DRUID society

Paper to be presented at
the DRUID16 20th Anniversary Conference
Copenhagen, June 13-15, 2016

Deadlines, Work Flows, Task Sorting, and Work Quality

Natarajan Balasubramanian
Syracuse University
Whitman School of Management
nabalasu@syr.edu

Jay Lee
Drexel University
Management
jayleej@gmail.com

Jagadeesh Sivadasan
University of Michigan
Ross School of Business
jagadees@umich.edu

Abstract
Deadlines are often used to manage the time of valuable human capital. In this multi-method paper, we propose a theoretical framework grounded in a formal model that encapsulates the key drivers and consequences of deadline-related time pressures on work flows, task sorting and work quality. We use large-scale data on patent filings along with insights from primary data collection to test our hypotheses. In line with our predictions, we find clustering of patent filings around month-ends, with month-end applications being more complex than those filed on other days. Consistent with time pressure reducing work quality, we find that work quality is lower for tasks completed at month-ends, more so for process measures of quality than for outcome measures. Calibrating our model to the data allows us to shed light on the benefits of deadlines, and suggests small levels of task acceleration but potentially larger working capital--related benefits for law firms.

Jelcodes:O31,M50

Deadlines, Work Flows, Task Sorting, and Work Quality

February 10, 2016

Abstract

Deadlines are often used to manage the time of valuable human capital. In this multi-method paper, we propose a theoretical framework grounded in a formal model that encapsulates the key drivers and consequences of deadline-related time pressures on work flows, task sorting and work quality. We use large-scale data on patent filings along with insights from primary data collection to test our hypotheses. In line with our predictions, we find clustering of patent filings around month-ends, with month-end applications being more complex than those filed on other days. Consistent with time pressure reducing work quality, we find that work quality is lower for tasks completed at month-ends, more so for process measures of quality than for outcome measures. Calibrating our model to the data allows us to shed light on the benefits of deadlines, and suggests small levels of task acceleration but potentially larger working capital–related benefits for law firms.

INTRODUCTION

A critical resource-allocation challenge for managers, particularly in knowledge-intensive, high-wage sectors, is optimally managing the time of valuable human capital. In organizational settings, there are potential trade-offs between speed and performance. While increasing speed or limiting time allocated to specific tasks may increase the quantity of output and help better coordinate across tasks, the costs of doing so could include potentially poorer quality or innovativeness of work due to increased time pressure (e.g., Amabile et al., 2002; Kelly and Karau, 1999).

Firms often manage these trade-offs by using "deadlines," or specific end-dates to complete tasks, with penalties for delaying work beyond those deadlines (Bluedorn and Denhardt, 1988; Locke and Latham, 1984; Schriber, 1986).

In this paper, we propose a theoretical framework (grounded in a formal model) that encapsulates the key drivers and consequences of deadline penalties. In our framework, set in a principal-agent context, it is optimal for the principal to impose a deadline penalty if the trade-offs seen by the agent when determining the optimal task duration are not congruent with the trade-offs perceived by the principal. Our model generates predictions for three important outcomes: (i) work flows, (ii) task complexity, and (iii) work quality. Specifically, we predict that in the presence of a deadline penalty, (i) tasks will be clustered at the period-end deadlines, (ii) tasks completed at the deadline will, on average, be of higher complexity, and (iii) work quality will be lower for tasks completed at the deadline, both unconditionally and conditional on task complexity. Our model also implies that (iv) the adoption and/or severity of deadline penalties will increase with firm size and the pace of innovation.

We then test and confirm these predictions using large-scale, high-frequency data on granted patents (about 3 million) and published patent applications (about 1.9 million) in the United States, supplemented with primary data collection efforts through interviews and surveys of inventors and patent attorneys. The unique nature of the data allows us to examine work flows, task complexity, and work quality at a level of detail that has not been possible in prior studies of deadlines and time pressure. Indeed, our setting meets the challenging data requirements for empirical analysis of the effect of deadlines in real-world settings: data on a comparable activity, and reasonably uniform measures of task complexity, and work quality across numerous firms and time periods. Because the data cover millions of patents and patent applications, we observe work flows for millions of fairly comparable tasks spread over several decades and across thousands of different firms. Our data also allow us to form a number of alternative measures for task complexity so that we can study how deadlines affect some of the tasks to be prioritized over others before the deadlines. Perhaps most importantly, the data allow us to construct numerous measures of different aspects of work quality, permitting a detailed and nuanced examination of the effects of time pressure close to deadlines.

Consistent with our prediction, patent filing work flows exhibit significant clustering of filings at month-ends (Figure 2). For instance, the last working day of a month accounts for 7.0% of all patent applications filed in that month, compared with an expected uniform rate of 4.8% in the sample, suggesting that nearly a third of the applications on the last working day are related to month-end clustering. Several empirical tests, analyses of matched primary-secondary data, and interviews with practitioners confirm the observed clustering to be due to deadlines.

Next, we use characteristics of applications to examine task sorting around deadlines. Specifically, we use six measures that are likely to be related to the complexity of the underlying idea: the log number of claims in the application, log number of drawings, log number of drawing sheets, log number of drawings per drawing sheet, log number of forward citations, and the probability of renewal of the patent. Consistent with our prediction, we find robust evidence that month-end applications are of higher complexity (Figure 3). For instance, our regression estimates are equivalent to a 4.5–6.2% difference between the number of claims in applications influenced by month-end deadlines and other applications. The corresponding differences for citations and probability of renewal are 3.8–5.6% and 1.0–2.1%, respectively.

The final but a central piece of our empirical analysis focuses on the implications of deadlines on work quality. Based on the suggestions received during our interviews with attorneys, we look at a total of 11 measures of work quality. Eight of these measures are transactions during the prosecution (examination) of the patent by the United States Patent and Trademark Office (USPTO) and relate to work-process quality—that is, they reflect or arise from errors or omissions in the application document (for instance, "Separate Inventor Oaths" indicates that the inventor oaths were not included in the initial filing; "Restriction Requirement" means that the application contained two distinct inventions rather than one). The other three measures — share of cites that are added by the examiner, probability of application approval, and log of review duration — are work-outcome measures; unlike the process measures, these are based on outcomes observable in the patent/application data.

We find that all work-process quality measures are considerably lower for applications filed at month-ends, with the strongest effects generally for applications filed on the last day of the month. Among the largest magnitude of effects are the propensities for Separate Inventor Oaths (18.6% higher for last day of the month), Application Incomplete notice (15.8%), and Additional Application Filing Fees (14.8%).

The work-outcome measures also indicate poorer quality for patent filings closer to month-ends, but the magnitude of the effects is more modest. Applications filed at month-ends have more examiner-added cites (by 0.69% for the last five days of the month), a lower approval rate (by 0.55% for the last five days, which translates to a 2.26% higher probability of rejection), and a longer review duration (by 1–2.3% over the last five days of the month). The more modest effects for the work-outcome measures suggest that the short-run corrections captured by the work-process measures moderate the long-run quality consequences of the acceleration of filings close to deadlines.

Having examined work flows, task complexity, and work-quality results for the sample as a whole, we next check if our key results vary with technology and firm characteristics as hypothesized. Because principals are likely perceive greater benefits from an earlier priority date, we may expect more widespread use of deadlines (and/or stronger deadline penalties) in technologies with a rapid pace of innovation. Indeed, we find that month-end clustering behavior is more pronounced in such technologies. Importantly, our results on complexity and work quality are also stronger in such technologies, consistent with our inference that these effects are driven by deadline penalties. Because monitoring and enforcing deadline penalties are likely to involve fixed costs, we expect firms with larger patent volume to be more likely to adopt deadlines. We find this to be true. As expected, we find stronger complexity and work-quality results for larger firms. These two sets of results significantly strengthen our interpretation of the baseline results for month-end clustering, complexity and work quality as being driven by month-end deadlines.

Finally, we undertake a calibration exercise that allows us to evaluate potential benefits of deadlines and perform interesting counterfactual analyses. Our analysis shows that for our baseline scenario (with penalties for billing deadlines every 22 days calibrated to match the observed amount of month-end clustering in the data), time savings from acceleration of jobs are quite modest (1.8%) relative to the increase in error rate (12%), but this small acceleration leads to a significant saving in terms of working capital requirements (31.6% saving in days of unbilled balances) for the accelerated tasks.

On the whole, our study makes several important contributions to a number of related fields of research, with implications for management. First, our analysis of work flows contributes to the literature in labor and personnel economics that has examined how incentives affect work outcomes including timing of effort. Although we are the first (to our knowledge) to examine the effects of deadlines on patent application work flows, our results are consistent with findings in this literature that document agents' modification of the timing of work efforts in response to incentives (e.g., Asch, 1990; Courty and Marschke, 1997; Oyer, 1998). For example, Oyer (1998) documents that manufacturing firm sales are higher at the end of the fiscal year and lower at the beginning.

