

Synchronization under Uncertainty*

Michael Ostrovsky Michael Schwarz
ostrovsky@gsb.stanford.edu mschwarz@berkeley.edu

June 25, 2005.

Abstract

A decision maker needs to schedule several activities that take uncertain time to complete and are only valuable together. Some activities are bound to be finished earlier than others, thus incurring waiting costs. We show how to schedule activities optimally, how to give independent agents performing them incentives that implement the efficient schedule, and how to form teams in the presence of uncertainty. The paper offers insights into important economic decisions such as planning large projects and coordinating product development activities. (JEL D78, L23, O22)

*We thank the Coeditor Andrew McLennan, Christopher Avery, Keith Chen, Katia Epshteyn, Drew Fudenberg, Paul Milgrom, Markus Möbius, Alvin Roth, Martin Weitzman and Muhamet Yildiz for comments and suggestions.

1 Introduction

This paper introduces and studies synchronization problems. In a variety of settings, several agents or components need to come together at a point in time. Often perfect coordination is impossible because some agents are bound to be ready earlier than others, thus incurring a waiting cost. For an example of a synchronization problem, consider hikers meeting at a parking lot. The exact time of the meeting does not matter, but everyone dislikes waiting for others. Define efficiency in the parking lot problem as minimizing the total waiting time. We characterize the optimal solution for this and other synchronization problems. A remarkable property of the optimal solution is that the probability that an agent arrives last is independent of the probability distribution representing the distribution of stochastic disturbances to agents' arrival times. This remains true even when agents face different distributions of stochastic disturbances. Important economic decisions such as planning large projects and coordinating product development activities can be modeled as synchronization problems.

Another example of a synchronization problem is the process of writing a handbook. Each of N contributors submits a survey of his subfield. The handbook is published as soon as all surveys are finished. Each contributor can control the time when he expects to finish his survey. We refer to this time as contributor's target arrival time. The actual time of finishing the survey is a random variable referred to here as the arrival time. We refer to the difference between the actual and the target arrival times as the disturbance term. A contributor can shift the distribution of arrival times to the right by postponing the time when he starts writing the survey.

Note that perfect coordination is not possible, there is always one survey that arrives last. The difference between the meeting time and the arrival time of a particular survey is the waiting time. Waiting is costly because readers value current surveys. We assume that waiting cost is linear in waiting time.

In Section 2.1 we consider socially efficient choices of target arrival times, i.e., the choices of target arrival times that minimize the total waiting time (we assume here that per-period cost of waiting is the same for all surveys). Note that this is purely a synchronization problem. The meeting time is unimportant, only minimization of the sum of waiting times matters for achieving efficiency.¹

Proposition 1 establishes that a necessary and sufficient condition for efficiency is that each contributor has the same probability of being last. This remains true even if contributors face different disturbance terms or if they can alter distributions of disturbance terms.

Proposition 4 establishes that for $N \geq 3$, target arrival times among agents with uniform disturbances are decreasing in variance. Perhaps surprisingly, this claim does not hold for $N = 2$. In fact, Corollary 3 shows that in a synchronization problem with two agents each with different but symmetric distribution of the disturbance terms both contributors will target the same arrival time.

¹There is a somewhat related literature on reliability in operations research. See, for example, Robert Harris (1970) and Doris L. Grosh (1989). The basic framework of reliability problem is the following: a system consisting of N components is useful until one of the components breaks. An engineer chooses the optimal expected lifetime of individual components of the system, given that increasing the expected life of the system is desirable but boosting the mean duration of each component is costly. Thus, there is a similar coordination problem. However, there is an essential difference—reliability literature is primarily concerned with estimating expected product life given limited data on reliability. In contrast, we assume that the distribution of disturbance terms is known and focus on characterizing the efficient strategy and implementing it.

The applications of synchronization problems reach beyond “the handbook coordination problem.” Design and manufacturing of products as dissimilar as cars and buildings can often be viewed as a synchronization problem similar to the above example. The waiting cost may correspond to depreciation of a component or to the cost of capital tied up in a component. Thus, per-period waiting costs are likely to be different for different components. In many instances, the total payoff may be higher if the project is completed sooner rather than later even if all waiting times are unchanged,² we can think of a cost of delay as a forgone profit that a project yields per period upon completion. Consequently, there is a cost of delay separate from the cost of waiting. In Section 2.2 we allow for positive delay costs and different waiting costs for different components. Fortunately, it is straightforward to generalize the results of the “handbook coordination problem” for this model (see Proposition 6).

In Section 3 we describe a natural mechanism that implements the socially efficient outcome when agents internalize their own waiting and delay costs, but not the costs of others. This mechanism implements the socially efficient outcome even if agents can change the variances of their arrival times at a cost. In Section 4 we show that the total waiting cost is submodular in standard deviations of the agents’ arrival times, and discuss implications. Section 5 shows that our stylized model is a valid approximation of a more realistic model with time discounting, provided that agents are sufficiently patient. Section 6 concludes.

²Adding a constant to all arrival times leaves all waiting times unchanged, while increasing delay.

2 The Model

The basic model of a synchronization problem is as following. There are N components (or agents). Target arrival time for component i is denoted by μ_i , and the difference between the actual arrival time (t_i) and the target arrival time (μ_i) is the disturbance term. Disturbances are drawn from continuous probability distribution $F_1 \times \cdots \times F_N$ (the disturbance terms are independent but not necessarily identical). The earliest possible target arrival time for component i is $m_i \geq 0$. The per-period cost of waiting associated with component i is c_i . The per-period cost of delay associated with component i is d_i , and $d \equiv \sum d_i$ is the aggregate per-period cost of delay. The decision maker selects target arrival times. The action space of the decision maker is given by $\{\boldsymbol{\mu} : (\boldsymbol{\mu} = (\mu_1, \dots, \mu_N) \in R^N, \mu_i \geq m_i \text{ for all } i)\}$.

The total payoff is given by $-\left(\sum_{i=1}^N c_i (t_* - t_i) + dt_*\right)$, where $t_* = \max_{i \in [1..N]} \{t_i\}$. The goal is to find the vector of target arrival times that maximizes the total payoff.

