



HIER

Harvard Institute of Economic Research

Discussion Paper Number 1969

Coordination Under Uncertainty

by

*Michael Ostrovsky
and
Michael Schwarz*

August 2002

Harvard University
Cambridge, Massachusetts

This paper can be downloaded without charge from the:

<http://post.economics.harvard.edu/hier/2002papers/2002list.html>

Coordination under Uncertainty

Michael Ostrovsky
Harvard University
mostrovsky@hbs.edu

Michael Schwarz*
Harvard University
mschwarz@fas.harvard.edu

June 5, 2001. Revised: August 12, 2002.

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, how to form teams, and how to optimally reduce uncertainty when it is possible to do so at a cost. The paper offers insights into important economic decisions such as planning large projects and coordinating product development activities.

*We thank 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

¹There is a somewhat related literature on reliability in operations research. See, for

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.

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 mechanisms that implement the socially efficient outcome when agents internalize their own waiting and delay costs, but not the costs of others. In Section 4 we show that the total waiting cost is submodular in standard deviations of the agents’ arrival times, and

example, Harris (1970) and 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.

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

discuss implications. Section 5 considers a model where agents can change the standard deviations of their arrival times at a cost. Section 6 shows that our stylized model is a valid approximation of a more realistic model with time discounting, provided that agents are sufficiently patient. Section 7 concludes.

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 $F_1 \times \dots \times F_N$ (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 $\{\mu : \mu = (\mu_1, \dots, \mu_N) \in R^N, \mu_i \geq m_i \text{ for all } i\}$.

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

Proposition 1 *There exists an N -tuple $\mu = \{\mu_1, \dots, \mu_N\}$ 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 identical up to adding the same constant to all target arrival times, and are therefore optimal.*

Proof. Existence is straightforward: the total waiting time is continuous in μ . 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 $\mu = 0$.

The proof of the second statement of the proposition is not hard either, but it is more interesting. Namely, suppose that for an optimal μ , for some

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

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 a small $\epsilon > 0$ 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 τ be a random variable drawn from $(\mu_1 + F_1) \times \cdots \times (\mu_N + F_N)$, and $\tau^\epsilon = \tau + (\epsilon, 0, \dots, 0)$. By construction, probability that in both τ and τ^ϵ agent 1 arrives last is p , probability that he arrives last in neither of them is $1 - q$, and 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 μ was not optimal.

To prove the converse, suppose that 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 as we move from μ to μ' , 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 is continuous and has connected support, then in the optimum agent 1 targets a strictly earlier time than agent 2.*⁴

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 aim at earlier times than agents with smaller variances.*

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 is zero and the per period waiting costs are the same for all components. Let us now consider the general model, where 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

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.

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 d is positive, then such minimizing N -tuple is unique and it is the only one that satisfies the property above; if $d = 0$, then 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. For an obvious reason, 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 the price of a component. 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.⁵ For the same reason, other things being equal, components based on rapidly evolving technology are more likely to be ready last.

⁵Of 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 Mechanisms for Implementing the Socially Optimal Outcome

If each agent internalizes his own waiting and delay costs, but not those of others, the self-interest of agents will fail to generate the efficient choice of target arrival times (unless the agents' costs of delay are so large relative to their costs of waiting that they all target $\mu_i = m_i$).⁶ To achieve the optimum, agents must have external incentives. In this section we present mechanisms which implement the optimal target arrival times.

It is easy to show by using Proposition 6 that Vickrey-Clark-Groves pivotal mechanism has a unique pure strategy equilibrium where players choose finite arrival times. This equilibrium implements efficient target arrival times. In this mechanism only the player who arrives last makes a non-zero transfer payment equal to the externality he imposes on others.

Namely, define game Γ' as following. There are N players. Each player's action is 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 t is equal to μ 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 VCG 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 equilibrium⁷, and this equilibrium is the socially optimal vector of target arrival times.

Proposition 7 *For a synchronization problem where c_i and d_i are positive⁸,*

⁶By Proposition 6, in the 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. See Ostrovsky and Schwarz (2002, in progress) for a more formal treatment of individual incentives of agents and resulting equilibria.

⁷For 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 expected payoff of each agent is negative infinity.

⁸If delay cost is zero game Γ' has multiple pure strategy equilibria that implement the

the corresponding game Γ' has exactly one pure strategy equilibrium. In this equilibrium each player chooses a socially optimal target arrival time.

