When Low Performance Takes Root: Deadwood and Delay at the U.S. Patent Office

Christina A. Nguyen*
Massachusetts Institute of Technology

November 1, 2025

(Please click here for the latest version)

Abstract

Organizations rely on accountability systems to maintain productivity. Yet when both sanctions for low performance and incentives for high performance are weak, employees may scale back their discretionary effort. This erosion of commitment undermines the informal norms of cooperation that support group performance. I develop and examine the concept of *deadwood*, a colloquial label for employees who shift from adequate productivity into a sustained period of low productivity but persist in the organization, often protected by tenure or other safeguards. Using administrative and personnel records from the U.S. Patent and Trademark Office and the U.S. Office of Personnel Management, I find that deadwood-heavy groups face longer processing delays among non-low performers and higher exit rates, particularly of highly educated examiners. These delays became especially pronounced after *Alice Corporation v. CLS Bank International*, an institutional shock that disrupted examination routines and exposed the vulnerability of deadwood-heavy groups. These results suggest that entrenched low performers not only depress peers' output, potentially through weakened morale, but also hinder adaptation when organizations are under pressure.

*MIT Sloan School of Management, 100 Main Street, Cambridge, MA 02142 (Email: canguyen@mit.edu). I am grateful to Pierre Azoulay and Scott Stern, as well as Jacquelyn Pless for their mentoring and support. I thank Robert Gibbons, Luca Gius, Danielle Li, David Papke, Nelson Repenning, Georg Rilinger, and participants of the Consortium on Competitiveness and Cooperation Doctoral Conference, MIT Economic Sociology Working Group, MIT Organizational Economics Lunch, MIT TIES Seminar, NBER Productivity Seminar, and AOM OMT Doctoral Student Consortium for helpful discussions and comments. I thank Deepak Hegde for sharing data and the staff of the U.S. Office of Personnel Management and the U.S. Patent and Trademark Office who helped fulfill my Freedom of Information Act data requests. Finally, I am appreciative to the many patent examiners, quality assurance specialists, administrative judges, patent attorneys, and patent agents who shared their valuable experiences and insights with me. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant Number 2141064.

1 Introduction

Sustaining employee effort over time is a fundamental challenge for organizations (Barnard 1938; Grant and Shin 2012; Khan 2025). When performance incentives are weak and sanctions for low performance are limited (Lazear 2000), even capable employees may reduce discretionary effort. These employees are colloquially referred to as *deadwood*: those who once performed adequately but whose effort and productivity have eroded while they remain in the organization. The accumulation of such employees results in organizational bloat, presenting two possible trajectories: it may erode productivity and morale (Felps, Mitchell, and Byington 2006), or it may preserve continuity and experience that organizations can draw on when conditions change. Examining the organizational impact of deadwood therefore offers insight into whether and how localized, persistent low performance may evolve into systemic inefficiency.

Although this problem is prevalent across sectors, research on declining performance remains scarce. Most empirical work has focused on how organizations elicit and reward effort (e.g., Lazear 2000; Bandiera, Barankay, and Rasul 2005) or on the positive externalities of high performers and "stars" (e.g., Azoulay, Zivin, and Wang 2010; Mas and Moretti 2009; Ichniowski and Preston 2014). High performance is visible, quantifiable, and formally acknowledged through rewards and promotions. Low performance, by contrast, is challenging to study because it is often obscured by team-based metrics, managerial discretion, or peers who compensate for weaker colleagues. Moreover, most settings and datasets capture only cross-sectional snapshots, making it difficult to observe how individual performance evolves over time. As a result, we know relatively little about how capable employees experience sustained performance declines, and how such declines affect the behavior and productivity of their peers. Because the emergence of deadwood and the spillover effects of their behavior unfold gradually, distinguishing entrenched low performance from temporary dips in effort requires longitudinal data that follow individuals within the same organization over time.

In this paper, I use longitudinal individual-level data from the U.S. Patent and Trademark Office (USPTO) to examine how deadwood examiners affect organizational outcomes. The analysis draws on two decades of linked administrative data (fiscal years 2000-2020) that combine detailed patent examination records with personnel data from the Office of Personnel Management (OPM), obtained through Freedom of Information Act requests.[1] Productivity is measured using weighted "production counts," the USPTO's primary and objective indicator of individual performance, adjusted for both technological complexity and examiner experience.[2]

Examiners' productivity typically follows a common life cycle: performance rises with experience, peaks, and eventually declines. Yet I find that not all slack is created equal. Some declines are mild and temporary, part of a normal easing after peak productivity, while others are steeper and persistent, signaling a lasting withdrawal of effort. In fact, early in their careers, some examiners who later become "deadwood" look indistinguishable from their productive peers—but over time, their performance curves diverge. To capture these dynamics, I classify examiners into three groups: deadwood employees, whose productivity declines and remains low after a period of adequacy; chronic low performers, whose productivity is consistently weak from the start; and non-low performers (neither deadwood nor chronic low performers), who maintain or improve their output. This distinction isolates low performance that reflects persistent decline rather than poor fit.

The USPTO provides an appropriate context for studying entrenched low performance, where examiners have openly remarked that "everyone knew who had a reputation for being deadwood" (IP Watchdog, July 6, 2021). Examiners perform highly structured, knowledge-intensive work governed by uniform procedures, enabling precise comparisons of productivity across individuals. At the same time, examiners are organized into technology-specific "workgroups" that evaluate applications of similar complexity, facilitating comparisons among peers performing equivalent tasks. Promotions follow a systematic process based on sustained adequate performance and tenure. The USPTO also features strong job protections and union representation that make dismissal rare, allowing performance declines to persist and be observed over long periods. Workgroups vary in composition, including seniority, experience, and the prevalence of deadwood, creating natural variation in exposure to low performers.

Many professional and public organizations share certain institutional features that allow deadwood to emerge and persist, including stable employment and procedural safeguards that make

[1] Some of these data were also shared with me by Deepak Hegde.

[2] Although patent examiners are also evaluated on other dimensions, including quality, most receive uniformly high ratings on these measures by default, making them less informative for distinguishing among examiners. The performance plan is discussed in more detail in Section 3.

termination rare. Large corporations, universities, and government agencies often combine accountability structures with protections intended to ensure fairness, but these same safeguards can inadvertently preserve low performers. Similar dynamics arise in professions where performance is difficult to quantify or evaluation relies heavily on managerial discretion. Yet in most of these settings, systematic performance data are limited or opaque, making persistent declines in effort difficult to observe. By contrast, the USPTO provides detailed, longitudinal records that capture individual productivity over time, offering a rare opportunity to examine how entrenched low performers persist and how their presence shapes peer behavior and organizational outcomes.

Several mechanisms may link deadwood to the behavior and outcomes of coworkers. First, workload redistribution can arise when peers compensate for low performers' reduced effort, boosting short-term productivity but risking overload and frustration over time (Karau and Williams 1991; Podsakoff et al. 2000). Second, social and behavioral spillovers may occur as effort norms adjust to the presence of disengaged colleagues: fairness concerns and perceptions of inequity can weaken motivation (Adams 1965; Kandel and Lazear 1992; Goerg, Himmler, and König 2024), while disengagement can spread through social contagion (Felps, Mitchell, and Byington 2006). Third, learning and knowledge-sharing mechanisms may be disrupted when low performers contribute less to collective problem solving or institutional memory (Huckman and Staats 2011). Finally, incentive structures may amplify or constrain these effects, depending on whether rewards emphasize individual or group performance (Lazear 2000; Bandiera, Barankay, and Rasul 2005). In the USPTO, where productivity is individually measured but examiners work within groups that handle shared backlogs, the social and workload mechanisms are likely to dominate.

I test these propositions using a panel design that exploits rich variation in examiner productivity and organizational structure over time. Workgroups of examiners within the same technological domain form the unit of analysis, enabling estimation of peer and group effects while holding fixed broader institutional factors.[3] I use the workgroup level as the unit of analysis because examiners can transition between performance states over time, appearing as non-low performers prior to a decline and as deadwood thereafter, allowing both sides of these peer relationships to be captured.

The dependent variable is first action pendency, the number of months that it takes to receive

[3]For organizational purposes, examiners are assigned to art units, small teams responsible for examining applications in a narrowly defined technical area, each led by a supervisor. Several related art units are grouped into a workgroup, which shares a broader technical domain and organizational oversight.

the first decision to an application.[4] Pendency is the USPTO's central organizational performance metric and proxy for responsiveness: longer pendency reflects a growing backlog of roughly 800,000 unexamined applications (U.S. Patent and Trademark Office 2025c), delaying innovation and frustrating applicants. To capture spillovers, I calculate workgroup pendency using only non-low performers, providing a clean measure of how entrenched low performers affect their peers' response time. I measure exposure to deadwood using an indicator for workgroups whose share of deadwood examiners exceeds the average across all groups. I also include a measure of chronic low performers (above-average percent of chronic low) to distinguish between the effects of different types of low performers.

I find that high-deadwood workgroups exhibit longer processing times even among their productive members, indicating negative peer spillovers. These spillovers are stronger than those associated with chronic low performers, suggesting that peers react more negatively to perceived shirking than to low ability: poor fits may evoke sympathy, while deadwood violate shared effort norms. These workgroups also experience higher exit rates, particularly among highly educated examiners with stronger outside options.

To reinforce my findings, I implement an alternative measure of deadwood, "deadwoodness," a measure that captures persistent examiner-level delays. Instead of defining deadwood based on a binary metric, this is a continuous measure that characterizes someone's likelihood to have delays in their post-peak period. The results remain consistent using this continuous specification.

While these spillovers matter in routine operations, they take on added significance under stress, when organizations depend on discretionary effort and adaptive coordination to absorb shocks and respond to change (Repenning and Sterman 2002; Kellogg 2009). Prior work suggests that weaknesses which remain hidden in routine operations often become more visible and more costly when organizations face disruption. Building on this idea, I examine the ability of workgroups to adapt to institutional change, exemplified by the U.S. Supreme Court's *Alice Corporation v. CLS Bank International* decision. The *Alice* ruling abruptly altered patentability standards in certain technologies, introducing substantial ambiguity into examination and requiring workgroups

[4] The USPTO measures first action pendency as the number of months between the filing date and the mailing of the first action. Instead of the filing date, I use the last docketed date of the application to capture the amount of time that the application is on the examiner's desk, so that it excludes the time that it takes to assign the application to an examiner.

to adapt their routines to new guidelines. Comparing affected and unaffected workgroups reveals how groups with differing concentrations of deadwood adapted to this disruption.

High-deadwood and *Alice*-affected workgroups became disproportionately slower following the court decision, suggesting that deadwood undermines adaptability under stress. Collectively, the results indicate that deadwood employees not only reduce their own output but also weaken the performance and stability of the workers around them. These effects are substantially larger than those associated with chronic low performers and hold under alternative ways of measuring deadwood.

Understanding the organizational consequences of entrenched low performers requires returning to a foundational insight: organizations run on cooperation. Formal authority and contractual incentives alone are insufficient to elicit sustained effort; performance depends on informal norms, shared purpose, and mutual accountability (Barnard 1938; McGregor 1960; Osterloh and Frey 2000). Deadwood employees violate these implicit norms. Their disengagement is not just passive but may reflect a shift toward self-interested behavior, doing only what benefits themselves while relying on others to uphold collective performance. Their continued presence signals that effort beyond minimal requirements is unnecessary, eroding the sense of fairness and collective responsibility that holds teams together. Over time, what begins as individual disengagement can evolve into a structural problem, reshaping team dynamics by undermining trust, reciprocity, and motivation (Barsade 2002; Felps, Mitchell, and Byington 2006; Rintamäki, Parker, and Spicer 2025; Usanova and Géraudel 2024).

It is widely acknowledged that deadwood employees exist and are less productive than their peers (Saez, Schoefer, and Seim 2023), yet their broader organizational consequences have been difficult to establish. I argue that deadwood are not isolated cases of individual low performance but an organizational condition that accumulates through institutional design.

The findings extend several research streams. They deepen understanding of organizational slack and inefficiency by distinguishing between forms of slack that buffer organizations and those that entrench decline (Cyert and March 1963; Repenning and Sterman 2002). They also advance research on accountability and personnel systems by showing how safeguards designed to promote fairness can inadvertently embed inefficiency. Finally, they connect these insights to theories of cooperation and adaptability, demonstrating how the erosion of effort norms undermines collective performance and amplifies fragility under stress.

More broadly, the study reorients human capital research toward the lower end of the performance distribution. Whereas most research on worker heterogeneity highlights the benefits (e.g., Azoulay, Zivin, and Wang 2010; Brox and Lechner 2024; Cohen-Zada, Dayag, and Gershoni 2023; Hendricks, Call, and Campbell 2022; Ichniowski and Preston 2014; Kamei and Ashworth 2023; Mas and Moretti 2009) and coordination challenges of strong performers (Tian et al. 2025; Zhou et al. 2025), much less is known about how persistently low performers affect their peers and organizations. Existing studies of poor-fit employees focus on short-term productivity losses (Langcawon 2025), but such cases are often transient because mismatched employees tend to exit (Bidwell 2011). Deadwood, by contrast, persist—their endurance allows them to reshape peer behavior and informal norms in ways that erode collective performance. To my knowledge, this is the first study to document the organizational costs of entrenched low performers and the magnitude of their negative spillovers on otherwise productive peers.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework for understanding deadwood employees in organizations, and Section 3 describes the institutional context of the USPTO. Section 4 introduces the data and presents descriptive evidence on the prevalence of deadwood examiners, while Section 5 outlines the empirical strategy and main results. Section 6 offers back-of-the-envelope calculations to situate the magnitude of deadwood's impact relative to the USPTO's organizational goals. Section 7 concludes by discussing the implications of these findings for managing entrenched low performance and sustaining organizational effectiveness.

2 Theoretical Framework

2.1 Conceptualizing Entrenched Low Performance

In its productive form, slack provides a buffer that enables adaptation and experimentation (Grigoriou and Rothaermel 2017; Repenning and Sterman 2002). Yet when that buffer becomes detached from high performance, it transforms from a resource into a liability. Deadwood employees represent this transition from adaptive slack to organizational bloat. Their persistence reflects the cumulative effects of stability-seeking systems, in which protections such as tenure designed to promote fairness and continuity also make it difficult to address sustained low performance (Saez, Schoefer, and Seim 2023).

Deadwood often begins as gradual disengagement rather than outright shirking. Employees may recalibrate effort in response to burnout, frustration, or shifting life priorities. Because sanctions and exits are costly, organizations often tolerate these slow declines. Over time, employees adapt through "satisficing" behaviors (Simon 1997), doing just enough to meet formal requirements while relying on others to sustain collective output. When performance management systems fail to distinguish temporary slowdown from persistent low performance, effort decay becomes self-reinforcing. Bounded rationality and institutional inertia favor preserving the status quo (March and Simon 1958), while organizational routines designed to ensure reliability further discourage risk-taking and learning. As evaluation thresholds stagnate, the meaning of acceptable performance drifts, and reduced effort becomes normalized within the organization's interpretive frame (Barnard 1938; McGregor 1960).

Deadwood thus differs fundamentally from chronic low performers. Chronic low performers stem from selection errors, where the organization hires individuals whose skills or motivation never aligned with job demands. Deadwood, by contrast, signals a retention and renewal problem, a failure of ongoing evaluation and motivational maintenance. These dynamics reflect a broader design trade-off between short-term productivity and long-term renewal. For instance, studies of production governance show that systems optimized for early efficiency often limit learning and sustained improvement (Novak and Stern 2008). Likewise, performance regimes that emphasize early achievement or rigid evaluation metrics can generate short-term gains at the cost of long-term engagement. The emergence of deadwood, therefore, is not merely a personal failure but an organizational outcome, rooted in designs that prioritize stability and immediate output over continuous renewal.

