American Economic Review 2014, 104(2): 609–623
http://dx.doi.org/10.1257/aer.104.2.609

Time Allocation and Task Juggling†

By Decio Coviello, Andrea Ichino, and Nicola Persico*

A single worker allocates her time among different projects which are progressively assigned. When the worker works on too many projects at the same time, the output rate decreases and completion time increases according to a law which we derive. We call this phenomenon "task juggling" and argue that it is pervasive in the workplace. We show that task juggling is a strategic substitute of worker effort. We then present a model where task juggling is the result of lobbying by clients, or coworkers, each seeking to get the worker to apply effort to his project ahead of the others'. (*JEL* J22, M12, M54)

This paper studies the way in which a worker allocates time across different projects, or equivalently, effort across different projects through time. We study, in particular, the phenomenon of *task juggling* (frequently called multitasking), whereby a worker switches from one project to another too frequently.

Task juggling is a first-order feature in many workplaces. Using time diaries and observational techniques, the managerial literature on time-use documents that knowledge workers (engineers, consultants, etc.) frequently carry out a project in short incremental steps, each of which is interleaved with bits of work on other projects. For example, in a seminal study of software engineers Perlow (1999, p. 64) reports that

> [...] *a large proportion of the time spent uninterrupted on individual activities was spent in very short blocks of time, sandwiched between interactive activities. Seventy-five percent of the blocks of time spent uninterrupted on individual activities were one hour or less in length, and, of those blocks of time,* 60 *percent were a half an hour or less in length.*

Similarly, in their study of information consultants, González and Mark (2005, p. 151) report that

> [...] *the information workers that we studied engaged in an average of about* 12 *working spheres per day.* […] *The continuous engagement with each working sphere before switching was very short, as the average working sphere segment lasted about* 10.5 *minutes.*

* Coviello: Institute of Applied Economics, HEC Montreal, 3000, chemin de la Cte-Sainte-Catherine, Montréal, Québec, Canada, H3T 2A7 (e-mail: decio.coviello@hec.ca); Ichino: Dipartimento di Scienze Economiche, Universitá di Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy, and Department of Economics European University Institute, via Piazzola 43, 50133 Firenze, Italy (e-mail: andrea.ichino@unibo.it); Persico: Kellogg School of Management, Northwestern University, 2001 Sheridan Road Evanston, IL 60208 (e-mail: nicola@nicolapersico.com). Thanks to Gad Allon, Canice Prendergast, Debraj Ray, and Lars Stole.

† Go to http://dx.doi.org/10.1257/aer.104.2.609 to visit the article page for additional materials and author disclosure statement(s).

The fact that much work is carried out in short, interrupted segments is, in itself, a descriptively important feature of the workplace. But what causes these interruptions? The time-use literature points to the "interdependent workplace," meaning an environment in which other workers can (and do) ask/demand immediate attention to joint projects which may distract the worker from her more urgent tasks. One of the workers interviewed by González and Mark (2005, p. 152) puts it this way:

> [...] *Sometimes you just get going into something and they* [*call*] *you and you have to drop everything and go and do something else for a while* [...] *it's almost like you are weaving through, it is like, you know, a river, and you are just kind of like: "Oh these things just keep getting in your way," and you are just like: "get out of my way" and then you finally get through some of the other tasks and then you kind of get back, get back along the stream, your tasks* [...].

The literature on human scheduling, instead, attributes task juggling to the cognitive limitations of individual human schedulers. Crawford and Wiers (2001, p. 34), for example, write:

> [...] *One way in which human schedulers try to reduce the complexity of the scheduling problem is by simplification* [...]. *However, a simplified scheduling model leads to the oversimplification of the real system to be scheduled, and this in turn creates unfeasible or suboptimal schedules.*

The physiological constraints on scheduling ability are explored in the medical literature.[1] The popular press, however, has already rendered its verdict: scheduling is a challenge for many workers for reasons both internal and external to the worker. Popular literature books such as Covey (1989) and Allen (2001) exhort (and attempt to help) the reader to prioritize better. In *The Myth of Multitasking: How "Doing It All" Gets Nothing Done* (Crenshaw 2008, p. 89), we find a list of suggestions designed to help people reduce multitasking on the job. The first two are (i) "resists making active [e.g., self-initiated] switches; "(ii) minimize all passive [e.g., other-initiated] switches."

Effects of Task Juggling on Productivity.—We are interested in task juggling insofar as it affects productivity. The next example illustrates a source of productivity loss which is inherent to task juggling.

EXAMPLE 1: *Consider a worker who is assigned two independent projects, A and B, each requiring ten days of undivided attention to complete. If she juggles both projects, for example working on A on odd days and on B on even days, the average duration of the two projects is equal to* 19.5 *days. If instead she focuses on each project in turn, she completes A on the tenth day and then takes the next ten days to complete B. In the second case, the average duration of both projects from the time of assignment is* 15 *days. Note that under the second work schedule project B does* not *take longer to complete, while A is completed much faster; in other words, avoiding task juggling results in a Pareto-improvement across projects durations.*

[1] See, e.g., Morris et al. (1993) and Baker et al. (1996).