To summarize the notation:

N number of components

$\mu_i \geq m_i$ target arrival time of component i

t_i actual arrival time of component i

F_i distribution of $t_i - \mu_i$

c_i per-period waiting cost associated with component i

d_i per-period cost of delay associated with component i

d the total per-period cost of delay, $\equiv \sum d_i$

t_* $\max_{i \in [1..N]} \{t_i\}$, i.e., arrival time of the last component

We begin by characterizing the socially efficient strategy profile—the profile that

maximizes the total payoff in the synchronization problem. There are three reasons for our interest in social efficiency: (1) a synchronization problem is often a single player decision problem. For example, a director of a project may be in a position to set target arrival times for the components of the project. (2) reputational concerns may compel agents to follow an efficient strategy even if it is not an equilibrium of a single period game.³ (3) in a variety of settings a mechanism of transfer payments may be used that leads to an efficient outcome, as we show in Section 3.

In the following section we consider the case where all per-period waiting costs c_i are identical and the cost of delay d is equal to zero. Then in Section 2.2 we relax these assumptions.

2.1 Social Optimum for the Case of Equal Waiting Costs and No Cost of Delay

Suppose all components' (or agents') per-period waiting costs are equal, and the cost of delay is zero. This corresponds to the case where the actual meeting time does not matter. In this case the objective becomes to minimize the total waiting time $\sum_{i=1}^N (t_* - t_i)$. Since there is no cost of delay, this is just a synchronization problem, where adding the same constant to all components' arrival times does not change the value of the objective function. Hence, we can ignore the restriction on the earliest arrival time of each component.

According to Proposition 1, a necessary and sufficient condition for an optimal

³A finite discretization of this game has a product structure and thus satisfies the conditions of Theorem 7.1 in Drew Fudenberg, David K. Levine, and Eric S. Maskin (1994).

strategy is for each agent or component to arrive last with probability $1/N$, *regardless of distributions of disturbances* $\{F_i(\cdot)\}$.

Proposition 1 *If $c_1 = c_2 = \dots = c_N > 0$ and $d = 0$, then for all F_1, \dots, F_N , there exists an N -tuple $\boldsymbol{\mu} = \{\mu_1, \dots, \mu_N\}$ of target arrival times that minimizes the expected total waiting time. For any such minimizing N -tuple the probability of each agent arriving last is $\frac{1}{N}$. Conversely, all N -tuples in which each agent arrives last with probability $\frac{1}{N}$ are optimal. Finally, if the supports of distributions of disturbances are connected, then all optimal N -tuples are identical up to adding the same constant to all target arrival times.*

Proof. Existence is straightforward: the total waiting time is continuous in $\boldsymbol{\mu}$. Without loss of generality we assume $\mu_1 = 0$, then the domain of interest becomes bounded—no μ_i can be greater than the expected total waiting time for $\boldsymbol{\mu} = 0$.

The proof of the second statement of the proposition is not hard either, but it is more interesting. Namely, suppose that for an optimal $\boldsymbol{\mu}$, for some agent (say, $i = 1$), the probability of arriving last is $p < \frac{1}{N}$. By assumption, all probability distributions are continuous, and if we let μ_1 grow arbitrarily large holding other agents' expected arrival times fixed, then the probability of agent 1 arriving last approaches 1. Therefore, we can increase μ_1 by some $\epsilon > 0$ such that the probability of agent 1 arriving last is some value q such that $p < q < \frac{1}{N}$ (of course, other agents' probabilities of arriving last will change also).

Now let $\boldsymbol{\tau}$ be a random variable drawn from $(\mu_1 + F_1) \times \dots \times (\mu_N + F_N)$, and $\boldsymbol{\tau}^\epsilon = \boldsymbol{\tau} + (\epsilon, 0, \dots, 0)$. By construction, the probability that in both $\boldsymbol{\tau}$ and $\boldsymbol{\tau}^\epsilon$ agent 1 arrives last is p , the probability that he arrives last in neither of them is $1 - q$,

and the probability that he arrives last in τ^ϵ but not in τ is $q - p$. In the first case, the total waiting time is increased in τ^ϵ vs. τ by $(N - 1)\epsilon$, in the second case it is increased by $-\epsilon$ (i.e., decreased), and in the last case it is increased by less than $(N - 1)\epsilon$. Hence, by adding ϵ to μ_1 , expected waiting time is increased by less than $p(N - 1)\epsilon - (1 - q)\epsilon + (q - p)(N - 1)\epsilon = (Nq - 1)\epsilon < 0$ since q was chosen to be less than $\frac{1}{N}$. Therefore, by adding ϵ to μ_1 we decreased expected total waiting time, and so $\boldsymbol{\mu}$ was not optimal.

To prove the converse, consider an optimal vector of target arrival times, $\boldsymbol{\mu}$, and another vector, $\boldsymbol{\nu}$, for which the probability of each player arriving last is equal to $\frac{1}{N}$. Since adding the same constant to all elements of a vector of target arrival times does not change either the expected waiting cost or the probabilities of being last, we can assume, without loss of generality, that $\min_{1 \leq i \leq N} \{\nu_i - \mu_i\} = 0$. If $\boldsymbol{\mu} = \boldsymbol{\nu}$, there is nothing left to prove, and we are done. Otherwise, consider the set of agents for whom $\nu_i = \mu_i$. Without loss of generality, we can assume that these agents are $1, \dots, j$ for some $j < N$, and so $\nu_i - \mu_i = 0$ for all $i \leq j$ and $\nu_i - \mu_i > 0$ for all $i > j$. Let $a = \min_{j < i \leq N} \{\nu_i - \mu_i\}$; let $\boldsymbol{\chi} = (0, \dots, 0, 1, \dots, 1)$ —the vector with the first j elements equal to 0 and the remaining $N - j$ elements equal to 1; and let $\boldsymbol{\nu}(x) = \boldsymbol{\nu} + (x - a)\boldsymbol{\chi}$.

Let $p_i(x)$ be the probability that player i is last, and let $W(x)$ be the expected total waiting time, under the vector of target arrival times $\boldsymbol{\nu}(x)$. Note first that for any i and for any $x \in [0, a]$, $p_i(x) = \frac{1}{N}$. That is because for each $i \leq j$ and any $x \in [0, a]$, $p_i(x)$ is at least as high as $p_i(a)$ and at most as high as the probability of player i being last under $\boldsymbol{\mu}$, both of which are equal to $\frac{1}{N}$; and for all $i > j$,

$p_i(x) \leq p_i(a) = \frac{1}{N}$ and so (since their sum is equal to $\frac{N-j}{N}$), $p_i(x)$ is also equal to $\frac{1}{N}$.

