Cutting Delays, Not Corners: The Role of Bureaucratic Performance in Housing Supply

Suvy Qin*
UC Berkeley

November 7, 2026

Job Market Paper

Click here for most recent version.

Abstract

Regulatory frictions are seen as a constraint on housing supply, yet the role of bureaucratic capacity in enforcing these rules remains largely unstudied. This paper examines how bureaucratic performance in the permitting process affects the timely provision of housing. Using novel administrative data from Los Angeles, I find that a one standard deviation faster bureaucrat issues permits in 46 fewer days on average. Faster bureaucrats are more likely to work voluntary overtime and to be hired in years with more competitive public sector pay. Higher-performing bureaucrats can manage a workload increase during busy periods without taking longer to review permits. I use an instrumental variables design to estimate how being assigned to a faster bureaucrat affects project quality, finding no measurable effect on the likelihood of a future code violation. Back-of-the-envelope estimates highlight the cost-effectiveness of increasing bureaucratic capacity through alternative personnel policies that improve the quality, quantity, or selection of plan checkers.

*Correspondence: suvyq@berkeley.edu. I am grateful to Max Auffhammer, Reed Walker, Guo Xu, Cailin Slattery, and Jim Sallee for their invaluable guidance and advice. I thank Pat Kline, Stephanie Kestelman, Antoine Levy, Kirill Borusyak, Tim McQuade, John Voorheis, participants at the 2026 North American Meeting of the UEA and WEAI GSW, and seminar participants at UC Berkeley for their helpful discussions and feedback. I thank Stephanie Park, Michael Pelham, Chad Doi, Todd Borzi, and Brian Ten for generously sharing their insights on institutional details. Joshua Kwong provided excellent research assistance. This work benefited from the generous support of the Fisher Center for Real Estate and Urban Economics and the Institute for Research on Labor and Employment Dissertation Fellowship. This paper was previously circulated under the title "Does Government Capacity Matter for Housing Production? Evidence from Los Angeles Bureaucrats." The views expressed are solely those of the author. Please do not cite without author's permission.

1 Introduction

Housing affordability has become one of the key policy challenges facing the United States, as one in three households spent more than 30% of their income on housing in 2024 (CRS, 2026). Policymakers and researchers argue that the lack of housing availability may be due to an increasingly restrictive regulatory environment, leading to costly permitting delays and lower productivity in the construction sector (Baum-Snow and Duranton, 2025; Garcia and Molloy, 2025; Soltas and Gruber, 2026). For example, in this paper's setting of Los Angeles, the average permit issuance time for new residential housing reached 400 days in 2019, which is more than double what it was in the early 2000s. This trend holds across various residential occupancy types and cities throughout the US (D'Amico et al., 2026).

Understanding the causes of permitting delays is important for addressing the lack of housing affordability. This paper focuses on a relatively understudied factor: the role of bureaucratic capacity. Past work argues that costly delays might be due to two aspects of the regulatory environment: strict land-use regulations that reduce housing supply (D'Amico et al., 2024; Gyourko and Molloy, 2015), and "not-in-my-backyard" (NIMBYism) pushback arising from the decentralized nature of these regulations (Mast, 2024; Brooks and Liscow, 2023; Ortalo-Magné and Prat, 2014). Additionally, housing regulations are notoriously complex, often requiring local bureaucrats to oversee approvals from an average of four different government agencies per project in a city like Los Angeles. Even as a growing literature has documented the importance of bureaucratic state capacity for delivering public goods and services at low cost (Liscow et al., 2024; Best et al., 2023), we know little about how bureaucrats tasked with *regulatory* authority contribute to the state's ability to effectively govern.

This paper evaluates whether and how bureaucratic state capacity impacts the effective enforcement of existing housing regulations. I study this at the municipal level, as local government agencies have regulatory discretion over housing development in their jurisdictions. Municipal building departments are typically tasked with ensuring that zoning, building safety, and performance standards are met during the permitting and construction phases of a project. I focus on the local bureaucrats responsible for verifying that each project complies with various housing regulations prior to permit issuance, who are known as *plan checkers*. Plan checkers are engineers who review building plans, issue corrections, and assign additional approval steps from other city agencies. I assemble novel administrative permit-level data from the Los Angeles Department of Building & Safety (LADBS) that allow me to observe the assigned plan checker and timestamped status updates for every new residential project over the past 25 years.

Using this newly assembled dataset, I provide new evidence on how bureaucratic performance affects a local government's regulation of housing. This paper answers three questions: (1) how much of the permitting timeline variation is due to the performance of the plan checker? (2) what do faster plan checkers do differently? (3) how does bureaucrat performance affect future project outcomes and quality? By linking permit data with payroll information of plan checkers, I also provide suggestive evidence of what predicts plan checkers' performance and illustrative estimates of the cost savings from shortening permitting timelines via alternative personnel policies.

In the first part of the paper, I use a two-way fixed-effects model to decompose the variation in permit issuance times into components attributable to the plan checker and the project applicant (i.e., the architect/engineer who submits the permit application). The permit-level data allow me to leverage a "switchers" design from project applicants assigned to different plan checkers on consecutive projects. After adjusting the variance components for sampling variability, I find that plan checkers explain 4% of the residual variation in permit issuance time, which implies that a one standard deviation faster plan checker on average issues permits in 15.4% less time (46 days) for new residential projects. The average permit issuance time is 296 days in my sample. The two-way fixed effects model relies on the assumption that the sequence of plan checkers an applicant is assigned to is unrelated to unobserved determinants of permit issuance time. I test for various forms of endogenous sorting between applicants and plan checkers that would bias the estimates of plan checker effects, and I do not find evidence of such sorting. I show that there is no sorting based on trends and that there are no performance gains from applicants being matched to a plan checker they have previously worked with.

Next, I ask what predicts being a faster plan checker. I use the plan checker fixed effects from the previous model as a measure of performance and correlate them with various individual attributes and project characteristics. I find that faster plan checkers were more likely to be hired in years with more competitive public sector pay relative to the private sector. Faster plan checkers are also more likely to choose to work overtime. In this setting, overtime is voluntary, and overtime work is distributed fairly and evenly among plan checkers who request to work overtime. I provide evidence that the plan checker fixed effects reflect speed through productivity, net of hours worked, suggesting that overtime and productivity are complementary, as overtime is more valuable to more productive plan checkers. To provide additional evidence of what a faster plan checker does differently, I implement a differences-in-differences design that leverages a workload shock due to increased permit filings from anticipated building code changes. Permits filed during the workload shock that are assigned to slower plan checkers take 20% longer to be issued than similar permits

assigned to faster plan checkers. In the weeks after the workload shock, faster plan checkers increase their overtime hours by an additional two hours per pay period relative to slower plan checkers, suggesting that the faster plan checkers are more willing to work overtime to adapt to the backlog during busy periods so that permit issuance times do not increase.

After documenting what determines plan checker speed, I estimate the causal impacts of being assigned a faster plan checker on project outcomes to examine whether there are adverse quality impacts of plan checker speed. For example, being assigned a faster plan checker might affect a project's future construction quality if faster plan checkers were less thorough during the permitting process. I proxy for project quality using future code enforcement cases and code violations. I use an instrumental variables design that exploits week-to-week variation in plan checker availability to deal with concerns of correlated permit-level shocks and measurement error in estimates of plan checker performance. I instrument for assigned plan checker performance using the availability of faster plan checkers in the same week and office that a permit was submitted in. This variation comes from staff leaves, rotations, and hiring that are unrelated to any individual project's characteristics. I provide evidence to support the validity of this instrumental variables design. I do not find strong evidence of negative quality impacts from being assigned a faster plan checker on building safety, as measured through future code violations. Being assigned a faster plan checker does not affect the likelihood that a project is completed within 3 years. Taken together, these results suggest that there are no measurable negative quality or extensive margin impacts from being assigned a faster plan checker; rather, faster plan checkers mainly affect permitting outcomes through time saved in the plan check process.

In the last section, I provide illustrative estimates of the value of adding bureaucratic capacity in the permitting process through three counterfactual scenarios that improve the quality, selection, and quantity of plan checkers. I estimate the days saved and the private value of the avoided delay time for each of the three alternative personnel policies. I convert the days saved into a monetary value and compute a return for each dollar spent (or ROI) on additional payroll costs.[1] First, I use the previous estimates of plan checker performance to calculate the days saved from replacing the 25% lowest-performing plan checkers. Replacing the 25% lowest-performing plan checkers with the median plan checker saves approximately 600,000 days and $263 million in private delay costs for builders, which translates to a ROI of $9.02. Secondly, I ask what the cost savings would be from increasing plan checker pay by 10%, using the prior estimates of how plan checker performance correlates with the public

[1] I convert the avoided delay time into a monetary value using reported project values and a simple accounting framework that aggregates financing costs, carrying costs, and foregone rents. Since LADBS operations, including salaries, are funded by permit fees, the ROI can be thought of as the private return from each additional dollar paid in permit fees.

sector pay gap to infer how higher pay might increase positive selection. I calculate a 10% increase in plan checker pay would save builders $125.8 million, with a ROI of $3.77. Finally, I simulate a queueing model from the operations literature to estimate the avoided delay time from increasing staffing by 5%. Simulations suggest that increasing plan checker staffing by 5% would save builders approximately 81,000 total days, with a ROI of $2.16. My results suggest that increasing bureaucratic capacity can be a complement to any policy reform that aims to streamline the permitting process.

This paper contributes to three separate, but related literatures. First, there is an increasing body of research studying how land-use regulations affect housing supply, as recently reviewed in Baum-Snow and Duranton (2025). Much of this literature argues that a burdensome regulatory environment, due to policies such as building height restrictions, minimum lot sizes, and zoning codes, reduces housing supply and raises housing costs (Glaeser and Gyourko, 2003; Glaeser et al., 2005; Glaeser and Gyourko, 2018; Gyourko and Krimmel, 2021; Gyourko et al., 2021; Gyourko and McCulloch, 2023; Bartik et al., 2025; Song, 2025). However, little attention has been devoted to understanding how bureaucratic state capacity, or governments' ability to enforce these regulations, might influence housing supply. This paper is the first to provide evidence of how the actions of individuals in local regulatory agencies affect the timely provision of housing.

Secondly, this paper contributes to the literature studying how bureaucrats affect public service provision and the overall effectiveness of the state. Much of this literature has focused on understanding the key inputs to bureaucratic performance, such as personnel selection, financial and career incentives, and task management (Best et al., 2023; James, 2024; Muñoz and Prem, 2024; Muñoz and Otero, 2025; Metcalfe et al., 2023; Dahis et al., 2025; Dal Bó et al., 2013; Aneja and Xu, 2024; Bertrand et al., 2020). Even though evidence has shown that the quality of personnel in government infrastructure projects can impact costs (Liscow et al., 2024), there is limited research on how the performance of bureaucrats affects *regulatory enforcement*. I provide new measurement of bureaucrat performance in a setting with high regulatory authority. I also provide suggestive evidence of voluntary overtime pay as a potential channel for encouraging higher performance in a public sector setting that features rigid pay scales and promotion patterns.

Finally, this paper contributes to a growing literature studying the costs and potential causes of regulatory delay in housing projects. Research has estimated the implicit costs of regulatory delay and its impacts on housing production (Wrenn and Irwin, 2015; Gabriel and Kung, 2026; Soltas and Gruber, 2026; D'Amico et al., 2026), finding that wait times are a main component of aggregate regulatory costs that increase the length and uncertainty of housing development. Recent work has also focused on evaluating the impacts of one part of

the permitting process (discretionary or design review by local planning departments) that some projects are subject to, finding that discretionary review increases the development timeline (Manville et al., 2023; Gabriel et al., 2026) but has modest impacts on aggregate housing supply (Kestelman, 2025; Gold, 2025). This paper adds to this literature by evaluating the potential role of bureaucratic capacity in permitting delays and its implications for policy efforts to streamline permitting.

2 Institutional Background

In Los Angeles, every construction project begins with an application for a building permit to the Department of Building & Safety (LADBS). Typically, the permit application is filed by the person who is responsible for designing the project and developing the building plans, either a licensed engineer or architect (hereafter, the applicant). Once an application is filed, it will be assigned to intake staff who check it for completeness. If it is deemed complete, the application is considered submitted and then a manager will assign it on a first-come-first-serve basis to a *plan checker*. The plan checker will review the application to ensure the project complies with various regulations, such as building codes, zoning codes, and other performance standards. The plan checker will then issue initial corrections for the applicant to address if needed and identify which additional approvals (clearances) the project needs to obtain from other city or county agencies (hereafter, approval agencies). Finally, before issuing the permit, the plan checker verifies that all clearances and corrections were addressed. Once the permit is issued, the project can begin construction. Throughout construction, the project will be inspected by staff from a separate LADBS division to ensure compliance.

Figure 1 shows the process during the plan check phase (the time between project submittal and permit issuance) in Los Angeles, while Figure 2 shows the entire permit and construction process. I refer to the time between project submittal and assignment to a plan checker as *queue time*, the time between assignment to a plan checker and initial corrections being issued as *initial review time*, the time between initial corrections being issued and plan check approval as *clearances and corrections time*. Depending on the project, an application could end up going through several rounds of back-and-forth with the plan checker during the clearances and corrections time (for example, if corrections are not adequately addressed or changes were required by other agencies). A plan checker has many concurrent project applications that they are responsible for during the initial review period and the clearances and corrections period.

Detailed permit records that include timestamped updates (at the daily level) allow

me to infer how long each phase takes and the name of the staff person responsible for verifying each respective update, as well as the clearances a project had to obtain and from which approval agencies (Figure A1). It is not uncommon for a project to be required to obtain half a dozen clearances from agencies such as the Bureau of Engineering, Bureau of Sanitation, Bureau of Street Services, City Planning, LADBS, and Fire Department before the building permit could be issued.

A key feature important for later identification is that applications are assigned to plan checkers based on the order in which they were received. According to LADBS staff, queue-based project assignment and office rotations of plan checkers are used to prevent developers from forming relationships with certain plan checkers. Managers in each office assign applications weekly, and they consider the project's basic attributes in who gets assigned which projects.[2] Re-assignments are rare.

Plan checkers are the key bureaucrats in this process, as they assign corrections to building plans and verify clearances that each project must obtain from other city agencies. While plan checkers are not responsible for how project applicants obtain clearances, plan checkers may suggest ways to modify building plans in order to comply with various codes. LADBS staff indicate that in theory, plan checkers do not have much discretion over what clearances are assigned, as plan checkers use a "Clearance Handbook" in determining which clearances apply for a given project based off of the project attributes (i.e. location, zoning, project size, etc). However, if there is a conflict in the various codes or a gray area, plan checkers are supposed to apply the rules according to the "spirit and intent" of the code.

Plan checkers are civil engineers by training, and their formal job title is "Structural Engineering Associate." Typically, an entry-level position is filled by recent college graduates who hold a Bachelor's degree in civil engineering. Over time, junior plan checkers can be promoted as they gain experience and seniority. Plan checkers' salaries are determined by city-wide pay scales negotiated by unions. Promotions are based on rules set by the city's civil service commission. A common career pathway is to begin as a plan checker in the "Structural Engineering Associate" job class (which has up to four pay grades) before being promoted to managerial roles that do not perform plan check review, as described in the city-wide career ladder shown in Figure 3.[3]

[2]For example, all plan checkers are equipped to review plans for low-density projects, such as single-family dwellings, duplexes, and accessory dwelling units, which are over 90% of all submitted projects in my sample. However, a new hire would not be assigned to an apartment building, for example, so controlling for project characteristics is important.

[3]Manager positions are those with the title "Structural Engineer" or "Building Civil Engineer." Plan checkers with a Structural Engineering Associate III title are considered experienced enough to be able to handle any kind of project. Structural Engineering Associate IV is meant for plan checkers with specialized expertise, so fewer plan checkers have this job title.

Every week, plan checkers can also request to work overtime, which pays 1.5 times their regular hourly pay. Not all overtime requests are granted depending on the backlog of permit applications. Overtime is not mandatory, and managers are instructed to assign projects as fairly as possible so that overtime requests are fulfilled proportionally among plan checkers who request it.[4] Managers cannot work overtime.

3 Data

I assemble several novel administrative datasets to establish descriptive facts and for use in the main empirical analyses. Appendix A provides additional detail on data construction.

Los Angeles building permits. LADBS tracks each status update of a building permit application, from when it was filed to its issuance. Importantly, each status update also includes the date of the update and the name of the individual plan checker who signed off on it. Over 92% of building permits in my sample also identify the engineer, architect, and/or contractor of the project; I designate as the *project applicant* the first non-missing individual in the trio of possible project contacts. I web scrape these permits to compile who was the assigned plan checker, project applicant information, additional permit characteristics, and the dates of each clearance and status update (i.e. the information shown in Figure A1 and additional inspection information not shown). This allows me to calculate the length of time each step takes and construct the workload of each plan checker. I also have additional information of each project's construction phase (after permit issuance) that details inspection trips and dates. To convert the clearance text into structured data, I reference the "Clearance Handbook" to identify the agency responsible for each clearance. The LADBS permit data reliably cover the years 1997-2025.

Los Angeles code enforcement cases and code violations. A separate division within LADBS is responsible for code enforcement cases and investigating code violations, which I web scrape and link to the previous data at the parcel-address level. Of the code enforcement cases that lead to a code violation, I web scrape the code violation descriptions to classify each code violation as being safety-related or not. I classify a code violation as safety-related if the description mentions words related to fire safety ("Missing smoke detectors"), electrical safety ("The building or premises is substandard due to hazardous electrical wiring"), structural

[4]If there are not enough permit applications to completely fulfill everyone's overtime requests, then applications will be assigned so that everyone has a similar fraction of overtime work. For example, if there are only enough applications to fulfill everyone's overtime requests at 50%, then everyone who requested overtime will be assigned applications at roughly 50%. This is based off both the memorandum of understanding negotiated between the City of Los Angeles and the Service Employees International Union, Local 721 (which represents the professional engineering and scientific unit) and conversations with LADBS managers.

hazard ("The floors and floor coverings are unsafe due to inadequate maintenance"), blocked entryways ("No means of egress"), general habitability ("The premises are substandard due to an accumulation of weeds, vegetation, junk"), as well as general un-permitted work for residential use ("Electrical work has been done without the required permits, inspections").

Salary & personnel records. Through public records requests to the Controller's Office at the City of Los Angeles and Los Angeles Open Data portal, I obtain pay-period and annual salary data. The biweekly data cover staff who were employed between 2005-2024 and report the person's name, job title, base pay, overtime pay, regular hours, overtime hours, hire date, separation date, and other pay. I use job titles to identify promotions and separations and to classify plan checkers as senior or not. For a subset of plan checkers in the LADBS permit data, I am able to find the public LinkedIn profiles of these individuals and obtain information on educational history, highest degree attained, and/or past experience. I am able to find 51% of plan checkers in my sample (out of 400 plan checkers), but only 41% have non-missing education information. I merge these data to the LADBS permit data using the names of the plan checkers.

Professional licensing data. I obtain public data on when licenses for project applicants were issued from the California Department of Consumer Affairs, which regulates the issuance of professional licenses for engineers, architects, and contractors. I use the license number and type to match project applicants in the LADBS building permit data, which allows me to observe project applicants across multiple permit submissions. Using the license as an identifier, I also construct a variable that tracks how many past projects a given project applicant has submitted.