Proof. See Appendix. ■

When there are no costs of delay, i.e. $\forall i, d_i = 0$, there is also a simple mechanism which guarantees that the sum of all payments is zero: the player who arrives last pays other players half of their waiting costs.⁹ Another way to get a mechanism where the sum of all payments is zero is to make every player pay the difference between the average total (delay + waiting) cost of all players and his own total cost (when this difference is negative, the player receives a payment).

Vickrey-Clark-Groves mechanism makes each player's payment dependent on the actual arrival times of other players. Such mechanisms may be hard to implement in practice—for example, a subcontractor would probably agree to a contract specifying a schedule of fines if he is late, but quite possibly would not sign a contract in which his fine depends on performance of other subcontractors. We can eliminate the above mentioned problems by modifying the mechanism so that transfer payments of each player depend only on his actual arrival time and not on the actual arrival time of other players. In the modified mechanism each player pays a fine equal to the expected externality that he caused computed based on his arrival time under the assumption that other players select socially optimal target arrival times. More formally, let game Γ'' be the same as game Γ' , except for the form of the transfer payment γ_i . The transfer payment is given by $\gamma_i(t_i) = -E[\{w \left(\sum_{j \neq i} (c_j + d_j) \right)$ if player i is last, 0 otherwise $\}|t_i]$, where w is the difference between t_i and the next-to-last player's arrival time, and expectation is taken assuming socially optimal behavior by other players.

Each player's individual incentives (in equilibrium) do not change, and so our previous equilibrium result holds. Interestingly, the resulting fine $-\gamma_i(t_i)$ is convex.

Proposition 8 *The transfer payment $-\gamma_i(t_i)$ in game Γ'' is convex. The game has one equilibrium, which is socially optimal.*

socially optimal outcome. Adding a constant to all target arrival times in one socially optimal Nash Equilibrium generates another socially optimal Nash Equilibrium profile.

⁹See Appendix for the proof and also for a generalization of this mechanism to the case where costs of delay may be positive.

Proof. Each player's optimal strategy is independent of other players, and, by construction, is the same as his equilibrium strategy in game Γ' . To show that $-\gamma_i(t_i)$ is convex, notice that for small ϵ , the difference $(-\gamma_i(t_i + \epsilon)) - (-\gamma_i(t_i))$ is equal to (probability that player i is last if he arrives at $t_i + \epsilon$) \times $\left(\sum_{j \neq i} (c_j + d_j)\right) + o(\epsilon)$. Hence, $(-\gamma_i'(t_i)) =$ (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. ■

Consequently, we may expect to find contracts with late penalty increasing faster than linearly in time.

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 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 Milgrom and 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 9 *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. ■

5 Social Optimum for the Case of Endogenous Variance

In the previous sections we consider a model where each agent's action consists of the choice of the target arrival time. We assumed that the distribution of the disturbance term is entirely outside of an agent's control. This need not necessarily be the case—in many situations agents can also control the shape of their distributions of disturbances, at some cost. In the present section we consider a particular case where agents can control variances of their arrival times.¹⁰

Suppose each agent can reduce the variance of his arrival time at a cost. Namely, 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.

¹⁰A common way of reducing variance at a cost is giving an employee several tasks with different deadlines. Then by adjusting the proportion of time spent on each of them the employee is able to reduce variance of completion time. On the other hand, constantly switching from project to project is quite taxing for many people—this is the cost of reducing variance.

¹¹Again, we 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.

As in Section 4, we assume that all agents' costs of delay are equal to 0 and costs of waiting are equal to 1. However, we now allow for n agents. Let $w(\sigma_1, \dots, \sigma_n)$ be the expected total waiting time when agents have standard deviations $\sigma_1, \dots, \sigma_n$ and choose their expected arrival times optimally. Notice that our restrictions on the families of disturbance distributions imply that function w is homogeneous of degree 1.

Definition 10 *Call family of distributions $F_\sigma(\cdot)$ submodular if for any n function $w(\sigma_1, \dots, \sigma_n)$ is submodular.*

Lemma 11 *When agents' distributions of disturbances are submodular, in the optimum they all choose the same variance and aim at the same arrival time.*

Proof. Suppose in the optimum agents choose $\sigma_1, \dots, \sigma_n$, which are not all the same. Take n^2 agents: n with variance σ_1 , n with variance σ_2 , etc. Since agents' distributions are submodular, $nw(\sigma_1, \dots, \sigma_n) > w(\sigma_1, \dots, \sigma_1) + \dots + w(\sigma_n, \dots, \sigma_n)$, and so $n[w(\sigma_1, \dots, \sigma_n) + g(\sigma_1) + \dots + g(\sigma_n)] > [w(\sigma_1, \dots, \sigma_1) + ng(\sigma_1)] + \dots + [w(\sigma_n, \dots, \sigma_n) + ng(\sigma_n)]$. But then for some i , $w(\sigma_i, \dots, \sigma_i) + ng(\sigma_i) < w(\sigma_1, \dots, \sigma_n) + g(\sigma_1) + \dots + g(\sigma_n)$. ■

An interesting feature of Lemma 11 is that it places no restrictions on the form of cost function $g(\cdot)$ —e.g. regardless of whether it is convex, concave, etc. the agents still (in the optimum) choose equal variances.