Second, the potential effect of deadlines on task sorting has, to our knowledge, not previously been investigated in the literature. Our finding of higher complexity of tasks completed close to deadlines provides a novel stylized fact, consistent with the model we propose. As we show in our model, this result is non-trivial; the prediction of higher complexity at the deadline hinges on the error rate not increasing too fast with complexity relative to how fast the deadline penalty changes with complexity.[1] In most modern organizations, workers and work groups simultaneously undertake multiple tasks of varying complexity, and managers are required to synchronize or coordinate across different tasks and work groups. Our finding provides suggestive evidence that stronger penalties are imposed on more-complex tasks, which is consistent with a number of the motivations we discuss for the imposition of the deadline penalty.

Third and perhaps most importantly, our findings directly speak to the literature on how time pressure affects work quality. While not focused on the effect of deadlines per se, parts of the rich literature in psychology, organizational behavior, and management on the effects of time pressure address the question of if and how time pressure affects work quality. The empirical work in this literature has yielded mixed results. Although some of the earlier work (e.g., Kelly and McGrath, 1985) found a negative effect of time pressure on performance quality, Isenberg (1981) and Karau and Kelly (1992) found evidence for a non-linear relationship, with the best outcomes achieved under moderate levels of time pressure. In contrast, a positive relationship between time pressure and work quality was found by Kelly and Karau (1999), and Andrews and Farris (1972) also found a positive long-run effect of time pressure on creativity. In more recent related work, Amabile et al. (2002) (as well as Amabile, Hadley, and Kramer, 2002), using a novel daily electronic questionnaire collected over 30 weeks from 177 individuals in seven companies, found a negative effect of higher time pressure on creativity. Thus, how time pressure affects work quality remains an important, yet open empirical question.

Our setting allows us to address this question in an extensively large sample context utilizing various measures of different dimensions of work quality. Our results suggest that increased time pressure close to deadlines does indeed have negative effects on work quality. These quality results are robust to numerous alternative tests (including using different sample periods and industry subsamples. Thus, across thousands of different firms over a fairly long time period, we find robust evidence for negative quality effects of time pressure, consistent with findings in some of the prior studies on time pressure (Amabile et al., 2002; Kelly and McGrath, 1985), but contrary to the

[1]In particular, if the error rate increased fast enough and more-complex tasks were penalized less, in equilibrium the complexity of tasks completed at the deadline could be *lower* than at other times.

positive effect found in other studies (Andrews and Farris, 1972; Kelly and Karau, 1999) and the null effect in Basett (1979).

Finally, our work contributes to the literature on innovation, by focusing attention on the relatively understudied application preparation part of the patenting process. Our work draws on and contributes to the smaller literature that uses information in the patent application/approval data to uncover underlying important features of the innovation and patenting process (Alcacer and Gittelman, 2006; Alcacer, Gittelman, and Sampat, 2009; Bessen 2008; Carley, Hegde, and Marco, 2014; Cockburn, Kortum, and Stern, 2002; Hegde, Mowery, and Graham, 2009; Lemley and Sampat, 2012; Sampat, 2010; Sukhatme and Cramer, 2014).

THEORETICAL MOTIVATION

We develop a formal mathematical model to motivate drivers and impacts of deadlines. Given the focus of this paper, however, we discuss the underlying intuition here, leaving the technical derivations to the Appendix (available on request).

The key elements of the model are summarized in Figure 1. Briefly, the model focuses on lawyers involved in the drafting, finalization, and filing of applications, as agents of a principal. The ultimate principal is the management of the inventing firm, but if the filing is completed by an external law firm, the partners of that firm also serve as a principal monitoring the attorney working on the filing. The agents complete tasks (that is, they file applications), possibly of varying "complexities," that arrive randomly and independently over time. The optimal time taken to complete a given task depends on the costs and the benefits to the agent from completing that task. These costs and benefits both increase, albeit at different rates, with "complexity" of the tasks but decrease with time spent on the task.

These agents face time pressure due to deadlines. Deadlines, which arise from considerations (left panel of Figure 1) described below, are defined as periodically occurring checkpoints when the agent faces a penalty for having incomplete tasks. This penalty is usually imposed directly by the principal—for instance, as lower performance evaluations or monetary incentives—but it can also be indirect, such as peer pressure associated with not meeting deadlines. We then analyze how agents respond to the deadline penalties, and study if that causes clustering of work flows, if they choose to accelerate some tasks more than others ("task sorting") depending on the complexity of the tasks, and the effect on observed work quality.

Need for Deadlines

In the context of our model, deadlines address the incongruity between the principal and agent with regard to the benefits and costs associated with adjusting the speed of completing a task. In general, as is well understood, the principal can try to align the agent's interests with theirs through suitably designed incentives. For instance, they can use performance measures to set bonuses, raises, or promotions for the attorneys who are undertaking the application work (e.g., Lazear and Gibbs, 2014, Ch. 9). However, as is also well known, there are numerous challenges to measuring performance and rewarding agents (e.g., Lazear and Gibbs, 2014, Ch. 9 and 10), so that even with a well-designed performance evaluation and incentive system, the perceived benefit and cost trade-offs from accelerating may not be the same for the principal and agent.

There are three key sources for such inconsistencies in our context. First, managers at the inventing firm (e.g., R&D managers) may have planning, reporting, and coordination considerations, which are not directly perceived by the attorneys. So the managers may require at least some subset of applications to be completed faster than what the attorneys would like to complete them in, leading the firm to impose deadlines. Second, principals may perceive a greater benefit from earlier filing than that perceived by the agent. For example, firms in more technologically competitive industries (i.e., with shorter cycle times) may wish to accelerate applications more than for the average firm across all sectors, which may not be internalized by the attorney working on the patent. Third, the partners of an external law firm may want to reduce working capital needs by minimizing un-billed balances. These working capital costs of unbilled hours may not directly enter the objective function for the attorney; then the benefit from accelerating seen by the principal is different from that perceived by the agent.

There may be additional sources of inconsistencies. For instance, the principal and agent may differ in the perceived opportunity cost of the agent's time (e.g., in terms of the importance of other tasks they need to work on, as in Holmstrom and Milgrom, 1991). Furthermore, procrastination tendencies (e.g., as discussed in Cadena et al., 2011) could lead the agents to delay completion of tasks, arising because the agents view the short-term costs of finishing the tasks as more salient than the benefits (e.g., O'Donaghue and Rabin, 2001). Taken together, these arguments imply that the principal is often likely to have a shorter optimal time for completing the task than the agent. In such a case, as shown in the appendix, it follows that imposing a deadline-related penalty will increase the principal's payoff.

Impact of Deadlines

When faced with a deadline penalty, agents weigh the cost of accelerating their task to meet the deadline with the benefits of such acceleration. For the agent, the key benefit of meeting the deadline is avoiding any deadline-related penalties. For instance, in our survey of attorneys, nearly 75% of the respondents said lower performance evaluation was a potential cost of missing a deadline, "sometime or more often". Sixty percent said the same for pressure regarding unbilled hours. The potential loss of repeat business due to client dissatisfaction or a damaged reputation is also a major cost, particularly for external attorneys.

However, the agent also faces costs associated with accelerating task completion. An important cost is the opportunity cost of the additional time the agent has to spend before the deadline completing the task (e.g., the agent may have to work late or over weekends). Another important cost relates to additional errors in the application due to rushing. In the worst-case scenario, the application may be fatally flawed, resulting in it being rejected entirely by the USPTO or a loss of priority date. However, given the reputation costs at stake for the agent, our interview and survey evidence suggest that is not very likely. More likely, though, are errors that lead to lack of clarity in the wording or cause additional transactions during the patent prosecution (review) process. Eight-four percent of our respondents cited lack of clarity in the "specification" (that is, the description of the invention being patented) as being a risk associated with rushing an application, "sometime or more often." Not surprisingly, nearly 65% of survey respondents rated the cost of an unclear specification—the loss of some claims—as a risk of rushing. Other risks include inadequate prior art disclosure and additional work and filing fees. Even though these costs are eventually borne by the inventing firm, they are passed to the agents, and hence affect their decision to (or not to) accelerate task completion.

Given the costs and benefits of task acceleration, we show that agents accelerate some tasks that would otherwise be completed "just after" the deadline, so that those tasks are completed at the deadline. To see the intuition, if a task would be completed on Apr 3 without acceleration, agents may complete it on Mar 31 if they face a modest deadline penalty. However, they are less likely to accelerate (to Mar 31) a task of the same complexity if it would normally be completed only on Apr 27. A much higher deadline penalty would be required to achieve that. Also, since there is no incentive to complete earlier than the deadline, all these accelerated tasks are completed at the deadline. Thus, it follows that:

Hypothesis 1: In the presence of a deadline penalty, there is clustering of task completions at the period-end deadline.