Geographical covariates. 90% of the LADBS building permits have geographical coordinates that I use to spatially intersect them with various city-wide datasets that contain information on land-use density, hillside zones, historical preservation overlay zones, and special plan areas to create indicator variables for whether a project is in one of these areas with greater regulatory stringency. I combine the spatially intersected land-use density information with a LADBS-provided variable to classify projects into density categories. I also use the coordinates to spatially intersect permits with a Digital Elevation Model to obtain elevation and slope. For a smaller subset of permits (17%), I am able to merge in ATTOM property tax deeds to identify the last parcel transaction of the empty/undeveloped lot.

US building permits. To provide descriptive facts about permitting trends, I use permit data purchased from a data vendor, Shovels, who aggregates permits from over 200 building permit-issuing departments across the United States. The data include the permit filing date, issuance date, completion date, address, and a description of the project.

3.1 Analysis sample

The main analysis sample comprises new residential permits that have an assigned plan checker, project applicant, and non-missing permit submittal and issuance dates. I exclude permits submitted by owner-builder applicants that do not have an architect, engineer, or contractor listed, as it is not possible to identify the same applicant over time. I also exclude a small number of supplemental permits (1% of all new residential permits), as these are not the main permit for a project and are often revisions of the main permits, so their permitting timelines may be compressed. Summary statistics comparing the main analysis sample and all new residential permits are shown in Table 1. Overall, the permits I analyze are relatively comparable to the main set, though the permits in my analysis sample spend an additional 9 days in the plan check phase and are slightly larger projects ($1.47 million vs $1.45 million) on average compared to the main set.

Table 2 presents summary statistics of plan checker characteristics. The first column includes all plan checkers for whom I observe different applicants being assigned to them (the analysis sample), while the second column includes all plan checkers. In the analysis sample, the plan checkers are more likely to be matched in the payroll data (83%) and to a LinkedIn profile (51%). Plan checkers who are assigned to different applicants have almost a year of prior work experience compared to 0.88 of a year for plan checkers in the full sample. Plan checkers who I observe separating from LADBS exit at around 12 years. Around 16% of plan checkers are promoted to manager roles. Plan checkers work an average of almost 12 hours of overtime biweekly, and their biweekly total pay is $5,532 in 2023 dollars. Demographic characteristics are also nearly identical between plan checkers in the analysis sample and full sample.

4 Descriptive Facts

Using the detailed permits data, I generate several descriptive facts about the permitting process in both Los Angeles and the United States. In Los Angeles, the average permit issuance time of new residential projects has more than doubled from the early 2000s (Figure 4). The average permit time is approximately 400 days for a permit submitted in 2019. This trend is observed nationally and not limited to Los Angeles. In the US, the average permit issuance time for new residential projects was 54 days in 2000, which has doubled to 120 days in 2019. In LA, the rise in permitting time is not driven by compositional changes in the types of projects that are being built, as the average clearances and permit issuance time have risen across four separate project types: single-family dwellings, duplexes, accessory dwelling units (ADUs), and condos/apartments (Figure A2). The permitting process also

involves more approval agencies and clearances than it once did (Figure A3). On average, new residential projects had to obtain 12 different clearances from 5 different agencies for permits submitted in 2019, compared to 9 clearances from 3 agencies for permits submitted prior to 2010.

These stylized facts highlight the significant variation across similar projects started at different times. To formally explore how much of the temporal variation is due to observable characteristics of the projects, I perform an Oaxaca-Blinder decomposition of the average difference in permit issuance time for new residential projects submitted during two time periods in my sample that correspond to two different policy regimes: pre-2011 and post-2011. Beginning in 2011, there were new municipal policies that increased the average number of clearances required per project. I decompose the mean difference in permit issuance times between the 2 permit groups as a component explained by differences in observable permit characteristics ("explained" effect) and a remaining component that is "unexplained" due to differential time impacts for the same observable permit characteristics. Table 3 displays the results for three different specifications that include more permit characteristics as one moves to the right of the table. In comparing columns 2 and 3, additionally controlling for the number of clearances assigned to a permit explains 49% of the average difference in log issuance times, but 51% of the gap remains unexplained. It is this unexplained component that the remainder of the paper focuses on.

5 The Effects of Plan Checkers on Permit Outcomes

This section explores how the performance of local bureaucrats affects the permitting process. I first develop a framework to quantify what share of permit timeline variation is due to the plan checker, adapting methods from the literature pioneered by Abowd et al. (1999). I discuss identification challenges and provide evidence to support the model specification that relies on project applicants being assigned different bureaucrats across consecutive projects. I perform a variance decomposition to evaluate the importance of plan checkers in the permitting process.

5.1 How much does plan checker performance matter for permitting outcomes?

The main job function of the plan checker can be thought of as allocating their time to verify each project satisfies building regulations. Thus, one key metric for evaluating the performance of a plan checker in this context is their impact on permit issuance time, which would be a function of the minimum necessary time a plan checker spent on the project to verify compliance with regulations. In this section, I leverage the assignment of plan checkers

to different projects to decompose permit issuance time into components attributable to the plan checker and the applicant.[5]

For a given development project i, I model the permit timeline outcome of log issuance time as a linear and additively separable function of project characteristics ($\mathbf{X}_i$), a component due to the assigned plan checker j ($\alpha_{j(i,t)}$), a component due to the applicant k (ϕ_k), time effects (τ_t), a vector of administrative attributes ($\boldsymbol{\nu}$), and an error term:

$$y_{it} = \mathbf{X}_i\beta + \alpha_{j(i,t)} + \phi_k + \tau_t + \boldsymbol{\nu} + \epsilon_{it} \tag{1}$$

in which $\mathbf{X}_i$ includes categorical variables for project height, project valuation, allowed density, clearances, and geographical fixed effects (planning zone, neighborhood, and council district), τ_t are year-month-office fixed effects, and $\boldsymbol{\nu}$ are fixed effects for plan check type and project use type. Equation (1) captures how permit issuance time may differ across two similar projects submitted by the same project applicant depending on the plan checker assigned to them, conditional on project attributes, seasonality, and macroeconomic shocks.[6] I interpret $\alpha_{j(i,t)}$ as the portable component of a plan checker's ability, which I will refer to as plan checker speed or plan checker performance.

Identification of the plan checker and applicant fixed effects relies on an exogenous mobility assumption: the sequence of plan checkers an applicant is assigned to is unrelated to permit-level shocks, conditional on project characteristics, the applicant, the office, and time period. Several features of this setting make this assumption plausible. Plan checkers do not determine which projects they receive, and project applicants do not determine which plan checkers they are assigned to. Rotation of plan checkers across offices and regular staffing changes from new hires and separations generate variation in the pool of plan checkers an applicant could be assigned to that is unrelated to project-level shocks. Note that the assumption of exogenous mobility does not require random assignment, as it allows for arbitrary correlation between applicants and plan checkers (i.e. better applicants assigned to the better plan checkers). Managers assigning projects based on any dimension of project difficulty captured by the observables or project applicant's portable component also do not violate the identifying assumptions.[7] I provide evidence to validate the model's assumptions

[5]The choice of who to interface with the LADBS plan checkers among the applicant team (engineer, contractor, or architect) depends on the project and what stage it is in when these professionals are brought on board, based on conversations with industry members.

[6]I model issuance time in logs because the distribution of issuance time has a long right tail and because plan checker effects are plausibly proportional, such that a slower plan checker lengthens every step rather than adding a fixed number of days.

[7]For example, if certain project applicants tend to submit poorly assembled building plans or more difficult projects, this would be captured in the project applicant effect, ϕ_k. In this setting, I argue that unobserved project-level shocks are realized after project assignment to a plan checker, which I discuss below.

in the next section.

5.1.1 Diagnostic checks

One typical concern with the model specified in Equation (1) is endogenous sorting between the plan checkers and applicants (i.e. sorting of better applicants to plan checkers whose performances are improving over time, for example). However, given the rotation of project assignment in my setting, I am able to leverage "switchers," or applicants who get assigned to different plan checkers on consecutive projects, as a way to test for endogenous mobility. One way to check for endogenous sorting is to plot how the outcome (average issuance time) changes as the same applicant is assigned to plan checkers of different levels of performance, based off of estimates of $\alpha_{j(i,t)}$ that are split up into quartiles. This is shown in Figure 5, which plots the average residualized issuance time for an applicant who was assigned a plan checker in the bottom or top quartile and then assigned to a different plan checker in the opposite quartile for their next project. The figure shows that being assigned a plan checker in the top quartile (i.e. the fastest plan checkers) after previously being assigned to one in the bottom quartile (i.e. the slowest plan checkers) decreases permit issuance time by more than 20%, as plotted in the solid black line. The symmetric effect between the two trajectories (of going from a lower- to higher-performing plan checker and vice versa, as seen when comparing the solid blue line with the solid black line) and the lack of pre-trends suggest that there is no systematic sorting between applicants and plan checkers that would bias estimates of α_j.

Additionally, I can use a similar set of switching applicants to see if relationship-based matching or expertise-based matching is important. To check for relationship-based matching, I first calculate how many and what share of an applicants' past projects were assigned to the same plan checker as the current project (hereafter, project shares). Then I classify switchers (applicants who are assigned to a different plan checker) as "high-contact" or "low-contact" based on whether the new plan checker is someone they have worked with at least 3 times on past projects *or* whether the currently assigned plan checker's project share is higher than the project share of the previously assigned plan checker. I do not find strong evidence of gains from relationship-based matching using both measures of high contact (Figure A4). While there seems to be a slight decline in permit issuance time after being switched to a plan checker with which an applicant has had "high-contact," this effect is both small and statistically insignificant.

To check for expertise-based matching, I create a measure that reflects how familiar a given plan checker is with a project type. I first group projects into 4 types (single-family, duplexes, ADUs, and condos/apartments), then I calculate for each project what share of

previous projects assigned to this plan checker were of the same type as the current one. I define switchers as being matched to a more experienced plan checker if that plan checker's share of past projects that were the same type is more than 50% and the previous plan checker's share was less than 50%. I do not find evidence of any matching effect from being assigned a more experienced plan checker (Figure A5). Rather, there might be a small increase in issuance time from being assigned a more experienced plan checker (though also statistically insignificant). Taken together, these exercises suggest that there is weak evidence of relationship-based and experience-based gains from matching.

Besides systematic sorting between plan checkers and applicants, there may remain other concerns of endogenous sorting between certain types of projects and plan checkers, since projects are not purely randomly assigned to plan checkers and certain unobserved project attributes (such as project complexity) might bias my estimates. However, I argue that any unobserved permit-level shock, such as difficulty, is realized *after* assignment and such shocks are mean zero across offices.[8] If there were other permit-level unobserved shocks that affected assignment towards certain plan checker types, then one would expect the ex-ante project covariates to be correlated with the plan checker fixed effects, $\hat{\alpha}_j$. This is not the case, as a joint test of significance yields a p-value of 0.10, indicating that there is not strong evidence of systematic selection in the assignment process (Figure A6). Similarly, if there was endogenous sorting of projects and plan checkers, then one might expect certain permit characteristics to be correlated with the plan checker assigned in the future period. I do not find evidence of plan checkers being assigned systematically different projects based on sorting from lagged permit characteristics (Table A1, column 2). I also check the sensitivity of the plan checker fixed effects in Equation (1) to different exclusion or inclusion of relevant covariates and controls. I show that once I control for time and administrative effects (τ_t and $\boldsymbol{\nu}$) and project characteristics ($\mathbf{X}_i$ excluding clearances), the estimated plan checker effects have a correlation of 0.94 with those from the baseline model (Table A2). Further controlling for plan checkers' weekly workload, years of experience, or their peers does not alter my estimates by much.[9] This suggests that the main sorting mechanism is via the timing of which plan checkers are working during different time periods and via permit types being assigned to certain plan checkers.

[8]The types of permit-level unobservables that would meaningfully affect permit issuance time are not likely to be obvious to the manager who assigns permits after a cursory review. These are attributes about the parcel (i.e. if a parcel has restrictive zoning constraints or difficult geography) that are encountered during the plan check review itself.

[9]Plan check review is neither a team nor a collaborative effort, so it is not surprising that the plan checker effects are not sensitive to controlling for the set of peers who were also receiving assignments in the same week as the focal plan checker.

5.1.2 Variance decomposition

To quantify what share of permit timeline variation is attributable to the plan checker, I perform a variance decomposition to assess the relative importance of plan checkers and project applicants. Using Equation (1) and suppressing subscripts, let $\alpha \equiv \alpha_{j(i,t)}$, $\phi \equiv \phi_{k(i)}$, and $x \equiv \mathbf{X}_i\boldsymbol{\beta} + \tau_t + \nu$ collect the observable controls and year effects. Then, the variance of log issuance time decomposes into:

$$\begin{aligned}\mathrm{Var}(y_{it}) = \; & \mathrm{Var}(\alpha) + \mathrm{Var}(\phi) + \mathrm{Var}(x) \\ & + 2\,\mathrm{Cov}(\alpha, \phi) + 2\,\mathrm{Cov}(\alpha, x) + 2\,\mathrm{Cov}(\phi, x) + \mathrm{Var}(\varepsilon)\end{aligned} \tag{2}$$

In this setting, there may be few "switchers" between pairs of plan checkers and applicants, which introduces sampling error in the estimates of $\mathrm{Var}(\alpha_j)$ and $\mathrm{Var}(\phi_k)$, known as "limited mobility bias." To mitigate concerns about such "limited mobility bias," I exclude single plan checker-applicant pairs, and I restrict my sample to the largest set of permit applications that are connected by applicants assigned to different plan checkers, known as the largest connected set. I also follow the literature by applying empirical Bayes shrinkage methods and the leave-out estimator proposed by Kline et al. (2020) to correct for "limited mobility bias" and sampling variability.[10]

Results from the variance decomposition on the largest connected set are shown in Table 4. The first group of columns reports the naïve plug-in estimates, the second group reports the "shrinkage" estimates, and the last group reports estimates using the leave-out estimator from Kline et al. (2020) (KSS). Given that limited mobility bias is a concern in this setting, I prefer the variance estimates corrected via the leave-out estimator from KSS. The results show that plan checkers explain 4% of the residual variation in permit timelines for new residential projects. Using the estimates in the last column, a one standard deviation faster plan checker issues permits in 15.4% less time on average. The average permit issuance time is 296 days in my sample, so a one standard deviation faster plan checker issued permits in 46 fewer days on average. Adjusting for sampling variability is important, as the plug-in estimates in the first column overstate the effects of plan checkers and applicants. The distribution of plan checker fixed effects is plotted in Figure 6. At any significance level, I can reject the null hypothesis that the plan checker fixed effects are jointly zero. Given that this is a regulatory setting where one might expect limited discretion from plan checkers who are tasked with consistent enforcement of existing codes, these results imply that differences

[10] I observe plan checkers working with an average of 52 unique applicants, and applicants with an average of 10 unique plan checkers. There are 23,114 pairs in my analysis sample, and there are an average of 2.3 permits per plan checker-applicant pair, so there are concerns about consistency of estimates.

in plan checker performance do explain a non-trivial share of the variation in permitting timelines.[11]

6 Understanding Determinants of Plan Checker Performance

This section examines what might predict being a faster plan checker and what actions differentiate a faster plan checker. First, I provide evidence on two channels, selection and on-the-job performance, that might explain why some plan checkers are faster than others. Second, I use a workload shock to explore mechanisms behind the faster performance of certain plan checkers to illustrate what faster plan checkers do differently on the job.

6.1 What predicts plan checker speed?

A faster plan checker might be faster due to two broad reasons: selection into the job, or performance on the job. To provide evidence of either channel, I regress the estimates of plan checker fixed effects on observable characteristics that reflect either selection or performance. For the characteristics that reflect selection, I include plan checkers' years of prior work experience, their education, and public sector pay gap in their hire year. For the characteristics that reflect performance on the job, I include variables related to the plan check process and variables related to promotion, overtime work, and pay. The results are shown in Figure 7. The coefficients represent the correlation between a 1 standard deviation increase in the project attribute and the *slowness* (as measured by time taken) of the plan checker on permit timelines, so the coefficients to the right of zero indicate these attributes are correlated with a plan checker being slower during the plan check phase on average.

The results suggest that faster plan checkers tend to take less time during their initial review and in between subsequent corrections rounds, on average (Figure 7). In the permit status updates, one can also see how many times an applicant returned to address corrections and the number of rounds of corrections.[12] Slower plan checkers tend to have more back-and-forth in the plan check process, as indicated by the positive coefficient for this variable. This could be due to plan checkers' actions if slower plan checkers are less thorough and miss

[11] There is a growing literature documenting how dispersion in workers' performance affects output in other settings, even if the skills and domain expertise are quite different from what I study. In public sector settings, Best et al. (2023) find that Russian procurement officers explain 11% of the variation in quality-adjusted prices, Fenizia (2022) finds that Italian social security managers explain 9% of the variation in worker productivity, and James (2024) estimates that managers of caseworkers for Texas public benefits programs explain 8% of the variation in the monthly caseworker productivity.

[12] The respective status updates are "Applicant returned to address corrections" and "Corrections issued." I combine these to construct the number of additional corrections beyond the first initial corrections that are issued.

errors initially, resulting in more corrections rounds subsequently.[13]

The bottom panel of Figure 7 explores two aspects of what could make for a faster plan checker: selection into the job and performance on the job. The first four coefficients look at variables related to selection: years of prior work experience, having a Master's degree, having a degree from a top engineering university (as ranked by US News & World Report), and the public sector wage gap in hire year. The negative coefficients on years of prior experience and having a Master's degree suggest that more experienced and more highly educated plan checkers are more likely to be faster on the job, though neither are statistically significant. The coefficient on public sector wage gap indicates that being a higher-performing plan checker is correlated with being hired in years where the wage gap between LADBS and the private sector is smaller.[14] This correlation is in line with previous work that has found better pay attracts higher quality public sector employees (Dal Bó et al., 2013; Krueger, 1988). Total annual pay and promotion to senior plan checker role does not strongly correlate with plan checker speed. Instead, faster plan checkers are ones who are more likely to work more overtime. Given the institutional practices around voluntary overtime being assigned proportionally in this setting, this correlation reflects faster plan checkers *choosing* to work more overtime. Overtime is desirable because it allows non-managers to earn pay that is comparable to what managers earn.[15] Taken together, these results are consistent with anecdotal evidence that there are two tracks for plan checkers to earn higher pay as they progress: either work more overtime or be promoted to a manager role (where they are not allowed to work overtime). While I am not able to definitely conclude *why* some plan checkers work more overtime than others, previous work finds that overtime work reflects bureaucrats' preferences and affects performance in other public sector settings (Chalfin and Gonçalves, 2026; Kohlhepp and McDonough, 2026). Thus, these results are consistent with both positive selection of higher-performing plan checkers and/or overtime as a motivating incentive to be more effective, especially in the public sector where the main way to increase one's salary is typically through promotion.

One might also be concerned that the relationship between faster plan checkers and

[13] Alternatively, this could also be explained by certain types of applicants who tend to ignore corrections, as LADBS staff have suggested anecdotally there may be some applicants who try to avoid or contest certain regulatory requirements. For a longer discussion of what project applicants contribute to the plan check process, see Appendix F.

[14] To compute the wage gap between LADBS plan checkers and their outside option (which I assume to be the private sector), I use the wage and income estimates in the 2005-2024 American Community Survey recorded for private-sector civil engineers working full-time in the Los Angeles metropolitan area. I match each plan checker to their outside option based on the year they started working at LADBS.

[15] Appendix D discusses the role of overtime in more detail and Table A8 examines what correlates with overtime hours, suggesting that higher base pay positively predicts overtime hours and promotion incentives may also factor into overtime hours worked.

more overtime hours is due to plan checkers "shirking" during regular hours, resulting in higher overtime hours. To explore the extent to which the performance measure α_j captures the true productivity of plan checkers, I test whether plan checkers who are higher-performing are working a different volume of residential permits. If plan checkers are higher-performing *because* they are able to work through more permits due to more overtime hours, then their annual volume of permits relative to total hours worked should be higher than that of lower-performing plan checkers. Using a project complexity-adjusted weekly productivity measure, I find that slower plan checkers tend to have a lower productivity measure as well (Figure A7). This correlation suggests that the main channel through which overtime increases performance is allowing plan checkers to process more permits.