2.2 From Individual Decline to Peer Spillovers

While deadwood originates as an individual problem, it becomes an organizational one through its social and structural consequences. Cooperative work depends on mutual accountability and the willingness to exert discretionary effort (Barnard 1938). When one member disengages, the effects reverberate: peers must absorb extra workload, coordination frays, and perceptions of fairness deteriorate. These responses are not purely functional but interpretive. Peers may respond differently depending on whether low performance appears rooted in inability or in effort withdrawal:

for example, deadwood may elicit frustration and resentment, whereas chronic low performers may be met with tolerance or sympathy.

Classic equity theory predicts that visible inequity erodes motivation (Adams 1965), while peer-monitoring models suggest that cooperative norms are sustainable only when members believe others are contributing their fair share (Kandel and Lazear 1992). Empirically, even a single persistently low performer can impose disproportionate coordination and morale costs on others (Felps, Mitchell, and Byington 2006). In interdependent settings, these effects arise because effort choices are complementary Cohen-Zada, Dayag, and Gershoni (2023). Translated to organizational contexts, the presence of a disengaged coworker lowers not only total output but also the equilibrium effort level of those around them. The demotivating influence of deadwood is likely nonlinear. Groups can often compensate for occasional low contributors, but as their prevalence grows, tolerance gives way to contagion. Each additional disengaged peer increases the perception that reduced effort is acceptable, gradually shifting collective norms. Once low performance becomes normalized, even motivated employees face diminishing returns to extra effort and may redirect their energy elsewhere.

Hypothesis 1: Non-low performers in groups with higher proportions of deadwood exhibit longer delays, with a stronger negative effect from deadwood than chronic low performers.

As burdens accumulate and reciprocity erodes, high performers reassess their attachment to the organization. When perceived fairness and future prospects deteriorate, the calculus of staying shifts toward exit. Classic turnover models suggest that declining efficacy and trust in the system lead employees to withdraw cognitively and, eventually, physically (March and Simon 1958; Podsakoff et al. 2000).

Hypothesis 2: Non-low performers in high-deadwood groups are more likely to exit the organization.

2.3 From Demotivation to Organizational Inertia

The persistence of deadwood not only depresses current performance but may also limits how groups adapt when facing change. Adaptation requires more than formal routines—it depends on members' discretionary willingness to modify those routines and invest effort in collective problem solving (Barnard 1938; Hannan and Freeman 1984). When reciprocity and trust have eroded, that willingness falters. What was once a reserve of flexible slack becomes inert capacity, available in principle but inaccessible in practice.

Research on sensemaking and relational coordination shows that organizational resilience depends on shared understanding, open communication, and mutual respect (Weick 1995; Gittell 2002). Deadwood weakens these relational foundations. Their disengagement signals that effort is optional, discouraging others from mobilizing when conditions shift. Over time, groups characterized by persistent disengagement display a form of dynamic rigidity: they maintain existing routines even as external demands change.

Managerial discretion can also moderate these dynamics. Proactive supervisors can buffer teams by reallocating tasks, clarifying expectations, and re-establishing effort norms. Passive or overloaded managers, by contrast, may allow low performance to spread unchecked, further entrenching inertia.

> **Hypothesis 3:** High-deadwood workgroups will exhibit lower adaptability in response to an institutional shock.

3 The U.S. Patent and Trademark Office

3.1 The Patent Examination Process

Understanding how examiners handle applications and make decisions provides important context for evaluating their performance and the organizational norms that develop. When a patent application is submitted to the USPTO, it arrives at the Office of Patent Application Processing (OPAP), where it receives a primary United States Patent Classification (USPC) based on the subject matter of its most comprehensive claim (Lu, Myers, and Beliveau 2017).[5] The application is then routed to the appropriate technology center and art unit, a group of examiners with expertise in related technological fields. Each art unit is overseen by a Supervisory Patent Examiner (SPE),

[5]The USPTO switched from using the United States Patent Classification (USPC) system to the Cooperative Patent Classification (CPC) system on January 1, 2013. The CPC system is a classification system developed with the European Patent Office. However, until 2022, the USPTO continued using the USPC system for assigning applications to examiners (Lu, Myers, and Beliveau 2017).

who assigns applications to individual examiners within the unit. Assignments are nominally random but can be adjusted to account for examiner workload, in order to balance caseloads and avoid excessive backlogs (U.S. Government Accountability Office 2016). Importantly, assignments are not intended to reflect the examiner's skill or the application's complexity, to ensure fairness and consistency (Lemley and Sampat 2012).[6]

The examiner begins by assessing whether the application claims more than one distinct invention. If so, the examiner may issue a restriction requirement, instructing the applicant to choose a single invention for examination ("prosecution on the merits"). The examiner then evaluates the selected claims for compliance with statutory requirements and searches for prior art (existing patents, publications, products, or other publicly known knowledge) to assess whether the invention is novel, non-obvious, and useful. If something similar already exists, there are grounds to reject the application.

The examiner then issues a First Action on the Merits (FAOM), which may either allow the application or reject some or all of the claims (non-final rejection). Of all filed applications, 86.4% receive a non-final rejection at this stage (Carley, Hegde, and Marco 2014). Applicants may respond with amendments, arguments, or by requesting an interview to discuss potential changes. Based on the response, the examiner may allow the application or issue a second office action, which may be another non-final or a final rejection. Unless the second office action raises new rejections not prompted by the applicant's amendments or late-submitted references, it is generally designated "final," limiting further amendments.

If the examiner ultimately determines that the application satisfies all legal criteria, the patent is granted. Otherwise, the applicant may still file a response after a final rejection, but the examiner is not obligated to accept it, especially if it would require a new prior art search or if amendments could have been made earlier. In such cases, the examiner issues an advisory action explaining why the application remains unallowable. At that point, the applicant may appeal, file a continuation application, or abandon the application.[7] These actions conclude the prosecution round, with a "disposal" if the application is allowed, abandoned, or appealed. If the applicant appeals an exam-

[6]Since 2022, the USPTO has automated the assignment process, incorporating workload and backlog considerations while preserving quasi-random allocation (U.S. Patent and Trademark Office 2022).

[7]Abandonment may result from the examination outcome or from external business considerations, such as shifting company priorities (Lemley and Sampat 2008).

iner's rejection (typically when they believe the rejection was unwarranted or overly stringent), the application goes to the Patent Trial and Appeal Board (PTAB) for a decision by an administrative patent judge who either affirms or reverses the examiner's decision. If the applicant files a Request for Continued Examination (RCE), a new round of prosecution begins.

3.2 Pendency is an Institutional Performance Priority

Rising application volume and complexity have contributed to a growing backlog, which refers to the queue of applications that remain unexamined or exceed the agency's processing capacity (Mitra-Kahn et al. 2013). It may also be exacerbated by persistent low performance among some examiners, whose low productivity further limits the agency's ability to keep pace with demand. Because of the backlog, reducing pendency (the time it takes for an application to receive a decision) has been a longstanding priority for the USPTO. Shorter pendency is important not only for applicants, who rely on timely decisions to make investment and commercialization choices, but also for the broader innovation system, where delays can create uncertainty and hinder technological diffusion. In response, the USPTO has built its internal performance management system, including production goals and examiner evaluation criteria, around the need to reduce pendency and improve throughput.

3.3 Patent Examiner Performance Appraisal Plan

To operationalize these performance priorities, the USPTO evaluates examiners using a structured Performance Appraisal Plan (PAP). Other than the one-year probationary period, examiners receive an overall annual performance rating based on a weighted average of their performance in four elements: production (35%), quality (35%), docket management (20%), and stakeholder interaction (10%).[8]

Production is a quota requirement that reflects the volume of office actions—formal written communications issued by the examiner that outline the legal and technical reasoning behind rejections, objections, or allowances—completed within the evaluation period. Docket management measures adherence to internal timeliness goals; specifically, it tracks whether the examiner takes

[8] In 2021, the weights were revised to 30% production, 30% quality, 30% docket management, and 10% stakeholder interaction. Stakeholder interaction also switched from being a non-critical element to a critical one.

timely action on assigned cases between applicant responses and is distinct from pendency, which captures the total time from filing to either first action or disposal. Quality assesses compliance with established quality-related procedures and review criteria. For example, there are Quality Assurance Specialists (QASes) who examine a sample of office actions to ensure that rejections and allowances are properly supported and procedurally correct. That said, this approach is insufficient because there are fewer QASes than there are examiners, and they stop reviewing once they meet their quota (former patent examiner, personal interview, June 5, 2024). The quality element is discussed further in Section 2.4. Stakeholder interaction, which accounts for the smallest share of the evaluation and is not considered a critical element, pertains to customer service and communication. For each of these elements, the examiner is given a rating of Outstanding, Commendable, Fully Successful, Marginal, or Unsatisfactory. Examiners generally need to achieve at least a rating of Fully Successful on each critical element to meet expectations.

The SPE assigns the examiner an overall rating, taking the weights into consideration. An examiner's overall rating can only be as high as her lowest rating among the critical elements. If performance in any critical element falls below Fully Successful, the overall rating is capped at that level, regardless of performance in the other areas. Production is a stringent measure of output and is where persistent low performance is most likely to appear.

3.4 "Production is the Only Thing That Matters"

Production is the primary element of examiner performance in the USPTO, carrying the greatest institutional weight and financial incentive (Patent Office Professional Association 2025). Examiners report that organizational pressure overwhelmingly focuses on meeting production quotas. As one examiner put it concisely, "production is the only thing that matters" (personal interview, June 5, 2024). Another explained, "We are a government agency accountable to the American people to do our job and examine applications in a timely manner...without quotas, innovation suffers, as it would take five or more years to even get a patent. By that time[,] they are useless as the technology has moved on" (Reddit, November 2024).

The USPTO treats low production as unequivocal evidence of poor performance, whereas high production (even when accompanied by lower quality) is often tolerated. As one administrative patent judge testified during a hearing on oversight, accountability, and quality, "a patent examiner never got fired for doing bad quality work, as long as they did a lot of it" (*Abuse of the USPTO's Telework Program: Ensuring Oversight, Accountability, and Quality* 2013). While quality still matters, the USPTO views production as the most meaningful and reliable indicator of persistent low performance, including deadwood and chronically low-performing examiners. Indeed, in the early 2000s, the then-Commissioner for Patents "threw quality out the window" and "categorized examiners as A, B, and C examiners" based almost completely on production count (former patent examiner, personal interview, July 23, 2025). I also investigate the behavior of deadwood examiners, some aspects of which may manifest as quality-related issues.

While detailed data on quality (i.e., error rates as determined by SPEs and QASes) and stakeholder interaction are unavailable, my interviews indicate that these two elements are often treated as formalities and exhibit little meaningful variation. Although production and quality each carry equal weight in the PAP, quality ratings are subjectively assigned by SPEs and tend to cluster at Fully Successful or above,[9] limiting their usefulness for identifying persistent low performers. As one examiner described it:

> The patent office's assurance quality tickets are like speeding tickets. The penalty of getting caught is not high enough to keep you from speeding. There are so few of them and so many things to do, they only have so much time to do a deep dive. It's like a slap on the wrist (former patent examiner, personal interview, June 6, 2024).

Similarly, one patent examiner noted, "The USPTO only cares about your numbers and that you maintain a passable minimum quality" (Reddit, December 2024).

Docket management measures whether examiners meet minimum timeliness requirements. Unlike production, which only credits substantive actions that move an application toward a decision, docket management can often be satisfied through non-substantive actions, such as issuing procedural notices or routine communications, that reset deadlines without materially advancing the case. Interviews suggest that most examiners meet these requirements without difficulty, and exceeding them offers minimal reward. As one examiner explained about their first year, "There is no production requirement the first year, but the SPEs will fudge quality so they meet up with production" (personal interview, June 4, 2024). Accordingly, while important for maintaining workflow, docket

[9] SPEs do not necessarily review all of an examiner's work or even a random sample of it, and often assign Fully Successful ratings even when errors are identified (U.S. Government Accountability Office 2025).

management exhibits less variation and institutional emphasis than production.

Stakeholder interaction is a minor, perfunctory component of the appraisal, with ratings typically defaulted to passing scores. The USPTO considers it as a non-critical element (whereas the other three are critical), and examiners themselves also view it as "not important" for performance reviews, though some enjoy having the interactions (personal interview, May 25, 2024).

Although the USPTO presents the PAP as a comprehensive system for accountability, in practice it focuses almost entirely on production, while the other elements function largely as formalities. This imbalance echoes March and Olsen's "garbage can" view of organizations, in which multiple formal rules coexist but many are irrelevant to actual decision-making, creating the appearance of thorough evaluation without its substance (March and Simon 1958).

3.5 Examiner Production Count System

Examiners receive a production rating based on the extent to which they meet or exceed their assigned production goals:

Rating	**Achievement of Goal**
Outstanding	more than 110% of goal
Commendable	105%-109% of goal
Fully Successful	95%-104%
Marginal	90%-94%
Unacceptable	less than 90% of goal

The production goal that examiners are expected to meet is calibrated according to their General Schedule (GS) grade level (proxy for experience and seniority) and the complexity of the technology they examine (assigned art unit). Junior examiners typically enter at GS-5 or GS-7 (depending on their education level and relevant prior experience) and may advance to GS-14 based on a combination of time in service and performance evaluations. A key milestone in this progression is promotion to primary examiner at GS-14, which grants the authority to sign office actions independently without requiring supervisory review. Primary examiners are responsible for training newer examiners and may also review work by juniors. Advancement beyond GS-14 generally requires taking on management responsibilities, most commonly through promotion to SPE. SPEs oversee a group of examiners within an art unit, approve junior examiners' work, conduct performance reviews, and provide technical and procedural guidance; they no longer have formal patent examining responsibilities and requirements.

Each GS level corresponds to a specific expected output measured in counts (explained in detail in Data and Methods section), which are weighted based on the type of action taken and the time allotted to complete it. More complex technologies carry lower production targets to account for their increased difficulty. Thus, an examiner's production goal is tailored to both their GS level and the technical complexity of their docket. These GS levels also determine base salary, with higher grades receiving greater pay and increased expectations.

4 Data and Descriptive Statistics

To identify low-performing patent examiners (deadwood and chronic low performers) and examine their impact on organizational outcomes, I combine two primary data sources: (1) patent application data from the USPTO, and (2) internal personnel and employment data on patent examiners from the USPTO and OPM, obtained through FOIA requests.

The application-level data span all patent applications filed between 2000 and 2020. These transaction-level records include detailed information on the timing, content, and outcomes of each application, documenting examiner actions, applicant responses, and procedural milestones throughout the examination process. From these data, I construct key measures of examiner performance (production count) and efficiency (first action pendency).

The examiner-level personnel data include administrative records on examiners' technology areas (art units and workgroups) and employment characteristics (e.g., hiring date, GS level, education, salary). Due to changes in data privacy policies, OPM was unable to fulfill FOIA requests for age data beyond the second quarter of 2014. As a result, age is available for only 53.2% of examiner–fiscal years. Missing data on the other variables is minimal, with less than 0.3% missing on any given measure.

Linking the application-level records to examiner identifiers enables the construction of longitudinal examiner-level performance measures and organizational aggregates. I classify deadwood and chronic low performers at the examiner-fiscal year level, which facilitates tracking longitudinal performance patterns and identifying sustained low performance. I then sum the number of low performers for organization-level analyses.