The example shows that a worker who juggles too many projects takes longer to complete each of them, than if she handled projects sequentially. The latter procedure corresponds to the "greedy algorithm," which is widely studied in the operations research literature.[2]

Outline of the Paper.—As a first step toward more complex models, in this paper we focus on *a single worker* who faces time allocation issues. In Section II we model a production process which may feature task juggling. Formally, the model is summarized in a system of four functional equations (1) through (4). Finding a solution to this system represents an original mathematical contribution which is offered in Theorem 1. Based on this solution, we demonstrate that effort and task juggling are *strategic substitutes*. This means that anything that makes workers juggle more tasks will also, indirectly, reduce the worker's incentives to exert effort.

Section III addresses the incentives that might generate task juggling. We model a lobbying game in which the worker allocates effort under pressure by her coworkers, superiors, or clients. This model is inspired by the idea of "interdependent workplace" discussed in the introduction. We fully characterize the equilibrium of the lobbying game and show that, no matter how low the cost of lobbying, in equilibrium there will be lobbying, which will induce task juggling. This model provides a microfoundation of task juggling.

I. Related Literature

What we call task juggling is viewed as an aberration in the queuing literature. The queuing literature prescribes algorithms ("greedy"-type algorithms, usually) that prevent task juggling. As we discussed in the introduction, we believe that this particular aberration is worth studying because it arises empirically, arguably as a predictable result of incentives. From a technical viewpoint, our model also departs from the queuing literature because that literature usually focuses on giving algorithms that keep queuing systems *stable*, that is, sufficient conditions under which queues can't ever get unacceptably or infinitely long.[3] Our model is by nature unstable because the arrival rate exceeds the worker's capacity (in our notation, $\alpha > \eta/X$). We believe that there is merit in going beyond stable queuing systems because stability requires the serving facility to be idle at least a fraction of their time, which is counterfactual in many environments.[4] Finally, our paper is distinct from most of the queuing literature in that the study of the incentives, such as the ones we examine, is largely absent from that literature.

In the economics literature, Radner and Rothschild (1975) discuss task prioritization by a single worker. They give conditions under which no element of a multidimensional controlled Brownian motion ever falls below zero. The control represents a worker's (limited) effort being allocated among several tasks, and the dimensions of the Brownian motion represent the satisfaction levels with which each task is performed.

[2] The name "greedy" refers to prioritizing those projects which are closest to completion (which project A is after day 1).

[3] An exception to the focus on stability is Dai and Weiss (1996) who do study the evolution of an unstable queuing network.

[4] In Coviello, Ichino, and Persico (2013), for example, we study the work organization of judges. Judges are never idle: in our data, they always have a backlog of cases that they should be working on.

Although broadly similar in its subject matter, that paper is actually quite different from the present one. Among other differences, it features no discussion of incentives.

Task juggling is studied in the sociological/management literature on time use (see Perlow 1999 for a good example and a review of the literature). This literature uses time logs and observations to document the patterns of uninterrupted work time, and the causes of the interruptions. This literature identifies "interdependent work" as the source of interruptions. The "lobbying by clients" model presented in Section III captures this effect. At a more popular level, there is a large time management culture which focuses on the dynamics of distraction and on "getting things done" (see e.g., Covey 1989; Allen 2001).[5]

The managerial "firefighting" literature (see Bohn 2000; Repenning 2001) documents the phenomenon whereby an organization focuses resources on unanticipated flaws in almost-completed projects (firefighting), and in so doing starves projects at earlier development stages of necessary resources, which in turn ensures that these projects will later require more firefighting, etc. This phenomenon is specular to the one we study because in our model the inefficiency is caused by too few, not too many, resources devoted to late-stage projects.

Dewatripont, Jewitt, and Tirole (1999) provide a model in which expanding the number of projects a worker works on will indirectly reduce the worker's incentives to exert effort. We get the same effect in Proposition 1. In their setup, the effect results from the worker's incentives to exert effort in order to signal his ability. This effect is different than the one analyzed in this paper.

II. The Production Process

In this section we introduce a dynamic production process which incorporates the possibility of multitasking in a very simple way. Imagine a worker who is assigned a stream of projects over time at rate α. Assuming the worker cannot deal with all the projects instantaneously, then the worker has to choose how to deal with the excess. We assume that, as cases are progressively assigned to the worker, she puts them in a queue of inactive cases. The worker draws from this queue at rate ν. A case drawn from the queue is "put in production": in our language, the case becomes active. All active cases receive an equal share of the worker's attention, a production process that generalizes the task juggling production process described in Example 1.

This modeling approach allows us to span the range between much task juggling (ν large, approaching α) and no task juggling, close to "greedy" (ν low). We will derive an exact formula for the production function which, given an effort rate, a degree of complexity of projects, and a level of task juggling, yields an output rate. Having an exact formula for the production function will allow us later to study strategic behavior pertaining to task juggling.