6.2 What do faster plan checkers do differently?

To further explore what behaviors are undertaken by faster plan checkers, I leverage an exogenous shock to workload to estimate how plan checkers respond. The shock I use is the triennial building code update process in California. Every three years, there is a spike in permit filings right before the effective date of the new building code, because the submittal date of the permit determines what code vintage a building must comply with. This rush is shown in Figure 8, which plots a histogram of total permits filed in Los Angeles in the weeks around a code change, separately for years with and without a code change.

I use this external demand shock to estimate how an increase in workload affects plan checker actions and plan check outcomes using a differences-in-differences design. A naïve comparison of permits submitted immediately prior to a code change against those submitted slightly earlier could be biased by unobserved, intra-year seasonal cyclicality in permit submissions. To account for this seasonality, I use the identical calendar weeks from the previous year in which there was no code change as the control group. I restrict the sample to an eight-week window prior to a code change date (and the corresponding eight-week placebo window in the previous year), because the permits submitted after a code change date would be subject to different building codes. I estimate the following model:

$$Y_{it} = \beta_1(\text{Rush}_{it} \times \text{CodeYear}_{it}) + \beta_2\text{Rush}_{it} + \beta_3\text{CodeYear}_{it} + \mathbf{X}_i'\gamma + \boldsymbol{\delta} + \varepsilon_{it} \tag{3}$$

in which Y_i is the outcome of interest (i.e. log issuance time, log initial review time, log clearances/corrections time, additional correction rounds) for permit application i. The variable Rush_{it} is a dummy variable equal to 1 if the permit is submitted in the "rush period" spanning weeks -4 to -1 prior to the code change date, and 0 if it is submitted in the "non-rush" period spanning relative weeks -8 to -5. The variable CodeYear_{it} is an

indicator equal to 1 if the permit is submitted in the year of an actual code change, and 0 if it is submitted in the placebo year (one year prior). The coefficient of interest is β_1, which captures the true anticipatory "rush" effect. It measures the differential change in outcomes during the four weeks immediately preceding a building code change, relative to the baseline period, net of standard seasonal variation observed during the same calendar weeks in non-code-change years. The vector $\mathbf{X}_i$ represents project-level covariates. To absorb time-invariant unobservables, I include a vector of covariates and high-dimensional fixed effects, denoted by $\boldsymbol{\delta}$. These include code change cohort-office fixed effects to control for office-specific macroeconomic trends, as well as the same project covariates, administrative fixed effects, and plan checker and applicant fixed effects from Equation (1). Standard errors are clustered at the plan checker and week level, allowing for arbitrary correlation in the error term for permits assigned weekly to a given plan checker. The identifying assumption of this model is that, in the absence of a building code change, the difference in outcomes for permits submitted in the four weeks just prior to the code change date and the preceding four weeks would have been similar to the difference observed across the exact same calendar weeks in the prior year. A balance table comparing the permits filed in the four time groups is shown in Table 5, which confirms that permits are comparable across the pre-period.

Table 6 reports the three β coefficients from Equation (3). Permits filed in the month before a code change take 14% longer on average to be issued compared to non-rushed permits in that same year, relative to the permits filed in the non-code change years. This is partly due to a 78% (or 57.6 log points) increase in pure wait time as the permit sits in the queue longer, waiting to be assigned to a plan checker. Moreover, there is a 15.4% (or 14.4 log points) increase in the time spent in the clearances and corrections phase and an additional 0.3 corrections rounds for permits submitted in the rush period in a code change year, compared to the control group on average.

While these results highlight the average effects of a workload spike on plan check outcomes, the subsequent question is to ask if certain plan checkers respond to the shock differently. To answer this, I estimate a modified version of Equation (3) separately for plan checkers above and below the median performance level (using estimates of plan checker speed $\hat{\alpha}_{j(i,t)}$ from re-estimating Equation (1) on the subsample of permits that exclude the "rush period").[16] Table 7 shows the results, with the above-median plan checkers on top panel and

[16]I exclude the "rush period" permits to avoid a mechanical bias from including the outcome variable in constructing estimates of $\hat{\alpha}_{j(i,t)}$. Note that the estimates of $\hat{\alpha}_{j(i,t)}$ with and without the "rush period" permits are highly correlated (correlation coefficient of 0.97), as shown in the bottom panel of Table A2. I also exclude the plan checker fixed effect in this specification. Another version of this exercise that yields similar estimates is to exclude the placebo year control group and simply compare above- and below-median plan checkers in the code change year, under the assumption that permits assigned to both types of plan checkers are comparable leading up to the "rush period." Table A3 and Table A4 show a balance test and

the below-median plan checkers on the bottom panel. The last row of the table indicates that the below-median plan checkers are slowest at issuing permits, since "rush" permits assigned to them in code change years are issued 20% slower relative to such permits in the placebo years. "Rush" permits assigned to the below-median plan checkers also face a 16.5% increase in the amount of time spent on the clearances and corrections phase and a higher number of corrections rounds (0.16) relative to the "rush" permits in pseudo-code change years, though the former effect is only statistically significant at the 10% level and the latter effect is not significant at conventional levels. In contrast, there are no similar treatment effects on plan check outcomes that involve plan checker actions for "rush" permits assigned to above-median plan checkers. These results suggest that faster plan checkers are better able to manage their higher workload without greater delays beyond the initial longer queue time, particularly during the phase after initial corrections are issued when plan checkers are juggling many projects returning for subsequent corrections.

Why are faster plan checkers better able to manage a higher workload? To explore what other aspects of job performance might characterize the faster plan checkers, I turn to the payroll data to examine how hours worked respond to this workload shock. I use the payroll data to estimate an event study specification of the previous model that compares above-median plan checkers to below-median checkers in the 5 pay periods around the code change date. I estimate the following equation:

$$Y_{jw} = \sum_{w=-4}^{5} \beta_w \left(\mathbf{1}[\tau_w = w] \times G_j \right) + \tau_w + \mu_{c(w)} \times G_j + \varepsilon_{jw} \tag{4}$$

in which Y_{jw} is overtime hours worked by plan checker j in pay period w, τ_w indicates event-time dummy variables for relative pay period w, G_j is an indicator variable equal to 1 for a plan checker j whose $\hat{\alpha}_j$ is above the median $\hat{\alpha}_j$, and $\mu_{c(w)}$ are code year cycle indicator variables. The event study coefficients are plotted in Figure A8, which suggests that overtime is the main channel through which plan checkers are able to adapt to fluctuating workloads. During the pay periods immediately *after* a new code change, the number of overtime hours worked by above-median plan checkers increases by 2 hours on average compared to below-median plan checkers. These results imply that faster plan checkers are able to work through a backlog during busy periods by voluntarily increasing their overtime hours.[17] The evidence so far suggests that overtime is one way to flexibly add capacity without having to hire new plan checkers.

results from this alternative specification.

[17]While it is beyond the scope of this paper to understand *why* certain plan checkers choose to work more overtime, the effect on *outcomes* is arguably the first-order concern and the focus of this paper.

7 The Quality Impacts of Plan Checker Performance

In this section, I evaluate whether there are any quality impacts of being assigned a faster plan checker on project outcomes. Being assigned a faster plan checker might lead to quality tradeoffs if, for example, faster plan checkers were worse at catching missed corrections during the plan check phase, contributing to lower quality construction or a longer construction phase. I proxy for project quality using future code violations and enforcement cases issued by a separate division within LADBS. To address concerns of attenuation bias due to noisy estimates of plan checker speed, I use an instrumental variables design that leverages variation in the availability of high-performing plan checkers. I provide validation for this design, and finally I present the results and robustness checks.

7.1 Instrumental variables design

To estimate the causal effect of plan checker performance on future project outcomes, I build on the previous framework that constructed a measure of plan checker performance, as described in Section 5.1.

Consider now the following model that relates how $\alpha_{j(i,t)}$ affects a given project's quality outcome V_{it}, conditional on $\mathbf{X}_i$, τ_t, and $\boldsymbol{\nu}$:

$$V_{it} = \rho\alpha_{j(i,t)} + \phi_k + \mathbf{X}_i\beta + \tau_t + \boldsymbol{\nu} + \eta_{it}$$

Simply using the estimates of $\alpha_{j(i,t)}$ from Section 5.1 as a measure of plan checker performance in a least squares regression warrants caution for three reasons. First, the estimates $\hat{\alpha}_{j(i,t)}$ are noisy, which could result in attenuation bias. Second, mechanical correlation arises from the estimates of $\alpha_{j(i,t)}$ from Section 5.1 which include own permit i when estimating $\hat{\alpha}_{j(i,t)}$. Third, the plan checker effects are estimated from a potentially small number of permits, so any permit-level shock could simultaneously affect the value of $\hat{\alpha}_{j(i,t)}$ for a particular plan checker and affect the future quality of the project.

To circumvent these concerns, I adopt a leave-out instrumental variables design. To avoid the mechanical correlation that arises when permit i is used to estimate the performance of its own assigned plan checker, I construct a leave-one-out (jackknife) measure of plan checker performance. Specifically, let $w(i)$ denote the calendar week in which permit i was assigned. I re-estimate Equation (1) excluding all permits assigned during week $w(i)$. The resulting fixed effect, $\tilde{\alpha}_{j,-w(i)}$, serves as an uncontaminated measure of plan checker j's performance for any permit assigned in that week.

While $\tilde{\alpha}_{j,-w(i)}$ removes mechanical bias, OLS estimation may still be biased due to

permit-level shocks and measurement error in $\tilde{\alpha}_{j,-w(i)}$. To overcome this, I employ an instrumental variables design that leverages exogenous variation in the timing of available plan checkers relative to permit submissions and show evidence supporting the validity of this approach.

The two-stage least squares (2SLS) model is estimated in the following form:

$$\text{First Stage:} \quad \tilde{\alpha}_{j(i,t),-w(i)} = \gamma Z_{jw} + \mathbf{X}_i\pi + \phi_k + \tau_t + \boldsymbol{\nu} + u_{it} \tag{5}$$

$$\text{Second Stage:} \quad V_{it} = \rho\tilde{\alpha}_{j(i,t),-w(i)} + \mathbf{X}_i\beta + \phi_k + \tau_t + \boldsymbol{\nu} + \eta_{it} \tag{6}$$

in which V_{it} is the outcome of interest, $\tilde{\alpha}_{j(i,t),-w(i)}$ is the jackknife plan checker performance measure, and Z_{jw} is the instrument, which is the average performance of the available pool of plan checkers that permit i could have been assigned to. Specifically, Z_{jw} is the average leave-one-week-out performance of the set of *other* plan checkers that were working in the same office and week as when permit i was assigned to plan checker j:

$$Z_{jw} = \frac{1}{|\mathcal{P}_{ow(i)} - 1|} \sum_{j' \in \mathcal{P}_{ow(i)}, j' \neq j} \tilde{\alpha}_{j',-w(i,t)} \tag{7}$$

in which $\mathcal{P}_{ow}$ denotes the set of plan checkers actively receiving assignments in permit office o and during week w. I exclude office-week cells with fewer than 3 plan checkers.

For my baseline controls, I include administrative fixed effects $\boldsymbol{\nu}$ (which include permit office, plan check division, and project use type) to allow for sorting of certain plan checkers across these dimensions and time effects τ_t to absorb seasonality and macroeconomic shocks. Given that there may be periods and offices that have longer backlogs, my baseline time effects τ_t include quarter-year×office fixed effects to control for office-specific seasonality effects. I also include applicant fixed effects ϕ_k because I do not want applicant effects to be attributed to plan checker performance.

The causal parameter of interest, ρ, is identified from week-to-week idiosyncratic variation in the composition of the available plan checker pool within a given office and quarter. This instrument isolates plausibly exogenous variation in the availability of high-performing plan checkers driven by vacation/leaves of absences, rotations, and workforce expansions and reductions which shift the expected performance of the assigned plan checker independent of a specific permit's characteristics. Said differently, this instrument exploits variation orthogonal to the permit assignment decision in order to identify ρ, so Z_{jw} is unaffected by sorting of certain permits within the pool of available plan checkers. To illustrate the variation in plan checker availability, Figure A9 plots the number of all plan checkers observed being assigned to a permit in a given week, separately for the three largest offices. The figure

demonstrates that there is quite a bit of weekly variation in the number of plan checkers who are receiving assignments. The main threat to this identification strategy is if plan checkers timed their availability in anticipation of permit-level shocks (unobserved by econometrician). In the prior example of sorting on unobserved permit difficulty, this would require certain types of plan checkers to respond to the anticipated difficulty of permits that have not been assigned yet, which is a much less plausible scenario. Table A5 provides evidence that there is no anticipatory selection in the availability of plan checkers by regressing the average performance of available plan checkers in the week after submission on various ex-ante permit characteristics. A joint test of significance indicates that there is little evidence of permit characteristics predicting future plan checker availability.

7.1.1 Identification

To interpret ρ as the LATE of being assigned to a higher-performing plan checker, I provide evidence to support the validity of the four identifying assumptions.

Relevance: For the instrument to be valid, a permit submitted in an office-week with a pool of plan checkers who are slower on average must be more likely to be assigned to a slower plan checker. The first stage results are shown in Figure 9, which plots a binscatter of the assigned plan checker's performance against the average performance of plan checkers in the same office. A one standard deviation increase in the average performance of the available plan checkers in the same office and week increases the assigned plan checker's performance by 0.18 standard deviations. The first stage F-statistic is 102.37, indicating a strong first stage.

Conditional Independence: To provide evidence that the availability of high-performing plan checkers is unrelated to a permit's potential outcomes, I correlate the instrument with various ex-ante permit characteristics, conditional on the baseline controls. Figure 10 plots bivariate regressions between the instrument and permit characteristics, finding that all coefficients are not statistically significant at the 5% level. A test of joint significance fails to reject at the 10% level, which suggests that the instrument is not systematically correlated with certain types of permits.

Exclusion: To provide evidence that the average availability of high-performing plan checkers only affects outcomes through the channel of being assigned a high-performing plan checker, I test for whether other plan checker actions and office-week conditions might be correlated with the instrument. Specifically, I show that plan checkers' average workload, permit queue depth, years of experience, and office-week permit filing volume (total and new residential), are not correlated with the average availability of high-performing plan checkers, Figure 11 plots these five variables against the instrument, which finds no statistically sig-

nificant correlation. This supports the argument that there are not other channels through which plan checker availability could affect outcomes (such as correlated office shocks across multiple weeks) besides increasing the likelihood of being assigned to a project.

Monotonicity: Finally, monotonicity requires that a permit assigned to a high-performing plan checker that was submitted in a week-office cell with fewer high-performing plan checkers would be assigned a high-performing plan checker if submitted in a cell with more high-performing plan checkers. To document support for this assumption, I show that the instrument and the jackknife plan checker performance measure are positively correlated for various subsamples defined by the median of permit characteristics (Figure A10).

7.2 Results

In this section, I discuss my main results from the 2SLS model. Table 8 and Table 9 report the OLS, reduced form, and IV estimates of the effect of plan checker performance on permit outcomes. Figure 12 plots the coefficients from Equation (6) for outcomes observed after the plan check phase (after permit issuance). The x-axis plots the assigned plan checker's leave-one-out performance measure (which is a measure of their speed).

I find that being assigned a 1 standard deviation faster plan checker decreases the amount of time a permit spends in the clearances and corrections phase by 17.3%, though this is not statistically significant. Why might this be? Part of the reason is that projects assigned to a 1 standard deviation slower plan checker have an additional round of corrections. These results support anecdotal evidence that higher-performing plan checkers are ones who are responsive and thorough, minimizing the number of back-and-forth interactions between applicants and plan checkers. Additionally, I do not find evidence that assignment to a faster or slower plan checker affects the likelihood of a permit being issued. This is perhaps unsurprising given the job function of plan checkers who are responsible for enforcing rules and pointing out what needs to be corrected in order for a project to be in compliance, so they are generally not rejecting projects outright.

Next, I estimate how being assigned a faster plan checker affects project outcomes after permit issuance. I find that being assigned a 1 standard deviation slower plan checker increases construction and inspection time by 2.4% and 13%, respectively, though these estimates are noisy and not statistically significant at conventional levels. I do not find evidence that there are extensive margin effects, as proxied by whether a project is completed within three years.[18] In the longer run, I do not find being assigned a faster plan checker affects a project's future likelihood of a code enforcement case, code violation, or safety-

[18]Since projects filed most recently are less likely to be completed, I restrict my analysis sample to permits whose issuance date was at least 3 years prior to the permit data collection date (October 23, 2025).

related code violation at standard significance levels. Altogether, these results indicate there are no measurable impacts on downstream quality and the extensive margin from being assigned a faster plan checker. These null effects on quality (as proxied by future code activity) can be interpreted as confirming that this is a regulatory enforcement context where plan checkers and inspectors are tasked with ensuring each project meets minimum safety requirements.

I also explore whether there are heterogeneous effects based on the seniority of the plan checker. I estimate a version of Equation (5) and Equation (6) where I include an interaction term for seniority:

$$y_{it} = \rho_1 \tilde{\alpha}_{j(i,t),-w(i)} + \rho_2 \tilde{\alpha}_{j(i,t),-w(i)} \times \text{Senior}_j + \rho_3 \text{Senior}_j + \mathbf{X}_i\beta + \phi_k + \tau_t + \boldsymbol{\nu} + \eta_{it} \quad (8)$$

I instrument for $\tilde{\alpha}_{j(i,t),-w(i)}$ and $\tilde{\alpha}_{j(i,t),-w(i)} \times \text{Senior}_j$ using Z_{jw} and $Z_{jw} \times \text{Senior}_j$. The results are noisy overall, but I do not find evidence that being assigned a senior plan checker who is high-performing differentially affects project outcomes (Figure A11).

I next describe a set of robustness checks that rule out concerns surrounding this instrumental variables design. One might be concerned that certain types of project applicants can better time their permit applications in smaller permit offices where they are more likely to get assigned to their desired plan checker, which would be a threat to identification since that would violate the assumption that the availability of plan checkers is unrelated to unobserved project attributes. As a robustness check, I re-estimate the 2SLS model (Equations (5) and (6)) but restrict my sample to the 3 largest permit offices (Metro, Van Nuys, and West Los Angeles), which handle 93% of all new residential permits. The estimates are similar to the main estimates shown previously, with smaller coefficients on the plan check outcomes (Figure A12). I also re-estimate the 2SLS model while excluding permits filed during the "rush" periods before a new building code takes effect (as discussed in Section 6.2), since the effect of the workload shock might be conflated with the impact of the plan checker. I also find similar estimates as those in the main results (Figure A13).

As another robustness check, I also implement a version of a split-sample instrumental variables (SSIV) design (Angrist and Krueger, 1995) that is described in more detail in Appendix G. I form two partitions of the data (denoted A and B) by sampling permits submitted in weeks associated with each plan checker-applicant combination. I estimate Equation (1) on both partitions, which yields two estimates of the plan checker effect (denoted $\hat{\alpha}^A_{j(i,t)}$ and $\hat{\alpha}^B_{j(i,t)}$) with independent errors. I use the out-of-sample estimate of a plan checker's effect as an instrument for own plan checker speed (i.e. I instrument for $\hat{\alpha}^A_{j(i,t)}$ using $\hat{\alpha}^B_{j(i,t)}$ for a permit in partition A and vice versa for a permit in partition B). While one

could also use a jackknife procedure in a standard judge/examiner design, the split sample design is more robust to correlated shocks within weeks in a given office.[19] I find similar null effects on the likelihood of a code enforcement case and code violation compared to those in my main estimates, while I find that a 1 standard deviation slower plan checker increases construction and inspection time by 3.8% and 7.2%, respectively. These estimates are qualitatively similar to the main estimates, which reinforces the result that being assigned a faster plan checker does not measurably affect future project quality and highlights how plan checker performance matters most for the time saved during a project.