Agents will not accelerate all tasks equally. In particular, more-complex tasks are likely to take longer to complete and are likely to have a higher error rate. Hence, accelerating such tasks may be more costly for the agents. On the other hand, more-complex tasks are also likely to have greater deadline-related penalties than less complex ones. For instance, unbilled balances on a 90-day task that is scheduled to end on April 3 has an 87-day unbilled balance on March 31, while a 30-day task has only a 27-day unbilled balance. Hence, the law firm partner may be more motivated to accelerate billing on the larger tasks. Similarly, the management of an inventing firm may care more about more valuable patents, and impose greater penalties (or may be more likely to impose penalties) for not completing the associated filing by the deadline. Also, if the incongruity arises from behavioral biases, agents may be more likely to procrastinate for complex tasks (O'Donaghue and Rabin, 2001). Hence, whether more or less complex tasks are chosen for acceleration depends on how complexity affects the costs of accelerating and the penalty of being incomplete at the deadline. This leads to the following proposition.[2]

[2]The conditions under which more complex tasks get accelerated conditions are set out formally in the Appendix.

Hypothesis 2: Under some conditions, the average complexity will be higher for tasks completed at the deadline.

As discussed in the introduction, the theoretical literature on the impact of time pressure on work quality is ambiguous, and potentially suggests an inverted-U-shaped relationship. On one hand, time pressure may improve performance by increasing search intensity, as suggested by March and Simon (1958, pp. 116, 154), or by narrowing attention on task-completion-focused activities (Karau and Kelly, 1992; Parks and Cowlin, 1995). On the other hand, having too little time could reduce performance quality (especially for complex tasks) by restricting the amount of information considered or the thoroughness with which information is evaluated (Kelly and Karau, 1999).

In our model, deadlines affect average work quality in two ways. First, conditional on quality, accelerating tasks unambiguously increases the error rate (our error function has an inverse square dependence on task duration). But a second, subtler effect kicks in due to task sorting, so that the equilibrium average error rate for accelerated tasks is ambiguous. In particular, if the marginal error cost is high relative to the marginal benefit from speeding, the duration choice will be high enough that the error rate is lower for more-complex tasks. Then, because more-complex tasks get sorted to month-ends, unconditional on quality, the error rate could be *lower* for the month-end. Hence, it follows that:

Hypothesis 3a: If the marginal effect of complexity on the error rate is low enough (relative to the marginal effect on gains from acceleration), average work quality is lower for tasks completed at the deadline.

Hypothesis 3b: For any given task complexity, work quality is lower for tasks completed at the deadline.

From a principal's point of view, though deadlines enable faster task completion, they are also likely to have fixed costs associated with monitoring performance relative to the deadlines. Hence, large firms with sufficient volumes are more likely to find it beneficial to use deadlines to motivate earlier patent filings. Furthermore, coordination considerations are likely to be more important in large firms, which further increases their incentive to use deadlines.

Technology characteristics may also affect the prevalence and severity of deadline penalties. Specifically, technologies vary in the speed of innovation; some technologies change slowly, while others change rapidly. As one of our interviewees stated "… *pharma firms want every day of [the] patent term.… computer tech[nology] change[s] very quickly, so, computer companies do not really worry that much about full term."* Not surprisingly, this speed of change will directly affect the deadline-related incentives faced by firms and their attorneys. In particular, the incentive to file early increases in areas where technological change is rapid, as the risk of being scooped by a rival is higher. Therefore, deadlines are likely to be more pervasive or more stringent, and consequently, complexity and quality effects at month-ends are likely to be stronger in sectors with shorter cycle times. In particular, firms are more likely to file earlier in technologies that are changing rapidly, so that they can maintain an advantage over their rivals. This, in turn, increases the benefit of using deadlines to motivate earlier filings. Taken together, it follows that:

Hypothesis 4: The impact of deadlines will be greater for larger firms and in sectors with a more rapid pace of technological innovation.

DATA

We use (i) secondary data from various datasets and (ii) primary data from surveys and interviews of practitioners.

Datasets

The primary source of our patent data is the USPTO. We purchased data on all utility patents granted between January 1976 and August 2009 (a total of 3,209,376 patents). This data included the patent number, U.S. classes and subclasses, the number of claims, and application year of each patent, as well as citations to a patent. For most of our analysis, we use only data on patents assigned to "organizations," identified by assignee codes "2" and "3." We excluded patents assigned to universities and multiple assignees, and eliminated patents that were not applied on a USPTO working day (0.6% of patents). We supplemented these data with the NBER Patent Data (2009), an updated version of Hall, Jaffe, and Trajtenberg (2001).

We also used application data that include successful and unsuccessful patent applications, purchased from Fairview Research LLC. The data include information on all U.S. Pre-Grant Applications published by the USPTO between January 1, 2001, and December 31, 2010. In addition to the application filing date and application number, the database includes application type (utility or design), number of claims, publication type, and importantly (to control for firm fixed effects), standardized assignee names. To merge the application data with other patent databases, we did some further standardization of the assignee names (code available on request).

To examine potential differences in work quality, we collected the transaction history of patent applications from Google Patents. To allow for sufficient time between filing and the final decision, we limited our sample to applications filed by U.S. firms during 2001–2004. As of January 2012, about 52% of the applications were available for download. Even with this subset, downloading entire documents posed technical challenges, as many of the files were very large in size. Trading off these challenges against diminishing benefits of increasing sample size, we downloaded a 25% random sample of the available applications.

To these, we added data on patent attorneys from the Patent Network Dataverse (Lai, D'Amour, and Fleming 2009) and data on examiner-added citations from Sampat (2012). Lai et al. also disambiguate the names of inventors and assign a unique identifier to each inventor that appears in patents granted between 1975 and 2008. In a robustness check, we use this information to control for inventor-level heterogeneity by using inventor-specific fixed effects. The various samples used in the main analyses are available on request.

Surveys and Interviews

We supplemented the above data with information from three surveys and 21 interviews with practitioners (inventors and attorneys) who had extensive firsthand experience in filing patent applications. Our first survey ("Inventor survey") focused on inventors and attempted to understand their role in the patent application process. Specifically, we administered a survey to global alumni of a globally known science and technological institute based in India, who were working in R&D departments, primarily as inventors. We obtained about 140 complete responses to our survey.

The second survey ("Law firm patent attorney survey") was administered to patent attorneys at law firms, and received about 50 complete responses. The third survey ("Corporate patent attorney survey") was administered to patent attorneys at inventing firms, and received about 13 complete responses. In addition to getting their opinion on the observed clustering of filings, these two surveys aimed to understand the incentives and penalties faced by patent attorneys in the patent filing and prosecution process. Details on these surveys are available on request.

In addition to these surveys, we conducted in-depth interviews with practitioners. The objective was to obtain detailed information on the patenting process inside firms, as well as to gain insights on the observed clustering of patent applications. The list of interviewees included: (i) legal staff (mostly IP attorneys) at seven large patent-intensive firms, (ii) legal staff from nine independent IP law firms, and (iii) four inventors, each of whom had multiple inventions patented at different global electronics and semiconductor firms. The interviews, conducted either via phone or in person, lasted from a minimum of 20 minutes to a maximum of two hours. All of these interviews were semi-structured; we first asked the interviewee several standardized questions based on our observations and then opened the conversation to unstructured responses. A summary of responses from the interviewees is available on request.

EMPIRICAL ANALYSIS

Deadlines and Work Flow Clustering

Figure 2 plots each calendar day's average share of granted applications (for the period 1976 to 2009) for corporates and individuals, and the difference in day-shares between corporates and individuals. The vertical lines denote month-ends. The figure shows a substantial upward spike in filings at the end of every month for corporates, but no spikes for individuals (except for September-end, which corresponds to the year-end for the USPTO, with a potential for an increase in fees beyond this date). For corporates, September and June show the highest month-end spikes, while the clustering is somewhat muted for May, November, and December, all of which have a U.S. holiday at the end of the month.We combine these two in the bottom row, which plots the difference in day-shares between corporates and individuals, and confirms that the clustering is much stronger for corporates, consistent with the patterns arising from periodic routines and associated deadlines.

To formally test for differences between corporates and individuals, we present in Table 1 results from the following regression specification: $y_d = \beta_1 D_{1-7} + \beta_2 D_{8-15} + \beta_3 D_{16-23} + \beta_4 D_k + \epsilon_d$, where $D_k, k \in 1, 3, 5$ indicates a dummy variable for the last k days of the month, and D_{m-n} is a dummy variable for days m to n of the month. The dependent variable in columns 1 to 3 of are day-shares for corporates defined for day d, $s_d^F \equiv \frac{\sum_{t=1976}^{2009} n_{dt}^F}{\sum_{t=1976}^{2009} n_t^F}$ where n_{dt}^F is the number of successful patents (or published applications) applied by corporates in day d of year t, and n_t^F is the number of successful patents (or published applications) applied by corporates in year t. In columns 4 to 6, the dependent variable is $\overline{s}_d^I$ is defined similarly for individuals.