4.1 Workgroup as the Organizational Unit of Analysis

Although most prior studies group examiners at the art unit level, I aggregate to the broader workgroup level to better capture organizational dynamics. Workgroups encompass multiple related art units within a Technology Center. For example, Art Units 3691 to 3696 comprise the 3690 or Finance Workgroup (U.S. Patent and Trademark Office 2025b), sharing a technical domain and managerial oversight. Compared to art units, workgroups provide larger and more stable group sizes, reducing random variation and measurement error in group-level measures. Examiner assignments also tend to be more stable at this level, as individuals typically remain within the same workgroup even when reassigned between art units to accommodate changing organizational needs or to better match their expertise with the subject matter. Finally, workgroups reflect a coherent technical and managerial environment that supports group-level norms and organizational dynamics more effectively than the narrower art unit level (Figure B.1). To account for remaining differences in workgroup composition over time, I control for workgroup size and age and include fixed effects.

4.2 Outcome Measures

The USPTO reports two standard measures of pendency: first action pendency (filing to mailing of first office action) and traditional pendency (filing to mailing of final disposition). I focus on first action pendency because it more directly reflects examiner-driven timeliness and excludes delays such as applicant inactivity or procedural delays,[10] which can obscure differences in examiner behavior. Pendency is the primary timeliness metric reported in USPTO dashboards and serves as a proxy for institutional responsiveness. Reducing pendency is viewed as a central indicator of agency efficiency.

While the USPTO's measure of first action pendency is computed as the number of months between the filing date and first substantive action, it does not take into account internal docket-

[10]For example, applicants are given up to three months to respond to an office action, and they can also pay a fee to extend this deadline for another three months (U.S. Patent and Trademark Office 2025a). In addition, if applicants do not explicitly abandon their applications, the USPTO considers the application abandoned only after six months have passed without a response to an office action (Kuznicki 2014). There can also be delays between when the examiner issues an office action and when it is actually mailed (and the mailing date is what the USPTO uses for measuring pendency).

ing events (assignments to the examiner). Specifically, applications can sometimes be artificially "re-docketed and given a fresh date so no case gets overdue," (former patent examiner, personal interview, May 28, 2024), referring to the internal the time frame that the USPTO uses to benchmark the timeliness of examination. To account for these administrative re-dockets, I construct a measure called "true pendency" that subtracts the time between the application's first and last docketed date from the USPTO's pendency measure. If the application was transferred, I consider the first and last docketed dates following the transfer request.

I compute the true first action pendency (hereafter referred to simply as "pendency") at the application level and then average it at either the individual- or workgroup–fiscal year level for the relevant analyses.

In addition to pendency, I include the number of examiner entries and exits within each workgroup–fiscal year as organizational outcome measures. Exit counts are based on the number of examiners who leave the data in that fiscal year (excluding fiscal year 2020, since departures cannot be reliably determined for that year). Entry counts are based on the number of examiners hired into the workgroup in that fiscal year.

To better understand who is leaving or entering, I examine exits and entries specifically among highly educated examiners, defined as those with professional and/or doctorate degrees.

4.3 Measuring Production

Examiners receive production counts for completing key actions in the patent review process, and their output is evaluated against individualized goals based on their General Schedule (GS) grade level (reflecting tenure and experience) and the art unit or workgroup (reflecting the technological complexity of their assignments).

The USPTO measures patent examiner production using a count system that assigns numerical values to specific examination actions. The production goal formula can be summarized as:

$$\frac{\text{Number of Examining Hours} \times \text{Seniority Factor}}{\text{Technology Complexity}} = \text{Number of PUs Needed for 100\% of Goal} \quad (1)$$

where one production unit equals two counts. More complex technologies are associated with lower production expectations to account for the additional time required per action. See Appendix A for details on coding production counts.

An examiner's production rating is determined by the ratio of her actual to expected counts (e.g., achieving at least 95% of the goal is required for a Fully Successful rating in the PAP). However, because the USPTO does not disclose the exact production targets by grade level or the complexity weights assigned to specific art units, this study uses percentile ranks within the same workgroup, GS grade level, and fiscal year to construct a proxy for the USPTO's production rating that reflects relative performance among examiners with comparable seniority and technological assignments. This approach follows the logic of Hegde, Ljungqvist, and Raj (2025).

4.4 Defining Low Performers

Building on the production measures described above, I identify low-performing examiners based on persistent low performance over time, using the combination of an absolute (raw production count) and a relative (percentile rank) measure. While the definition of low performance is inherently ambiguous, this framework provides a consistent, data-driven way to capture sustained declines in production. I further distinguish between two types of low performers: deadwood and chronic low performers.

Examiner production trends typically follow an initial learning curve: output rises early in the career as examiners gain experience and then stabilizes at a certain level. Low performers deviate from this expected trajectory in two distinct ways (Figure 1). Deadwood examiners match the expected pattern early on, showing a clear rise in production, but then experience a sharp decline and settle into a prolonged period of persistently low output. In the beginning, especially during the probationary period (one year), deadwood are on the same trajectory as non-low performers. Chronic low performers, by contrast, perform well below peers from the outset and remain low even at their peak, consistent with a mismatch between skills or interests and the requirements of the job. Examiners who do not exhibit either pattern are classified as non-low performers.

To operationalize these trends, I define deadwood as examiners who peak above the median but subsequently experience an average post-peak production at least 25% below their peak and consistently remain below the median thereafter (and over at least two fiscal years). Chronic low performers are defined as examiners who remain consistently below the median throughout their tenure. To ensure reliable identification of these patterns, examiners must have at least five consecutive years of performance data and have been hired before fiscal year 2000.[11] If an examiner has the same peak multiple times, I consider the last peak as the defining peak. Examiners who are promoted to SPE are excluded from the sample at the time of promotion, as they are no longer responsible for examining applications.

4.5 Measuring Deadwoodness

As an alternative to the binary and proportion measures of deadwood, I construct a continuous "deadwoodness" measure that captures the average tendency toward delay across all examiners in their post-peak years. It is defined as the examiner fixed effect $\hat{\alpha}_i$, extracted from a regression of pendency on examiner and fiscal year fixed effects:

$$\text{Pendency}_{it} = \alpha_i + \gamma_t + \varepsilon_{it} \tag{2}$$

Here, pendency is measured in months; α_i represents the examiner fixed effect and γ_t captures fiscal year fixed effects. The estimated examiner fixed effect $\hat{\alpha}_i$ reflects each examiner's average contribution to pendency, net of common year-specific shocks. This measure captures time-invariant examiner-level differences in processing speed, with higher values indicating a persistent tendency to delay (greater deadwoodness). Workgroup fixed effects are excluded from this specification in order to preserve both individual-level and field-specific sources of variation in pendency.

To reduce estimation noise, particularly for examiners with fewer observations, I apply Bayesian shrinkage to the $\hat{\alpha}_i$ estimates, drawing them toward the global mean based on the precision of each estimate. This procedure improves the reliability of the resulting measure by accounting for differential uncertainty across examiners (Morris 1983; Chetty, Friedman, and Rockoff 2014). I then aggregate the shrunken estimates by computing the average deadwoodness across all individual examiners within each workgroup, yielding a workgroup–level measure of persistent delay tendencies.

This approach leverages all examiners and not just those crossing a binary threshold, providing another view of low discretionary effort and its spillovers and reflecting the full distribution of post-peak performance within a group. In this way, deadwoodness summarizes the overall drag from low performance while recognizing that declines can vary in intensity across individuals.

[11] Restricting to only examiners who were hired after fiscal year 2000 (the beginning of the sample) is to ensure that examiners are not misclassified due to missing earlier years of performance data.

4.6 Measuring an Institutional Shock

Having measured the extent of deadwood in each workgroup, I now turn to measuring an institutional shock that tests a workgroup's ability to perform under pressure. This shock is useful both conceptually and empirically: it creates a sudden, externally imposed demand that makes it harder for groups to maintain performance (whether through inertia or slack), and it reveals underlying differences in adaptability. Observing groups in these high-stress moments allows me to evaluate whether the presence of deadwood compromises the group's ability to respond effectively.

A significant institutional shock arrived with the *Alice Corporation v. CLS Bank International* decision on June 19, 2014 (Supreme Court of the United States 2014), which changed the standards of examination in certain technologies. This ruling fundamentally changed examination standards in certain technology areas by holding that abstract ideas implemented on a computer are not patent-eligible unless they include an "inventive concept" that transforms the idea into a patentable invention.

While the USPTO has used U.S. Patent Classifications to identify *Alice*-affected applications (Toole and Pairolero 2020), my focus is at the workgroup level. Because workgroups often span multiple classifications, I identify *Alice*-affected groups based on their concentration in software-intensive and abstract idea–prone domains.

For these workgroups, *Alice* created a substantive shock by introducing ambiguity and complexity into their daily work. Examiners were suddenly required to apply unfamiliar legal criteria, justify rejections differently, and respond to heightened scrutiny and applicant resistance. This shock is conceptually valuable because it did not increase the volume of applications but rather raised the cognitive demands of the work, requiring judgment, learning, and adaptability. It tests a different dimension of organizational capacity: not just throughput under pressure, but also the ability to internalize new standards and exercise discretion effectively. Workgroups with a higher share of deadwood were likely at a disadvantage in rising to these new demands.

To identify the technologies most affected by *Alice*, I classify workgroups concentrated in domains such as business methods and e-commerce (3620–3690), computer science and software (2120, 2150–2160, 2190), networking and communications (2420–2490), and biotechnology and bioinformatics (1640–1650). These groupings correspond to areas where examiners faced heightened scrutiny and ambiguity after the decision.

4.7 Analytic Sample and Descriptive Statistics

The final sample consists of 6,927 examiners across 92 workgroup–fiscal years from 2000–2020. I restrict the sample to patent examiners who could be matched to their corresponding examination and employment records, who had at least five consecutive years of examination data, and who worked in workgroup–fiscal years with at least five examiners. The five-year minimum tenure excludes individuals who left during the one-year probationary period, as well as those who may have viewed the USPTO as a short-term stepping stone "just to say they did it on their resume" (personal interview, May 25, 2024), and also ensures sufficient data to observe meaningful performance trends over time. Requiring at least five examiners per workgroup–fiscal year helps ensure reliable measurement of group-level characteristics and minimizes the influence of idiosyncratic outliers in very small groups.

4.7.1 Examiner-Level Descriptive Statistics

Of 76,056 examiner–fiscal years, 5.4% are classified as deadwood, 5.5% as chronic low performers, and the remaining 89.1% as non-low performers (Table 1). This is broadly consistent with prior evidence on poor performance in the federal workforce. An OPM (1999) study estimated that about 3.7% of federal employees were poor performers, based on supervisors' assessments of their work units. In OPM's study, supervisors defined poor performers as employees they were "seriously disappointed" in, who could not be relied upon to complete their work properly, often required significant rework or reassignment, and were not "pulling their weight." While the OPM measure relied on subjective judgments and does not distinguish between types of poor performance, the prevalence of low-performing examiners identified here (whether deadwood or chronic low performers) is of a comparable magnitude, suggesting that the USPTO is not an outlier in the broader federal context.

Compared to chronic low performers and non-low performers, deadwood examiners are older, with only 26.6% under 30 years old compared to 34.4-36.9% in the other categories. On average, they have slightly lower educational attainment (bachelor's or lower degree), longer tenure, and higher salaries, suggesting they enter the agency with lower baseline qualifications but stay long enough to accumulate seniority, higher pay, and institutional protection despite dropping to persistent low performance. On average, deadwood have 8.1 years of tenure. Deadwood also slow down application processing: on average, they take 13.2 months (SD = 5.2) to issue a first action, compared to 11.7 months (SD = 4.6) for non-low performers (Figure 2a).

Chronic low performers, by contrast, resemble non-low performers in age but differ in several notable ways. A larger percentage have professional degrees (likely law degrees), are more likely to reside in the Washington, D.C. area, and have shorter tenure and lower salaries than both deadwood and non-low performers. These patterns suggest the USPTO may hire chronic low performers as a short-term strategy by bringing in educated but less experienced and lower-cost employees to meet immediate workload demands without fully considering their long-term contributions. Their presence in the Washington, D.C. area indicates that they were either already local or willing to relocate, making them readily available for on-site work and face-to-face interaction if necessary (many examiners work remotely through the USPTO's telework program). Yet this accessibility does not necessarily translate into a good fit for the role or the organization. Despite their weaker performance, chronic low performers also tend to remain in the agency, likely because the stability, pay, and benefits of the job outweigh the drawbacks. Chronic low performers also take longer than non-low performers to issue the first action (Figure 2b).

4.7.2 Workgroup-level Descriptive Statistics

To examine the organizational implications of deadwood and chronic low performers, I aggregate examiner-level data to the workgroup–fiscal year level and construct descriptive statistics. I measure deadwood in three ways: the count of examiners classified as deadwood, the proportion of deadwood examiners (normalized by group size), and a deadwoodness score, defined as the average examiner fixed effect $\hat{\alpha}_i$ from a regression of pendency on examiner and fiscal year fixed effects. I also construct a count of chronic low performers to capture a distinct form of low performance.

There is an average of 64 examiners (SD = 44) in a workgroup-fiscal year (Figure 3), with about 30% of workgroups with no deadwood (Figure 4a). There is an average of 4.4% deadwood examiners in a workgroup in any given fiscal year. While many workgroups may be able to avoid deadwood, nearly all workgroups have at least one chronic low performer in any given fiscal year (Figure 4b), with an average of 6.2%.

Of 1,187 workgroup–fiscal years, 584 (49.2%) have an above average share of deadwood, and 603 (50.8%) have a below average share (Table 2). To complement these measures, Figure B.3 presents the distribution of the deadwoodness score, which behaves differently from the binary deadwood measures.

Workgroups with an above-average share of deadwood exhibit broadly similar processing times among their non-low performers, with average pendency of 11.5 months (SD = 3.7) compared to 11.6 months (SD = 3.2) in groups with a below-average share of deadwood. At first glance, this suggests that groups with more deadwood are not inherently slower overall. However, high-deadwood groups differ along other dimensions: they experience more examiner exits (mean = 2.0 vs. 0.7, SD = 1.9 vs. 1.1) and fewer entries (mean = 1.1 vs. 6.3, SD = 2.6 vs. 13.5), as well as larger and somewhat older workforces (79.6 vs. 49.2 examiners; 12.5 vs. 6.7 years). These patterns indicate that high- and low-deadwood groups differ systematically in their size and stability, features that may mask underlying performance dynamics explored in the empirical analyses.

While non-low performers appear to have similar first-decision times in high-deadwood workgroups (Figure 5), they are noticeably slower when deadwood is measured continuously using deadwoodness (Figure 6). This contrast highlights how the prevalence of deadwood may affect peers in more subtle ways than raw group averages suggest.

5 The Impact of Deadwood on Organizational Outcomes

5.1 Presence of Deadwood

To examine how the presence of deadwood examiners relates to organizational delays, I estimate linear regressions of the form:

$$\text{Pendency}_{jt} = \beta_1 \text{Deadwood}_{jt} + \beta_2 \text{Chronic Low}_{jt} + X'_{jt}\delta + \phi_j + \gamma_t + \varepsilon_{jt} \tag{3}$$

where Pendency_{jt} represents the average first action pendency (in months) of all the non-low performers in workgroup j during fiscal year t, and Deadwood_{jt} is a measure of deadwood: percent of deadwood, above average percent of deadwood (0 otherwise), or deadwoodness (average of the individual examiner fixed effects or $\bar{\alpha}_{jt}^{\text{shrunk}}$ from Equation 2). The vector X_{jt} includes time-varying workgroup characteristics, including workgroup size (in seven size categories: 5-25, 26-50, 51-75, 76-100, 101-125, 126-250, and 151+ examiners) and age (number of years in the data, in five age categories: <2, 2-5, 6-10, 11-15, 16+). The term ϕ_j denotes workgroup fixed effects, which control for time-invariant differences across groups. Standard errors are clustered at the workgroup to account for within-group correlation over time.