8 The Value of Bureaucratic Capacity in the Permitting Process

The evidence presented so far establishes several results about plan checker performance: faster plan checkers save time for projects assigned to them, without substantively worsening the project quality, and that these higher-performing plan checkers choose to work more voluntary overtime. This section discusses how I quantify the value of shorter permitting timelines from alternative personnel policies that increase plan checker capacity. However, the cost implications of these policies depend also on the relationship between performance and overtime, since hiring a high-productivity plan checker who works little overtime is lower cost than hiring a low-productivity plan checker who works more overtime. A simple framework for plan checker behavior is outlined in Appendix E that highlights under what conditions productivity and overtime are complementary, which is what the empirical evidence has shown support for. For the following counterfactual scenarios, I highlight how this empirical relationship affects cost-effectiveness of each alternative personnel policy.

The estimates presented in this section are suggestive, as the counterfactual scenarios abstract away from other extensive margin changes (from increased demand, for example) and other costs (such as hassle costs) by only considering the time saved for new residential projects. In an ideal situation, I would like to estimate the costs of the avoided delay time for each scenario. However, absent data on builders' pre-construction loan values, land purchase prices, and the prior use of parcels for every project in my sample, I present estimates that are best-guess approximations meant to illustrate the potential costs of permitting delay. For these approximations, I convert the avoided delay time into a monetary value using the reported project value and an accounting exercise that sums across financing costs (the cost of capital from a pre-development loan), carrying costs (the cost of sitting on an empty lot, such as property taxes, insurance, etc), and foregone rents (delayed rental income). I assume

[19] In the standard judge/examiner design, leaving out a particular permit when calculating the plan checker performance measure for that permit may be invalid if there were other week-level shocks that affect both that week's plan checker assignment and outcomes.

a loan-to-value ratio of 0.7, interest rate of 6%, land costs are 25% of total project value, a 1% property tax rate, and a capitalization rate of 4%.[20] To calculate the net benefit-to-cost ratio, I compute the ratio of private builders' cost savings to the additional personnel costs using plan checkers' total annual salaries, which restricts these estimates to the years 2005-2024.[21] LADBS operations (including salaries) are funded through permit fees and other service charges collected in its own Enterprise Fund, not from the city's General Fund (City of Los Angeles Controller, 2019). Thus, this ratio of private savings to personnel costs can be thought of as the return of an additional dollar in permit fees paid by builders that is spent on increasing plan checker capacity.

First, I consider the counterfactual policy of *improving the lowest-performing 25% of plan checkers* to be like the median plan checker, using estimates from the previous two-way fixed effects model. I find that if the lowest-performing 25% of plan checkers were to perform like the median plan checker, this would save approximately 598,380 days and $263 million for 8,689 affected projects, which is approximately $30,264 per project. Using the actual annual salaries of plan checkers, I find that the return from each dollar spent in this scenario is $9.02. These estimates assume that the replacement plan checkers work overtime according to their observed overtime in the data, so I also estimate the savings and costs if one assumes that the replacement plan checkers' overtime hours were fixed to the level of the replaced group's annual overtime hours. I find that the return from each dollar spent in this scenario is $8.11, partly because there are now fewer days saved and lower private savings. I also consider alternative personnel policies in Table 10, which reports the private savings, city payroll costs, net benefits, and return from each dollar spent under various scenarios for substituting the low-performing plan checkers (bottom 25% or 10%) with either the top 25th percentile or the median plan checker. Across every scenario, the return from each dollar spent is greater than 1, highlighting the cost-effectiveness of investing in bureaucratic capacity in this setting.

Another relevant counterfactual scenario would be to ask: *what is the value from increasing plan checker pay by 10% through the selection channel?* While I cannot directly estimate how changes to financial incentives affect the quality of plan checkers, I use the wage gap elasticity estimates from Section 6.1 to impute how higher pay might shift the distribution of plan checker performance ($\hat{\alpha}_{j(i,t)}$).[22] I find that a 10% increase in plan checker

[20]This approximation also assumes homogeneous builders who take out their pre-development loan when they submit their permit applications. More details about this calculation are in Appendix H, and sensitivity to these assumptions is shown in Figure A22.

[21]The annual salary includes base pay, overtime pay, and benefits pay. I only have benefit pay for 2017-2024, so I impute benefit pay in other years to have the same average share of base pay as in 2017-2024.

[22]The elasticity between the log wage gap and standardized plan checker fixed effect that I estimate is -1.16, so a 10% increase in plan checker wages would reduce the log wage gap by 0.0953 log points, which

pay would save approximately 239,160 days of waiting time and save builders over $125.8 million or $3,093 per permit, with an ROI of $3.92. I also estimate this scenario while holding overtime fixed to plan checkers' existing levels and calculate an ROI of $3.77, highlighting that most of the gains from higher-performing plan checkers are not due to faster plan checkers being willing to work more overtime. These back-of-the-envelope exercises are illustrative and demonstrate that the net benefits of additional bureaucratic capacity are potentially large and economically meaningful across a variety of counterfactual personnel policy changes.

In the last counterfactual exercise, I provide estimates of the value from *increasing staffing by 5%*. To estimate how adding plan checker capacity would affect permit issuance times, I adapt a simple queueing framework from the operations literature (Kingman, 1961; Little, 1961) to simulate how increased staffing would affect permitting delay. I simulate the plan check process as a multi-server queue in which offices are staffed by their observed annual plan checker headcount, applications arrive on their submittal dates and are assigned weekly to the next available plan checker, and plan check initial review times are a function of the speed of the assigned plan checker ($\hat{\alpha}_{j(i,t)}$). Two parameters are not observed directly: the maximum number of concurrent applications a plan checker can be assigned, and what amount of influence plan checkers have during the clearances/corrections phase (or what share of the clearances/corrections phase still occupies a plan checkers' capacity). I calibrate them via simulated method of moments so that the simulated queue reproduces the mean, dispersion, and upper tail of observed queue times and the increase in queue time during the filing rush before a building code change. Appendix I describes the model, estimation, and the model's fit. The fit is imperfect, as the model comes close to matching the 90th percentile of queue times but slightly understates their mean and overstates their dispersion and the size of the rush, and it is rejected by an overidentification test. In the model, an office that receives more applications than it can review falls further behind, whereas actual offices likely absorb part of such surges through overtime, for instance (Section 6.2). I therefore treat the results of this exercise as illustrative. At the calibrated parameters, increasing each office's capacity by 5%, with new plan checkers of median performance, saves 80,956 days of queue time in total for permits submitted in 2005-2024. Using the accounting framework above, builders would save $19.7 million, and the return on each dollar spent on new plan checkers is $2.16. Note that this estimate is lower than the previous counterfactual scenarios because this counterfactual only computes days saved during the queueing phase of the plan

would shift the distribution of $\hat{\alpha}_{j(i,t)}$ by -0.11 standard deviations towards faster plan checkers. The wage gap elasticity I estimate can be confounded with other time-varying labor market shocks, so these estimates should be interpreted with caution.

check process, which likely understates time savings.

What do these estimates imply for housing supply? Across the three scenarios that change plan checker capacity, the average percentage reduction in permit issuance time ranges from 0.6% for the 5% staffing change, 1.8% for the 10% pay increase, to 4.4% for the replacement of the 25% lowest-performing plan checkers with the median plan checker. To translate these time savings into housing units, I follow Gabriel and Kung (2026) and simulate how many projects that were under construction as of October 23, 2025 would have been completed by that date under each personnel policy (Appendix J).[23] Replacing the lowest-performing 25% of plan checkers with the median plan checker would have completed an additional 899 units (235 projects) by the end of my sample, or 2.7% of the 32,975 units in outstanding projects. Increasing plan checker pay by 10% would have completed an additional 790 units (109 projects), and increasing staffing by 5% an additional 26 units (10 projects). Relative to all units completed during my sample, these percentage gains are below 1%. While these estimates may appear small, these magnitudes are in line with the size of the time savings from each of the three changes to personnel capacity. The implied elasticity of completed units with respect to issuance time is 0.14 for replacing the lowest-performing plan checkers and 0.31 for the 10% pay increase, compared to an elasticity of 0.5 in Gabriel and Kung (2026). Their average approval time is 652 days (because it includes an entitlement review stage prior to permit submittal and is limited to multi-family projects), compared to 296 days of issuance time in my sample, so a 1% reduction in issuance time corresponds to fewer days saved in my setting. Scaling their elasticity by this ratio of the two averages yields an elasticity of 0.23, which is comparable to the elasticities implied by the two plan checker performance scenarios.

9 Conclusion

In the US, bureaucrats are typically responsible for regulatory decisions of local government agencies that affect the supply of housing, especially those related to the permitting of new housing units in any jurisdiction. This paper presents the first evidence of how the performance of one key bureaucrat in the permitting process can impact the timely provision of new housing.

I use newly assembled administrative data to shed light on how plan checkers affect building permit issuance, a required step for every development project but prone to delays and bottlenecks, especially in high-demand urban areas like Los Angeles, which is the main

[23] These estimates are arguably lower bounds, since they only capture the "pull-forward" effect (i.e. projects already under construction that are completed sooner) and exclude the "incentive" effect (i.e. reductions in issuance time incentivizing new projects), both as discussed in Gabriel and Kung (2026).

setting of this paper. I first show how permit issuance times have increased steadily over the last twenty years across the US and in Los Angeles. Then, using a two-way fixed effects model, I document that a one standard deviation faster plan checker takes 15.4% less time (46 fewer days) to issue permits on average. I do not find strong evidence of endogenous sorting between plan checkers and applicants based on past experience or past relationships. I find that faster plan checkers are ones who are more likely to work overtime, since overtime allows non-manager plan checkers to earn comparable pay to managers. Plan checkers hired in years with more competitive public sector pay are more likely to be faster plan checkers. These results are consistent with both positive selection and a greater incentive to perform due to overtime pay. Using a differences-in-differences design that exploits a shock to workload, I show that faster plan checkers are better at managing their workload of concurrent projects, likely because they voluntarily work more overtime after the workload shock.

To estimate the causal effect of a faster plan checker on project outcomes, I adopt an instrumental variables approach that leverages variation in the availability of faster plan checkers in a given week and office. I find that there are no detectable effects of plan checker performance on the likelihood of a future code violation or code enforcement, suggesting that concerns of a quality-quantity tradeoff are minimal in this setting. In the last section of the paper, I compare the cost-effectiveness of various counterfactual personnel policies and find that the net value of time saved from improving plan checker performance or increasing staffing is substantial and economically meaningful.

Overall, these results speak to the larger question of what makes for an effective bureaucrat in a regulatory agency. This paper highlights that a combination of both effort and selection influences performance even in a public sector setting where bureaucrats are supposed to enforce regulations "by the books" with limited discretion. The availability of overtime may be a way to complement the promotion-based incentives typically used in bureaucratic agencies. The evidence shown here suggests that an incentive structure with both overtime pay and promotions may induce positive selection of high-performers *and* increase the motivation to work harder. Moreover, the descriptive evidence so far suggests that there are costly organizational bottlenecks that arise from requiring applicants to navigate across an average of four different city agencies, but absent policy changes to organizational structure, concrete evidence is scarce. Future research should examine how varying incentive structures, such as overtime and promotions, or organizational structures affect the selection and performance of bureaucrats with regulatory authority.

These results have implications for ongoing policy debates about housing affordability and regulatory burden, especially as policymakers clamor for permit streamlining reforms. Regarding external validity, the results should be interpreted as representing the impacts of

bureaucratic capacity given the existing regulatory environment. While policy changes to the regulatory environment are important for addressing housing supply, changing regulations themselves may potentially be politically less feasible compared to policies targeting bureaucratic capacity. My findings suggest that the latter type of policy change may be cost-effective and feasible for reducing housing timelines. Moreover, my results highlight *how* permit streamlining is implemented matters. For any reform that stipulates additional verifications or requirements before a project can be expedited, the key question for policymakers is: does streamlining the permitting process require investing in additional bureaucratic capacity?

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

Anderson, T. W. and Herman Rubin, "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 1949, *20* (1), 46–63.

Aneja, Abhay and Guo Xu, "Strengthening State Capacity: Civil Service Reform and Public Sector Performance during the Gilded Age," *American Economic Review*, August 2024, *114* (8), 2352–2387.

Angrist, Joshua D. and Alan B. Krueger, "Split-Sample Instrumental Variables Estimates of the Return to Schooling," *Journal of Business & Economic Statistics*, 1995, *13* (2), 225–235.

Bartik, Alexander, Arpit Gupta, and Daniel Milo, "The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement," August 2025.

Baum-Snow, Nathaniel and Gilles Duranton, "Chapter 6 - Housing supply and housing affordability," in Dave Donaldson and Stephen J. Redding, eds., *Handbook of Regional and Urban Economics*, Vol. 6 of *Handbook of Regional and Urban Economics*, Elsevier, 2025, pp. 353–461.

Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu, "The Glittering Prizes: Career Incentives and Bureaucrat Performance," *The Review of Economic Studies*, 2020, *87* (2 (313)), 626–655.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi, "Individuals and Organizations as Sources of State Effectiveness," *American Economic Review*, August 2023, *113* (8), 2121–2167.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, April 2023, *15* (2), 1–30. Working Paper.

Chalfin, Aaron and Felipe Gonçalves, "Professional Motivations in the Public Sector: Evidence from Police Officers," *American Economic Review*, August 2026, *116* (8), 2962–95.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff, "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates," *American Economic Review*, September 2014, *104* (9), 2593–2632.

City of Los Angeles Controller, "Strengthening Oversight of the City's IT Contracts," Technical Report 2019.

Congressional Research Service (CRS), "Housing Cost Burdens in 2024: In Brief," 2026.

Dahis, Ricardo, Laura Schiavon, and Thiago Scot, "Selecting Top Bureaucrats: Admission Exams and Performance in Brazil," *The Review of Economics and Statistics*, March 2025, *107* (2), 408–425.

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi, "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service*," *The Quarterly Journal of Economics*, August 2013, *128* (3), 1169–1218.

D'Amico, Leonardo, Edward L. Glaeser, Joseph Gyourko, William R. Kerr, and Giacomo A.M. Ponzetto, "Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation," Working Paper 33188, National Bureau of Economic Research November 2024.

D'Amico, Leonardo, Evan Soltas, and Audrey Wang, "Where Is Housing Slow to Build, And Is It Getting Slower?," 2026. Working Paper.

Fenizia, Alessandra, "Managers and Productivity in the Public Sector," *Econometrica*, 2022, *90* (3), 1063–1084.

Gabriel, Stuart and Edward Kung, "Development approval times and new housing supply: Evidence from Los Angeles," *Journal of Urban Economics*, August 2026, *154*, 103872.

_ , M. J. Histen, and Edward Kung, "How Bureaucratic Capacity Shapes Regulatory Outcomes: Evidence from Land Use in Los Angeles," 2026. Working Paper.

Garcia, Daniel and Raven Molloy, "Reexamining lackluster productivity growth in construction," *Regional Science and Urban Economics*, June 2025, *113*, 104107.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, February 2018, *32* (1), 3–30.

Glaeser, Edward L and Joseph Gyourko, "The Impact of Building Restrictions on Housing Affordability," *Economic Policy Review*, 2003, *9* (June), 21–39.

_ , _ , and Raven E Saks, "Why Have Housing Prices Gone Up?," *American Economic Review*, April 2005, *95* (2), 329–333.

Gold, Daniel E, "Regulatory Hurdles and Costly Delay in Housing Development," 2025. Working Paper.

Gyourko, Joe and Jacob Krimmel, "The impact of local residential land use restrictions on land values across and within single family housing markets," *Journal of Urban Economics*, November 2021, *126*, 103374.

Gyourko, Joseph and Raven Molloy, "Chapter 19 - Regulation and Housing Supply," in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, Vol. 5 of *Handbook of Regional and Urban Economics*, Elsevier, 2015, pp. 1289–1337.

_ **and Sean McCulloch**, "Minimum Lot Size Restrictions: Impacts on Urban Form and House Price at the Border," Working Paper 31710, National Bureau of Economic Research September 2023.

_ **, Jonathan S. Hartley, and Jacob Krimmel**, "The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index," *Journal of Urban Economics*, July 2021, *124*, 103337.

Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 1982, *50* (4), 1029–1054.

Iaria, Alessandro, Carlo Schwarz, and Fabian Waldinger, "Gender Gaps in Academia: Global Evidence Over the Twentieth Century," July 2024. Working Paper.

James, Erik, "The Effect of Managers on Public Service Provision: Evidence from Medicaid, SNAP, and TANF," 2024. Working Paper.

Kestelman, Stephanie, "Effects of Transit-Oriented Incentive Programs: Evidence from Los Angeles," 2025. Working Paper.

Kingman, J. F. C., "The single server queue in heavy traffic," *Mathematical Proceedings of the Cambridge Philosophical Society*, October 1961, *57* (4), 902–904.

_ , "On Queues in Heavy Traffic," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1962, *24* (2), 383–392.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten, "Leave-Out Estimation of Variance Components," *Econometrica*, 2020, *88* (5), 1859–1898. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA16410.

Kohlhepp, Jacob and Robert McDonough, "The Wheel of (Over)Time: Efficient Overtime Exchange via Coworker Networks," 2026. Working Paper.

Krueger, Alan B., "The Determinants of Queues for Federal Jobs," *ILR Review*, July 1988, *41* (4), 567–581.

Lee, David S., Justin McCrary, Marcelo J. Moreira, and Jack Porter, "Valid t-Ratio Inference for IV," *American Economic Review*, October 2022, *112* (10), 3260–90.

Liscow, Zachary, Cailin Slattery, and Will Nober, "State Capacity and Infrastructure Costs," 2024.

Little, John D. C., "A Proof for the Queuing Formula: L = λW," *Operations Research*, June 1961, *9* (3), 383–387.

Manville, Michael, Paavo Monkkonen, Nolan Gray, and Shane Phillips, "Does Discretion Delay Development?: The Impact of Approval Pathways on Multifamily Housing's Time to Permit," *Journal of the American Planning Association*, July 2023, *89* (3), 336–347.

Mast, Evan, "Warding off Development: Local Control, Housing Supply, and NIMBYs," *The Review of Economics and Statistics*, May 2024, *106* (3), 671–680.

Metcalfe, Robert D, Alexandre B Sollaci, and Chad Syverson, "Managers and Productivity in Retail," Working Paper 31192, National Bureau of Economic Research April 2023.

Muñoz, Pablo and Cristóbal Otero, "Managers and Public Hospital Performance," *American Economic Review*, November 2025, *115* (11), 4040–74.

_ **and Mounu Prem**, "Managers' Productivity and Recruitment in the Public Sector," *American Economic Journal: Economic Policy*, November 2024, *16* (4), 223–53.

Ortalo-Magné, François and Andrea Prat, "On the Political Economy of Urban Growth: Homeownership versus Affordability," *American Economic Journal: Microeconomics*, 2014, *6* (1), 154–181.

Soltas, Evan and Jonathan Gruber, "How Costly Is Permitting in Housing Development?," 2026. Working Paper.