We see strong clustering of filing in the last few days of the month for corporates (columns 1 to 3), whereas the clustering is much more muted for individuals (columns 4 to 6). For readability, all coefficients and standard errors in this and other tables have been scaled by 1,000. The month-end effects are clearest when looking at the differences between the month-end dummy and the dummy for days 1–7 (in row 7 of Table 1); this shows that for corporates the share of filings in the last day of the months are about 0.09 percentage points (column 1) more than that for an average day in the first week of the month, while the same difference is only 0.028 percentage points for individuals (column 5), so that there is excess clustering for corporates (relative to individuals) of about 0.065 percentage points (column 7). This effect is significant in magnitude, considering that if patent filing patterns were uniform throughout the year, we would expect 0.27% (1/366) of annual patents on any given day for both individual inventors and firms. Thus, the excess month-end filing share is about 23.6% of the expected uniform filing rate.

Together, these results strongly support Hypothesis 1.

Deadlines and Task Complexity

Recall that in our model, complexity has the following characteristics: the benefits for completing more-complex tasks earlier are higher, and the cost of errors is higher for more-complex tasks, so that more time (intuitively, more effort) is required to complete a more "complex" task. In our empirical analyses, we examine six different observed variables that are likely to be correlated with the complexity of an application in the aforementioned sense. Specifically, we use the number of claims (available for both patent and application data), the number of drawings, the number of drawing sheets, and the number of drawings per drawing sheets (all available only in the patent data) as measures of complexity. Applications with more claims and drawings are likely to take longer to complete and often result in higher fees for attorneys (70% of our attorneys in our surveys said their fees depend on these variables). In addition, we used the number of citations to a patent within the five years from its application date and the probability of renewal at 3.5 years after the grant as additional measures. Though these measures are considered a proxy for patent value (e.g., Hall, Jaffe, and Trajtenberg 2005; Pakes 1986), more valuable patents are likely to involve more care and hence more effort in drafting the application, greater benefits to completing the application, and higher costs of any errors. Hence, citations and renewal are likely to indirectly reflect the deadline-related trade-offs being studied here.

Before we present formal regression estimates, we show results for three of the measures in the top panel of Figure 3. Across three key measures of complexity, there is a sharp increase toward the end of the month; patents or applications filed close to month-ends contain more claims, are cited more frequently, and are more likely to be renewed. We test the significance of this pattern for all six of our measures using the following specification:

$$C_{pjtm} = \beta_D . D^k_{pjltm} + \alpha . y_{jt} + \tau_t + \nu_m + f_j + \widetilde{\epsilon}_{pjtm} \tag{1}$$

where D^k_{pjltm} is a dummy defined earlier for the last k working days of the month, C_{pjtm} is a measure of complexity, y is a measure of firm size defined as the number of patents (or applications) belonging to firm j that were applied in year t, τ_t and ν_m are application year and application month fixed effects, and f_j denotes firm fixed effects.

For brevity, only the coefficients on the month-end dummy are presented in Table 2. The number of claims is significantly larger for month-end applications: 0.9% larger for applications filed in the last three days of the month, and 1.7% larger for those filed on the last working day. Similarly, the number of drawings and drawing sheets are higher by about 1.2% and 1.4%, respectively, for applications filed on the last working day of a month. Consistent with more-complex drawings being larger, the number of drawings per drawing sheet is lower by about 1.1% for applications filed on the last working day of the month. The two indirect measures of complexity also show a similar pattern. The five-year citation count is greater for month-end patents: the citation count is about 1.5% higher for patents filed in the last working day of the month. Month-end patents are also more likely to be renewed, by about 0.2%. Together, these tests strongly confirm Hypothesis 2, that applications of higher complexity are more prone to be completed toward month-ends.

Deadlines and Work Quality

We measure work quality of a patent application using two types of measures: (i) errors and additional work needed during the patent prosecution process ("work process–based measures") and (ii) the quality of outcomes associated with the application ("work outcome–based measures"). With regard to the former, we rely on specific transactions that occurred during the prosecution. For instance, the USPTO sends an "Application Incomplete" notice if the application is missing some documents, and records "Additional Application Filing Fees" if the applicant had to pay such fees at some point during the prosecution (typically due to delays in filing documents). Specifically, we included the following in the first group of measures: (i) application incomplete notice—filing date assigned; (ii) a statement by one or more inventors satisfying the requirement under 35 USC 115, oath of the applicant, is missing; (iii) additional application filing fees; (iv) new or additional drawing filed; (v) information disclosure statement (IDS) filed after the initial filing date; (vi) non-final rejection; (vii) request for extension of time - granted; and, (viii) restriction/election requirement.

We identified six of these transactions based on the list of transactions that appear in the data, and independently confirmed their validity through interviews with three different practicing patent attorneys. One measure (IDS filed after the initial filing date) was added based on the suggestion of an interviewee. In addition, we performed a factor analysis (PCA with varimax rotation) of these measures and identified a single factor (eigenvalue 2.23 vs. 0.53 for the next factor), which we added as another measure of process-based work quality. The practitioners confirmed that all of these measures are more likely to occur if the filing is rushed, and hence can be considered good indicators of a poorly prepared application. Although they are unlikely to result in the loss of a priority date, these transactions require rework and re-filing by the attorneys, and hence impose costs on the filing firm. Our interviewees indicated that the extra time and effort to address some of these transactions could be a sizable fraction of the cost of the initial application. Thus, in general, it is in the best interests of the inventing firm and its attorneys to ensure that the application is as well prepared as possible.

In addition to these eight work process measures, we chose three measures that are more closely related to the outcome of the work: (i) the share of citations added by examiners, (ii) application success, and (iii) log review time. The share of citations added examiners reflects the extent of prior art disclosure in the initial filing. Poorly prepared applications may tend to have fewer citations to prior art, requiring the examiner to add some of his or her own during the prosecution (examination) process. Application success is an indicator of whether the application resulted in a patent or not, and is a reflection of the overall outcome of the task (i.e., the filing of the application). Finally, the review time is also reflective of the overall outcome, though it can also be viewed as an indicator of the overall process quality.

To examine whether month-end applications are more likely to have a lower work quality (Hypothesis 3a), we estimate a regression model similar to that in Equation 1, except that the dependent variable is q_{pjtm}, a measure of work quality. For each of the eight work process measures, it is defined as 1 if the transaction history of that application contained that transaction and 0 otherwise. Application success was defined similarly depending on whether the application resulted in a patent or not. For the other two outcome measures, we used the actual values as the dependent variable. To test Hypothesis 3b, we add log number of claims as a control.

Table 3 presents the coefficients on the month-end dummy. They are positive and significant for all work process measures. Hence, month-end applications are much more likely to encounter these transactions than other applications. Specifically, applications filed on the last working day of the month are 4.6% more likely to receive an incomplete notice from the USPTO. The corresponding numbers for the last three and five working days of the month are 3.6% and 3.0%, respectively. Given that only 29% of applications in the sample receive such a notice, these differences are substantial (a 15.8%, 12.4%, and 10.3% higher probability relative to the baseline rate). Similarly, applications filed on the last working day are about 4.9% more likely to submit the inventor oaths separately, 4.7% more likely to pay additional filing fees, 1.5% more likely to submit an additional drawing, and 2.2% more likely to file an information disclosure statement after the initial filing. Compared with the respective baseline rates of these transactions, these are significant (18.6%, 14.8%, 11.0%, and 3.7% higher than their respective baseline rates). The effects for the other three measures are similar: 2.3% higher probability relative to the baseline for a non-final rejection, 4.8% higher for a request for extension of time, and 12.9% higher for a restriction requirement.

Turning to the work outcome measures, the effects are in the same direction as the process measures but are more modest in magnitude. Compared with the mean, applications filed at month-ends have a 4.3–6.9% greater share of examiner added citations, take about 0.9–2.2% longer to review, and are about 0.1–0.5% less likely to be approved.

Columns 5 to 8 repeat the analysis conditioning on complexity (measured as log number of claims); consistent with Hypothesis 3b, we find lower month-end work quality in this case also.

The bottom panel of Figure 3 presents a subset of these results in a graphical format. Specifically, the figures present a predicted fractional polynomial fit of the residuals of the dependent variable (de-meaned of firm fixed effects) on working day of the month. It is clear from the figures that there is a discernible change in work quality as we move from the first working day of a month to the last. The magnitude of the differences is similar to that documented in Table 3.

Together, these results clearly support Hypotheses 3a and 3b, and strongly suggest that the work quality of the applications completed on the last few days of the month is distinctly lower. Interestingly, the effect on the immediate work process is larger than on the longer-term work outcomes, which is consistent with firms expending additional effort on addressing errors in month-end applications during the prosecution process in order to avoid longer-term negative outcomes.

Heterogeneity in the Deadline Effects

Hypothesis 4 relates clustering patterns to heterogeneity in size across firms (with larger firms more likely to incur fixed costs associated with deadline routines) and speed of technology change across industries (with sectors with shorter technology cycle time more likely to provide incentives for rushing by using deadlines).

