

Time Allocation and Task Juggling*

Decio Coviello
HEC Montreal

Andrea Ichino
Bologna

Nicola Persico
Northwestern

November 2012

Abstract

A single worker is assigned a stream of projects over time. We provide a tractable theoretical model in which the worker allocates her time among different projects. When the worker works on too many projects at the same time, the output rate decreases and the time it takes to complete each project grows. We call this phenomenon “task juggling,” and we argue that this phenomenon is pervasive in the workplace. We show that task juggling is a strategic substitute of worker effort. We then present an augmented model, in which task juggling is the result of lobbying by clients, or co-workers, each of whom seeks to get the worker to apply effort to his project ahead of the others’. We take the model to new data on judicial productivity.

1 Introduction

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

*Thanks to Gad Allon, Canice Prendergast, Debraj Ray, Lars Stole. The paper refers to an Online Appendix, where some formal results are derived, which can be downloaded at the following link: http://nicolapersico.com/files/appendix_nopub_continuous.pdf
Email: nicola@nicolapersico.com; andrea.ichino@unibo.it; decio.coviello@hec.ca.

of which is interleaved with bits of work on other projects. For example, in a seminal study of software engineers Perlow (1999) 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 Gonzalez 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.”

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 Gonzalez 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 sub-optimal 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*, we find a list of suggestion designed to help people reduce multitasking on the job. The first two are:

- Resists making active [e.g., self-initiated] switches.
- Minimize all passive [e.g., other-initiated] switches.

(Cited from Crenshaw 2008, p. 89).

1.1 Effects of task juggling on productivity

We are interested in task juggling insofar as it affects productivity. The next example illustrates a productivity loss which is mechanical, and inherent to task juggling.

Example 1 *Consider a worker who is assigned two independent projects, A and B, each requiring 10 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 projects in turn, she completes A on the 10-th 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 projects 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.*

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

In addition to this mechanical slowdown, there may be “human” effects related to interruption: the worker may forget what she was doing before being interrupted, which impacts the speed or quality of the worker’s output. Or the worker may enjoy the variety and, conceivably, work harder if tasks are alternated.

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

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

1.2 Need for a model of task juggling

That task juggling must decrease productivity is well known from the operations research literature. But that literature has a *normative approach*: it tells us that workers *shouldn't* juggle and should do “greedy” instead. And yet the evidence overwhelmingly shows that workers actually do juggle. If we see this behavior in the workplace, and it is prevalent, then as economists we want to understand it. Accordingly, we are interested in the following *positive questions*:

1. How can we tell how much do workers juggle, if at all? (That is, establish an appropriate metric on juggling which can be taken to data).
2. How large, quantitatively, are the productivity consequences of juggling? (Spell out the relationship between juggling and productivity, for given effort and ability of the worker)
3. Who/what makes workers juggle? And when workers are made to juggle, do they do it in the specific way that our model assumes?
4. How does task juggling impact the workers' incentives to exert effort?

Some of these questions have to do with incentives and, therefore, are properly in the domain of economics as opposed to (classical) operations research. All of these questions are based on the premise that, contrary to the prescriptions of operations research, workers juggle tasks. To answer these questions, therefore, we need a new model, one where workers can juggle tasks. A useful model must be richer than the “toy model” in Example 1, because we want the model to match time series/panel data patterns like the ones presented in Section 6 below. It must also be tractable, because we want the model to answer questions 1-4 above. This paper provides such a model.

1.3 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 3 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 4 addresses the incentives that generate task juggling. We model a lobbying game in which the worker allocates effort under pressure by her co-workers, 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. We also show that, when worker effort is non-contractible, more intense lobbying which makes workers juggle more tasks will also, indirectly, reduce the worker’s incentives to exert effort. This indirect, strategic effect compounds the direct effect of task juggling.

Section 5 extends the analysis to the case in which the worker pays a cost from switching between projects. In Section 6 we take the model to the data, using original data on the productivity of Italian judges. We demonstrate that task juggling is prevalent among these judges, and that their choice of effort and scheduling behavior is well summarized, from an empirical viewpoint, by Theorem 1.

2 Related Literature

What we call task juggling is viewed as an aberration in the queuing literature. The focus of the queuing literature is to provide 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 the technical viewpoint, our model also departs from the queuing literature because that literature 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 capacity of the worker (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 (judges always have a backlog of cases that they should be working on, for example). 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

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

being allocated among several tasks, and the dimensions of the Brownian motion represent the satisfaction levels with which each task is performed. Although broadly similar in its subject matter, that paper is actually quite different from the present one. Among other differences, it focuses on optimality and 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 4 captures this effect. There is also a small literature devoted to measuring the *disruption cost* of interruptions, i.e., the additional time to reorient back to an interrupted task after the interruption is handled (see e.g. Mark *et al.* 2008, who review the literature). We introduce this cost in Section 5. At a more popular level, there is large time management culture which focuses on the dynamics of distraction and on “getting things done” (see e.g. Covey 1989, Allen 2001).⁴

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 *et al.* (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 4. In their setup, the effect results from the worker’s incentives to exert effort in order to signal his ability. Clearly, that effect is quite different than the one analyzed in this paper.