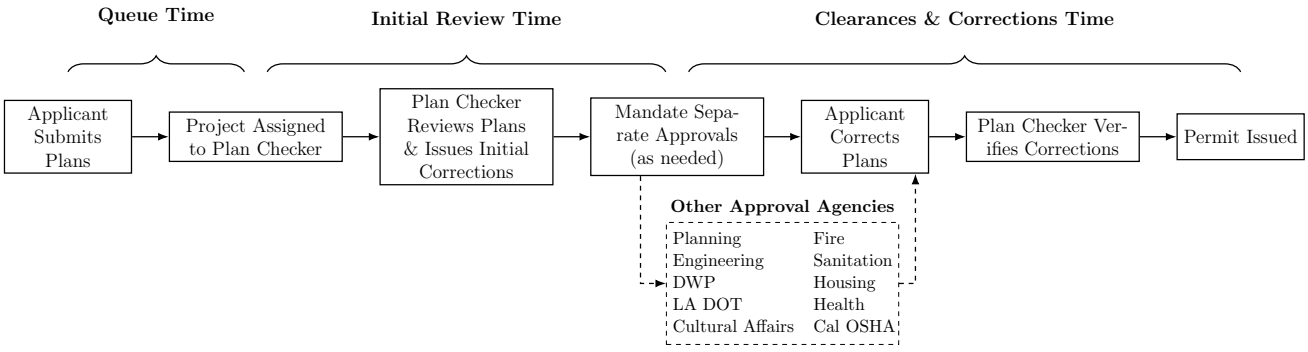
Song, Jaehee, "The effects of residential zoning in U.S. housing markets," *Journal of Urban Economics*, September 2025, *149*, 103784.

Whitt, Ward, "Approximations for the GI/G/m Queue," *Production and Operations Management*, 1993, *2* (2), 114–161.

Wrenn, Douglas H. and Elena G. Irwin, "Time is money: An empirical examination of the effects of regulatory delay on residential subdivision development," *Regional Science and Urban Economics*, March 2015, *51*, 25–36.

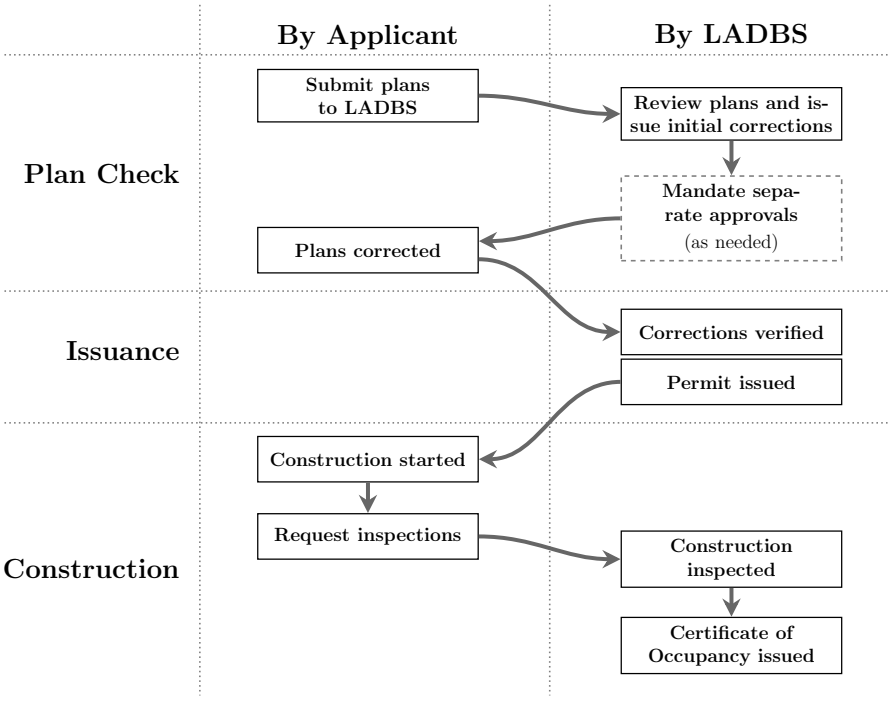
Figures and Tables

Figure 1: Timeline of LA Building Permit Process



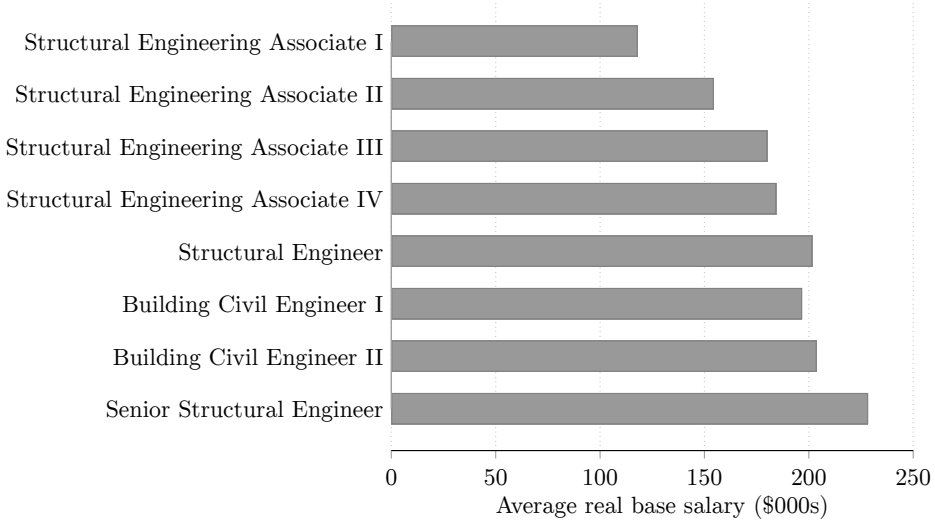
NOTES: This flowchart depicts the permitting process at Los Angeles' Department of Building and Safety (LADBS). The three bolded time phases are the definitions I use throughout the paper to refer to their respective steps. [Back to text]

Figure 2: Full LA Building Permit and Construction Process

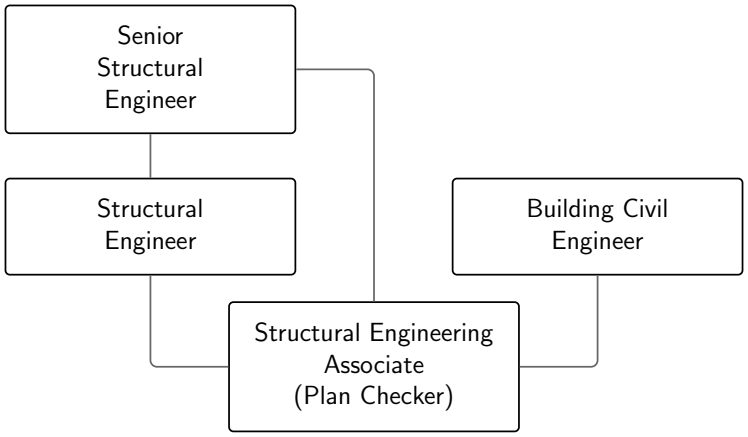


NOTES: This timeline depicts the entire permitting and construction phases in Los Angeles. The first two phases (Plan Check and Issuance) summarize the permitting timeline illustrated in Figure 1. This flowchart is adapted from the Los Angeles Department of Building & Safety' website. [Back to text]

Figure 3: Plan checkers' salaries and career ladder



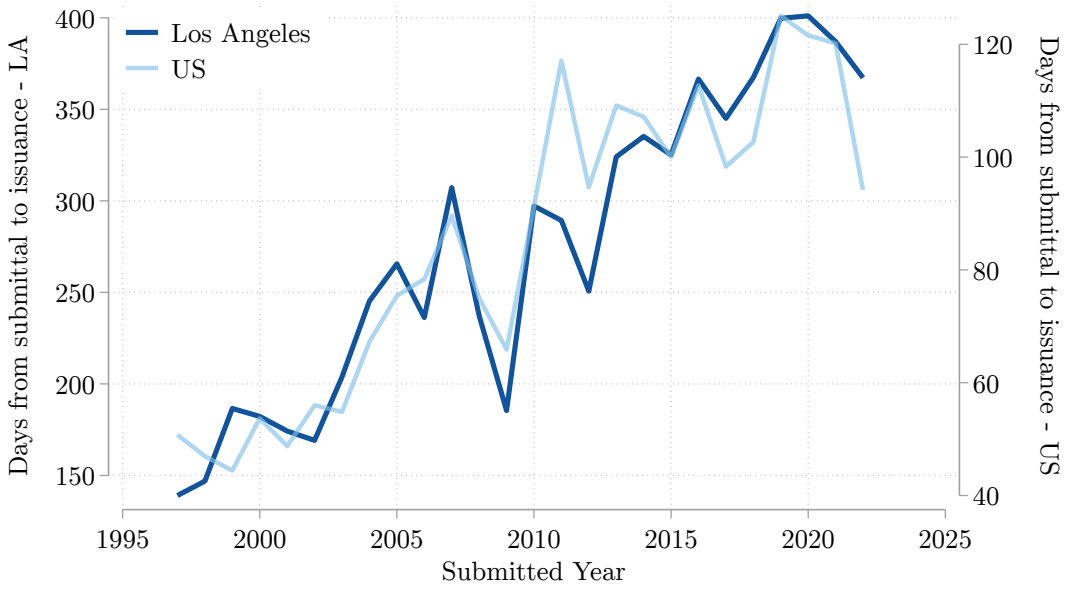
(a) Average base salary



(b) Plan checker career ladder

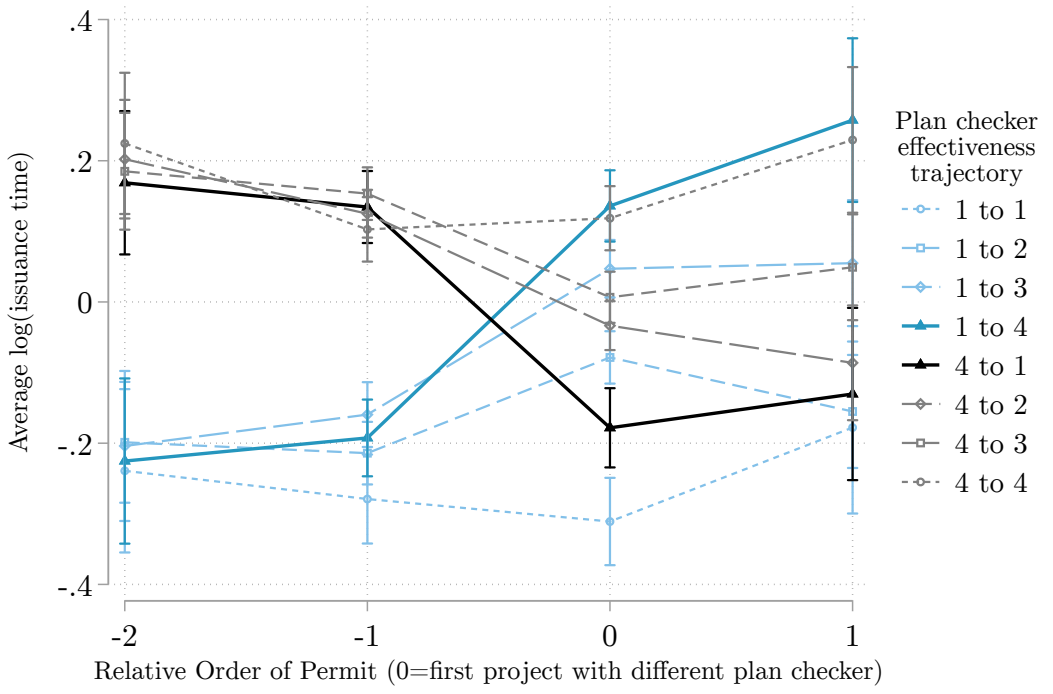
NOTES: The top panel plots the average base salary (in 2023 $) of various plan checker and manager positions, excluding benefits, overtime, or other additional pay. Plan checkers have the job title of "Structural Engineering Associate." The bottom panel plots the career ladder for plan checkers, adapted from the LA Department of Personnel's website. [Back to text]

Figure 4: Average issuance time for new residential permits



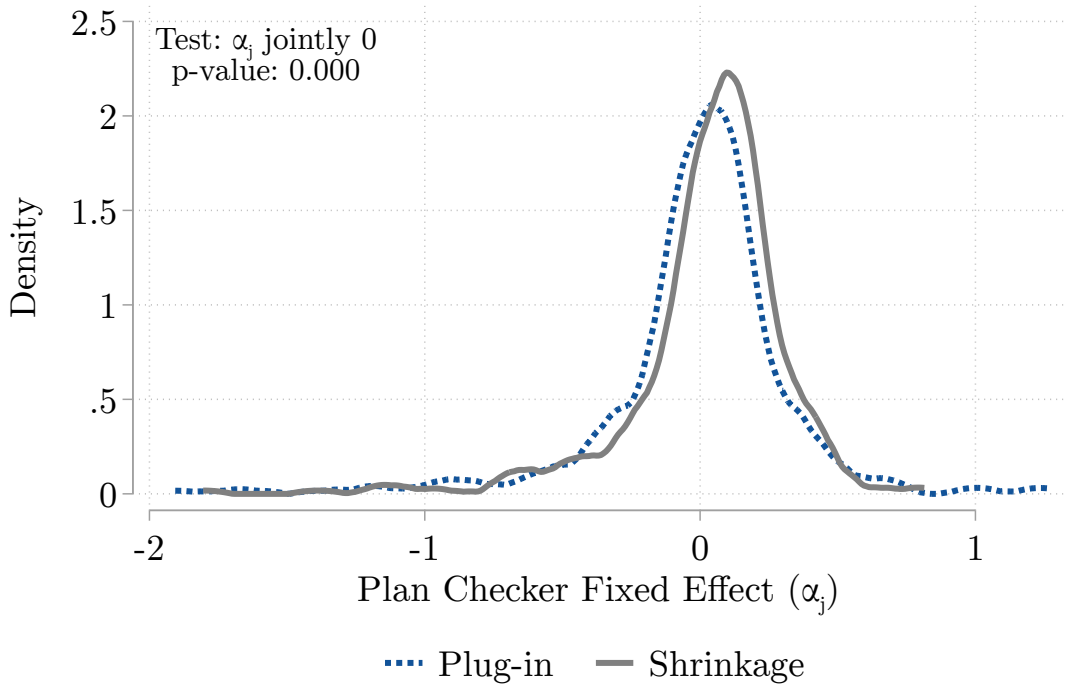
NOTES: This figure plots average permit issuance times for new residential permits, separately for Los Angeles and for all of the US permit-issuing jurisdictions available in the building permits data from Shovels. [Back to text]

Figure 5: Average issuance time around applicants that switch plan checkers



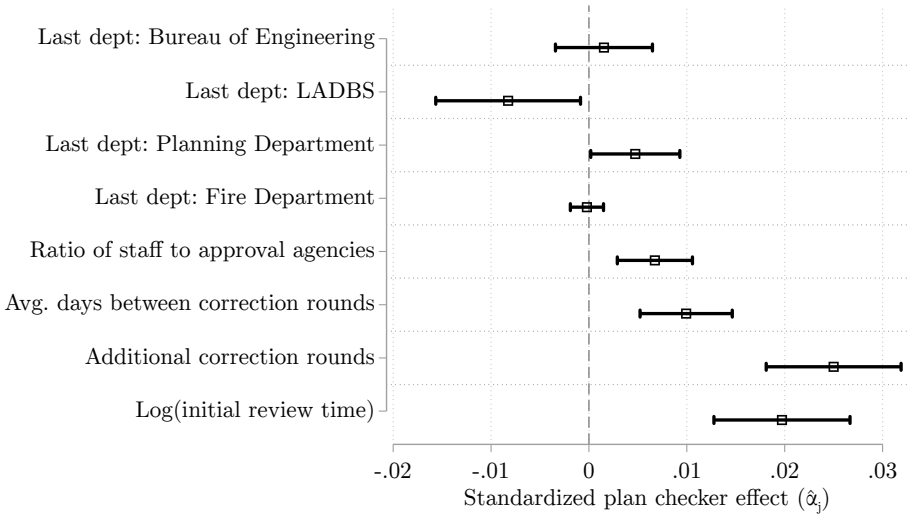
NOTES: This figure plots the average residual permit issuance time for project applicants that submit multiple permits which are assigned to a different plan checker on their subsequent permit submission. Plan checkers are classified into quartiles based on the "covariance shrinkage" estimates of $\alpha_{j(i,t)}$. Permit issuance time is residualized on the vector of time fixed effects, administrative fixed effects, and permit covariates. [Back to text]

Figure 6: Distribution of plan checker fixed effects

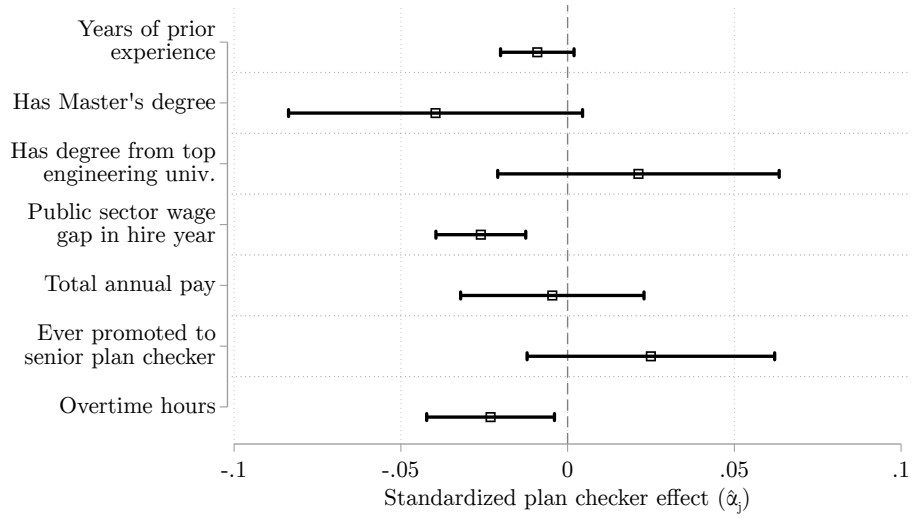


NOTES: This figure shows the distribution of plan checker fixed effects ($\alpha_{j(i,t)}$) as estimated from Equation (1), separately for the plug-in estimates and the shrinkage estimates. The p-value in the figure tests the null hypothesis that plan checker effects are jointly zero. [Back to text]