[5] For a review of the academic literature on this subject see Bellotti et al. (2004).

A. *The Model*

The model lives in continuous time, starting from $t = 0$. At time 0 the worker has no active projects. There is a continuum of projects. Projects are assigned at an exogenous rate α.

Each project takes X steps to complete. A project is characterized, at any point in time, by its degree of completion $x \in [0, X]$, which measures how far away the project is from being completed. We call a project *completed* when $x = 0$. Note that, because x is a continuous variable, we are assuming that there is a continuum of steps for each project. X can be interpreted as measuring the complexity of the project, or the worker's ability.

As soon as the worker starts working on a project, we say that the project becomes *active*. All projects remain active until they are completed. At any time t, the worker has A_t active projects, in various degrees of completion. The distribution $\varphi_t(x)$ denotes the mass of active projects which are exactly x steps away from being done. By definition, the number of active projects at time t is

$$A_t = \int_0^X \varphi_t(x)\ dx. \tag{1}$$

We assume that all active projects are moved toward completion at a rate η_t/A_t, where η_t is the rate at which effort is exerted. Informally, this means that in the time interval between t and $t + \Delta$, the worker shaves off approximately $(\eta_t/A_t)\,\Delta$ steps from *each* active project.[6] This formulation captures the idea that the worker divides a fixed amount of working hours equally among all projects active at time t. This procedure means that the worker is working "in parallel" on all active projects. If all active projects proceed at the same speed, then after Δ has elapsed, the distribution $\varphi_t(x)$ is translated horizontally to the left (refer to Figure 1), and so for Δ "small enough" we can write intuitively

$$\varphi_{t+\Delta}\left(x - \frac{\eta_t}{A_t}\Delta\right) = \varphi_t(x).$$

To express this condition rigorously, bring $\varphi_t(x)$ to the right-hand side, divide by Δ and let $\Delta \to 0$ to get

$$\frac{\partial \varphi_t(x)}{\partial t} - \frac{\partial \varphi_t(x)}{\partial x}\frac{\eta_t}{A_t} = 0. \tag{2}$$

This partial differential equation embodies the assumption of perfectly parallel work on the active projects.

The projects that fall below 0 (grey mass in Figure 1) are the ones that get completed within the interval Δ. These are the projects whose x at t is smaller than $\frac{\eta_t}{A_t}\Delta$. Therefore, the mass of output between t and $t + \Delta$ is approximately

$$\int_0^{\frac{\eta_t}{A_t}\Delta} \varphi_t(x)\ dx.$$

[6] Note that this formulation requires $A_t > 0$.

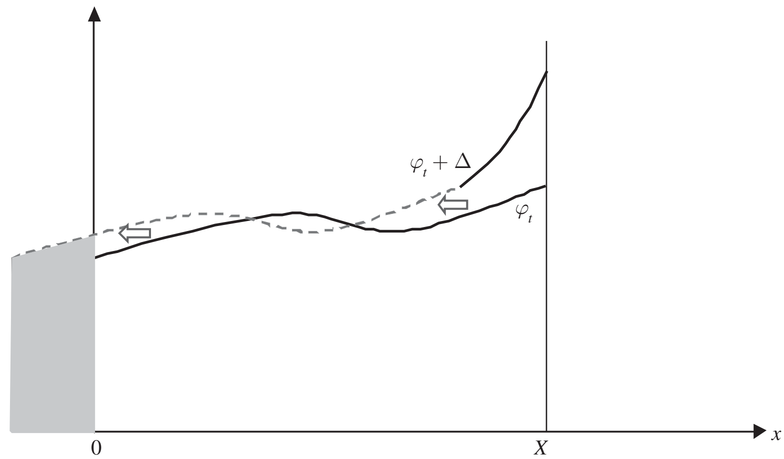

Figure 1. The φ_t Function

Note: The figure depicts the relationship between input and output rates on a growth path.

To get the *output rate* ω_t, divide this expression by Δ and let $\Delta \to 0$ to get

$$\omega_t = \lim_{\Delta \to 0} \frac{1}{\Delta} \int_0^{\frac{\eta_t}{A_t}\Delta} \varphi_t(x)\, dx = \frac{\eta_t}{A_t} \varphi_t(0). \tag{3}$$

The worker is not required to open projects as soon as they are assigned. Rather, we allow the worker to open new projects at a rate ν_t. A larger ν_t will mean more task juggling—more projects being worked on simultaneously. This ν_t is seen either as a control variable: depending on the specific environment, either a choice on the part of the worker, or determined by lobbying, or else imposed by some regulation. For Δ small, the change in the mass of projects active at t is approximately

$$A_{t+\Delta} - A_t = \nu_t \cdot \Delta - \omega_t \cdot \Delta.$$

Divide both sides by Δ and let $\Delta \to 0$ to get the formally correct expression

$$\frac{\partial A_t}{\partial t} = \nu_t - \omega_t. \tag{4}$$

Graphically, the mass of newly opened projects is squeezed in at the back of the queue in Figure 1, just to the left of X, in whatever space is vacated on the horizontal axis by the progress made in Δ on the preexisting open projects.