This specification isolates the within–workgroup relationship between the presence of deadwood and first-action pendency, netting out unobserved heterogeneity across workgroups and fiscal-year-specific variation. The primary coefficient of interest, β_1, captures the association between average examiner-level delay tendencies and workgroup-level pendency outcomes.

In the raw data, non-low performers in high- and low-deadwood groups display similar average pendency, suggesting that groups with more deadwood are not systematically slower overall. Once controlling for group and year fixed effects, however, pendency among non-low performers increases with deadwood prevalence. This contrast implies that the raw similarity masks important within-group dynamics: when a given workgroup experiences more entrenched underperformance, even its productive examiners take longer to issue decisions.

Workgroups with an above average share of deadwood show higher pendency for non-low performers, even after accounting for group size, experience composition, and fixed differences across workgroups. Estimates from Table 3 indicate roughly a 0.5 to 0.8-month increase in pendency, equivalent to about 14–24 additional days (2–4 weeks) per application relative to comparable groups with fewer deadwood examiners. The magnitude, while modest at the case level, compounds across applications and examiners, contributing to delays that hinder the agency's progress on backlog reduction.

The distinction between deadwood and chronically low-performing examiners further supports this interpretation. In the same specification, the share of chronic low performers shows smaller and statistically weaker effects. Deadwood thus appear more disruptive than workers who were never high performers to begin with, consistent with the idea that a decline from previously higher effort can demoralize peers or erode shared performance norms.

When the percent of deadwood is treated as a continuous measure, the relationship becomes weaker and statistically insignificant once workgroup fixed effects are included (Table B.1). This attenuation suggests that the greater presence of deadwood (rather than its exact percentage) is what matters most, possibly because salience and social comparison amplify their impact on peer behavior.

The contrast between the continuous and threshold specifications points to a nonlinear relationship between deadwood and workgroup performance. The continuous share of deadwood largely reflects stable differences across groups. Some art units consistently have more or less deadwood due to enduring managerial or technological factors, so its association with pendency disappears once these characteristics are held constant. In contrast, crossing a threshold such as the organizational average may mark a socially visible shift in group composition, where deadwood becomes noticeable enough to influence peer norms and coordination. This pattern aligns with theories of normative visibility and peer monitoring, in which low effort becomes contagious only after it reaches a salient proportion within a team.

By contrast, using the continuous deadwoodness measure, which captures the average post-peak delay tendency across all examiners, yields consistent and statistically significant results (Table 4). Moving from the 25th to the 75th percentile of deadwoodness is associated with an increase of roughly two months in pendency, or around nine weeks longer per application. Because this measure reflects the overall degree of post-peak slowdown rather than a threshold count, it highlights that even modest, pervasive reductions in examiner effort collectively translate into substantial processing delays. Together, these findings suggest that the performance drag from deadwood is not solely a function of how many low performers exist, but of how entrenched and widespread their slowdown becomes within a team's workflow.

5.2 Exits and Entries

I next assess whether deadwood presence relates to turnover by estimating quasi-maximum likelihood (QML) Poisson regressions of workgroup-level exits and entries per year, with workgroup size included as an exposure term. Let Y_{jt} denote the number of exits (or entries) among examiners in group j and year t. I assume:

$$Y_{jt} \sim \text{Poisson}(\lambda_{jt})$$

where the expected count λ_{jt} satisfies:

$$\log(\lambda_{jt}) = \log(\text{Examiners}_{jt}) + \beta_1 \,\text{Deadwood}_{jt} + \beta_2 \,\text{Chronic Low}_{jt} + \delta \,\text{Age}_{jt} + \gamma_t + \epsilon_{jt}. \qquad (4)$$

Here, Examiners_{jt} is the total number of examiners in group j in year t, included as an exposure term. Its logarithm enters the equation with a fixed coefficient of 1, effectively modeling the rate of exits (or entries) per examiner rather than the raw count.

I estimate β by QML. Since Deadwood_{jt} is a binary indicator equal to 1 if the workgroup has at least the median number of deadwood (four) and 0 otherwise, the estimated coefficient β can be interpreted as the log of the ratio of expected exit (or entry) rates between groups with high deadwood versus without:

$$\frac{\lambda_{jt} \text{ when } \text{Deadwood}_{jt} = 1}{\lambda_{jt} \text{ when } \text{Deadwood}_{jt} = 0} = \exp(\hat{\beta}).$$

In other words, $\exp(\hat{\beta})$ represents the factor by which the expected number of exits (or entries), adjusted for group size, changes when a workgroup has high deadwood. For example, if $\exp(\hat{\beta}) = 1.2$, then the expected number of exits per examiner in workgroups with high deadwood is 20% higher than in those without. The same logic is used for the chronic low indicator.

Workgroups with high deadwood are associated with significantly higher (30.5%) overall rates of examiner exits (Table 5). The association is especially pronounced among highly educated examiners (defined as those with a professional or doctorate degree), whose exit rates are 62.1% higher in high-deadwood groups. These patterns suggest that deadwood undermines morale and imposes additional burdens, making the environment less attractive for capable employees who have better outside options. As one indicator of these dynamics, several former examiners shared in interviews that they or their colleagues attended law school or studied for the bar while at the USPTO, or left the agency to join a law firm. By contrast, exits of high performers are not associated with the prevalence of deadwood, indicating that examiners who thrive in the job appear largely insulated from the negative peer environment: deadwood may discourage those weighing alternative career paths, but it does not drive out the most productive examiners.

There is an asymmetry when it comes to entry rates. While entry rates tend to decline in high-deadwood workgroups, the estimates are small and statistically insignificant across all groups. This suggests that deadwood is more clearly associated with higher turnover among incumbents than with reduced inflows of new hires. In other words, the main organizational cost appears to come from pushing out employees who might otherwise have stayed, rather than deterring recruitment at the outset.

This pattern may reflect both applicant behavior and organizational allocation decisions. Although applicants cannot choose their specific workgroup (assignments are based on fit and organizational need), perceptions of the USPTO as a less attractive employer, perhaps due to the prevalence of deadwood (and thus potentially more work for new hires), may deter some potential candidates from applying at all. More importantly, the lower entry rates likely reflect how the USPTO allocates its workforce: the agency prioritizes placing new hires in groups it identifies as understaffed. As one patent examiner explained, "Anyone being hired will, by definition, go into an understaffed [art unit]" (Reddit, 2023). High-deadwood groups may not appear understaffed on paper, since they already meet nominal headcount levels, even though much of their capacity is ineffective; this misrepresentation could guide the agency to direct new hires elsewhere.

These results suggest that the presence of deadwood is associated with unfavorable dynamics in group composition: higher turnover among better-qualified incumbents without compensating inflows of strong new hires. Together, the findings point to a dynamic in which deadwood not only depresses immediate productivity but also undermines group capacity over time by driving away stronger incumbents and failing to replenish them with equally capable or better replacements.

5.3 *Alice Corp v. CLS Bank International*

Finally, I examine whether the association between deadwood and performance changes after an institutional change to see how workgroups with high deadwood adapt when facing increased stress. I estimate the effect of *Alice*:

$$
\begin{aligned}
\log(\text{Pendency}_{jt}) = & \beta \text{Deadwood}_{jt} + \theta \text{Post-Alice}_{jt} + \eta \text{Alice-Affected}_j \\
& + \mu(\text{Deadwood}_{jt} \times \text{Post-Alice}_{jt}) \\
& + \lambda(\text{Deadwood}_{jt} \times \text{Alice-Affected}_{jt}) \\
& + \rho(\text{Deadwood}_{jt} \times \text{Post-Alice}_{jt} \times \text{Alice-Affected}_j) \\
& + X'_{jt}\delta + \phi_j + \gamma_t + \varepsilon_{jt}
\end{aligned}
\tag{5}
$$

where Post-Alice_{jt} is 2015 and after, to account for a lag in the decision in June 2014, Alice-Affected_j is omitted because it is collinear since it is defined at the workgroup level. Interactions with chronic low are also included in the model, but are omitted from the equation for brevity.

The *Alice* decision exposed and amplified the fragility of deadwood-heavy workgroups. Even before the ruling, these groups exhibited slower processing, with non-low performers taking roughly 0.7 months longer on average than those in low-deadwood groups. After *Alice*, the gap widened substantially. Pendency rose sharply in high-deadwood groups operating in *Alice*-affected technologies: the triple-interaction coefficient in Table 6 indicates an additional 1.3-month (about 6-week) increase in first-action pendency among their non-low performers.

By contrast, workgroups with a high share of chronic low performers showed a different pattern. Although these groups are generally slower overall (about 0.7 months higher pendency), their performance did not worsen further after *Alice*. The absence of a post-shock effect suggests that while chronic low performers impose a steady drag, they do not fundamentally undermine the group's ability to adapt when external conditions shift. In other words, deadwood introduces fragility, whereas chronic low performers create persistent but predictable inefficiency.

Patterns in staffing further highlight these dynamics. High-deadwood and *Alice*-affected groups were associated with increased entries after the shock, suggesting that management redirected hiring toward units under stress. Yet these inflows did little to offset the internal slowdown.

Taken together, the results show how deadwood becomes disproportionately costly under institutional shocks. *Alice* demanded conceptual and procedural change—precisely the kind of adjustment that exposes the limits of entrenched routines. As one SPE noted, "Every time there is a law change, it's not easy to teach an old dog new tricks. It is hard for them to readjust" (personal interview, June 26, 2024). An examiner may feel "on the edge already. They're then asked to do duplicate work or revise work to do something" because of case-law changes like *Alice* (former patent examiner, personal interview, July 23, 2025). *Alice* did not merely add work; it required judgment, learning, and adaptation. In those moments, the presence of entrenched low performers left groups slower and more vulnerable.

5.4 Exploring Examiner Behaviors

Beyond production and timeliness, deadwood may also differ in how they conduct examinations. It is possible that examiners who are low in production are not necessarily low performers on other dimensions. To test this, I examine this through four dimensions: discretionary effort (Panel A), decision style and risk orientation (Panel B), responsiveness to production incentives (Panel C), and examination quality (Panel D).

Discretionary effort captures the extent to which examiners limit patent scope and document prior art. The first measure, *claim reductions*, is the average change in independent claims between filing and grant per application for each examiner–fiscal year. The second, *examiner-inserted citations*, is the average number of citations added by the examiner to foreign patents, U.S. patent applications, and granted U.S. patents, divided by the total number of applications.

Decision style and risk orientation is reflected in *leniency*, the cumulative share of applications allowed, following Sampat and Williams (2019). Unlike their implementation, I do not leave one observation out, as I do not use this measure as an instrumental variable.

I introduce two new measures that reflect responsiveness to production incentives, which considers when examiners complete their work and how end-of-period incentives may shape decisions: the *fraction of annual production completed in the last fiscal month* and the *allowance rate in the last fiscal month*.

Examination quality is measured by the *appeal affirmation rate* or the share of appealed decisions upheld. The fact that appeals are decided with some lag is part of what makes them useful: examiners may have to "wait 2.5 years for the board to tell [them] whether [they're] right or wrong" (former patent examiner, personal interview, June 5, 2024), so the prospect of appeal does not influence day-to-day behavior. Because appeals are rare, I calculate this rate cumulatively, as with leniency, to capture more stable patterns.

For each examiner, I calculate these measures for the pre- and post-peak periods and take the difference. These behavioral measures complement production and pendency by showing whether deadwood examiners not only produce less but also change how they apply examination standards, with implications for patent quality and applicant experience.

The results show that deadwood and non-low performers follow some similar trends but diverge in the magnitude of their behavioral shifts after peak production (Figure 7; Table 7). Deadwood show a pronounced decline in examiner-inserted citations, suggesting less attention to one of the more discretionary and time-intensive aspects of the job. Yet on other dimensions, such as leniency and end-of-year allowances, their changes are modest, suggesting they are not as responsive to incentives.

Non-low performers, by contrast, display many of the same shifts but with greater magnitude. They reduce substantially more claims, become more lenient, and are much more likely to allow cases at the end of the fiscal year. These adjustments suggest that non-low performers are more responsive to production incentives, altering their practices in ways that help sustain throughput. These patterns show that post-peak decline takes different forms, with deadwood examiners neither uniformly disengaged nor uniformly diligent.

6 The Organizational Costs of Deadwood

These results document that deadwood examiners impose significant delays on their non-low performing colleagues and on the organization as a whole. This section provides back-of-the-envelope calculations of the organizational burdens imposed by deadwood, considering not only their contribution to the application backlog but also the personnel costs of retaining them and the potential gains from restoring their productivity. Though simplified, these calculations provide an intuitive sense of the scale of inefficiencies that even a modest share of deadwood can generate in a large, process-driven institution like the USPTO.

6.1 Backlog Contribution

As previously discussed, the USPTO's backlog of pending patent applications remains a central organizational challenge, holding back innovation and potentially rendering patents useless if the process takes too long. Deadwood examiners contribute to delays both through their own slower output and through negative spillovers onto peers. To quantify the total burden on the agency's backlog, consider that 5.4% of examiner–fiscal years (about 196 examiners annually[12]) are classified as deadwood. On average, each examiner processes 90 patent applications annually, meaning that deadwood examiners collectively handle roughly 41,940 applications per year.

My estimates indicate that each application handled by a deadwood examiner is associated with an additional 1.5 months of pendency, relative to applications handled by peers. Multiplying these figures suggests that deadwood examiners contribute roughly 62,910 application-months of additional delay annually. If an examiner's average *true* first action pendency (as defined in this

[12] This is a conservative estimate because it is based on the study's sample, which has several strict individual and workgroup-level inclusion criteria.

paper) is 11.7 months, this estimate translates to 2,260 additional applications sitting in the backlog at any given time.

Their presence also slows down their non-low-performing peers: in groups with high deadwood prevalence, pendency among non-low performers rises enough to generate a comparable backlog effect. Applying the same conversion, these peer spillovers contribute an additional 11,926 applications stuck in the queue each year. The combined direct and indirect effects of deadwood add up to roughly 14,186 extra applications annually, underscoring how a small number of entrenched low performers can reshape the performance of the entire system.

6.2 Personnel Costs of Deadwood

Beyond contributing to delays, deadwood examiners also represent a substantial personnel cost. Based on personnel records, the average salary of a deadwood examiner is approximately $107,604 per year. With approximately 196 deadwood examiner–fiscal years annually, the agency spends roughly $21.0 million each year compensating deadwood examiners. This is $1.5 million more annually than it would cost to pay non-low performers in those same positions. Moreover, deadwood examiners have an average tenure of 8.1 years, which means their salary premium alone accumulates to about $61,317 per examiner over the course of their career, not counting the productivity losses or backlog impact they impose. These resources could otherwise fund additional examiners, improve training or supervision, or support other organizational priorities. Viewed differently, the agency not only bears the productivity losses and backlog costs imposed by deadwood, but also pays a salary premium to retain them over long careers.

6.3 Unrealized Potential of Deadwood

The USPTO operates below its potential because deadwood examiners have previously demonstrated much higher productivity at their peaks. While all examiners' production tends to decline over time, the magnitude of this decline varies substantially across groups (Table B.2). Deadwood examiners currently operate at about 60% of their peak output, compared to non-low performers who sustain approximately 79% of peak (a 21% drop). This gap highlights that the costs of low performance are not simply about having less productive employees, but about allowing some employees to decline to drastically low levels compared to their demonstrated potential.