Now take any $x \in (0, a)$. Take a small ϵ , either positive or negative, and consider $W(x + \epsilon) - W(x)$. For a particular realization of disturbances, the probability that the last player to arrive under both vectors is one of the first j players is equal to $\frac{j}{N} + O(\epsilon)$, the probability that the last player to arrive under both vectors is one of the remaining $N - j$ players is equal to $\frac{N-j}{N} + O(\epsilon)$, and the probability that different players arrive last under these two vectors is of the order $O(\epsilon)$. Hence, similarly to the analysis in the first part of the proof, $W(x + \epsilon) - W(x) = \epsilon j(\frac{N-j}{N} + O(\epsilon)) - \epsilon(N - j)(\frac{j}{N} + O(\epsilon)) + \epsilon O(\epsilon) = o(\epsilon)$. But this means that for any $x \in (0, a)$, $W'(x) = 0$, which, together with the continuity of W on $[0, a]$, implies that $W(0) = W(a)$, i.e., the expected waiting time under $\nu(0)$ is equal to the expected waiting time under the original vector ν . We can now apply the same procedure to vector $\nu(0)$, and after several iterations we will get to vector μ . Therefore, the expected waiting time is the same under μ and ν , and so ν is an optimal vector.

Finally, suppose the distributions of disturbances have connected supports, and suppose there are two N -tuples, μ and μ' , in which each agent arrives last with probability $\frac{1}{N}$. Suppose not all $(\mu'_i - \mu_i)$ are the same. Let j be an agent with the largest $(\mu'_i - \mu_i)$. Then, relative to him, some agents target the same arrival time in μ' as in μ , and some target an earlier arrival time. But this implies that the probability of agent j arriving last increases strictly as we move from μ to μ' (because the distributions of realized arrival times have connected supports, and overlap for any pair of players), which contradicts our assumption that it is $\frac{1}{N}$ in both cases. ■

Corollary 2 *If distributions F_1, \dots, F_N are identical, then in the optimum all agents*

target the same arrival time. If for some subset $[1..I]$ of agents all F_1, \dots, F_I are identical, then in any optimum agents $1..I$ target the same arrival time.

Corollary 3 *If $N=2$ and distributions F_1, F_2 are symmetric around zero (but not necessarily identical), then it is optimal for both agents to target the same arrival time.*

Corollary 3 is contrary to the intuition that high variance components should target earlier arrival time. It is not hard to reconcile the result of Corollary 3 with this intuition. When $N = 2$, one of the two components will be the last one to arrive. Note that the social cost of arriving early vs. arriving late is symmetric. In contrast, for $N \geq 3$ “arriving late is socially more costly than arriving early,” because at least two components wait for the component that is the last to arrive. Thus, for $N \geq 3$ the conclusion of Corollary 3 does not hold. The following proposition provides a sufficient condition under which components with higher variance of the disturbance term target earlier arrival times.

Proposition 4 *Suppose there are at least three agents, and two of them have zero-mean uniform distributions of disturbances F_1 and F_2 . Suppose $V(F_1) > V(F_2)$. Then in the optimum agent 1’s target arrival time is less than or equal to agent 2’s. If, additionally, we assume that each agent’s distribution of disturbances has connected support, then in the optimum agent 1 targets a strictly earlier time than agent 2.⁴*

⁴The following example shows that the connectedness of other agents’ distributions is a necessary condition for the strict inequality. There are 3 agents. Agent 1 has uniform distribution of disturbances on $[-1, 1]$. Agent 2 has uniform distribution on $[-2, 2]$. Agent 3’s distribution consists of two parts—uniform on $[200, 202]$ with probability $\frac{1}{3}$ and uniform on $[-101, 100]$ with probability $\frac{2}{3}$. In this case it is optimal for all agents to aim at the same time 0.

Proof. See Appendix. ■

The following is a straightforward corollary of Proposition 4.

Corollary 5 *If all agents have zero-mean uniform distributions of disturbances, in the optimum agents with larger variances will target earlier times than agents with smaller variances.*

The assumption that the distributions of disturbances are uniform is important for the conclusion of Proposition 4. There may be other classes of distributions for which this result holds, but, in general, it is not necessarily the case that in the optimum, an agent with a higher variance will target an earlier arrival time than an agent with a lower variance, even if the former's distribution of disturbances is a mean-preserving spread of the latter's, and both are symmetric. The following example illustrates the possibility of the opposite ordering of target arrival times.

Consider the following three distributions of disturbances, F_1 , F_2 , and G . Note that F_2 is a mean-preserving spread of F_1 , and both are symmetric around zero.

- $F_1(x)$: with probability $\frac{1}{12}$, x is equal to -5 ; with probability $\frac{1}{12}$, x is equal to 5 ; and with probability $\frac{5}{6}$, x is distributed uniformly on $[-1, 1]$.
- $F_2(x)$: with probability $\frac{1}{12}$, x is equal to -10 ; with probability $\frac{1}{12}$, x is equal to 10 ; and with probability $\frac{5}{6}$, x is distributed uniformly on $[-2, 2]$.
- $G(x)$: with probability $\frac{28}{41}$, x is equal to -13 ; and with probability $\frac{13}{41}$, x is equal to 28 .

Now consider three agents with these distributions of disturbances (f_1 , f_2 , and g), and suppose they target arrival times 0 (f_1), $\frac{1}{35}$ (f_2), and 12 (g). Then each agent

arrives last with probability $\frac{1}{3}$, and so this is the optimal vector of target arrival times, even though f_2 targets a later arrival time than does f_1 .⁵

2.2 Social Optimum for the Case of Different Waiting Costs and Positive Cost of Delay

In the previous section we assumed that the cost of delay were zero and the per-period waiting costs were the same for all components. Let us now consider the general model, in which the cost of delay is allowed to be positive and the waiting costs are allowed to differ across agents ($c_i > 0$ for all i). Note that when the per-period cost of delay is positive, other things being equal an earlier meeting time is preferable. Consequently, the constraint $\mu_i \geq m_i$ becomes binding for some agents. In the following proposition we show that since per-period costs of waiting are different, agents for whom the above constraint is not binding will arrive last with unequal probabilities proportional to their waiting costs.

Proposition 6 *There exists an N -tuple μ of target arrival times minimizing the expected total waiting cost subject to $\mu_i \geq m_i, \forall i$. For any such minimizing N -tuple the following property is satisfied: for every i , if $\mu_i > m_i$, then the probability of agent i arriving last is equal to $\frac{c_i}{\sum c_j + d}$, and if $\mu_i = m_i$, then the probability of agent i arriving last is greater than or equal to $\frac{c_i}{\sum c_j + d}$. If the distributions of disturbances have connected supports, then for a positive d such minimizing N -tuple is unique and*

⁵We should point out that this example is inconsistent with the assumption of our model that all distributions of disturbances are continuous. Note, however, that if we replace all point masses in distributions F_1 , F_2 , and G with uniform distributions on small intervals around these point masses, then by continuity the optimal vector of target arrival times will be close to $(0, \frac{1}{35}, 12)$.

it is the only one that satisfies the property above; and for $d = 0$, all minimizing N -tuples are identical up to adding a constant and are the only ones satisfying the property.