Size-related Heterogeneity. The first row of Table 4 (columns 1 to 3) presents the coefficients on log firm-year size (columns 1-4) from a regression of a month-end dummy (D1, D3, or D5) on these variables, and application year and month fixed effects. The positive and strongly statistically significant coefficient in all three cases reveals that clustering is indeed strongly positively correlated with firm size. In the rest of the rows, we regress measures of work complexity and quality on an interaction of the month-end dummy with firm-year size, the corresponding direct terms, and application year and month fixed effects. Work complexity, as measured by claims and citations, at month-ends shows a strong positive correlation with firm size. Similarly, work process quality at month-ends is also distinctly lower for larger firms. We do not find any strong interaction effects on the work outcome measures, perhaps due to the larger resources available to large firms to address errors in the work process. Taken together, our results support the use (and/or intensity) of month-end deadlines increasing in firm size.

Speed of Technological Change-related Heterogeneity. We adopt the same approach as above, except that we use "technology cycle time" (instead of firm size), defined as the mean time gap between a patent application and its backward citations. This measure has been used in the literature to represent the speed of technological changes in a field (e.g., Balasubramanian and Lee, 2008). The results, presented in columns 5 to 7 of Table 4, show that technology cycle time is indeed significantly negatively correlated with month-end clustering (row 1). This strongly supports our hypothesis that technologies with shorter lead times are more likely to exhibit clustering. Results in rows 2 to 4 show that, consistent with our hypothesis, complexity by all three measures (claims, citations, and renewal probability) for month-end filings appears to be negatively correlated with cycle time (though the results are not statistically significant for renewal probability).

Rows 5–10 also reveal that work quality reduction is more modest in industries with longer cycle times (or equivalently that work quality at month-ends is relatively worse in industries with faster innovation), consistent with our hypothesis. Taken together, these results are consistent with the speed of innovation increasing the incentive to engage in rushing behavior.

Because the expected sources of heterogeneity are related to the expected pervasiveness or stringency of deadline penalties, the consistency of the results for clustering with those for complexity and quality provides strong additional evidence that all of the effects — greater clustering, higher task complexity, and lower work quality at month-ends — are indeed driven by month-end deadlines.

Robustness Checks

We performed numerous checks based on our primary and secondary data to reinforce the conclusion that clustering patterns are caused by deadlines. Details are available on request. Briefly, evidence from interviews and surveys of practitioners strongly support this conclusion (e.g., not a single one of the 80 respondents in our interviews and surveys mentioned a non-deadline related reason for the clustering pattern; respondents who stated missing deadlines was more costly exhibited stronger clustering patterns) as does the analysis of fiscal year changes, which show a corresponding change in the patent filing patterns. Further, the surveys and interviews ruled out inventors being the primary cause of clustering (e.g., there is a large time gap between inventor disclosure and filing). Additional analyses also ruled out that month-end clustering is due to economizing on common costs, due solely to priority date considerations, and due to special applications such as provisionals and continuations.

BENEFITS FROM RUSHING: A CALIBRATION EXERCISE

Our empirical analysis shows degradation of work quality associated with rushing around deadlines. However, as discussed in our theoretical framework (and related model in the Appendix), the principal may reap offsetting benefits from the acceleration of work. Since the data do not allow us to directly observe potential benefits of rushing, we adopt an alternative approach by calibrating the theoretical model to our data. This allows us to estimate two sources of benefits: days saved, and days of un-billed balances saved, and undertake counterfactual exercises altering both the strength of, and the intervals between, the deadline penalty.

To calibrate, we impose additional structure on the baseline model and calibrate our γ parameter (which proxies intensity of the deadline penalty) to a value that yields a last-day share equal to that in our sample of applications. We standardize this value of γ to be $\gamma_s = 1$. In results available on request, we show that our calibration replicates the empirical results well. Next, we estimate the same penalty parameter (γ_s) for each of the top 200 patenting firms in our sample so that the last-day share from the model matches the average last day share for each of these firms within a tolerance of 0.005%, subject to $\gamma > 0$. We then use Lowess smoothing to non-parametrically map the associated model predictions regarding additional error at the month-end to the actual mean additional error rates (as measured by propensity for "Application Incomplete" notices) of these firms. This mapping, along with the set of γ_s, forms the basis for our exercise.

Table 5 presents the results of our counterfactual analyses. Unless stated otherwise, the table presents standardized values to make comparison across different scenarios meaningful. We present four statistics for each scenario: the standardized period-end share (i.e., the percentage of period-end filings that are accelerated), the percentage error differential, the percentage duration reduction, and the percentage reduction in days of unbilled balances for accelerated tasks relative to other filings.

Focusing on task acceleration, an aspect that cannot be directly observed from the data even for the baseline case, the calibrated model suggests that the average task acceleration of rushed tasks in the sample is about 1.5 days, or about 1.8% of the time taken to complete the task, were it not accelerated. The extent of acceleration, though increasing in penalty costs, remains small across all the scenarios. Even in the high-penalty scenario, acceleration of rushed tasks is never more than 2.2%. Hence, our calibration exercise suggests that benefits in workdays saved from task acceleration is fairly small.

However, even the small acceleration in duration leads to substantial savings in working capital (days of unbilled balances), as this pushes accelerated jobs to be billed one full cycle earlier than they would be otherwise. Another key result is that the working capital savings move proportionately with the deadline intervals (assumed here to be the billing cycle length). This is because, for a job that gets completed just after a billing deadline, the working capital costs are larger if the billing intervals are longer. The counterfactual simulations show that gains (from working capital savings) are muted, and costs (from errors) are larger when the billing intervals are shorter, so that smaller penalties are likely to be more optimal when inter-billing intervals are shorter.

DISCUSSION AND CONCLUSION

Deadlines are ubiquitous, and a particularly critical tool for managers to optimize time allocation of valuable human capital among tasks. Not surprisingly, deadlines have been shown to have real effects, affecting work flows in many contexts (e.g., Carpenter et al., 2012, Asch, 1990). What has remained ambiguous is how the time pressure related to deadlines influences performance quality (e.g., Cadena et al., 2011, Karau and Kelly, 1992, Kelly and McGrath, 1985). What has also remained unexamined empirically is how deadlines influence sorting of tasks for completion. Our study offers robust evidence for a substantial decline in work process quality associated with the clustering of work flows around deadlines. We are also the first, to our knowledge, to document strong sorting of more-complex tasks close to deadlines.

We do so by examining filing of patent applications, a context that offers an opportunity to study a single, relatively uniform process across thousands of firms over a long period of time, and allows us to construct numerous measures of task complexity and work quality. We bring to bear three different methods—formal mathematical modeling combined with a calibration exercise, large-scale statistical analysis, and primary data collection and analyses—on our research question. The insights from these methods complement each other and provide a well-rounded picture of the causes and consequences of deadlines in workplaces. The mathematical model allows us to develop precise hypotheses linking deadline-related time pressure, work flows, task sorting and work quality. Conventional statistical analyses help test these hypotheses, providing us a view of the "costs" of deadline-related time pressure. The mathematical model is also realistic enough to allow us to calibrate key parameters to the data. This allows us to show that the "benefits" from accelerating tasks could indeed be substantial (in terms of savings of working capital), which can be contrasted with the costs due to errors from such acceleration to meet the deadline. Underpinning all of this are insights from our several interviews and surveys.

We find economically significant clustering of patent applications at month-ends. The excess filing at month-ends accounts, on average, for nearly a third (31.4%) of all applications filed on the last working day of the month. Several pieces of evidence strongly indicate that this clustering is due to deadlines. To reiterate our findings on clustering, information from surveys of inventors, law firm and corporate attorneys, and detailed interviews with patent attorneys and inventors directly indicate that routine-related deadlines and associated penalties are the most likely cause of the clustering patterns. Several tests (based on these surveys), using information about frequency of deadlines, concerns about penalties for missing deadlines, and concerns about quality, confirm that clustering patterns are correlated with indicated deadlines, stronger for firms where missing deadlines are perceived badly, and lower where attorneys express concerns about quality perceptions. Several other tests reinforce a strong role for corporate deadlines. Applications by individual assignees, in strong contrast to corporate assignees, generally do not show sharp spikes near month-ends. The propensity to file during month-ends is significantly higher among firms with higher patent volume. Changes in fiscal year-ends are systematically correlated with changes in the clustering patterns of patent filings. We also verified that month-end clustering is unlikely to be driven by the work behavior of inventors, by the incentive to economize on costs by bunching tasks, solely for the seeking of earlier priority dates, or due to the use of non-standard applications such as continuing and provisional applications. Lastly, the results on effect-heterogeneity (Table 4) also provide strong support for the role of deadlines.

Consistent with our hypothesis, we find that clustering of work flows is *not* driven by marginal applications; rather, applications filed near month-ends are of higher complexity. Based on Table 3, applications filed near month-ends have more claims (0.8–1.7%), contain more drawings (0.3–1.1%), and use more drawing sheets (0.4–1.4%). Conditional on approval, the five-year citation count for patents filed near month-ends is also higher (0.4–1.5%), as is the probability of renewal

(0.2–0.3%). These estimates are differences in averages. The effects are substantially higher when we compare rushed filings with those that are not rushed. For instance, the coefficient on claims can be interpreted as a 4.5–6.2% difference between the number of claims in applications influenced by the month-end deadline and other applications.[3] Similarly, patents influenced by the month-end deadlines have 3.8–5.6% more citations, and are 1.0–2.1% more likely to be renewed than other patents. This finding is particularly relevant because, in many modern organizations, workers and work groups often simultaneously undertake multiple tasks of varying complexity, and managers are required to synchronize/coordinate across different tasks and groups. This finding as well as the calibration results suggest that deadlines could have a bigger impact on completion times of more-complex tasks that involve longer time frames.