The description of the production process is now complete. In the production process, two variables are interpreted (for now) as given: η_t and ν_t. The first describes *how much* the worker works, the second *how* she works—how many projects she keeps open at the same time. These two variables will determine, through the process described mathematically by equations (1) through (4), the output rate ω_t which is the key variable of interest. This variable, in turn, will determine how long a project takes to complete. Our first major task is to uncover the law through which η_t and ν_t determine ω_t. We turn to this next.

B. *Derivation and Characterization of the Production Function*

To build some intuition about how η_t and ν_t determine ω_t, let us start with the "greedy" input rate. Fix a constant effort level $\eta_t = \eta$. The "greedy" input rate is $\nu_t = \eta/X$. At this input rate, in every short time interval Δ the worker starts work on $\Delta \cdot (\eta/X)$ new projects. At the beginning of time, between $t = 0$ and Δ, there are no preexisting projects and so the only active projects are the newly started ones. Formally, $A_0 = \Delta \cdot (\eta/X)$. According to our specification of the production process, in this first time interval the worker's effort shaves off approximately $(\eta/A_0)\,\Delta = X$ from each active project. This means that by time $t = \Delta$ all active projects have been completed. Therefore, the throughput rate during this first interval equals the input rate, and all projects are completed almost instantaneously (to be exact, within Δ of being started). Let us now turn to the second time interval $(\Delta, 2\Delta)$. This second time interval is exactly identical to the first one, and so the same conclusions apply: the throughput rate is equal to the input rate and projects are completed instantaneously. The same logic applies to all successive time periods.

What goes wrong when the input rate exceeds the greedy level? In this case the worker is not able to complete within the first time period all the projects which were started. These projects will need additional work during the second time period, which will divert effort from projects started in the second period. Thus, projects started in the second period will receive less attention during period 2, than first-period projects received in period 1. Therefore period 2-projects will be even less complete when they reach period 3. This effect snowballs down to all future projects. Soon, a period is reached where the worker is simultaneously working on many vintages of projects, some of which will only be completed far in the future. This means that a fraction of the current period's worker effort will not pay off today, but only in the future. This observation suggests that today's throughput should be smaller, relative to the greedy case. However this is not obvious, because it is also true that some of yesterday's effort pays off today. Nevertheless, in this section we prove that when the input rate exceeds the greedy level, throughput is smaller than its greedy level.

DEFINITION 1: *Fix X. We say that input and effort rates* ν_t, η_t generate *output rate* ω_t *if the quintuple of positive real functions* $[\nu_t, \eta_t, \varphi_t(x), A_t, \omega_t]_{x\in[0,X]}^{t\in(0,\infty)}$ *satisfies* (1), (2), (3), (4), *and* $A_0 = 0$.

The next theorem identifies the law through which ν_t and η_t generates ω_t. Implicitly, then the theorem identifies the production function. The theorem restricts attention to the case in which ν_t and η_t are constant and equal to ν and η respectively.

THEOREM 1 (Production Function): *The pair of constant functions* $[\nu_t = \nu,\ \eta_t = \eta]$ *generate* $\omega_t \equiv \omega$ *if the triple* ν, η, ω *solves*

$$\omega\,\frac{X}{\eta} - \log(\omega) = \nu\,\frac{X}{\eta} - \log(\nu). \tag{5}$$

PROOF:

We start by guessing a functional form for $\varphi_t(x)$ and A_t. Fix η and pick any two real numbers ν and $\omega > \nu$. Let

$$\varphi_t^*(x) = \frac{(\nu - \omega)}{\eta} \omega \, t \, e^{\frac{\nu-\omega}{\eta} x},$$

and

$$A_t^* = (\nu - \omega) \, t.$$

One can verify directly that for any K, λ, the pair $\varphi_t(x) = Kte^{\frac{\lambda}{\eta}x}$, $A_t = \lambda t$ solves (2) above. Moreover, for any λ the triple $\varphi_t(x) = Kte^{\frac{\lambda}{\eta}x}$, $A_t = \lambda t$, ω_t satisfies (3) if and only if $K = \frac{\lambda}{\eta} \omega_t$, which implies $\omega_t = \omega$. Finally, the triple ν_t, A_t, ω satisfies (4) if and only if $\lambda = \nu_t - \omega$, which implies $\nu_t = \nu$. This shows that, for any ν, ω, the quadruple $[\nu, \varphi_t^*(x), A_t^*, \omega]$ satisfies all the required equalities except (1). We now show that the pair $\varphi_t^*(x) = Kte^{\frac{\lambda}{\eta}x}$, $A_t^* = \lambda t$ solves (1) if and only if equation (5) holds. This equation implicitly identifies which values of ν, η, and ω are compatible with each other.

