundedness whenever X is a bounded set.

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 that is not orthogonal to any axis. In Appendix A.9, we will provide an example that demonstrates the need for condition (b) for our main results to hold.

Histories and Strategies

A history at $-t \in (-T, 0]$ where players observed k (≥ 0) offers[14] in $[-T, -t)$ consists of

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

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 $\tilde{\mathcal{H}} = \bigcup_{-t\in[-T,0]} \tilde{\mathcal{H}}_t$.[15] Let

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

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.[16]

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 $u_i(\sigma \mid h)$ be the expected continuation payoff of player i given that

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

[15] We let $\tilde{\mathcal{H}}_T^0$ be $\{(T, \varnothing)\}$ and $\tilde{\mathcal{H}}_T^k$ be an empty set for each $k \geq 1$.

[16] The function $u_i(\sigma)$ is well-defined for the following reason: $\mathcal{H}^k := \bigcup_t \mathcal{H}_t^k$ is seen as a subset of $\mathbb{R}^{(2n+1)k+(n+1)}$ because every history $h \in \mathcal{H}^k$ is written as $h = \big((t^1, x^1, b^1), \ldots, (t^k, x^k, b^k), (t, x)\big)$ where for each $l = 1, \ldots, k$, $(t^l, x^l) \in \mathbb{R} \times \mathbb{R}^n$ and $b^l \in \{\text{accept}, \text{reject}\}^n$ which is equivalent to $\{0, 1\}^n \subset \mathbb{R}^n$, and $(t, x) \in \mathbb{R} \times \mathbb{R}^n$. Thus, $\mathcal{H}^k$ is 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 continuous-time games. The definition of behavioral strategies implicitly assumes that the random variables given by the distributions of responses at distinct k's are mutually independent. An analogue of Kuhn's theorem holds under this assumption, so the use of behavioral strategies is without loss of generality. See Aumann (1964) for more discussions on the definition of behavioral strategies and a proof of Kuhn's theorem in extensive-form games with uncountably many nodes.

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 every history $h \in \mathcal{H}_t$ at every $-t$, $\sigma_i(h)$ 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 perfect equilibrium and σ_i is a Markov strategy for all $i \in N$. We will later show that players play Markov perfect equilibrium actions (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 ε-constrained 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 an *extensive-form* (or agent normal-form) *trembling-hand perfect equilibrium* (henceforth, *trembling-hand equilibrium* for short) if there exists a sequence $(\varepsilon^m)_{m=1,2,\ldots}$ and a sequence of strategy profiles $(\sigma^m)_{m=1,2,\ldots}$ such that $\varepsilon^m > 0$ for all m, $\lim_{m\to\infty} \varepsilon^m = 0$, $\sigma^m \in \Sigma^{\varepsilon^m}$, 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 $h \in \mathcal{H}$ with respect to the pointwise convergence in histories.[17]

3 Preliminary Results

In this section, we present preliminary results which will become useful in the subsequent analyses. 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 strategy 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 time length from the deadline with the same ratio. Finally, by examining the differential equation, the limit equilibrium payoff profile as $\lambda \to \infty$ is shown to be weakly efficient.

The next proposition shows that trembling-hand equilibrium is essentially unique.

Proposition 1. *Suppose that σ and σ' are both trembling-hand equilibria. Then for almost all $t \in [0, T]$ and almost all histories $h, h' \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$, $u_i(\sigma| \, h) = u_i(\sigma'| \, h')$ for all $i \in N$.*[18]

That is, regardless of the history, any two trembling-hand equilibria give rise to the same continuation payoff at time $-t$. Four remarks are in order: First, we ruled out histo-

[17] The ε^n does not need to be the same across all histories. All the results in the paper remain true as long as there is a sequence $(\bar{\varepsilon}_i^m)_{i\in N, m}$ of maps with $\bar{\varepsilon}_i^m : \mathcal{H} \to (0, 1/2)$ such that for each $i \in N$, (i) $\inf_{h\in\mathcal{H}} \bar{\varepsilon}_i^m(h) > 0$ for each m, and (ii) $\bar{\varepsilon}_i^m(h) \to 0$ as $m \to \infty$ for all $h \in \mathcal{H}$.

[18] We assume that the set of histories $\tilde{\mathcal{H}}_t^k \setminus \mathcal{H}_t^k$ with k realizations is endowed with a measure induced by the product of the Lebesgue measure (for t's), the measure μ (for x's), and the counting measure (for accept/reject choices), and that $\tilde{\mathcal{H}}_t \setminus \mathcal{H}_t = \bigcup_k \tilde{\mathcal{H}}_t^k \setminus \mathcal{H}_t^k$ is endowed with the measure given by the countable sum of the measures. "Almost all" histories are considered with respect to this measure.

ries in $\mathcal{H}_t$, because different payoffs realized at $-t$ clearly give rise to different continuation payoffs if these payoffs are high enough for players to accept. Second, the proposition in particular implies that the equality holds with probability one if h and h' are induced as a result of the play by σ and σ', respectively. Third, since agents move simultaneously, there exists a subgame perfect equilibrium in which all players reject any allocations. Trembles rule out such a trivial equilibrium. 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. Fourth, and relatedly, there exist subgame perfect equilibria in which players accept low offers but reject high offers, and have strict incentives at almost all histories in $\mathcal{H}$.[19] Our trembling-hand equilibrium also rules out such equilibria. Since we work with continuous time, essential uniqueness is not an easy consequence of finite horizon. Indeed, an analogous proof would be necessary even when the moves at each opportunity were sequential. We detail the intuition behind the result in Section 6.

A Markov strategy σ_i of player $i \in N$ is a *cutoff strategy* if there exists a function $c_i : [0,T] \to \mathbb{R}_+$ such that player i who is to respond at time $-t$ accepts allocation $x \in X$ whenever $x_i \geq c_i(t)$, and rejects it otherwise. We call $c_i(t)$ the cutoff of σ_i at time $-t$. For a payoff profile $v = (v_1, \ldots, v_n) \in \mathbb{R}^n_+$, we define a set of payoff profiles that each agent i finds weakly better than v_i 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 $c(t) := \big(c_1(t), \ldots, c_n(t)\big)$, we call $A(c(t))$ an "acceptance set" as they agree with an allocation x at time $-t$ if and only if $x \in A(c(t))$. We often denote this acceptance set by $A(t)$ with a slight abuse of notation when the cutoff profile in consideration is not ambiguous.

If players play a Markov strategy profile σ and there is no Poisson arrival at time $-t$, then player i has an expected payoff $u_i(\sigma|\, h)$ at $-t$ that does not depend on history $h \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ realized at time $-t$. To simplify the notation, as long as a Markov strategy profile in consideration is clear, we hereafter omit to explicitly refer to it, and denote by $v_i(t)$ the continuation payoff of i conditional on histories in $\tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ when it is played.

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

Proposition 2. *There exists a trembling-hand equilibrium that consists of (Markov) cutoff strategies such that (i) the cutoff profile at each time $-t \in [-T, 0]$ equals the equilibrium continuation payoff profile $v(t) = (v_1(t), \ldots v_n(t))$, (ii) $v(t)$ is differentiable in t, and (iii) $v(t)$ is given by a solution of the following ordinary differential equation (ODE)*

$$v'(t) = \lambda \int_{A(t)} \big(x - v(t)\big)\, d\mu \tag{1}$$

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

[19] An example similar to the one in Cho and Matsui (2013, Proposition 4.4) can be used to show this result.

This proposition is shown by the following argument. First, suppose that players play a cutoff strategy profile σ whose cutoff profile coincides with the continuation payoff profile $v(t)$ at each time $-t$. Given the cutoff strategy σ_i, player i accepts the drawn payoff profile $x \in X$ if and only if player i finds x_i to be no worse than her continuation payoff. Since such a behavior clearly maximizes her continuation payoff at each time $-t$, σ is a Markov perfect equilibrium.

We can next show that there exists a cutoff strategy profile whose cutoff profile coincides with the continuation payoff profile $v(t)$ at each time $-t$ if $v_i(t)$ satisfies the following recursive expression for each $i \in N$ and each $-t \in [-T, 0]$:

$$\begin{aligned} v_i(t) &= \int_0^t \left(\int_{X \setminus A(\tau)} v_i(\tau)\, d\mu + \int_{A(\tau)} x_i \, d\mu \right) \lambda e^{-\lambda(t-\tau)}\, d\tau \\ &= \int_0^t \left(v_i(\tau) + \int_{A(\tau)} \big(x_i - v_i(\tau)\big)\, d\mu \right) \lambda e^{-\lambda(t-\tau)}\, d\tau. \end{aligned} \tag{2}$$

This is because of the following: Suppose that the cutoff profile is $c(t)$ at each time $-t$. After time $-t$, players receive the first Poisson opportunity at time $-\tau \in (-t, 0)$ with probability density $\lambda e^{-\lambda(t-\tau)}$. If all players accept the offer x, i.e., $x \in A(c(\tau))$, then an agreement is reached with x which gives x_i to each i. If some player rejects x, then search continues with continuation payoff $v_i(\tau)$ for each i. Recalling that we defined $A(\tau) = A(v(\tau))$, if $v_i(t)$ satisfies (2) for each $i \in N$ and each $-t \in [-T, 0]$, then there exists a cutoff strategy profile with cutoff profile $c(t) = v(t)$ for each $-t$ such that the continuation payoff profile is $v(t)$ for each $-t$.

We can show that Bellman equation (2) implies differentiability of $v_i(t)$.[20] Multiplying both sides of (2) by $e^{\lambda t}$ and differentiating with respect to t yield the ordinary differential equation (1) of continuation payoff profile $v(t)$ defined in $\hat{X}$.

Now, a standard argument for ordinary differential equations shows that ODE (1) has a solution whenever Assumption 1 holds.[21] The above argument only shows that ODE (1) has a solution, and if a cutoff strategy profile employs a cutoff profile given by a solution of (1), it is a Markov perfect equilibrium. In Appendix D.3, we will show that it is in fact a trembling-hand equilibrium. By Proposition 1 and the above argument, the solution of ODE (1) is essentially unique, and since a solution $v(t)$ must be continuous, ODE (1) has a unique solution.

Therefore the game has an essentially unique trembling-hand equilibrium for any given μ satisfying Assumption 1. Let us denote the unique solution of (1) by $v^*(t; \lambda)$, which is 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

[20] We show this in Appendix D.3.

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

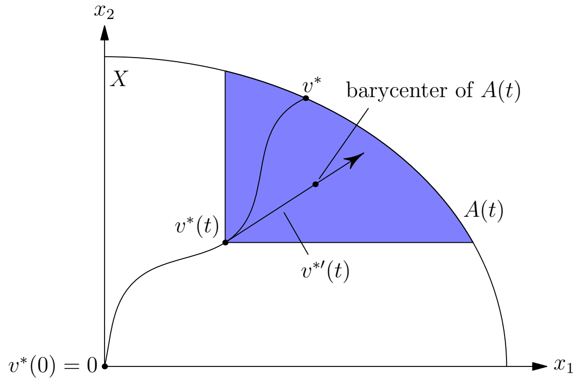

Figure 3: The path and the velocity vector of ODE (1)

a realized allocation at time $-t$ on the equilibrium path conditional on the event that an opportunity arrives at $-t$, i.e., $\mu(A(v^*(t)))$, is referred to as the "acceptance probability" at time $-t$. Because this is unique for each $-t$, the expected search duration is uniquely defined. This uniqueness will be helpful in conducting unambiguous comparative statics.

Let us make three observations about ODE (1). First, Figure 3 describes an illustration of a typical path and the velocity vector that appear in this ODE for $n = 2$. The shaded area shows the acceptance set $A(t)$, whose barycenter with respect to the probability measure μ is $\int_{A(t)} x d\mu / \mu(A(t))$. 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 (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 $\mu(A(t)) = 0$. Thus, for each $i \in N$, the continuation payoff $v_i^*(t)$ grows as t increases, and eventually either converges to the limit payoff, or diverges to infinity.

Second, since the right hand side of ODE (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 3. *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 $v^* \in \mathbb{R}^n_+$. In the next section, we sometimes deal with these two limits interchangeably. Note that the equality implies that $\lim_{\lambda\to\infty} v^*(T; \lambda)$ does not depend on $T > 0$.

Third, we argue weak Pareto efficiency of the limit allocation. Suppose that $v^* := \lim_{\lambda\to\infty} v^*(T; \lambda) = \lim_{T\to\infty} v^*(T; \lambda)$ 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. Since $A(v^*)$ contains an open neighborhood of x, $\mu\big(A(v^*)\big) > 0$. This implies that the right hand side of ODE (1) is strictly positive, contradicting the starting assumption that $v^* = \lim_{\lambda\to\infty} v^*(t; \lambda) = \lim_{t\to\infty} v^*(t; \lambda)$. Hence we obtain the following proposition:

Proposition 4. *Fix $t > 0$, and suppose that the solution $v^*(t;\lambda)$ of equation* (1) *converges to $v^* \in \hat{X}$ as $\lambda \to \infty$. Then v^* is weakly Pareto efficient.*[22]

We note that if $n = 1$, weak Pareto efficiency immediately implies Pareto efficiency. For the case with $n \geq 2$, we will have further discussions about efficiency in Section 5, where we will show in particular that the limit allocation v^* is Pareto efficient for all convex X.

4 Duration of Search

In this section, we discuss the duration of search in our model. Our argument consists of three steps: In Section 4.1 we show that even under quite mild conditions, search takes a positive amount of time even in the limit as the friction vanishes. In Section 4.2, we investigate the agents' incentives when they face the opponents, and explore the implications of these incentives on equilibrium behaviors. In Section 4.3, we demonstrate that the limit expected duration is "close" to the expected 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 understanding the duration in reality.

Let $D(\lambda)$ be the expected duration in the equilibrium for given λ when $T = 1$.[23] Note again that this is uniquely defined due to the essential uniqueness that we proved in Propositions 1 and 2. Since we have $v^*(T;\lambda) = v^*(1;\lambda T)$ as discussed in Section 3, the expected duration for general T with arrival rate λ is written as $D(\lambda T)T$. We define $D(\infty) = \lim_{\lambda\to\infty} D(\lambda)$ whenever the limit is well-defined.[24]

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 μ satisfying fairly mild 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 1) and detail the intuition. Then we generalize the result to the case of an arbitrary number of players (Theorem 2). Nonstationarity of strategies will be important in both proofs.

[22]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 payoff profiles that strictly Pareto-dominate the limit expected payoff profile, while such allocations cannot be achieved under a trembling-hand equilibrium. See Example 4 in Section 5 for a further discussion on this.

[23]We do not express dependence on n and μ explicitly as long as there is no room for confusion.

[24]In Section 4.2.2, we will be explicit about the conditions that guarantee existence of the limit.

4.1.1 Single Agent

Roughly, there are two effects of having a higher arrival rate. On one hand, for any given (small) time interval, there are an increasing number of opportunities, thus it becomes more likely 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 because the value of the second effect is difficult to quantify. This difficulty is caused by the fact that agents need to make decisions before observing all future options, and in such an environment we need to investigate the trade-off between an incremental gain in payoffs from waiting for a future opportunity to arrive and an increased probability of reaching the deadline.

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 as a function of time $-t < 0$ is $O(\frac{1}{\lambda t})$ then the limit expected duration is strictly positive.[25] Then we show that the acceptance probability must be indeed $O(\frac{1}{\lambda t})$.

Suppose that the acceptance probability is $O(\frac{1}{\lambda t})$. Then, the probability that an agreement does not take place until 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 $-t$. 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 the continuation payoff which must be equal to the cutoff, and this gives us an upper bound of the acceptance probability:

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

[25] In general, 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 $y \geq \bar{y}$. In this expression, λt is substituted for y.

To see what type of future strategies will generate an interesting bound, we note that a good future strategy would satisfy two features. The first is that the payoff conditional on acceptance is high, and the second is that the probability of acceptance over the whole time horizon is high. In order to examine the tradeoff between these two, below we examine two scenarios with constant future cutoffs. In each of them, either one of the two features fails. Then we examine an alternative future strategy that takes the best of both worlds and gives us the desired bound.

As the first scenario, suppose that the cutoff is $1 - O(\frac{1}{\lambda t})$ at any time $-s$ after $-t$. 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 1 (the best possible payoff) and 0 irrespective of λ, which means that the acceptance probability at $-t$ is bounded from below by a strictly positive number that is independent of λ, and thus cannot be $O(\frac{1}{\lambda t})$.

In the second scenario, consider a future strategy such that at any time $-s$ after $-t$, the cutoff is such that the player accepts with a probability of 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 high future cutoffs gives her a high payoff conditional on acceptance, but reduces the acceptance probability. On the other hand, setting low future cutoffs results in a low payoff conditional on acceptance but raises the acceptance probability. This suggests that a 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 deadline to ensure a high enough acceptance probability. Specifically, consider a *non-stationary* cutoff plan $1 - \frac{2}{\lambda s + 2}$ for each time $-s$ after time $-t$. This plan has a feature that at any 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 proba-

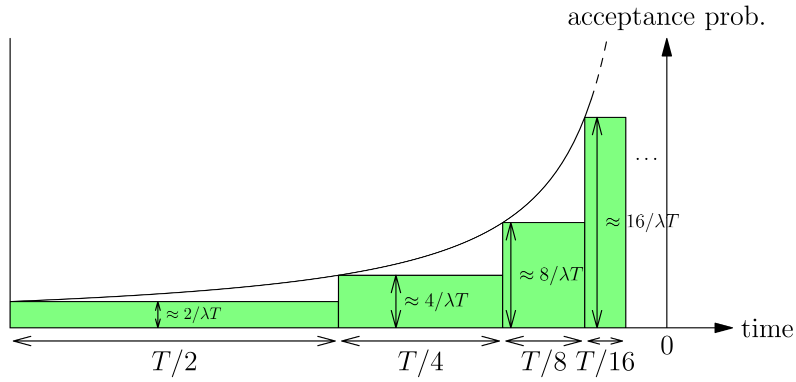

Figure 4: A computation of a lower bound of the agreement probability.

bility of 1, as the probability for no acceptance can be calculated as:

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

We now provide rough intuition for why this can achieve the limit acceptance probability of 1. In Figure 4, the curve represents the acceptance probability when the agent plays the above strategy. The height of the rectangle in subinterval $[-\frac{T}{2^{k-1}}, -\frac{T}{2^k}]$ for fixed k approximates $2^k/(\lambda T)$ when λ is sufficiently large. Thus the area of the rectangle is approximately $1/\lambda$, which is independent of k. Since the expected number of opportunities in a subinterval is proportional to λ, the agreement probability in a subinterval is constant in k, and the disagreement probability over all subintervals decreases with an exponential speed as the number of subintervals increases.

More precisely, we can show that for any small $\varepsilon > 0$, the acceptance probability is higher than $1-\varepsilon$ for a large enough λ. Let $K \geq -2\log\varepsilon$ and $\lambda \geq 2^K/T$. Let us consider division $[-\frac{T}{2^{k-1}}, -\frac{T}{2^k})$ for $k = 1, 2, \ldots, K$ of the interval $[-T, -\frac{T}{2^K})$. The strategy with cutoff $1 - \frac{2}{\lambda s+2}$ makes an acceptance in each subinterval with probability at least

$$1 - e^{-\lambda \cdot \frac{2}{\lambda(T/2^{k-1})+2} \cdot \frac{T}{2^k}} = 1 - e^{-\frac{\lambda T}{\lambda T + 2^k}} \geq 1 - e^{-1/2} > 0.$$

Then the probability of reaching an agreement within the K subintervals is higher than $1 - e^{-K/2} \geq 1 - \varepsilon$.[26]

Finally, under the strategy given, the continuation payoff is $1 - O\left(\frac{1}{\lambda t}\right)$ for all $t > 0$. Thus, conditional on accepting, the loss from the best payoff is only of the order of $O\left(\frac{1}{\lambda t}\right)$ for any $t > 0$.

Overall, a lower bound of the continuation payoff at time $-t$ with the given future strategy is $1 - O\left(\frac{1}{\lambda t}\right)$, which implies that $O\left(\frac{1}{\lambda t}\right)$ is an upper bound of the acceptance probability at time $-t$. Hence, we conclude that, when $X = [0,1]$ and the distribution μ

[26] In the kth subinterval, agents disagree with probability at most $e^{-1/2}$ for every k. Therefore the probability that agents disagree over all K subintervals is at most $e^{-K/2}$.

is uniform, the limit expected duration is strictly positive. This argument is generalized to the cases of any distributions satisfying Assumption 1 and the following assumption. Let $F(x)$ be the cumulative distribution function of μ.

Assumption 2. There exists a concave function φ such that $1 - \varphi(x)$ is of the same order as $1 - F(x)$ in X.[27]

To see what this assumption means, consider two separate cases—bounded X and unbounded X. If X is bounded, the assumption follows from a simple condition that there exists a neighborhood of the supremum payoff such that F is differentiable and has a bounded slope. If X is unbounded, a sufficient condition is that there exists $\tilde{x}$ such that F is concave on $(\tilde{x}, \infty)$, or equivalently, there exists a nonincreasing density function f on $(\tilde{x}, \infty)$.

Recall that $D(\lambda)$ is the expected duration in the equilibrium for the arrival rate λ and $T = 1$. We obtain the following:

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

Remark 1. (a) In Assumption 2, we assumed that the cumulative distribution function F is approximated by a concave function. In the proof of Theorem 1, concavity of φ lets us invoke the Jensen's inequality to bound the cumulative acceptance probability. To understand the role of concavity, consider μ with a concave cumulative distribution function F over $\mathbb{R}_+$. Notice that this search problem is equivalent to the search problem in which the monetary offer m is distributed uniformly over the support $[0, 1]$, and the agent has a utility function $u(m) = F^{-1}(m)$. Since F is concave, u is convex, i.e., the agent is *risk-loving* in this new problem. What we discussed before stating Assumption 2 can be interpreted as showing that a risk-neutral agent spends a positive amount of time on search in the limit as $\lambda \to 0$, when monetary offers are distributed uniformly over $[0, 1]$. It would be intuitive that risk-loving agents spend more time on search than risk-neutral agents on average under the same distribution of monetary offers.[28]

(b) In Appendix D.4, we prove Theorem 1 under a weaker assumption than Assumption 1 that there exists $\alpha > 0$ such that $\big(1 - \varphi(x)\big)^{\alpha}$ is of the same order as $1 - F(x)$. This generalization is possible because one can show that the limit expected search duration is positive under $F(x) = x^{\alpha}$ with support $[0, 1]$ by the same technique as the one that we used in the discussion before stating Assumption 2. In the context of the interpretation in Remark 1 (a) with the uniform distribution of monetary offers, this

[27]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 $\tilde{y} < \sup(Y)$ such that $c|h(y)| \leq |g(y)| \leq C|h(y)|$ for all $y \geq \tilde{y}$.

[28]Nachman (1972) shows that risk-loving agents spend more time on search than risk-neutral agents in his discrete-time model.

implies that the result of positive duration holds with a wide class of utility functions including those exhibiting risk aversion. □

4.1.2 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 $X = \{x \in \mathbb{R}^2_+ \mid x_1 = x_2 \leq 1\}$ 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, an upper bound of the acceptance probability at each time over $\left[-t, -\left(t - \frac{1}{\sqrt{\lambda t}}\right)\right]$ 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.[29]

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 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 cutoffs. 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 $p(s)$. Such a strategy exists because the marginal cumulative distribution of payoffs is continuous by Assumption 1. Notice that, since each opponent's cutoff is decreasing, player i's marginal payoff distribution conditional on acceptance at each time $-s$ first-order stochastically dominates the one with fixed cutoffs for $-i$, while at each moment the acceptance probability does not change. 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 $D(\lambda)$ is the expected duration in the equilibrium for given arrival rate λ and $T = 1$. For any given $v \in \hat{X}$ such that $\mu(A(0, v_{-i})) > 0$, let F_i^v be the marginal cumulative distribution function of player i's payoff conditional on the event $A(0, v_{-i})$.

[29] There also exist strategies for player 2 that are independent of λ and still give rise to a low enough cutoff for player 1 so that the limit expected duration is zero, such as $v_2(t) = e^{-t}$.

Assumption 2′. There is $i \in N$ such that for all $v \in \hat{X}$ 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 $1 - F_i^v(x_i)$ in $\{x_i \in \mathbb{R} \,|\, (x_i, v_{-i}) \in X\}$.[30]

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

4.2 Step 2: Comparative Statics of Limit Expected Durations

The second step of our argument conducts comparative statics that arises when the agents face the opponents. *Two effects, which we call the "ascending acceptability effect" and the "preference heterogeneity effect," determine the way the search duration is affected by changes in the number of players and the distribution of payoffs, respectively.* We explain these effects in turn.

4.2.1 Ascending Acceptability Effect

In Section 4.1.2 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 in which we add players whose preferences are independent of those of the existing players. Specifically, consider the following three models with $T = 1$:

(i) the n-player model with probability measure μ on $\mathbb{R}_+^n$;

(ii) the m-player model with probability measure γ on $\mathbb{R}_+^m$; and

(iii) the $(n+m)$-player model with product probability measure $\mu \times \gamma$ on $\mathbb{R}_+^{n+m}$.

Suppose that μ and γ satisfy Assumption 1. Let $D_{\mu\gamma}(\lambda)$ be the expected duration in model (iii) when the arrival rate is λ.

Theorem 3. *If the limit expected duration D_μ exists for model (i) and the probability of agreement before the deadline goes to one as $\lambda \to \infty$ in model (ii), then* $\liminf_{\lambda \to \infty} D_{\mu\gamma}(\lambda) \geq D_\mu$.

Thus, by adding problem (ii) to problem (i), the limit expected duration weakly increases.

Remark 2. Two remarks are in order.

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

(a) There are two premises in the theorem. The first premise, that the limit expected duration in model (i) exists, is not essential to the result—we assume it only to simplify the exposition.[31] For the second premise, we note that the agreement probability is one at the deadline if the limit payoff profile is Pareto efficient by Proposition 8 in Section 5.[32] We exploit this premise in our proof. We will indicate below when we use this premise.

(b) If models (i) and (ii) satisfy a mild technical condition that we provide later to ensure the existence of the limit expected durations,[33] this result turns out to be a corollary of stronger results that we prove later. First, in Section 4.2.2, we will actually provide the explicit formula for the limit probability distribution of the duration. The formula in particular implies that the limit distribution of the duration in model (i) is first-order stochastically dominated by that of model (iii), which implies Theorem 3. The dominance is strict, rather than weak as indicated in Theorem 3, if $D_\mu < 1$ and the limit expected duration D_γ in model (ii) is strictly positive. Next, the formula for the probability distribution also gives us the formula for the limit expected duration, and it turns out that this gives us the formula of the limit expected duration $D_{\mu\gamma}$ in model (iii) as a function of D_μ and D_γ, which again in turn implies Theorem 3.[34] □

There is a simple reasoning behind Theorem 3. Note first that for any t, there exists t' such that the locus of the path in model (iii) in the time interval $[-t, 0]$ projected on X is identical to the one in model (i) in $[-t', 0]$ because, by (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 the 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 are increasing chances for desirable draws to be accepted over time ("ascending acceptability"). This is roughly why we expect a longer duration with more players. Another way to put this is that the increase in the acceptance probability caused by the 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. Consistent with this intuition, under a mild technical assumption (Assumption 3 in Section 4.2.2), there exist the limit expected durations D_γ and $D_{\mu\gamma}$ in models (ii) and (iii), respectively, and the ratio of remaining times $\frac{1-D_{\mu\gamma}}{1-D_\mu}$ can be shown to be strictly increasing in D_γ. This is intuitive: Higher D_γ implies that the

[31] If the limit does not exist, then we can replace "D_μ" in the right hand side of the inequality in the theorem with "$\liminf_{\lambda\to\infty} D_\mu(\lambda)$," where $D_\mu(\lambda)$ denotes the expected duration in model (i) when the arrival rate is λ.

[32] In Section 5, we argue that the limit payoff profile is Pareto efficient under quite general environments.

[33] Assumption 3 in Section 4.2.2.

[34] The formula is $D_{\mu\gamma} = 1 - \frac{(1-D_\mu)(1-D_\gamma)}{1-D_\mu D_\gamma}$.

acceptance probability given by the additional players does not increase so much until the deadline comes close enough. Thus players in model (i) have more incentives to wait than in the case with a lower D_γ.

Formally, let J and $-J$ be the sets of players in models (i) and (ii), respectively. By Propositions 1 and 2, the probability that all the players in $-J$ accept at time $-t$ in model (iii) depends only on t and λ. Let this probability be $p^*_{-J}(t;\lambda)$. By Proposition 2, this is decreasing in t (i.e., p^*_{-J} increases as time passes). Thus, each player in J in model (iii) chooses an (essentially unique) equilibrium strategy in model (i) given a *time-dependent* arrival rate $\lambda p^*_{-J}(t;\lambda)$. Since this is decreasing, it follows that the expected duration in this game is strictly higher than that of the game with a *constant* arrival rate $l(\lambda) := \int_0^1 \lambda p^*_{-J}(t;\lambda)\, dt$ because these two games become identical after rescaling the measurement of time. Note that this arrival rate $l(\lambda)$ diverges to infinity as $\lambda \to \infty$ since the agreement probability is one in the limit.[35]

Now, take a sequence of arrival rates $\{\lambda_k\}_{k=1}^{\infty}$ such that $\lim_{k\to\infty} D_{\mu\gamma}(\lambda_k) = \liminf_\lambda D_{\mu\gamma}(\lambda)$. The above argument shows that $D_{\mu\gamma}(\lambda_k) > D_\mu(l(\lambda_k))$ for each k, where $D_\mu(\lambda')$ denotes the limit expected duration given arrival rate λ' in model (i). This inequality implies that $\lim_{k\to\infty} D_{\mu\gamma}(\lambda_k) \geq \lim_{k\to\infty} D_\mu(l(\lambda_k))$.[36] Hence we obtain $\liminf_{\lambda\to\infty} D_{\mu\gamma}(\lambda) \geq D_\mu$ because (a) the arrival rate $l(\lambda)$ diverges to infinity, (b) D_μ exists, so $\lim_{k\to\infty} D_\mu(l(\lambda_k)) = \lim_{\lambda'\to\infty} D_\mu(\lambda')$, and (c) $\lim_{k\to\infty} D_{\mu\gamma}(\lambda_k) = \liminf_\lambda D_{\mu\gamma}(\lambda)$ by the definition of the sequence $\{\lambda_k\}_{k=1}^{\infty}$. This ends the proof of Theorem 3.