Proof. See Appendix. ■

As an example, consider design of a new high tech product. A product can go into production as soon as the design of each module is complete. Delay in releasing a product is costly. Finishing a component early is often costly as well, because over time better and better designs become feasible. For example, design of a new car normally takes several years. For many components of a car, waiting costs are negligible, because the technological progress is relatively slow. Little is lost if the design of seats is complete two years earlier than necessary. In contrast, an opportunity to use a more current technology is lost, if the design of an on board computer is finished two years earlier than necessary. In the case of product design we can think of waiting costs as being determined by the cost of components and the rate of technological progress. Consequently, if the rate of technological progress is the same for two components, then per-period waiting costs are proportional to their prices. Proposition 6 predicts that, other things being equal, the probability that a design of a particular component is finished last is proportional to the price of the component.⁶ For the same reason, other things being equal, components based on rapidly evolving technology are more likely to be ready last.

⁶Of course, this statement only applies to components for which target arrival times do not correspond to a corner solution. There is a simple test that identifies components for which target arrival time is not at a corner solution. The target arrival time does not correspond to a corner solution for any component for which development starts later than the earliest possible time (i.e., the time when product development starts).

3 Implementation of the Socially Optimal Outcome

If each agent internalizes his own costs, but not those of others, the self-interest of agents will fail to generate the efficient choice of target arrival times, unless the costs of delay are so large relative to the costs of waiting that all agents target $\mu_i = m_i$ (See Ostrovsky and Schwarz, 2005, for a detailed treatment of individual incentives of agents and resulting equilibria).⁷ To achieve the optimum, agents must have external incentives. In this section we present a natural mechanism that implements optimal target arrival times. To assure that the optimal vector of target arrival times is unique, we assume for the rest of the section that each player's distribution of disturbances has full support on some interval.

The mechanism is analogous to the Vickrey-Clark-Groves pivotal mechanism: only the player who arrives last makes a non-zero transfer payment, which is equal to the externality he imposes on others. It is, however, not the same as VCG, since the latter implements the efficient outcome in a setting where players have private information and need incentives to reveal it truthfully, whereas in our setting all information is public, and players need incentives to choose socially optimal actions. Proposition 7 shows that the game induced by this mechanism has a unique pure strategy equilibrium. This equilibrium implements efficient target arrival times.

Namely, define game Γ' as follows. There are N players. Each player's action is

⁷By Proposition 6, in the social optimum, unconstrained agent i arrives last with probability $\frac{c_i}{\sum c_j + d}$. Using similar logic, his individual optimization implies that he arrives last with probability $\frac{c_i}{c_i + d_i}$. For these two probabilities to be equal, waiting and delay costs of all other agents have to equal 0, which of course does not hold generally.

his target arrival time $\mu_i \geq m_i$. The payoff of player i is given by the expected value of $-c_i(t_* - t_i) - d_it_* + \gamma_i$, where $t_* = \max_{i \in [1..N]} \{t_i\}$, vector \mathbf{t} is equal to $\boldsymbol{\mu}$ plus a random vector of disturbances drawn from continuous probability distribution $F_1 \times \dots \times F_N$, and γ_i is the transfer to agent i within the mechanism. When player i is last, $\gamma_i = -w \left(\sum_{j \neq i} (c_j + d_j) \right)$, where w is the difference between his arrival time and the next-to-last player's arrival time; when player i is not last, $\gamma_i = 0$. It is easy to see that Γ' has one pure strategy equilibrium, and this equilibrium is the socially optimal vector of target arrival times.⁸

Proposition 7 *For a synchronization problem in which c_i and d_i are positive⁹, the corresponding game Γ' has exactly one equilibrium with finite payoffs. In this equilibrium, each player chooses a socially optimal target arrival time.*

Proof. See Appendix. ■

The mechanism imposes a fine on the last agent to arrive. Let $\varphi_i(t_i)$ be the expected fine paid by agent i if he arrives at time t_i , i.e.,

$$\varphi_i(t_i) = E\left[\left\{w \left(\sum_{j \neq i} (c_j + d_j) \right) \text{ if player } i \text{ is last, } 0 \text{ otherwise}\right\} | t_i\right],$$

where w is the difference between t_i and the next-to-last player's arrival time, and

⁸For the sake of completeness, we should mention that this mechanism also has a set of equilibria in mixed strategies. It is easy to see that a strategy profile is a mixed strategy Nash Equilibrium if and only if at least two agents randomize among target arrival time in such a way that their expected arrival times are equal to infinity. In any mixed strategy equilibrium, the expected payoff of each agent is negative infinity.

⁹If delay cost is zero, game Γ' has multiple pure strategy equilibria that implement the socially optimal outcome. Adding a constant to all target arrival times in one socially optimal Nash Equilibrium generates another socially optimal Nash Equilibrium profile.

the expectation is taken assuming socially optimal behavior by other players.¹⁰ Then the expected penalty for being last increases faster than linearly in time.

Proposition 8 *The expected fine $\varphi_i(t_i)$ in game Γ' is convex, assuming equilibrium behavior by the players.*

Proof. For a small ϵ , the difference $(\varphi_i(t_i + \epsilon) - \varphi_i(t_i))$ is equal to (probability that player i is last if he arrives at $t_i + \epsilon$) $\times \left(\sum_{j \neq i} (c_j + d_j)\right) \times \epsilon + o(\epsilon)$. Hence, $(\varphi'_i(t_i)) =$ (probability that player i is last if he arrives at t_i) $\times \left(\sum_{j \neq i} (c_j + d_j)\right)$, which is increasing in time. ■

The mechanism implements socially optimal actions by the agents even if they can control higher moments of the distribution of noise, not just the expected arrival time. In particular, suppose each agent i chooses both target arrival time μ_i and standard deviation σ_i of his distribution of disturbances $F_i(\cdot)$.¹¹ His payoff is $-c_i(t_* - t_i) - d_i t_* - g(\sigma_i)$, where t_* is the latest arrival time of all agents, t_i is agent i 's arrival time, and $g(\sigma_i)$ is the cost of variance of agent i . We assume that function $g(\sigma)$ is differentiable and $g'(\sigma) < 0$. The social planner's objective is to choose vectors μ and σ so as to maximize the expected sum of the agents' payoffs. Notice that Proposition 1 still holds. Let Γ^2 be the game induced by the mechanism in this setting, and assume all c_i and d_i are positive.