Condition (1) reads

$$A_t^* = \int_0^X \varphi_t^*(x) \, dx.$$

Substituting for $\varphi_t^*(x)$ and A_t^* yields

$$\lambda t = \int_0^X Kte^{\frac{\lambda}{\eta}x} \, dx$$
$$= \frac{\eta}{\lambda} Kt \left[e^{\frac{\lambda}{\eta}X} - 1 \right].$$

Now substitute for $K = \frac{\lambda}{\eta} \omega$ and $\lambda = \nu - \omega$ and rearrange to get

$$\frac{\nu}{\omega} = e^{\frac{(\nu-\omega)}{\eta} X}.$$

Taking logs yields equation (5).

Finally, the last condition in Definition 1 is satisfied because $A_0^* = 0$. Therefore, Theorem 1 is proved.

Equation (5) implicitly yields the production function we are seeking. The equation is most easily interpreted as follows: given an effort rate η and degree of task juggling ν, the implicitly identified ω represents the generated output rate. A convenient result also proved by Theorem 1 is that given constant effort and input rates, a constant output rate is generated. This is actually a subtle result, as we discuss on page 1 in the online Appendix.

We will now study the properties of the implicit production function. Before we start, however, an observation. The functions $\varphi_t^*(x)$, A_t^* identified in Theorem 1 are only well defined if the input rate ν exceeds the output rate ω. Expressed in terms of primitives, this condition is equivalent to $\nu > \eta/X$. (This equivalence is proved in the online Appendix). The limiting case η/X represents the "greedy" input rate, the smallest input rate at which the worker is never idle.[7] So our analysis is restricted to input rates such that the worker is never idle. From now on, we implicitly maintain this "non-idleness" assumption.

PROPOSITION 1 (Comparative Statics on the Production Function): *For each pair* $(\nu, \eta/X)$ *denote by* $\Omega(\nu; \eta/X)$ *the unique* $\omega < \nu$ *that is generated by* ν, η *through* (5). *Then we have:*

(*i*) $\Omega(\nu; \eta/X)$ *is decreasing in* ν.

(*ii*) $\Omega(\nu; \eta/X)$ *is increasing in* η/X.

(*iii*) $\frac{\partial\Omega(\nu; \eta/X)}{\partial\nu\partial\eta} < 0$, *which means that* ν *and* η *are strategic substitutes in the production of* ω.

(*iv*) *The function* $\Omega(\cdot; \cdot)$ *is homogeneous of degree* 1.

(*v*) $\Omega(\eta/X; \eta/X) = \eta/X$.

PROOF:

See the online Appendix. Part (v) is proved in Proposition 5 in the online Appendix.

Part (i) captures the effect of task juggling: increasing the input rate ν reduces output. Therefore setting ν as small as possible, provided that the worker is not idle, produces the maximum feasible output rate. Maximal output is therefore achieved when $\nu = \eta/X$. In that case, part (v) shows that the output rate equals η/X. This policy corresponds to the "greedy algorithm," and gives rise to a steady state which is analyzed in Proposition 5 in the online Appendix.

Part (ii) simply says that if a worker works more then the output rate is larger.

Part (iii) deals with the complementarity of inputs in the production of the output rate. It says that the returns to effort decrease when ν increases. Intuitively, this is because A_t is larger and so an increase in effort needs to be spread over a greater number of projects.

Part (iv) is a constant-returns-to-scale result: if we scale both inputs by the same parameter r, output increases by the same amount. The parameter r can be interpreted as governing the pace at which the system operates. Setting $r > 1$ means that the entire system is working at a faster pace: per unit of time, we have more input, more effort, and more output, all in the same proportion.

Part (iii) has implications for the scenario in which effort is chosen endogenously, rather than being exogenously given. Suppose effort $\eta^*(\widehat{\nu})$ is determined as the solution to the problem

$$\max_{\eta} \Omega\left(\widehat{\nu}; \frac{\eta}{X}\right) - c(\eta), \tag{6}$$

[7] The case $\nu \leq \eta/X$ is treated in Proposition 5 in the online Appendix.

where the input rate $\widehat{\nu}$ is exogenously given and the function $c(\cdot)$ represents the cost of effort. Problem (6) represents the problem of a worker choosing how much to work, given the constraint that she needs to put projects into production at rate $\widehat{\nu}$. Such constraints might be determined by the hierarchical organization of the workplace (how many coworkers can, or choose to, pressure the worker for their work to be done, as in the "interdependent workplace" modeled in the next section), or they can be mandated by regulation (Italian judges are required to start work on a case within 60 days of the case being assigned to them). Suppose that $c'(0) = 0$, which guarantees that the optimally chosen effort is positive. Then the following implication holds true.