4.2.2 Preference Heterogeneity Effect

Theorem 3 considers the case where preferences of players in model (i) are independent of those of players in model (ii). In many relevant cases, however, players' preferences are not independent; they are often heterogeneous. We now analyze how heterogeneity in preferences, captured by the change in X and μ, affects the duration. Specifically, we find that there are two channels through which preference heterogeneity affects the search duration. In this subsection we first give two examples (Examples 1 and 2) to intuitively explain these two channels, and then provide a general duration formula which depends on two terms that correspond to the two channels. Then we use this formula to analyze a specific class of search problems that the literature of multi-agent search with infinite horizon has extensively analyzed (Example 3).

Example 1 (Preference heterogeneity implies a larger probability of an "extra region.").

[35] This is because the probability of no agreement until the deadline in model (ii) is given by $e^{-l(\lambda)}$.

[36] As we noted earlier, this does not show that a *strict* inequality holds. This is because here we do not pin down in what way $p^*_{-J}(t,\lambda)$ approaches a "flat" curve with respect to t. The precise argument to show the strict inequality needs to identify the order of p^*_{-J} in λ and its coefficient when λ is large. A rigorous proof for the strict inequality can be made once we impose Assumption 3 that will be stated in Section 4.2.2, where the detail of this argument is provided.

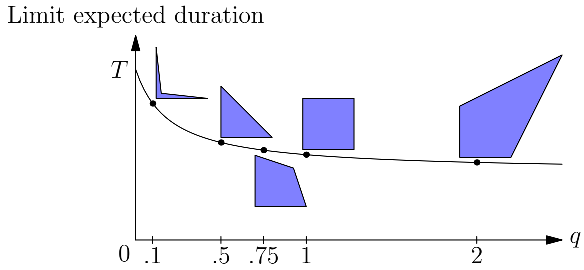

Figure 5: The limit expected duration for the quadrilateral with vertices $(0,0),(1,0),(q,q),(0,1)$.

We consider the two-player case with the uniform distribution on a quadrilateral with vertices $(0,0)$, $(1,0)$, (q,q), and $(0,1)$ ($q>0$). Here, we intend to capture the idea of preference heterogeneity by the parameter q: As q grows, the kink of the boundary at the limit payoff profile becomes sharper, which we interpret as preferences becoming less heterogeneous.

By applying the formula that we will obtain in Theorem 4, one can show that the limit expected duration is $\frac{2q+1}{5q+1}$ when $T=1$, which is decreasing in q. The limit expected duration is depicted in Figure 5 as a function of q. The intuition for this comparative statics is as follows: The ascending acceptability effect states that a player's opponent accepts more offers in the future by lowering the cutoff, so expanding the acceptance region. The significance of this effect is determined in part by the probability assigned to such an "extra region" of the acceptance set. For a large q, the "extra region" does not contain relatively favorable allocations for the player, while for a small q, the region contains relatively favorable allocations for the player, so has a large probability.

The problem (X,μ) in this example itself is special, but it captures a wide range of problems because the acceptance sets in many applications approximate the shape proportional to X for some q in this example in the limit as $t\to\infty$. For example, suppose a given problem satisfies Assumption 4 that we will postulate in Example 3 (essentially requiring the limit acceptance set to be bounded and well-behaved) while the smoothness condition need not be satisfied. Then, the Pareto frontier of X may have a kink at the limit expected payoff profile, and the conditional distribution on the acceptance set when t is large can be approximated by the uniform distribution on X for some q after rescaling on each axis. □

Example 2 (**Preference heterogeneity implies a higher conditional gain.**).

Consider 2-player symmetric $X=\mathbb{R}_+^2$ and μ which is associated with a density function f_σ parameterized by $\sigma>0$ as follows:

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

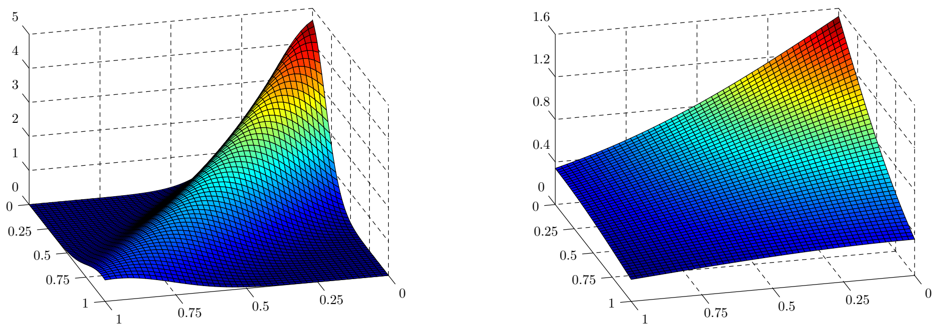

Figure 6: The probability density functions for $\sigma = .2$ (left) and $\sigma = 1$ (right).

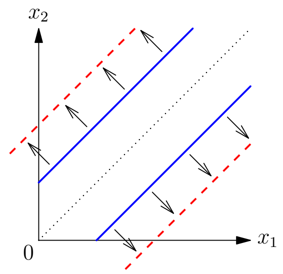

Figure 7: Change of the density when σ increases.

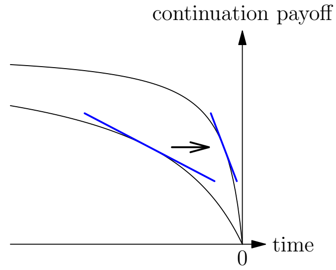

Figure 8: If the cutoff falls fast, the cutoff is high until reaching close to the deadline.

where $M_\sigma = \int_0^\infty \frac{1}{\sqrt{2\pi\sigma^2}} e^{-y} e^{-\frac{y^2}{2\sigma^2}}\, dy$ is a constant ensuring that the total probability is one. That is, we consider an exponential distribution in the direction of the 45 degree line, and a normal distribution with variance σ^2 in the direction of the 135 degree line with a restriction to $\mathbb{R}_+^2$. Here, we intend to capture the idea of preference heterogeneity by using the parameter σ. Figure 6 illustrates how the distribution given by f_σ becomes more heterogeneous as σ increases. Notice that the limit distribution as $\sigma \to 0$ is the (degenerate) exponential distribution over the 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 $\sqrt{2}$. The limit expected durations in these two limit cases can be solved for analytically. In the former case the problem becomes isomorphic to that of one-player case and the limit expected duration goes to $\frac{1}{2}$, and in the latter case it approaches $\frac{2}{3}$. The limit expected durations in the intermediate values of σ are numerically computed, and are increasing in σ. We note that the "probability of the extra region" that we discussed in the previous example is invariant with respect to σ in the current example. Thus this example features another channel through which preference heterogeneity affects search duration. The channel here is the expected gain relative to the cutoff conditional on acceptance, and this varies with σ. Specifically, the expected gain rises with σ from $\frac{1}{\sqrt{2}}$ (when $\sigma \to 0$) to $\sqrt{2}$ (when $\sigma \to \infty$).

The intuition for why preference heterogeneity implies such an increase of the expected gain and the gain in turn affects the search duration, is as follows: First, the more heterogeneous the preferences are, the more realizations of payoffs are scattered outside of the acceptance set.[37] In Figure 7, the density on the solid line moves to that on the dashed line as σ increases. Notice that the payoff profiles with small sums of payoffs on the dashed line move out of the acceptance set, while those with large sums stay in the acceptance set.

This implies that higher heterogeneity (larger σ) leads to higher expected payoffs conditional on acceptance. Thus, if preferences are more heterogeneous, the gain relative to the continuation payoff conditional on acceptance is higher. This means that the loss from a unit time passing is larger for any given cutoff, so the player will decrease the cutoff faster for any given cutoff. Thus, as Figure 8 suggests, the cutoff is higher when the deadline is close. Hence, the cutoff does not fall much until reaching close to the deadline, implying a longer expected duration. □

Now we derive the duration formula. Let us define a variable r as follows:

$$r = \lim_{t \to \infty} \sum_{i \in N} d_i(v^*(t)) \cdot b_i(v^*(t)) \tag{3}$$

where

$$b_i(v) = g_i(A(v)) - v_i, \qquad d_i(v) = -\frac{\partial \mu(A(v))/\partial v_i}{\mu(A(v))},$$

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.

The assumptions we have imposed do not imply that $\partial \mu(A(v))/\partial v_i$ in the definition of d_i exists, or the limit in the definition of r exists. But the existence of these things holds quite generally. A sufficient condition for the existence of $\partial \mu(A(v))/\partial v_i$ is that μ is associated with a locally bounded density function over X. In the single-agent case, the limit in the definition of r exists if there exists some $\tilde{x} < \sup X$ such that the hazard rate $f(x)/(1 - F(x))$ is weakly concave in the interval $(\tilde{x}, \sup X)$. An analogous condition is sufficient for the multiple-agent cases.

For expositional simplicity, in the main text we limit attention to the case in which r exists. We deal with the fully general case in Appendix D.5.

Assumption 3. (a) The variable r is well-defined, i.e., the partial derivative in the definition of d_i and the limit in (3) exist.

(b) The probability measure μ is absolutely continuous with respect to the Lebesgue

[37] The total probability on the acceptance set is independent of σ whenever v is on the 45 degree line because for all $\delta > 0$ and all $(x_1, x_2) \in \mathbb{R}^2_+$, we have $f_\sigma(x_1 + \delta, x_2 + \delta)/f_\sigma(x_1, x_2) = e^{-2\delta}$.

measure on $\mathbb{R}^n_+$.[38,39]

We denote the acceptance probability by $p(t;\lambda) = \mu(A(v^*(t;\lambda)))$, which induces the probability of no agreement until time $-t$, denoted by $P(t;\lambda)$.[40] Recall that $D(\infty)$ is the limit expected duration when $T = 1$. Now we can show that $P(t;\infty) := \lim_{\lambda\to\infty} P(t;\lambda)$ and $D(\infty)$ can be written in the following way:

Theorem 4. *Under Assumptions 1 and 3, for all $-t \in [-T, 0]$, the limits $P(t;\infty)$ and $D(\infty)$ exist, and*

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

if $r > 0$, and $P(t;\infty) = \mathbf{1}_{\{t=T\}}$ and $D(\infty) = 0$ if $r = 0$.

Proof Sketch. Here, we consider the case with $r > 0$. A formal proof is given in Appendix D.5. To show the result, we prove

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

where $p(t) = \mu(A(v^*(t)))$. To see this, notice that by the ODE (1), $v^{*\prime}(t) = \lambda(g(A(v^*(t))) - v^*(t)) \cdot p(t)$. Since $\frac{\partial \mu(A(v))}{\partial v_i}$ exists by Assumption 3 (b), $p(t) = \mu(A(v^*(t)))$ is differentiable, and we obtain

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

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 $t \to \infty$, which exists by Assump-

[38]Assumption 3 (b) is not crucial. We argue in Appendix D.5 that a more elaborate definition of r enables us to rule out Assumption 3 (b).

[39]If Assumption 3 (b) fails, the duration formula we will present in Theorem 4 may fail even when r is well-defined in (3). For example, consider the one-dimensional uniform distribution on $X = \{(x_1, x_2) \in \mathbb{R}^2 \mid 0 \le x_1 = x_2 \le 1\}$. Since this is equivalent to a single-agent problem with the uniform distribution on $[0, 1]$, we can obtain $D(\infty) = 1/3$. However, according to the definition of r in (3), we would obtain $r = 1$, with which the duration formula would yield a false value of $D(\infty) = 1/2$.

[40]This probability is computed by $P(t;\lambda) = \exp\Big(-\int_t^T \lambda p(s;\lambda)\, ds\Big)$.

tion 3 (a). Thus, for any $\varepsilon > 0$, there exists $\bar{t}$ such that $t \geq \bar{t}$ implies

$$r - \varepsilon \leq -\frac{p'(t)}{\lambda p(t)^2} \leq r + \varepsilon. \tag{4}$$

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 $\lim_{\lambda \to \infty} p(t) \cdot \lambda t = \frac{1}{r}$. This equality enables us to compute the approximated probability of disagreement P, and show that $P(t; \infty) = \left(\frac{t}{T}\right)^{\frac{1}{r}}$ and $D(\infty) = \frac{1}{1+r^{-1}}$.[41] $\square$

In the above discussion, we used the fact that the expected duration is determined by the evolution of the acceptance probability, and this probability p is a function of the continuation payoff profile v whose evolution (i.e., v') we know from ODE (1) is determined in equilibrium by the acceptance probability p itself. This means that, in equilibrium, the evolution of the acceptance probability (i.e., p') depends on the acceptance probability p—this led to the idea of using ODE (1) to derive an ODE with respect to p.

Theorem 4 immediately implies the following: If $r > 0$, the limit probability density of agreement $P'(t; \infty)$ exists and is strictly positive for all $-t \in [-T, 0)$, and (a) if $0 < r < 1$, $P'(t; \infty)$ is strictly increasing in t and $\lim_{t \to 0} P'(t; \infty) = 0$, (b) if $r = 1$, $P'(t; \infty) = \frac{1}{T}$ for all $-t \in [-T, 0)$, and (c) if $r > 1$, $P'(t; \infty)$ is strictly decreasing in t and $\lim_{t \to 0} P'(t; \infty) = \infty$.

An intuitive explanation of the duration formula.

In Examples 1 and 2 preceding Theorem 4, we argued that there are two channels through which preference heterogeneity affects the limit expected duration. The first channel (the probability of the extra region) corresponds to $d_i(v^*(t))$ and the second channel (the expected gain relative to the cutoff conditional on acceptance) corresponds to $b_i(v^*(t))$. Formula (1) implies that the cutoff falls with the speed proportional to such an expected gain. These facts imply that (3) can be expressed as follows:

$$r \propto \lim_{t \to \infty} \sum_{i \in N} \left(\left[\begin{array}{c} \text{the probability} \\ \text{of the extra region} \end{array} \right] \times \left[\begin{array}{c} \text{the speed of} \\ \text{cutoff decrease} \end{array} \right] \right). \tag{5}$$

Now, the product of the marginal increase of the probability of the extra region (with the unit of measurement ($\frac{\text{probability}}{\text{distance}}$)) and the speed of cutoff decrease (with the unit of measurement ($\frac{\text{distance}}{\text{time}}$)) is equal to the speed with which the probability of acceptance

[41] The computation is given in Lemma 21 in Appendix D.1.

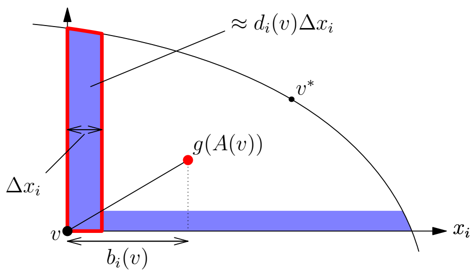

Figure 9: Density term and barycenter term

increases (with the unit of measurement ($\frac{\text{probability}}{\text{time}}$)). It is intuitive that this speed determines the search duration. □

The graphical intuition for the formula in Theorem 4 is depicted in Figure 9. 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. Recall that the speed with which the cutoff moves towards the limit point is determined by this distance, by equation (1). Hence the formula for r in equation (3) measures the speed with which the acceptance probability falls. This is consistent with the fact that the duration formula in equation (3) is increasing in $r > 0$ because if the acceptance probability falls quickly for a given cutoff level, then players reject with high probability for a long time, resulting in a long duration (as in Figure 8).

Theorem 4 also explains the reasoning behind the formula in footnote 34. Under Assumption 3, r is well-defined in (3) in each of models (i) and (ii). Let r_μ and r_γ be associated with models (i) and (ii), respectively. Then, since model (iii) considers the product measure, $r_{\mu\gamma}$ in model (iii) exists and equals $r_\mu + r_\gamma$. Hence the limit expected duration in model (iii) exists and is $\frac{1}{1+r_{\mu\gamma}^{-1}} = \frac{1}{1+(r_\mu+r_\gamma)^{-1}}$. Rearranging terms, we obtain the formula in footnote 34, which leads to the conclusion of Theorem 3 under Assumption 3.

Now we use the formula given in Theorem 4 to analyze a specific class of games to understand the preference heterogeneity effect further.

Example 3 (Applying the Duration Formula).

Here we impose assumptions employed often in the literature on multi-agent search (Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013)).

Assumption 4. (a) X is a convex and compact subset of $\mathbb{R}^n_+$, and has a smooth Pareto frontier.[42]

(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 strictly positive and continuous on X.

[42] We say that the Pareto frontier is smooth if it can be defined by an implicit function that is continuously differentiable.

Assumption 4 allows us to explicitly compute the limit expected duration as follows. To make the dependence of the expected duration on n explicit, here we denote by $D(\infty; n)$ the limit expected duration as $\lambda \to \infty$ given n:

Proposition 5. *Under Assumptions 1 and 4,* $D(\infty; n) = \dfrac{n^2}{n^2 + n + 1}$.

Corollary 6. *Under Assumptions 1 and 4,* $D(\infty; n)$ *is increasing in* n.

The solution of the limit expected duration provided in Proposition 5 implies that, if only two players are involved in search, the limit expected duration is $\frac{4}{7}T$, and it monotonically increases to approach T as n gets larger. The second row of Table 1 in the Introduction shows the limit expected duration for several values of n when $T = 1$.

Let us give a proof idea.[43] When the continuation payoff profile v is sufficiently close to the Pareto frontier, the acceptance set can be approximated by an n-dimensional pyramid by the assumption of the smooth Pareto frontier, and the distribution over this acceptance set can be approximated by the uniform distribution due to the assumption of the strictly positive and continuous density. Therefore, we can compute the limit expected duration by computing that in the n-dimensional pyramid with the uniform distribution. So for a fixed $v \in X$, assume for now that $A(v)$ is the exact pyramid and the conditional distribution over $A(v)$ is uniform. Let $\hat{v}_i = \max\{x_i \mid (x_i, v_{-i}) \in X\}$. We can compute $b_i(v)$ and $d_i(v)$ as follows: Since $b(v)$ is the vector from v to the barycenter of the n-dimensional pyramid, $b_i(v) = \dfrac{\hat{v}_i - v_i}{n+1}$. By the definition of d_i,

$$d_i(v) = -\frac{\partial \mu(A(v))/\partial v_i}{\mu(A(v))} = \frac{f(v)\prod_{j \neq i}(\hat{v}_j - v_j)}{f(v) \cdot \frac{1}{n}\prod_{j \in N}(\hat{v}_j - v_j)} = \frac{n}{\hat{v}_i - v_i}.$$

Therefore, $r = \sum_{i \in N} d_i(v) b_i(v) = \dfrac{n^2}{n+1}$, so $D(\infty; n) = \dfrac{1}{1 + \left(\frac{n^2}{n+1}\right)^{-1}} = \dfrac{n^2}{n^2 + n + 1}$.

Under Assumption 4, the ascending acceptability effect can be seen in Figure 9 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 piece of payoff regions that will be accepted in the future. The probability density in the extra region conditional on the acceptance set increases not only because the number of segments increases, but also because the length of each segment increases. This happens precisely because players' preferences become heterogeneous so the density of the marginal distribution is larger for low payoffs than for high ones in the acceptance set. 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 decreases due to this preference heterogeneity

[43] In the proof in Appendix D.6, we show this result under more general assumptions (Assumptions 1 and 8).

as well, the overall effect is positive. We call this effect the preference heterogeneity effect. Note that the role of the preference heterogeneity effect is to determine the magnitude of the ascending acceptability effect. Mathematically, the preference heterogeneity effect affects the values of summands in (5), while the ascending acceptability effect increases the number of summands in (5). □

The duration formula also applies to Examples 1 and 2. In Example 1, let $\hat{v}_i = \max\{x_i \mid (x_i, v_{-i}) \in X\}$. Note that symmetry implies $v_1^*(t) = v_2^*(t)$ for all t. Suppose that $v \in X$ satisfies $v_1 = v_2$. Then, a straightforward computation shows $b_i(v) = \frac{2q+1}{6}(\hat{v}_i - v_i)$, and $d_i(v) = \frac{1}{q(\hat{v}_i - v_i)}$. Since $d_i(v)b_i(v) = \frac{2q+1}{6q}$ and this is constant in v, $r = \lim_{v_1=v_2\to q}\sum_{i\in N} d_i(v)b_i(v) = \sum_{i\in N}\frac{2q+1}{6q} = \frac{2q+1}{3q}$, and $D(\infty) = 1/\big(1 + \big(\frac{2q+1}{3q}\big)^{-1}\big) = \frac{2q+1}{5q+1}$. In Example 2, again symmetry implies $v_1^*(t) = v_2^*(t)$ for all t. As discussed in Example 2, $b_i(v)$ is increasing in σ for each $v = (v_1, v_2)$ with $v_1 = v_2$.[44] Since the conditional distribution on $A(v)$ (with a normalized origin at v) is independent of v with $v_1 = v_2$, we have $d_i(v) = d_i(0) = \frac{1}{M_\sigma}\int_0^\infty \frac{1}{\sqrt{2\pi\sigma^2}} e^{-x_i} e^{-\frac{x_i^2}{2\sigma^2}}\, dx_i = \frac{M_\sigma}{M_\sigma} = 1$ for all σ. Therefore r is increasing in σ, giving us the comparative statics that we explained in Example 2.[45]

4.3 Step 3: Finite Arrival Rates

To evaluate the significance of the effects that we identify in the previous discussion, we now consider cases with finite arrival rates. We will show that *the expected duration converges to its limit fast*, providing evidence that our limit analysis contains economically-meaningful content—so the effects in Steps 1 and 2 are the keys to understanding the duration in reality.

First, we show that the convergence speed of the expected 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 5. *Under Assumptions 1 and 3,* $|D(\lambda) - D(\infty)| = O(\frac{1}{\lambda})$.

This is a fast rate of convergence. In particular, it means that, although the expected number of offers until acceptance diverges to infinity as $\lambda \to \infty$, the expected number of offers with arrival rate λ in the time interval between $D(\lambda)$ and $D(\infty)$ is bounded above by a finite number. When payoffs realize upon agreement and there is a positive discount rate (with a finite horizon as in Appendix A.1 or with an infinite horizon), $|D(\lambda) - D(\infty)|$ is of the same order as $\frac{1}{\lambda^{\frac{1}{n+1}}}$ under Assumptions 1 and 4.

Moreover, as in Remark 2 (b) in Section 4.2.1, we can show a stronger conclusion than in the theorem. Specifically, for any $q \in [0, 100]$, the q-percentile of the duration

[44] A computation shows that $b_i(0) = \dfrac{\sqrt{\sigma^2}}{2e^{\frac{\sigma^2}{2}}\int_\sigma^\infty e^{-\frac{z^2}{2}}\,dz} + \dfrac{1-\sigma^2}{2}$.

[45] The limit distribution as $\sigma \to 0$ does not have a density function on $\mathbb{R}^n$ and violates Assumption 3 (b). However, we provide a fully general formula in Appendix D.5 to cover such a case.

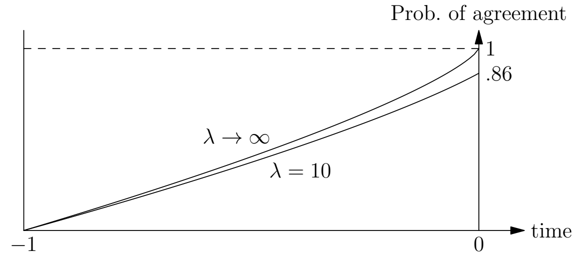

Figure 10: A numerical example of the cumulative probability of agreement in Case 1 with $n = 2$.

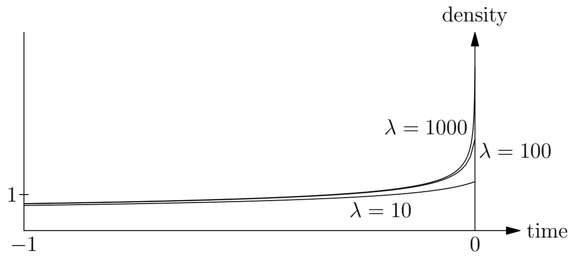

Figure 11: The probability density of the time of agreement in Case 1 with $n = 2$.

distribution with arrival rate λ is different from its limit by $O(\frac{1}{\lambda})$. That is, not only the expected duration but also the duration distribution converges fast.

We further support our claim numerically through a number of examples. We find that the limit expected duration of Theorem 4 is not far away from those with finite λ in many cases. The differential equation (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 integrations. For this reason, we solve the differential equation and integration numerically to obtain the values of $D(\lambda)$ for specific values of λ.[46] We considered the following distributions standard in the literature with $T = 1$.

Case 1: μ is the uniform distribution over $X = \{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i \leq 1\}$.

Case 2: μ is the uniform distribution over $X = \{x \in \mathbb{R}^n_+ \mid \sum_{i \in N} x_i^2 \leq 1\}$.

Case 3: μ is the product measure over $X = \mathbb{R}^n_+$ where each marginal corresponds to an exponential distribution with parameter $a_i > 0$.

In the apartment search example in the Introduction, if the couple has ten weeks before the deadline and a broker provides information of an apartment once per week on average (a very *in*frequent case), the situation corresponds to $\lambda = 10$. Figure 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 11 shows the probability density function of the time of agreement in that case (i.e., $P(t; \lambda) \cdot p(t; \lambda)$). In Table 2, we provide the computed values for selected parameter values. We provide the complete description of all the computed values in Appendix C.[47]

According to our calculation, $D(\lambda)$ is within 10% difference from $D(\infty)$ except for a single case where the difference is 19.4%, which happens in Cases 1 and 2 with $n = 1$. Generally, the percentage falls as the number of agents becomes larger and the arrival

[46] Some cases can be computed analytically, e.g., Cases 1 and 3 below.

[47] In Appendix C, we also consider other cases: the uniform distribution over a cube and the log-normal distribution.

			λ					
			10	20	30	100	1000	∞
Case 1	$n = 1$	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	0.608	0.591	0.585	0.576	0.572	0.571
		Percentage (%)	6.48	3.44	2.35	0.731	0.0716	0
Case 2	$n = 1$	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 3	$n = 1$	Expected duration	0.545	0.524	0.516	0.505	0.500	0.5
		Percentage (%)	9.09	4.76	3.23	0.990	0.0999	0
	$n = 2$	Expected duration	0.693	0.681	0.676	0.670	0.667	0.667
		Percentage (%)	3.91	2.11	1.45	0.465	0.0489	0

Table 2: Expected durations for finite arrival rates. "Percentage" is defined by 100×(the duration difference)/(the limit expected duration).

rate goes up.[48] For example, if we add another player in Case 1, the difference falls down dramatically to 6.5%. 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%. Notice that we predict "over-shooting" of the expected duration in Case 2. 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 as t diverges from 0 (precisely, the preference heterogeneity effect *would* be smaller than in the case of a triangle *if* the limit shape of the acceptance set *were* proportional to that of X). This suggests that convexity of the set of available allocations, which is often assumed in the literature, facilitates a fast convergence. Case 3 considers distributions with an unbounded X. As in Case 1, the distribution conditional on the acceptance set is a geometric translation of the original distribution irrespective of the cutoff profile. In fact, Table 2 shows that the computed values for Case 3 present the same trend as in Case 1.

5 Welfare Implications

In Section 3, we showed that the limit expected payoff profile must be weakly Pareto efficient if the limit exists. In this section we seek further welfare implications. We will

[48]The monotonicity with respect to arrival rates can be analytically proven in many cases, e.g. 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 expected duration cannot be decreasing in λ. Note also that in Case 2 with $n = 2$, after the "overshooting" that we will explain shortly, the expected duration comes back to the limit. Thus $D(\lambda)$ is nonmonotonic also in this case.

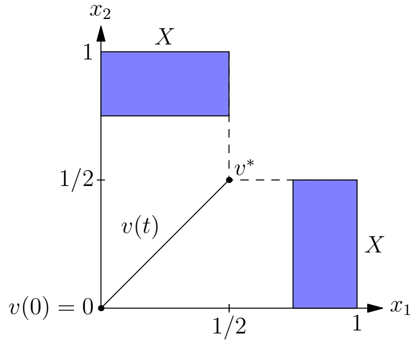

Figure 12: A path that converges to a weakly Pareto efficient allocation.

prove that the limit payoff profile is Pareto efficient in a wide class of payoff distributions, and argue that a further prediction is hard to obtain. Let us impose the following assumption to rule out cases that are not interesting for welfare analysis:

Assumption 5. (a) X is a compact subset of $\mathbb{R}^n$.

(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 strictly positive and continuous on X.

Condition (a) in Assumption 5 is a standard assumption when we consider welfare implications. Note that we do not assume convexity here. Condition (b) rules out irregularity involving subsets with zero Lebesgue measure.

In general, v^* is not necessarily Pareto efficient in X even if it exists. There is an example of a distribution μ satisfying Assumptions 1 and 5 in which $v^*(t)$ converges to an allocation that is not Pareto efficient.

Example 4. Let $n = 2$, $X = \big([0, 1/2] \times [3/4, 1]\big) \cup \big([3/4, 1] \times [0, 1/2]\big)$, and suppose f is the uniform density function on X, which is shown in Figure 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.[49] □

Note that v^* is weakly Pareto efficient (as implied by Proposition 4), and that X is a non-convex set in this example. In fact, we can show that v^* is Pareto efficient if X is convex.

[49] 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 limit expected payoff profiles supported by subgame perfect equilibria. However, we rule out such subgame perfect equilibria in a view that rejecting everything after a deviation is not a credible threat if a player expects the others to accept with a small probability.

Proposition 7. *Suppose that X is convex. Under Assumptions 1 and 5, v^* is Pareto efficient in X.*

For example, when $n = 2$ and $X = [0,1]^2$, this proposition says that $v^*(t)$ cannot converge to, e.g., $(1/2, 1)$. This is because if $v^*(t)$ is sufficiently close to $(1/2, 1)$, the slope of $v^{*\prime}(t) \le \frac{f_H(1-v_2^*(t))}{f_L(1-v_1^*(t))}$ is significantly smaller than the slope of the line connecting $v^*(t)$ and $(1/2, 1)$ $(= \frac{1-v_2^*(t)}{1/2-v_1^*(t)})$, where $f_L := \min_{x\in X} f(x)$ and $f_H := \max_{x\in X} f(x)$ satisfying $0 < f_L \le f_H < \infty$ by Assumption 5. Such comparison of slopes generalizes to the case with general distributions satisfying Assumptions 1 and 5 with convex X. Even when X is not convex, we can show that the limit payoff profile is Pareto efficient for "generic" distributions when $n = 2$. Appendix A.5 formalizes what we mean by genericity, and explains the intuition for why such a result holds.