3 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 project over time. Assuming the worker cannot deal with all the projects instantaneously (a reasonable assumption), then the worker has to choose how to deal with the excess. We assume that, as cases are progressively assigned to the worker at rate α , she puts them in a queue of inactive cases. The worker draws from this queue at rate ν . When a case is drawn from the

⁴For a review of the academic literature on this subject see Bellotti *et al.* (2004).

queue it is “put in production” along with all other already active cases. Finally, we assume that the worker’s attention is divided in a perfectly equal fashion among all active cases, in a process that parallels the “perfect task juggle” in Example 1. This modeling strategy 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. Of note, in this section we abstract from the possibility that multitasking might cause the worker to forget; this additional effect of multitasking will be introduced in Section 5.

3.1 The Model

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

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**. The project stops being active when it is 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 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 towards 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’s work shaves off approximately $(\eta_t/A_t) \Delta$ steps from *each* active project.⁵ 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.” If all active projects proceed at the same speed, then after

⁵Note that this formulation requires $A_t > 0$.

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 $\Delta \rightarrow 0$ to get

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

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

Figure 1: The φ_t function



Note. The function φ_t is translated horizontally to the left as time passes. Newly opened cases are added to the right. The grey mass of cases to the left of zero are completed.

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

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

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

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

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, *ceteris paribus*, mean more task juggling—more projects being worked on simultaneously. This ν_t is seen either as a choice on the part of the worker, or as 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 $\Delta \rightarrow 0$ to get the formally correct expression

$$\frac{\partial A_t}{\partial t} = \nu_t - \omega_t. \quad (4)$$

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 pre-existing open projects.

This completes the description of the production process. In the model, two variables are interpreted (for now) as exogenously 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 key variable of interest, the output rate ω_t . This variable, in turn, will determine the duration of a project and its completion time.⁶ Our first major task is to uncover the law through which η_t and ν_t determine ω_t . We turn to this next.

⁶The two (endogenous) functions A_t and $\varphi_t(x)$ are, perhaps, of merely instrumental interest: they describe the state of the worker's docket at any point in time—how many projects he has open, and the degree of completeness of each.

3.2 Derivation and Characterization of the Production Function

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 $[\nu_t, \eta_t, \varphi_t(x), A_t, \omega_t]_{\substack{t \in (0, \infty) \\ x \in [0, X]}}$ satisfies (1), (2), (3) and (4).

The next theorem identifies the law through which ν_t and η_t determines ω_t . Implicitly, then the theorem identifies the production function. The result 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

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

Proof. We start by guessing a functional form for $\varphi_t(x)$ and A_t . 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 equalities except (1). However, we do not yet know which values of ν and ω are compatible with each other along a growth path. We now show that the pair $\varphi_t^*(x) = Kte^{\frac{\lambda}{\eta}x}$, $A_t^* = \lambda t$ solves (1) if and only if $X \frac{\nu - \omega}{\eta} = \log(\nu) - \log(\omega)$. Condition (1) reads

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

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

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

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

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

Taking logs yields

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

Therefore, Theorem 1 is proved. ■

Equation (5) implicitly yields the production function we are seeking. A convenient feature of the production function is that the output rate is constant (this is actually a subtle result, as we discuss on page 1 in the appendix). We will now be studying the properties of the 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 $\nu \geq \eta/X$. (This equivalence is proved in Appendix 1.1.) The threshold η/X represents the “greedy” input rate, the smallest input rate at which the worker is never idle. So our analysis is restricted to input rates such that the worker is never idle. (See Appendix 1.1 for what happens when $\nu < \eta/X$). 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:*

- a) $\Omega(\nu; \eta/X)$ is decreasing in ν .
- b) $\Omega(\nu; \eta/X)$ is increasing in η/X .
- c) $\frac{\partial \Omega(\nu; \eta/X)}{\partial \nu \partial \eta} < 0$, which means that ν and η are strategic substitutes in $\Omega(\nu; \eta/X)$.
- d) The function $\Omega(\cdot; \cdot)$ is homogeneous of degree 1.
- e) $\Omega(\eta/X; \eta/X) = \eta/X$.

Proof. See the Appendix. Part e) is proved in Proposition 2 in Appendix 1.1. ■

Part a) 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 e) 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 2.