COROLLARY 1: *Suppose effort* $\eta^*(\widehat{\nu})$ *solves* (6). *If the input rate grows to* $\widehat{\nu}' \geq \widehat{\nu}$ *then optimal effort* $\eta^*(\widehat{\nu}')$ *decreases relative to* $\eta^*(\widehat{\nu})$.

PROOF:
A direct consequence of Proposition 1 (iii).

This proposition highlights another dimension of inefficiency associated with task juggling. Not only does task juggling slow down projects, but it also induces output-motivated workers to slack off.

We now define two measures of durations: they are the measures employers or policy makers often care about.

DEFINITION 2: *For a project assigned at t we define the* duration D_t *as the time which elapses between t and the completion of the project. For a project opened at t* (*and thus assigned at a time before t*), *we define* completion time C_t *as the time which elapses between t and the completion of the project.*

The next result translates results about output rates into results about durations. The main takeaway is that because durations are decreasing in the output rate, task juggling increases durations.

PROPOSITION 2:

(i) Fix ω, ν, η. Then $C_t = \frac{(\nu - \omega)}{\omega}\, t$ and $D_t = \frac{(\alpha - \omega)}{\omega}\, t$.

(ii) Fix η, and let ω be generated by $[\nu, \eta]$. Then C_t and D_t are increasing in ν.

PROOF:
See the online Appendix.

III. Strategic Determination of Degree of Task Juggling and Endogenous Effort

In the previous sections we have assumed that ν_t, the exogenous input rate, is constant through time and, furthermore, that it exceeds the duration-minimizing "greedy" rate η/X. We have not discussed how such a ν_t might come about. In this section we microfound such a ν_t by introducing a game in which the input rate is determined endogenously as an equilibrium phenomenon. In the equilibrium of this game ν_t will in fact turn out to be constant through time, and to exceed η/X.

Therefore, this section microfounds the time-use behavior which was taken to be exogenous in the previous section.

The setup is that each project is owned by a different coworker, supervisor, or *client* who in each instant can lobby the worker to devote a fraction of effort to his project, regardless of its order of assignment. For the client, the private benefit of lobbying is to avoid his own project waiting inactive. But such lobbying has a negative externality on all other projects, because it increases the number of active projects which, as shown in the previous section, slows down all projects. This externality, which is not internalized by the lobbyists, gives rise to an excessively high input rate.

The model is as follows. The worker's effort η is constant through time and fixed exogenously (we will relax the second assumption later). Lobbying is modeled as a technology whereby, at any instant t, a client can pay $\kappa \cdot \Delta$ and force activity on his project during the interval $(t, t + \Delta)$. Activity on the project means that the project moves forward by $(\eta/A_t) \cdot \Delta$. The rate κ is interpreted as the per-unit of time cost of lobbying. If κ is not paid then the project sits idle at some x until either lobbying is restarted or the never-lobbied projects of its vintage (those assigned at the same time) *catch up* to x, at which time the project becomes active again and stays active without any need of, or benefit from, further lobbying. In every instant, $\underline{\nu}$ never-lobbied projects are opened, in the order they were assigned. Once a never-lobbied project is opened, it forever remains active whether or not it is lobbied. The rate $\underline{\nu}$ represents the input rate that would prevail in the absence of any lobbying by the clients.[8] In this section A_t denotes the mass of all projects active in instant t and it is composed of the two type of projects: all those that are lobbied in that instant, and some that are not.[9]

We assume that clients minimize B times the duration of their project, from assignment to completion, plus κ times the time spent lobbying. B represents the rate of loss experienced by a client whose project is not completed. We assume no discounting for simplicity.

In this model, clients are not allowed to use a variable amount of resources to lobby; rather, the cost of lobbying per unit of time is assumed to be fixed exogenously. We interpret this fixed cost as a sort of cost of supervision, the cost of stopping by and asking "how are we doing on my project?" or of exerting other kinds of pressures. We believe this formulation best captures the process that goes on within organizations, where monetary transfers of this kind are not allowed. Also, this type of lobbying process might take place after several principals have signed separate contracts with an agent, for example after several homeowners have contracted for

[8] One could be concerned that in equilibrium there might not be enough never-lobbied projects to open, and that therefore it would be more precise to state that in every instant the worker opens the minimum of $\underline{\nu}$ never-lobbied cases and the balance of the never-lobbied projects. However, we will see that in equilibrium the balance of never-lobbied projects never falls below $\underline{\nu}$.

[9] Under these rules, for a case that has been lobbied in the past, two scenarios are possible in instant t. First, the case may have been caught up by the never-lobbied cases of its own assignment vintage; in other words, the case was lobbied in the past, but then the lobbying lapsed and the case is now at the same stage of advancement (same x) as its never-lobbied assignment vintage. Such a case is worked on without the need for further lobbying and proceeds at speed η/A_t. The second scenario is that the case has not been caught up at time t. In this scenario the case is worked on in the interval Δ and makes $\eta\Delta/A_t$ progress if $\kappa\Delta$ is spent; otherwise, the case does not proceed.

the services of a single building contractor and now each is pushing and cajoling the contractor to finish his home first.