Proposition 9 *In game Γ^2 , it is an equilibrium for each worker to choose socially optimal variance and target arrival time.*

¹⁰Note that charging each agent a fine equal to $\varphi_i(t_i)$, instead of the actual externality he imposes on others given all realizations of disturbances, would also implement the social optimum.

¹¹We assume that changing the standard deviation is equivalent to "stretching" the distribution along the time dimension: $F_\sigma(x) = F_1(x/\sigma)$ and that $E[F] = 0$. Zero-mean normal and uniform families of distributions satisfy these requirements.

Proof. Take player i . By construction, if all other players choose socially optimal target arrival times and variances, player i 's individual optimization is collinear with social welfare optimization, since the mechanism forces him to internalize all externalities he imposes on others. Therefore, his best response is the socially optimal target arrival time and variance, and therefore these socially optimal actions form an equilibrium in the game Γ^2 . ■

It is also possible to implement efficient target arrival times with budget-balanced mechanisms, in which the sum of transfers to all agents is equal to zero in every state of the world. The following mechanism is budget-balanced, and also has the property that only the last agent to arrive has to pay. The rest of the agents receive non-negative transfers.

When there are no costs of delay, i.e., $\forall i, d_i = 0$, this mechanism is particularly simple: the player who arrives last pays other players half of their waiting costs. When costs of delay are positive, the payments are more complicated: the player who arrives last has to pay every player j an amount equal to $w_j r_j + t_* s_j$, where t_* is the time of arrival of the last player, $w_j = t_* - t_j$, $r_j = \frac{c_j}{2}$, and $s_j = d_j - \frac{\sum d_k}{2N-2}$.

Proposition 10 *With positive costs of waiting and delay, the game induced by the above mechanism has a unique equilibrium with finite payoffs, and this equilibrium is socially optimal.*

Proof. Let q_i be player i 's desired probability of being last in this setup. We need to show that this q_i is the same as in the optimum, i.e., is equal to $c_i/(\sum(c_j + d_j))$. Analyzing player i 's FOC we get $q_i(d_i + \sum r_j - r_i + \sum s_j - s_i) = (1 - q_i)(c_i - r_i) \Leftrightarrow q_i(d_i + \sum \frac{c_j}{2} - \frac{c_i}{2} + \frac{N-2}{2N-2} \sum d_j - d_i + \frac{1}{2N-2} \sum d_j) = (1 - q_i)(\frac{c_i}{2}) \Leftrightarrow q_i(\sum c_j - c_i + \sum d_j) =$

$(1 - q_i)c_i \Leftrightarrow q_i(\sum c_j + \sum d_j) = c_i$. Thus, in any equilibrium, for players who do not target the earliest possible arrival times the FOC holds, whereas for players who do target the earliest possible arrival times the probability of being last is higher than $c_i/(\sum(c_j + d_j))$. But then, as shown in the proof of Proposition 7, players can not mix in any equilibrium with finite payoffs, and there can only exist one vector of equilibrium arrival times. ■

4 Forming Teams of Agents with Different Distributions of Disturbances

We have shown how a decision maker should schedule different agents' target arrival times, and how he can give them incentives to stick to that schedule. Another task that planners often face is how to put agents in teams to work on several different projects. In this section we show that expected waiting time is submodular in standard deviations of agents' arrival times, and so it is better to put agents with similar variances together, rather than mix. For manufacture of complex products consisting of many subcomponents, this result is similar in spirit to Michael Kremer's (1993) O-Ring model, although along a different dimension. The O-Ring model of production function describes a process consisting of several tasks; workers implementing these tasks sometimes make mistakes. Higher probability of mistakes by a worker decreases productivity of other workers by reducing the value of the composite good. As a result, it is optimal to put high-quality (low probability of mistakes) workers with other high-quality workers.

Another important economic implication is that reducing uncertainty in one part of the production process makes reducing uncertainty in another part more profitable. This is parallel to Paul Milgrom and John Roberts (1990), who list several variables which determine a firm's profitability, and find complementarities among them by showing that the profit function is supermodular in these variables. Since waiting cost is submodular in standard errors of agents' arrival times, the profit function is supermodular (we can model the profit function as a constant minus the total cost of waiting).

Formally, suppose there are two agents, $i = 1, 2$. Their distributions of disturbances belong to a family of zero-mean distributions parameterized by standard deviation, i.e., $F_\sigma(x) = F_1(x/\sigma)$. Also, for simplicity, assume that agents' costs of delay are equal to 0 and costs of waiting are equal to 1. Let $w(\sigma_1, \sigma_2)$ be the expected total waiting time when agents have standard deviations σ_1, σ_2 and choose their expected arrival times optimally.

Proposition 11 *For uniform and normal families of distributions of disturbances, reducing σ_1 is complementary to reducing σ_2 , i.e., function $w(\sigma_1, \sigma_2)$ is submodular.*

Proof. See Appendix. ■

Now suppose a manager has four workers, A, B, C , and D , and needs to form two 2-person teams. The workers are identical, except for the variances of the times it takes them to complete a task: a, b, c , and d , respectively. Without loss of generality, $a \leq b \leq c \leq d$. Assume also that either all distributions of disturbances are uniform, or all are normal. Then it is optimal for the manager to group workers with similar variances together.

Corollary 12 *Among all possible divisions of the four workers into two groups of two, putting together A with B and C with D will result, in expectation, in the lowest total waiting cost.*

Proof. Consider any other partition of the four workers into groups of two, (X, Y) and (Z, T) , with variances x, y, z , and t . Without loss of generality, assume that $x \leq y$ and $z \geq t$. By Proposition 11, function w is submodular, and so $w(x, y) + w(z, t) \geq w((x, y) \wedge (z, t)) + w((x, y) \vee (z, t)) = w(a, b) + w(c, d)$. ■

5 Robustness to Time Discounting

In the previous sections we considered stylized models where agents were infinitely patient, i.e., their time discount coefficient β was equal to 1. This allowed us to derive simple and intuitive results. The current section shows that our results are robust to small changes in β , i.e., for a sequence of discount factors converging to 1, the optimal arrival times converge to the ones derived in Section 2. Roughly, this also says that if disturbances are small, a planner can ignore time discounting.