Let's now see what happens to each agent's expected waiting time in the optimum as the number N of agents goes to infinity.

Proposition 12 *Suppose agents' distributions of disturbances are submodular. If they are unbounded (on the right), then as the number of agents N goes to infinity, in the optimum the expected waiting time of each agent goes to $-\lim_{\sigma \rightarrow 0^+} \sigma g'(\sigma)$. If distribution of disturbances is bounded, each agent's expected waiting time goes to a positive constant.*

Proof. See Appendix. ■

For example, if $g(\sigma) = -\frac{\sigma^a}{a}$, then as the number of agents N goes to infinity, each agent's expected waiting time goes to 0 (if $a > 0$) or to $+\infty$ (if $a < 0$). If $g(\sigma) = -\ln(\sigma)$, each agent's expected waiting time goes to 1 (in fact, as is clear from the proof of the proposition, it is equal to 1 for any number of agents).

It is clear that as N increases, it makes sense to reduce variance of individual components. Proposition 12 implies that (depending on the cost function $g(\cdot)$) it may pay to reduce variance so much that the average waiting time of an individual component also goes down as the number of components goes up.

6 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 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. ■

7 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 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, a non-cooperative model 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). We explore non-cooperative models of synchronization in the context of adoption of standards in Ostrovsky and Schwarz (in progress).

References

- [1] Fudenberg, Drew, David Levine, and Eric Maskin (1994), The Folk Theorem in Repeated Games with Imperfect Public Information, *Econometrica*, 62, 997-1039.

- [2] Grosh, Doris Lloyd (1989). *A Primer of Reliability Theory*, John Wiley & Sons.
- [3] Harris, Robert (1970). An Application of Extreme Value Theory to Reliability Theory, *Annals of Mathematical Statistics*, 41(5), 1456-65.
- [4] Kremer, Michael (1993). The O-Ring Theory of Economic Development, *Quarterly Journal of Economics*, 108(3), 551-575.
- [5] Milgrom, Paul and John Roberts (1990). The Economics of Modern Manufacturing: Technology, Strategy, and Organization, *American Economic Review*, 80(3), 511-528.

A Proof of Proposition 4

Before we begin the proof of Proposition 4, let's prove the following

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 $P(X_1 = \max\{X_1, X_2, z\}) \geq P(X_2 = \max\{X_1, X_2, z\})$, 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.

$$P_1 = P(X_1 = \max\{X_1, X_2, 0\}) = \frac{b_1 - b_2}{b_1 - a_1} + \frac{b_2}{b_1 - a_1} \left(\frac{1}{2} \cdot \frac{b_2}{b_2 - a_2} + \frac{-a_2}{b_2 - a_2} \right).$$

$$P_2 = P(X_2 = \max\{X_1, X_2, 0\}) = \frac{b_2}{b_2 - a_2} \left(\frac{1}{2} \cdot \frac{b_2}{b_1 - a_1} + \frac{-a_1}{b_1 - a_1} \right).$$

$$P_1 - P_2 = \frac{b_1 - b_2}{b_1 - a_1} - \frac{a_2 b_2}{(b_1 - a_1)(b_2 - a_2)} + \frac{a_1 b_2}{(b_1 - a_1)(b_2 - a_2)} = \frac{(b_1 - b_2)(b_2 - a_2) + (a_1 - a_2)b_2}{(b_1 - a_1)(b_2 - a_2)}.$$

Rearranging the inequality for the means, $a_1 - a_2 \geq -(b_1 - b_2)$, so

$$P_1 - P_2 \geq \frac{(b_1 - b_2)(b_2 - a_2) - (b_1 - b_2)b_2}{(b_1 - a_1)(b_2 - a_2)} = \frac{-a_2(b_1 - b_2)}{(b_1 - a_1)(b_2 - a_2)} > 0.$$

■

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 other agents distributions' supports 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.