Since our goal is to explain why lobbying makes the input rate ν inefficiently large, let's tie our hands by stipulating that the input rate of never-lobbied projects $\underline{\nu}$ is "low," that is, it belongs to the interval $\left[0, \frac{\eta}{X}\right]$. This choice of baseline ensures that any slowdown in the output rate cannot be attributed to an excessively large $\underline{\nu}$.

Projects are indexed by the time τ they are assigned and by an index a that runs across the set of the α projects assigned at time τ. We now introduce the notion of lobbying strategy and lobbying equilibrium.

DEFINITION 3: *A* lobbying strategy *for project (a, τ) is a measurable indicator function $S_{a\tau}(t)$ defined on the interval $[\tau, \infty)$ which takes value* 1 *if project a is lobbied in instant t, and is zero otherwise. A* lobbying equilibrium *is a set of strategies such that, for each project (a, τ), the strategy $S_{a\tau}(t)$ minimizes κ times the time spent lobbying plus B times the project's duration.*

Equilibrium strategies could potentially be quite unwieldy, featuring complex patterns of activity interspersed with periods of no lobbying. Lemma 3 in the online Appendix characterizes equilibrium strategies, achieving considerable simplification. Based on that result, we conjecture (and show existence below) of simple equilibria in which a time-invariant fraction z of the α newly assigned projects is never lobbied, and the remaining fraction $(1 - z)\alpha$ is lobbied immediately upon assignment and then continuously until they are done. We will call these equilibria *constant-growth lobbying equilibria*. Note that the definition of constant-growth lobbying equilibrium does not restrict the strategy space.

If players follow the strategies of a constant-growth lobbying equilibrium, the input rate $\nu(z)$ is determined by z via the identity

$$\nu(z) = \underline{\nu} + (1 - z)\alpha.$$

The percentage of lobbyists $(1 - z^*)$, and hence the input rate $\nu(z^*)$, are determined in equilibrium.

The equilibrium construction is delicate. In every instant each client has a choice to lobby or not, and so in equilibrium each client has to opt to follow the equilibrium prescription. Moreover, every newly assigned client must be indifferent between lobbying and not. The cost of lobbying is proportional to the time the project is expected to require lobbying, which is the time that active projects take to get done. The drawback of not lobbying is the additional delay incurred from not "skipping the line."

PROPOSITION 3: *Suppose $\alpha > \frac{\eta}{X}$. Then, for any $\underline{\nu}$ and any cost of lobbying κ,*

(*i*) *a constant-growth lobbying equilibrium exists;*

(*ii*) *in any constant-growth lobbying equilibrium $\nu(z^*) > \frac{\eta}{X}$, i.e., the equilibrium input rate exceeds the duration-minimizing one;*

(*iii*) *the constant-growth lobbying equilibrium is unique;*

(*iv*) *the fraction* $(1 - z^*)$ *of projects that are lobbied in equilibrium is increasing in* $\frac{\alpha}{\underline{\nu}}$ *and* $\frac{\eta}{X}$, *and decreasing in* $\frac{\kappa}{B}$*;*

(*v*) *the equilibrium input rate* $\nu(z^*)$ *is decreasing in* $\frac{\kappa}{B}$ *and increasing in* $\frac{\alpha}{\underline{\nu}}$ *and* $\frac{\eta}{X}$.

PROOF:

See the online Appendix.

Part (i) can be viewed as providing a microfoundation for the behavioral assumption of constant ν_t which was maintained through Section II. What was previously a behavioral assumption about the worker is now the outcome of lobbying equilibrium where, in principle, ν_t need not be constant.

Part (ii) of the proposition says that, no matter how large the cost of lobbying, input rates will always exceed the "greedy" rate, and so we will have task juggling in equilibrium. The intuition is clear: if input rates were efficient, say $\nu \leq \eta/X$, then completion time would be zero.[10] This means that the cost of lobbying would be zero and, also, that a project which is lobbied would be completed instantaneously. Therefore lobbying is a dominant strategy, which would give rise to an input rate $\nu = \alpha > \eta/X$. Thus an equilibrium input rate ν cannot be smaller than η/X.

Part (v) of the proposition says that if a worker is less susceptible to lobbying, which we can model as κ being larger, then the worker will have a smaller input rate and a larger output rate. Moreover, there is more lobbying when the assignment rate is larger, which is intuitive because then a non-lobbying client anticipates waiting longer for his project to be opened. Finally, harder-working workers and easier projects will give rise to more lobbying. Intuitively, this is because then the completion time gets shorter relative to the duration of a non-lobbied project.

A few words of comment. Social inefficiency in this model results not only from the wasted cost of lobbying, but also from the decrease in the output rate. Thus, the inefficiency goes beyond that in a "common pool" model where a number of agents expend resources lobbying for a share of a fixed pie.