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

Part c) 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 d) is a constant-returns-to-scale result: if we scale both inputs by the same parameter r , output increases by the same amount. 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 e) identifies the “greedy” rate of input. Given η and X , that rate is $\nu = \eta/X$. At this rate, output is $\omega = \eta/X$, the highest achievable output rate (given effort and ability).

We now define two measures of durations.

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 the time which elapses between t and the completion of the project.*

The next result translates results about output into results about durations. The main message is that task juggling increases durations.

Proposition 2 (a) *Fix ω, ν, η . Then $C_t = \frac{(\nu-\omega)}{\omega}t$ and $D_t = \frac{(\alpha-\omega)}{\omega}t$.*

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

Proof. See Appendix 1.1. ■

4 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 “micro-found” ν_t by introducing a game in which the input rate is determined endogenously as an equilibrium phenomenon. In this game ν_t will in fact turn out to be constant through time, and to

exceed η/X . Therefore, this section microfound the time-use behavior which was assumed to be exogenous in the previous section.

The basic setup is that each project is “owned” by a different co-worker, 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. The private benefit of lobbying is that the client avoids its project waiting unopened and gets the worker working on it immediately. The social cost of lobbying is that the worker distributes her effort among more projects. This will increase the number of active projects, which slows down all projects. This externality, which is not internalized by the lobbyists, gives rise to an inefficiency.

Clients are not allowed to use money 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 bribes 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 the services of a single building contractor and now each is pushing and cajoling the contractor to finish her home first.

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.⁷ Here 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.⁸

⁷One 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 $\underline{\nu}$.

⁸Under 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

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.

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 $[0, \frac{\eta}{\bar{x}}]$. 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 Appendix 1.2 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 adopt 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.

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.

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 ,*

- a) a constant-growth lobbying equilibrium exists;*
- b) in any constant-growth lobbying equilibrium $\nu(z^*) > \frac{\eta}{X}$, i.e., the input rate exceeds the duration-minimizing one;*
- c) the constant-growth lobbying equilibrium is unique;*
- d) 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 $\frac{\kappa}{B}$;*
- e) 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 Appendix. ■

Part a) can be viewed as providing a microfoundation for the behavioral assumption of constant ν_t which was maintained through Section 3. What was previously a behavioral assumption about the worker is now the outcome of lobbying equilibrium in which, in principle, ν_t need not be constant.

Part b) 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. The intuition is clear: if input rates were efficient, say $\nu \leq \eta/X$, then completion time would be zero. 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 e) 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 the time spent waiting for one’s project to be opened becomes larger. 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 on the causes of inefficiency. The source of a slowdown in output is that, if an additional project is lobbied, that project is able to obtain a small fraction of

effort, taking it away from other active projects. In this respect, our model is analogous to models of common resource extraction (“common pool” models) where utilizers cannot be excluded from the pool. We think this is a natural modeling assumption in many cases.

Finally we turn to the case in which η is chosen by the worker, rather than being exogenously given. Suppose η is determined as the solution to the problem

$$\max_{\eta} \Omega \left(\nu(z^*); \frac{\eta}{X} \right) - c(\eta). \quad (6)$$

According to this formulation, the worker chooses η by trading off the output rate (increasing in η) against a cost of effort $c(\eta)$. Note that since z^* is taken as given in problem (6), the worker does *not* behave as a Stackelberg leader. This assumption reflects the idea that the worker cannot commit to maintain a given level of effort regardless of lobbying.

Definition 4 *A lobbying equilibrium with endogenous effort is a lobbying equilibrium in which effort η^* solves (6).*

To ensure that the equilibrium effort level is greater than zero and smaller than αX we assume $c'(0) = 0$, and $c'(\alpha X) = \infty$.

Proposition 4 *Consider a lobbying equilibrium with endogenous effort. If κ increases, then the input rate decreases and the worker’s effort increases.*

Proof. See the Appendix. ■

This proposition highlights another dimension of inefficiency associated with lobbying. Not only does lobbying slow down projects, but it also induces the worker to slack off. The intuition behind this result lies in the “strategic substitutes” property stated in Proposition 1 c).

5 Extension: Forgetful Worker

In this section we deal with the case in which, as completion time grows and any open project is worked on less and less frequently per unit of time, the worker progressively forgets about the details of each individual project. Thus, every time the worker picks up a project again, she needs to spend some additional effort to “remind herself” of where she left off before she can make progress.

We model this phenomenon by assuming that in the time interval between t and $t + \Delta$, the worker's effort shaves off approximately $\frac{\eta}{A_t + F_t} \Delta$ steps from each active project. The factor $F_t > 0$ captures a “forgetfulness penalty.” We assume that F_t becomes larger over time; its exact form will be specified later. The presence of forgetfulness requires amending equations (2) and (3) from Section 3. The two amended equations read

$$\frac{\partial \varphi_t(x)}{\partial t} - \frac{\partial \varphi_t(x)}{\partial x} \frac{\eta}{A_t + F_t} = 0, \quad (7)$$

and

$$\omega_t = \frac{\eta}{A_t + F_t} \varphi_t(0). \quad (8)$$

Equations (1) and (4) remain unchanged.