If deadwood examiners could be restored even halfway to their peak productivity, their output would increase by about 20%, unlocking meaningful organizational gains without additional hiring. Moreover, the results show that workgroups with high deadwood are associated with fewer new examiners, and since new examiners have a one-year probationary period and take nearly six years to reach peak productivity, investing in existing employees may offer a faster and more effective path to improving organizational performance.

These calculations, while illustrative and based on simplifying assumptions, underscore the significant organizational burden posed by deadwood examiners. They assume linear and additive effects, uniform spillovers, and no behavioral or structural adjustments by the agency. Nevertheless, they demonstrate the potential magnitude of inefficiencies created by persistent low performers and highlight the value of organizational interventions to address deadwood employees.

7 Conclusion

This paper examines how entrenched low performers affect organizational outcomes in the USPTO. Using linked administrative and personnel data from 2000–2020, I show that even a modest share of deadwood generates delays and inefficiencies for the organization, not only through their own output but also through negative spillovers onto peers and teams.

Qualitative evidence and organizational theory help explain why the costs of deadwood exceed their individual production shortfall. Colleagues note that low performers often retain seniority and pay, which demoralizes higher performers and distorts workplace norms. One former patent examiner described the stereotypical deadwood examiner as someone who "does the bare minimum to meet [their quota] and then disappears for the rest of the time" (personal interview, May 29, 2024). Another remarked that it was discouraging to work alongside colleagues who were "paid the same or more" despite contributing less (personal interview, May 28, 2024). Such perceptions echo classic accounts of cooperation as grounded in fairness and reciprocity (Barnard 1938), and resonate with Roy's account of "banana time" disruptions when shared rhythms collapse as some workers withdraw (Roy 1959).

While the USPTO prioritizes production as its main metric, it is possible that deadwood or low-producing examiners are misunderstood as slackers. Some interviews suggest that certain slower examiners are more meticulous rather than disengaged. As one former examiner explained, “There are people who struggle because they are trying to send out the best they can. You do run into them sometimes, and you can tell they’re trying to do a good job” (personal interview, June 5, 2024). Others may be quietly resisting an incentive system that prizes production and speed over rigor:

> When an examiner said, “We don’t have enough time to do a good job,” his manager put it this way: “You have to understand we’re not being paid to put a piano finish on an orange tree.” *This* is what patent examining is about (former patent examiner, personal interview, May 29, 2024).

These quotes highlight the tradeoff between the pressure to meet production quotas and the rigor of review. Because my analyses of examiner behaviors do not strongly suggest that deadwood examiners systematically produce lower-quality examination than their peers, they suggest that productivity and other aspects of examination do not necessarily move together. Both deadwood and non-low performers adjust their behaviors after peak productivity, but deadwood scale back modestly while non-low performers make sharper shifts that reflect stronger responsiveness to production incentives. This nuance shows why persistent low performance is best understood not only as an individual trajectory but also as an organizational outcome shaped by accountability systems.

Altogether, this study establishes a new area of research on low performance by developing a conceptual and empirical lens on deadwood employees as a distinct form of persistent low performance. This perspective deepens our understanding of performance interdependence by showing how entrenched low performers can affect not only their own output but also the productivity of their peers. More broadly, the findings show how accountability systems can, paradoxically, sustain low performance—sometimes even legitimizing it—until an external shock reveals its full cost.

Addressing low performance requires rethinking reactive, individual-focused interventions. Performance improvement plans or forced rankings are poorly suited to identifying entrenched decline, since they focus on short-term dips or relative rankings. Recognizing deadwood as a structural, team-embedded issue points to alternative strategies. For example, longitudinal monitoring, more balanced performance metrics, and redesigned incentives can surface sustained disengagement without discouraging discretion or quality. For organizations in both the public and private sectors, it is crucial to design systems that foster productivity and resilience through aligned incentives and shared norms, while also guarding against entrenched disengagement.

References

Adams, J. Stacy. 1965. "Inequity in social exchange." In *Advances in Experimental Social Psychology,* edited by Leonard Berkowitz, 2:267–299. New York, NY: Academic Press.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang. 2010. "Superstar Extinction." *The Quarterly Journal of Economics* 125 (2): 549–589.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data." *The Quarterly Journal of Economics* 120 (3): 917–962.

Barnard, Chester I. 1938. *The Functions of the Executive.* Cambridge, MA: Harvard University Press.

Barsade, Sigal G. 2002. "The Ripple Effect: Emotional Contagion and its Influence on Group Behavior." *Administrative Science Quarterly* 47 (4): 644–675.

Bidwell, Matthew. 2011. "Paying More to Get Less: The Effects of External Hiring Versus Internal Mobility." *Administrative Science Quarterly* 56 (3): 369–407.

Brox, Enzo, and Michael Lechner. 2024. "Teamwork and Spillover Effects in Performance Evaluations." *Working Paper.*

Carley, Michael, Deepak Hegde, and Alan Marco. 2014. "What is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Technology* 17.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review* 104 (9): 2633–2679.

Cohen-Zada, Danny, Itay Dayag, and Naomi Gershoni. 2023. "Effort Peer Effects in Team Production: Evidence from Professional Football." *Management Science* 70 (4): 2355–2381.

Cyert, Richard Michael, and James G March. 1963. *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice-Hall.

Felps, Will, Terence R. Mitchell, and Eliza Byington. 2006. "How, When, and Why Bad Apples Spoil the Barrel: Negative Group Members and Dysfunctional Groups." *Research in Organizational Behavior* 27:175–222.

Gittell, Jody Hoffer. 2002. "Relationships between Service Providers and Their Impact on Customers." *Journal of Service Research* 4 (4): 299–311.

Goerg, Sebastian J, Oliver Himmler, and Tobias König. 2024. "Norm violations and behavioral spillovers—Evidence from the lab and the field." *European Economic Review* 166:104776.

Grant, Adam M, and Jihae Shin. 2012. "Work Motivation: Directing, Energizing, and Maintaining Effort (and Research)." In *Oxford handbook of motivation,* 1st ed., edited by Richard M Ryan, 505–519. New York, NY: Oxford University Press.

Grigoriou, Konstantinos, and Frank T Rothaermel. 2017. "Organizing for knowledge generation: internal knowledge networks and the contingent effect of external knowledge sourcing." *Strategic Management Journal* 38 (2): 395–414.

Hannan, Michael T, and John Freeman. 1984. "Structural Inertia and Organizational Change." *American Sociological Review* 49 (2): 149–164.

Hegde, Deepak, Alexander Ljungqvist, and Manav Raj. 2025. "Race and Slower Promotions for Equal Work." *Swedish House of Finance Research Paper,* Paper No. 21–09.

Hendricks, Jonathan L, Matthew L Call, and Elizabeth M Campbell. 2022. "High Performer Peer Effects: A Review, Synthesis, and Agenda for Future Research." *Journal of Management* 49 (6): 1997–2029.

Huckman, Robert S, and Bradley R Staats. 2011. "Fluid Tasks and Fluid Teams: The Impact of Diversity in Experience and Team Familiarity on Team Performance." *Manufacturing and Service Operations Management* 13 (3): 281–418.

Ichniowski, Casey, and Anne Preston. 2014. "Do Star Performers Produce More Stars? Peer Effects and Learning in Elite Teams." *NBER Working Paper,* 20478.

Kamei, Kenju, and John Ashworth. 2023. "Peer learning in teams and work performance: Evidence from a randomized field experiment." *Journal of Economic Behavior & Organization* 207 (1): 413–432.

Kandel, Eugene, and Edward P. Lazear. 1992. "Peer Pressure and Partnerships." *Journal of Political Economy* 100 (4): 801–817.

Karau, Steven J, and Kipling D Williams. 1991. "Social loafing and social compensation: the effects of expectations of co-worker performance." *Journal of Personality and Social Psychology* 65 (4): 681–706.

Kellogg, Katherine C. 2009. "Operating Room: Relational Spaces and Microinstitutional Change in Surgery." *American Journal of Sociology* 115 (3): 657–711.

Khan, Muhammad Yasir. 2025. "Mission Motivation and Public Sector Performance: Experimental Evidence from Pakistan." *American Economic Review* 115 (7): 2343–75.

Kuznicki, Kathleen. 2014. "How Long Does It Take to Get a Patent?" Exit Promise, April 3, 2014. Accessed May 10, 2025. https://exitpromise.com/patent-approval-process/.

Langcawon, Maria Judesa C. 2025. "The Impact of Job Mismatch on Workplace Productivity at Mariners' Polytechnic Colleges Inc." *Cognizance Journal of Multidisciplinary Studies* 5 (3): 340–35.

Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review* 90 (5): 1346–1361.

Lemley, Mark A, and Bhaven Sampat. 2008. "Is the Patent Office a Rubber Stamp?" *Emory Law Journal* 58 (181).

———. 2012. "Examiner Characteristics and Patent Office Outcomes." *The Review of Economics and Statistics* 94 (3).

Lu, Qiang, Amanda F. Myers, and Scott Beliveau. 2017. "USPTO Patent Prosecution Research Data: Unlocking Office Action Traits." *SSRN.*

March, James, and Herbert Simon. 1958. *Organizations.* 2nd ed. New York: Wiley-Blackwell.

Mas, Alexandre, and Enrico Moretti. 2009. “Peers at Work.” *American Economic Review* 99 (1): 112–45.

McGregor, Douglas. 1960. *The Human Side of Enterprise.* New York, NY: McGraw–Hill.

Mitra-Kahn, Benjamin, Alan Marco, Michael Carley, Paul D’Agostino, Paul Evans, Carl Frey, and Nadiya Sultan. 2013. *Patent Backlogs, Inventories and Pendency: An International Framework.* Intellectual Property Office.

Morris, Carl. 1983. “Parametric Empirical Bayes Inference: Theory and Applications.” *American Statistical Association* 78 (381): 47–55.

Novak, Sharon, and Scott Stern. 2008. “How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry.” *Management Science* 54 (12): 1963–1979.

Osterloh, Margit, and Bruno S Frey. 2000. “Motivation, Knowledge Transfer, and Organizational Forms.” *Organization Science* 11 (5): 538–550.

Patent Office Professional Association. 2025. “Awards.” Accessed May 9, 2025. http://www.popa.org/about/awards/.

Podsakoff, Philip M, Scott B MacKenzie, Julie Beth Paine, and Daniel G Bachrach. 2000. “Organizational Citizenship Behaviors: A Critical Review of the Theoretical and Empirical Literature and Suggestions for Future Research.” *Journal of Management* 26 (3): 513–563.

Repenning, Nelson P, and John D Sterman. 2002. “Capability traps and self-confirming attribution errors in the dynamics of process improvement.” *Administrative Science Quarterly* 47 (2): 265–295.

Rintamäki, Jukka, Simon Parker, and André Spicer. 2025. “Institutional Parasites.” *Academy of Management Review* 50 (3): 612–631.

Roy, Donald F. 1959. “”Banana Time” Job Satisfaction and Informal Interaction.” *Human Organization* 18 (4): 158–168.

Saez, Emmanuel, Benjamin Schoefer, and David Seim. 2023. "Deadwood Labor: The Effects of Eliminating Employment Protection." *SSRN.*

Sampat, Bhaven, and Heidi L. Williams. 2019. "How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome." *American Economic Review* 109 (1): 203–236.

Simon, Herbert A. 1997. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations.* 4th ed. New York: Free Press.

Tian, Zhifan, Cheng (Colin) Zeng, Chaofan Li, and Yi Wu. 2025. "Peer effects of star-analysts' departure: New evidence from China." *Journal of Corporate Finance* 94:102844.

Toole, Andrew A, and Nicholas A Pairolero. 2020. *Adjust to Alice: USPTO patent examination outcomes after Alice Corp v. CLS Bank International.* 3. Alexandria, VA: U.S. Patent and Trademark Office, April.

U.S. Government Accountability Office. 2016. *Patent Office Should Strengthen Search Capabilities and Better Monitor Examiners' Work.* Report to the Chairman, Committee on the Judiciary, House of Representatives GAO-16-479.

———. 2025. *Patent Office Should Strengthen Its Efforts to Address Persistent Examination and Quality Challenges.* April.

U.S. Office of Personnel Management. 1999. "Poor Performers in Government: A Quest for the True Story." *Report of a Special Study.*

U.S. Patent and Trademark Office. 2022. "909.01(b) Classification and Assignment to Examiner in CPC." In *Manual of Patent Examining Procedure.* Accessed April 23, 2025. https://www.uspto.gov/web/offices/pac/mpep/s909.html#ch900_d33654_1dd68_1ab.

———. 2025a. "Response time period." Response time period. Last Modified: 2023-03-24T14:38:09Z. Accessed May 10, 2025. https://www.uspto.gov/trademarks/apply/response-time-period.

———. 2025b. "Subject matter descriptions." Accessed July 9, 2025. https://www.uspto.gov/patents/basics/essentials/business-methods/subject-matter-descriptions.

U.S. Patent and Trademark Office. 2025c. *Utility, Plant, and Reissue (UPR) unexamined patent application inventory.* Accessed July 5, 2025. https://www.uspto.gov/dashboard/patents/#uprinventory.

Usanova, Ksenia, and Mickaël Géraudel. 2024. "Talent turnover and the contagion effect: the case of the aerospace industry." *Human Resource Development International* 27 (4): 458–480.

USPTO Joint Labor and Management Count System Task Force. 2010. "Overview of Count System Initiatives and Changes." March 8, 2010. Accessed June 26, 2024.

Weick, Karl E. 1995. *Sensemaking in Organizations.* Vol. 3. Foundations for Organizational Science. SAGE Publications, Inc.

Zhou, Jinyi, Ning Li, Shiyong Xu, and Wei Chi. 2025. "High performers in the shadow: The adverse effect of star employees on their peers." *Journal of Applied Psychology* 110 (6): 776–797.

Figure 1: Distinguishing between types of workers

(a) Production Count

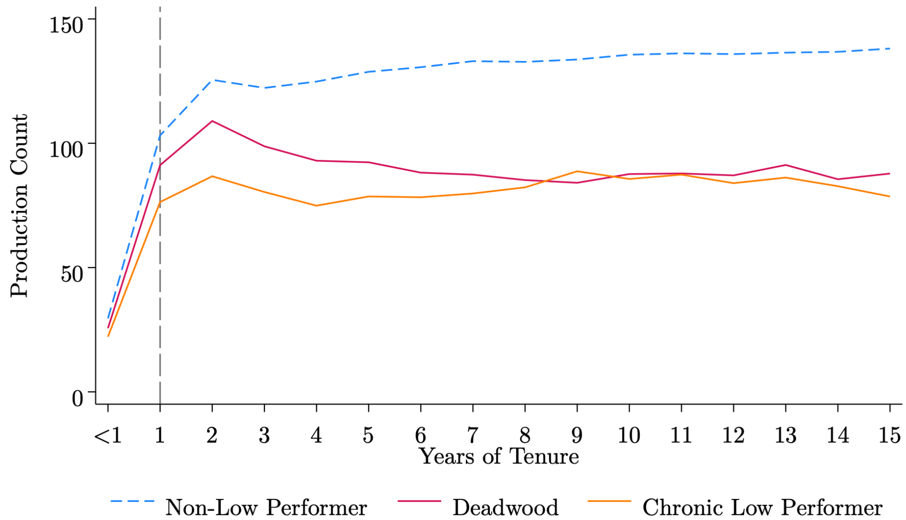

(b) Percentile Compared to Peers

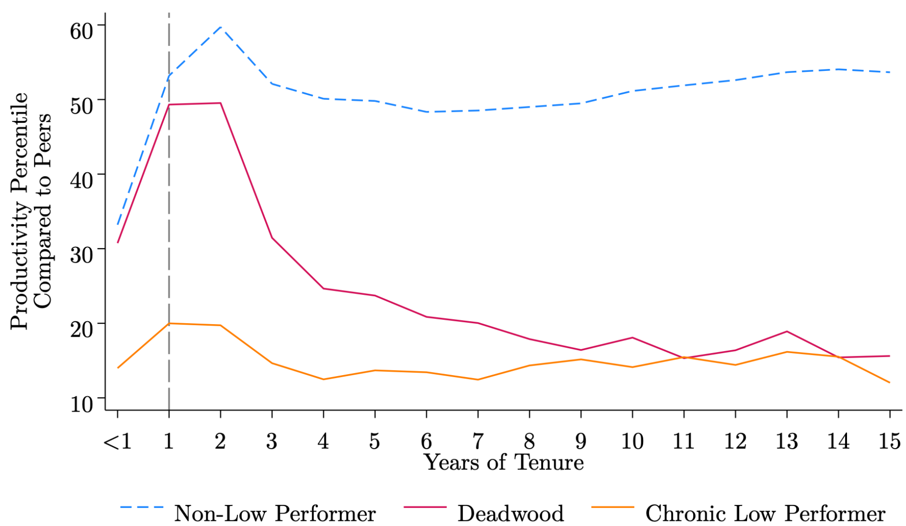

Notes: This figure shows the trends in performance over time by type of worker. Worker types are based on a combination of an absolute (production count) and a relative (percentile compared to peers with the same technology complexity and seniority level) measure. A deadwood examiner is defined as having a post-peak average of at least 25% below her peak and consistently below the median thereafter. A chronic low performer is an examiner who never achieves performance above the median. Examiners who do not fit the definitions of either deadwood or chronic low are considered non-low performers. Any year as a manager is excluded.