Figure 7: Correlates of plan checker performance



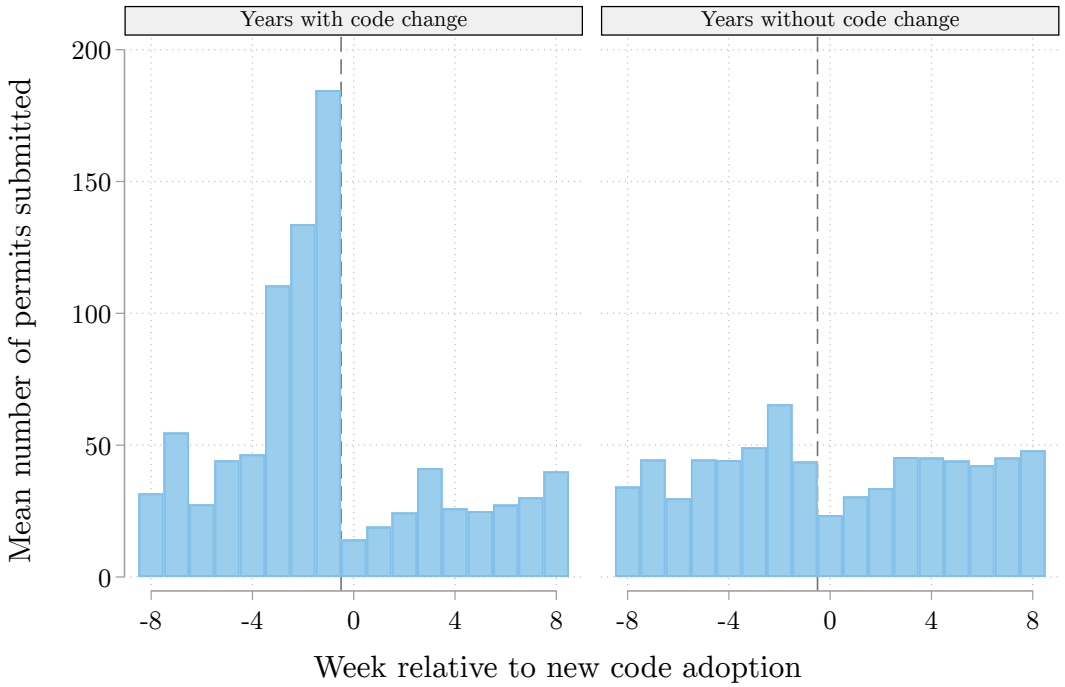
(a) Plan check process



(b) Plan checker characteristics

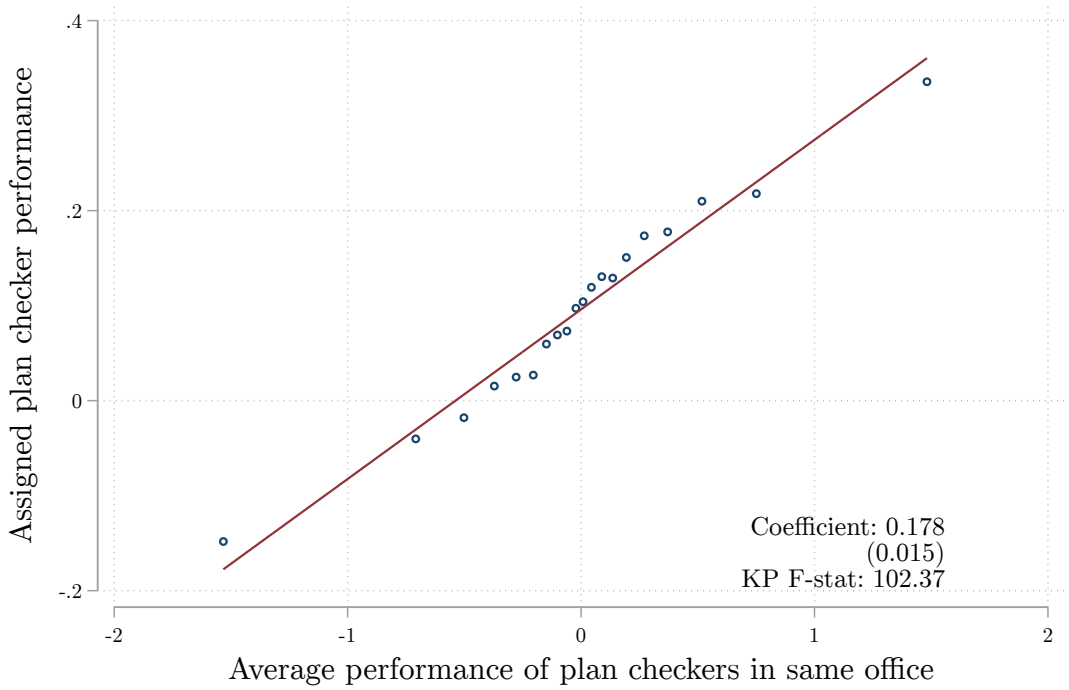
NOTES: The figure plots the coefficients of bivariate regressions between the estimated plan checker fixed effect ($\alpha_{j(i,t)}$) and the listed covariate. $\alpha_{j(i,t)}$ and the covariates are standardized to have mean zero and unit variance. For plan checker covariates, the regressions are permit-weighted. For the permit-level regressions, each covariate is first residualized on the control vector (administrative and time fixed effects from Equation (1)). Confidence intervals are computed using the heteroskedasticity-robust standard errors proposed by Kline et al. (2020). [Back to text]

Figure 8: Permit filing behavior by code cycle years



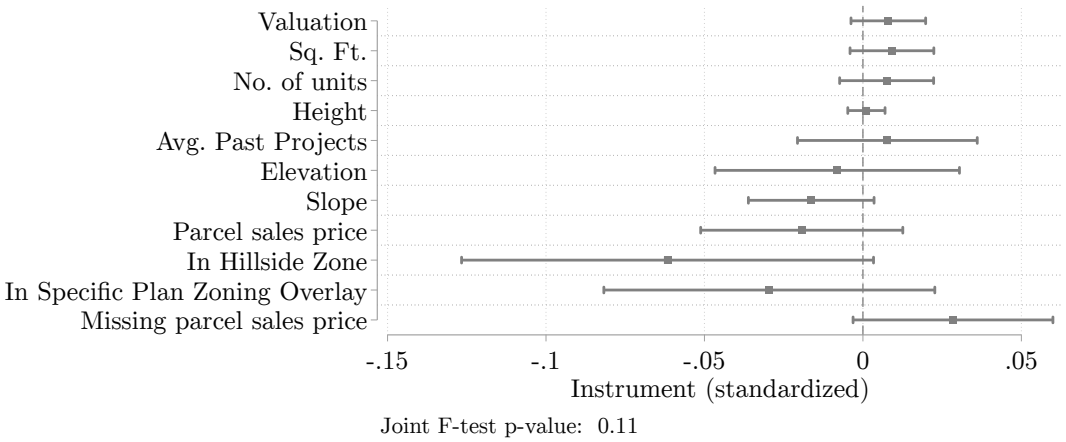
NOTES: This figure plots the average number of permits filed in each week relative to the date of when a new building code takes effect, separately for the years with a code change and the years prior to a code change. [Back to text]

Figure 9: Graphical first stage



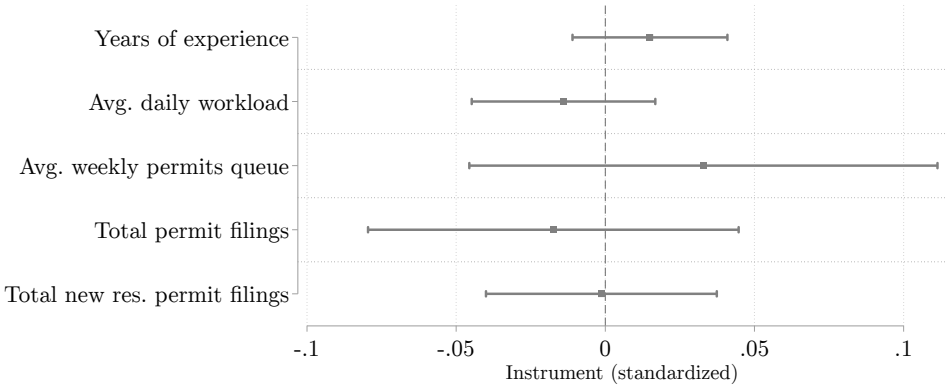
NOTES: This figure plots a binscatter of the assigned plan checker performance against the average performance of plan checkers in the same office and week of a given permit, after residualizing both variables on time fixed effects (month-year-office), administrative fixed effects (plan check division, project type), and applicant fixed effects and adding the means of both variables back to its residuals. The coefficient reports the estimated slope of the line-of-best-fit and the standard error is reported in parentheses. The standard error is clustered at the plan checker and week level. [Back to text]

Figure 10: Covariate balance



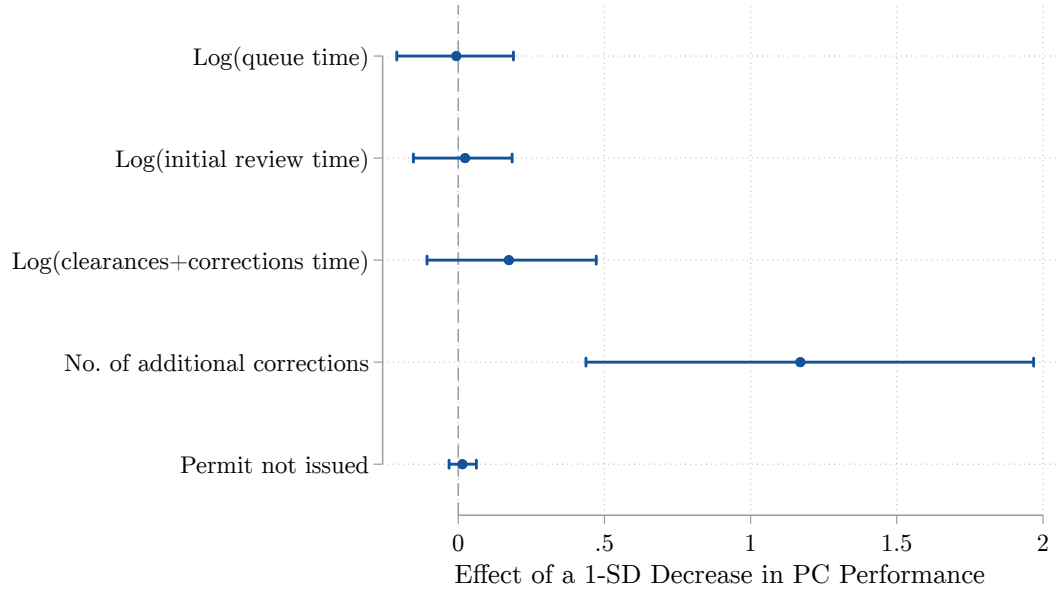
NOTES: This figure plots coefficients from regressions of the instrument and ex-ante permit characteristics, while including baseline controls (time and administrative fixed effects). The F-statistic for a test of joint significance is reported on the bottom. [Back to text]

Figure 11: Balance in other plan checker and office characteristics

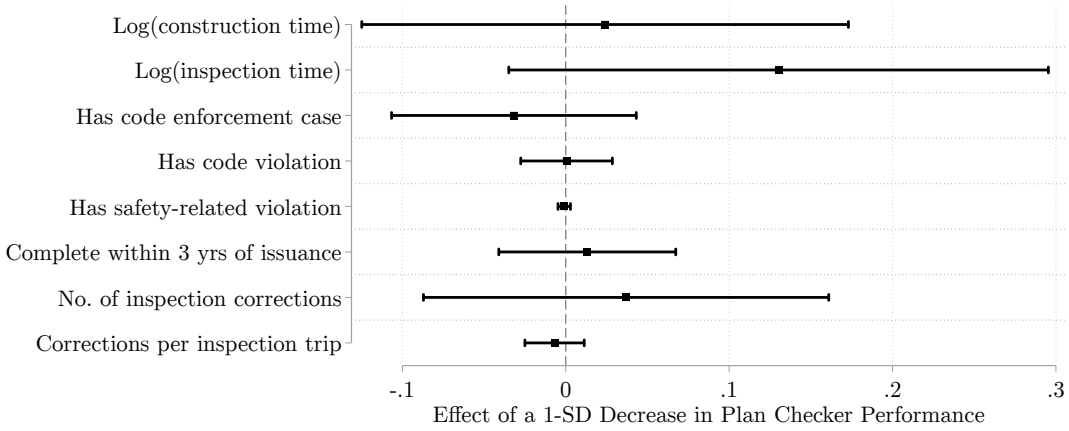


NOTES: This figure plots regression coefficients between the instrument (average performance) and the listed characteristics of the available plan checkers in the same office, conditional on the baseline controls and fixed effects. Standard errors are clustered at the plan checker and week level. [Back to text]

Figure 12: IV estimates of plan checker performance on permit outcomes



(a) Plan check outcomes



(b) Post-plan check outcomes

NOTES: This figure plots the second-stage coefficients from Equation (6) for each outcome variable as listed. Positive values on the x-axis represent being assigned a slower plan checker. Standard errors are clustered at the plan checker and week level. [Back to text]

Table 1: Sample statistics of new residential building permits

	Full sample		Analysis sample			Full sample		Analysis sample	
	Mean	SD	Mean	SD		Mean	SD	Mean	SD
Plan check phase					**Project characteristics**				
Total issuance time	287.1	(302.4)	296.1	(306.2)	Valuation (2023 $000s)	1454	(8564)	1467	(8681)
Queue time	18.0	(24.8)	18.1	(23.2)	Height (feet)	29.5	(23.4)	30.0	(24.4)
Active review time	274.9	(293.3)	279.8	(294.9)	Sq. ft.	6314.1	(28639.5)	6351.8	(28425.1)
Initial review time	9.4	(18.7)	9.7	(18.7)	Slope	4.4	(7.5)	4.6	(7.8)
Clearances+corrections time	267.2	(291.0)	271.1	(292.4)	Elevation (m)	181.2	(133.9)	183.4	(135.8)
Additional correction rounds	1.5	(2.6)	1.7	(2.7)	Uses a standard plan	0.1	(0.3)	0.1	(0.3)
Construction phase					**Clearances**				
Construction time	680.4	(662.5)	681.3	(646.4)	Total clearances	10.7	(5.4)	11.1	(5.2)
Inspection time	553.6	(573.3)	560.7	(564.9)	Approval agencies	4.3	(1.7)	4.4	(1.7)
No. of inspection corrections	3.7	(4.7)	3.8	(4.7)	Planning dept. clearances	2.4	(2.7)	2.5	(2.7)
Corrections per inspection trip	0.2	(0.2)	0.2	(0.2)	Engrg. dept. clearances	5.1	(1.8)	5.3	(1.7)
Post-construction outcomes					LADBS clearances	1.0	(1.1)	1.0	(1.1)
Has code enforcement case	0.13	(0.34)	0.12	(0.33)					
Has code violation	0.03	(0.17)	0.03	(0.16)	**Observations**	64,119		53,744	
Has safety-related code violation	0.00	(0.04)	0.00	(0.04)	Number of plan checkers	432		400	
Complete within 3 yrs of issuance	0.70	(0.46)	0.71	(0.46)	Number of applicants	3850		2041	

NOTES: This table reports means and standard deviations of permit-related characteristics for the full sample of new residential permits and the analysis sample as described in Section 3.1. Active review time refers to the time between plan check assignment and plan check approval. Approval agencies refers to the number of distinct city and county agencies a permit has to receive clearances from. [Back to text]

Table 2: Plan checker characteristics

	(1) Analysis sample mean	(2) Full sample mean
Prior years of exper.	0.94	0.88
Has Master's degree, %	0.43	0.44
Attended top engr. university, %	0.39	0.39
Years of exper. at exit	12.17	12.33
Ever promoted to manager, %	0.16	0.16
Exit at manager-level, %	0.07	0.07
Avg. biweekly overtime hours	11.96	11.88
Avg. biweekly total pay (2023 $)	5537.75	5537.97
Avg. biweekly base pay (2023 $)	4394.79	4403.18
Female, %	0.20	0.20
White non-Hispanic, %	0.30	0.30
Hispanic, %	0.24	0.23
Payroll match rate	0.83	0.78
LinkedIn profile match rate	0.51	0.48
Observations	400	432

NOTES: This table reports the summary statistics of plan checker characteristics. Column (1) reports averages for the analysis sample used in Section 5.1, and column (2) reports averages for the full sample. Educational and work history are obtained via LinkedIn profiles, and performance and demographic characteristics are obtained via payroll data. [Back to text]

Table 3: Oaxaca-Blinder decomposition of mean permit issuance time: pre- vs post-2011

	(1) Baseline FEs	(2) Add Covariates	(3) Add Clearances
Pre-2011 Mean	4.846***	4.846***	4.846***
	(0.028)	(0.026)	(0.026)
Post-2011 Mean	5.464***	5.464***	5.464***
	(0.016)	(0.016)	(0.016)
Total Difference	-0.618***	-0.618***	-0.618***
	(0.032)	(0.030)	(0.031)
Explained by Observables	0.010	-0.026	-0.300***
	(0.012)	(0.029)	(0.041)
Unexplained	-0.628***	-0.592***	-0.318***
	(0.032)	(0.031)	(0.041)
N	53744	53744	53744

NOTES: The outcome variable is log permit issuance time. The fixed effects include indicator variables for month-of-filing, project type, plan check type, and permit office. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside ordinance areas, hillside zones, historical preservation overlay zones, and special plan areas), clearances, height, project valuation, elevation, and slope. The final column includes the total number of clearances assigned as an additional covariate. [Back to text]

Table 4: Variance Decomposition of New Residential Log Permit Issuance Time

	Naïve (Plug-in)		Shrinkage		KSS estimator	
	Std. Dev.	Variance	Std. Dev.	Variance	Std. Dev.	Variance
Plan Checker	0.197	0.039	0.183	0.034	0.154	0.024
	(0.025)					
Applicant	0.413	0.171	0.373	0.139	0.345	0.119
	(0.009)					
Cov(Plan Checker, Applicant)		0.002				0.004
Log issuance time	1.036	1.073				
Log issuance time| covariates	0.824	0.680				
Number of plan checkers	400					
Number of applicants	2041					
Observations (largest connected set)	53744					
Observations (leave-out connected set)	51453					
Mean no. of applicants/plan checker	52					
Mean issuance time	296					

NOTES: This table reports variance decomposition estimates from Equation (2). Naïve (Plug-in) refers to the baseline plug-in fixed effects. Bootstrap standard errors are reported in parentheses. Shrinkage refers to the empirical Bayes adjusted estimates. KSS estimator refers to the heteroskedasticity-robust estimation of variance components on the leave-out connected set. [Back to text]

Table 5: Balance table around code changes

	Non-rushed permits (not treated)		Rushed permits (treated)		Diff. in means
	(1)	(2)	(3)	(4)	(5)
	Prior year	Code change year	Prior year	Code change year	Pre-trends test
	Mean/(SD)	Mean/(SD)	Mean/(SD)	Mean/(SD)	Coef./SE
Valuation (2023 $)	1402540.98	2304706.29	1638500.61	1478566.55	235959.63
	(5535822.07)	(11328165.63)	(9828704.10)	(5310855.80)	(346237.12)
Project sq. ft.	6366.89	9939.80	6526.18	6530.84	159.29
	(22776.31)	(43705.95)	(29249.75)	(18782.22)	(1080.74)
No. of units	4.60	7.67	4.63	4.62	0.03
	(19.89)	(41.04)	(23.43)	(18.41)	(0.78)
Height (ft.)	28.02	28.87	28.93	30.57	0.91
	(10.87)	(18.23)	(14.50)	(11.41)	(0.98)
Avg. past projects	178.72	224.09	190.48	221.44	11.76
	(266.58)	(331.56)	(262.19)	(279.15)	(40.49)
Elevation	162.02	181.27	152.50	210.75	-9.52
	(124.37)	(138.51)	(113.39)	(144.02)	(15.86)
Slope	3.56	4.90	2.99	5.60	-0.57
	(6.57)	(8.05)	(5.84)	(9.07)	(0.70)
In hillside zone	0.18	0.23	0.14	0.29	-0.04
	(0.39)	(0.42)	(0.35)	(0.45)	(0.04)
In specific plan zoning overlay	0.34	0.41	0.30	0.39	-0.04
	(0.47)	(0.49)	(0.46)	(0.49)	(0.04)
Observations	918	948	1216	2853	2134

NOTES: This table reports average permit characteristics for the four groups described in Equation (3) and coefficients from regressing permit characteristics on treatment status for the prior-year permits. Standard errors in column 5 are clustered at the plan checker and week level. Non-rushed permits are defined as permits filed either: in the second-to-last month before a code change takes effect (code change year) or in the second-to-last month before a pseudo code change (year prior to code change year). Rushed permits are defined similarly, except they are the permits filed in the last month before a code change (or pseudo code change for permits filed in the prior year). [Back to text]