Definition 5 Fix the function F_t . We say that the pair ν, η generate ω if the quintuple $[\nu, \eta, \varphi_t(x), A_t, \omega]$ satisfies conditions (1), (4), (7) and (8).

Now let us specify F_t . We want to capture the notion of “time elapsed between the accomplishment of two consecutive steps,” even though in our model steps are continuous and so strictly speaking there are no two consecutive steps. In our model, instead, we can think about the time that elapses between the accomplishment of given percentiles of completion, say between 20% and 30% of completion. A large completion time C_t corresponds to bigger stretches of time elapsing between the achievement of any two percentiles of completion, so we assume that the “forgetfulness penalty” F_t is proportional to the completion time C_t according to a factor of proportionality f . Formally, we assume

$$\begin{aligned} F_t &= f \cdot C_t \\ &= f \frac{\nu - \omega}{\omega} t. \end{aligned}$$

where the second equality follows from Proposition 2, and the third is simply a definition of the real number F . Note that, by making F_t proportional to C_t , we have made F_t endogenous.

Proposition 5 Suppose the worker is forgetful. Then:

(a) The pair ν, η generates ω if the following equation is satisfied:

$$\log(\nu) - \log(\omega) = \frac{X}{\eta} \left(\nu - \omega + f \frac{\nu - \omega}{\omega} \right). \quad (9)$$

(b) If $\nu > \frac{\eta}{X}$ and the worker is forgetful then the output rate is smaller than in the one described in Section 3.2.

Proof. See the Appendix. ■

6 Empirical Application

In this section we take to data the theoretical model of Section 3. The goal is to see whether the main result in that section, namely equation (5), fits the data at hand. How could such a “mechanical” model not fit the data? Although the model does not involve maximization, it nevertheless relies on a number of assumptions about the environment and the worker’s behavior. For example, it is assumed that certain choice variables (worker effort, opened cases) are stationary over time; that all open cases receive the same amount of worker effort; and that all cases require the same effort to dispose. Clearly, these assumptions can’t hold *literally* in any real-world workplace. In this sense, the model is an approximation of reality, both in terms of worker’s actions and in terms of the work environment. Is the approximation valid “in the wild,” in the sense that equation (5) fits the data? This section addresses this question.

The data refer to the production process of a panel of Italian labor law judges in the Milan labor court.⁹ Adjudication in this court takes place in a series of distinct steps (involving motions, hearings, etc.). The interleaving of these steps across different trials raises the possibility that a judge may, as in Example 1, work on too large a number of cases at a point in time. The possibility that task juggling may significantly reduce judicial productivity, which in Italy is a major policy concern, gives this empirical application considerable policy relevance.¹⁰