Table 1: Examiner-fiscal year sample statistics

	Deadwood		Chronic low		Non-low	
		SD		SD		SD
Age						
Percent <30 years old	26.55	44.18	36.87	48.26	34.41	47.51
Percent 30-44 years old	55.83	49.67	44.01	49.65	46.94	49.91
Percent 45-59 years old	13.81	34.52	14.48	35.20	15.67	36.35
Percent 60+ years old	3.80	19.13	4.64	21.04	2.99	17.03
Education						
Percent bachelor's or lower degree	61.50	48.67	57.59	49.43	60.86	48.81
Percent master's degree	20.73	40.54	23.67	42.51	20.83	40.61
Percent professional degree	8.27	27.54	10.63	30.83	7.48	26.31
Percent doctorate degree	9.51	29.34	8.11	27.30	10.83	31.08
Highest degree year	2003	6.89	2003	7.21	2003	6.88
Percent in Washington, D.C. area	77.25	41.93	80.09	39.94	77.66	41.65
Percent primary examiner	43.43	49.57	20.63	40.47	36.48	48.14
Tenure (years)	8.07	3.76	5.23	3.83	6.48	4.42
Adjusted basic pay ($)	107,604	21,233	91,741	22,332	100,005	25,699
Number of examiner-fiscal years	4,101	4,101	4,158	4,158	67,848	67,848

Notes: This table shows the sample statistics at the examiner-fiscal year level. Age was only available for 2000-2013, and for examiners who were employed during the first two quarters of 2014. Based on data release policies changing over time and availability, age data was available for 53.2% of examiner-fiscal years. For the other variables, either zero or minimal data (<0.3%) was missing. Primary examiner status is defined as being at GS-14 or above. Washington, D.C. area includes District of Columbia, Maryland, and Virginia. Adjusted basic pay is annual salary adjusted for cost of living and to be competitive with private sector equivalents.

Figure 2: Pendency (in months) by type of low performer

(a) Deadwood

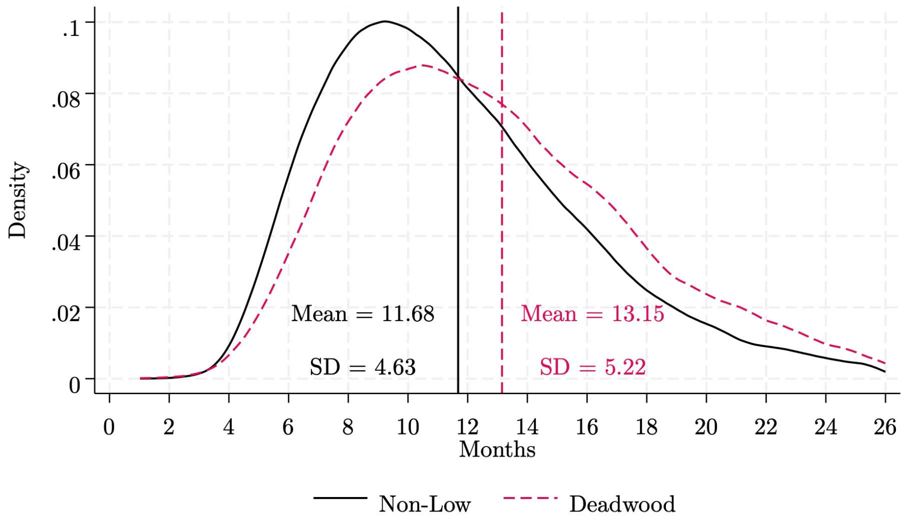

(b) Chronic low

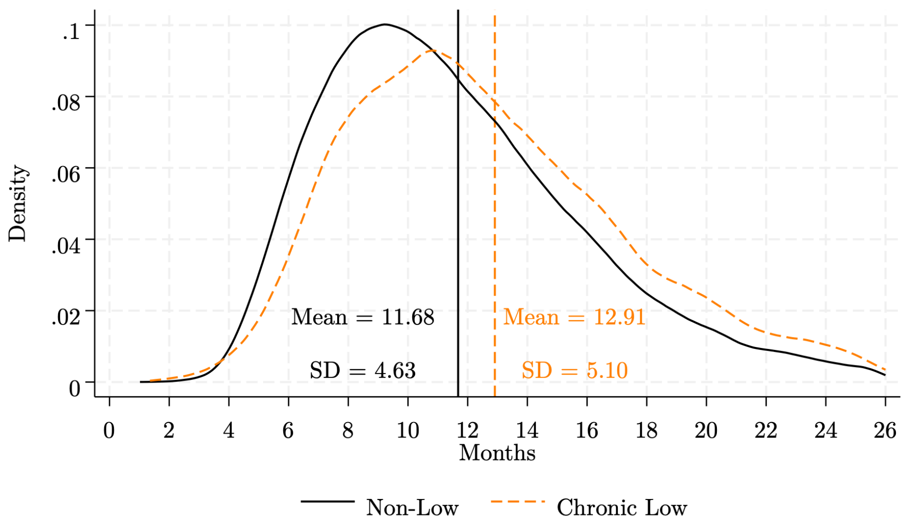

Notes: This figure shows the distributions of "true" first action pendency (in months) by type of low performer. This measure of pendency subtracts the time between the application's first and last docketed date from the USPTO's pendency measure. This adjustment accounts for internal re-docketing events that prolongs the examination time.

Figure 3: Distribution of workgroup size

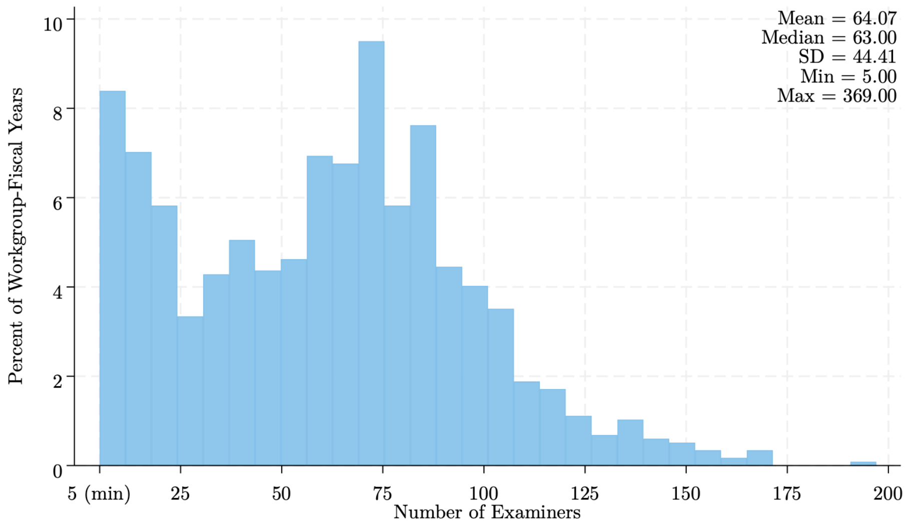

Figure 4: Percent of low performers in a workgroup-fiscal year

(a) Histogram of deadwood

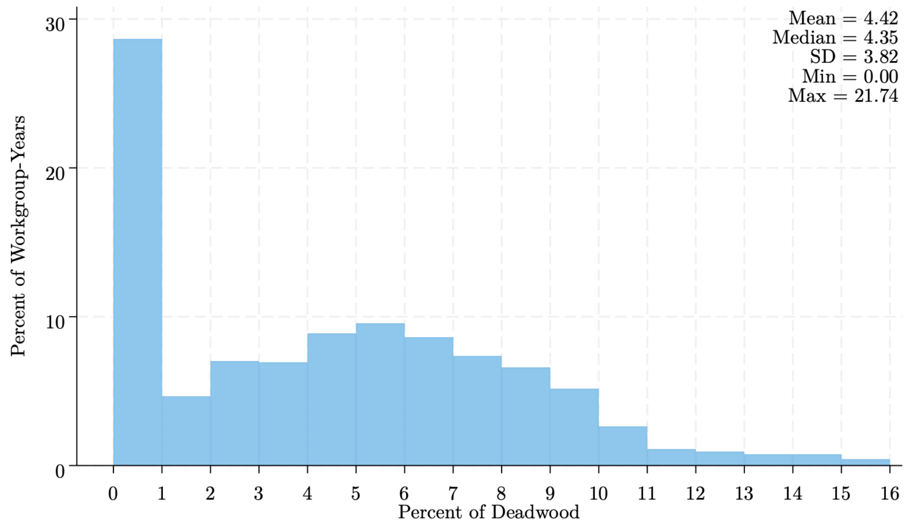

(b) Histogram of chronic low

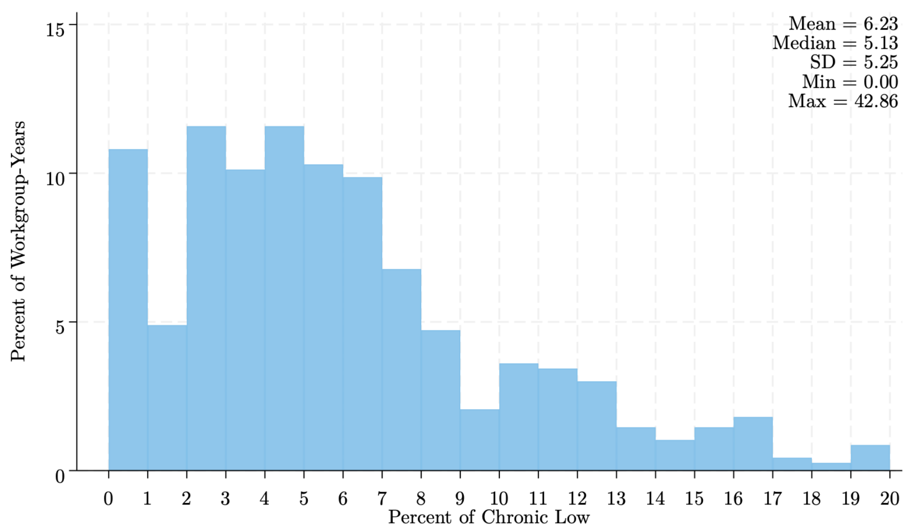

Notes: These figures show the distribution of the percent of deadwood and chronic low performers in a workgroup in a given fiscal year. Deadwood is defined as examiners whose production peak above the median (50th percentile) compared to their peers (other examiners in the same workgroup, grade, and fiscal year) and then plateau, where their post-peak average indicates a 25% or greater decline from their peak and every post-peak year is below the median. Chronic low performers are examiners who never achieved production above the median.

Table 2: Workgroup-fiscal year sample statistics

	High deadwood			Low deadwood		
	Mean	Median	SD	Mean	Median	SD
Pendency of non-low performers (months)	11.45	10.77	3.71	11.58	11.38	3.21
Exits						
Number of overall exits	2.05	2.00	1.92	0.67	0.00	1.11
Number of high performer exits	0.14	0.00	0.40	0.04	0.00	0.22
Number of high education exits	0.37	0.00	0.71	0.10	0.00	0.32
Entries						
Number of overall entries	1.06	0.00	2.57	6.28	2.00	13.46
Number of high performer entries	0.34	0.00	1.03	1.63	0.00	3.77
Number of high education entries	0.28	0.00	0.99	1.13	0.00	2.98
Deadwood Measures						
Number of deadwood	5.99	6.00	3.25	1.02	0.00	1.62
Proportion of deadwood	7.67	7.25	2.53	1.29	0.00	1.57
Deadwoodness	0.24	0.13	1.43	-0.09	-0.20	1.82
Number of chronic low	4.08	3.00	3.27	2.94	2.00	2.83
Proportion of chronic low	5.18	4.57	3.48	7.24	6.00	6.37
Workgroup Characteristics						
Number of examiners	79.58	76.00	38.69	49.16	39.00	44.48
Workgroup age	12.53	13.00	5.07	6.68	6.00	5.05
Number of workgroup-fiscal years	582	582	582	605	605	605

Notes: This table shows the sample statistics at the workgroup-fiscal year level. High (low) deadwood includes workgroups with an above (at or below)-average percent of deadwood. Pendency is true first-action pendency of non-low performers. High performers are those who have a professional and/or doctorate degree. Deadwoodness is the individual examiner fixed effects extracted from a linear regression of the true first action pendency on fiscal year and examiner fixed effects, only including each examiner's post-peak years. The individual examiner fixed effects are then Bayesian shrunk, and the average is computed at the workgroup-fiscal year level.