Turning to the managerial implications of our study, our findings strongly suggest that deadlines can be costly. We find that work quality is indeed significantly lower around month-ends, both unconditionally and conditional on complexity of the application. The effect is stronger for measures that likely reflect errors in the last stage of the filing process. For instance, based on Table 3, applications filed closer to month-ends are about 3.0–4.5% more likely to be considered incomplete by the patent examiners. In relative terms, applications rushed to meet the month-end deadline are 15.5–24.7% more likely to receive such a notice than those that are not. This negative effect on work quality is consistently evident across different measures of short-term consequences, with varying magnitudes of effect. Furthermore, our results suggest deadlines may affect process–related costs and outcome–related costs differently. The impact on longer-term work-outcome measures is smaller; applications filed on month-ends have a 0.25–0.41% lower share of examiner added cites, tend to be approved 0.14–0.46% less often than those filed during other days, and take 1.0–2.4% longer to review. These correspond to a difference of 0.5–3.7% lower rate of approval and 7.7–8.3% longer review time (or about 60 days) between filings rushed to meet the month-end deadline and other filings. The lower magnitudes for longer-term measures suggest that firms are able to fix larger work quality issues reflected in the work-process measures in such a way as to

[3]((9.434/1000 + 0.207)/0.207)-1=4.5%.

contain longer-term fallout. Being able to obtain such a nuanced view of the costs of deadlines is certainly an advantage of our context.

A rough estimation exercise suggests that the monetary cost of the work process errors from rushed filing can be sizable (Appendix Table A1). The estimated extra cost due to rushed filing is between \$614 and \$782 per rushed patent, or about 7–8% of the initial filing cost. This estimate is likely to be conservative, as it only considers the direct cost of filing errors, ignoring other, potentially bigger costs due to longer review duration and a lower approval probability due to such errors. Thus, this exercise suggests that the productivity loss associated with deadline-related time pressure is not trivial.

In the context of the inverted-U relation between work quality and time pressure posited in the literature (e.g., Bluedorn and Denhardt, 1988), our results suggest that agents in our setting are not likely operating in a zone of "too low" time pressure, where increasing time pressure through deadlines would improve work quality. Rather, they are more likely to be in a zone of "too high" time pressure, where increasing time pressure reduces work quality. This is likely, in part, due to the nature of our data, which pertain to real firms in moderate to intensely competitive sectors rather than to students in a college or in a controlled experiment. In our context, the pressure to innovate is likely to be high enough that firms are already using their human capital close to capacity. We speculate that this is likely to be the case for a significant majority of firms in the private sector, where entry is largely unregulated—that Carpenter et al. (2012) find strikingly large work quality effects when examining deadlines at a government agency is not entirely inconsistent with this conjecture. Thus, our theoretical framework and empirical results are likely to apply more broadly than just within the narrow context of patent filings.

Although the preceding discussion focuses on the costs of deadlines, we do not intend to suggest that deadlines are unambiguously bad. Indeed, our theoretical framework, formal model, and the calibration exercise all specifically incorporate potential gains from deadlines. In our theoretical framework and model, firms impose deadlines because they perceive some benefits from doing so. In line with this, we find larger deadline-related effects for larger firms (for whom the benefits are likely to outweigh the fixed costs of deadline monitoring) and in sectors with faster technological changes (where the principal is likely to perceive stronger benefits from obtaining an earlier priority date than the agent). Further, imposing deadlines increases the quantity of output by accelerating task completion. Our calibration exercise suggests, within the bounds of its assumptions, that this acceleration is likely to be small (about two days or 1.8%) in our context. On the other hand, from the patent attorney's standpoint, deadlines may be more beneficial; our exercise suggests that working capital savings from billing earlier can be relatively large (around 30% reduction in days of unbilled balances).

In conclusion, our study reiterates the need to carefully evaluate the gains from reduced task duration against costs from poorer work quality when designing and implementing routines with deadlines. One way to understand the trade-offs and arrive at an optimal design would be to experiment with alternative deadline periods and the strength of incentives. We attempted something along these lines in our calibration exercise, and find that increasing billing frequency while reducing deadline penalties could reduce error rates while moderating working capital pressures. This suggests that separating patent-related deadlines from other organizational deadlines such as accounting, billing, or reporting deadlines may be beneficial to patenting outcomes. In our context, reexamining contractual arrangements with law firms (e.g., pay-per-filing vs. hourly billing), along with the routine-generated deadlines may also be useful. Finally, survey evidence suggests that lawyers who perceive negative consequences from sloppy work tend to cluster work less. So, managers could consider tracking such work process measures and using these as part of key performance indicators for agents to improve outcomes.