Table 6: DID estimates on plan check outcomes

	Log(issuance time)	Log(initial review time)	Log(queue time)	Log(clearances+corrections time)	Additional corrections
Rushed	0.074	-0.041	0.025	0.099*	-0.195
	(0.049)	(0.056)	(0.061)	(0.048)	(0.119)
Code change year	0.013	0.142	-0.367**	0.012	-0.231
	(0.057)	(0.092)	(0.114)	(0.060)	(0.155)
Rushed × Code change year	0.139*	-0.100	0.576**	0.144*	0.300+
	(0.068)	(0.140)	(0.113)	(0.068)	(0.163)
N	5219	5177	5194	5173	5219

NOTES: The table reports coefficients from estimating Equation (3). The fixed effects include indicator variables for year-office, project type, plan check type, project applicant, and assigned plan checker. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside zones, and special plan areas), clearances, height, project valuation, land use density, elevation, and slope. Standard errors are clustered at the plan checker and week level. [Back to text]

Table 7: DID estimates on plan check outcomes: heterogeneity by plan checker effectiveness

	Log(issuance time)	Log(initial review time)	Log(queue time)	Log(clearances+corrections time)	Additional corrections
Panel A: Above median plan checkers					
Rushed	0.073	0.006	0.071	0.126+	0.234+
	(0.066)	(0.067)	(0.081)	(0.074)	(0.137)
Code change year	0.099	0.118	-0.316**	0.166+	0.396*
	(0.074)	(0.072)	(0.091)	(0.088)	(0.180)
Rushed × Code	0.068	-0.039	0.532**	-0.035	-0.359
change year	(0.092)	(0.094)	(0.118)	(0.114)	(0.218)
N	2366	2340	2345	2338	2366
Panel B: Below median plan checkers					
Rushed	0.104+	-0.004	-0.058	0.178**	0.050
	(0.063)	(0.077)	(0.065)	(0.063)	(0.192)
Code change year	0.026	0.198	-0.446**	0.049	-0.116
	(0.079)	(0.135)	(0.169)	(0.074)	(0.164)
Rushed × Code	0.199*	-0.244	0.689**	0.165+	0.164
change year	(0.093)	(0.215)	(0.164)	(0.091)	(0.189)
N	2659	2646	2656	2645	2659

NOTES: The table reports coefficients from estimating Equation (3) separately for plan checkers above and below the median $\hat{\alpha}_{j(i,t)}$. The fixed effects include indicator variables for year-office, project type, plan check type, and project applicant. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside zones, and special plan areas), clearances, height, project valuation, land use density, elevation, and slope. Standard errors are clustered at the plan checker and week level. [Back to text]

Table 8: Effects of plan checker performance on plan check outcomes

	Log(queue time)	Log(initial review time)	Log(clearances+corrections time)	Additional correction rounds
Panel A: OLS				
Jackknife staff	0.116**	0.096**	0.195**	0.507**
performance (std.)	(0.025)	(0.025)	(0.030)	(0.122)
Panel B: Reduced Form				
Avg. performance of	-0.001	0.003	0.028	0.218**
available PC	(0.015)	(0.012)	(0.023)	(0.069)
Panel C: First Stage				
Avg. performance of	0.150**	0.144**	0.164**	0.186**
available PC	(0.013)	(0.013)	(0.016)	(0.015)
Panel D: IV using availability instrument				
Jackknife staff	-0.007	0.023	0.173	1.170**
performance (std.)	(0.100)	(0.085)	(0.145)	(0.385)
tF standard error	(0.100)	(0.085)	(0.145)	(0.385)
AR 95% confidence interval	[-0.210, 0.189]	[-0.154, 0.184]	[-0.107, 0.472]	[0.437, 1.968]
$\hat{r}$	0.122	0.261	-0.170	-0.259
First-stage F	142.32	119.68	109.92	149.12
N	41985	41043	33631	45510
notes	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates

NOTES: This table displays the OLS, reduced form, and IV estimates of the impact of plan checker performance on various permit outcomes during the plan check phase. "tF standard error" is the adjusted standard error using the adjustments in Lee et al. (2022). "AR 95% confidence interval" is the 95% confidence interval using Anderson and Rubin (1949). $\hat{r}$ is the empirical correlation between the two-stage least squares main equation residual $\hat{\eta}$ and the first stage residual $\hat{u}$. The fixed effects include indicator variables for month-year-by-office, project type, plan check type, and project applicant. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside zones, and special plan areas), clearances, height, project valuation, land use density, elevation, and slope. Standard errors are clustered at the plan checker and week level. [Back to text]

Table 9: Effects of plan checker performance on permit outcomes

	Log(construction time)	Log(inspection time)	Has code enforcement case	Has code violation	Has safety-related violation	Complete within 3 years of issuance	No. of inspection corrections	Ratio of corrections per inspection trip
Panel A: OLS								
Jackknife staff	0.035*	0.059**	-0.011	-0.009**	-0.001	-0.002	0.055***	0.004
performance (std.)	(0.015)	(0.019)	(0.009)	(0.003)	(0.001)	(0.006)	(0.016)	(0.003)
Panel B: Reduced Form								
Avg. performance of	0.005	0.025	-0.006	0.000	-0.000	0.002	0.007	-0.001
available PC	(0.015)	(0.016)	(0.007)	(0.003)	(0.000)	(0.005)	(0.012)	(0.002)
Panel C: First Stage								
Avg. performance of	0.203**	0.192**	0.186**	0.186**	0.186**	0.186**	0.191**	0.191**
available PC	(0.021)	(0.017)	(0.015)	(0.015)	(0.015)	(0.015)	(0.017)	(0.017)
Panel D: IV using availability instrument								
Jackknife staff	0.024	0.130	-0.032	0.000	-0.001	0.013	0.037	-0.007
performance (std.)	(0.076)	(0.084)	(0.038)	(0.014)	(0.002)	(0.028)	(0.063)	(0.009)
tF standard error	(0.076)	(0.084)	(0.038)	(0.014)	(0.002)	(0.028)	(0.063)	(0.009)
AR 95% confidence interval	[-0.126, 0.176]	[-0.037, 0.297]	[-0.108, 0.043]	[-0.026, 0.031]	[-0.005, 0.003]	[-0.038, 0.071]	[-0.088, 0.163]	[-0.026, 0.011]
$\hat{r}$	-0.028	0.009	0.053	-0.397	0.018	-0.395	-0.038	0.149
First-stage F	97.62	123.54	149.12	149.12	149.12	149.12	132.07	132.07
N	40931	42908	45510	45510	45510	45510	43379	43379
notes	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates	FEs+covariates

NOTES: This table displays the OLS, reduced form, and IV estimates of the impact of plan checker performance on various permit outcomes after the plan check phase. "tF standard error" is the adjusted standard error using the adjustments in Lee et al. (2022). "AR 95% confidence interval" is the 95% confidence interval using Anderson and Rubin (1949). $\hat{r}$ is the empirical correlation between the two-stage least squares main equation residual $\hat{\eta}$ and the first stage residual $\hat{u}$. The fixed effects include indicator variables for month-quarter-by-office, project type, plan check type, and project applicant. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside zones, and special plan areas), clearances, height, project valuation, land use density, elevation, and slope. Standard errors are clustered at the plan checker and week level. [Back to text]

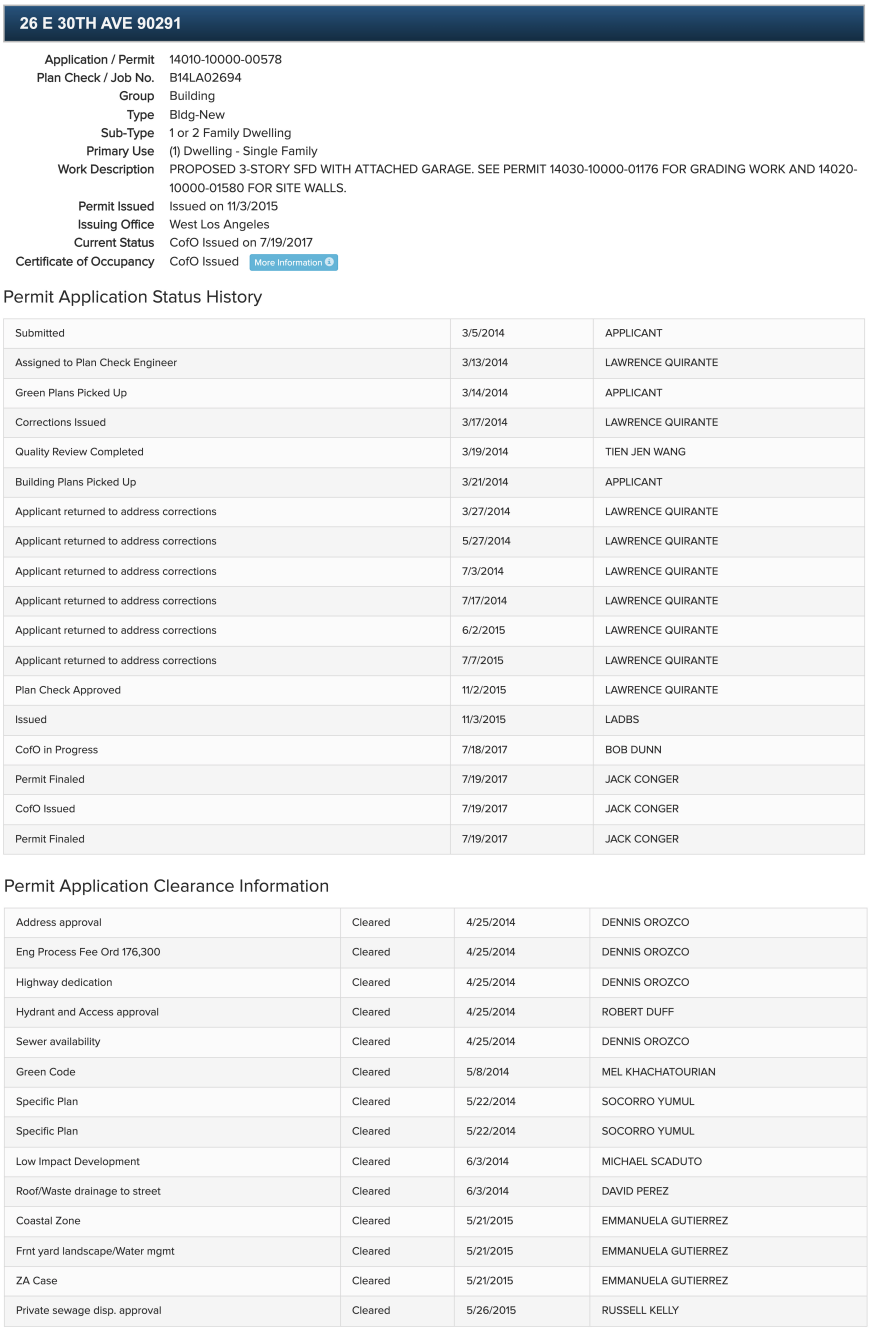
Table 10: Alternative personnel policies

Replaced Group	Replacement Level	Private Savings ($ M)	Days Saved (000's)	City Cost ($ M)	City Cost From Overtime ($ M)	Net Benefit ($ M)	ROI
Lowest 25%	Median	262.97	598.38	29.14	8.43	233.82	9.02
hold hours fixed		146.46	467.67	18.06	-0.44	128.40	8.11
Lowest 25%	Top 25%	573.43	1,268.35	21.66	5.99	551.77	26.47
hold hours fixed		457.64	1,112.40	13.11	-0.85	444.53	34.90
Lowest 10%	Median	95.43	175.79	9.51	3.84	85.92	10.03
hold hours fixed		60.16	141.08	4.75	0.03	55.41	12.67
Lowest 10%	Top 25%	162.83	275.96	7.37	3.14	155.46	22.10
hold hours fixed		128.49	238.30	3.40	-0.04	125.09	37.84
10% pay raise	Better selection	125.81	239.16	32.07	7.69	93.74	3.92
hold hours fixed		117.62	225.25	31.16	6.77	86.46	3.77

NOTES: This table displays the private cost savings, city payroll costs, net benefits, and return on each dollar invested from various personnel policy counterfactuals. The private cost savings are computed as follows: use the estimated $\hat{\alpha}_{j(i,t)}$ to predict the days saved from the given policy change (denoted d), then converted to a dollar amount by aggregating across the financing costs (the project value in 2023$ $\times$ the loan-to-value ratio $\times$ the annual interest rate $\times$ $d/365$), carrying costs (the project value $\times$ the fraction of total costs for land $\times$ property tax rate $\times$ $d/365$), and foregone rents (the project value $\times$ capitalization rate $\times$ $d/365$). I assume a loan-to-value ratio of 0.7, interest rate of 6%, land costs are 25% of total project value, a 1% property tax rate, and a capitalization rate of 4%. The city payroll costs are computed as the change in annual pay from a given personnel policy change across all pay years, plus a frictional hiring cost equal to 25% of the average annual pay of new hires. For the first four scenarios, I also estimate the savings and costs while holding overtime hours fixed, by using an estimate of $\hat{\alpha}_{j(i,t)}$ that controls for hours worked and assuming that the replacement plan checkers work the same average annual overtime hours as the plan checkers they replaced. For the 10% pay raise scenario, I use the wage gap elasticity from Section 6.1 to impute how higher pay might shift the distribution of plan checker performance to arrive at an estimate of days saved, and I also estimate a scenario that excludes the additional overtime pay from hiring faster plan checkers. [Back to text]

Appendix Figures and Tables

Figure A1: An example of a LADBS building permit and its status updates



26 E 30TH AVE 90291

Application / Permit	14010-10000-00578
Plan Check / Job No.	B14LA02694
Group	Building
Type	Bldg-New
Sub-Type	1 or 2 Family Dwelling
Primary Use	(1) Dwelling - Single Family
Work Description	PROPOSED 3-STORY SFD WITH ATTACHED GARAGE. SEE PERMIT 14030-10000-01176 FOR GRADING WORK AND 14020-10000-01580 FOR SITE WALLS.
Permit Issued	Issued on 11/3/2015
Issuing Office	West Los Angeles
Current Status	CofO Issued on 7/19/2017
Certificate of Occupancy	CofO Issued More Information

Permit Application Status History

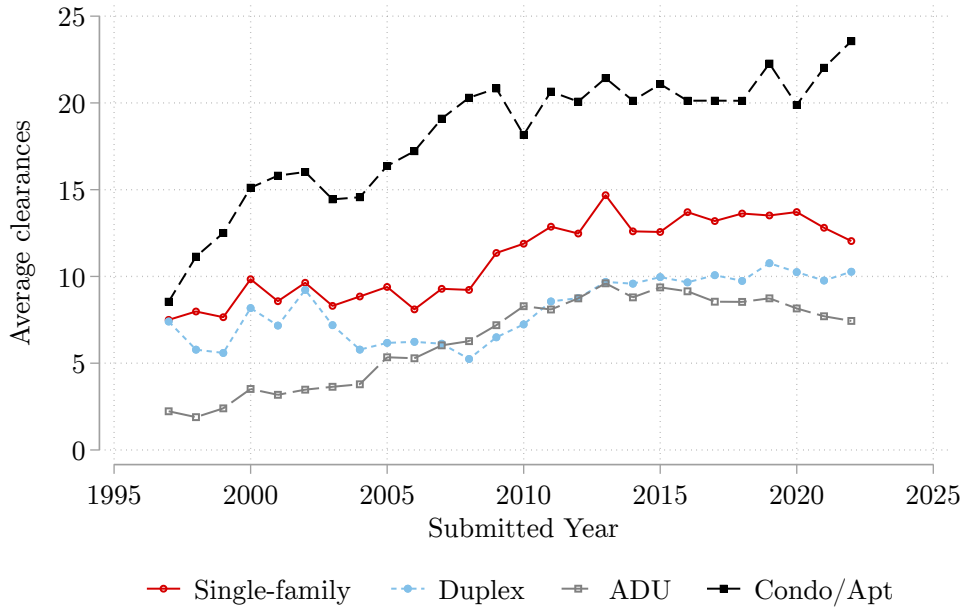
Submitted	3/5/2014	APPLICANT
Assigned to Plan Check Engineer	3/13/2014	LAWRENCE QUIRANTE
Green Plans Picked Up	3/14/2014	APPLICANT
Corrections Issued	3/17/2014	LAWRENCE QUIRANTE
Quality Review Completed	3/19/2014	TIEN JEN WANG
Building Plans Picked Up	3/21/2014	APPLICANT
Applicant returned to address corrections	3/27/2014	LAWRENCE QUIRANTE
Applicant returned to address corrections	5/27/2014	LAWRENCE QUIRANTE
Applicant returned to address corrections	7/3/2014	LAWRENCE QUIRANTE
Applicant returned to address corrections	7/17/2014	LAWRENCE QUIRANTE
Applicant returned to address corrections	6/2/2015	LAWRENCE QUIRANTE
Applicant returned to address corrections	7/7/2015	LAWRENCE QUIRANTE
Plan Check Approved	11/2/2015	LAWRENCE QUIRANTE
Issued	11/3/2015	LADBS
CofO in Progress	7/18/2017	BOB DUNN
Permit Finaled	7/19/2017	JACK CONGER
CofO Issued	7/19/2017	JACK CONGER
Permit Finaled	7/19/2017	JACK CONGER

Permit Application Clearance Information

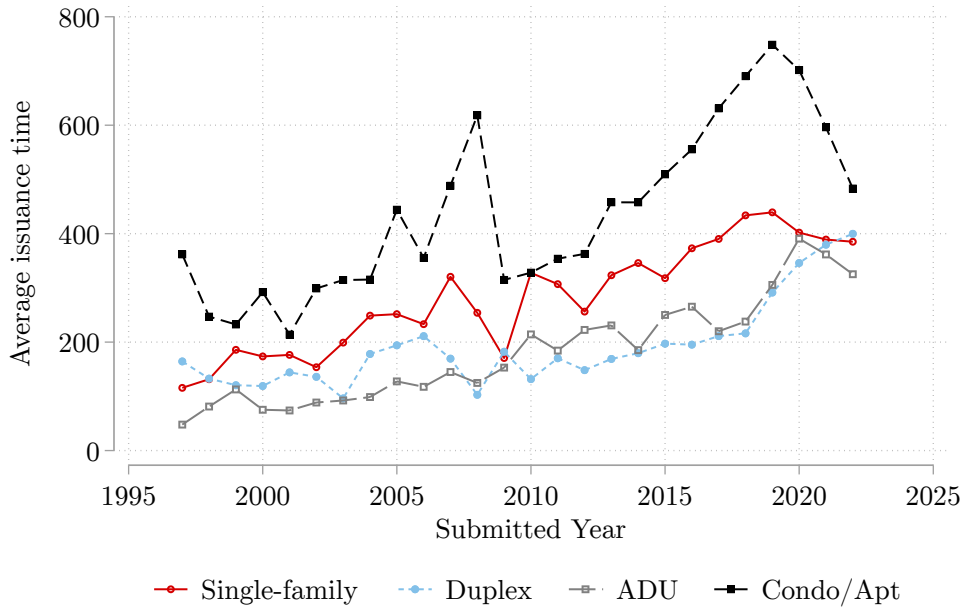
Address approval	Cleared	4/25/2014	DENNIS OROZCO
Eng Process Fee Ord 176,300	Cleared	4/25/2014	DENNIS OROZCO
Highway dedication	Cleared	4/25/2014	DENNIS OROZCO
Hydrant and Access approval	Cleared	4/25/2014	ROBERT DUFF
Sewer availability	Cleared	4/25/2014	DENNIS OROZCO
Green Code	Cleared	5/8/2014	MEL KHACHATOURIAN
Specific Plan	Cleared	5/22/2014	SOCORRO YUMUL
Specific Plan	Cleared	5/22/2014	SOCORRO YUMUL
Low Impact Development	Cleared	6/3/2014	MICHAEL SCADUTO
Roof/Waste drainage to street	Cleared	6/3/2014	DAVID PEREZ
Coastal Zone	Cleared	5/21/2015	EMMANUELA GUTIERREZ
Frnt yard landscape/Water mgmt	Cleared	5/21/2015	EMMANUELA GUTIERREZ
ZA Case	Cleared	5/21/2015	EMMANUELA GUTIERREZ
Private sewage disp. approval	Cleared	5/26/2015	RUSSELL KELLY