Pareto efficiency of the limit payoffs implies that players reach an agreement with probability close to one 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 $-t$. By definition, $\pi(t)$ is nondecreasing and bounded in λ, so $\lim_{\lambda\to\infty} \pi(t)$ exists. The expected continuation payoff profile $v^*(t)$ 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 that $(1/\pi(t))v^*(t) \in \{x \in co(\hat{X}) \mid x \ge v^*(t)\}$. Since v^* is Pareto efficient in X which is closed, the set $\{x \in co(\hat{X}) \mid x \ge v^*(t)\}$ shrinks to a singleton $\{v^*\}$ as $v^*(t)$ goes to v^*. Therefore we have $(1/\pi(t))v^*(t) \to v^*$ as $\lambda \to \infty$. This implies $\lim_{\lambda\to\infty} \pi(t) = 1$ for all $t > 0$ because $\lim_{\lambda\to\infty} v^*(t) = v^*$ for all $t > 0$. That is, we have the following proposition:

Proposition 8. *Suppose that Assumption 1 holds. If v^* is Pareto efficient, then the probability of agreement before the deadline converges to one as $\lambda \to \infty$.*

We note that the conclusion of this proposition fails if v^* is only weakly Pareto efficient. In Example 4, players reach no agreement before the deadline with positive probability: Since the limit expected payoff profile is $(1/2, 1/2)$, for a fixed λ and any $\varepsilon > 0$, there exists $\bar{t}$ such that at any time $-t \le -\bar{t}$, the cutoff of each agent is higher than $1/2 - \varepsilon$. Thus, a player's expected payoff conditional on agreement at time $-t \le -\bar{t}$ is larger than $\big(7/8 + (1/2 - \varepsilon/2)\big)/2 = 11/16 - \varepsilon/4$. Suppose that $\lim_{T\to\infty} \pi(T) = 1$ for λ. Since $\pi(\bar{t}) < 1$ for λ, the probability of agreement in subinterval $[-T, -\bar{t}]$ must converge to one as $T \to \infty$. Therefore each player's limit unconditional expected payoff is larger than $11/16 - \varepsilon/4$ for any $\varepsilon > 0$. This contradicts the fact that the limit expected payoff is $1/2$. Hence we must have $\lim_{T\to\infty} \pi(T) < 1$, implying that the limit probability of agreement as $\lambda \to \infty$ does not converge to 1 either.

The reader may wonder which point on the Pareto frontier the continuation payoff profile converges to. In fact, this question is difficult to answer: An example is shown in Appendix A.5 in which player i's marginal distribution of a probability measure μ is first-order stochastically dominated by that of another probability measure γ, and i's limit expected payoff under μ exceeds the one under γ. This happens because the limit payoff profile is determined by the *joint* distribution over X, and player j's marginal distribution may have improved under γ compared to μ, to a larger degree than i's improvement. This suggests difficulty of characterizing a general property of limit payoff profiles. In fact, we can show that most points on the Pareto frontier can be reached in the limit under some probability measure. Formally, for any Pareto efficient payoff profile w in X which is not at the edge of the Pareto frontier,[50] we show that there exists a density f that satisfies Assumptions 1 and 5 such that the limit of the solution $v^*(t)$ of equation (1) is w.

Proposition 9. *Suppose that $X \subseteq \mathbb{R}^n_+$ satisfies Assumption 5 (a). 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 5 hold, and $\lim_{\lambda\to\infty} v^*(t) = w$ for all $t \in (0, T]$.*

In the proof, we introduce a family of probability density functions $\{f_y\}$ having a large weight near each Pareto efficient $y \in X$ in the neighborhood of $w \in X$, and construct a function that maps the difference between w and y to the difference between w and the limit point given f_y. We then use a fixed point theorem to show existence of y such that the value of the constructed function is zero. Note that Proposition 9 is not so obvious as it may appear because the result pertains to the *limit* payoff profile, so continuity of the solution of an ODE in its parameter cannot be used for the proof. Indeed, we will see in Appendix A.1 that the limit is independent of density f if there is a positive discount rate $\rho > 0$, as long as certain assumptions hold.

The lesson here is that, if the payoffs realize at the deadline, then the limit allocations depend on distributions so much so that any Pareto-efficient allocation is possible, under the assumptions given. Despite such indeterminancy, this section has provided a general prediction such as Pareto efficiency under a wide range of distributions.

6 Discussions

The analysis in the main sections illuminated key incentive issues in the presence of multiple players and a finite horizon. Under the presence of these two features, there are many ways to extend and/or modify the model. Appendix A discusses a number of topics on such extensions/modifications. Here we preview the most interesting parts of such discussions, by which we aim to invite the reader to the exploration that we pursue in the Appendix.[51] These discussions not only help understand the key assumptions and techniques used in the main analysis, but also highlight the wideness of the variety of questions we can ask in the context of multi-agent search with deadline.

[50] We formally define this property in Appendix D.10.

[51] Among others, Appendices A.1–A.4 present full versions of the discussions on the topics that we consider in this section.

The Payoffs Realizing upon Agreement

Economic search situations are nontrivial to analyze because agents face a tradeoff between deciding now and doing so later. The benefit of deferring the acceptance of an offer lies in a prospect of getting a better payoff in the future, while the cost comes with some form of penalty on an act of deferring. There are two types of such penalty that economic agents face in reality: discounting (or time costs) and deadlines. The traditional infinite horizon models and our finite horizon model are at two extreme points: the former assumes discounting or time costs without deadlines, while the latter (our main sections) assumes the risk of reaching the deadline without assuming discounting or time costs.[52] We believe the reality is at somewhere in the middle of these two extreme cases, and which model is a better approximation of the reality depends on the particular applications that the modeler is interested in.[53] To understand these "middle" situations, we think it a necessary step to analyze both of the two polar cases. We already know much about one of these cases from the literature, and the other case is what we analyzed in the main sections.

Here we present a way to formally connect these two cases. Specifically, in Appendix A.1 we consider the case where the payoffs realize as soon as an agreement is reached, as opposed to assuming that the payoffs realize only at the deadline. Before considering "middle" cases, we first consider one of the polar cases in which we fix the discount rate $\rho > 0$ and then take the limit as the arrival rate λ tends to infinity. We show that the path of the continuation payoff profile is close to that in the case of our main model when λ is sufficiently high and the deadline is relatively close, while it diverges from such a path when the deadline is far away. Under technical assumptions, the limit payoff profile is shown to be an element of the "Nash set," a generalization of the Nash bargaining solution. This result shows a certain robustness of the limit payoffs on distributional assumptions under discounting. At the first glance, this may look at odds with our results for the case with no discounting where the limit equilibrium payoffs are sensitive to the payoff distribution. However, when the Nash set consists of multiple points, we argue that *which point in the Nash set becomes the limit equilibrium payoff profile depends on the distribution of offers.* In determining a point in the Nash set, the resemblance of the path under $\rho > 0$ with the one under $\rho = 0$ becomes useful. In contrast, in the infinite-horizon model, all points in the Nash set arise as limits of equilibrium payoff profiles.[54]

We also show that the limit expected search duration is zero when we fix $\rho > 0$ and

[52] We will review the literature on the infinite horizon models in Appendix B.

[53] No discounting may be the best model when the payoffs realize at the deadline as in our motivating example of apartment search, although one may still be able to imagine a possibility of small time costs.

[54] Although Wilson (2001) considers only convex feasible payoff sets, his argument can be generalized to the non-convex cases. Cho and Matsui (2013, Proposition 4.3) prove a related result with non-convex feasible payoff sets.

then let $\lambda \to \infty$. One might present a critique attributing this to lack of robustness of our result that the limit expected duration as $\lambda \to \infty$ is strictly positive for fixed $\rho = 0$ and T. Our reaction would be that the limit as $\lambda \to \infty$ with fixed $\rho > 0$ implies that there are an increasingly large number of offers "before payoffs are substantially discounted" and thus the deadline is not at the relevant future, so the situation is getting more and more similar to the one represented by the infinite horizon model. However, if we consider a situation where the payoffs are not so much discounted even at the deadline, i.e., ρ is relatively small compared to λ, then the deadline is in the relevant future, so the result of our finite horizon model is more relevant.

In order to understand what these "situations" would refer to, we consider the "middle" cases. Specifically, we consider a general model in which ρ depends on λ, and quantify the speeds of simultaneous convergence as $\lambda \to \infty$ and $\rho \to 0$ such that our results are relevant. We show that the limit payoffs and durations depend on the limit of $\lambda\rho^n$. Since this is a continuous function in both λ and ρ, the result implies that the relevance of the two extreme analyses—our main analysis and the analysis in which $\lambda \to \infty$ with fixed $\rho > 0$—depends on what situations we want to analyze. Specifically, if ρ is relatively large compared to $1/\lambda^{1/n}$, then the deadline is not in the relevant future. On the other hand, if ρ is relatively small compared to $1/\lambda^{1/n}$, it is in the relevant future so the analyst may want to refer to the result of our main analysis.

A further critique to our reaction to distinguish between different situations would be to say why our limit result is relevant when the justifying argument concerns finite λ. This is where our step 3 kicks in, which states that the convergence speed of the expected search duration is high with respect to λ.

All these results suggest a wide applicability of our results for the case in which payoffs realize at the deadline, and we believe the analysis of the case in which payoffs realize upon agreement complements our main analysis.

Market Designer's Problem

In the main sections, we took a search environment as given and analyzed equilibrium behaviors of the players. Let us now step back and consider problems faced by a market designer who has a control over certain parameters of the model. In particular, we consider two ways by which the designer can affect the search environment. The first is an adjustment of the horizon length T. That is, the designer commits to a horizon length T first, and then players start searching until the length T of time passes. The second is the possibility in which the designer can instead affect the probability distribution over potential payoff profiles, by "holding off" some offers. Formally, given μ, we let the designer choose a measure μ' such that $\mu'(Y) \leq \mu(Y)$ for all Borel subsets $Y \subseteq X$.[55]

In Appendix A.2, we argue that these two ways of designing the market create different

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

effects on the equilibrium behaviors depending on the timing of the payoff realizations. This implies that it would be of the designer's interest to look at the specific application at hand to see what options, if any, to adjust the environment should be utilized. The discussion in Appendix A.2 gives a recipe for such an adjustment.

General Voting Rules

In the main sections, we considered the case when players use the *unanimous rule* for their decision making. This is a reasonable assumption in many applications such as the apartment search, but there are other applications in which different voting rules (e.g., majority rules) may fit the reality better.

In order to model general voting rules, we suppose that $\mathcal{C} \subseteq 2^N$ is the set of winning coalitions, and the object of search is accepted if and only if there is a winning coalition $C \in \mathcal{C}$ in which every player says "accept" upon its arrival. A minimal winning coalition is a coalition $C \in \mathcal{C}$ such that if $C' \subseteq C$ and $C' \in \mathcal{C}$, then $C' = C$. We assume that any player can be pivotal, i.e., for all $i \in N$, there exists a minimal winning coalition $C \in \mathcal{C}$ with $i \in C$. The voting rule naturally induces a coalitional-form game with non-transferable utility, (N, V), where the characteristic function V is defined as $V(C) = X$ if $C \in \mathcal{C}$, otherwise $V(C) = \{0\}$. The core is defined in the standard manner, and in particular it equals the weak Pareto frontier of X in the case of the unanimous rule ($\mathcal{C} = \{N\}$).

In this setting, we argue that the welfare implication and the search duration critically depend on the nature of the induced coalitional-form game. In particular, in Appendix A.3, we show the following result: if X is compact and convex, the core is nonempty if (and only if) there exists μ with support X satisfying Assumption 1 such that the limit expected duration is positive and v^* is weakly Pareto efficient. The intuition is that if the limit payoff profile is not in the core, then some players forming a winning coalition can expect better payoffs than the limit payoffs by accepting those better payoff profiles. Since these "better payoff profiles" are agreed upon with positive probability conditional on having an opportunity at any time, the limit expected duration is zero, and the convexity of X implies Pareto inefficiency of the limit expected payoff profile.

Intuition for Essential Uniqueness of Trembling-Hand Equilibrium

The (essential) uniqueness of trembling-hand equilibrium (Proposition 1) is important as it facilitates unambiguous comparative statics. Its proof, however, is nontrivial as we work with continuous time so the standard backward-induction argument does not apply. In a single-player search model, if two strategies give rise to two different continuation payoffs at some time $-t$ then the strategy with the lower continuation payoff is obviously suboptimal, and this trivially implies uniqueness. This proof cannot be used for the case of two or more players. For example, it might be the case that in one equilibrium player 1 is picky and player 2 is generous, while in another equilibrium the opposite happens, and these two are both equilibria as both imply reasonable levels of acceptance probabilities at each point in time.

The proof consists of two steps. In the first step, we bound the supremum difference of continuation payoffs at time $-t$ across all the trembling-hand equilibria using those for time $-\tau \in (-t, 0]$. Then we show in the second step that such bounds at all time $-t$ imply that such differences are zero at all time $-t$.

The key idea for why these differences are zero is as follows: In order to create a difference in the current continuation payoff, one needs an enough variation in the future continuation payoffs. But if the remaining time is short, the variation should be large in absolute term, which is impossible because of our first step, that is, we do not have full flexibility to vary their strategies due to the "trembling-hand" restriction.[56] Appendix A.4 explains the idea in more detail.

It is an open question whether a similar method applies in the infinite-horizon problem to rule out sequential equilibria that do not use cutoff strategies. The second step of the proof would not immediately apply due to lack of a deadline, but still discounting or time costs that would be present in infinite-horizon models might perform the role of a deadline. Certainly, the contrived sequential equilibria constructed in Cho and Matsui (2013, Section 4.7) would not constitute a trembling-hand equilibrium even under infinite-horizon models.

7 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 determinants of the positive duration that we observe in reality. We first showed that the (well-defined) expected 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. Second, when there are multiple agents, this limit expected duration increases as a result of two effects: the ascending acceptability 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 ones for the player due to heterogeneity of preferences. Third, we showed that the convergence speed of the expected duration as the friction vanishes is high, and numerically demonstrated that expected

[56]In this sense, the proof idea is somewhat similar to that of Proposition 11 in Kamada and Kandori (2011), in which they show that in revision games there is a unique equilibrium if the payoff function satisfies a certain regularity condition. See footnote 72 in Appendix A.4 for more details on this issue.

durations with positive frictions are reasonably close to the limit expected 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 understanding the duration in reality.

We also conducted a welfare analysis, and showed that the limit expected payoff profile is Pareto efficient under a wide class of payoff distributions, and depends on the distribution of offers. Lastly, we provided a wealth of discussions both in the main text and in Appendix A to examine the robustness of our main conclusions and to analyze a variety of alternative specifications of the model.

The combination of multiple agents and a finite horizon introduced technical difficulties that we needed to handle. In order to obtain a meaningful insight in such a challenging environment, we took two approaches to the problem. To show that the positive duration holds generally, we did not use Bellman equations but considered bounds of payoffs to partially identify equilibrium behaviors. This approach enabled us to understand a wide range of distributional environments in a single proof. We then closely examined the differential equation that characterizes the equilibrium continuation payoffs to derive a key measure of duration r. This measure r is easy enough to compute in many applications, and we used it to derive the duration formula. The use of differential equations is inherent under the non-stationary environment with finite horizon. The comparative statics based on these analyses has meaningful contents because we show that the equilibrium is unique—the uniqueness was obtained by using the "trembling-hand" refinement, which is new to this sort of setting. The proof is nontrivial as there can be an indefinite sequence of punishments in our continuous-time setting, and it is potentially applicable to other settings involving indefinite sequences of punishments such as infinite-horizon problems.[57]

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 always be T, but the duration until the first agreement is not obvious. This is 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, we analyze this case and find that under certain assumptions, the expected duration until the first acceptance is positive even in the limit as the friction vanishes. In that work, 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 randomly matched and some payoff

[57] Kamada and Sugaya (2014) use our technique to prove uniqueness of the equilibrium in their model.

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 an overlapping generation model with a constant inflow of agents where different agents face different deadlines. In our ongoing research, we solve for a steady-state equilibrium strategy and characterize the 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 distribution of payoffs must change 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 result carries over.[58] Our result on time-varying distributions in Appendix A.11 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 (that perhaps resolves over time) about the distribution over outcomes or the opponents' preferences. We conjecture such uncertainty would increase search durations because such uncertainty adds an option value of waiting. 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. We hope the current paper serves as a basis for such future work and provides useful insights and techniques for the analyses in such work.

References

Abdelaziz, F. and S. Krichen (2007): "Optimal stopping problems by two or more decision makers: a survey," *Computational Management Science*, 4, 89–111.

Albrecht, J., A. Anderson, and S. Vroman (2010): "Search by committee," *Journal of Economic Theory*, 145, 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, 377–394.

Alpern, S., S. Gal, and E. Solan (2010): "A sequential selection game with vetoes," *Games and Economic Behavior*, 38, 1–14.

[58]Baughman (2014) considers a search model with two-sided large populations in which every agent shares a common deadline. His model has a payoff structure different from ours, and is not a generalization of the model we propose here, and vice versa.

Ambrus, A., J. Burns, and Y. Ishii (2014): "Gradual bidding in ebay-like auctions," Mimeo.

Ambrus, A. and S.-E. Lu (2010): "Legislative bargaining with long finite horizons," Mimeo.

——— (2014): "A continuous-time model of multilateral bargaining," Forthcoming in *American Economic Journal: Microeconomics*.

Aumann, R. J. (1964): "Mixed and Behavior Strategies in Infinite Extensive Games," in *Advances in Game Theory*, ed. by M. Dresher, L. S. Shapley, and A. W. Tucker, Princeton University Press, vol. 52 of *Annals of Mathematics Studies*, 627–650, reprinted at `http://www.ma.huji.ac.il/%7Eraumann/pdf/Mixed%20and%20Behavior.pdf`.

Baughman, G. (2014): "Deadlines and Matching," Mimeo.

Bergemann, D. and J. Välimäki (2011): "Efficient Search by Committee," Mimeo.

Calcagno, R., Y. Kamada, S. Lovo, and T. Sugaya (2014): "Asynchronicity and coordination in common and opposing interest games," *Theoretical Economics*, 9, 409–434.

Cho, I.-K. and A. Matsui (2013): "Search Theory, Competitive Equilibrium, and the Nash Bargaining Solution," *Journal of Economic Theory*, 148, 1659–1688.

Coddington, E. A. and N. Levinson (1955): *Theory of Ordinary Differential Equations*, McGraw-Hill.

Compte, O. and P. Jehiel (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, 189–221.

Ferguson, T. S. (1989): "Who solved the secretary problem?" *Statistical Science*, 4, 282–296.

——— (2005): "Selection by Committee," in *Advances in Dynamic Games*, ed. by A. S. Nowak and K. Szajowski, Birkhäuser Boston, vol. 7 of *Annals of the International Society of Dynamic Games*, 203–209.

Fudenberg, D. and D. Levine (1983): "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31, 251–268.

Gomes, A., S. Hart, and A. Mas-Colell (1999): "Finite Horizon Bargaining and the Consistent Field," *Games and Economic Behavior*, 27, 204–228.

Heckman, J. J. and B. Singer (1984): "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 52, 271–320.

Herings, P. J.-J. and A. Predtetchinski (2014): "Voting in Collective Stopping Games," Mimeo.

Herrero, M. J. (1989): "The Nash Program: Non-Convex Bargaining Problems," *Journal of Economic Theory*, 49, 266–277.

Imai, H. and H. Salonen (2012): "A characterization of a limit solution for finite horizon bargaining problems," *International Journal of Game Theory*, 41, 603–622.

Ishii, Y. and Y. Kamada (2011): "The Effect of Correlated Inertia on Coordination," Mimeo.

Kamada, Y. and M. Kandori (2009): "Revision Games," Mimeo.

——— (2011): "Asynchronous Revision Games," Mimeo.

Kamada, Y. and N. Muto (2011): "Multi-Agent Search with Deadline," Mimeo.

Kamada, Y. and T. Sugaya (2010): "Asynchronous Revision Games with Deadline: Unique Equilibrium in Coordination Games," Mimeo.

——— (2014): "Valence Candidates and Ambiguous Platforms in Policy Announcement Games," Mimeo.

Kiefer, N. M. (1988): "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 26, 646–679.

Kiefer, N. M. and G. R. Neumann (1979): "An Empirical Job-Search Model, with a Test of the Constant Reservation-Wage Hypothesis," *Journal of Political Economy*, 87, 89–107.

Lippman, S. A. and J. J. McCall (1976): "The Economics of Job Search: A Survey," *Economic Inquiry*, 14, 155–189.

Ma, C.-T. A. and M. Manove (1993): "Bargaining with Deadlines and Imperfect Player Control," *Econometrica*, 61, 1313–1339.

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, Cambridge University Press, 321–330.

McCall, J. J. (1970): "Economics of Information and Job Search," *The Quarterly Journal of Economics*, 84, 113–126.

Moldovanu, B. and X. Shi (2013): "Specialization and partisanship in committee search," *Theoretical Economics*, 8, 751–774.

Mortensen, D. T. (1970): "Job Search, the Duration of Unemployment, and the Phillips Curve," *The American Economic Review*, 60, 847–862.

Nachman, D. C. (1972): "On Risk Aversion and Optimal Stopping," Discussion paper No.26, Department of Managerial Economics and Decision Sciences, Graduate School of Management, Northwethern University.

Rogerson, R., R. Shimer, and R. Wright (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43, 959–988.

Romm, A. (2013): "Building Reputation at the Edge of the Cliff," Mimeo.

Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

Sakaguchi, M. (1973): "Optimal Stopping in Sampling from a Bivariate Distribution," *Journal of the Operations Research Society of Japan*, 16, 186–200.

——— (1978): "When to Stop: Randomly Appearing Bivariate Target Values," *Journal of the Operations Research Society of Japan*, 21, 45–58.

Smith, L. (1999): "Optimal Job Search in a Changing World," *Mathematical Social Sciences*, 38, 1–9.

Stinchcombe, M. B. (1992): "Maximal Strategy Sets for Continuous-Time Game Theory," *Journal of Economic Theory*, 56, 235–265.

van den Berg, G. J. (1990): "Nonstationarity in Job Search Theory," *The Review of Economic Studies*, 57, 255–277.

Wilson, C. A. (2001): "Mediation and the Nash bargaining solution," *Review of Economic Design*, 6, 353–370.

A Appendix: Additional Discussions

A.1 The Payoffs Realizing upon Agreement

Economic search situations are worthy of analysis because agents face a tradeoff between deciding now and doing so later. The benefit of deferring the acceptance of an offer lies in a prospect of getting a better payoff in the future, while the cost comes with some form of penalty on an act of deferring. There are two types of such penalty that economic agents face in reality: discounting (or time costs) and deadlines. The traditional infinite horizon models and our finite horizon model are at two extreme points: the former assumes discounting or time costs without deadlines, while the latter (our main sections) assumes the risk of reaching the deadline without assuming discounting or time costs.[59] We believe the reality is at somewhere in the middle of these two extreme cases, and which model is a better approximation of the reality depends on the particular applications that the analyst is interested in.[60] To understand these "middle" situations we think it a necessary step to analyze both of the two polar cases. We already know much about one of these cases from the literature, and the other case is what we analyzed in the main sections.

Here we present a way to formally connect these two cases. Specifically, we consider the case where the payoffs realize as soon as an agreement is reached, as opposed to assuming that the payoffs realize only at the deadline. Before considering "middle" cases, we first consider one of the polar cases in which we fix the discount rate $\rho > 0$ and then take the limit as the arrival rate λ tends to infinity. We show that the path of the continuation payoffs is close to that in the case of our main model when the deadline is relatively close and λ is sufficiently high for a fixed $\rho > 0$, while it diverges from such a path when the deadline is far away. The limit payoff profile is shown to be an element of the "Nash set," a generalization of the Nash bargaining solution. We also show that the limit expected search duration is zero.

After analyzing such an extreme case, we consider general convergence as $\lambda \to \infty$ with ρ depending on λ. We show that the limit payoffs and durations depend on the limit of $\lambda\rho^n$. Since this is a continuous function in both λ and ρ, the result implies that the relevance of the two extreme analyses—our main analysis and the analysis in which $\lambda \to \infty$ with fixed $\rho > 0$—depends on what situations we want to analyze. The latter analysis considers the case in which there are an increasingly large number of offers before payoffs are substantially discounted and thus the deadline is not at the relevant future, so the situation is getting more and more similar to the one described by the infinite horizon model. However, if we consider a situation where the payoffs are not so much discounted even at the deadline, i.e. ρ is relatively small compared to $1/\lambda^{1/n}$, then the

[59]We will review the literature on the infinite horizon models in Appendix B.

[60]No discounting may be the best model when the payoffs realize at the deadline as in our motivating example of apartment search, although one may still be able to imagine a possibility of small time costs.

deadline is in the relevant future, so the result of our main analysis is more relevant.

Formally, suppose that if a payoff profile $x = (x_1, \ldots, x_n)$ 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.[61] 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. Thus in this section, we focus on the case with $\rho > 0$. Under Assumption 1, essential uniqueness of trembling-hand equilibrium is obtained by a proof analogous to the one for Proposition 1. A straightforward computation shows that the following differential equation characterizes the (unique) continuation payoff $v(t)$ of the trembling-hand equilibrium:[62]

$$v'(t) = -\rho v(t) + \lambda \int_{A(t)} \big(x - v(t)\big)\, d\mu \tag{A.1}$$

with an initial condition $v(0) = (0, \ldots, 0) \in \mathbb{R}^n$.[63] The second term in the right hand side is the same as the right hand side of (1). The new addition is the first term, which corresponds to discounting (thus it is accompanied by a negative sign).

Suppose Assumptions 1 and 5 hold. Let $v^*(t; \rho, \lambda)$ be the (unique) solution of ODE (A.1). If λ is large, the right hand side of equation (A.1) can be approximated by the right hand side of equation (1) when the value of the integral is not too small. Therefore, $v^*(t; \rho, \lambda)$ is close to the solution of equation (1) in the case of $\rho = 0$, for λ large relative to ρ. This resemblance of trajectories holds until $\mu(A(t))$ approaches 0, which is when the value of the integral is small. In particular, we can show that for high enough λ, there exists $\bar{t} > 0$ such that the locus of $v^*(t; \rho, \lambda)$ in time interval $[-\bar{t}, 0]$ approximates that of $v^*(t; 0, \lambda)$ in $[-T, 0]$.

Proposition 10. *For all $\rho > 0$ and $\varepsilon > 0$, there exists $\bar{\lambda} > 0$ such that for all $\lambda \geq \bar{\lambda}$, there exists $\bar{t}$ such that for all t,*

$$\left| v^*(t; 0, \lambda) - v^*\big(\min\{t, \bar{t}\}; \rho, \lambda\big) \right| \leq \varepsilon.$$

The proposition shows that, even when there is discounting, the equilibrium dynamics represented by the trajectory of continuation payoffs (which are equal to the cutoffs) can be analyzed by what we identified in our main sections. This is helpful because the analysis of (1) is easier than that of (A.1): for example, the result implies that $v^*(t; \rho, \lambda)$

[61] This entails a loss of generality, but setting a nonzero threat-point payoff leads only to minor modifications of the statements of our results.

[62] Note that an argument similar to that of Proposition 1 shows that there exists a trembling-hand equilibrium consisting of cutoff strategies with cutoffs $v(t)$.

[63] In footnote 11 in Section 2, we noted that it is without loss of generality to assume that the disagreement payoff profile x^d is 0 if $\rho = 0$. This is not the case when $\rho > 0$. However, even if $x^d \neq 0$, the subsequent analyses go through with the initial condition for (A.1) being changed to $v(0) = x^d$.

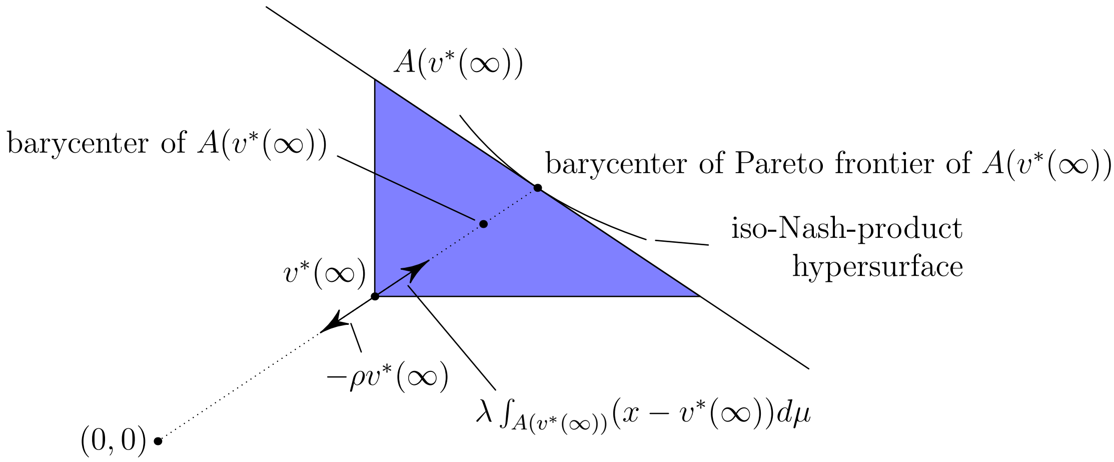

Figure 13: Vectors when $t \to \infty$.

comes close to the weak Pareto frontier if λ is sufficiently large, and makes it easier to apply local conditions of μ that we will assume later (Assumption 6). Moreover, even though this result itself does not say anything about the locus of the continuation payoffs when the deadline is far away, it will determine the limit expected payoffs under certain circumstances. We will be clearer on this in Remark 4 (b) after presenting Lemma 11.

Remark 3 (Difference between $\lambda \to \infty$ and $t \to \infty$)**.**

Before analyzing $v^* = \lim_{\lambda\to\infty} v^*(t;\rho,\lambda)$, let us consider another limit $v^*(\infty) := \lim_{t\to\infty} v^*(t;\rho,\lambda)$. Here we assume existence of these limits, which will be shown later in the proof of Lemma 11.

Since the right hand side of equation (A.1) is not proportional to λ, these two limits do not coincide for positive $\rho > 0$. If the limit $v^*(\infty)$ exists,[64] by continuity of the right hand side of (A.1) in v, we have $\lim_{t\to\infty} v^{*\prime}(t) = 0$, and $v^*(\infty)$ must satisfy

$$\rho v^*(\infty) = \lambda \int_{A(v^*(\infty))} \left(x - v^*(\infty)\right) d\mu. \tag{A.2}$$

For $\rho > 0$, equality (A.2) shows $\mu(A(v^*(\infty))) > 0$, which implies that $v^*(\infty)$ is Pareto *in*efficient in X. This contrasts with efficiency of $v^* = \lim_{\lambda\to\infty} v^*(t;\rho,\lambda)$ that we will show in Lemma 11.

Equality (A.2) also implies that the vector $v^*(\infty)$ is parallel to the vector from $v^*(\infty)$ to the barycenter of $A(v^*(\infty))$, as shown in Figure 13 for the two-player case. □

We impose the following assumption in addition to Assumptions 1 and 5. Except for minor points, they are satisfied in the literature (Wilson (2001), Compte and Jehiel (2010), Cho and Matsui (2013)).