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

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 μ exists because the total cost is continuous in μ and the "relevant" range is compact. Take an optimal μ . 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 μ satisfying this condition, suppose there are two, μ^1 and μ^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 μ^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 Mechanisms

C.1 Proof of Proposition 7

By construction, socially optimal arrival times form an equilibrium. Now, suppose there are two different equilibrium vectors of target arrival times, μ^1 and μ^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 μ^1 and (b) his probability of being last is strictly higher in μ^1 than μ^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. ■

As a side note, there are many mixed equilibria in this game. In this 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.

C.2 Generalization of the “Pay (Waiting Cost)/2” Mechanism

Proposition 15 *In addition to our basic setup, suppose that the player who arrives last has to pay every player j 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 $s_j = d_j - \frac{\sum d_k}{2N-2}$.*

Then this game has a unique equilibrium, 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 $q_i(d_i+\sum r_j-r_i+\sum s_j-s_i) = (1-q_i)(c_i-r_i) \Leftrightarrow q_i(d_i+\sum \frac{c_j}{2}-\frac{c_i}{2}+\frac{N-2}{2N-2}\sum d_j-d_i+\frac{1}{2N-2}\sum d_j) = (1-q_i)(\frac{c_i}{2}) \Leftrightarrow q_i(\sum c_j - c_i + \sum d_j) = (1-q_i)c_i \Leftrightarrow q_i(\sum c_j + \sum d_j) = c_i$. ■

D Proof of Proposition 9

D.1 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-Roberts 1990, Theorem 2). This derivative is equal to $-ab(a^2 + b^2)^{-3/2} \leq 0$. ■

D.2 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 , $a \leq b$. $w(a, b) = \frac{a}{b} \cdot \frac{a}{3} + \frac{b-a}{b} \cdot \frac{a+b}{4} = \frac{1}{12} \cdot \frac{a^2}{b} + \frac{b}{4}$. $w''_{ab}(a, b) = -\frac{1}{6} \frac{a}{b^2} \leq 0$. ■

E Proof of Proposition 12

Let $w_N(\sigma) = w(\sigma, \dots, \sigma)$ with N agents. Then the total cost, with N agents, is $Ng(\sigma) + w_N(\sigma)$. The first order condition on σ implies $Ng'(\sigma) + w'_N(\sigma) = 0$. Since w is homogeneous of degree 1, $w'_N(\sigma) = w'_N(1)$. Expected waiting time of one agent is $w_N(\sigma)/N = \sigma(w_N(1)/N) = -\sigma g'(\sigma)$. If the distribution is

unbounded, then as N goes to infinity, $w_N(1)/N$ goes to infinity, and so σ goes to zero. If the distribution is bounded, then as N goes to infinity, $w_N(1)/N$ goes to a positive constant, and so σ goes to a constant and $\sigma(w_N(1)/N)$ goes to a constant. ■

F 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

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

if $\mu_i > m_i$ and

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

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

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

if $\mu_i > m_i$ and

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

if $\mu_i = m_i$.

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

$$\begin{aligned} c_i \int_{-\infty}^{\infty} f_i(t_i - \mu_i) dt_i &= \sum (c_j + d_j) \int_{-\infty}^{\infty} \text{Prob}(t_i \geq t_j \forall j | t_i) f_i(t_i - \mu_i) dt_i \\ &\Updownarrow \\ \frac{c_i}{\sum (c_j + d_j)} &= \text{Prob}(i = \text{last}), \end{aligned}$$

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

Now suppose $\mu^*(\beta)$ does not go to $\mu^*(1)$ as β goes to 1. Then there exists a subsequence $\{\beta^n\}$ converging to 1 such that $\mu^*(\beta^n)$ converges to some $\tilde{\mu} \neq \mu^*(1)$ (set of $\mu^*(\beta)$ is bounded as $\beta \rightarrow 1$). Then by continuity, $\tilde{\mu}$ 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 $\mu^*(1)$.

(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,

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

Since $\mu_1 < \mu_2$,

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

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

Let t_i^* be such that $\sum (c_j + d_j) \text{Prob}(t_i^* = \text{last}) - c_i = 0$. 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,

$$\begin{aligned} & \int \beta_2^{t_i} (\sum (c_j + d_j) \text{Prob}(t_i = \text{last}) - c_i) f(t_i - \mu_2) dt_i \geq \\ & \left(\frac{\beta_2}{\beta_1}\right)^{t_i^*} \int \beta_1^{t_i} (\sum (c_j + d_j) \text{Prob}(t_i = \text{last}) - c_i) f(t_i - \mu_2) dt_i > 0. \end{aligned}$$

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