Define synchronization problem $\Gamma(\beta)$ for $0 < \beta < 1$ as follows. There are N agents. Each agent i can choose his target arrival time $\mu_i \geq m_i$. His actual arrival time t_i is equal to μ_i plus a random disturbance drawn from continuous probability distribution F_i independent of other agents' disturbances. Once he arrives, he has to pay his waiting cost c_i until the last agent arrives. When all agents arrive, agent i starts getting benefit d_i forever. (An equivalent way to think about this model is to say that once the agent arrives, he has to pay c_i forever, and when all agents arrive,

he starts getting gross benefit $(c_i + d_i)$, also forever. The total instantaneous cost of delay for the group, d , is equal to $\sum d_i$.) The agent's payoff $\Pi_i = E[\int_{t_i}^{\infty} \beta^t \pi_i(t) dt]$, where $\pi_i(t)$ is the sum of cost and benefit received at time t , and the total payoff is the sum of all agents' payoffs.

Define $\Gamma(\beta = 1)$ as the general synchronization problem from Section 2, with no discounting. The following proposition says that as $\beta \rightarrow 1$, the strategy profile that maximizes the total payoff in $\Gamma(\beta)$ goes to the strategy profile that maximizes the total payoff in $\Gamma(1)$. Also, it says that as β increases, optimal target arrival times decrease, i.e., as agents become more patient, it is optimal for them to arrive earlier.

Proposition 13 *Consider decision problem $\Gamma(\beta)$ and let $\mu^*(\beta)$ be the vector of target arrival times maximizing the total expected payoff. Then (i) as $\beta \rightarrow 1$, $\mu^*(\beta) \rightarrow \mu^*(1)$ and (ii) for any $0 < \beta_1 \leq \beta_2 \leq 1$, $\mu^*(\beta_1) \geq \mu^*(\beta_2)$.*

Proof. See Appendix. ■

6 Concluding Remarks

This paper introduces and studies “synchronization problems.” In a variety of settings, several agents or components need to come together at a point in time. The actual time of arrival of each component is its target arrival time plus a noise term. In this case perfect coordination is impossible, because some agents are bound to be ready earlier than others, thus incurring waiting and delay costs. We show that each agent's optimal probability of being last is independent of the probability density

function of the noise term. We provide a simple operable formula for comparing optimal probabilities of being late for components with different costs of waiting.

Our model readily lends itself to empirical testing. An econometrician only needs to observe the frequency with which each component in a large project is finished last, the cost of waiting for each component, and the cost of delay. He does not need to estimate the distribution of disturbances, which would presumably be very hard. Neither does he need to observe the components' planned completion dates. Alternatively, if in a particular application it is reasonable to assume that arrival times are chosen optimally, the model can be used to estimate relative magnitudes of delay and waiting costs of different components in a project.

The present paper investigates synchronization problems in a cooperative setting, and so our results are directly applicable to synchronization problems that emerge within a firm, such as management of large projects or coordination of product development activities. When coordination of actions of different firms is concerned, non-cooperative models of synchronization may be more appropriate (if "arrival times" are not contractible and the need for coordination is not recurrent, agents will act in their own interests, which generally leads to inefficient outcomes). The analysis of such non-cooperative synchronization problems can be found in Anderson, Goeree, and Holt (2001) and Ostrovsky and Schwarz (2005).

A Proof of Proposition 4

First, let us prove the following lemma.¹²

Lemma 14 *Let X_1 and X_2 be independent random variables distributed uniformly on positive-length intervals $[a_1, b_1]$ and $[a_2, b_2]$, respectively, and suppose $E[X_1] \geq E[X_2]$ and $V[X_1] \geq V[X_2]$. Take any point z . Then $P(X_1 = \max\{X_1, X_2, z\}) \geq P(X_2 = \max\{X_1, X_2, z\})$, and the inequality becomes strict if $E[X_1] > E[X_2]$.*

Proof. Without loss of generality assume $z = 0$. Our conditions on means and variances of X_1 and X_2 can be reformulated in equivalent terms as $a_1 + b_1 \geq a_2 + b_2$ and $b_1 - a_1 \geq b_2 - a_2$. This in turn implies that $b_1 > b_2$ (if $b_1 = b_2$ then we necessarily have $a_1 = a_2$ and the claim becomes obvious). If b_2 is negative or equal to 0, the claim of the lemma is clearly true. If at least one of a_1 or a_2 is non-negative, then 0 is never equal to $\max\{X_1, X_2, 0\}$, and the claim of the lemma is also true since it becomes equivalent to $P(X_1 > X_2) \geq (>) \frac{1}{2} \Leftrightarrow E[X_1] \geq (>) E[X_2]$.

So the only interesting case left to prove is when a_1 and a_2 are negative and b_1 and b_2 are positive. Without loss of generality, we can assume that $a_1 + b_1 = a_2 + b_2$, since the general result follows once one shows it with X_1 replaced by $X_1 - (b_1 - a_1 - b_2 + a_2)/2$. Then $a_1 \leq a_2 < z = 0 < b_2 \leq b_1$ and $b_1 - b_2 = a_2 - a_1$. For $i = 1, 2$, let $E_i = \{X_i = \max\{X_1, X_2, 0\}\}$ and $F_i = E_i \cap \{a_2 \leq X_1 \leq b_2\}$. Note that the transformation $(x_1, x_2) \mapsto (x_2, x_1)$ is a measure-preserving bijection between F_1 and F_2 , so $P(F_1) = P(F_2)$. Now, $P(E_1) = P(F_1) + P(X_1 > b_2)$ and $P(E_2) =$

¹²We are grateful to Andrew McLennan for the simplified proof of the lemma.

$P(F_2) + P(X_1 < a_2 \text{ and } X_2 > 0)$, so it suffices to observe that

$$P(X_1 > b_2) = \frac{b_1 - b_2}{b_1 - a_1} \geq \frac{a_2 - a_1}{b_1 - a_1} \frac{b_2 - a_0}{b_2 - a_2} = P(X_1 < a_2 \text{ and } X_2 > 0).$$

■

Let us now prove Proposition 4.

Proof. We know from Proposition 1 that in the optimum all agents should arrive last with equal probabilities. Let G be the probability distribution of the latest arrival time of agents $3, \dots, N$ ($x = \max\{t_3, t_4, \dots, t_N\}$, $x \sim G(\cdot)$). Then $P_i = P(\text{agent } i \text{ is last}) = \int P(\text{agent } i \text{ is last} \mid \text{latest arrival time of agents } 3, \dots, N = x) dG(x)$, where $i \in \{1, 2\}$.