REFERENCES AVAILABLE ON REQUEST

Figure 1: Model Framework – Causes and Consequences of Deadlines

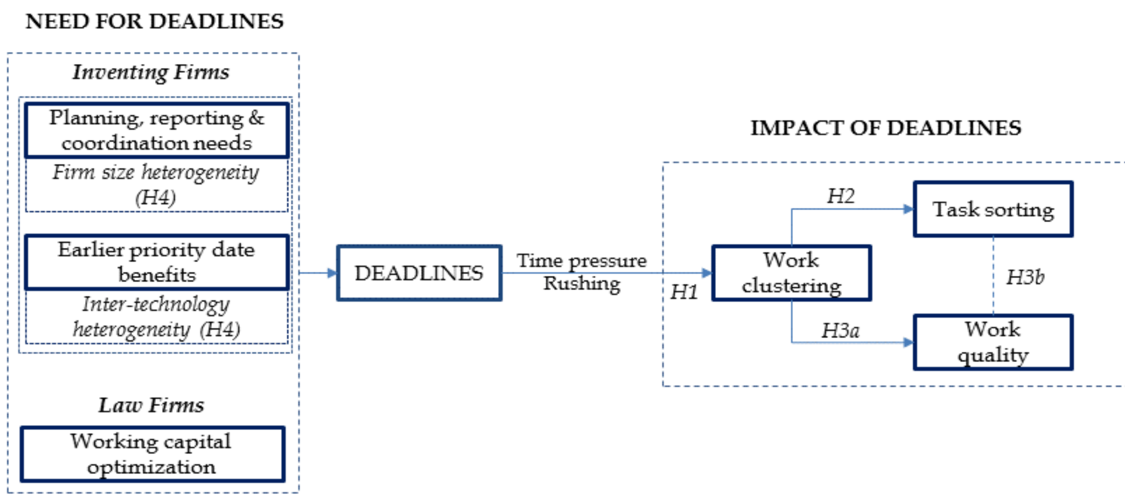

Figure 2: Patent Filings Day Share of Year Total

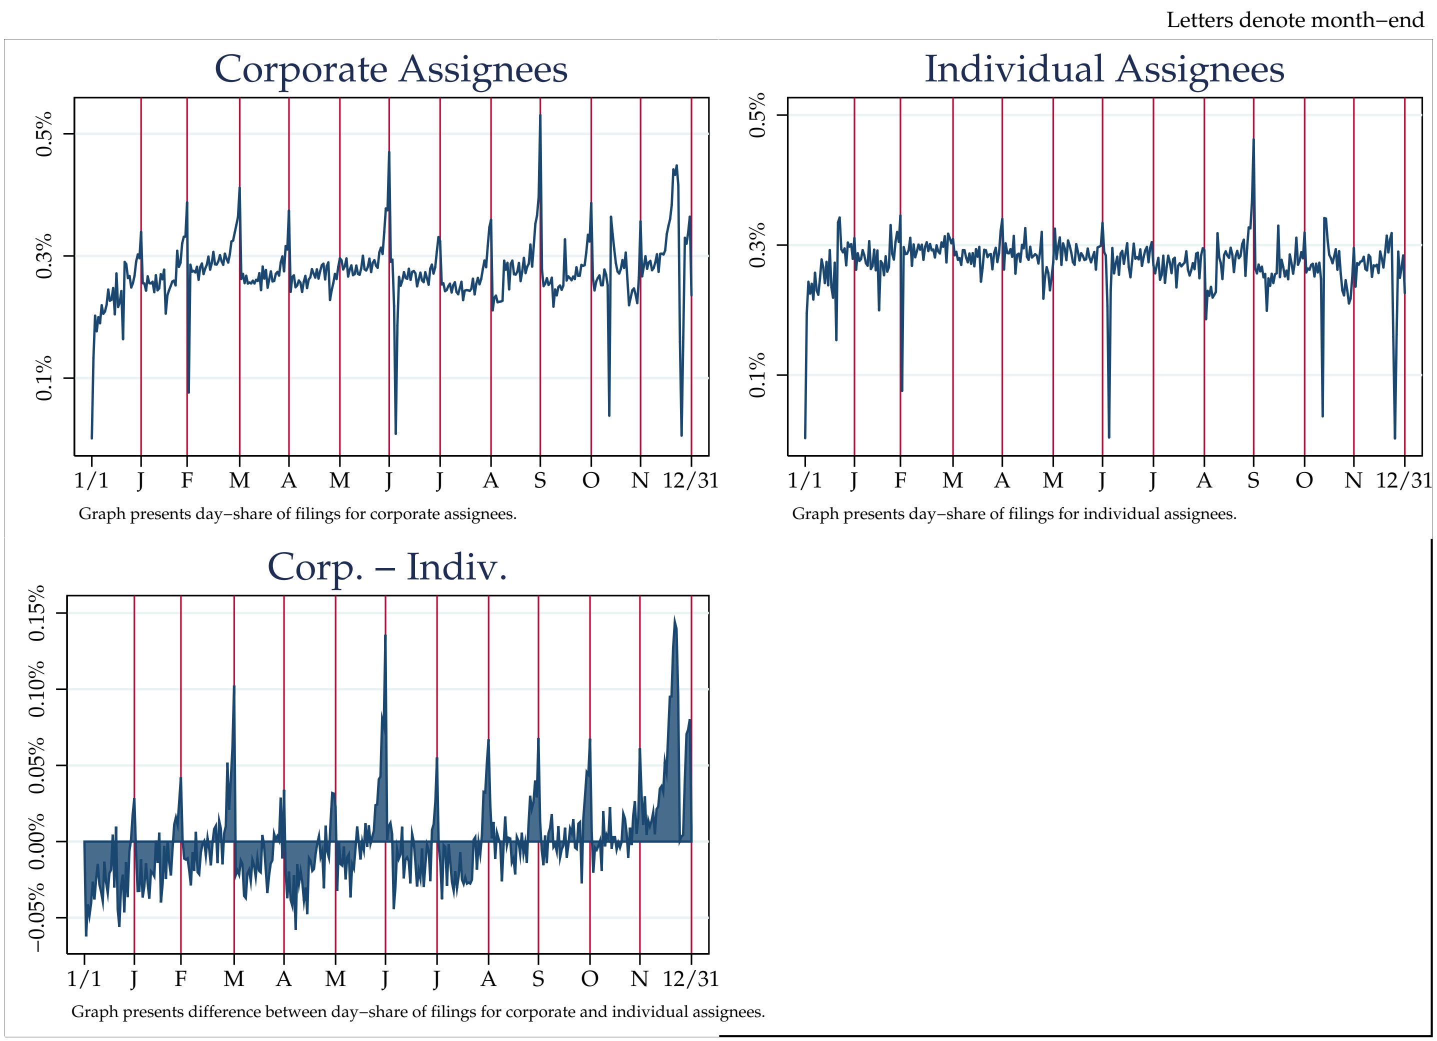

Figure 3: Month-ends, Application Complexity, and Application Quality

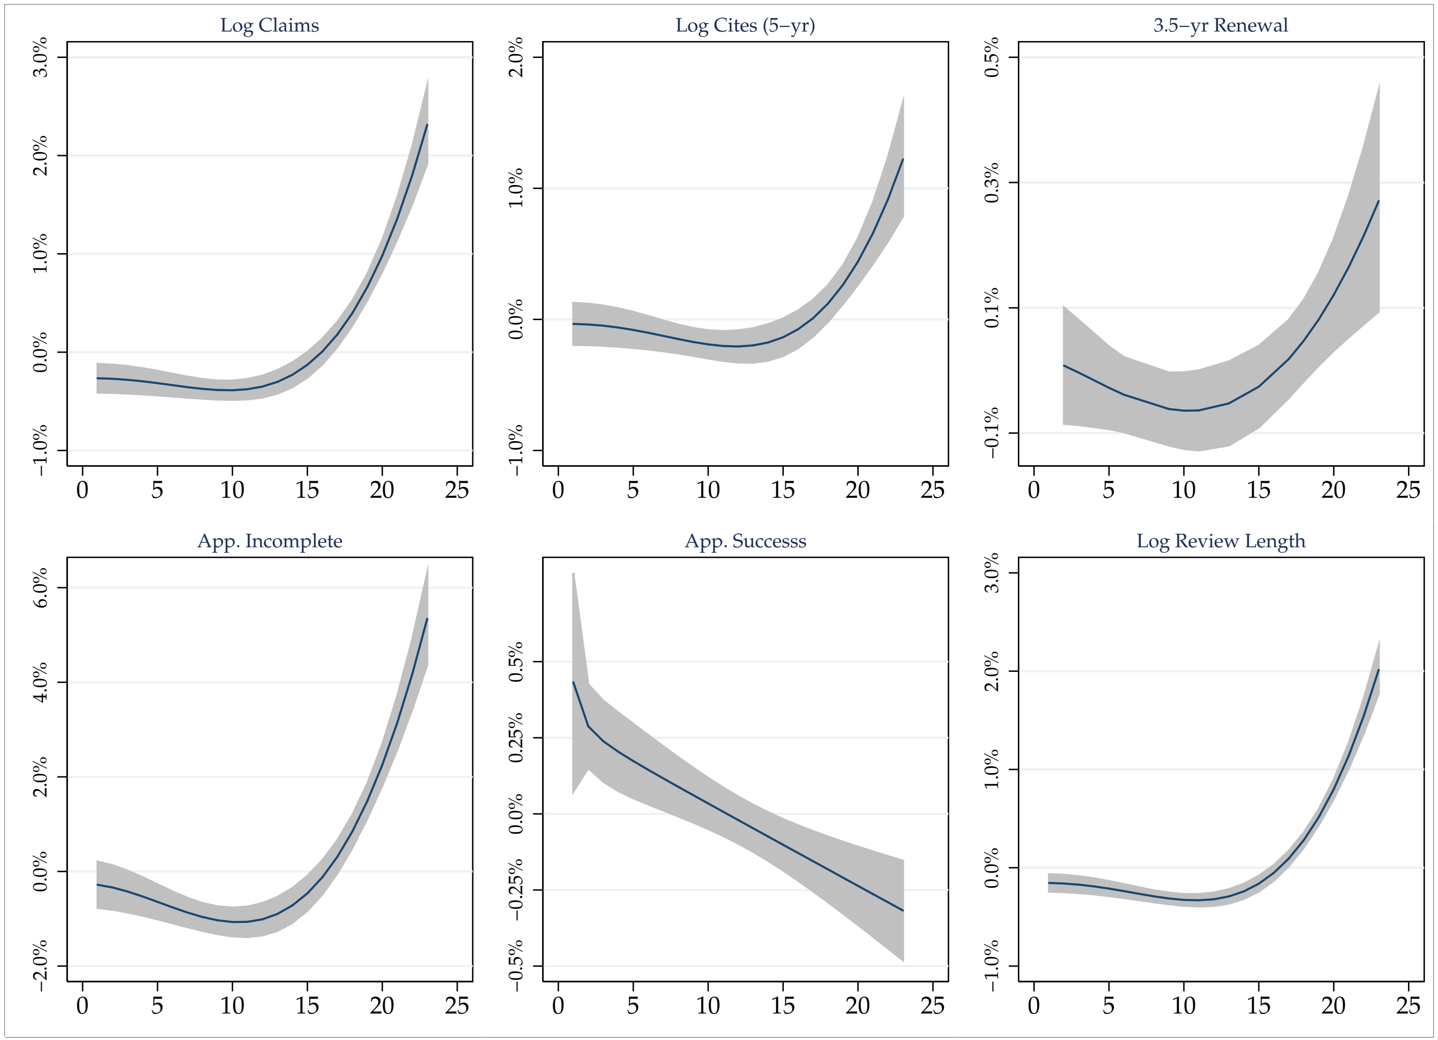

Table 1: Work Flow Clustering Around Month-ends

	Corporates			Individuals			Corporate – Individual		
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Days 1-7	-0.375**	-0.214*	-0.053	-0.120	-0.048	0.031	-0.256**	-0.166**	-0.084*
	(0.091)	(0.097)	(0.126)	(0.079)	(0.090)	(0.125)	(0.039)	(0.039)	(0.039)
Days 8-15	-0.283**	-0.122	0.040	-0.029	0.043	0.122	-0.254**	-0.165**	-0.082*
	(0.073)	(0.080)	(0.114)	(0.063)	(0.076)	(0.116)	(0.035)	(0.035)	(0.035)
Days 16-23	0.0133+	0.028	0.190	0.022	0.094	0.173	-0.016**	-0.066	0.017
	(0.076)	(0.083)	(0.116)	(0.057)	(0.071)	(0.112)	(0.046)	(0.046)	(0.046)
D1	0.055+			0.164			0.389**		
	(0.032)			(0.246)			(0.108)		
D3		0.587**			0.234*			0.353**	
		(0.140)			(0.114)			(0.058)	
D5			0.595**			0.259*			0.336**
			(0.135)			(0.126)			(0.049)
D*n*-Days 1-7	0.928**	0.802**	0.647**	0.284	0.282*	0.228**	0.645**	0.520**	0.420**
	(0.032)	(0.138)	(0.105)	(0.248)	(0.110)	(0.085)	(0.107)	(0.056)	(0.046)
N	366	366	366	366	366	366	366	366	366
D*n*-Days 1-7	0.753*	0.688**	0.554**	0.106	0.160	0.127	0.647**	0.527**	0.427**
(excl. September)	(0.313)	(0.137)	(0.106)	(0.023)	(0.105)	(0.084)	(0.116)	(0.060)	(0.050)