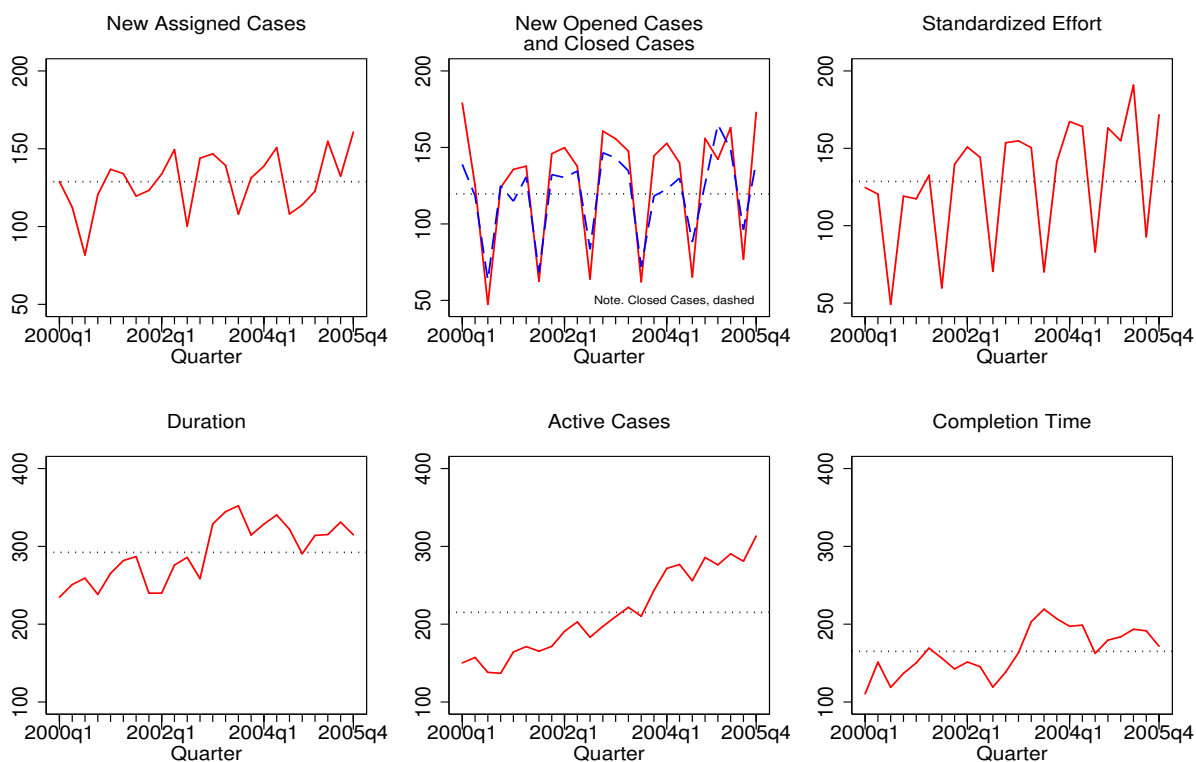
Figure 2 suggests visually that judges arrange their work process in a manner consistent with task juggling. Averaging across judges, the stock of **Active Cases** (cases which have received a first hearing but are not yet closed; see the bottom central panel of the figure) grows at a steady pace, from about 150 in the year 2000 to more than 300 in 2005. In contrast, the duration-minimizing “greedy algorithm” described in Example 1 would lead to a constant number of active cases. Furthermore, consistent with the model’s predictions, we see a steady increase of **Completion Time** C_t (computed as the number of days elapsing between the date in which the first hearing is held and the date in which the case is completed; see bottom right panel in the figure). That is, it takes judges longer and longer to work through a given case. Again, this finding suggests that judges do not follow the “greedy algorithm,”

⁹For details on this data see Appendix 2 and Coviello et al. (2012).

¹⁰In Italy in 2009, civil trials lasted on average 960 days in the court of first instance, and an additional 1,509 days in court of appeals (if appealed). Such durations place Italy at n. 88 in the world in “speed of enforcing contracts” as measured by the Doing Business survey of the World Bank—behind Mongolia, the Bahamas, and Zambia.

because under that procedure completion times would not increase over time. In sum, Figure 2 suggests that these judges do engage in task juggling.

Figure 2: Temporal evolution of the main stock and flow variables



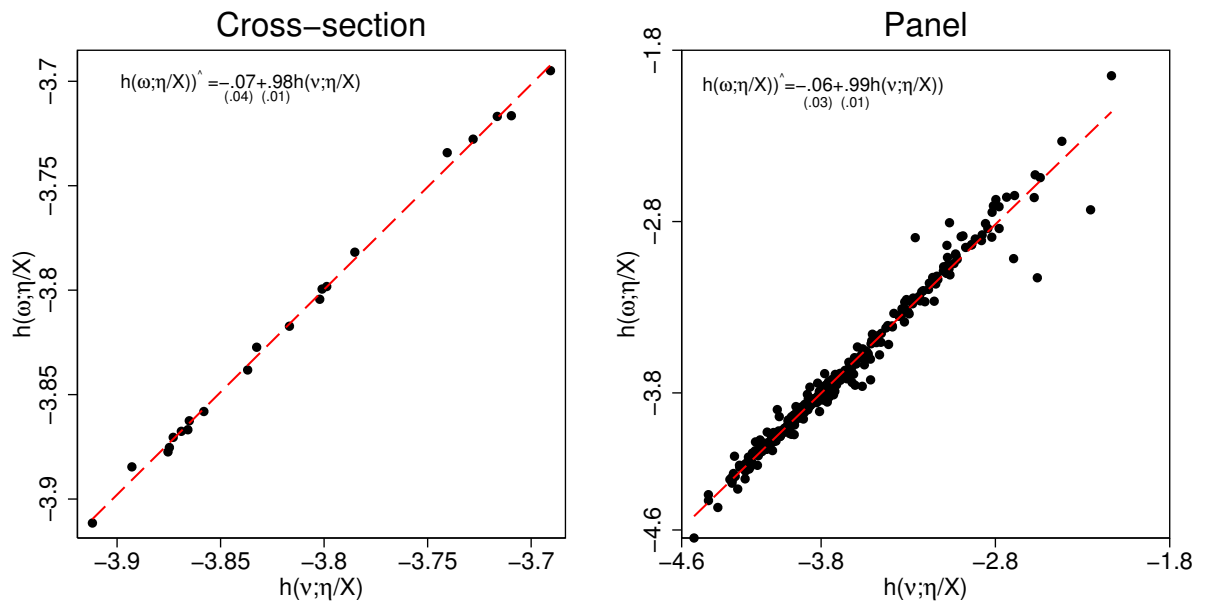
Note. All variables are defined as judges quarterly averages. *New Opened Cases* are defined as the number of cases that receive the first hearing in the quarter. *Standardized effort* is defined as the ratio between the number of hearings held by the judge in a quarter for any case (independently of the time of assignment) and the number of hearings that were necessary to decide the cases assigned to the judge in the quarter. *Duration* is defined as the number of days elapsing between the filing date and the date in which the case is completed. *Active Cases* are defined as the number of cases which have received a first hearing but are not yet closed. *Completion Time* is defined as the number of days elapsing between the date in which the first hearing is held and the date in which the case is completed. Dotted lines represent sample averages.