Figure 5: Pendency of non-low performers in workgroup-fiscal years with above average percent of deadwood

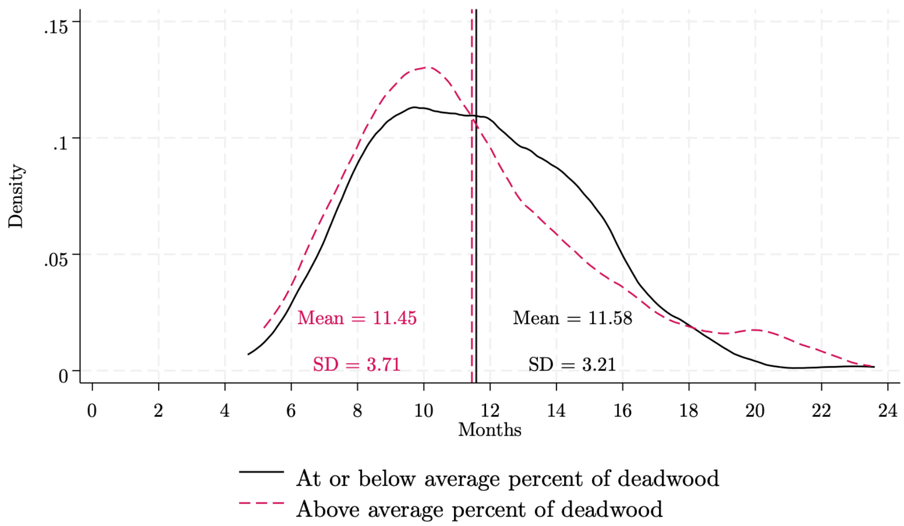

Figure 6: Pendency of non-low performers in workgroup-fiscal years with above average deadwoodness

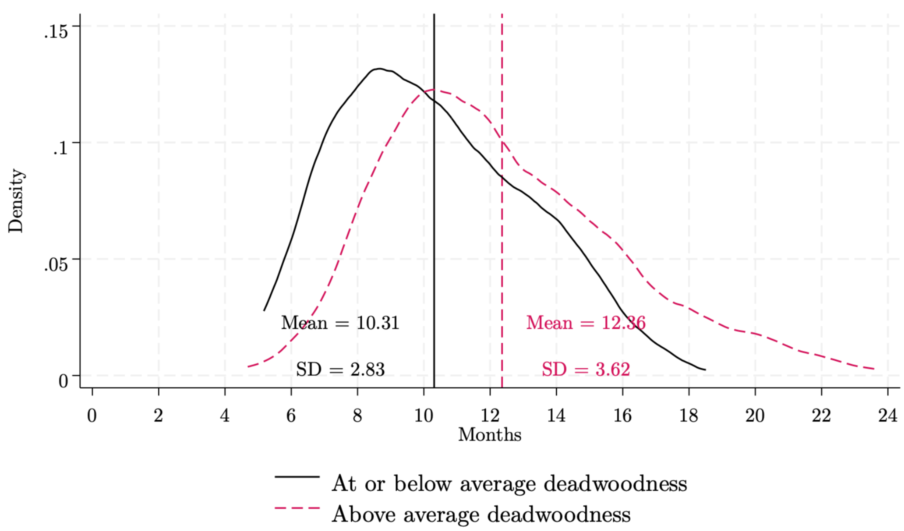

Notes: Deadwoodness is the individual examiner fixed effects extracted from a linear regression of the true first action pendency on fiscal year and examiner fixed effects, only including each examiner's post-peak years. The median deadwoodness in a workgroup-fiscal year is 0.061.

Table 3: Deadwood examiners have a disproportionate effect on pendency of non-low performers compared to chronic low performers

	Pendency of non-low performers			
	(1)	(2)	(3)	(4)
Above average percent of deadwood	0.798***	0.824***	0.439**	0.455**
	(0.278)	(0.270)	(0.206)	(0.201)
Above average percent of chronic low	0.433*	0.411*	0.263	0.267
	(0.225)	(0.228)	(0.169)	(0.173)
Fiscal year FEs	Yes	Yes	Yes	Yes
Workgroup FEs	No	No	Yes	Yes
Workgroup size and age controls	No	Yes	No	Yes
R-squared	0.649	0.652	0.805	0.808
Mean(Pendency of non-low performers)	11.517	11.517	11.517	11.517
Number of workgroups	79	79	79	79
Number of observations	1,187	1,187	1,187	1,187

Notes: Pendency is a workgroup's average true first action pendency of non-low performers in a given fiscal year. A positive coefficient indicates an increase in the amount of time it takes for an applicant to receive the first decision from the examiner. High prevalence of deadwood is a workgroup with at least four deadwood examiners; high prevalence of chronic low is a workgroup with at least four chronic low performers. Observations from singleton groups (i.e., workgroups with only one observation or no within-group variation) and those collinear with fixed effects are excluded by the estimator due to perfect collinearity. Robust standard errors clustered at the workgroup level in parentheses. *Abbreviation:* FEs = Fixed Effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Deadwoodness

Panel A: Deadwoodness is associated with higher pendency of non-low performers

	Pendency of non-low performers			
	(1)	(2)	(3)	(4)
Deadwoodness	1.047***	1.048***	1.014***	1.008***
	(0.050)	(0.050)	(0.105)	(0.106)
Fiscal year FEs	Yes	Yes	Yes	Yes
Workgroup FEs	No	No	Yes	Yes
Workgroup size and age controls	No	Yes	No	Yes
R-squared	0.845	0.847	0.868	0.869
Mean(Pendency of non-low performers)	11.517	11.517	11.517	11.517
Number of workgroups	70	70	70	70
Number of observations	974	974	974	974

Notes: Pendency is a workgroup's average true first action pendency of non-low performers in a given fiscal year. A positive coefficient indicates an increase in the amount of time it takes for an applicant to receive the first decision from the examiner. Deadwoodness is the individual examiner fixed effects extracted from a linear regression of the true first action pendency on fiscal year and examiner fixed effects, only including each examiner's post-peak years. The individual examiner fixed effects are then Bayesian shrunk, and the average is computed at the workgroup-fiscal year level. Observations from singleton groups (i.e., workgroups with only one observation or no within-group variation) and those collinear with fixed effects are excluded by the estimator due to perfect collinearity. Robust standard errors clustered at the workgroup level in parentheses. *Abbreviation:* FEs = Fixed Effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Interquartile effect of deadwoodness on pendency of non-low performers

	(1)	(2)	(3)	(4)
Change in pendency of non-low performers	2.04	2.04	1.96	1.97

Notes: This table reports the estimated percent change in pendency from an interquartile shift in deadwoodness (25th to 75th percentile), calculated as $100 \times \left(e^{\hat{\beta}\times\mathrm{IQR}} - 1\right)$. The interquartile range of deadwoodness is 1.95.

Table 5: Above average percent of deadwood is associated with more exits, especially highly educated examiners

	Exits			Entries		
	(1) Overall	(2) High Performer	(3) High Education	(4) Overall	(5) High Performer	(6) High Education
Above average percent of deadwood	0.266**	0.199	0.483***	-0.118	-0.091	-0.187
	(0.104)	(0.314)	(0.182)	(0.107)	(0.130)	(0.148)
Above average percent of chronic low	0.173**	-0.006	-0.059	-0.104	-0.129	-0.196
	(0.084)	(0.273)	(0.144)	(0.080)	(0.117)	(0.145)
Pseudo R-squared	0.239	0.125	0.214	0.689	0.546	0.567
Log likelihood	-1309.314	-280.740	-482.870	-2006.594	-989.740	-767.909
Number of workgroups	70	53	64	76	75	71
Number of observations	944	694	874	907	900	865

Notes: Estimates are from quasi-maximum likelihood Poisson regressions, with total examiners as an exposure variable. High performer are those who were consistently above the median in the previous (following) three years for exits (entries). High education is defined as those with a professional and/or doctorate degree. All models include fiscal year and workgroup fixed effects, as well as workgroup age as a control. Observations from singleton groups (i.e., workgroups with only one observation or no within-group variation) and those collinear with fixed effects are excluded by the estimator due to perfect collinearity. For interpretation, $\exp(\hat{\beta})$ represents the factor by which the expected number of exits (or entries), adjusted for group size, changes when a workgroup has an above average percent of deadwood. For example, if $\exp(\hat{\beta}) = 1.2$, then the expected number of exits per examiner in workgroups with an above average percent of deadwood is 20% higher than in those without. Robust standard errors clustered at the workgroup level in parentheses. *Abbreviation:* FEs = Fixed Effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: *Alice* disproportionately affects pendency of non-low performers in workgroups with above average percent of deadwood

	(1) Pendency of non-low performers	(2) Overall Exits	(3) Overall Entries
Above average percent of deadwood	0.672**	0.183	-0.009
	(0.285)	(0.148)	(0.127)
Above average percent chronic low	0.746***	0.006	-0.062
	(0.240)	(0.116)	(0.068)
Post-*Alice*	0.949	-0.146	-3.426***
	(1.226)	(0.282)	(0.745)
Post-*Alice* × *Alice*-affected	-0.347	-0.342	-1.498***
	(0.498)	(0.312)	(0.539)
Post-*Alice* × Above average percent of deadwood	-0.330	0.133	-0.964***
	(0.357)	(0.211)	(0.279)
Affected by *Alice* × Above average percent of deadwood	-0.655	-0.000	0.181
	(0.549)	(0.245)	(0.236)
Post-*Alice* × *Alice*-affected × Above average percent of deadwood	1.351**	0.364	0.997*
	(0.644)	(0.272)	(0.557)
Post-*Alice* × Above average percent chronic low	-0.293	0.166	-0.200
	(0.353)	(0.157)	(0.288)
Affected by *Alice* × Above average percent chronic low	-1.093**	0.330*	-0.329
	(0.444)	(0.174)	(0.360)
Post-*Alice* × *Alice*-affected × Above average percent chronic low	-0.441	0.186	0.980
	(0.968)	(0.237)	(0.665)
R-squared	0.816		
Pseudo R-squared		0.244	0.684
Log likelihood		-1300.353	-2041.748
Number of workgroups	79	70	76
Number of observations	1,187	944	907

Notes: Pendency is a workgroup's average true first action pendency of non-low performers in a given fiscal year. A positive coefficient indicates an increase in the amount of time it takes for an applicant to receive the first decision from the examiner. Column 1 is a linear regression, while Columns 2 and 3 are quasi-maximum likelihood (QML) Poisson regressions, with total examiners as an exposure variable. All regressions include fiscal year and workgroup fixed effects, as well as workgroup size and age controls (no size control for QML regressions with exposure variable). Alice-affected is omitted because it is fully captured by workgroup fixed effects, given that the Alice shock is defined at the workgroup level. All regressions include fiscal year and workgroup fixed effects, as well as workgroup size and age controls (no size control for QML regressions with exposure variable). Observations from singleton groups (i.e., workgroups with only one observation or no within-group variation) and those collinear with fixed effects are excluded by the estimator due to perfect collinearity. Robust standard errors clustered at the workgroup level in parentheses. *Abbreviation:* FEs = Fixed Effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 7: Examiner Behaviors

Panel A: Discretionary Effort

(a) Claim Reductions

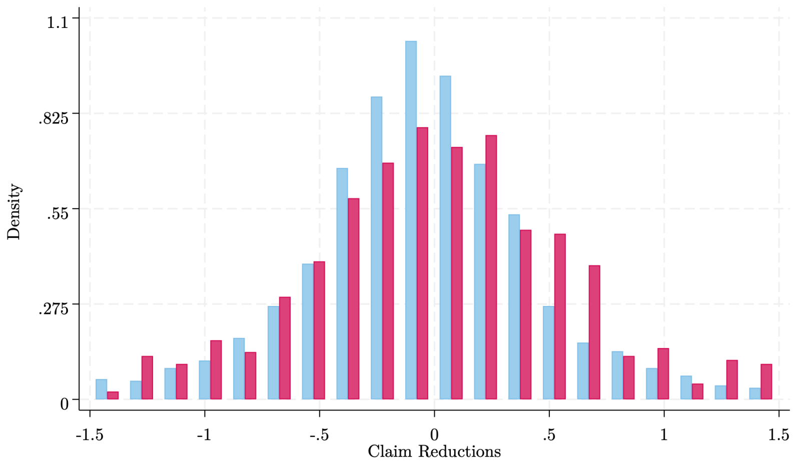

(b) Examiner-Inserted Citations

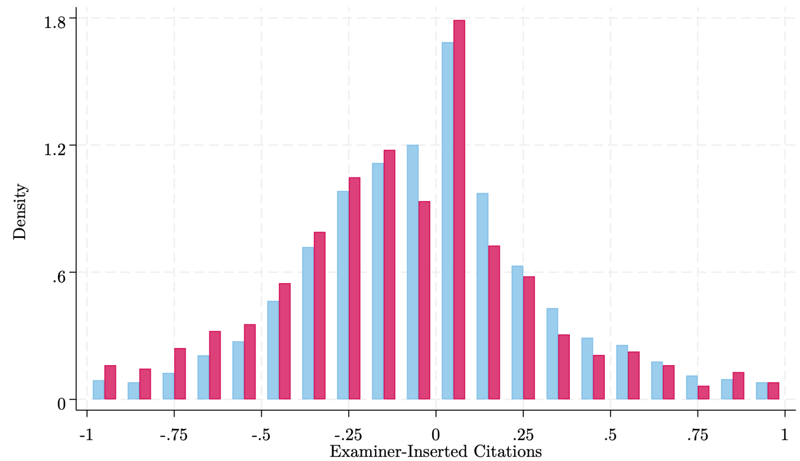

Panel B: Decision Style and Risk Orientation

(c) Leniency

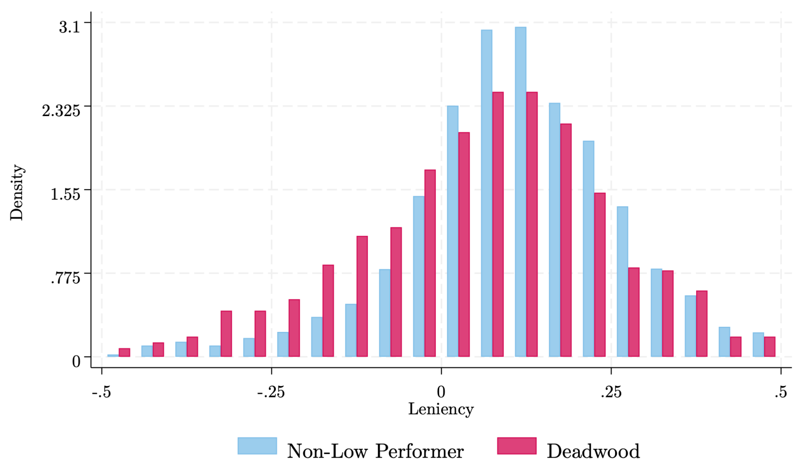

Panel C: Responsiveness to production incentives

(d) Fraction of Production in Last Fiscal Month

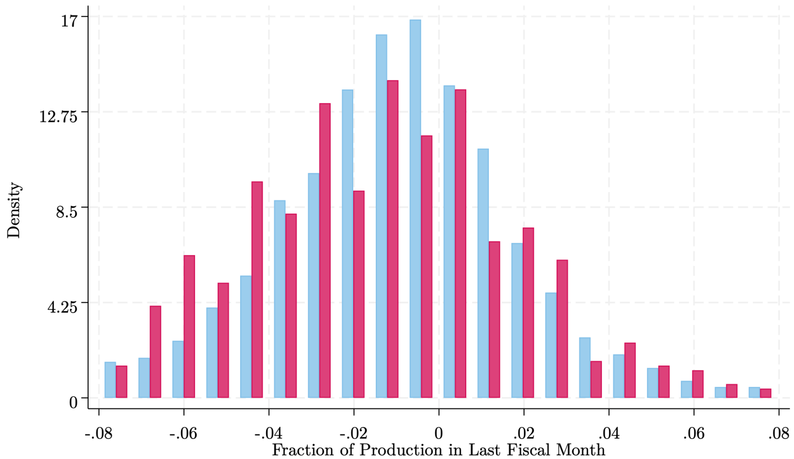

(e) Allowance Rate in Last Fiscal Month

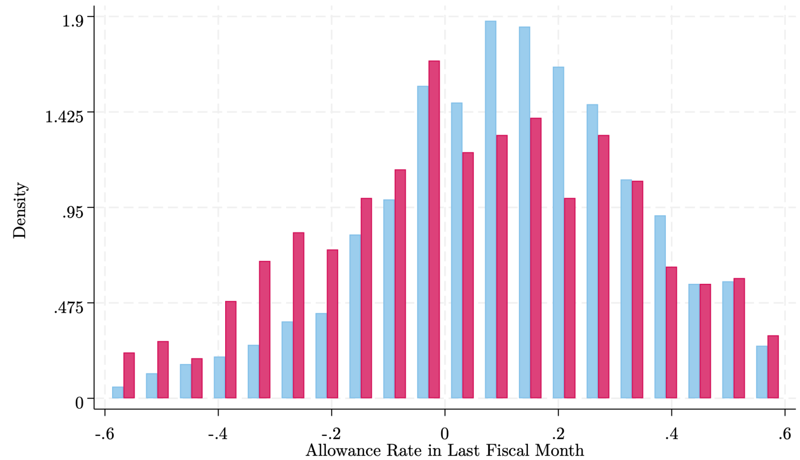

Panel D: Examination Quality

(f) Appeals Affirmation Rate

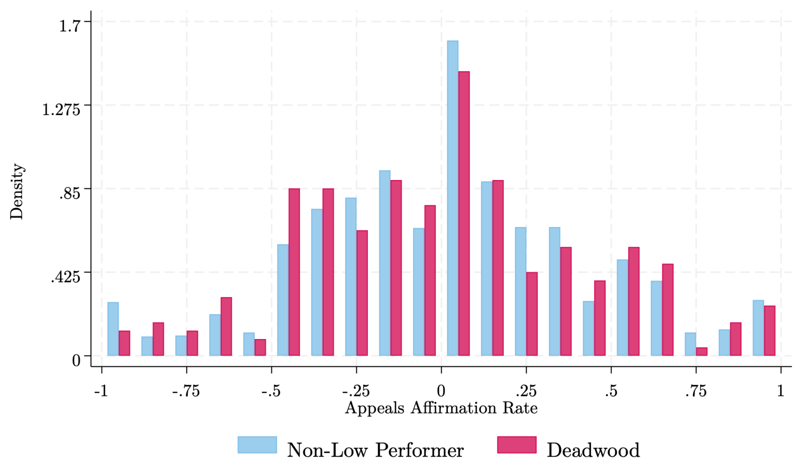

Notes: This figure shows the distributions of the difference between the post-peak and pre-peak means for deadwood and non-low performing examiners.