[64] In the proof of Lemma 11, we show that $v^*(\infty)$ indeed exists.

Assumption 6. (a) At any point on the weak Pareto frontier in $\hat{X}$, the frontier is smooth and every component of the normal vector is strictly positive.

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

In Condition (a), the smoothness ensures the existence of a normal vector at any point on the weak Pareto frontier. The assumption that all the normal vectors are strictly positive implies that any weak Pareto efficient payoff profile is Pareto efficient. Condition (b) ensures existence of a "thick" Pareto frontier of X.

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 3. Assumptions 5 and 6, respectively, ensure that the density f is approximately uniform in $A(v^*(\infty))$ if $A(v^*(\infty))$ is a set with a very small volume and that $A(v^*(\infty))$ approximates a small n-dimensional pyramid. The vector in the right hand side of equality (A.2) is parallel to the vector from $v^*(\infty)$ to the barycenter of $A(v^*(\infty))$. We use this property to show that, in the interior of the Pareto frontier of $A(v^*(\infty))$, $A(v^*(\infty))$ 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* (Maschler et al. (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, and this is the standard Nash bargaining solution.

The above observation can be formalized with additional technical conditions to show the following lemma.

Lemma 11. *Suppose that Assumptions 1, 5, and 6 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 for all $t > 0$. If X is convex, this limit coincides with the Nash bargaining solution.*

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

Remark 4. (a) It is possible that the limit payoff in the model with $\rho = 0$ is isolated from a Nash set. In such a case, the lemma and Proposition 10 together imply that *the continuation payoff may not be monotone over time* in general, and in particular we expect to observe the rise of continuation payoffs over time when the deadline is not very close. We will discuss this nonmonotonicity more in Appendix A.12.

(b) The proof of Lemma 11 in particular proves that when λ is large enough, the Nash product is nondecreasing in t. To see the implication of this property, consider X in which there are two isolated local maximizers $\tilde{v}$ and $\hat{v}$ of the Nash product (which are elements of the Nash set). Consider two distributions with support X, $\tilde{\mu}$ and $\hat{\mu}$, such that in the model with $\rho = 0$ the continuation payoff of the former converges to $\tilde{v}$ and that of the latter converges to $\hat{v}$. Then, Proposition 10 implies that the continuation payoff can approach arbitrarily close to $\tilde{v}$ under $\tilde{\mu}$, and then Lemma 11 the nondecreasingness of the Nash product imply that the continuation payoff does converge to $\tilde{v}$ after that. A parallel statement holds for $\hat{v}$ under $\hat{\mu}$. That is, *which point in the Nash set becomes the limit equilibrium payoff profile depends on the distribution of offers.* Hence, the analysis of the case with $\rho = 0$ helps determine a rough position of the limit payoff, while the exact point is determined by the properties of Nash products on the boundary of X. In contrast, in the infinite-horizon model, all the Nash points arise as limits of equilibrium payoff profiles.[65] □

The key idea behind the proof is to first show that the probability of the acceptance set $A(v(t))$ shrinks to zero as $\lambda \to \infty$, and then use this fact to approximate $A(v(t))$ by a polyhedron (by the smoothness of the boundary assumed in Assumption 6 (a)). Once we do this, we can compute the approximate direction of the continuation payoffs and use that to show that the Nash product is nondecreasing in t when $A(v(t))$ is small enough. This leads to the desired result.

The fact that $A(v(t))$ shrinks to zero as $\lambda \to \infty$ is not trivial because the motion of $v(t)$ may not be monotone. Indeed, this was easy to show in the case without discounting because $v'(t)$ was nonnegative. Also, if we consider a stationary equilibrium in an infinite horizon problem, we can simply set $v'(t) = 0$ and observe that the acceptance set needs to be small enough when λ is high. In the current setting, $v'(t)$ may be negative and it is ex ante unclear if it goes to zero as $\lambda \to \infty$. We overcome this problem by using the fact that the negative term $-\rho x$ is bounded independently of λ, so the degree to which nonmonotonicity kicks in is limited.

In showing that the Nash product is increasing in t, we construct a Lyapunov function $L(t)$ that is a logarithm of the Nash product. Then we note that when the payoff for a player is higher, the discounting term in the ODE (A.1) makes the increase of the payoff slower the speed of increase of the payoff more slowly. This relation enables us to apply Chebyshev's sum inequality to show that $L'(t)$, defined by the sum of $v_i'(t)/v_i(t)$ across all i, is nonnegative.

In contrast to Theorems 1 and 2 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.

Lemma 12. *Suppose that Assumptions 1, 5, and 6 hold. If $\rho > 0$, then the limit expected search duration is zero.*

These lemmas, which concern the case with fixed $\rho > 0$, can be used to show the general result when ρ depends on λ and λ diverges to infinity.

[65] Although Wilson (2001) considers only convex feasible payoff sets, his argument can be generalized to the non-convex cases. Cho and Matsui (2013, Proposition 4.3) prove a related result with non-convex feasible payoff sets.

Proposition 13. *Suppose that Assumptions 1, 5, and 6 hold, and the discounting rate* ρ_λ *depends on* λ *and is bounded in* λ*. The limit payoff profile* $v^* = \lim_{\lambda\to\infty} v^*(t;\rho_\lambda,\lambda)$ *and the limit expected duration as* $\lambda \to \infty$ *satisfy the following claims: (i) If* $\lambda\rho_\lambda^n \to 0$*, then* $v^* = \lim_{\lambda\to\infty} v^*(t;0,\lambda)$*, and the limit expected duration is positive, which are the limits analyzed in Sections 4 and 5. (ii) If* $\lambda\rho_\lambda^n \to \infty$*, then* $v^* = \lim_{\lambda\to\infty} v^*(t;\bar{\rho},\lambda)$ *for any* $\bar{\rho} > 0$*, and the limit expected duration is zero, which are the limits shown in Lemmas 11 and 12.*

The proof idea is as follows: The limit of the expected payoffs depends on whether or not the first term in ODE (A.1) is negligible 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. Since ρ_λ is bounded in λ, $z(t;\rho,\lambda)$ is continuous in ρ and λ, and $z(t;\rho_\lambda,\lambda)$ is close to zero whenever λ is sufficiently large for each fixed $\rho \geq 0$, we have that $z(t;\rho_\lambda,\lambda)$ is close to zero for sufficiently large λ. Then we can apply an analogous argument to the one provided in the discussion in the proof sketch of Theorem 4 to show that the acceptance probability is approximately proportional to λ^{-1}, and since $\mu(A(t))$ approximates $z(t;\rho_\lambda,\lambda)^n$ times some constant, $z(t;\rho_\lambda,\lambda)$ is approximately proportional to $\lambda^{-1/n}$. Since the length of the vector from $v^*(t;\rho_\lambda,\lambda)$ to the barycenter of $A(t)$ is proportional to $z(t;\rho_\lambda,\lambda)$, the second term in ODE (A.1) is of order $\lambda \cdot \lambda^{-1/n} \cdot \lambda^{-1} = \lambda^{-1/n}$. Therefore if $\lambda\rho_\lambda^n \to 0$, then the first term, which approximates $-\rho_\lambda v^*$, is negligible because ρ_λ vanishes more rapidly than $\lambda^{-1/n}$. Thus the limit in this case is the same as in Sections 4 and 5. If $\lambda\rho_\lambda^n \to \infty$, the first term is significant because ρ_λ does not vanish rapidly compared to $\lambda^{-1/n}$. This leads to the limit in Lemma 11. An analogous argument can be made for the limit expected durations, which is given in Lemma 12.

Proposition 13 is used in Appendix A.7 to show that in the model of infinite horizon with discounting, the limit expected duration is zero as arrival rate λ tends to infinity. Combined with the "continuity at infinity" argument analogous to Fudenberg and Levine (1983), our result makes it possible to understand under a unified framework the equilibrium and its properties such as the limit expected duration in the finite horizon model and those of the infinite horizon model.

To sum up, the different timing of payoff realization can bring in a different set of results. However, the extent to which such results are relevant depends on the situation that we are interested in; here we identified the situation under which each analysis is more relevant than the other. Moreover, even in the case in which discounting is prominent, the equilibrium behavior is similar to the case with non-prominent discounting, when the deadline is close. This in particular implied that the continuation payoff may not be monotone over time, and the payoff distribution determines a rough position of the limit payoff profile. All these results suggest a wide applicability of our results for the case in which payoffs realize at the deadline, and we believe the analysis of the case in which payoffs realize upon agreement compliments our main analysis.

A.2 Market Designer's Problem

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

First, consider the case in which the payoffs realize at the deadline, and the designer can tune the horizon length T. In this case, if the designer is interested in efficiency of allocations, there is no point in making the horizon shorter because the continuation payoff profile $v(t)$ 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, we let the designer choose a measure μ' such that $\mu'(Y) \leq \mu(Y)$ for all Borel subsets $Y \subseteq X$.[66] 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 argument analogous to Proposition 9). On the other hand, however, changing the distribution will decrease the expected number of offer arrivals in the finite horizon. Agents then face more risks of disagreement, which detrimentally affect the payoffs $v(T)$. 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 basically if the horizon length T is high then reducing probabilities would not lead to too much loss, and thus the market designer has a high degree of freedom to choose from nearly efficient outcomes. As Proposition 9 shows, the freedom in choosing among the payoffs on the Pareto frontier is quite high.

Notice that proportional change in μ, i.e., $\mu'(Y) = a\mu(Y)$ for some constant $a \in (0, 1)$ for all measurable $Y \subseteq X$, has the same effect as shortening the horizon length from T to aT. This is because such a change in μ is equivalent to changing the arrival rate from λ to $a\lambda$, which we know is equivalent to the change in T when the payoffs realize at the deadline. However, there are at least two reasons that tuning of T is interesting. First, in some circumstances the designer may be able to lengthen T (and it is beneficial when the payoffs realize at the deadline, as discussed above). Since the probability cannot be integrated up above 1, there is no corresponding change in μ. Second, the designer may not have an access to all the possible tuning strategies of μ (or maybe one has no way to change μ), and if some desirable way of tuning of μ is not available then one may want to tune T instead when he or she has an easier access to tuning of T.

Next, we consider the case with payoffs realizing upon agreement as in Section A.1. In such a case, the equivalence of the changes in λ and in T no longer holds, so the effect of changes in μ and T are not equivalent, even when μ is changed proportionally.

Specifically, for a long enough horizon length, the tuning of T has little effect on the welfare while tuning of μ can have a large effect.[67] The reason is precisely that

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

[67] To see the difference clearly, suppose that $n = 2$ and μ is the uniform distribution over $X = \{x \in$

tuning of μ corresponds to tuning of λ, which means that the amount of discounting between opportunities is on average high. Below we consider tuning of T and then non-proportional tuning of μ.

First, notice that 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 profile at time $-T$ is less close to the limit payoff profile as $T \to \infty$ which is close to the Pareto frontier. However, this does not necessarily imply that the resulting expected payoff profile is less socially desirable. Let us assume convexity of X, and consider the solution for the case in which payoffs realize at the deadline, $v^*(T; 0, \infty)$. If it is more 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; recall that by Proposition 10 the expected payoffs for a certain intermediate range of time $-t$ is close to $v^*(T; 0, \infty)$).

On the other hand, non-proportional tuning of the distribution has a small effect in contrast to the case with payoffs realizing at the deadline when X is convex, as we know that the payoffs eventually converge to the Nash bargaining solution. However, since $v^*(T; 0, \infty)$ depends on the distribution, Proposition 10 implies that the direction from which the payoff converges varies as the designer varies the distribution. Also, without convexity, tuning of the distribution may have a large effect. This is because there may exist multiple points in the Nash set, and different distributions may lead to different points in the Nash set the continuation payoff profile converges to.

To wrap up, the market designer may have different ways to tune the parameters of the model, and they create different effects depending on the timing of the payoff realizations.

A.3 General Voting Rules

In the main sections we considered the case when players use the *unanimous rule* for their decision making. This is a reasonable assumption in many applications such as the apartment search, but there are other applications in which different voting rules (e.g., majority rules) may fit the reality better. This section is devoted to the analysis of such cases.

Let us consider a general voting rule in which $\mathcal{C} \subseteq 2^N$ is the set of winning coalitions; The object of search is accepted if and only if there is a winning coalition $C \in \mathcal{C}$ in which every player says "accept" upon its arrival. A minimal winning coalition is a coalition $C \in \mathcal{C}$ such that if $C' \subseteq C$ and $C' \in \mathcal{C}$, then $C' = C$. We assume that any player can be

$\mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$. Fix $\rho > 0$. Let $v^a(T)$ for $a \in (0, 1]$ be the payoff profile at time $-T$ under the measure μ' such that $\mu'(Y) = a\mu(Y)$ for all measurable $Y \subseteq \mathbb{R}^2_+$. Then, for any $a \in (0, 1)$ and $\varepsilon > 0$, there exists $\bar{\lambda}$ such that for all $\lambda > \bar{\lambda}$, there exists $\bar{T}$ such that for all $T > \bar{T}$, (i) $\left|(\frac{1}{2}, \frac{1}{2}) - v^1(aT)\right| \leq (1+\varepsilon)\left|(\frac{1}{2}, \frac{1}{2}) - v^1(T)\right|$; and (ii) $\left|(\frac{1}{2}, \frac{1}{2}) - v^a(T)\right| \geq \left((\frac{1}{a})^{\frac{1}{3}} - \varepsilon\right)\left|(\frac{1}{2}, \frac{1}{2}) - v^1(T)\right|$.

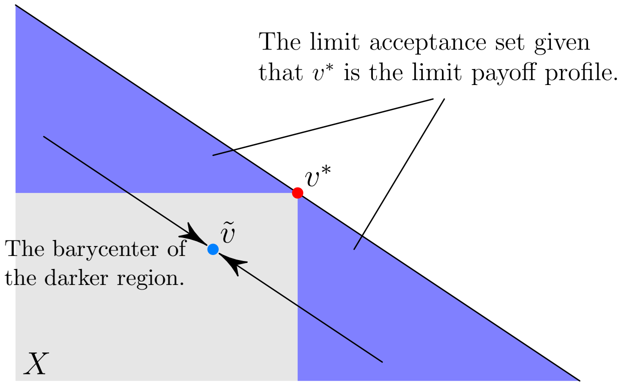

Figure 14: Equilibrium continuation payoffs under a majority rule: If the point v^* were the limit payoff profile, then the acceptance set should have contained the darker region, and the barycenter of the darker region, $\tilde{v}$, is in the interior of X as X is convex.

pivotal, i.e., for all $i \in N$, there exists a minimal winning coalition $C \in \mathcal{C}$ with $i \in C$. It is straightforward to check that if μ satisfies Assumption 1, Propositions 1 and 2 (existence of an essentially unique trembling-hand equilibrium and the use of cutoff strategies) carry over to this case.

The voting rule naturally induces a coalitional-form game with non-transferable utility, (N, V), where the characteristic function V is defined as $V(C) = X$ if $C \in \mathcal{C}$, otherwise $V(C) = \{0\}$. The core of (N, V) is the set of all payoff profiles $x \in X$ such that there are no $C \subseteq N$ and $y \in V(C)$ with $y_i > x_i$ for all $i \in C$. Note that in the case of the unanimous rule ($\mathcal{C} = \{N\}$), the core equals the weak Pareto frontier of X.

Suppose that X is compact and convex, and has a nonempty interior. First, suppose that the limit expected payoff profile v^* is not in the core. This occurs for example if the core is empty. By convexity of X, $v^* \in X$. If v^* is not in the core, there exists $C \in \mathcal{C}$ such that $\{x \in X \,|\, x_i > v_i^* \text{ for all } i \in C\}$ is nonempty. Since this set is an open subset in X, the probability that the payoff profile realizes in the subset is positive. Therefore the limit expected duration must be zero. Furthermore, we can show that v^* cannot be weakly Pareto efficient. To see this, suppose that v^* is weakly Pareto efficient. An intuitive explanation can be seen in Figure 14 which describes the case where any single player out of two can decide the final outcome. One can find 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 v^* must be an interior point as well. This contradicts the assumption that the limit payoff profile is weakly Pareto efficient. Just as in the case when v^* is not in the core, we can show that v^* not being weakly Pareto efficient implies that the limit expected duration is zero.[68]

Next, suppose that the core is nonempty. Then we can show that the limit expected

[68]This discussion is parallel to that of Compte and Jehiel (2010, Proposition 7) who consider majority rules in a discrete-time infinite-horizon search model.

duration is positive for some probability measure μ with sufficiently high density near the core. We summarize our findings as follows:

Proposition 14. *Suppose that $X \subseteq \mathbb{R}^n_+$ is compact and convex, and has a nonempty interior. Then, if a probability measure μ over X satisfies Assumption 1 and the limit expected payoff profile v^* under μ is not in the core, then v^* is not weakly Pareto efficient and the limit expected duration is zero. In addition, the core is nonempty if and only if there exists μ with support X satisfying Assumption 1 such that the limit expected duration is positive and v^* is weakly Pareto efficient.*

A.4 Intuition for Essential Uniqueness of Trembling-Hand Equilibrium

The (essential) uniqueness of trembling-hand equilibrium (Proposition 1) is nontrivial as we work with continuous time so the standard backward-induction argument does not apply. Here we explain the key idea of the proof. In a single-player search model, if two strategies give rise to two different continuation payoffs at some time $-t$ then the strategy with the lower continuation payoff is obviously suboptimal, and this trivially implies uniqueness. This proof cannot be used for the case of two or more players. For example, it might be the case that in one equilibrium player 1 is picky and player 2 is generous, while in another equilibrium the opposite happens, and these two are both equilibria as both imply reasonable levels of acceptance probabilities at each point in time.

Here we present the proof idea for a two-player case to keep the notation simple. The idea can be generalized to the cases with more than two players.

The proof consists of two steps. In the first step, we bound the supremum difference of continuation payoffs at time $-t$ across all the trembling-hand equilibria using those for time $-\tau \in (-t, 0]$. Then we show in the second step that such bounds at all time $-t$ imply that such differences are zero at all time $-t$.

So consider the problem with players 1 and 2. Let $\bar{q}_i(t)$ and $\underline{q}_i(t)$ be the supremum and the infimum, respectively, of continuation payoffs for player i across all trembling-hand equilibria.[69] Let $w_i(t) = \bar{q}_i(t) - \underline{q}_i(t)$ be the difference. Our goal is to first bound $w_i(t)$ using $w_1(\tau)$ and $w_2(\tau)$ for $\tau \in [0, t)$, and then to use such a bound to show that $w_i(t) = 0$ for all t. For simplicity, assume that $w_i(t)$ is increasing in t for each i. This needs a proof, but would be intuitive because as the remaining time shrinks there is less and less room for equilibrium strategies to make a difference in terms of payoffs.

[69] Formally, let $Q_i(t;\sigma)$ be the set of continuation payoffs for player i at time $-t$ at which no acceptance has occurred in the time interval $[-T, -t]$ (note that this interval includes time $-t$), under the strategy profile σ. Let Σ^* be the set of trembling-hand equilibria (which is non-empty due to Proposition 2). We define $\bar{q}_i(t) = \sup_{y \in Q_i(t;\sigma), \sigma \in \Sigma^*} y$ and $\underline{q}_i(t) = \inf_{y \in Q_i(t;\sigma), \sigma \in \Sigma^*} y$.

The key of the first step is to observe that, given the "trembling-hand" restriction at time $-t$, player i having an opportunity at time $-\tau$ accepts all the payoffs strictly above $\bar{q}_i(\tau)$ and rejects all the payoffs strictly below $\underline{q}_i(\tau)$. This is true for every future time $-\tau \in (-t, 0]$. So if any two equilibria give rise to two different continuation payoffs at time $-t$, the difference should be attributed to that of the agents' behavior at some future time $-\tau$ when the payoff realization is between $\bar{q}_i(\tau)$ and $\underline{q}_i(\tau)$.[70] The probability with which the offer falls in this "ambiguous region" is at most $L \cdot (w_1(\tau) + w_2(\tau))$ where L is the maximum across players of the supremums of the marginal densities of payoffs, which exist by Assumption 1 (b). When there is a payoff difference, the difference can be in expectation at most the expected payoff from the distribution, which is finite by Assumption 1 (a). Denote the maximum across players of these expected payoffs by $\bar{x}$. These facts imply, for each $i \in I$,

$$\begin{aligned} w_i(t) &\leq \int_0^t \left(\begin{bmatrix} \text{Probability of the} \\ \text{offer falling in the} \\ \text{“ambiguous region”} \\ \text{at time } -\tau \end{bmatrix} \times \begin{bmatrix} \text{The maximum} \\ \text{payoff change} \\ \text{on average} \end{bmatrix} \right) \lambda e^{-\lambda(t-\tau)} \, d\tau \\ &\leq \int_0^t \big(L(w_1(\tau) + w_2(\tau)) \times \bar{x} \big) \lambda e^{-\lambda(t-\tau)} \, d\tau. \end{aligned}$$

Letting $M = 2\lambda L \bar{x}$ and summing across players implies that

$$w_1(t) + w_2(t) \leq \int_0^t M\big(w_1(\tau) + w_2(\tau)\big) \, d\tau. \tag{A.3}$$

The intuition for the derivation is the following: Under trembling-hand equilibria, the actions taken when the payoffs are either very high or very low are uniquely determined, and the only thing that could possibly depend on a particular choice of trembling-hand equilibrium is the actions taken when receiving an offer in the ambiguous region. The probability of receiving such an offer at time $-\tau$ is at most proportional to the difference of payoffs $w_1(\tau)+w_2(\tau)$, because (i) the density of payoffs is bounded and (ii) the expected payoffs are finite.

The second step is to show that inequality (A.3) with the initial condition $w_i(0) = 0$ has a unique solution $w_i(t) = 0$ for all t. The actual proof for this is analogous to the standard textbook-proof for uniqueness of the solution of a differential equation, but here let us try to provide a detailed intuition referring to the structure of our game.

Recall that $w_i(t)$ is increasing in t. Inequality (A.3) implies

$$w_1(t) + w_2(t) \leq Mt\big(w_1(t) + w_2(t)\big).$$

[70] For the sake of argument, we ignore the effect that future trembles contribute to the difference in the current expected payoffs. The proof in Appendix D.3 does not ignore this effect.

If t is small, the only way to satisfy this inequality is to have $w_1(t) + w_2(t) = 0$ which implies $w_i(t) = 0$.[71] This is true for all $t \leq \frac{1}{M}$. But given this, we can rewrite inequality (A.3) for $t \geq \frac{1}{M}$ as follows:

$$w_1(t) + w_2(t) \leq \int_{\frac{1}{M}}^{t} M\big(w_1(\tau) + w_2(\tau)\big)\, d\tau.$$

Then we can iterate the same argument for all $t \in [\frac{1}{M}, \frac{2}{M}]$, and this goes on indefinitely.

Let us summarize: In order to create a difference in the current continuation payoff, it needs an enough variation in the future continuation payoffs. But if the remaining time is small, the variation should be large in absolute term, which is impossible because of our first step, that is, we do not have full flexibility to vary their strategies due to the "trembling-hand" restriction.[72]

The formal proof in Appendix D.3 considers the ε-constrained game explicitly, does not hinge on the assumption that w_i is increasing, and deals with the n-player case.

A.5 Additional Welfare Implications

Pareto Efficiency under Generic Distributions

In Proposition 7 in Section 5, we showed that the limit payoff profile v^* is Pareto efficient whenever X is convex. Even when X is not convex, we can argue that the limit payoff profile v^* is Pareto efficient under "generic" probability measures μ on X. To formalize what "generic" means, let $\mathcal{F}$ be the set of density functions that satisfy Assumptions 1 and 5. We consider a topology on $\mathcal{F}$ defined by the pointwise convergence. To exclude uniteresting cases, we focus on the two-agent case in which there exist single-variable functions $g, h : \mathbb{R}_+ \to \mathbb{R}_+$ such that X is expressed as $X = \big\{(x_1, x_2) \in \mathbb{R}^2_+ \,|\, x_2 \in [0, g(x_1)], x_1 \in [0, h(x_2)]\big\}$. This set consists of all the points $x \in \mathbb{R}^n_+$ that lie below the graph of g and on the left of the graph of h as illustrated in Figure 15. We say that a function is *piecewise continuous* if it is continuous at all but a finite number of points.

Proposition 15. *Consider a two-agent model with both g and h being piecewise continuous and quasiconcave. Under Assumptions 1 and 5, the set $\{f \in \mathcal{F} \,|\, v^*$ is Pareto efficient in $X\}$ is open and dense in $\mathcal{F}$.*

In the proof of openness, we use the fact that when the density changes continuously, the limit point v^* changes continuously if it is Pareto efficient. Given the general property of v^* identified in the proof of Proposition 7, the piecewise-continuity condition implies that any Pareto efficient payoff profile x has a neighborhood $B_\varepsilon(x)$ such that each point

[71] This is because, by definition, $w_i(t) \geq 0$ for each t and i.

[72] The regularity condition imposed by Kamada and Kandori (2011), which we discussed in footnote 56 in Section 6, is analogous to having a term proportional to $w_i(\tau)$ for each τ in the integrand of the right hand side of inequality (A.3).

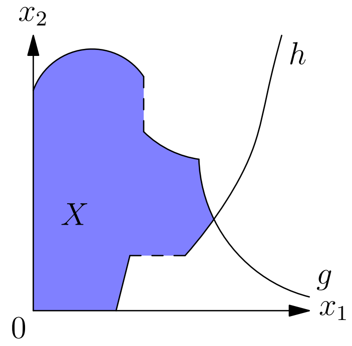

Figure 15: An example of the domain X defined by functions g and h, both of which satisfy piecewise continuity and quasiconcavity assumed in Proposition 15.

in $B_\varepsilon(x)$ is either (i) Pareto efficient, or (ii) not a limit of $v^*(t)$ as $t \to \infty$ for any $f \in \mathcal{F}$. Therefore, if v^* under f is Pareto efficient, any density sufficiently close to f leads to a Pareto efficient limit point. This ensures openness.[73]

To show denseness, suppose that the limit payoff profile v^* is not Pareto efficient, g is discontinuous at v_1^*, and $\lim_{x_1 \nearrow v_1^*} g(x_1) > v_2^* \geq \lim_{x_1 \searrow v_1^*} g(x_1)$. Consider a modified density function $\bar{f}$ with the same support but a small amount of density added near a Pareto efficient profile $\bar{y} = (v_1^*, \lim_{x_1 \nearrow v_1^*} g(x_1))$, and assume that the limit payoff profile remains the same. Then the slope of the trajectory under $\bar{f}$ must be steeper near v^*. Since we assumed that the limit payoff profile is the same, the trajectory near the limit point under $\bar{f}$ comes below the one under f. However, since the slope under the modified density is steeper at every point on the trajectory under f, the trajectory under $\bar{f}$ comes above the one under f.[74] This is a contradiction, implying that the two limits are not the same. By an argument similar to the proof of Proposition 7, we can show that these two limit are distant from each other, which implies the denseness of the set of density functions with which the limit point is Pareto efficient. We assumed quasiconcavity to guarantee existence of the limit of the slope of the trajectories.[75]

Difficulty of Comparative Statics

Before Proposition 9, we argued that it is difficult to conduct comparative statics regarding the limit payoff profile. To illustrate such difficulty, we present an example in which player i's marginal distribution of a probability measure μ is first-order stochastically dominated by that of another probability measure γ, and i's limit expected payoff under μ exceeds the one under γ.

[73] In Appendix D.9, we will present a pathological counterexample in which the set is not open in $\mathcal{F}$, and g or h violates piecewise continuity.

[74] Such an argument cannot be generalized to the case with $n \geq 3$, in which a trajectory does not divide X into two, and there is no clear way to define a region "above" or "below" a trajectory.

[75] In Appendix D.9, we will present a pathological counterexample in which the set is not dense in $\mathcal{F}$, and g or h violates quasiconcavity.

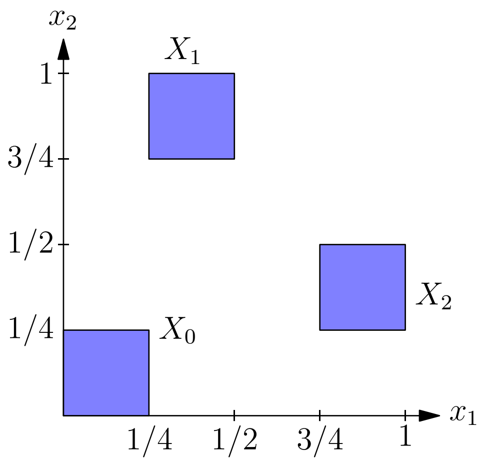

Figure 16: The support $X = X_0 \cup X_1 \cup X_2$ of density functions f and g.

Example 5. Let $n = 2$, $X_0 = [0, 1/4]^2$, $X_1 = [1/4, 1/2] \times [3/4, 1]$, $X_2 = [3/4, 1] \times [1/4, 1/2]$, and $X = X_0 \cup X_1 \cup X_2$ as presented in Figure 16. Let f_0, f_1, and f_2 be the uniform density function on X_0, X_1, and X_2, respectively.

Suppose that probability measure μ has a density function f defined by $f(x) = 0.5f_0(x) + 0.2f_1(x) + 0.3f_2(x)$. Since f has a larger density on X_2 than on X_1, the limit payoff profile under μ as $\lambda \to \infty$ is $(1, 1/2)$.