As interpreted through the theory in Section 3, therefore, Figure 2 suggests that the outcome of interest, **duration** D_t (computed as the number of days elapsing between the filing date and the date in which the case is completed; see bottom left panel in the figure) might actually be reduced by more judicious work practices *without adding manpower*. If true, this would have important policy consequences. However, a skeptical critical reader might look at Figure 2 and find some discrepancies with the theory in Section 3. For example, there is considerable seasonality of the “flow” variables in the top panel and a slight upward trend in them. So the model’s assumptions are, strictly speaking, violated. Similarly, Duration and

Completion time in the bottom panel, while trending up, are not exactly straight lines as predicted by Proposition 2. None of these discrepancies ought to be unexpected: no theory can be expected to match the data *exactly*. And yet, these discrepancies raise the interesting question of whether the “production function” given by equation (5), which is the main result of Section 3, is nevertheless a good description of the data. We address this question directly now.

In Figure 3 we portray graphically equation (5). We denote the left-hand side of equation

Figure 3: Graphical Representation of the Equilibrium Equation (5)



Note. The Figure portrays graphically $h(\omega; \frac{\eta}{X}) = h(\nu; \frac{\eta}{X})$, where $h(\omega; \frac{\eta}{X}) \equiv (\frac{X}{\eta}) \omega - \log(\omega)$, and $h(\nu; \frac{\eta}{X}) \equiv (\frac{X}{\eta}) \nu - \log(\nu)$, which is equation (5) in the paper. In the left panel, each dot represents the average of the quarter observations for each judge. In the right panel, instead, all quarters observations for each judge are considered separately, so that each dot represents a judge in a quarter. The dashed lines (text boxes) report the predicted linear fit (estimated coefficients and standard errors in parentheses) of equation $h(\omega; \frac{\eta}{X}) = \alpha + \beta h(\nu; \frac{\eta}{X}) + \varepsilon$, in the data.

(5) by $h(\omega; \eta/X)$ and its right-hand side by $h(\nu; \eta/X)$. The figure plots the left hand side on the vertical axis and the right hand side on the horizontal axis. If the equation held exactly, all the points should lie on the 45 degree line. In the left panel of Figure 3, each dot represents the average of the quarter observations for each judge. In the right panel, instead, all quarters observations for each judge are considered separately, so that each dot

represents a judge in a quarter. The caption of the figure describes how the variables are constructed in each panel. In both panels, the theoretical relationship appears to hold very tightly. The points are very closely aligned with the 45 degree line. In fact, when we estimate the regression

$$h(\omega; \eta/X) = \alpha + \beta h(\nu; \eta/X) + \varepsilon \quad (10)$$

and test the joint hypotheses that $\alpha = 0$ and $\beta = 1$, the null can never be rejected independently of whether the regression line is fitted on cross sectional data, or on panel data with and without judge fixed effects. As shown in Table 1 the P-values of the tests are 0.22, 0.10 and 0.12, respectively for these three specifications. An alternative testing strategy delivering similar results is a simple test for the equality of the means of the random variables on the two sides of equation (5). Also in these tests we can never reject the null, with *p-values* that are respectively equal to 0.44 and 0.42, respectively for the cross-sectional dataset of 21 judges and for the panel dataset of 380 quarter-judge observations.

Table 1: Test of the equilibrium equation (5)

Data	Cross-section (1)	Panel-OLS (2)	Panel-FE (3)
$\beta : h(\nu; \frac{\eta}{X})$	0.98 (0.01)	0.99 (0.01)	0.98 (0.01)
$\alpha : \text{Constant}$	-0.07 (0.04)	-0.06 (0.03)	-0.06 (0.03)
$H_0 : \alpha = 0 \ \& \ \beta = 1$	0.22	0.12	0.10
R^2	0.99	0.98	0.98
Number of judges	21	21	21
Observations	21	380	380

Notes. Each column reports coefficients (and standard errors in parentheses) estimated using regressions of the form:

$$h(\omega; \frac{\eta}{X}) = \alpha + \beta h(\nu; \frac{\eta}{X}) + \varepsilon$$

where $h(\omega; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \omega - \log(\omega)$, and $h(\nu; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \nu - \log(\nu)$. $H_0 : \alpha = 0 \ \& \ \beta = 1$ is the *p-value* of the test of the joint hypothesis that $\alpha = 0$ and $\beta = 1$.