Table 7: T-tests comparing deadwood and non-low performer examiner behaviors

	Deadwood		Non-Low		
	N	Mean	N	Mean	Difference
Panel A: Discretionary Effort					
Claim Reductions	647	-0.026	4,250	-0.155	0.129*
Examiner-Inserted Citations	647	-0.120	4,250	-0.030	-0.090***
Panel B: Decision Style and Risk Orientation					
Leniency	779	0.061	4,717	0.116	-0.054***
Panel C: Responsiveness to Production Incentives					
Fraction of Production in Last Fiscal Month	651	-0.018	4,241	-0.011	-0.007***
Allowance Rate in Last Fiscal Month	651	0.067	4,241	0.132	-0.065***
Panel D: Examination Quality					
Appeals Affirmation Rate	235	-0.002	1,940	0.013	-0.015

Notes: This table shows the means and corresponding t-tests between deadwood and non-low performers for each of the measures of examiner behavior. Means are of examiners' post-peak minus pre-peak performance. The samples are restricted to examiners who had both pre- and post-peak data for all of the measures in the corresponding panel. Difference is the deadwood mean minus non-low performer mean. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix A. Production Count

For cleaning the data and constructing prosecution rounds, I follow the assumptions described in Sampat and Williams (2019). For constructing the production counts, I refer to the USPTO Joint Labor and Management Count System Task Force's guidelines (USPTO Joint Labor and Management Count System Task Force 2010):

Until February 14, 2010, examiners earned 1.0 count for the first office action and 1.0 count for a disposal, which could be either a final rejection or an allowance, regardless of whether it occurred during the initial examination or a subsequent round.

The USPTO then implemented changes to reduce the credit awarded for continued examination. Specifically, the count allocation for actions following the FAOM and for RCEs was decreased. As a result, applications entering second or later rounds of prosecution yielded fewer than 2.0 counts per round, lowering examiner credit for work on previously examined applications.

Before February 14, 2010:

Original Case (Non-RCE)			1st RCE			2nd & Subsequent RCEs			Counts	
FAOM	Final	All/Abn	FAOM	Final	All/Abn	FAOM	Final	All/Abn		
1.00		1.00							2	Original
1.00		1.00	1.00		1.00				2	1st RCE
1.00		1.00	1.00		1.00	1.00		1.00	2	2nd & Subsequent RCEs

After February 14, 2010:

Original Case (Non-RCE)			1st RCE			2nd & Subsequent RCEs			Counts	
FAOM	Final	All/Abn	FAOM	Final	All/Abn	FAOM	Final	All/Abn		
1.25	0.25	0.5							2.00	Original
1.25	0.25	0.5	1.00	0.25	0.5				1.75	1st RCE
1.25	0.25	0.5	1.00	0.25	0.5	0.75	0.25	0.5	1.50	2nd & Subsequent RCEs

where RCE = Request for Continued Examination, FAOM = First Action on the Merits, Final = Final Office Action, and All/Abn = Allowance, Abandonment, or any other Disposal.

Appendix B. Tables and Figures

Figure B.1: Number of examiners in art unit-fiscal year

(a) Histogram of art unit size

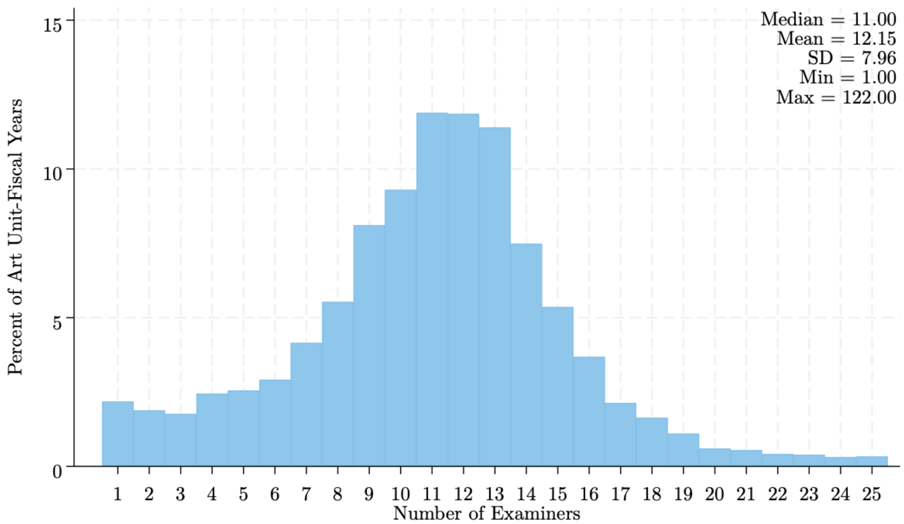

(b) Kernel density of art unit size

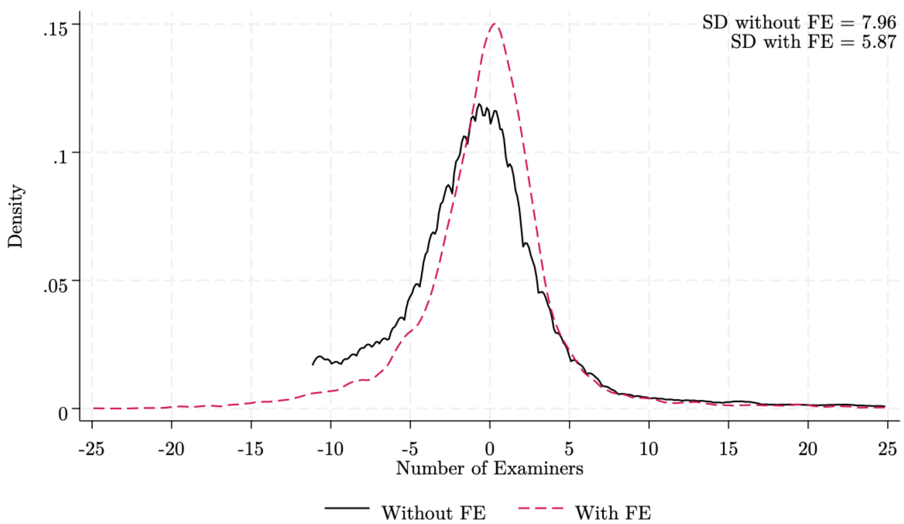

Notes: These figures show the distributions of the number of patent examiners in each art unit in a given fiscal year. Art units in art group-fiscal years with fewer than 5 examiners are dropped. The histogram shows the raw number of examiners in an art unit-fiscal year. The kernel density shows the number of examiners in an art unit-fiscal year, centered at the mean, in blue ("without FEs") and the residuals from a linear regression of the number of examiners in an art unit-fiscal year on art unit and fiscal year fixed effects in red ("with FEs"). *Abbreviation:* FE = Fixed effects.

Figure B.2: Number of examiners in workgroup-fiscal year

(a) Histogram

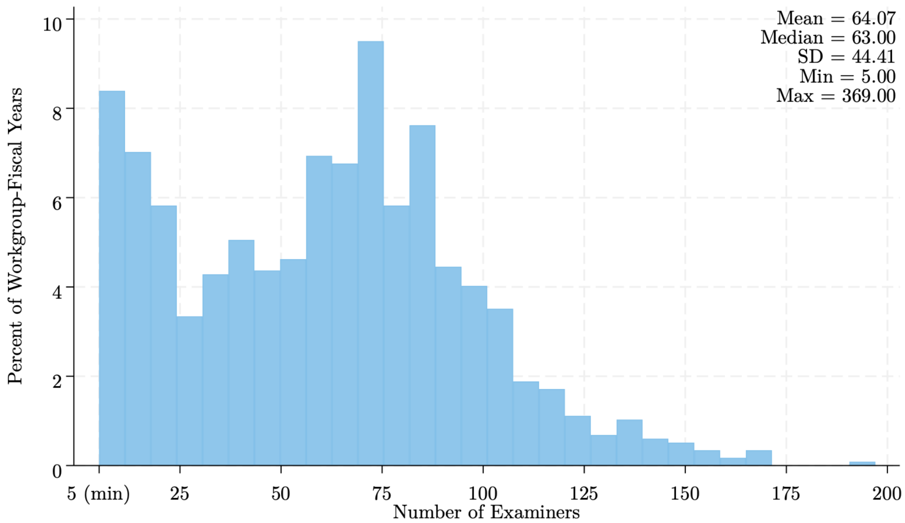

(b) Kernel density

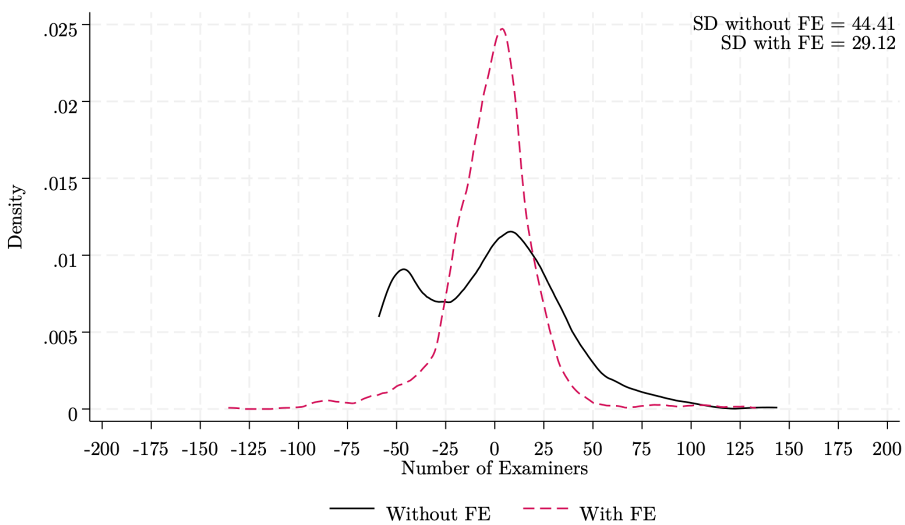

Notes: This figure shows the distributions of the number of patent examiners in each workgroup in a given fiscal year. Workgroup-fiscal years with fewer than five examiners are dropped. The histogram shows the raw number of examiners in an workgroup-fiscal year. The kernel density shows the number of examiners in a workgroup-fiscal year, centered at the mean, in blue ("without FEs") and the residuals from a linear regression of the number of examiners in a workgroup-fiscal year on workgroup and fiscal year fixed effects in red ("with FEs"). *Abbreviation:* FE = Fixed effects.

Figure B.3: Deadwoodness of a workgroup-fiscal year

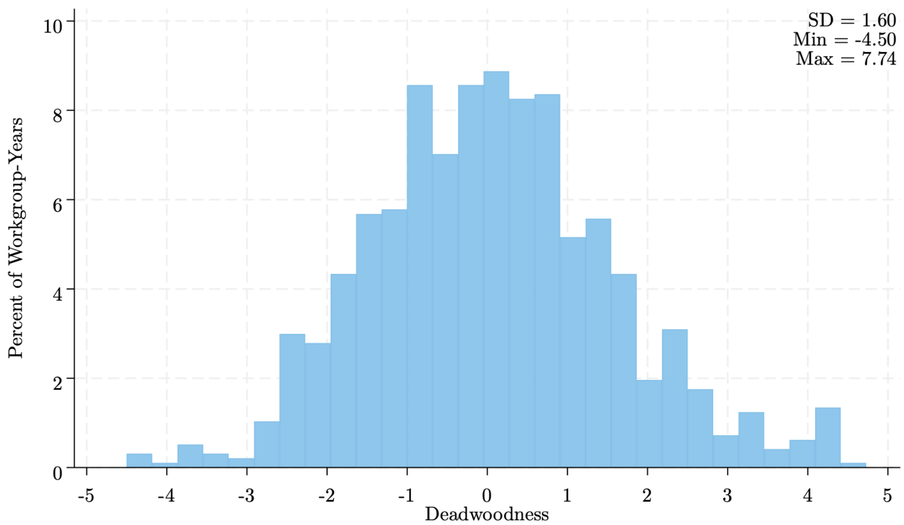

Notes: This figure shows the distribution of average deadwoodness in a workgroup in a given fiscal year. Deadwoodness is the individual examiner fixed effects extracted from a linear regression of true first-action pendency on fiscal year and examiner fixed effects, including only post-peak years for each examiner. The individual examiner fixed effects are then Bayesian shrunk, and the average is computed at the workgroup-fiscal year level.

Table B.1: Percent of deadwood examiners in workgroup is associated with higher pendency for non-low performers

	Pendency of non-low performers			
	(1)	(2)	(3)	(4)
Percent of deadwood	0.095**	0.099**	0.018	0.020
	(0.039)	(0.038)	(0.030)	(0.029)
Percent of chronic low	0.026	0.024	0.026	0.028
	(0.018)	(0.019)	(0.018)	(0.019)
Fiscal year FEs	Yes	Yes	Yes	Yes
Workgroup FEs	No	No	Yes	Yes
Workgroup size and age controls	No	Yes	No	Yes
R-squared	0.645	0.648	0.803	0.807
Mean(Pendency of non-low performers)	11.517	11.517	11.517	11.517
Number of workgroups	79	79	79	79
Number of observations	1,187	1,187	1,187	1,187

Notes: Pendency is a workgroup's average true first action pendency of non-low performers in a given fiscal year. A positive coefficient indicates an increase in the amount of time it takes for an applicant to receive the first decision from the examiner. Observations from singleton groups (i.e., workgroups with only one observation or no within-group variation) and those collinear with fixed effects are excluded by the estimator due to perfect collinearity. Robust standard errors clustered at the workgroup level in parentheses. *Abbreviation:* FEs = Fixed Effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table B.2: Peak and Post-Peak Productivity by Examiner Type

Performer Type	N	Years to Peak	Peak	Post-Peak	Percent Drop
Deadwood	1,083	4.06 (2.90)	136.96 (42.42)	81.62 (29.35)	40.40
Chronic Low	685	3.82 (2.85)	109.27 (40.23)	74.42 (36.27)	31.91
Non-Low	7,031	5.90 (3.70)	169.33 (56.25)	133.76 (44.51)	21.01

Notes: Less than 1% of deadwood, 6.0% of poor fits, and 11.6% of non-low performers do not have a post-peak period (i.e., last year of performance data is the peak).