If $E(F_1) > E(F_2)$, then from the previous lemma we know that $\forall x$, $P(\text{agent 1 is last} \mid x) > P(\text{agent 2 is last} \mid x)$. But then $P_1 > P_2$, which contradicts our assumption of all agents arriving last with equal probabilities, and so $E(F_1) \leq E(F_2)$.

Now suppose that the supports of other agents' distributions are connected, and suppose $E(F_1) = E(F_2)$. Since uniform distributions are also connected, this and the fact that every agent can arrive last with positive probability implies that there exists an interval such that each agent's probability density function is positive on this interval. Consequently, G' is positive on this interval. Take points a, b inside the interval.

$$(A.1) \quad P_1 - P_2 \geq \int_a^b (P(\text{agent 1 is last} \mid x) - P(\text{agent 2 is last} \mid x)) dG > 0,$$

which contradicts our assumptions, and so $E(F_1) < E(F_2)$. ■

B Proof of Proposition 6

The proof is completely analogous to the proof of Proposition 1, so we'll be brief. An optimal $\boldsymbol{\mu}$ exists because the total cost is continuous in $\boldsymbol{\mu}$ and the “relevant” range is compact. Take an optimal $\boldsymbol{\mu}$. If μ_i is greater than m_i , then the idea that a small disturbance doesn't change the expected total cost (up to the first order) gives us $q_i(\sum_{j \neq i} c_j + d) = (1 - q_i)c_i$, where q_i is the probability that agent i arrives last. If $\mu_i = m_i$, small increase in μ_i can not decrease the total cost, and so $q_i(\sum_{j \neq i} c_j + d) \geq (1 - q_i)c_i$. To show that there is no more than one $\boldsymbol{\mu}$ satisfying this condition, suppose there are two, $\boldsymbol{\mu}^1$ and $\boldsymbol{\mu}^2$. Take i such that $\mu_i^1 > \mu_i^2$ and $\mu_i^1 - \mu_i^2$ is the biggest such increase (if such i does not exist, there exists one such that $\mu_i^2 > \mu_i^1$, which case can be dealt with completely analogously). Then at $\boldsymbol{\mu}^1$ the probability of agent i being last is strictly greater than the bound given in proposition, and also μ_i^1 is greater than m_i —contradiction.

C Proof of Proposition 7

First, let us show that players cannot mix in an equilibrium with finite payoffs. Indeed, suppose player i is indifferent between target arrival times t and $t' > t$; both are best responses to other players' (possibly mixed) strategies. From the player's first order conditions, it has to be the case that he is the last one to arrive with probability $\frac{c_i}{\sum(c_j + d_j)}$ at both target arrival times (this probability cannot be greater

than that at t' , which is strictly greater than the earliest possible target time m_i , and therefore also cannot be greater than that at t , because the probability of being last while targeting arrival time t' is at least as high as the probability of being last while targeting arrival time t).

By assumption, all distributions of disturbances have full supports on some intervals. Suppose the distribution of disturbances of player i , F_i , has support on some interval $[a, b]$. Since the probability of player i being last is the same whether he targets t or t' , it has to be the case that the probability of the last one of the remaining players arriving between $t + a$ and $t' + b$ is equal to zero. But other players also have distributions of disturbances with connected supports, and so for each pure strategy played by any other player (say, player j) in this equilibrium, that player is either guaranteed to arrive before player i or guaranteed to arrive after player i . But in the former case player j , playing that strategy, arrives last with probability zero, which violates his first-order condition, since by assumption c_j is positive. Hence, for all of the strategies that player j plays in this equilibrium, he is guaranteed to arrive after player i , which in turn implies that player i arrives last with probability zero, which contradicts his own first-order condition.

Let us now focus on pure-strategy equilibria. By construction, socially optimal arrival times form an equilibrium. Now, suppose there are two different equilibrium vectors of target arrival times, $\boldsymbol{\mu}^1$ and $\boldsymbol{\mu}^2$. By construction, the system of first order conditions for the players' choices of target arrival times is the same as in the optimum, and the probabilities of being last are also the same (for the unconstrained players). Take player i for whom $\mu_i^1 - \mu_i^2$ is the highest among all players and is

positive (this is w.l.o.g.—if necessary, consider $\mu_i^2 - \mu_i^1$ instead). Then (a) he is unconstrained in vector $\boldsymbol{\mu}^1$ and (b) his probability of being last under $\boldsymbol{\mu}^1$ is greater than under $\boldsymbol{\mu}^2$, and is therefore strictly greater than $\frac{c_i}{\sum(c_j+d_j)}$. This contradicts his first order condition, which has to bind for the unconstrained players.

It is interesting to note that the above proof can be reinterpreted as saying that the potential of the modified game, $E[-\sum_i c_i(t_* - t_i) - \sum_i d_i t_*]$, has a unique local maximum, and also to note that the mechanism does not depend on the form of noise. Online appendix to Ostrovsky and Schwarz (2005) shows that the loss of potential structure in synchronization games is related to the loss of robustness to noise. The connection between potentiality and robustness to noise is also explored in Blume (2003) and Frankel, Morris, and Pauzner (2003).

We should also note that there are many mixed equilibria in this game; in all of these equilibria players get infinite negative payoffs. For example, it is an equilibrium for each player to come at time 2 with probability 1/2, 4 with probability 1/4, etc. Restricting strategy spaces to large but finite intervals of target arrival times would bring the model in line with typical potential games, and would restore equilibrium uniqueness.

D Proof of Proposition 11

D.A Normal Disturbances

Suppose we have agents with normal distributions of disturbances with standard deviations a and b . From Corollary 3 we know that in the optimum in each group

both agents target the same arrival time. Hence, the difference between their arrival times is distributed normally with variance $a^2 + b^2$. The expected waiting time is just the expectation of the absolute value of the difference between the arrival times, which is equal to a constant multiplied by the standard deviation of the distribution, i.e., $c\sqrt{a^2 + b^2}$. So, to show that waiting time is submodular, we need to show that the function $\sqrt{a^2 + b^2}$ has a non-positive 2nd cross-derivative (Milgrom and Roberts 1990, Theorem 2). This derivative is equal to $-ab(a^2 + b^2)^{-3/2} \leq 0$.

D.B Uniform Disturbances

In the uniform case, it is easier to parameterize the distribution by the support size rather than the standard error; notice that it's just a proportional rescaling. Suppose the agents' distributions have supports of sizes a and b , $a \leq b$. $w(a, b) = \frac{a}{b} \times \frac{a}{3} + \frac{b-a}{b} \times \frac{a+b}{4} = \frac{1}{12} \times \frac{a^2}{b} + \frac{b}{4}$. So, $w''_{ab}(a, b) = -\frac{1}{6} \frac{a}{b^2} \leq 0$.