*Note: **: p<0.01, *: p<0.05, +: p<0.1. Robust standard errors in parentheses. Coefficients and standard errors multiplied by 1000.*

Table 2: Task-Sorting (“Complexity”) at Month-ends

Measure	D1	D3	D5	N
Log number of claims	17.354* (3.156)	9.434* (2.259)	7.953* (2.155)	1,801,602
Log number of drawings	11.170** (3.274)	5.165* (2.070)	2.864+ (1.729)	2,829,361
Log number of drawing sheets	14.444** (3.091)	7.444** (1.910)	4.469** (1.594)	2,829,361
Log number of drawings per sheet	-11.483** (2.596)	-8.051** (1.661)	-4.337** (1.396)	2,829,361
Log number of cites (5 years)	15.018* (3.018)	6.116* (1.968)	3.897+ (1.682)	2,633,488
Renewal (3.5 years)	2.823+ (1.253)	2.051+ (0.859)	2.161* (0.765)	1,938,296

*Note: **: p<0.01, *: p<0.05, +: p<0.1.*

Table 3: Work Quality around Month-ends

Measure	Mean (s.d.)	D1	D3	D5	D1	D3	D5	N
			Unconditional			**Conditional on task "complexity"**		
Work process-- measures								
Application Incomplete Notice	0.290	45.820**	36.129**	29.983**	45.203**	35.897**	29.773**	92, 533
	(0.454)	(10.737)	(7.063)	(5.550)	(10.697)	(7.050)	(5.525)	
Separate Inventor Oaths	0.262	48.661**	37.288**	32.278**	48.175**	37.105**	32.112**	92,533
	(0.440)	(8.605)	(6.208)	(4.871)	(8.578)	(6.194)	(4.866)	
Additional App. Filing Fees	0.314	46.607**	36.410**	30.267**	45.563**	36.220**	30.095**	92,533
	(0.471)	(9.403)	(6.745)	(5.405)	(9.383)	(6.738)	(5.401)	
New/Additional Drawings	0.133	14.600*	15.278**	9.535**	14.376**	15.194**	9.459**	92,533
	(0.340)	(6.109)	(3.960)	(3.520)	(6.109)	(3.955)	(3.505)	
Info. Discl. Stmt. After Initial Filing	0.594	22.210**	14.435**	14.209**	20.441**	14.209**	13.608**	92,533
	(0.491)	(7.820)	(5.502)	(4.845)	(7.810)	(4.846)	(4.859)	
Non-Final Rejection	0.834	18.840**	8.239*	5.178	18.347**	8.053*	5.009	92, 533
	(0.372)	(5.493)	(4.071)	(3.632)	(5.489)	(4.056)	(3.615)	
Request for Extension of Time	0.387	18.515*	7.948	8.980*	17.672**	7.631	8.611+	92, 533
	(0.487)	(7.798)	(5.371)	(4.447)	(7.788)	(5.363)	(4.481)	
Restriction/Election Requirement	0.149	19.175**	12.913**	16.843**	16.567**	11.933**	15.958**	92, 533
	(0.356)	(5.601)	(4.384)	(3.674)	(5.540)	(4.233)	(3.576)	
Quality Factor 1	-0.000	109.180**	83.961**	71.542**	107.791**	83.438**	71.069**	92, 533
	(0.945)	(19.413)	(13.826)	(10.93)	(19.336)	(13.793)	(10.882)	
Work outcome--measures								
Examiner Add. Cite Share	0.596	0.254	4.140+	4.095*	0.780	4.440*	4.236*	627,209
	(0.374)	(3.388)	(1.980)	(1.788)	(3.383)	(2.005)	(1.805)	
Application Approved	0.803	-1.441	-4.578*	-4.456*	-1.626	-4.672**	-4.549**	785,051
	(0.398)	(2.239)	(1.497)	(1.153)	(2.237)	(1.497)	(1.151)	
Log Review Duration	6.907	23.139**	14.246**	9.011**	20.369**	12.985**	7.817**	557,383
	(0.479)	(4.015)	(2.526)	(2.403)	(3.819)	(2.587)	(2.311)	

*Note: **: p<0.01, *: p<0.05, +: p<0.1. Errors are clustered by firm. Coefficients and standard errors multiplied by 1000.*

TABLE 4: Firm Size, Speed of Innovation and Deadline Effects

	Firm Size Interaction				Technology Cycle Time Interaction			
	D1	**D3**	**D5**	**N**	**D1**	**D3**	**D5**	**N**
Month-end clustering	2.842**	5.533**	6.542**	2,878,229	-26.690**	-35.072**	-32.279**	2,865,282
	(0.582)	(0.702)	(0.676)		(3.729)	(4.777)	(4.892)	
Log claims	7.306*	7.678**	6.006**	1,856,804	-66.849**	-47.644**	-27.366+	886,979
	(3.148)	(2.130)	(1.745)		(24.326)	(17.235)	(16.055)	
Log cites (5 yrs.)	5.726**	4.855**	4.178**	2,633,488	-12.040+	-13.218**	-8.026*	2,620,553
	(1.879)	(1.279)	(0.986)		(7.220)	(4.202)	(3.538)	
Renewal (3.5 yrs.)	0.445	0.575	0.408	1,938,295	-2.769	-3.772	-7.006	1,938,182
	(0.625)	(0.395)	(0.326)		(5.550)	(3.509)	(3.124)	
App. incomplete	10.250**	6.949**	7.046**	92,533	-108.792*	-84.390**	-50.440*	67,817
	(3.835)	(2.516)	(2.000)		(44.137)	(28.932)	(24.352)	
Sep. Inv. Oaths	5.475+	4.731*	5.677**	92,533	-58.797	-61.550*	-55.630**	67,817
	(3.154)	(2.224)	(1.724)		(36.795)	(25.159)	(19.623)	
Add. Filing Fees	2.886	4.002+	4.955**	92,533	-47.047	-47.259+	-51.404*	67,817
	(3.428)	(2.411)	(1.887)		(38.072)	(26.624)	(21.188)	
Exam. Add. Cites	1.511	0.069	0.565	627,209	-18.043+	0.622	-0.798	1,214,859
	(1.857)	(1.165)	(1.012)		(10.969)	(7.070)	(5.393)	
App. Approved	-0.066	0.813	0.721	785,159				
	(1.882)	(1.045)	(0.811)					
Log review time	3.520	2.722+	2.491*	557,493	-5.361+	-4.256*	-1.262	2,865,186
	(2.368)	(1.518)	(1.245)		(2.799)	(1.881)	(1.427)	

*Note: **: p<0.01, *: p<0.05, +: p<0.1. Errors are clustered by firm. Coefficients and standard errors multiplied by 1000.*

TABLE 5: Benefits of Deadlines: Results from the Calibration Exercise

	Std. Penalty Parameter=1				Std. Penalty Parameter=2				Std. Penalty Parameter=1/2			
	Per. End Share	Error Diff	Dur. Accel	Unbill. Accel	Per. End Share	Error Share	Dur. Accel	Unbill. Accel	Per. End Share	Error Share	Dur. Accel	Unbill. Accel
Model Calibration												
Sample Actuals	7.0%	4.1%										
Calibrated Model	7.0%	4.2%	**1.5 d**	**7.5 d**								
Standardized Values												
Sample Actuals	31.3%	13.1%										
Calibrated Model	34.7%	12.0%	**1.8%**	**31.6%**								
Predicted Counterfactuals Based on Model Calibration (Standardized Values)												
Weekly deadline (5 days)	36.4%	17.6%	1.6%	8.3%	46.9%	20.9%	2.1%	8.3%	25.5%	11.3%	1.3%	7.9%
Bi-weekly deadline (11 days)	36.4%	13.1%	1.6%	17.6%	47.3%	15.7%	2.1%	18.0%	26.3%	10.9%	1.3%	16.8%
Monthly deadline (22 days)					46.2%	13.8%	2.2%	31.9%	24.4%	6.7%	1.5%	30.9%
Semi-quarterly (33 days)	34.7%	10.0%	1.7%	41.7%	46.3%	13.5%	2.2%	41.9%	23.2%	1.9%	1.3%	41.2%
Quarterly deadline (66 days)	33.6%	10.4%	1.6%	60.7%	45.3%	12.5%	2.1%	60.3%	21.2%	2.0%	1.2%	61.7%