It is important to realize that, using the same tests, we instead always strongly reject the null hypotheses that other functions of ν and ω are equal, like for example, $\nu = \omega$ and

$\left(\frac{X}{\eta}\right) \nu = \left(\frac{X}{\eta}\right) \omega$. Only the specific functional form suggested by our theoretical model, i.e. equation (5), fits the data well.

Table 2: Are judges forgetful?

Data	Cross-section (1)	Panel-OLS (2)	Panel-FE (3)
$\delta : h(\nu; \frac{\eta}{X})$	0.97 (0.01)	1.00 (0.01)	0.99 (0.01)
$f : q(\nu, \omega; \frac{\eta}{X})$	3.18 (1.44)	3.24 (0.80)	2.59 (0.82)
$\gamma : \text{Constant}$	-0.11 (0.04)	-0.02 (0.03)	-0.03 (0.03)
$H_0 : \gamma = 0 \ \& \ \delta = 1$	0.06	0.55	0.44
R^2	0.99	0.98	0.98
Number of judges	21	21	21
Observations	21	380	380

Notes. Each column reports coefficients (and standard errors in parentheses) estimated using regressions of the form:

$$h(\omega; \frac{\eta}{X}) = \gamma + \delta h(\nu; \frac{\eta}{X}) + f q(\nu, \omega; \frac{\eta}{X}) + \mu$$

where $h(\omega; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \omega - \log(\omega)$, and $h(\nu; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \nu - \log(\nu)$, and $q(\nu, \omega; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \left(\frac{\nu-\omega}{\omega}\right)$ is the indicator of forgetfulness. $H_0 : \gamma = 0 \ \& \ \delta = 1$ is the p -value of the test of the joint hypothesis that $\gamma = 0$ and $\delta = 1$.

We experimented also with a slightly richer specification of the production function, the one derived in equation (9). This specification allows judges to be forgetful. The forgetfulness parameter f was estimated using the following modified version of equation (10):

$$h(\omega; \eta/X) = \gamma + \delta h(\nu; \eta/X) + f q(\nu, \omega; \eta/X) + \mu \quad (11)$$

where $q(\nu, \omega; \frac{\eta}{X}) \equiv \left(\frac{X}{\eta}\right) \left(\frac{\nu-\omega}{\omega}\right)$. Also in this specification we still cannot reject the hypothesis that $\gamma = 0$ and $\delta = 1$, as the theory predicts. Moreover the value of the forgetfulness parameter f is estimated to be positive and significant. Therefore, we cannot reject the hypothesis that judges forget a case's details long after it was last treated. The estimates are reported in Table 2.

In sum, this section shows that the model of Section 3 fits the data quite well in this particular application, in the sense that equation (5) holds tightly in the data. True, this empirical analysis cannot answer the question of what leads judges to juggle hearings.¹¹ Nevertheless, it is important to know that the model has empirical validity. If this model can be similarly validated in other empirical settings, the model could become a benchmark in a future field of the measurement of worker time use.

7 Conclusion

Task juggling is prevalent in the workplace. We have developed a theory of a worker who deals with overload by choosing 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 investigated an “interdependent workplace” environment which will lead the worker to behave in this inefficient way. Moreover, we have shown that task juggling and effort are strategic substitutes, suggesting that when effort is not contractible, whatever worsens task juggling will also indirectly decrease effort.

A noteworthy feature of our model is that, unlike queuing theory, we look at an environment in which the worker is never idle. We do this for the purpose of realism: judges in Italy are never in a situation of no backlog, and the same is true of many other workers.¹² An implication of this assumption is that in our model the worker’s backlog and duration grow without bound as time goes on. We do not think of this trend as realistic in the long run, as long as the workforce can costlessly be expanded. But we interpret the model as a description of short and medium run congestion effects, such as those portrayed in Figure 2.