Consider another probability measure γ with a density function g defined by $g(x) = 0.2f_0(x) + 0.5f_1(x) + 0.3f_2(x)$. Since any point in the interior of X_1 strictly Pareto dominates any point in X_0, for each $i = 1, 2$, i's marginal distribution of μ is first-order stochastically dominated by that of γ. The limit payoff profile is, however, $(1/2, 1)$ because g has a larger density on X_1 than on X_2. Hence, for large enough λ, player 1's expected payoff under μ is strictly larger than that under γ. $\square$

A.6 Non-Poisson Arrival Processes

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

Specifically, for each integer $m \geq 1$, consider dividing the time horizon of length T into small subintervals each with length Δ_m (so there are $\frac{T}{\Delta_m}$ periods in total) with $\lim_{m\to\infty} \Delta_m \to 0$. At the end of each subinterval, players obtain an opportunity with probability $\lambda_m \Delta_m$. Notice that Poisson processes correspond to the case when λ_m is constant with respect to m and we let $\Delta_m \to 0$ as $m \to \infty$. Here we allow for general sequences $(\lambda_m)_m$, such as $\lambda_m = a/\Delta_m$ or $\lambda_m = a/\sqrt{\Delta_m}$ for some constant $a > 0$. Under Assumption 1, backwards induction implies that for every period, all trembling-hand equilibria yield the same continuation payoff profile for almost all histories, and among them, there exists a Markov perfect equilibrium. In the search problem with subintervals with length Δ_m, for each $k = 0, 1, 2, \ldots, \frac{T}{\Delta_m}$, let $v_i(k; m)$ be the (unique) continuation payoff for player i at time $-k\Delta_m$ after rejecting an offer if any. Then,

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

Hence,

$$v_i\Big(\frac{t}{\Delta_m} + 1; m\Big) - v_i\Big(\frac{t}{\Delta_m}; m\Big) = \lambda_m \Delta_m \int_{A(v(\frac{t}{\Delta_m}; m))} \Big(x_i - v_i\Big(\frac{t}{\Delta_m}; m\Big)\Big)\, d\mu. \tag{A.4}$$

Notice that if we set $\lambda_m = \lambda$ being constant and take the limit as $m \to \infty$, the left hand side divided by Δ_m converges to $v_i'(t)$ in the Poisson model with arrival rate λ and the right hand side divided by Δ_m converges to $\lambda \int_{A(v(t))} (x_i - v_i(t)) d\mu$, consistent with equation (1), and it turns out that the expected search duration as $\lambda_m \to \infty$ and $\Delta_m \to 0$ converges to the one that we derived in our model of continuous time. Thus, we can show the following:

Proposition 16. *Under Assumptions 1 and 4,* $\lim_{m\to\infty} \lambda_m = \infty$ *implies that the limit expected duration as* $m \to \infty$ *is* $\frac{n^2}{n^2+n+1} T$.

Note that this result is consistent with Proposition 5 where we consider the Poisson process and take a limit as $\lambda \to \infty$. This shows robustness of our result to move structures.

A.7 Approximating an Equilibrium of an Infinite-Horizon Game

Although we consider a finite-horizon model, our convergence result in Lemma 11 is suggestive of that in infinite-horizon models such as Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013), 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 models are similar to a finite-horizon model with $T \to \infty$ if $\rho > 0$.[76] 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 and the assumptions imposed in Proposition 13 hold. To see this, note that by Proposition 13, $\lim_{\rho \to 0} v^*(1; \rho^{1-a}, \rho^{-a}\lambda)$ is the Nash bargaining solution for all $a \in (n/(n+1), 1)$ and all λ. Since enlarging T is equivalent to raising both λ and ρ in the same ratio by the form of ODE (A.1), $v^*(1; \rho^{1-a}, \rho^{-a}\lambda) = v^*(\rho^{-a}; \rho, \lambda)$ holds for each $\rho \in (0, 1)$. Thus, for sufficiently small $\rho > 0$, there exists a large T such that $v^*(T; \rho, \lambda)$ is sufficiently close to the Nash bargaining solution. For the same reason, the expected duration in the limit as λ goes to ∞ in the infinite-horizon model is zero, which is analogous to our Lemma 12 in which we send λ to ∞ while $\rho > 0$ is fixed. Therefore, in the infinite-horizon search model, the expected duration in a stationary equilibrium converges to zero as $\lambda \to \infty$.

A.8 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 expected search durations with reasonable parameter values are close to the limit expected 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 (1) is modified in the following way:

$$v_i'(t) = -c + \lambda \int_{A(t)} \big(x_i - v_i(t)\big)\, d\mu \tag{A.5}$$

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

The analysis of this differential equation is similar to the one in Section A.1, with an exception that under Assumptions 1, 4, and 6, 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. Let $T = 1$ and denote by $D(\lambda; c)$ the expected duration when the arrival rate is λ and the time cost is $c \geq 0$. A proof similar to the one for Proposition 13 shows the following:

Proposition 17. *Suppose that Assumptions 1, 4, and 6 hold, and c_λ depends on λ, and bounded in λ. Then, (i) If $\lambda c_\lambda^n \to 0$ as $\lambda \to \infty$, then $\lim_{\lambda\to\infty} D(\lambda; c_\lambda) > 0$, and $\lim_{\lambda\to\infty} v^*(t; c_\lambda, \lambda) = \lim_{\lambda\to\infty} v^*(t; 0, \lambda)$, which are the limits analyzed in Sections 4 and*

[76] An analogous proof as in Fudenberg and Levine (1983) shows that the limit of any subgame perfect equilibrium in our model is a subgame perfect equilibrium of the corresponding infinite-horizon game as the horizon grows long because their "continuity at infinity" condition holds in our model.

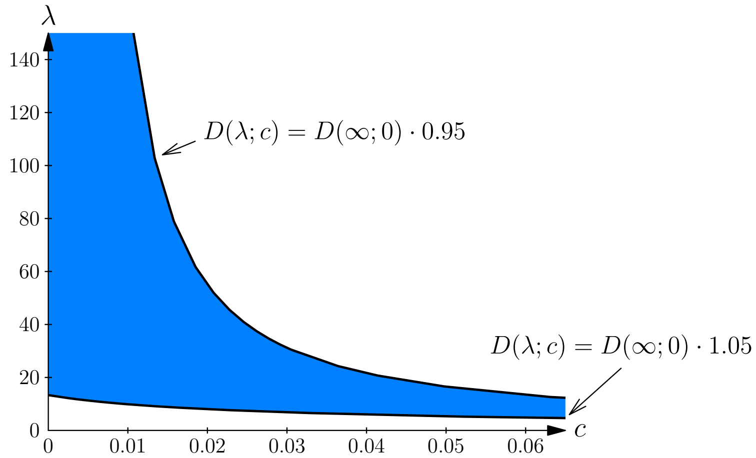

Figure 17: Time costs and arrival rates. The shaded region describes the set of pairs (c, λ) with which the expected duration is within 5% difference from the limit expected duration.

5. (ii) If $\lambda c_\lambda^n \to \infty$ *as* $\lambda \to \infty$*, then* $\lim_{\lambda\to\infty} D(\lambda; c_\lambda) = 0$ *and* $\lim_{\lambda\to\infty} v^*(t; c_\lambda, \lambda) = v^S$*.*

The proposition suggests that for a high arrival rate λ, the expected 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. We can 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 the 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 expected duration. As shown in Figure 17, such a range contains a wide variety of pairs of parameter values (note that the limit expected 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 expected duration is more than 100, and that of 105% is less than 10.[77] These results suggest that the limit argument that we conducted in Steps 1 and 2 of the main sections is economically meaningful in a wide range of problems.

A.9 Counterexamples of Positive Duration of Search

In Theorems 1 and 2, we showed that the limit expected duration of search is positive if certain assumptions hold. In this section, we present examples of distributions under

[77] In our continuation work, we explore this case more and show that the expected duration is positive even in the limit as $\lambda \to \infty$, exhibiting a stark contrast to Lemma 12.

which some of these assumptions are not satisfied and the expected duration converges to zero as $\lambda \to \infty$ for a certain sequence of equilibria.

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.[78] 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 does, there are multiple trembling-hand equilibria for each λ under μ, and for some sequence of equilibria the limit expected duration vanishes as $\lambda \to \infty$.

Example 6. Consider $X = \big(\{0,1\} \times [1,2]\big) \cup \big([1,2] \times \{0,1\}\big)$ and let μ be a uniform distribution over this X. First, consider a strategy profile in which an agent accepts an offer until time $-t^*$ if and only if it gives her a payoff strictly above 1, and accepts one after $-t^*$ if and only if it gives her a strictly positive payoff, where t^* satisfies the indifference condition at $-t^*$:

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

or $t^* = \frac{2}{\lambda}\ln(5)$. 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.[79]

However, there exist other equilibria. For example, consider a strategy profile which is exactly the same as the above one except that both players accept offers with payoffs larger than or equal to 1 whenever $-t \in [-T, -t^*]$ and each player has accepted every previous offer with a payoff larger than or equal to 1 for her. Since the continuation payoff after a player rejects an offer with a positive payoff at $-t \in [-T, -t^*]$ is 1 as we have argued, this also constitutes a trembling-hand equilibrium. Thus there are multiple trembling-hand equilibria. Moreover, the limit expected duration in the latter is zero because the agreement probability on the equilibrium path at any time $-t \in [-T, -t^*]$ is $1/2$ independently of λ. This suggests the need for Assumption 1 (b) for Theorem 2 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.

[78] In this case, Kamada and Sugaya (2010)'s "three-state example" can be interpreted as our model with a probability measure μ that assigns equal probabilities to $(2,1)$ and $(1,2)$. They show that multiple subgame perfect equilibria exist, and the limit expected duration can be zero in a subgame perfect equilibrium.

[79] This is because the continuation payoff is 1 even with trembles since the unconditional expected payoff of μ is 1.

Assumption 1 (b) rules out such a situation. □

Next, we show that even if μ's marginal has no point masses, the limit expected duration may be zero when μ violates Assumption 2 (and Assumption 7 in Appendix D.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 $f(x) = \frac{1}{(1-x)(\ln(1-x))^2}$ for $x \in [1 - e^{-1}, 1)$. Recalling formula (3), the density term is

$$d_1(v) = \frac{f(v)}{1 - F(v)} = -\frac{1}{(1-v)\ln(1-v)},$$

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 4, the limit expected 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 $x \in [1 - \varepsilon, 1]$. That is, $F(x)$ converges to 1 as $x \to 1$ at a speed slower than any polynomial functions, so in a sense F is very close to a discrete distribution. In such a case, the above computation shows that the limit expected duration can be zero, which is the same as the case with discrete distributions. □

A.10 The Effect of a Slight Change in the Distribution

The limit result in Proposition 5 depends crucially on the assumption of smooth Pareto frontier 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 distributions over $\mathbb{R}^n_+$ which may or may not have full support, and introduce a notion of distance between two distributions, $d(\mu, \gamma) = \sup_{\text{Borel } A \subseteq \mathbb{R}^n_+} |\mu(A) - \gamma(A)|$.

A standard argument on ordinary differential equations shows the following:

Proposition 18. *Under Assumption 1, for any λ, the expected duration is continuous with respect to the distribution of payoff profiles μ.*

That is, for any finite arrival rate, the expected duration is not substantially affected by a slight change in distribution. Combined with the result that our limit result approximates the situations with finite but high arrival rates, this suggests that our limit expected duration is relevant even for the distributions that do not satisfy (but are not too far from a distribution that satisfies) our assumptions (Assumptions 1 and 5).

A.11 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 certain situations it might be more realistic that the distribution changes over time. In this section, we examine whether the positive duration result in Theorem 1 (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 (history-independent) cumulative distribution function of the payoff at time $-t$ satisfying Assumptions 1 and 2.

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: 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.

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 distribution shrinks proportionally with an exponential speed as time passes, then the analysis of the duration becomes equivalent to that for the case with discounting, in which Lemma 12 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 1 provided in Appendix D.4, 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 for all t, $1 - \varphi(x)$ is of the same order as $1 - F_t(x)$ in $\{x \in \mathbb{R} \mid F_t(x) < 1\}$.

Notice that we require the existence of φ that is applicable to all F_t.

Proposition 19. *Suppose $n = 1$. Under Assumptions 1 and 2″*, $\liminf_{\lambda\to\infty} D(\lambda) > 0$.

When $n \geq 2$, one can obtain the result of positive duration under an assumption parallel to Assumption $2'$.

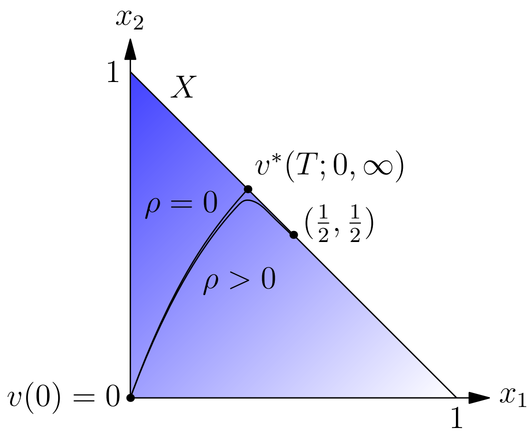

Figure 18: Paths of continuation payoffs. The probability density is low near $(1, 0)$, and high near $(0, 1)$.

A.12 Dynamics of the Bargaining Powers

Consider the case where $X = \{x \in \mathbb{R}^2_+ \mid x_1 + x_2 \le 1\}$ 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 Appendix A.1, and the discount rate $\rho > 0$ is very small. In this case, the limit of the solution of ODE (1) with $\rho = 0$, denoted $v^*(T; 0, \infty)$, locates at the boundary of X by Proposition 7, 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 (A.1). Hence, by Proposition 10, 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 18 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.

A.13 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-agent infinite-horizon search model minimal so that the analysis isolates 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 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 limit expected 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.[80] 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.[81] Specifically, we consider the case with costly transferable utility: Suppose that $X = \{x \in \mathbb{R}^2_+ \,|\, x_1 + x_2 \leq 1\}$ 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 examining the geometric properties of the Nash bargaining solution, we can compute the limit expected duration in this environment. Let $T = 1$.

[80] 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.

[81] We expect that nothing substantial would change even if we extended the argument to the cases of three or more players.

Proposition 20. *Under Assumptions 1, 5, and 8, the limit expected duration as* $\lambda \to \infty$ *in the game with negotiation is* $\dfrac{4+4a^2}{7+6a+7a^2}$ *for* $a \in [0,1)$.

Notice that the limit expected duration becomes shorter in the presence of negotiation (it is decreasing in a). 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 limit expected duration must be strictly positive even with negotiation, and it converges to 4/7 as $a \to 0$, which is the same duration as we found in Proposition 5. That is, our main result is robust to the introduction of negotiation.

Note that the proposition excludes the extreme case in which utilities are perfectly transferable, i.e., $a = 1$. In this case, the acceptance probability is of the order of $1/|v^* - v^*(t)|$, and the limit expected duration is 1/3. Thus, the duration formula in Proposition 20 is discontinuous at $a = 1$.

A.14 Preference Heterogeneity with Finite Arrival Rates

Here we present a further discussion that supports our argument that preference heterogeneity implies long search durations.

Example 8 (Change in the shape of X under Assumptions 1 and 4). 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

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

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 $\bar{\mu}^a\big(\{y^a(x) \mid x \in C\}\big) = \mu(C)$ for any Borel set $C \subseteq X$.

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 19 for a graphical description in the case of two players. Proposition 5 shows that as long as Assumptions 1 and 4 are met, expected duration is unaffected by the specificality of distribution μ. This is because, in both cases, the distribution is still uniform around the limit point and the Pareto frontier is smooth even under $\bar{\mu}^a$, so exactly the same calculation as in the case with μ suggests that the limit expected 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 3 shows the effect of preference heterogeneity. As preferences become less heterogeneous (smaller q and smaller a), the expected duration becomes shorter. □

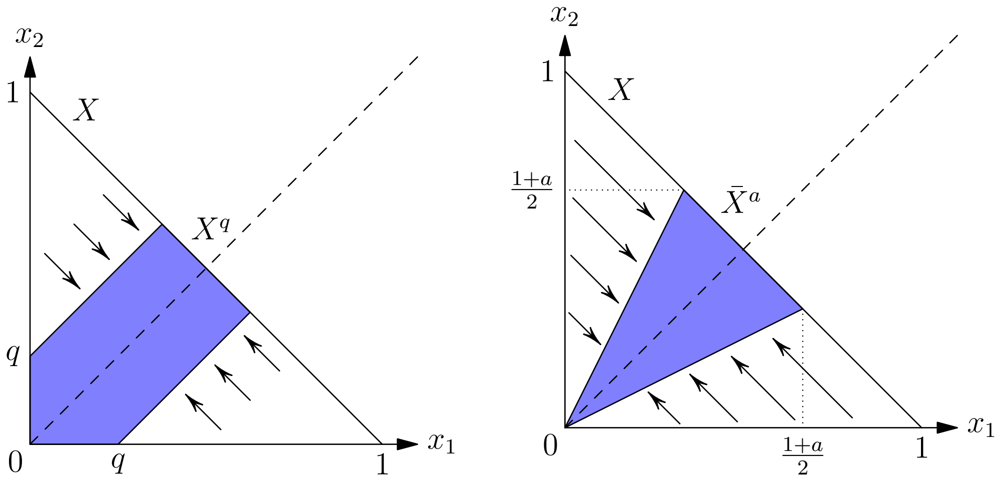

Figure 19: Transfer of allocations in the negotiation for μ^q (left) and $\bar{\mu}^a$ (right), where the original distribution is uniform over $\{(x_1, x_2) \in \mathbb{R}^2_+ \mid x_1 + x_2 \leq 1\}$.

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

Table 3: Preference heterogeneity effect under Assumptions 1 and 4. q and a measure heterogeneity of preferences.

A.15 Decomposition

The duration formula in Theorem 4 gives a way for a market designer to examine the effect of alternative policies, such as an increase of the number of players, or a change of distribution. To give a concrete measure of what policy influences the duration in what way, decomposing the effects that determine the search duration can be helpful. There would be many methods to do so; here we provide one of them.

For example, consider 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 $\lambda = 10$, which is 0.608. The limit expected duration as $\lambda \to \infty$ in this case is $\frac{4}{7}$, so the difference is 0.037. These durations are illustrated in Figure 20. The limit expected duration in general is computed from a key variable r determined by the details of the model (X, μ). The larger the r is, the longer the expected duration is. In this example, the limit expected duration $\frac{4}{7}$ is calculated from 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 expected duration is $\frac{1}{3}$, and the number r is

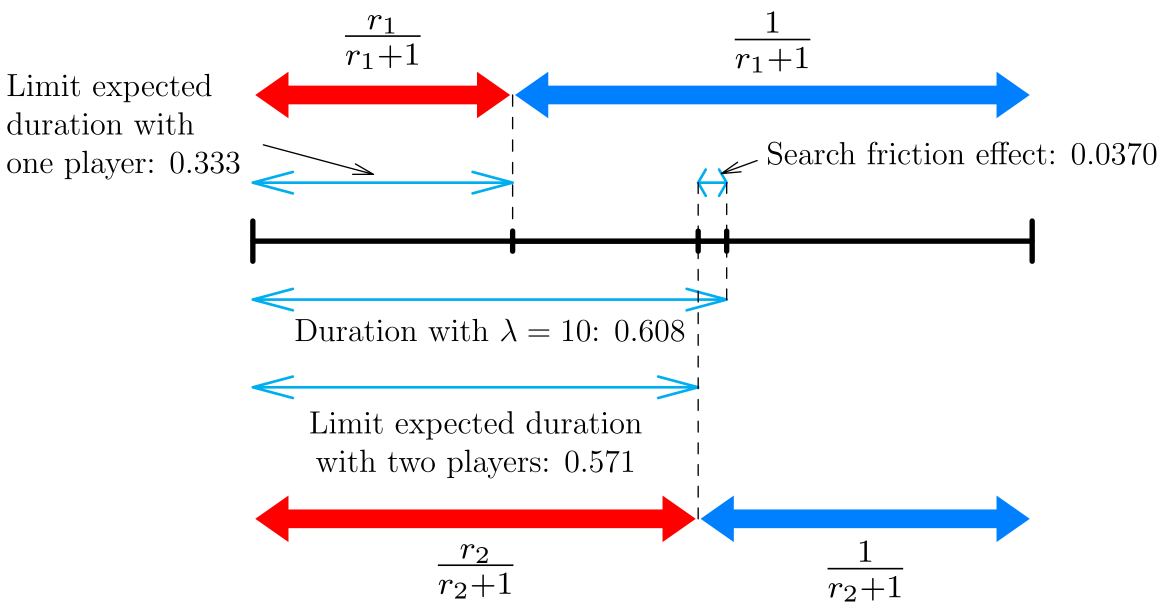

Figure 20: Decomposition of expected search durations: The case with uniform distribution over the space depicted in Figure 1 and the horizon length of 1. The one-player duration is computed by assuming uniform distribution over the unit interval. r_1 and r_2 are illustrated in Figure 21.

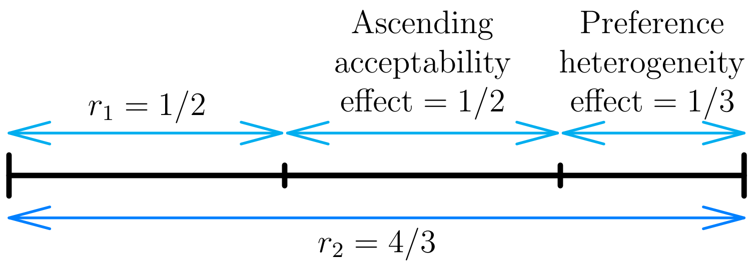

Figure 21: Decomposition of $r_2 - r_1$.

$r_1 := \frac{1}{3}$. 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 limit expected duration and r in this case are $\frac{1}{2}$ and $r_{\text{aa}} := 1$, respectively, and the difference in terms of r is given by $r_{\text{aa}} - r_1 = 1 - \frac{1}{2} = \frac{1}{2}$. Now the preference heterogeneity effect is the change in r caused by the change in the distribution from this product measure to X. This is given by $r_2 - r_{\text{aa}} = \frac{4}{3} - 1 = \frac{1}{3}$. In this example, the former effect is larger than the latter. Figure 21 illustrates these values.

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 $(X \times [0, 1]^m, \mu \times (U[0, 1])^m)$. In fact, this difference is $\frac{m}{2}$ by Theorem 4. 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).[82] Since r is additive by the definition of r in (3), 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 $l < m$.

B Appendix: Comprehensive 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.[83] Some recent papers in game theory discuss infinite-horizon search models in which a group of decision-makers determine when to stop. 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.

Wilson (2001), Compte and Jehiel (2010), and Cho and Matsui (2013) consider search models in which a unanimous agreement is required to accept an alternative, and show that the equilibrium payoff profile 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 Appendix A.1 where payoffs realize as soon as an agreement is reached and there is discounting. In Appendix A.7, we will discuss a common logic behind these convergence results. Wilson (2001) and Cho and Matsui (2013) show that an agreement is reached immediately in the limit as the frequency of offers goes to infinity. 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 payoff profiles. Albrecht et al. (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. The most related to ours is their result on the unanimity case, in which they show that the expected search duration increases in the number of agents. The logic is that the cutoff decreases in the number of agents while

[82]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.

[83]See Rogerson et al. (2005) for a survey.

the expected gain conditional on future acceptance does not change so much due to their distributional assumption, and hence the equilibrium condition implies that the expected wait time until the acceptance has to increase. We show the same comparative statics with respect to the number of agents. However, as we explain in Section 4.2.1, our logic does not rely on any distributional assumptions but relies on nonmonotonicity of cutoffs which is not present in their analysis where stationary equilibrium is assumed.

Alpern and Gal (2009), and Alpern et al. (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 (2013) analyze an infinite-horizon multi-agent search problem with interdependent preferences with respect to private signals of the payoffs independently realized in every period. They also show that the expected duration becomes longer if the number of decision-makers increases from one to two while retaining the information structure.[84] 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. Herings and Predtetchinski (2014) analyze an infinite-horizon search model with or without discounting in which alternatives are chosen according to general voting rules. They show that for each alternative in the core, there exists a stationary subgame perfect equilibrium implementing that alternative, and even if the core may be empty, there exists a subgame perfect equilibrium that sustains a stationary play on the path. Importantly, in the equilibria that all of these papers consider, the expected search durations converge to zero as the frequency of offer arrivals tends to infinity.

Multi-agent search with finite horizon.

A few papers consider multi-person search problems with finite horizon (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,[85] and then computed equilibrium strategies in several specific examples.[86]

[84] Moldovanu and Shi (2013) 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 expected duration is shorter with a larger conflict.

[85] 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.

[86] 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 expected search duration is trivially zero, our results imply that all cases have positive limit expected durations. In all of these examples, the feasible payoff set is unbounded or the Pareto frontier is a singleton, so an analysis of the limit expected

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 partly implied by this equation.

Ferguson (2005)'s main interest is in existence and uniqueness of the subgame perfect equilibrium sustained by Markov cutoff strategies in models with discrete time, general voting rules, varying distributions over time, and presence of fixed costs of search.[87,88] The sufficient condition for uniqueness that he obtains is different from ours.[89]

Single-agent search with finite horizon.
Search with finite horizon is a special class of a problem with search under nonstationarity. In the literature on job search theory, van den Berg (1990), Smith (1999), and others consider a single-agent search problem under nonstationary environments. These two papers analyze comparative statics with respect to changes in the primitives over time in a single-agent setting. In contrast, we consider a specific nonstationary environment, i.e., finite horizon, and analyze how the limit expected duration is affected by the number of agents or the payoff distributions.

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 problem 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 expected search duration. The reason for the 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. The tradeoff behind the positive duration is that the agent wants to wait until she gathers enough information, while she wants to have enough candidates in the future in order to decrease the probability of reaching the deadline without a good candidate. In our setting the first part of this tradeoff is absent, and instead the gain from waiting is only attributed to the future opportunities,

payoff profile was not an issue. In our general setup, such an analysis is meaningful because our model subsumes the case of bounded feasible payoff sets with non-singleton Pareto frontiers.

[87] He mentions the idea of trembling-hand equilibrium only verbally, and does not introduce it in his model.

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

[89] 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.

at which she may get a very good draw. 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 i.i.d. draw of payoffs over an infinite horizon. These models are extended in many directions.[90] 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 that 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 expected search duration as the friction goes away in models of this line of the literature is zero. 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 expected duration" as the friction vanishes.

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 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.[91]

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

[91] Cho and Matsui (2013) 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

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 the expected payoff profile cannot be weakly Pareto inefficient in the limit as the search friction vanishes.

Multi-agent search vs. bargaining with finite horizon.
Ambrus and Lu (2014), Gomes et al. (1999) and Imai and Salonen (2012) consider a bargaining model with finite horizon, in which players obtain an opportunity to propose a share distribution of the surplus at asynchronous times, having full control over proposals, and analyze the equilibrium payoffs.[92] 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 expected duration until the agreement can become arbitrarily small.[93] 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 et al. (1999) and Imai and Salonen (2012) shows that in some cases the limit equilibrium is the Nash bargaining solution. Although our result in Appendix A.1 about equilibrium payoff profiles is reminiscent of these results, the results are different. Since the proposer has full control over what to propose in their models, an agreement is reached at the first opportunity. On the other hand, in our model the average number of opportunities necessary to reach an agreement tends to infinity in the limit they consider. This causes the difference in the conditions under which the limit payoff profile is the Nash bargaining solution.[94]

interpretation, every player is "active" although the "activeness" is embedded in the model.

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

[93]A finite-horizon version of Rubinstein (1982)'s model with Poisson opportunities is a special case of Ambrus and Lu (2014)'s model, so the limit expected duration is zero in such a model.

[94]See Remark 2 in the previous version of this paper (Kamada and Muto (2011)) 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.

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 (2009, 2011) show the possibility of cooperation in such a setting,[95] and Calcagno et al. (2014) and Ishii and Kamada (2011) examine the effect of asynchronous timing of revisions on the equilibrium outcome in revision games. Kamada and Sugaya (2014) apply the revision games setting to election campaigns. Romm (2013) analyzes the implication of introducing a "reputational type" in the model of Calcagno et al. (2014). 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).[96]

C Appendix: Numerical Results for Finite Arrival Rates

We present five cases below, among which Cases 1–3 are the examples we discussed in Section 4.3.

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 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 4: 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 $\lambda = 10, 20, 30, 100$.

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 $\lambda = 10, 20, 30, 100$.

[95]See Ambrus et al. (2014) for a related work on an analysis of eBay-like auctions.

[96]For example, Example 4 in this paper is reminiscent of the "three state example" in Kamada and Sugaya (2010).

C.1 Uniform Distribution over Multi-Dimensional Triangle (Case 1)

Consider the distribution given by the uniform distribution over $\{x \in \mathbb{R}_+^n \mid \sum_{i\in N} x_i \leq 1\}$.

		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	0.608	0.591	0.585	0.576	0.572	0.57143
	Percentage (%)	6.48	3.44	2.35	0.731	0.0716	0
$n = 3$	Expected duration	0.716	0.705	0.701	0.695	0.693	0.692
	Percentage (%)	3.35	1.82	1.26	0.404	0.0430	0

C.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 following. Note that the limit duration for $n = 1$ is the same as in the case of uniform distribution over $\{x \in \mathbb{R}_+^n \mid \sum_{i\in N} x_i \leq 1\}$.

		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

C.3 Exponential Distribution (Case 3)

Consider the exponential distribution with parameter a_i for each player i.

		λ					
		10	20	30	100	1000	∞
$n = 1$	Expected duration	0.545	0.524	0.516	0.505	0.500	0.5
	Percentage (%)	9.09	4.76	3.23	0.990	0.0999	0
$n = 2$	Expected duration	0.693	0.681	0.676	0.670	0.667	0.667
	Percentage (%)	3.91	2.11	1.45	0.465	0.0489	0
$n = 3$	Expected duration	0.767	0.759	0.756	0.752	0.750	0.75
	Percentage (%)	2.27	1.24	0.864	0.284	0.0310	0
$n = 10$	Expected duration	0.912	0.911	0.910	0.910	0.909	0.909
	Percentage (%)	0.370	0.206	0.145	0.0499	0.00602	0

C.4 Uniform Distribution over a Cube (Case 4)

Consider the distribution given by the uniform distribution over $\{x \in \mathbb{R}^n_+ \mid \max_{i \in N} x_i \leq 1\}$.

		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	0.545	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

C.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 as follows:

		λ			
		10	20	30	100
$n = 1$	$\sigma = \frac{1}{4}$	0.449	0.462	0.469	0.484
	$\sigma = 1$	0.612	0.595	0.588	0.575
	$\sigma = 4$	0.961	0.952	0.946	0.926

D Appendix: Proofs of the Results

D.1 Computation of the Limit Expected Durations

We first present a method to compute the expected duration. Recall that $P(t; \lambda)$ is the probability that there is no agreement until time $-t$:

$$P(t; \lambda) = e^{-\int_t^T \lambda p(s;\lambda)\, ds}. \tag{D.1}$$

The expected duration $D(\lambda T)T$ is computed from $p(t;\lambda)$ and $P(t;\lambda)$ by integration by parts:

$$\begin{aligned} D(\lambda T)T &= T \cdot \underbrace{P(0;\lambda)}_{\substack{\text{The probability of no agreement} \\ \text{until time } 0}} \\ &\quad + \int_0^T \underbrace{(T-t)}_{\substack{\text{The duration when the search} \\ \text{ends at time } -t}} \cdot \underbrace{P(t;\lambda)}_{\substack{\text{The probability that} \\ \text{the search does not end until } -t}} \cdot \underbrace{\lambda p(t;\lambda)}_{\substack{\text{The probability density of} \\ \text{agreement at time } -t}} dt \\ &= T \cdot P(0;\lambda) + [(T-t)P(t;\lambda)]_0^T + \int_0^T P(t;\lambda)\, dt \\ &= \int_0^T P(t;\lambda)\, dt. \end{aligned} \tag{D.2}$$

The integral (D.2) has a direct interpretation: Since $P(t)$ is the probability that the duration is greater than $T-t$, we can compute the integral by integrating $T-t$ times $dP(t)$, which amounts to the area below the graph of $P(t)$. This is why the expression in (D.2) measures the expected duration.