E Proof of Proposition 13

For simplicity, assume that distributions of disturbances F_i are bounded.

(i) Suppose $\beta < 1$. By the same “marginal delay” reasoning as in Proposition 1, agent i 's FOC for choosing μ_i when we optimize social welfare is

$$(E.1) \quad c_i \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(i \neq \text{last} | t_i) f_i(t_i - \mu_i) dt_i = \left(\sum_{j \neq i} c_j + \sum d_j \right) \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(i = \text{last} | t_i) f_i(t_i - \mu_i) dt_i$$

if $\mu_i > m_i$ and

$$(E.2) \quad c_i \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(i \neq \text{last}|t_i) f_i(t_i - \mu_i) dt_i \leq \\ \left(\sum_{j \neq i} c_j + \sum d_j \right) \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(i = \text{last}|t_i) f_i(t_i - \mu_i) dt_i$$

if $\mu_i = m_i$.

By adding $c_i \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(i = \text{last}|t_i) f_i(t_i - \mu_i) dt_i$ to both sides, we get the equivalent FOC

$$(E.3) \quad c_i \int_{-\infty}^{\infty} \beta^{t_i} f_i(t_i - \mu_i) dt_i = \\ \sum (c_j + d_j) \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(t_i \geq t_j \forall j | t_i) f_i(t_i - \mu_i) dt_i$$

if $\mu_i > m_i$ and

$$(E.4) \quad c_i \int_{-\infty}^{\infty} \beta^{t_i} f_i(t_i - \mu_i) dt_i \leq \\ \sum (c_j + d_j) \int_{-\infty}^{\infty} \beta^{t_i} \text{Prob}(t_i \geq t_j \forall j | t_i) f_i(t_i - \mu_i) dt_i$$

if $\mu_i = m_i$.

Crucially, both sides are continuous in β and for $\beta = 1$ become (we only write the equation for $\mu_i > m_i$)

$$(E.5) \quad c_i \int_{-\infty}^{\infty} f_i(t_i - \mu_i) dt_i = \sum (c_j + d_j) \int_{-\infty}^{\infty} Prob(t_i \geq t_j \forall j | t_i) f_i(t_i - \mu_i) dt_i,$$

and so

$$(E.6) \quad \frac{c_i}{\sum (c_j + d_j)} = Prob(i = last),$$

which is the FOC for the social optimum with no discounting.

Now suppose $\boldsymbol{\mu}^*(\beta)$ does not go to $\boldsymbol{\mu}^*(1)$ as β goes to 1. Then there exists a subsequence $\{\beta^n\}$ converging to 1 such that $\boldsymbol{\mu}^*(\beta^n)$ converges to some $\tilde{\boldsymbol{\mu}} \neq \boldsymbol{\mu}^*(1)$ (set of $\boldsymbol{\mu}^*(\beta)$ is bounded as $\beta \rightarrow 1$). Then by continuity, $\tilde{\boldsymbol{\mu}}$ satisfies the FOC with $\beta = 1$ and is therefore an optimum of decision problem $\Gamma(1)$. But we know that $\Gamma(1)$ has only one optimum, equal to $\boldsymbol{\mu}^*(1)$.

(ii) Take $\beta_1 < \beta_2$, and suppose for some i , $\mu_1 = \mu_i^*(\beta_1) < \mu_2 = \mu_i^*(\beta_2)$. Without loss of generality, assume that $i = \arg \max_j \{\mu_j^*(\beta_2) - \mu_j^*(\beta_1)\}$. By FOC,

$$(E.7) \quad \sum (c_j + d_j) \int \beta_1^{t_i} Prob(t_i = last) f(t_i - \mu_1) dt_i \geq c_i \int \beta_1^{t_i} f(t_i - \mu_1) dt_i.$$

Since $\mu_1 < \mu_2$,

$$(E.8) \quad \sum (c_j + d_j) \int \beta_1^{t_i} Prob(t_i = last) f(t_i - \mu_2) dt_i > c_i \int \beta_1^{t_i} f(t_i - \mu_2) dt_i.$$

$$(E.9) \quad \int \beta_1^{t_i} (\sum (c_j + d_j) Prob(t_i = last) - c_i) f(t_i - \mu_2) dt_i > 0.$$

Let t_i^* be such that $\sum (c_j + d_j) Prob(t_i^* = last) - c_i = 0$. The integrand is negative for $t_i < t_i^*$ and positive for $t_i > t_i^*$. $\beta_2 > \beta_1$, and so $\left(\frac{\beta_2}{\beta_1}\right)^t$ is an increasing function.

Therefore,

$$(E.10) \quad \int \beta_2^{t_i} (\sum (c_j + d_j) Prob(t_i = last) - c_i) f(t_i - \mu_2) dt_i \geq \left(\frac{\beta_2}{\beta_1}\right)^{t_i^*} \int \beta_1^{t_i} (\sum (c_j + d_j) Prob(t_i = last) - c_i) f(t_i - \mu_2) dt_i > 0.$$

But this, together with $\mu_2 > \mu_1 \geq m_i$, is a violation of the FOC for the optimum.

References

- [1] Anderson, S.P., Goeree, J.K., and Holt, C.A. "Minimum-Effort Coordination Games." *Games and Economic Behavior*, 2001, 34, pp. 177-199.
- [2] Blume, L.E. "How noise matters," *Games and Economic Behavior*, 2003, 44(2):251-271.
- [3] Frankel, D., Morris, S., and Pauzner, A. "Equilibrium Selection in Global Games with Strategic Complementarities," *Journal of Economic Theory*, 2003, 108, 1-44.
- [4] Fudenberg, D., Levine, D.K., and Maskin, E.S. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 1994, 62, pp. 997-1039.
- [5] Grosh, Doris Lloyd. *A Primer of Reliability Theory*. Wiley, 1989.

- [6] Harris, Robert. "An Application of Extreme Value Theory to Reliability Theory." *Annals of Mathematical Statistics*, 1970, 41(5), pp. 1456-65.
- [7] Kremer, Michael. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics*, 1993, 108(3), pp. 551-75.
- [8] Milgrom, Paul and Roberts, John. "The Economics of Modern Manufacturing: Technology, Strategy, and Organization." *American Economic Review*, 1990, 80(3), pp. 511-28.
- [9] Ostrovsky, Michael and Schwarz, Michael. "Adoption of Standards under Uncertainty," 2005, *RAND Journal of Economics*, forthcoming.