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

¹¹One reason is that judges are required by law to open all cases within 60 days of their being assigned, as we show in our companion paper Coviello et al. (2012). This rule, we believe, is intended to ensure that judges don’t shirk. In principle judges could satisfy this requirement by only holding the first hearing and then leaving the case aside. If they did so, then the theory would not yield equation (5). Since we find empirically that equation (5) does describe the data, we conclude that judges in fact are behaving as in our model: once a case is open, it is not left aside. Our analysis, therefore, demonstrates that judges do not interpret the legal requirement “cleverly.” The question is why. Perhaps the lobbying model in Section 4 can help provide the answer. While the judges under study are not lobbied explicitly, in our conversations they told us that they do feel a psychological pressure to be seen by the parties as working on each case. By pressuring workers to work on all assigned cases, this psychological pressure operates in ways analogous to the lobbying model of Section 4. That being said, the present section abstracts from the reasons that may lead a judge to juggle tasks, and hence this section does not speak directly to the theory developed in Section 4.

¹²Public health care systems in many countries, for example, also have permanent queues for various treatments.

incentives. If so, then could task juggling be eliminated? We think not. Casual observation (see the Introduction) and empirical work (see Section 6) suggest that, often, task juggling takes place. Why can't incentives work? We think the case of the judiciary is paradigmatic. In the judicial setting, there is a long tradition/culture of suspicion of explicit or implicit incentives. Promotions and pay raises among Italian judges, for example, are largely determined by seniority and almost not at all by productivity or quality measures. Similarly, the president of the Tribunal typically interprets his role in a very hands-off way. This institutional culture reflects the idea that strong incentives might lead judges to distort their decisions. The very notion that judges might be subject to a principal is seen as possibly erosive of judicial independence. In the language of economic theory, we would say that the Holmstrom-Milgrom multitasking distortions are so severe, that principals and incentives simply need to recede into the background. We believe that this is an important lesson from the case study, and that these same distortions apply quite broadly, if not always with the same power. Whenever knowledge workers have a monopoly over expertise (as do physicians, scientific researchers, etc.), Holmstrom-Milgrom multitasking kicks in, and providing incentives becomes dicey. Under this set of circumstances, weak incentives are often optimal and so, frequently, agents end up being relatively unconstrained in their ability to organize their own work schedule. In this case agents can 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 micro-data to quantitatively evaluate the inefficiencies caused by task juggling, and to perform counterfactual calculations. In this paper we have presented a way of using the data to validate the model. In our companion paper (Coviello et al. 2012) we use this 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

- [1] Allen, David (2001) *Getting Things Done*. Viking.
- [2] 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), pp. 515-526.
- [3] Victoria Bellotti, 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*, pp.735-742, April 24-29, 2004, Vienna, Austria.
- [4] Bohn, Roger (2000) "Stop Fighting Fires," *Harvard Business Review*, 78(4) (July-Aug 2000), pp. 83-91 .
- [5] Covey, Stephen (1989). *The seven habits of highly effective leaders*. New York: Simon & Schuster.
- [6] Coviello, Decio, Andrea Ichino and Nicola Persico (2012) "The Inefficiency of Worker Time Use." Manuscript, NYU.
- [7] Crawford, S., and V.C.S. Wiers (2001) "From anecdotes to theory: reviewing the knowledge of the human factors in planning and scheduling." In: B.L. MacCarthy and J.R. Wilson (Eds.), *Human Performance in Planning and Scheduling*, Taylor and Francis, London, 15-43.
- [8] Crenshaw, Dave (2008) *The Myth of Multitasking: How "Doing It All" Gets Nothing Done*. San Francisco, John Wiley and Sons.
- [9] Dai, J. G and G. Weiss (1996) "Stability and Instability of Fluid Models." *Mathematics of Operations Research* 21(1), February 1996
- [10] Dewatripont, Mathias, Ian Jewitt, and Jean Tirole (1999) "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies." *The Review of Economic Studies* 66(1) pp. 199-217.
- [11] Doing Business (2009) "Comparing Regulation in 181 Economies." The International Bank for Reconstruction and Development / The World Bank.
- [12] González, Victor M. and Gloria Mark (2005) "Managing Currents of Work: Multi-tasking Among Multiple Collaborations." In H. Gellersen et al. (eds.), *ECSCW 2005*:

Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work, 18-22 September 2005, Paris, France, 143–162.

- [13] Holmstrom, Bengt and Paul Milgrom (1991) “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design.” *Journal of Law, Economics, & Organization*, Vol. 7.
- [14] Mark, G., D. Gudith, and U. Klocke. (2008) “The Cost of Interrupted Work: More Speed and Stress.” In *CHI '08: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 107–110, Florence, Italy, April 2008. ACM Press.
- [15] 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), pp. 1367-1378 .
- [16] Perlow, Leslie (1999) “The time famine: Toward a sociology of work time.” *Administrative Science Quarterly*, 44 (1), pp. 57–81.
- [17] Radner, Roy and Michael Rothschild (1975) “On the allocation of effort.” *Journal of Economic Theory* 10 (1975), pp. 358–376.
- [18] Repenning, Nelson (2001) “Understanding fire fighting in new product development.” *Journal of Product Innovation Management* 1, pp. 85–300.