IV. Conclusion

Task juggling is prevalent in the workplace. We have developed a theory of a worker who chooses how many projects to work on simultaneously. Working on too many projects at the same time reduces the worker's output, for given effort and ability. We have derived the production function that describes the slowdown in output. We have shown that task juggling and effort are strategic substitutes in the production function, suggesting that when effort is not contractible, whatever worsens task juggling will also indirectly decrease effort. We have also modeled an

[10] That completion time is zero under the greedy rate was intuitively discussed on page X. Formally, the result follows from $A_0 = 0$ and Proposition 5 in the online Appendix.

"interdependent workplace" environment which will lead the worker to work on too many projects.

Our analysis does not touch on the possible counter-measures that might reduce task juggling. A principal, for example, might want to control an agent's task juggling through productivity-based incentives. If so, then could task juggling always be eliminated? We think not. When evaluating productivity is difficult, such as when knowledge workers have a monopoly over expertise as do physicians, scientific researchers, etc., strong productivity-based incentives may be counterproductive (Holmström-Milgrom multitasking).[11] In these cases weak incentives can be optimal and agents can easily fall prey to task juggling.

We view the single-worker model presented here as a building block for future research of two types. First, empirical work, which might take advantage of increasingly available workplace microdata to quantitatively evaluate the inefficiencies caused by task juggling, and to perform counterfactual calculations. In our companion paper (Coviello, Ichino, and Persico 2013) we use a related framework to estimate the causal effect of an exogenously-induced increase in parallel working. We find that the slowdown in output resulting from task juggling is large.

Second, we foresee the possibility of theoretical work extending this analysis to a multi-worker hierarchical workplace.

REFERENCES

Allen, David. 2001. *Getting Things Done: The Art of Stress-Free Productivity*. New York: Viking.

Baker, S. C., R. D. Rogers, A. M. Owen, C. D. Frith, R. J. Dolan, R. S. J. Frackowiak, and T. W. Robbins. 1996. "Neural Systems Engaged by Planning: A PET Study of the Tower of London Task." *Neuropsychologia* 34 (6): 515–26.

Bellotti, Victoria, Brinda Dalal, Nathaniel Good, Peter Flynn, Daniel G. Bobrow, and Nicolas Ducheneaut. 2004. "What a To-Do: Studies of Task Management Towards the Design of a Personal Task List Manager." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, April 24–29, Vienna, Austria,* 735–42. New York: Association for Computing Machinery.

Bohn, Roger. 2000. "Stop Fighting Fires." *Harvard Business Review*, July-Aug 2000, 83–91.

Covey, Stephen R. 1989. *The 7 Habits of Highly Effective Leaders.* New York: Simon & Schuster.

Coviello, Decio, Andrea Ichino, and Nicola Persico. 2013. "The Inefficiency of Worker Time Use." Unpublished.

Crawford, Sarah, and Vincent C. S. Wiers. 2001. "From Anecdotes to Theory: A Review of Existing Knowledge on Human Factors of Planning and Scheduling." In *Human Performance in Planning and Scheduling*, edited by B. L. MacCarthy and J. R. Wilson, 15–42. London: Taylor & Francis.

Crenshaw, Dave. 2008. *The Myth of Multitasking: How "Doing It All" Gets Nothing Done.* San Francisco: John Wiley and Sons.

Dai, J. G., and G. Weiss. 1996. "Stability and Instability of Fluid Models for Reentrant Lines." *Mathematics of Operations Research* 21(1): 115–34.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies." *Review of Economic Studies* 66 (1): 199–217.

González,Victor M., and Gloria Mark. 2005. "Managing Currents of Work: Multi-Tasking Among Multiple Collaborations." *In ECSCW 2005: Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work, 18–22 September 2005, Paris, France,* edited by Hans Gellersen, Kjeld Schmidt, Michel Beaudouin-Lafon, and Wendy Mackay, 143–62. Netherlands: Springer.

Holmström, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." Special Issue, *Journal of Law, Economics, & Organization* 7 (3): 24–52.

[11] See Holmström and Milgrom (1991).

Morris, R. G., S. Ahmed, G. M. Syed, and B. K. Toone. 1993. "Neural Correlates of Planning Ability: Frontal Lobe Activation During the Tower of London Test." *Neuropsychologia* 31 (12): 1367–78.

Perlow, Leslie A. 1999. "The Time Famine: Toward a Sociology of Work Time." *Administrative Science Quarterly* 44 (1): 57–81.

Radner, Roy, and Michael Rothschild. 1975. "On the Allocation of Effort." *Journal of Economic Theory* 10 (3): 358–76.

Repenning, Nelson P. 2001. "Understanding Fire Fighting in New Product Development." *Journal of Product Innovation Management* 18 (5): 285–300.

This article has been cited by:

1. David A. Miller, Kareen Rozen. 2014. Wasteful Sanctions, Underperformance, and Endogenous Supervision. *American Economic Journal: Microeconomics* **6**:4, 326-361. [Abstract] [View PDF article] [PDF with links]