Now we 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 21. *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 (D.1), 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 (D.2), $D(\lambda T)T = \int_0^T P(t;\lambda)dt$ is bounded as follows:

$$\int_{\varepsilon}^{T} \left(\frac{t}{T}\right)^{C} dt \leq D(\lambda T)T$$
$$\frac{T^{1+C} - \varepsilon^{1+C}}{(1+C)T^{C}} \leq D(\lambda T)T.$$

Since the above inequality is satisfied for all $\varepsilon > 0$ and sufficiently large λ, we have $\liminf_{\lambda\to\infty} D(\lambda T) \geq \frac{T^{1+C}}{(1+C)T^{C}\cdot T} = \frac{1}{1+C}$.

Next, a parallel argument shows (ii). Finally, (i) and (ii) together imply (iii). □

D.2 Proof of Proposition 1

Suppose that there exists at least one trembling-hand equilibrium. We show that the continuation payoffs of player i at time $-t$ are the same as each other 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), if there are two or more players, we can find a Lipschitz constant L_i for $i \in N$ such that $\mu(\{x \in X \mid x_i \in [x'_i, x''_i]\}) \leq L_i|x'_i - x''_i|$ for all x'_i, x''_i in the above domain of payoffs. Let $L_{-i} = \max_{j\neq i} L_j$.

Fix $\varepsilon \in (0, \frac{1}{2})$. For each $i \in N$ and each $-t \in [-T, 0]$, let $\overline{v}_i^{\varepsilon}(t)$ and $\underline{v}_i^{\varepsilon}(t)$ be the supremum and the infimum, respectively, of player i's expected payoff across every Nash equilibrium in every subgame starting at time $-t$ in the ε-constrained game Σ^{ε}. (Note that Assumption 1 (a) ensures boundedness of the continuation payoffs for finite t.) Since all subgames after time $-t$ are the same in our model, for all $\delta \in (0, T-t]$ and all $\eta > 0$, there exists a Nash equilibrium in which every player i obtains a continuation payoff larger than $\overline{v}_i^{\varepsilon}(t) - \eta$ (or smaller than $\underline{v}_i^{\varepsilon}(t) + \eta$) in each subgame after no players have an opportunity to play actions in time interval $[-(t+\delta), -t]$. Since the probability of arrival of an offer in $[-(t+\delta), -t]$ uniformly vanishes as $\delta \to 0$, $\overline{v}_i^{\varepsilon}(t)$ and $\underline{v}_i^{\varepsilon}(t)$ are continuous in t. Let $w_i^{\varepsilon}(t) = \overline{v}_i^{\varepsilon}(t) - \underline{v}_i^{\varepsilon}(t)$, and $\overline{w}^{\varepsilon}(t) = \max_{i\in N} w_i^{\varepsilon}(t)$. Note that $\overline{w}^{\varepsilon}(0) = 0$ for all ε.

For each $i \in N$ and each Nash equilibrium σ in the ε-constrained game Σ^{ε}, let $\overline{\overline{v}}_i^{\varepsilon}(t, \sigma)$ and $\underline{\underline{v}}_i^{\varepsilon}(t, \sigma)$ be the essential supremum and the essential infimum, respectively, of continuation payoffs $u_i(\sigma|\, h)$ across histories $h \in \tilde{\mathcal{H}}_t \setminus \mathcal{H}_t$ at time $-t$.[97] Let $\overline{\overline{v}}_i^{\varepsilon}(t) = \sup_{\sigma\text{: Nash eq. in } \Sigma^{\varepsilon}} \overline{\overline{v}}_i^{\varepsilon}(t, \sigma)$ and $\underline{\underline{v}}_i^{\varepsilon}(t) = \inf_{\sigma\text{: Nash eq. in } \Sigma^{\varepsilon}} \underline{\underline{v}}_i^{\varepsilon}(t, \sigma)$. We would like to show that for all $i \in N$ and all $-t \in [-T, 0]$, $\overline{\overline{v}}_i^{\varepsilon}(t) = \underline{\underline{v}}_i^{\varepsilon}(t)$. Since by the definitions, $\overline{v}_i^{\varepsilon}(t) \geq \overline{\overline{v}}_i^{\varepsilon}(t) \geq \underline{\underline{v}}_i^{\varepsilon}(t) \geq \underline{v}_i^{\varepsilon}(t)$, it suffices to show that $\overline{w}^{\varepsilon}(t) = 0$ for all $-t \in [-T, 0]$.

Let us consider strategies in the ε-constrained game. Suppose that a payoff profile

[97] Suppose that f is a real-valued function on a measure set $(Y, \mathcal{Y}, \gamma)$ which is measurable and bounded. The essential supremum of f is $\inf\{a \in \mathbb{R} \mid \gamma(f^{-1}((a, \infty))) = 0\}$, and the essential infimum of f is $\sup\{a \in \mathbb{R} \mid \gamma(f^{-1}((-\infty, a))) = 0\}$.

$x \in X$ is realized at time $-t$. If player i accepts x, 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 $\tilde{v}_i^\varepsilon(t) - \underline{v}_i^\varepsilon(t) = \frac{1}{\varepsilon^{n-1}} w_i^\varepsilon(t)$ by the definition of $w_i^\varepsilon(t)$.

Let $X_i^1(t) = \{x \in X \,|\, x_i > \tilde{v}_i^\varepsilon(t)\}$, $X_i^m(t) = \{x \in X \,|\, \underline{v}_i^\varepsilon(t) \le x_i \le \tilde{v}_i^\varepsilon(t)\}$, and $X_i^0(t) = \{x \in X \,|\, x_i < \underline{v}_i^\varepsilon(t)\}$. Then $\mu(X_i^m) \le \frac{L_i}{\varepsilon^{n-1}} w_i^\varepsilon(t)$ if $n \ge 2$. 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 $X = \big(\bigcup_{j\in N} X_j^m(t)\big) \cup \big(\bigcup_{(s_1,\dots,s_n)\in\{0,1\}^n} \bigcap_{j\in N} X_j^{s_j}(t)\big)$ (where $X_j^m(t)$'s have a nonempty intersection). Then

$$\begin{aligned}\overline{v}_i^\varepsilon(t) \le \int_0^t \Big(&\int_{X_i^m(\tau)} \tilde{v}_i^\varepsilon(t)\, d\mu + \sum_{j\neq i} \int_{X_j^m(\tau)} \overline{x}_i \, d\mu \\ &+ \sum_{(s_1,\dots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big((1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)} x_i \\ &\quad + (1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}) \overline{v}_i^\varepsilon(\tau)\big)\, d\mu\Big) \lambda e^{-\lambda(t-\tau)} d\tau,\end{aligned}$$

and

$$\begin{aligned}\underline{v}_i^\varepsilon(t) \ge \int_0^t \Big(&\int_{X_i^m(\tau)} \underline{v}_i^\varepsilon(t)\, d\mu + \sum_{j\neq i} \int_{X_j^m(\tau)} 0 \, d\mu \\ &+ \sum_{(s_1,\dots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big((1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)} x_i \\ &\quad + (1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}) \underline{v}_i^\varepsilon(\tau)\big)\, d\mu\Big) \lambda e^{-\lambda(t-\tau)} d\tau.\end{aligned}$$

Therefore $w_i^\varepsilon(t) = \overline{v}_i(t) - \underline{v}_i(t)$ is bounded as follows:

$$
\begin{aligned}
w_i^\varepsilon(t) \leq & \int_0^t \Big(\int_{X_i^m(\tau)} \frac{1}{\varepsilon^{n-1}} w_i^\varepsilon(t)\, d\mu + \sum_{j\neq i} \int_{X_j^m(\tau)} \overline{x}_i\, d\mu \\
& + \sum_{(s_1,\dots,s_n)\in\{0,1\}^n} \int_{\bigcap_{j\in N} X_j^{s_j}(\tau)} \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) w_i^\varepsilon(\tau) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau \\
\leq & \int_0^t \Big(\frac{1}{\varepsilon^{n-1}} w_i^\varepsilon(t) + \sum_{j\neq i} \overline{x}_i \frac{L_{-i}}{\varepsilon^{n-1}} w_j^\varepsilon(\tau) \\
& + \sum_{(s_1,\dots,s_n)\in\{0,1\}^n} \int_X \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) w_i^\varepsilon(\tau) d\mu \Big) \lambda e^{-\lambda(t-\tau)} d\tau \\
\leq & \int_0^t \Big(\frac{1}{\varepsilon^{n-1}} + \sum_{j\neq i} \max_{k\in N}\{\overline{x}_k\} \frac{L_{-i}}{\varepsilon^{n-1}} \\
& + \sum_{(s_1,\dots,s_n)\in\{0,1\}^n} \big(1-(1-\varepsilon)^{\sum_{j\in N} s_j} \varepsilon^{\sum_{j\in N}(1-s_j)}\big) \Big) \overline{w}^\varepsilon(\tau) \lambda e^{-\lambda(t-\tau)} d\tau.
\end{aligned}
$$

Since the above inequality holds for all $i \in N$, there exists a constant $M > 0$ such that the following inequality holds:[98]

$$
\overline{w}^\varepsilon(t) \leq \int_0^t M \overline{w}^\varepsilon(\tau) e^{-\lambda(t-\tau)} d\tau.
$$

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 $\frac{d}{dt}\big(W^\varepsilon(t)e^{-Mt}\big) = \big(W^{\varepsilon\prime}(t) - MW^\varepsilon(t)\big)e^{-Mt} \leq 0$ for all $t \geq 0$. 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 $t \geq 0$. Since $\overline{w}^\varepsilon(t) \geq 0$ for all $t \geq 0$ by definition, $\overline{w}^\varepsilon(t) = 0$ for all $t \geq 0$. Since $\varepsilon \in (0, 1/2)$ was arbitrary, any trembling-hand equilibria yield the same continuation payoff profile after almost all histories at time $-t \in [-T, 0]$.

D.3 Proof of Proposition 2

We show that a solution $v^*(t)$ of ODE (1) characterizes a trembling-hand equilibrium. For $s_i \in \{+,-\}$ and $v_i \in [0,\infty)$, let

$$
I_i^{s_i}(v_i) = \begin{cases} [v_i, \infty) & \text{if } s_i = +, \\ [0, v_i) & \text{if } s_i = -, \end{cases}
$$

[98] We note that M does not depend on L_i in the single-player case.

$p^+ = 1-\varepsilon$ and $p^- = \varepsilon$. For $\varepsilon > 0$, let us write down a Bellman equation similar to (2) with respect to a continuation payoff profile $v^\varepsilon(t)$ in the ε-constrained game:

$$v_i^\varepsilon(t) = \int_0^t \Big(\sum_{s\in\{+,-\}^n} \int_{(I_1^{s_1}(v_i^\varepsilon(\tau))\times\cdots\times I_n^{s_n}(v_i^\varepsilon(\tau)))\cap X} \big(p^{s_1}\dots p^{s_n}\cdot x_i + (1-p^{s_1}\dots p^{s_n})v_i^\varepsilon(\tau)\big)\, d\mu\Big) \cdot \lambda e^{-\lambda(t-\tau)}\, d\tau.$$

Since the right hand side is continuous in t, v_i^ε is continuous in t and thus bounded within $[0,T]$. By Assumption 1 (a), the summands in the right hand side are all bounded, and thus the right hand side is Lipschitz continuous in t. Therefore v_i^ε is differentiable in t almost everywhere.

Multiplying both sides of the above equality by $e^{\lambda t}$ and differentiating yield

$$v_i^{\varepsilon\prime}(t) = \lambda \sum_{s\in\{+,-\}^n} p^{s_1}\dots p^{s_n} \int_{(I_1^{s_1}(v_1^\varepsilon(t))\times\cdots\times I_n^{s_n}(v_n^\varepsilon(t)))\cap X} \big(x_i - v_i^\varepsilon(t)\big)\, d\mu$$

almost everywhere. This ODE has a unique solution because the right hand side is Lipschitz continuous in v_i^ε by Assumption 1 (b). (We show this Lipschitz continuity at the end of this proof. Note that an analogous argument shows that the Bellman equation (2) implies differentiability of $v_i(t)$.) Let $v^\varepsilon(t)$ be this solution, which is a cutoff profile of a Nash equilibrium in the ε-constrained game by construction. Since the acceptance set at $-t$ in the ε-constrained game is $(I_1^+(v_1^\varepsilon(t))\times\cdots\times I_n^+(v_n^\varepsilon(t)))\cap X$ for all $\varepsilon>0$, and $p^+\to 1$ and $p^-\to 0$ as $\varepsilon\to 0$, ODE (1) is obtained at almost all t by letting $\varepsilon\to 0$. Therefore $v^\varepsilon(t)$ converges to $v^*(t)$ at almost all t as $\varepsilon\to 0$ because the above ODE is continuous in ε.[99] Hence the cutoff strategy profile with cutoffs $v^*(t)$ is a trembling-hand equilibrium.

Finally, we prove that $\int_{(I_1^{s_1}(v_1)\times\cdots\times I_n^{s_n}(v_n))\cap X}(x-v)\, d\mu$ is Lipschitz continuous in v under Assumption 1. If $n\ge 2$, this is obvious because of Assumption 1 (b). Suppose $n=1$. Then if $s_1=+$, the change of the integral when v_1 varies is bounded as follows:

$$\begin{aligned}
&\left|\int_{I_1^{s_1}(v_1+\Delta v_1)\cap X}\big(x_1-(v_1+\Delta v_1)\big)\, d\mu - \int_{I_1^{s_1}(v_1)\cap X}(x_1-v_1)\, d\mu\right| \\
&= \left|\int_{[v_1+\Delta v_1,\infty)}\big(x_1-(v_1+\Delta v_1)\big)\, d\mu - \int_{[v_1,\infty)}(x_1-v_1)\, d\mu\right| \\
&= \left|-\int_{[v_1+\Delta v_1,\infty)}\Delta v_1\, d\mu - \int_{[v_1,v_1+\Delta v_1)}(x_1-v_1)\, d\mu\right| \\
&\le \left|\int_{[v_1+\Delta v_1,\infty)}\Delta v_1\, d\mu\right| + \left|\int_{[v_1,v_1+\Delta v_1)}\Delta v_1\, d\mu\right| \\
&= |\Delta v_1|\big(\mu([v_1+\Delta v_1,\infty)) + \mu([v_1,v_1+\Delta v_1))\big) \le 2|\Delta v_1|.
\end{aligned}$$

[99]See, e.g., Coddington and Levinson (1955, Theorem 7.4 in Chapter 1).

A similar computation shows that the same bound applies when $s_1 = -$. Therefore the integral is Lipschitz continuous with a Lipschitz constant 2.

D.4 Proof of Theorem 1

We prove the theorem under a weaker assumption than Assumption 2:

Assumption 7. There exist a concave function φ, constants $\kappa \geq 1$ and $\alpha > 0$, and $\tilde{x} < \sup X$ such that for all $x \in X$ with $x \geq \tilde{x}$,

$$\big(1 - \varphi(x)\big)^{\alpha} \leq 1 - F(x) \leq \kappa\big(1 - \varphi(x)\big)^{\alpha}.$$

We note that Assumption 2 is equivalent to Assumption 7 with $\alpha = 1$.

Proof of Theorem 1. Suppose that φ, κ, α, and $\tilde{x}$ are given in Assumption 7. We can assume that $\varphi(x)$ is increasing in $[0, \sup X)$ without loss of generality. Let $\beta = \dfrac{1+\alpha}{\alpha}$ (> 1). Let us consider a cutoff strategy with the cutoff $w(t)$ satisfying

$$F\big(w(t)\big) = 1 - \frac{\beta}{\lambda t + \beta},$$

namely, the strategy with acceptance probability $\frac{\beta}{\lambda t+\beta}$ at each time $-t$. Let $\tilde{t}$ be such that $w(\tilde{t}) = \tilde{x}$. By Assumption 7, we have $w(t) \geq \varphi^{-1}\big(1 - (\frac{\beta}{\lambda t+\beta})^{\beta-1}\big)$ for all $t \geq \tilde{t}$. Let $P(t)$ be the probability that the search does not stop before time $-t$ when $w(t)$ is played. Then for all $t \geq \tilde{t}$,

$$\begin{aligned} P(t) &= \exp\Big(-\int_t^T \frac{\beta}{\lambda\tau + \beta} \cdot \lambda d\tau\Big) \\ &= \Big(\frac{\lambda t + \beta}{\lambda T + \beta}\Big)^{\beta}. \end{aligned}$$

For all $t \geq \tilde{t}$, the continuation payoff obtained from this strategy is larger than

$$\int_t^{\tilde{t}} w(\tau) d\Big(\frac{1 - P(\tau)}{P(t)}\Big) \geq \int_{\tilde{t}}^t \varphi^{-1}\left(1 - \Big(\frac{\beta}{\lambda\tau + \beta}\Big)^{\beta-1}\right) \frac{dP(\tau)}{P(t)}.$$

Let $W(t)$ be the payoff on the right hand side. By concavity of φ, $\varphi(W(\tau))$ is bounded as follows: For all $t \geq \tilde{t}$,

$$\begin{aligned}
\varphi(W(t)) &= \varphi\left(\int_{\tilde{t}}^{t} \varphi^{-1}\Big(1 - \Big(\frac{\beta}{\lambda\tau + \beta}\Big)^{\beta-1}\Big)\frac{dP(\tau)}{P(t)}\right) \\
&\geq \int_{\tilde{t}}^{t} \varphi\left(\varphi^{-1}\Big(1 - \Big(\frac{\beta}{\lambda\tau + \beta}\Big)^{\beta-1}\Big)\right)\frac{dP(\tau)}{P(t)} \\
&= \int_{\tilde{t}}^{t} \left(1 - \Big(\frac{\beta}{\lambda\tau + \beta}\Big)^{\beta-1}\right) d\Big(\Big(\frac{\lambda\tau + \beta}{\lambda t + \beta}\Big)^{\beta}\Big) \\
&= 1 - \frac{\lambda\beta^{\beta}(t - \tilde{t}) + (\lambda\tilde{t} + \beta)^{\beta}}{(\lambda t + \beta)^{\beta}} \\
&\geq 1 - \frac{(\lambda t + \beta)\beta^{\beta} + (\lambda\tilde{t} + \beta)^{\beta}}{(\lambda t + \beta)^{\beta}}.
\end{aligned}$$

Let $\bar{t} \geq \tilde{t}$ be sufficiently large so that $(\lambda t + \beta)\beta^{\beta} \geq (\lambda\tilde{t} + \beta)^{\beta}$ for all $t \geq \bar{t}$. Then for all $t \geq \bar{t}$,

$$\varphi(W(t)) \geq 1 - \frac{2\beta^{\beta}}{(\lambda t + \beta)^{\beta-1}}.$$

Let $Q(t)$ be the probability that the search does not stop before time $-t$ when the player plays a cutoff strategy with cutoff $W(t)$. Then for all $t \geq \bar{t}$,

$$\begin{aligned}
Q(t) &= \exp\left(-\int_{t}^{T} \big(1 - F(W(\tau))\big)\lambda d\tau\right) \\
&\geq \exp\left(-\int_{t}^{T} \kappa\big(1 - \varphi(W(\tau))\big)^{\alpha}\lambda d\tau\right) \qquad \text{(by Assumption 7)} \\
&\geq \exp\left(-\int_{t}^{T} \kappa\Big(\frac{2\beta^{\beta}}{(\lambda t + \beta)^{\beta-1}}\Big)^{\alpha}\lambda d\tau\right) \\
&= \exp\left(-\int_{t}^{T} 2\kappa\beta^{\alpha}\Big(\frac{\beta}{\lambda t + \beta}\Big)\lambda d\tau\right) \\
&= \left(\frac{\lambda t + \beta}{\lambda T + \beta}\right)^{2\kappa\beta^{\alpha+1}}
\end{aligned}$$

which is bounded away from zero irrespective of λ for all $-t \in [-T, 0)$ whenever $T \geq \bar{t}$. 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 probability that the search does not stop before time $-t$ when the player plays an optimal strategy is smaller than or equal to $Q(t)$. Since $\inf_{\lambda>0} Q(t) > 0$ for all $-t \in [-T, 0)$, the search stops with probability strictly lower than 1 before time $-t \in [-T, 0)$ even in the limit as $\lambda \to \infty$. This proves Theorem 1. □

D.5 A Generalized Version of Theorem 4 and Its Proof

Here we state a generalized version of Theorem 4 and provide its proof. We first define a few pieces of notation to deal with the case when the limits that are assumed to exist in Theorem 4 in the main text does not necessarily exist.

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

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

where

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

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. We note that if $n = 1$, $\underline{d}_i(v)$ and $\overline{d}_i(v)$ can be infinity.

Recall that $P(t;\lambda)$ is the probability of no agreement until time $-t$, and $D(\lambda)$ is the expected duration when $T = 1$. Now we can show that $P(t;\infty) = \lim_{\lambda\to\infty} P(t;\lambda)$ and the limit expected duration $D(\infty) = \lim_{\lambda\to\infty} D(\lambda)$ can be written in the following way: For each $r > 0$, let $P^r(t) = \left(\frac{t}{T}\right)^{1/r}$, and let $P^0(t) = 0$ if $-t \in (-T, 0]$ and $P^0(T) = 1$.

Theorem 4′. Under Assumption 1, for all $-t \in [-T, 0]$

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

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

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

Proof. First, suppose that $\underline{r} > 0$. By ODE (1), $v_i^{*\prime}(t) = \lambda b_i(v^*(t)) \cdot p(t)$ for each $i \in N$. Since $\mu(A(v))$ is continuous in v by Assumption 1 (b) if $n \geq 2$, we have

$$\liminf_{\varepsilon\to 0} \frac{p(t) - p(t+\varepsilon)}{\varepsilon} = \liminf_{\varepsilon\to 0} \frac{\mu(A(v^*(t))) - \mu(A(v^*(t) + \varepsilon v^{*\prime}(t)))}{\varepsilon}$$

for all t. Note that this holds even when $n = 1$. Let $\varepsilon' = \frac{v_1^{*\prime}(t)}{b_1(v^*(t))}\varepsilon$ $(= \frac{v_i^{*\prime}(t)}{b_i(v^*(t))}\varepsilon$ for all

$i \in N$ because $v^{*\prime}(t)$ is parallel to $b(v^*(t))$). Then, for all t,

$$\begin{aligned}
&\liminf_{\varepsilon\to 0} \frac{\mu(A(v^*(t))) - \mu(A(v^*(t) + \varepsilon v^{*\prime}(t)))}{\varepsilon} \\
&= \liminf_{\varepsilon\to 0} \sum_{i\in N} \frac{\mu(A((v_j^*(t) + \varepsilon v_j^{*\prime}(t))_{j<i}, (v_j^*(t))_{j\geq i})) - \mu(A((v_j^*(t) + \varepsilon v_j^{*\prime}(t))_{j\leq i}, (v_j^*(t))_{j>i})))}{\varepsilon} \\
&= \liminf_{\varepsilon'\to 0} \sum_{i\in N} \frac{\mu(A((v_j^*(t) + \varepsilon' b_j(v^*(t)))_{j<i}, (v_j^*(t))_{j\geq i})) - \mu(A((v_j^*(t) + \varepsilon' b_j(v^*(t)))_{j\leq i}, (v_j^*(t))_{j>i})))}{\frac{b_i(v^*(t))}{v_i^{*\prime}(t)}\varepsilon'} \\
&\geq \sum_{i\in N} \liminf_{\varepsilon'\to 0} \frac{\mu(A((v_j^*(t) + \varepsilon' b_j(v^*(t)))_{j<i}, (v_j^*(t))_{j\geq i})) - \mu(A((v_j^*(t) + \varepsilon' b_j(v^*(t)))_{j\leq i}, (v_j^*(t))_{j>i})))}{\frac{b_i(v^*(t))}{v_i^{*\prime}(t)}\varepsilon'} \\
&= \sum_{i\in N} \big(\underline{d}_i(v^*(t))p(t) \cdot v_i^{*\prime}(t)\big) \\
&= \sum_{i\in N} \big(\underline{d}_i(v^*(t))p(t) \cdot \lambda b_i(v^*(t))p(t)\big).
\end{aligned}$$

By the definition of $\underline{r}$, for all $\eta > 0$, there exists $\bar{t}$ such that for all $t \geq \bar{t}$,

$$\frac{\liminf_{\Delta t\to 0} \frac{p(t)-p(t+\Delta t)}{\Delta t}}{\lambda p(t)^2} \geq \underline{r} - \eta.$$

Integrating the both sides from 0 to t, we have

$$\frac{p(t)}{1-p(t)} \cdot \lambda t \leq (\underline{r} - \eta)^{-1}$$

and letting $\lambda \to \infty$ and $\eta \to 0$, we have

$$\limsup_{\lambda\to\infty} p(t) \cdot \lambda t \leq \underline{r}^{-1}.$$

By Lemma 21, we obtain

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

If $\underline{r} = 0$, we obviously have $\liminf_{\lambda\to\infty} P(t;\lambda) \geq P^0(t)$.

Next, suppose that $\overline{r} > 0$. An analogous argument shows that

$$\liminf_{\lambda\to\infty} p(t) \cdot \lambda t \geq \overline{r}^{-1}.$$

By Lemma 21, we obtain

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

Finally, suppose that $\underline{r} = \overline{r} = 0$. Then an analogous argument shows that

$$\lim_{\lambda\to\infty} p(t)\cdot\lambda t = \infty.$$

By Lemma 21, this implies that $P(t;\infty) \geq \left(\frac{t}{T}\right)^a$ for all $a > 0$. Thus, $P(t;\infty) = P^0(t)$. In such a case, $D(\infty) = 0$. □

D.6 Proof of Proposition 5

We prove Proposition 5 in an environment more general than Assumption 4.

Assumption 8. (a) The limit $v^* = \lim_{\lambda\to\infty} v^*(t)$ exists and 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 $i \in N$.[100]

(d) For all $\eta > 0$, there exists $\varepsilon > 0$ such that $\{x \in \mathbb{R}^n_+ \,|\, |v^* - x| \leq \varepsilon, \alpha\cdot(x - v^*) \leq -\eta\}$ is contained in X, where "$\cdot$" denotes the inner product in $\mathbb{R}^n$.

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

Proposition 5′. Under Assumptions 1 and 8, $\lim_{\lambda\to\infty} D(\lambda) = \dfrac{n^2}{n^2+n+1}$.

Proof. Let $f_H(t) = \sup_{x\in A(t)} f(x)$, and $f_L(t) = \inf_{x\in A(t)} f(x)$. Since f is continuous and strictly positive, both $f_H(t)$ and $f_L(t)$ are continuous and converge to $f(v^*)$ (> 0) as $t \to \infty$. For each $\eta > 0$ and each t, let

$$\underline{A}(t) = \{x \in \mathbb{R}^n_+ \,|\, x \geq v^*(t), \alpha\cdot(x - v^*) \leq -\eta\},$$
$$\overline{A}(t) = \{x \in \mathbb{R}^n_+ \,|\, x \geq v^*(t), \alpha\cdot(x - v^*) \leq \eta\}.$$

The volume of $\underline{A}(t)$ (with respect to the Lebesgue measure on $\mathbb{R}^n$) is

$$V(\underline{A}(t)) = \frac{1}{n!}\prod_{j\in N}\Big(\frac{\alpha\cdot(v^* - v^*(t)) - \eta}{\alpha_j}\Big),$$

and the volume of $\overline{A}(t)$ is

$$V(\overline{A}(t)) = \frac{1}{n!}\prod_{j\in N}\Big(\frac{\alpha\cdot(v^* - v^*(t)) + \eta}{\alpha_j}\Big).$$

[100] If X is convex as assumed in Assumption 4, the normal vector at v^* must satisfy this property. This can be shown by a similar technique to the one in the proof of Proposition 7.

By Assumption 8 (b), (c), and (d), for each $\eta > 0$, there exist $\varepsilon > 0$ and $\bar{t}$ such that for all $t \geq \bar{t}$, $|v^* - v^*(t)| \leq \varepsilon$ and $\underline{A}(t) \subseteq A(t) \subseteq \overline{A}(t)$.

For all $t \geq \bar{t}$, the barycenter term is bounded as

$$\frac{f_L(\bar{t}) \int_{\underline{A}(t)} \big(x_i - v_i^*(t)\big) dx}{f_H(\bar{t}) V(\overline{A}(t))} \leq b_i(t) \leq \frac{f_H(\bar{t}) \int_{\overline{A}(t)} \big(x_i - v_i^*(t)\big) dx}{f_L(\bar{t}) V(\underline{A}(t))},$$
$$\frac{f_L(\bar{t}) V(\underline{A}(t)) \big(\frac{\alpha \cdot (v^* - v^*(t)) - \eta}{(n+1)\alpha_i}\big)}{f_H(\bar{t}) V(\overline{A}(t))} \leq b_i(t) \leq \frac{f_H(\bar{t}) V(\overline{A}(t)) \big(\frac{\alpha \cdot (v^* - v^*(t)) + \eta}{(n+1)\alpha_i}\big)}{f_L(\bar{t}) V(\underline{A}(t))}.$$

For all $t \geq \bar{t}$, $\mu(A(v^*(t)))$ is differentiable in v_i by Assumption 8 (e), and the density term is bounded as

$$\frac{f_L(\bar{t}) n V(\underline{A}(t)) \big(\frac{\alpha \cdot (v^* - v^*(t)) - \eta}{\alpha_i}\big)^{-1}}{f_H(\bar{t}) V(\overline{A}(t))} \leq d_i(t) \leq \frac{f_L(\bar{t}) n V(\overline{A}(t)) \big(\frac{\alpha \cdot (v^* - v^*(t)) + \eta}{\alpha_i}\big)^{-1}}{f_H(\bar{t}) V(\underline{A}(t))}.$$

Then $r(t) := \sum_{i \in N} b_i(t) d_i(t)$ is bounded as

$$\sum_{i \in N} \frac{n f_L(\bar{t})^2 V(\underline{A}(t))^2}{(n+1) f_H(\bar{t})^2 V(\overline{A}(t))^2} \leq r(t) \leq \sum_{i \in N} \frac{n f_H(\bar{t})^2 V(\overline{A}(t))^2}{(n+1) f_L(\bar{t})^2 V(\underline{A}(t))^2}$$

for all $t \geq \bar{t}$. By Assumption 8 (b), $V(\underline{A}(t))/V(\overline{A}(t))$ converges to 1 as $\eta \to 0$ and $t \to \infty$. Letting $\eta \to 0$ and $t \to \infty$, therefore, we have

$$r = \lim_{t \to \infty} r(t) = \sum_{i \in N} \frac{n}{n+1} = \frac{n^2}{n+1}.$$

Finally, Theorem 4 implies that

$$D(\infty) = \lim_{\lambda \to \infty} D(\lambda) = \frac{1}{1 + \big(\frac{n^2}{n+1}\big)^{-1}} = \frac{n^2}{n^2 + n + 1}.$$

□

D.7 Proof of Theorem 5

We can follow the discussion in the proof sketch of Theorem 4, and obtain inequality (4). Since this inequality holds for any λ, for all $\varepsilon > 0$, there is a large $\bar{t}$ such that for all $t \geq \bar{t}$

$$-(r + \varepsilon) p(t; 1)^2 \leq \frac{\partial p}{\partial t}(t; 1) \leq -(r - \varepsilon) p(t; 1)^2.$$

For any $\lambda > 0$, let $\eta = \bar{t}/\lambda$. Since $p(t; \lambda) = p(t/\lambda; 1)$ for all t and λ, we have

$$-(r + \varepsilon) p(t; \lambda)^2 \leq \frac{\partial p}{\partial t}(t; \lambda) \leq -(r - \varepsilon) p(t; \lambda)^2$$

for all $t \geq \eta$. We now simplify the notation as $p(t) = p(t;\lambda)$. Solving this with an initial condition at η, for all $\lambda > 0$ and $t \geq \eta$,

$$\frac{1}{(r-\varepsilon)\lambda(t-\eta)+p(\eta)^{-1}} \leq p(t) \leq \frac{1}{(r+\varepsilon)\lambda(t-\eta)+p(\eta)^{-1}}.$$

By formula (D.1), for all $\lambda > 0$ and $t \geq \eta$, we have

$$e^{-\int_t^T \frac{1}{(r+\varepsilon)(s-\eta)+p(\eta)^{-1}/\lambda}\,ds} \leq P(t) \leq e^{-\int_t^T \frac{1}{(r-\varepsilon)(s-\eta)+p(\eta)^{-1}/\lambda}\,ds}$$

$$\Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r-\varepsilon)^{-1}} \leq P(t) \leq \Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r+\varepsilon)^{-1}}.$$

By formula (D.2), for all $\lambda > 0$ and $t \geq \eta$, we have

$$\int_\eta^T \Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r-\varepsilon)^{-1}} dt \leq D(\lambda T)T \leq \eta + \int_\eta^T \Big(\frac{r\lambda(t-\eta)+p(\eta)^{-1}}{r\lambda(T-\eta)+p(\eta)^{-1}}\Big)^{(r+\varepsilon)^{-1}} dt$$

$$\frac{1}{1+(r-\varepsilon)^{-1}}\Big(T-\eta+\frac{1}{r\lambda p(\eta)}-\frac{1}{r\lambda p(\eta)(T-\eta)+1}\Big)$$
$$\leq D(\lambda T)T \leq \eta + \frac{1}{1+(r+\varepsilon)^{-1}}\Big(T-\eta+\frac{1}{r\lambda p(\eta)}-\frac{1}{r\lambda p(\eta)(T-\eta)+1}\Big).$$

Since the above inequalities are satisfied for all $\varepsilon > 0$, all $\lambda > 0$, and $\eta = \bar{t}/\lambda$, and $D(\infty) = \frac{1}{1+r^{-1}}$, by letting $\varepsilon \to 0$, we have

$$-\frac{\bar{t}}{1+r^{-1}} + \frac{1}{1+r^{-1}}\Big(\frac{1}{rp(\eta)} - \frac{1}{rp(\eta)(T-\eta)+1}\Big)$$
$$\leq \lambda\big(D(\lambda T)T - D(\infty)T\big) \leq \frac{\bar{t}}{1+r} + \frac{1}{1+r^{-1}}\Big(\frac{1}{rp(\eta)} - \frac{1}{rp(\eta)(T-\eta)+1}\Big).$$

Since we showed $\lim_{\lambda\to\infty} p(t)\cdot\lambda t = 1/r$ in the proof sketch of Theorem 4, $p(\eta) = p(\bar{t}/\lambda) \to 1/(r\bar{t})$ as $\lambda \to \infty$. Hence, $\lambda\big(D(\lambda T)T - D(\infty)T\big)$ is bounded for all λ. When $T = 1$, we obtain $|D(\lambda) - D(\infty)| = O(\frac{1}{\lambda})$.

D.8 Proof of Proposition 7

Let $f_L = \inf_{x\in X} f(x) > 0$, $f_H = \sup_{x\in X} f(x)$, and $\overline{x}_i = \max\{x_i \,|\, x \in X\}$ for $i \in N$. Assumption 5 ensures existence of these values.

We first prove a lemma under Assumptions 1 and 5, together with a couple of conditions: Given a fixed payoff profile $w \in \hat{X}$ that is weakly Pareto efficient but not Pareto efficient, let $I = \{i \in N \,|\, {}^\exists y \in X \text{ such that } y_i > w_i, y_j \geq w_j\, {}^\forall j \in N\}$ and $J = N \setminus I$. Since w is not Pareto efficient, I is nonempty. For such I and J, and $v \in \hat{X}$, let $f_J(v) = \int_{\{x_I \geq v_I\}} f(x_I, v_J)\, dx_I$ be the marginal density of J in $A(v)$. For $\varepsilon \geq 0$ and $v \in \hat{X}$, let $Y_J^\varepsilon(v) = \{x_J \in \mathbb{R}_+^{|J|} \,|\, x_J \geq v_J, f_J(v_I, x_J) > \varepsilon\}$. For a bounded set $Y_J \subseteq \mathbb{R}^{|J|}$ with a positive Lebesgue measure, let $b(Y_J) = \int_{Y_J} x_J \, dx_J / \int_{Y_J} dx_J$ be the barycenter of Y_J with respect to the Lebesgue measure.

Lemma 22. *Suppose that Assumptions 1 and 5 hold. For each $w \in \hat{X}$ that is weakly Pareto efficient but not Pareto efficient, if (i) $A(w)$ is convex, (ii) $f_J(w) > 0$, and (iii) there exist $\alpha \in \mathbb{R}^n_+ \setminus \{0\}$, $\tilde{\delta} > 0$, and $M > 0$ such that $\alpha_i = 0$ for all $i \in I$, and $\alpha \cdot \big((v_I, b(Y^0_J(v))) - v\big) \leq M\big(\alpha \cdot (w - v)\big)$ for all $\delta \in (0, \tilde{\delta}]$ and all $v \in \hat{X}$ with $w_i - v_i \in (0, \delta]$ for all $i \in N$, then w cannot be the limit v^* of the solution $v^*(t)$.*

Proof of Lemma 22. By condition (i), there exists $y \in X$ such that $y_i > w_i$ for all $i \in I$. Since w is weakly Pareto efficient, $I \neq N$, namely, $J \neq \emptyset$.

Assume on the contrary that $v^* = w$, which is assumed to be not Pareto efficient. Since $I \neq \emptyset$, we can fix $\bar{i} \in I$ arbitrarily. Let $z(t) = \alpha \cdot \big(v^* - v^*(t)\big)$. We will show that there exists $\bar{t}$ such that for all $t \geq \bar{t}$,

$$-z'(t) \leq \frac{z(\bar{t})}{2\big(v^*_{\bar{i}} - v^*_{\bar{i}}(\bar{t})\big)} v^{*\prime}_{\bar{i}}(t).$$

If this inequality is shown, integrating both sides yields

$$-z(t) + z(\bar{t}) \leq \frac{z(\bar{t})}{2\big(v^*_{\bar{i}} - v^*_{\bar{i}}(\bar{t})\big)} \big(v^*_{\bar{i}}(t) - v^*_{\bar{i}}(\bar{t})\big)$$

for all $t \geq \bar{t}$. By letting $t \to \infty$, we have $0 < z(\bar{t}) \leq z(\bar{t})/2$, and a contradiction follows.

Let $\varepsilon = f_J(v^*)/4$, which is positive by condition (ii). By Assumption 5, there exists $\delta \in (0, \tilde{\delta}]$ such that $Y^{2\varepsilon}_J(v^*_I, v^*_J - (\delta, \ldots, \delta)) = Y^0_J(v^*_I, v^*_J - (\delta, \ldots, \delta))$. Let δ be sufficiently small so that $\delta^{|I|} f_L \leq \varepsilon/2$. Then $Y^{\varepsilon}_J(v^*_I, v^*_J - (\delta, \ldots, \delta)) \supseteq Y^{2\varepsilon}_J(v^* - (\delta, \ldots, \delta))$. Let $\bar{\delta} \in (0, \delta]$ be sufficiently small so that $\bar{\delta} \leq \frac{\varepsilon^2}{4(f_H)^2 M \prod_{i \in I} \overline{x}_i}$.

Since $v^*(t)$ converges to v^* as $t \to \infty$, there exists $\bar{t}$ such that $\max_{i \in N}(v^*_i - v^*_i(t)) \leq \bar{\delta}$. By equation (1), for all $t \geq \bar{t}$,

$$\begin{aligned}
-z'(t) &= \lambda \sum_{j \in J} \alpha_j \int_{A(t)} \big(x_j - v^*_j(t)\big) \, d\mu \\
&\leq \lambda \sum_{j \in J} \alpha_j \int_{Y^0_J(v^*(t))} \big(x_j - v^*_j(t)\big) \int_{\prod_{i \in I}[0, \overline{x}_i]} f_H \, dx_I \, dx_J \\
&= \lambda f_H \Big(\prod_{i \in I} \overline{x}_i\Big) \sum_{j \in J} \alpha_j \int_{Y^0_J(v^*(t))} \big(x_j - v^*_j(t)\big) \, dx_J \\
&\leq \lambda f_H \Big(\prod_{i \in I} \overline{x}_i\Big) M z(t) \int_{Y^0_J(v^*(t))} dx_J \qquad \text{(by condition (iii))} \\
&\leq \lambda f_H \Big(\prod_{i \in I} \overline{x}_i\Big) M z(\bar{t}) \int_{Y^0_J(v^*(t))} dx_J.
\end{aligned}$$

By equation (1) again, for all $t \geq \bar{t}$,

$$\begin{aligned}
v_{\bar{i}}^{*\prime}(t) &= \lambda \int_{A(t)} \big(x_{\bar{i}} - v_{\bar{i}}^*(t)\big)\, d\mu \\
&\geq \lambda \int_{\{(x_I,x_J)\,|\,x_I \geq v_I^*, x_J \in Y_J^{\varepsilon}(v_I^*, v_J^*(t))\}} (x_{\bar{i}} - v_{\bar{i}}^*)\, d\mu \\
&\geq \lambda \int_{Y_J^{\varepsilon}(v_I^*, v_J^*(t))} \frac{\varepsilon}{2f_H} \cdot \varepsilon\, dx_J \\
&\geq \frac{\lambda \varepsilon^2}{2f_H} \int_{Y_J^0(v^*(t))} dx_J.
\end{aligned}$$

Then for all $t \geq \bar{t}$,

$$\begin{aligned}
-\frac{z'(t)}{v_{\bar{i}}^{*\prime}(t)} \cdot \frac{v_{\bar{i}}^* - v_{\bar{i}}^*(\bar{t})}{z(\bar{t})} &\leq \frac{2(f_H)^2 M \big(\prod_{i \in I} \overline{x}_i\big)\big(v_{\bar{i}}^* - v_{\bar{i}}^*(\bar{t})\big)}{\varepsilon^2} \\
&\leq \frac{v_{\bar{i}}^* - v_{\bar{i}}^*(\bar{t})}{2\bar{\delta}} \leq \frac{1}{2}.
\end{aligned}$$

This proves the lemma. □

Proposition 7 is now easily shown.

Proof of Proposition 7. Suppose that X is convex, and $w \in \hat{X}$ is weakly Pareto efficient but not Pareto efficient. It suffices to show that w satisfies three conditions (i)–(iii) in Lemma 22. Condition (i) is obvious. To show Condition (ii), observe that the definition of I and convexity of X imply existence of $y \in X$ such that $y_i > w_i$ for all $i \in I$ and $y_j \geq w_j$ for all $j \in J$. By concavity of X, $f_J(w) \geq \frac{f_L}{|I|!} \prod_{i \in I}(y_i - w_i)$, which is strictly positive.

We show condition (iii). Since X is convex, $\hat{X}$ is also convex. Then there exists $\alpha \in \mathbb{R}_+^n \setminus \{(0, \ldots, 0)\}$ such that $\alpha_i = 0$ for all $i \in I$, and $\alpha \cdot (w - v) \leq 0$ for all $v \in \hat{X}$. Since $x_J \in Y_J^0(v)$ implies that $\alpha \cdot \big(w - (x_I, x_J)\big) \leq 0$ for all $x_I \in \mathbb{R}^{|I|}$, condition (iii) is satisfied for this α, any $\bar{\delta} > 0$, and $M = 1$. □

D.9 Examples of Pareto Inefficiency and Proof of Proposition 15

Before proving the proposition, we present two counterexamples in which the condition assumed in the statement of Proposition 15 fails. First, in the following example, X is written by functions which are not piecewise continuous, and openness fails.

Example 9. Let $n = 2$ and $Y = \{x \in \mathbb{R}_+^2 \,|\, x_1 + x_2 \leq 1\}$. Let $(a^k)_{k=1,2,\ldots}$ be a sequence in the interior of Y such that a_1^k is decreasing in k, a_2^k is increasing in k, and $\lim_{k\to\infty} a^k = (1/2, 1/2)$. We consider $X = Y \setminus \bigcup_{k=1}^{\infty} \{x \in Y \,|\, x_i > a_i^k \text{ for } i = 1, 2\}$. Then $X = \big\{(x_1, x_2) \in \mathbb{R}_+^2 \,|\, x_2 \in [0, g(x_1)], x_1 \in [0, h(x_2)]\big\}$ where g and h have countably infinite discontinuous points.

For each $k = 1, 2, \dots$, one can construct a density $f^k \in \mathcal{F}$ such that $v^*(t)$ converges to a^k as in Example 4. Then, for $f = \lim_{k\to\infty} f^k \in \mathcal{F}$, $v^*(t)$ converges to $\lim_{k\to\infty} a^k = (1/2, 1/2)$, which is Pareto efficient in X. Therefore $\{f \in \mathcal{F} \mid \lim_{t\to\infty} v^*(t) \text{ is Pareto efficient}\}$ is not open. $\square$

Second, in the next example, if X is written by the form $X = \big\{(x_1, x_2) \in \mathbb{R}^2_+ \mid x_2 \in [0, g(x_1)], x_1 \in [0, h(x_2)]\big\}$, then h cannot be quasiconcave, and denseness fails.

Example 10. Let $n = 2$ and μ be the uniform distribution on $X = X_1 \cup X_2 \cup X_3$ where $X_1 = [0, 1]^2$, $X_2 = \big\{(x_1, x_2) \in \mathbb{R}^2_+ \mid x_1 \in [0, 1], x_2 \in [2 - (1 - x_1)^3, 2]\big\}$, and $X_3 = \big\{(x_1, x_2) \in \mathbb{R}^2_+ \mid x_2 \in [0, 1], x_1 \in [2 - (1 - x_2)^3, 2]\big\}$. By symmetry with respect to the 45 degree line, we have $\lim_{t\to\infty} v^*(t) = (1, 1)$, which is not Pareto efficient in X.

Let us consider a probability measure $\tilde{\mu}$ with the same support X which assigns slightly higher probability than μ near $(1, 2) \in X_2$. Since $\mu\big(X_2 \cap A(1-\varepsilon, 1-\varepsilon)\big) = O(\varepsilon^4)$, and $\int_{X_1 \cap A(1-\varepsilon,1-\varepsilon)} \big(v_2 - (1 - \varepsilon)\big)\, d\mu$ is of order ε^3, the slight change of probability near $(1, 2) \in X_2$ does not affect the limit outcome. A similar argument shows that the slight change of probability near $(2, 1) \in X_1$ does not affect the limit outcome. Therefore, the density function $f \in \mathcal{F}$ of the uniform distribution belongs to the interior of the complement of $\{f \in \mathcal{F} \mid \lim_{t\to\infty} v^*(t) \text{ is Pareto efficient}\}$. This implies that this set is not dense. $\square$

Proof of Proposition 15. In the proof, we denote by $v^*(t; f)$ the solution of ODE (1) for density function $f \in \mathcal{F}$, and $v^*(f) = \lim_{\lambda\to\infty} v^*(t; f) = \lim_{t\to\infty} v^*(t; f)$. We want to show that $\mathcal{F}^e := \{f \in \mathcal{F} \mid v^*(f) \text{ is Pareto efficient in } X\}$ is open and dense in $\mathcal{F}$.

Proof of openness: Before proving openness, we show that the mapping $v^*(\infty; \cdot) : \mathcal{F} \to \mathbb{R}^n$ is continuous at all $f \in \mathcal{F}^e$, i.e., for all $f \in \mathcal{F}^e$, all $\eta > 0$, and any sequence $f_k \in \mathcal{F}$ $(k = 1, 2, \dots)$ with $f_k \to f$ $(k \to \infty)$, there exists $\bar{k}$ such that

$$|v^*(f_k) - v^*(f)| \le \eta$$

for all $k \ge \bar{k}$.

Since $\lim_{t\to\infty} v^*(t; f) = v^*(f)$, for all $\delta > 0$ there exists $\bar{t} > 0$ such that $|v^*(f) - v^*(t; f)| \le \delta$ for all $t \ge \bar{t}$. By Pareto efficiency of $v^*(f)$, let $\delta > 0$ be sufficiently small so that $A\big(v^*(\bar{t}; f) - (\delta, \delta, \dots, \delta)\big)$ is contained in the η-ball centered at $v^*(f)$. Since the right hand side of ODE (1) is continuous in v by Assumption 1 (b), the unique solution of (1) is continuous with respect to parameters in (1). Therefore, for a finite time interval $[0, T]$ including $\bar{t}$, there exists $\bar{k}$ such that $|v^*(t; f_k) - v^*(t; f)| \le \delta$ for all $t \in [0, T]$ and all $k \ge \bar{k}$. This implies that $v^*(\bar{t}; f_k) \in A\big(v^*(\bar{t}; f) - (\delta, \delta, \dots, \delta)\big)$, thereby $v^*(f_k) \in A\big(v^*(\bar{t}; f) - (\delta, \delta, \dots, \delta)\big)$. Therefore we have $|v^*(f_k) - v^*(f)| \le \eta$. Hence $v^*(f)$ is continuous at all $f \in \mathcal{F}^e$.

Now we prove openness. By the assumption on X, a weakly Pareto efficient payoff profile $w \in \hat{X}$ violates the assumptions in Lemma 22 only if g is discontinuous at w_1 and there exists $\varepsilon > 0$ such that $h(y_2) = h(w_2)$ for all $y_2 \in [w_2 - \varepsilon, w_2]$, or h is continuous at w_2 and there exists $\varepsilon > 0$ such that $g(y_1) = g(w_1)$ for all $y_1 \in [w_1 - \varepsilon, w_1]$. Since g and h are piecewise continuous, the cardinality of the set of weakly Pareto efficient payoff profiles w satisfying the above property is finite. Therefore, there exists $\eta > 0$ such that if w is Pareto efficient in X, w' is weakly Pareto efficient, g is discontinuous at w'_1, and h is discontinuous at w'_2, then $|w - w'| > \eta$. Since v^* is continuous at all f such that $v^*(f)$ is Pareto efficient, $v^*(\tilde{f})$ must be Pareto efficient whenever $\tilde{f}$ is sufficiently close to f. Hence $\mathcal{F}^e$ is open in $\mathcal{F}$.

Proof of denseness: By the definition of X, and Lemma 22, if $w := v^*(f)$ is not Pareto efficient, then g is discontinuous at w_1 or h is discontinuous at w_2 (or both). Since g and h are piecewise continuous, there exists $\bar{t}$ such that for all $t \geq \bar{t}$, $A(v^*(t; f)) \setminus \{w\}$ contains no point at which g and h are discontinuous. Thus we can assume without loss of generality that g is continuous on $[0, w_1) \cup (w_1, \overline{x}_1]$ and h is continuous on $[0, w_2) \cup (w_2, \overline{x}_2]$. Under this assumption, for any $\bar{f} \in \mathcal{F}$, the limit payoff profile $v^*(\bar{f})$ is Pareto efficient, or $v^*(\bar{f}) = w$.

Suppose that w is not Pareto efficient, and assume without loss of generality that h is discontinuous at w_2. For each $\bar{f} \in \mathcal{F}$, each $v \in \hat{X}$, and $i \in N$, let $b_i(v; \bar{f})$ be the barycenter term with respect to $\bar{f}$. By quasiconcavity of g and h, and continuity of g on $[0, w_1)$ and of h on $[0, w_2)$, the limit $\lim_{v \nearrow w} b_i(v; \bar{f})$ exists for each $i \in N$. Let $\beta(t, \bar{f}) := \dfrac{b_2(v^*(t; \bar{f}); \bar{f})}{b_1(v^*(t; \bar{f}); \bar{f})}$. We note that for all t, $\dfrac{v_2^{*\prime}(t'; \bar{f})}{v_1^{*\prime}(t; \bar{f})} = \beta(t, \bar{f})$ by ODE (1).

Let $\bar{y} \in X$ be the Pareto efficient payoff profile such that $\bar{y}_1 > w_1$ and $\bar{y}_2 = w_2$, and let $y = \big(\lim_{v \nearrow w} b_1(v; f), w_2\big)$. Let $y = (\bar{y} + y)/2$. By quasiconcavity of h, and the definition of X, we have $[w_1, \bar{y}_1] \times \{w_2\} \subset X$, and by continuity of h on $[0, w_2)$, there exists $\varepsilon \in (0, y_1 - w_1]$ such that $[y_1 - \varepsilon/2, y_1 + \varepsilon/2] \times [y_2 - \varepsilon, y_2] \subset X$.

Since $v_i'^*(t; f) > 0$ for all $i = 1, 2$ and all t, $v_2'^*(0; f)/v_1'^*(0; f) > 0$. This together with the definition of X imply that for each t, there exists $z(t)$ in the interior of X such that $z_i(t) > v_i^*(t; f)$ for all t and $i = 1, 2$, $\dfrac{\big(z_2(t) - v_2^*(t; f)\big)}{\big(z_1(t) - v_1^*(t; f)\big)} < \beta(t, f)$ for all t, and $z(t) \to y$ as $t \to \infty$. For small $\varepsilon > 0$, the ε-neighborhood Z of the trajectory of $z(t)$ is contained in X. Let $\tilde{f}$ be a continuous density function on Z that is positive in the interior of Z, assigns zero on the boundary of Z, and satisfies $\dfrac{b_2\big(v^*(t; f); \tilde{f}\big)}{b_1\big(v^*(t; f); \tilde{f}\big)} < \beta(t, f)$ for all t. We note that this inequality holds if the density function is constructed by taking densities decreasing exponentially as $v^*(t; f)$ diverges from the origin. We define $f_k \in \mathcal{F}$ by

$$f_k(x) = \frac{k}{k+1} f(x) + \frac{1}{k+1} \tilde{f}(x)$$

for each $x \in X$ and $k \geq 1$. Note that this f_k satisfies Assumption 4 for all $k \geq 1$. Since we assumed that $\dfrac{b_2\big(v^*(t;f);\tilde{f}\big)}{b_1\big(v^*(t;f);\tilde{f}\big)} < \beta(t,f)$ for all t, we have $\dfrac{b_2\big(v^*(t;f);f_k\big)}{b_1\big(v^*(t;f);f_k\big)} < \beta(t,f)$ for all t and all $k \geq 1$.

We assume that there exists $\bar{k}$ such that $\lim_{t\to\infty} v^*(t;f_k) = w$ for all $k \geq \bar{k}$, and derive a contradiction. Since $\lim_{t\to\infty} \beta(t,f_k) \leq \lim_{t\to\infty} \beta(t,f)$, and $\dfrac{b_2(v;f_k)}{b_1(v;f_k)} < \dfrac{b_2(v;f)}{b_1(v;f)}$ for all v such that v is close to w, and $v_1 < w_1$ and $v_2 < w_2$, there exists t_k such that $v^*(t_k;f_k)$ comes above the trajectory of $v^*(t;f)$. By continuity of the trajectory, there exist t and t' such that $v^*(t;f_k) = v^*(t';f)$ and $\dfrac{v_2^{*\prime}(t;f_k)}{v_1^{*\prime}(t;f_k)} > \dfrac{v_2^{*\prime}(t';f)}{v_1^{*\prime}(t';f)}$. However, by ODE (1), this inequality never holds because we have $\beta(t,f_k) < \beta(t',f)$, implying $\dfrac{v_2^{*\prime}(t;f_k)}{v_1^{*\prime}(t;f_k)} < \dfrac{v_2^{*\prime}(t';f)}{v_1^{*\prime}(t';f)}$. This is a contradiction.

Therefore for all $\bar{k}$, there exists $k \geq \bar{k}$ such that $v^*(f_k) \neq w$. Hence $\mathcal{F}^e$ is dense in $\mathcal{F}$.

□

D.10 Proof of Proposition 9

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 $D = \{z \in \mathbb{R}^n \,|\, w \cdot z = 0\}$. For $\xi > 0$, let D_ξ be an $(n-1)$-dimensional disk defined by

$$D_\xi = \{z \in D \,|\, |z| \leq \xi\}.$$

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 $w_z \in X$.

Proof of Proposition 9. Let $B_\varepsilon(y) = \{x \in X \,|\, |y - x| \leq \varepsilon\}$ for $y \in X$ and $\varepsilon > 0$. Let g_ε be an arbitrary continuous function on $\mathbb{R}^n$ such that $g_\varepsilon(x) > 0$ if x is in the interior of the n-dimensional ball centered at $0 \in \mathbb{R}^n$ with radius ε, and $g_\varepsilon(x) = 0$ otherwise. Let $\tilde{f}$ be the uniform density function on X. For a Pareto efficient payoff profile $y \in X$, we define a probability density function f_y on X by

$$f_y(x) = \eta \tilde{f}(x) + (1-\eta)\frac{g_\varepsilon(y-x)}{\int_{x' \in B_\varepsilon(y)} g_\varepsilon(y-x')\,dx'}$$

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 (1) with density f_{w_z}, and define a function φ from D_ξ to D by $\varphi(z) = \tilde{\varphi}(z) - \delta w \in D$ where $\delta \in \mathbb{R}$ is chosen to satisfy $\tilde{\varphi}(z) - \delta w \in D$ depending on z. Thus, φ measures the projection to D of the vector from w to the limit payoff profile under f_w. By the form of ODE (1), the solution of (1) with density f_{w_z} changes continuously if z moves continuously. Since w is not located 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. In the first step, we will show that there exists $\xi > 0$ such that $\psi(D_\xi) \subseteq D_\xi$ where $\psi(z) := z - \varphi(z)$. In the second step, we will use the first step to apply the fixed point theorem to ψ and show that there exists z such that $\varphi(z) = z$.

Step 1: We show that there exist $\xi > 0$, $\varepsilon > 0$ and $\eta > 0$ such that $|\varphi(z) - z| \leq \xi$ for all $z \in D_\xi$. Take any $\xi > 0$ such that φ is continuous in D_ξ. 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 payoff profile under 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 payoff profile 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, we can take such small $\varepsilon > 0$ and $\eta > 0$ independent of z.

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. □

D.11 Proof of Proposition 10

Fix any $\varepsilon > 0$. Since $\lim_{t\to\infty} v^*(t; 0, \lambda)$ exists for all λ, there exists $\bar{t} > 0$ such that for all $t \geq \bar{t}$,

$$|v^*(t; 0, 1) - v^*(\bar{t}; 0, 1)| \leq \varepsilon/2. \tag{D.3}$$

Since the right hand side of ODE (A.1) is continuous in ρ and λ, and uniformly Lipschitz continuous in v, the unique solution $v^*(t; \rho, \lambda)$ is continuous in ρ and λ for all $t \in [0, \bar{t}]$. Therefore by continuity in ρ, there exists $\bar{\lambda} > 0$ such that for all $\lambda \geq \bar{\lambda}$ and all $t \in [0, \bar{t}]$,

$$|v^*(t; 0, 1) - v^*(t; \rho/\lambda, 1)| \leq \varepsilon/2. \tag{D.4}$$

By (D.3) and (D.4), for all $\lambda' \geq \bar{\lambda}$ and all $t \geq 0$

$$|v^*(t; 0, 1) - v^*(\min\{t, \bar{t}\}; \rho/\lambda', 1)| \leq \varepsilon.$$

Recalling that $v^*(t;\rho,\alpha\lambda) = v^*(\alpha t;\rho/\alpha,\lambda)$ for all $\alpha > 0$, we have

$$|v^*(t/\lambda;0,\lambda) - v^*(\min\{t,\bar{t}\}/\lambda';\rho,\lambda')| \leq \varepsilon$$

for all $\lambda,\lambda' \geq \bar{\lambda}$ and all $t \geq 0$. Let $\tilde{t} = t/\lambda$ and $t' = \min\{t,\bar{t}\}/\lambda'$. Then for all $\lambda,\lambda' \geq \bar{\lambda}$ and all $\tilde{t} \geq 0$, there exists t' such that

$$\left|v^*(\tilde{t};0,\lambda) - v^*(t';\rho,\lambda')\right| \leq \varepsilon.$$

If $\lambda = \lambda'$, we have $t' = \min\{t/\lambda,\bar{t}/\lambda\} = \min\{\tilde{t},\bar{t}/\lambda\}$. By replacing $\bar{t}/\lambda$ by $\bar{t}$, we have shown that for all $\varepsilon > 0$, there exists $\bar{\lambda}$ such that for all $\lambda,\lambda' \geq \bar{\lambda}$, there exists $\bar{t} > 0$ such that for all $\tilde{t}$,

$$\left|v^*(\tilde{t};0,\lambda) - v^*(\min\{\tilde{t},\bar{t}\};\rho,\lambda')\right| \leq \varepsilon.$$

Letting $\lambda' = \lambda$, we obtain the desired inequality.

D.12 Proof of Lemma 11

For each $i \in N$, let $\overline{x}_i = \max\{x_i \,|\, (x_i,x_{-i}) \in X\}$, which exists by Assumption 5 (a). To simplify notations, let $v(t)$ be the solution of ODE (A.1) for given ρ and λ. Fix $\bar{t} > 0$ arbitrarily. The proof consists of five steps. In the first step, we show that the acceptance probability becomes small as λ gets large. This enables approximation of the distribution conditional on the acceptance set for large λ. This approximation is used in the second step to identify the approximate direction of the vector from $v(t)$ to the barycenter. Applying this approximation, we show that the Nash product is increasing in t in the third and fourth steps. Finally, in the fifth step, we show that $v(t)$ converges to a Nash point independent of t.

Step 1: We show that $\mu(A(\bar{t})) \to 0$ as $\lambda \to \infty$. If not, there exist a positive value $\varepsilon > 0$ and an increasing sequence $(\bar{\lambda}_k)_{k=1,2,\ldots}$ such that $\mu(A(\bar{t})) \geq \varepsilon$ for all $\bar{\lambda}_k$. Since $v_i'(t') \geq -\rho\overline{x}_i$ for each $i \in N$ and all t' by equation (A.1), $v_i(\tilde{t}) \leq v_i(\bar{t}) + \rho\overline{x}_i(\bar{t}-\tilde{t})$ for all $i \in N$ and all $\tilde{t} < \bar{t}$. Then for all $\tilde{t} < \bar{t}$,

$$\begin{aligned}
\mu\big(A(v(\bar{t})) \setminus A(v(\tilde{t}))\big) &\leq \sum_{i\in\{i\,|\,v_i(\tilde{t})>v_i(\bar{t})\}} \mu\Big([v_i(\bar{t}),v_i(\tilde{t})] \times \prod_{j\neq i}[0,\overline{x}_j]\Big) \\
&\leq f_H \sum_{i\in N} \rho\overline{x}_i(\bar{t}-\tilde{t}) \prod_{j\neq i} \overline{x}_j \\
&= (\bar{t}-\tilde{t}) f_H \rho n \prod_{j\in N} \overline{x}_j.
\end{aligned}$$

Let $\tilde{t} = \bar{t} - \dfrac{\varepsilon/2}{f_H\rho n\prod_{j\in N}\overline{x}_j}$. Then we have $\mu(A(t)) \geq \varepsilon/2$ for all $t \in [\tilde{t},\bar{t}]$. Since for all

$t \in [\tilde{t}, \bar{t}]$ and all $\eta > 0$,

$$\mu\big(A(v(t)) \setminus A(v(t) + (\eta, \dots, \eta))\big) \leq \sum_{i \in N} \mu\Big([v_i(t), v_i(t) + \eta] \times \prod_{j \neq i} [0, \overline{x}_j]\Big)$$
$$\leq \sum_{i \in N} \Big(f_H \eta \prod_{j \neq i} \overline{x}_j\Big),$$

we have $\mu\big(A(v(t) + (\eta, \dots, \eta))\big) \geq \varepsilon/4$ by letting $\eta = \dfrac{\varepsilon/4}{f_H \sum_{i \in N} \prod_{j \neq i} \overline{x}_j}$. For this η, the integral in ODE (A.1) is bounded from below as

$$\int_{A(t)} \big(x_i - v_i(t)\big) d\mu \geq \int_{A(v(t)+(\eta,\dots,\eta))} \big(x_i - v_i(t)\big) d\mu$$
$$\geq \int_{A(v(t)+(\eta,\dots,\eta))} \eta d\mu$$
$$\geq \eta\varepsilon/4$$

for all $t \in [\tilde{t}, \bar{t}]$. By ODE (A.1), for all $t \in [\tilde{t}, \bar{t}]$ and all $k \geq 1$,

$$v_i'(t) \geq -\rho \overline{x}_i + \bar{\lambda}_k \eta \varepsilon/4.$$

Therefore, $v_i(\bar{t}) - v_i(\tilde{t}) \geq (-\rho \overline{x}_i + \bar{\lambda}_k \eta \varepsilon/4)(\bar{t} - \tilde{t}) > \overline{x}_i$ for sufficiently large $\bar{\lambda}_k$. This contradicts the definition of $\overline{x}_i$ because $v_i(\tilde{t}) \geq 0$.

Step 2: We will consider approximation of $v(t)$ by approximating the shape of $A(v(t))$ when t is in a small time interval including $\bar{t}$. We note that this approximation is valid in a time interval that is independent of λ. This is because $v_i'(t) \geq -\rho \overline{x}_i$ for all $i \in N$ and all t by equation (A.1), so even when $v(t)$ moves away from the Pareto frontier, its speed is bounded independent of λ. We will compute the direction of $\int_{A(t)} \big(x - v(t)\big) d\mu$ in the limit as $\lambda \to \infty$ based on this approximation.

By Step 1, the boundary of X contains all accumulation points of $\{v(\bar{t}) \,|\, \lambda > 0\}$. Fix an accumulation point $v^*(\bar{t})$, and take an increasing sequence $(\lambda_k)_{k=1,2,\dots}$ with $v^*(\bar{t}) = \lim_{k \to \infty} v(\bar{t})$. By Assumption 6, there exists a unit normal vector of X at $v^*(\bar{t})$, which we denote by $\alpha \in \mathbb{R}_{++}$.

Step 1 implies that $v(\bar{t})$ is very close to the boundary of X when λ_k is very large. By Assumption 6 (a), when t is close to $\bar{t}$, $A(t)$ is approximated by a polyhedron defined by convex hull of $\{v(t), v(t) + (z_1(t), 0, \dots, 0), v(t) + (0, z_2(t), 0, \dots, 0), \dots, v(t) + (0, \dots, 0, z_n(t))\}$ where for each $i \in N$, $z_i(t) > 0$ is chosen such that $v(t) + (0, \dots, 0, z_i(t), 0, \dots, 0)$ is on the boundary of $\hat{X}$. This vector $z(t)$ is approximately parallel to $(1/\alpha_1, \dots, 1/\alpha_n)$. Let $\zeta(t)$ be the ratio between the length of $z(t)$ and $(1/\alpha_1, \dots, 1/\alpha_n)$, i.e., $\zeta(t) = z_1(t)\alpha_1 = \cdots = z_n(t)\alpha_n$.

Since density f is continuous by Assumption 5 (b), probability measure μ conditional on $A(t)$ is approximated by the uniform distribution on $A(t)$ if λ_k is large. Then the integral $\int_{A(t)} \big(x - v(t)\big) d\mu$ is approximated by $\mu\big(A(t)\big)$ times the vector from $v(t)$ to the barycenter of the polyhedron, namely, $\mu\big(A(t)\big) z(t)/(n+1)$. Therefore $\int_{A(t)} \big(x - v(t)\big) d\mu$ is approximately parallel to $(1/\alpha_1, \dots, 1/\alpha_n)$ when λ_k is large.

Step 3: We show that $\sum_{i \in N} \alpha_i v_i'(\bar{t}) \geq 0$ for large λ. Let $(\lambda_k)_{k=1,2,\dots}$ be the sequence defined in Step 2. For large λ_k and t close to $\bar{t}$, $A(t)$ is again approximated by a polyhedron, and μ conditional on $A(t)$ is approximated by the uniform distribution on $A(t)$. By Step 2, for each $i \in N$, the ODE near $v_i(\bar{t})$ is approximated by

$$v_i'(t) = -\rho v_i(t) + \lambda_k \frac{z_i(t)}{n+1} \cdot \mu(A(t)). \tag{D.5}$$

Note that $v_i(t)$ is close to $v_i^*(\bar{t})$ because t is close to $\bar{t}$, and $\mu(A(t))$ is of order n of the length of $z(t)$. By replacing the above equation by $\zeta(t)$, ODE (D.5) approximates

$$\zeta'(t) = \rho a - \lambda_k b \zeta(t)^{n+1} \tag{D.6}$$

for some constants $a, b > 0$.

Suppose that $\mu(A(t))$ is not decreasing in t at $\bar{t}$. Then $\zeta'(\bar{t}) \geq 0$. By (D.6), this implies that there exists $\tau > 0$ such that $v(t)$ is closer or equally close to the Pareto frontier for any $t \in [\bar{t} - \tau, \bar{t}]$. Since the approximation explained in Step 2 is valid as long as $v(t)$ is close the Pareto frontier, τ can be taken arbitrarily large. Therefore, the approximation of (D.6) holds for all $t \in [0, \bar{t}]$, and thus $v(t)$ must be closer or equally close to the Pareto frontier to any $t \in [0, \bar{t}]$. This contradicts the fact that $v(0) = 0$. Therefore $\mu(A(t))$ is decreasing in t near $\bar{t}$. For large λ_k, this implies that the distance from $v(t)$ to the Pareto frontier, which is proportional to $\alpha \cdot \big(v^*(\bar{t}) - v(t)\big)$, is decreasing, and thus

$$\alpha \cdot v'(\bar{t}) = \sum_{i \in N} \alpha_i v_i'(\bar{t}) \geq 0.$$

Step 4: We show that the Nash product is nondecreasing in t at $\bar{t}$ if λ is large. By ODE (D.5), we have

$$\alpha_i v_i'(\bar{t}) = -\rho \alpha_i v_i(\bar{t}) + \beta \tag{D.7}$$

where $\beta = \lambda_k \mu(A(\bar{t}))/(n+1)$ independent of i. Let us assume without loss of generality that $\alpha_1 v_1'(\bar{t}) \geq \cdots \geq \alpha_n v_n'(\bar{t})$. Then we must have $1/(\alpha_1 v_1(\bar{t})) \geq \cdots \geq 1/(\alpha_n v_n(\bar{t}))$.

Let $L(t) = \sum_{i \in N} \ln v_i(t)$ be a logarithm of the Nash product. By Chebyshev's sum inequality,

$$L'(\bar{t}) = \sum_{i \in N} \frac{v_i'(\bar{t})}{v_i(\bar{t})} = \sum_{i \in N} \frac{\alpha_i v_i'(\bar{t})}{\alpha_i v_i(\bar{t})} \geq \frac{1}{n} \Big(\sum_{i \in N} \alpha_i v_i'(\bar{t}) \Big) \Big(\sum_{i \in N} \frac{1}{\alpha_i v_i(\bar{t})} \Big) \geq 0.$$

Therefore, the Nash product is nondecreasing in t at $\bar{t}$ if λ_k is large. Moreover, the derivative is zero if and only if $\alpha_1 v_1'(\bar{t}) = \cdots = \alpha_n v_n'(\bar{t})$ or $\alpha_1 v_1(\bar{t}) = \cdots = \alpha_n v_n(\bar{t})$.

Step 5: We show that $\bar{t}$, $v(\bar{t})$ converges to a point in the Nash set as $\lambda \to \infty$, and the limit is irrelevant to the choice of $\bar{t}$. Since Step 1 implies that $L'(\bar{t})$ converges to zero as $\lambda_k \to \infty$, we have $\alpha_1 v_1'(\bar{t}) = \cdots = \alpha_n v_n'(\bar{t})$ or $\alpha_1 v_1(\bar{t}) = \cdots = \alpha_n v_n(\bar{t})$ in the limit as $\lambda_k \to \infty$. If the former case holds, then ODE (D.7) implies that the latter case also holds. Therefore the latter case always holds in the limit as $\lambda_k \to \infty$, i.e., $\alpha_1 v_1(\bar{t}) = \cdots = \alpha_n v_n(\bar{t})$. This implies that the boundary of X at $v(\bar{t})$ is tangent to the hypersurface $\left\{y \in \mathbb{R}^n_+ \,\middle|\, \prod_{i\in N} y_i = \prod_{i \in N} v_i(\bar{t})\right\}$. Hence any accumulation point $v^*(\bar{t})$ at every $\bar{t}$ belongs to the Nash set.

Since we assumed that the Nash set consists of isolated points, the accumulation point $v^*(\bar{t})$ is isolated in the Nash set. Suppose that $v(\bar{t})$ does not converge to $v^*(\bar{t})$ as $\lambda \to \infty$. Then for any $\delta > 0$ and any $\bar{\lambda}$, there exists $\lambda \geq \bar{\lambda}$ such that $|v(\bar{t}) - v^*(\bar{t})| \geq \delta/2$. Let $\delta > 0$ be small such that $\{x \in X \mid |v^*(\bar{t}) - x| \leq \delta\} \setminus \{v^*(\bar{t})\}$ has no intersection with the Nash set. Since $v(\bar{t})$ is continuous with respect to λ and $v^*(\bar{t})$ is an accumulation point, for any $\bar{\lambda}$, there exists $\lambda > \bar{\lambda}$ such that $\delta/2 \leq |v(\bar{t}) - v^*(\bar{t})| \leq \delta$. Since $\{x \in X \mid \delta/2 \leq |v^*(\bar{t}) - x| \leq \delta\}$ is compact, $v(\bar{t})$ must have an accumulation point in this set. This contradicts the fact that any accumulation point is contained in the Nash set. Since the choice of $\bar{t}$ was arbitrary, we have shown that for all t, $v^*(t) = \lim_{\lambda\to\infty} v(t)$ exists and belongs to the Nash set.

Finally, we show that $v^*(t)$ is independent of t. Suppose not. Then there exist t_1 and t_2 such that $v^*(t_1) \neq v^*(t_2)$. Since $v(t)$ converges to $v^*(t)$ as $\lambda \to \infty$ for each $t \in [t_1, t_2]$, for any $\delta > 0$, there exists $\bar{\lambda} > 0$ such that $|v^*(t) - v(t)| \leq \delta/3$ for all $\lambda \geq \bar{\lambda}$ and all $t \in [t_1, t_2]$. Since $v(t)$ is continuous with respect to t, for any $\delta > 0$ and $\lambda = \bar{\lambda}$, there exists $\varepsilon > 0$ such that if $t, t' \in [t_1, t_2]$ and $|t - t'| \leq \varepsilon$, then $|v(t) - v(t')| \leq \delta/3$. Let $\lambda = \bar{\lambda}$. These imply that for all $\delta > 0$, there exists $\varepsilon > 0$ such that if $t, t' \in [t_1, t_2]$ and $|t - t'| \leq \varepsilon$, then $|v^*(t) - v^*(t')| \leq \delta$. If $\delta > 0$ is such that $\{x \in X \mid |v^*(t_1) - x| \leq \delta\} \setminus \{v^*(t_1)\}$ has no intersection with the Nash set, this implies that there exists $\varepsilon > 0$ such that if $t, t' \in [t_1, t_2]$ and $|t - t'| \leq \varepsilon$, then $v^*(t) = v^*(t')$. Thus $v^*(t_1) = v^*(t_2)$, which is a contradiction. Hence $v^*(t)$ is constant with respect to t.

D.13 Proof of Lemma 12

The ODE (D.6) is rewritten as

$$\zeta'(t) = -\lambda b \left(\zeta(t) - \left(\frac{\rho a}{\lambda b}\right)^{\frac{1}{n+1}} \right) \left(\sum_{k=0}^{n} \zeta(t)^{n-k} \left(\frac{\rho a}{\lambda b}\right)^{\frac{k}{n+1}} \right)$$

where λ_k is simply written as λ which is sufficiently large. Since $\zeta'(t) < 0$ by Step 3 in the proof of Lemma 11 and $\zeta''(t) > 0$ by ODE (D.6), we have $\zeta(t - \varepsilon) - \zeta(t) > -\varepsilon\zeta'(t)$ (> 0)

for any small $\varepsilon > 0$. This implies $\lim_{\lambda\to\infty} \zeta'(t) = 0$ for all $t > 0$ because $\lim_{\lambda\to\infty} \zeta(t) = 0$ for all $t > 0$ by Lemma 11. Thus, given $t > 0$, $\zeta'(t)$ must be close to zero for sufficiently large λ.

Therefore for sufficiently large λ, $\zeta(t)$ is approximated by $\zeta(t) = \left(\frac{\rho a}{\lambda b}\right)^{\frac{1}{n+1}}$ because the term in the second parentheses is greater than $\left(\frac{\rho a}{\lambda b}\right)^{\frac{n}{n+1}} > 0$. Since $\mu(A(t))$ is approximated by a function proportional to $\zeta(t)^n$, $\mu(A(t)) = c\rho^{\frac{n}{n+1}}\lambda^{-\frac{n}{n+1}}$ where $c > 0$ is a constant. The probability that an agreement is reached before time $-(T-s)$ is

$$1 - e^{-\int_{T-s}^{T}\mu(A(t))\lambda\,dt} = 1 - e^{-sc\rho^{\frac{n}{n+1}}\lambda^{\frac{1}{n+1}}},$$

which converges to one as $\lambda \to \infty$.

D.14 Proof of Proposition 13

The approximated ODE (D.6) for large λ in the proof of Proposition 11 can be rearranged as follows:

$$\lambda^{\frac{1}{n}}\zeta'(t) = \lambda^{\frac{1}{n}}\rho_\lambda a - b\cdot\left(\lambda^{\frac{1}{n}}\zeta(t)\right)^{n+1}.$$

If $\lambda^{\frac{1}{n}}\rho_\lambda \to 0$, this ODE is approximated as

$$\begin{aligned}\lambda^{\frac{1}{n}}\zeta'(t) &\approx -b\left(\lambda^{\frac{1}{n}}\zeta(t)\right)^{n+1}\\ \zeta'(t) &\approx -\lambda b\zeta(t)^{n+1}.\end{aligned} \tag{D.8}$$

An argument similar to the proof of Theorem 4 shows that the above approximation (D.8) is applied also to ODE (1) whenever $\mu(A(t))$ is close to zero. If $\mu(A(t))$ is far away from zero, since ρ_λ is bounded, Proposition 10 shows that the solution of ODE (A.1) is approximated by that of ODE (1). Therefore, in both cases, the solution of ODE (A.1) is approximated by that of ODE (1), and thus $v^* = \lim_{\lambda\to\infty} v^*(t; 0, \lambda)$. Since (D.8) yields $\zeta(t) = O\left(\frac{1}{(\lambda t)^{\frac{1}{n}}}\right)$, an argument similar to Theorem 4 shows the result of positive duration.

On the other hand, if $\lambda^{\frac{1}{n}}\rho_\lambda \to \infty$, the ODE is approximated as

$$\begin{aligned}\lambda^{\frac{1}{n}}\zeta'(t) = &\left((\lambda^{\frac{1}{n}}\rho_\lambda a)^{\frac{1}{n+1}} - b^{\frac{1}{n+1}}(\lambda^{\frac{1}{n}}\zeta(t))\right)\\ &\cdot\left((\lambda^{\frac{1}{n}}\rho_\lambda a)^{\frac{n}{n+1}} + (\lambda^{\frac{1}{n}}\rho_\lambda a)^{\frac{n-1}{n+1}}\cdot b^{\frac{1}{n+1}}(\lambda^{\frac{1}{n}}\zeta(t)) + \cdots + b^{\frac{n}{n+1}}(\lambda^{\frac{1}{n}}\zeta(t))^n\right)\\ \approx &(\lambda^{\frac{1}{n}}\rho_\lambda a) - (\lambda^{\frac{1}{n}}\rho_\lambda a)^{\frac{n}{n+1}}\cdot b^{\frac{1}{n+1}}(\lambda^{\frac{1}{n}}\zeta(t)).\end{aligned}$$

Since $\zeta'(t) \to 0$ as $\lambda \to \infty$ for all $t > 0$ by the argument in the proof of Lemma 12, $\zeta(t)$ is approximated by $\zeta(t) = \left(\frac{\rho_\lambda a}{\lambda b}\right)^{\frac{1}{n+1}}$. Thus, by the proof of Lemma 12, the limit expected duration is zero.

D.15 Proof of Proposition 14

In the main text, we have already shown the first statement, which immediately implies the "if" part in the second statement. Thus, we only prove the "only-if" part in the second statement, i.e., that if the core is non-empty, then there exists a probability measure μ with support X satisfying Assumption 1 such that the limit expected duration is positive under μ. Let $B^\varepsilon = \{x \in X \,|\, |c - x| \leq \varepsilon \text{ for some } c \text{ in the core}\}$. By the assumptions, there exist a Pareto efficient core allocation $c \in X$ and $\delta \in (0, 1)$ such that $Y(\delta) \cap B^\varepsilon \neq \emptyset$ for all $\varepsilon > 0$, where $Y(\delta) = \{x \in X \,|\, c_i - x_i \geq \delta(c_j - x_j) > 0 \text{ for all } i, j \in N\}$. For such c and $\delta \in (0, 1)$, let f be a density function on X such that

$$f(x) \propto \begin{cases} |c - x|^{-\alpha} & \text{if } x \in Y(\delta), \\ f_L & \text{otherwise.} \end{cases}$$

where $\alpha > 0$ and $f_L > 0$ are positive constants. Let $A(v) = \{x \in X \,|\, \text{there is } C \in \mathcal{C} \text{ such that } x_i \geq v_i \text{ for all } i \in C\}$ be the acceptance set in this voting rule. Since f represents a pointwise distribution at c in the limit as $\alpha \to \infty$ and $f_L \to 0$, the expected payoff profile conditional on agreement $\int_{A(v)} x f(x) dx / \int_{A(v)} f(x) dx$ continuously approaches c as α becomes large and f_L vanishes. Thus one can easily show that there is large $\alpha > 0$ such that $\int_{A(v)} (x_i - v_i) f(x) dx > 0$ for all $v \in Y(\delta)$ and $i \in N$. This implies that for such α, $v_i^{*\prime}(t) > 0$ for all t and all $i \in N$ if $v^*(t) \in Y(\delta)$. Since $v^*(t) \in Y(\delta)$ for sufficiently large t and sufficiently small f_L, we have that for all $\varepsilon > 0$, there exist α, f_L, $\bar{t}$ such that $v^*(t) \in Y(\delta) \cap B^\varepsilon$ for all $t \geq \bar{t}$.

Then the limit expected duration can be computed with an approximated density function $\tilde{f}$ with support $Y(\delta)$ defined as $\tilde{f}(x) \propto f(x)$ in $Y(\delta)$, and $\tilde{f}(x) = 0$ otherwise. Let $\tilde{\mu}$ be the probability measure corresponding to $\tilde{f}$. For $v \in Y(\delta)$, let

$$b_i(v) = \tilde{g}_i(A(v)) - v_i, \qquad d_i(v) = -\frac{\partial \tilde{\mu}(A(v))/\partial v_i}{\tilde{\mu}(A(v))},$$

where $\tilde{g}(Y) = (\tilde{g}_1(Y), \ldots, \tilde{g}_n(Y))$ denotes a barycenter of the set $Y \subseteq \mathbb{R}^n$ with respect to $\tilde{\mu}$.

A similar argument to the proof of Theorem 4 shows that the limit expected duration is $\lim_{t\to\infty} \dfrac{r(t)}{1 + r(t)}$ where

$$r(t) = \sum_{i \in N} d_i(v^*(t)) b_i(v^*(t)).$$

Since it is straightforward to show that $\lim_{t\to\infty} r(t) > 0$, the limit expected duration must be strictly positive.

D.16 Proof of Proposition 16

By equation (A.4), $v_i^*(\frac{t}{\Delta t})$ is a nondecreasing sequence. Since X is bounded and convex, $v^*(\frac{t}{\Delta t})$ converges to a Pareto efficient payoff profile as $\Delta t \to 0$. Let $v^*(\frac{t}{\Delta t})$ be the solution of equation (A.4), and $v^* = \lim_{\Delta t \to 0} v^*(\frac{t}{\Delta t})$ for $t > 0$.

We borrow ideas from the proof of Proposition 5. Let $f_H(\frac{t}{\Delta t}; m) = \sup_{x \in A(v^*(\frac{t}{\Delta_m}; m))} f(x)$, and $f_L(\frac{t}{\Delta_m}; m) = \inf_{x \in A(v^*(\frac{t}{\Delta_m}; m))} f(x)$. With the same notations as in the proof of Proposition 5, for all $i \in N$, all $\eta > 0$, all $t \geq \bar{t}$, and all m,

$$
\begin{aligned}
&\lambda_m \Delta_m f_L\Big(\frac{\bar{t}}{\Delta_m}; m\Big) \int_{\underline{A}\big(v^*(\frac{t}{\Delta_m}; m)\big)} \Big(x_i - v_i^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\, d\mu \\
&\quad \leq v_i^*\Big(\frac{t}{\Delta_m} + 1; m\Big) - v_i^*\Big(\frac{t}{\Delta_m}; m\Big) \leq \lambda_m \Delta_m f_H\Big(\frac{\bar{t}}{\Delta_m}; m\Big) \int_{\overline{A}\big(v^*(\frac{t}{\Delta_m}; m)\big)} \Big(x_i - v_i^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\, d\mu, \\
&\lambda_m \Delta_m f_L\Big(\frac{\bar{t}}{\Delta_m}; m\Big) V\Big(\underline{A}\Big(v^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\Big)\Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta_m}; m)) - \eta}{(n+1)\alpha_i}\Big)\, d\mu \\
&\qquad \leq v_i^*\Big(\frac{t}{\Delta_m} + 1; m\Big) - v_i^*\Big(\frac{t}{\Delta_m}; m\Big) \\
&\qquad\qquad \leq \lambda_m \Delta_m f_H\Big(\frac{\bar{t}}{\Delta_m}; m\Big) V\Big(\overline{A}\Big(v^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\Big)\Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta_m}; m)) + \eta}{(n+1)\alpha_i}\Big)\, d\mu.
\end{aligned}
$$

Therefore we have $\lim_{m \to \infty} \dfrac{v_j^*(\frac{t}{\Delta_m} + 1; m) - v_j^*(\frac{t}{\Delta_m}; m)}{v_i^*(\frac{t}{\Delta_m} + 1; m) - v_i^*(\frac{t}{\Delta_m}; m)} = \dfrac{\alpha_i}{\alpha_j}$. For large m and $t \geq \bar{t}$, we have an approximation

$$
\begin{aligned}
&\Big(v_i^* - v_i^*\Big(\frac{t}{\Delta_m} + 1; m\Big)\Big) - \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta_m}; m\Big)\Big) \\
&\qquad \approx -\lambda_m \Delta_m f\Big(\frac{\bar{t}}{\Delta_m}; m\Big) V\Big(A\Big(v^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\Big)\Big(\frac{\alpha \cdot (v^* - v^*(\frac{t}{\Delta_m}; m))}{(n+1)\alpha_i}\Big)\, d\mu \\
&\qquad \approx -\frac{\lambda_m \Delta_m f(\frac{\bar{t}}{\Delta_m}; m)}{n(n+1)} \Big(\prod_{j \neq i} \frac{\alpha_i}{\alpha_j}\Big) \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t}; m\Big)\Big)^{n+1}.
\end{aligned}
$$

Then we can show that for all $i \in N$ and all $t \geq \bar{t}$,

$$
\lim_{m \to \infty} \alpha_i \Big(v_i^* - v_i^*\Big(\frac{t}{\Delta t}; m\Big)\Big) \cdot (\lambda_m t)^{\frac{1}{n}} = \Big(\frac{n+1}{f(v^*) n^{n+1}} \prod_{j \in N} \alpha_j\Big)^{\frac{1}{n}}.
$$

Here we used the fact that λ_m is large if m is large, to ignore the constant derived from an initial condition. Since by the above equality,

$$
\begin{aligned}
p(t) &\approx f(v^*) V\Big(A\Big(v^*\Big(\frac{t}{\Delta_m}; m\Big)\Big)\Big) \\
&\approx \frac{f(v^*)}{n \lambda_m t} \prod_{i \in N} \left(\frac{n}{\alpha_i}\Big(\frac{n+1}{f(v^*) n^{n+1}} \prod_{j \in N} \alpha_j\Big)^{\frac{1}{n}}\right) = \frac{n+1}{n^2(\lambda_m t)},
\end{aligned}
$$

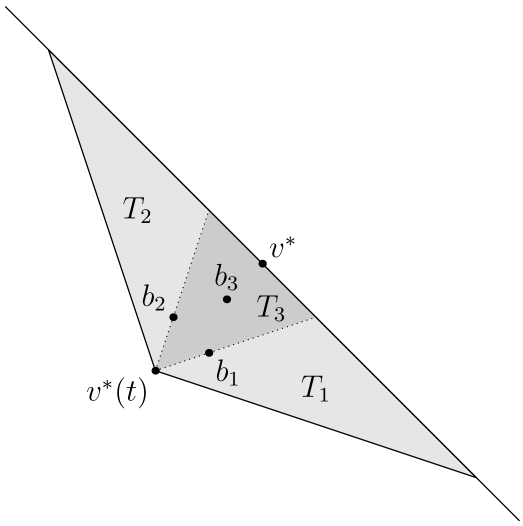

Figure 22: The set of realized allocations that the players accept

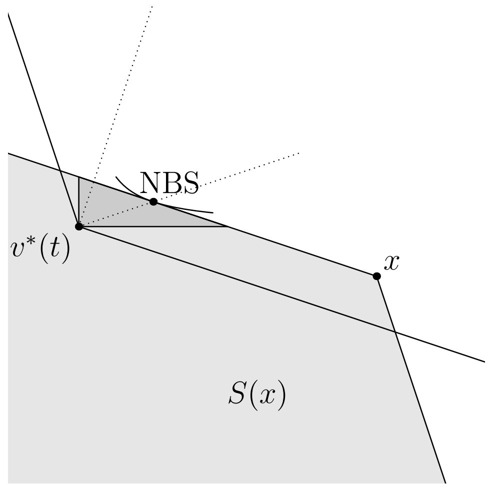

Figure 23: The set of feasible allocations when $x \in T_1$ is realized

we have

$$\lim_{m\to\infty} p(t) \cdot (\lambda_m t) = \frac{n+1}{n^2}.$$

By Lemma 21, the limit expected duration is $\dfrac{n^2}{n^2+n+1}$.

D.17 Proof of Proposition 20

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 $z(t) \le \frac{1-a}{2(1+a)}$. 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 22, where the slopes of four line segments starting from $v^*(t)$ are $-a, a, 1/a$ and $-1/a$, respectively, from southeast to northwest.

Suppose that allocation x belongs to the triangle T_1 in Figure 22. Then the set $S(x)$ of feasible allocations is described in Figure 23. 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 $-a$. A 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 accordingly.

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 and b_3. By ODE (1),

$$\begin{aligned}
z'(t) &= -v_1^{*\prime}(t) \\
&= -\lambda\big((b_1 - v(t))\mu(T_1) + (b_2 - v^*(t))\mu(T_2) + (b_3 - v(t))\mu(T_3)\big) \\
&= -\lambda \cdot \frac{8(1+a^2)}{3(1-a^2)} z(t)^3.
\end{aligned}$$

Since $p(t) = \mu(T_1) + \mu(T_2) + \mu(T_3) = \dfrac{4(1+a)}{1-a} z(t)^2$,

$$p'(t) = \frac{8(1+a)}{1-a} z(t) z'(t) = \lambda \cdot \frac{4(1+a^2)}{3(1+a)^2} p(t)^2.$$

Therefore the constant r defined in Section 4.2.2 is $\frac{4(1+a^2)}{3(1+a)^2}$. By Theorem 4, the limit expected duration is

$$\frac{1}{1 + \left(\frac{4(1+a^2)}{3(1+a)^2}\right)^{-1}} = \frac{4 + 4a^2}{7 + 6a + 7a^2}.$$