NOTES: This figure depicts an example of a building permit that was web-scraped. The permit data contains the status updates, assigned plan checker names, and additional permit information. [Back to text]

Figure A2: Average clearances and issuance time by project type in LA



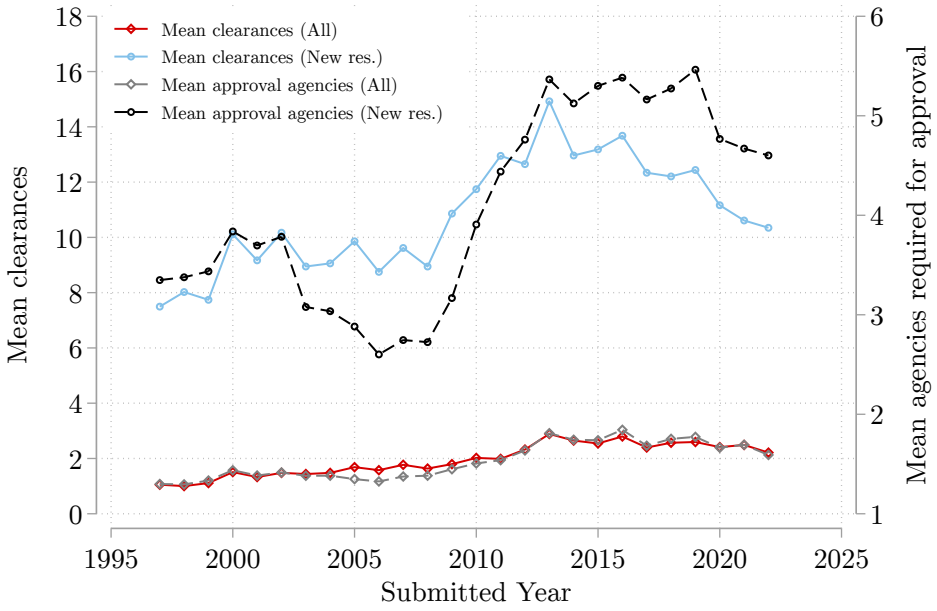
(a) Average clearances by project type in LA



(b) Average issuance time by project type in LA

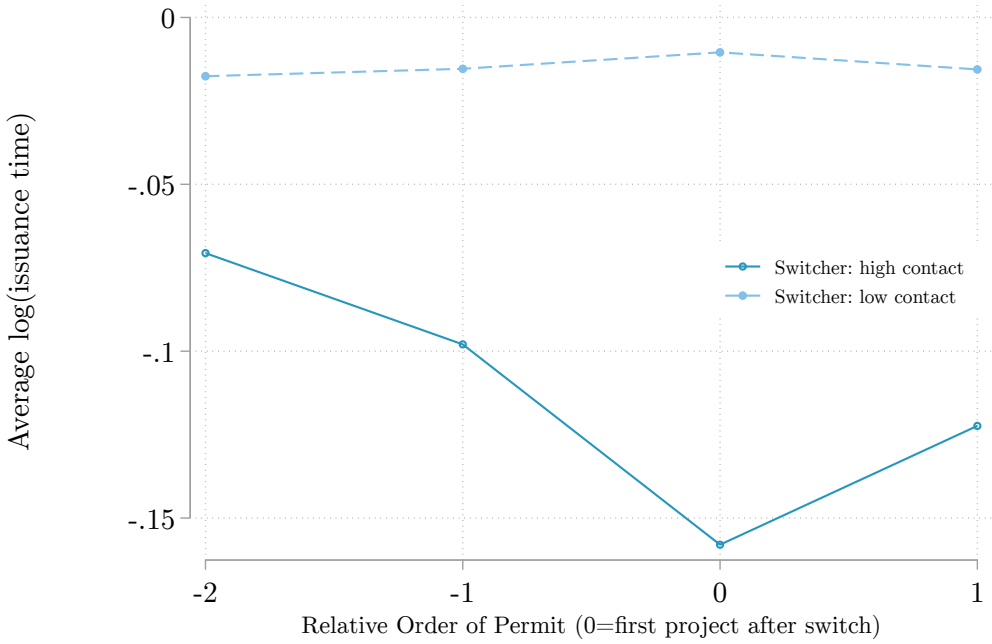
NOTES: These figures plot the average clearances assigned and issuance times for new residential permits in LA, separately by project type categories in LADBS permit data. [Back to text]

Figure A3: Average clearances and approval agencies in Los Angeles

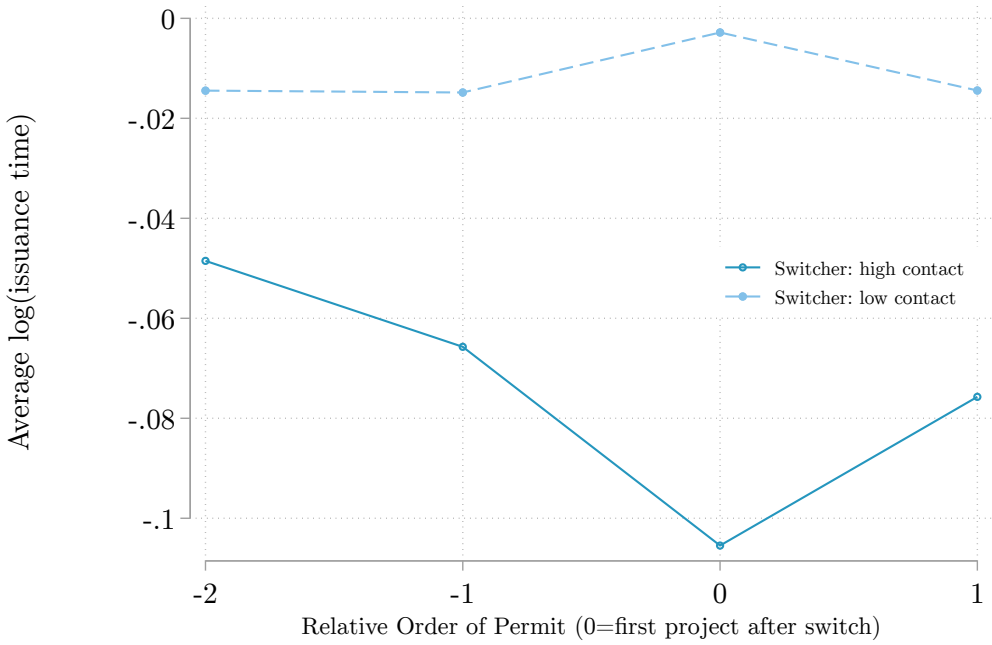


NOTES: The figure plots the average number of clearances (solid lines) and approval agencies (dashed lines) assigned in LADBS data by submittal year, separately for all permits and new residential permits. Clearances are plotted on the left y-axis, and approval agencies are plotted on the right y-axis. [Back to text]

Figure A4: Average issuance time around applicants that switch plan checkers based on high-contact vs low-contact



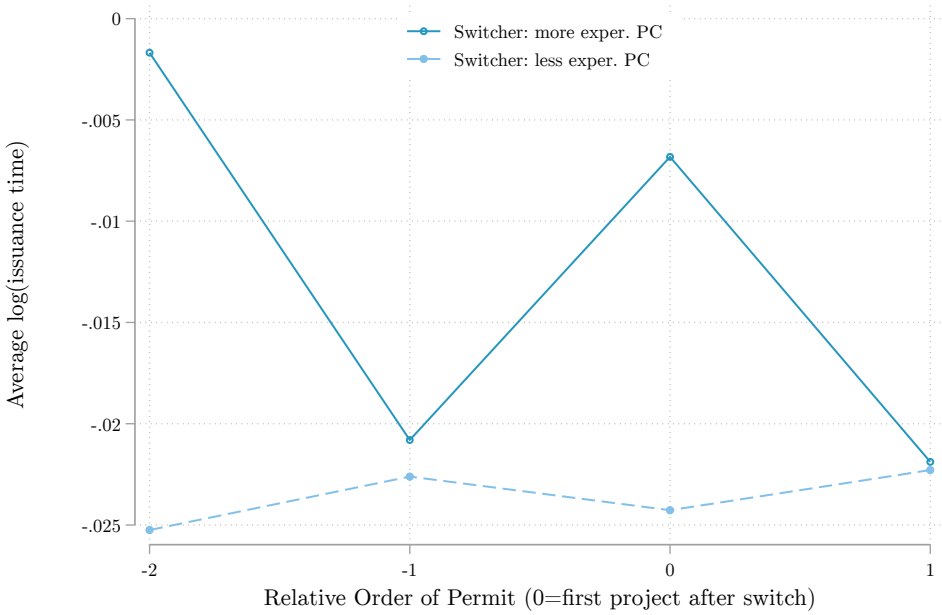
(a) High contact: >2 past projects with same plan checker after switch



(b) High contact: higher project share after switch

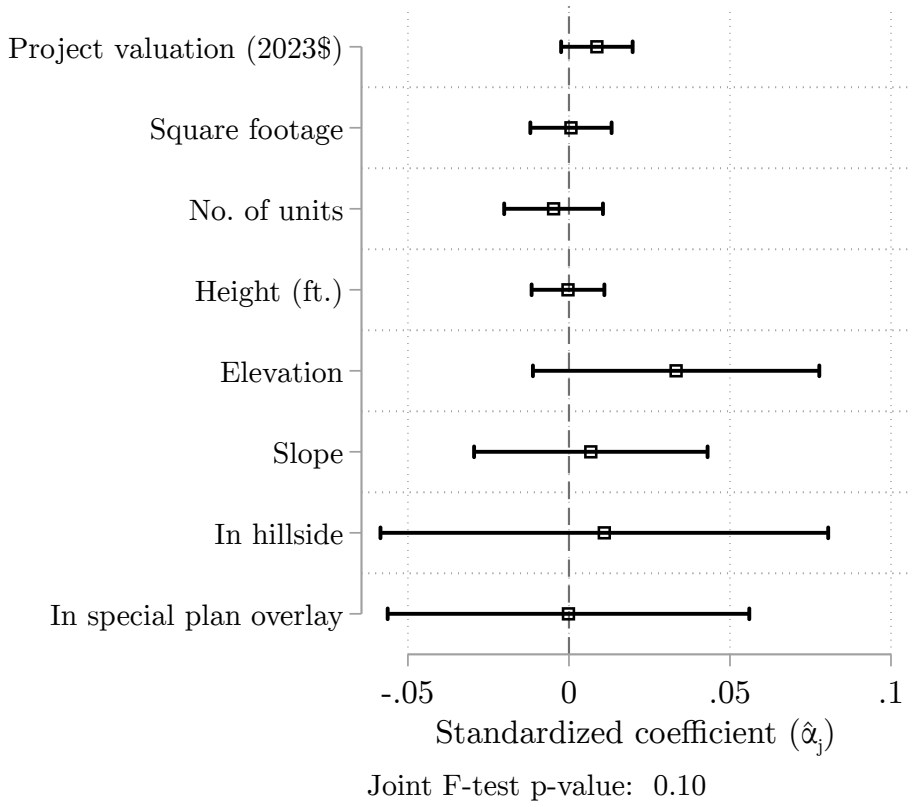
NOTES: This figure plots the average residual permit issuance time for project applicants that submit multiple permits which are assigned to a different plan checker on their subsequent permit submission ("switchers"). The top panel plots the average issuance time for switchers who are assigned to a plan checker who they have worked with at least 3 times on past projects. The bottom panel plots the average issuance time for switchers who are assigned to a plan checker with whom the switcher has had more experience with, compared to the previous plan checker that this switcher was assigned to. Permit issuance time is residualized on the vector of time and administrative fixed effects. [Back to text]

Figure A5: Average issuance time around applicants that switch plan checkers based on plan checkers' past expertise



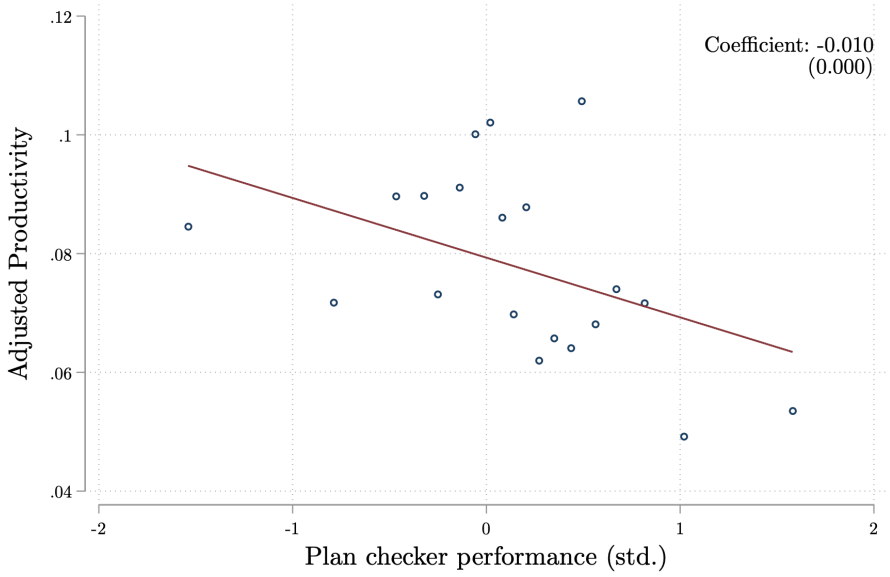
NOTES: This figure plots the average residual permit issuance time for project applicants that submit multiple permits which are assigned to a different plan checker on their subsequent permit submission ("switchers"), separately by whether the switcher is assigned to a plan checker who has worked on the same project type more than 50% of the time, relative to the previous plan checker that the switcher applicant was assigned to. [Back to text]

Figure A6: Selection test: covariate balance



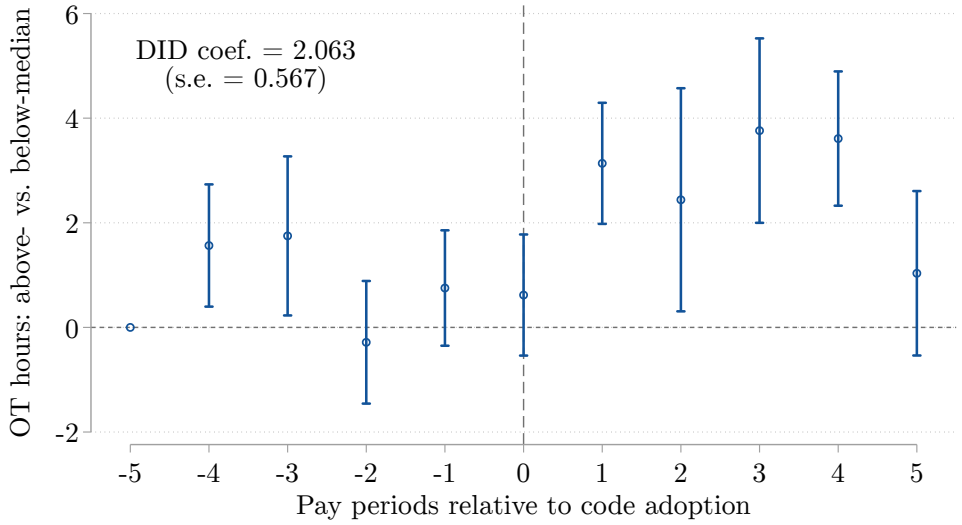
NOTES: This figure plots coefficients from regressions of the plan checker effect ($\hat{\alpha}_j$) on each ex-ante permit characteristic, while including the same administrative and time controls from Equation (1). This regression includes permits excluded from the analysis sample. The p-value from a joint test of significance is reported below. [Back to text]

Figure A7: Binscatter of plan checker performance and weekly productivity



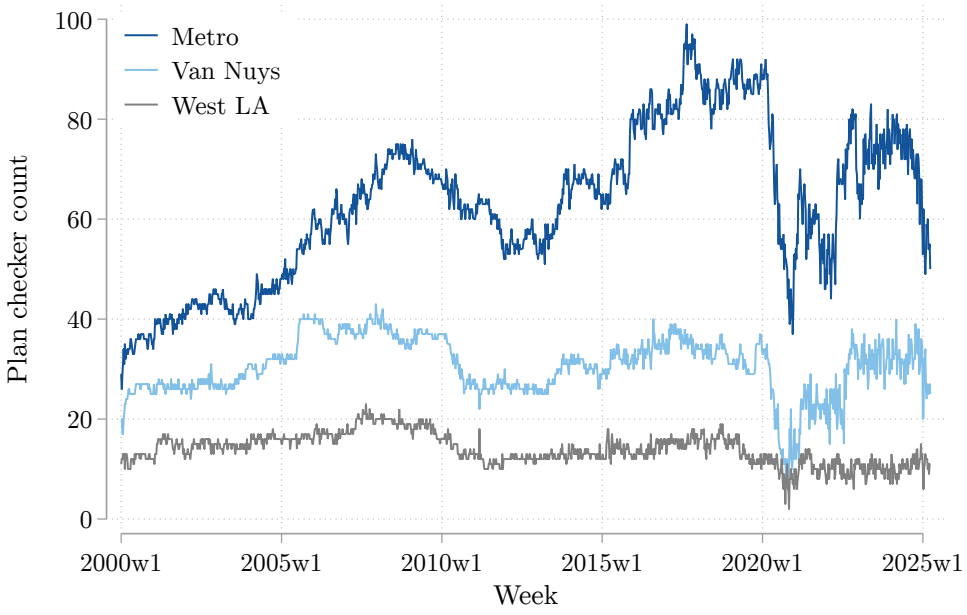
NOTES: This figure plots a binscatter of plan checker performance against the weekly adjusted productivity of that plan checker. The adjusted productivity measure is constructed by estimating each permit's expected issuance time based on its characteristics, prorating the expected time across weeks the permit was active, and dividing each plan checker's weekly sum of prorated expected time by their total hours worked (regular plus overtime). The coefficient reports the estimated slope of the line-of-best-fit and the robust standard error is reported in parentheses. [Back to text]

Figure A8: Average overtime hours around code cycle changes by above/below median plan checker $\hat{\alpha}$



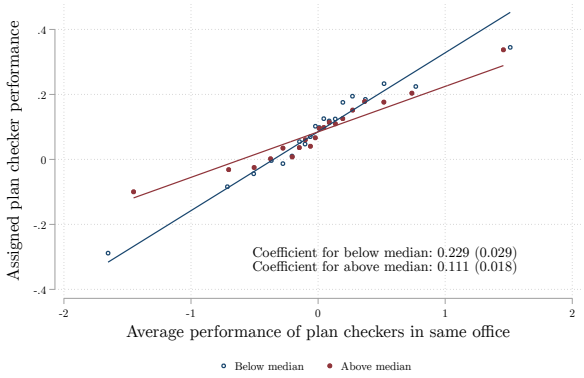
NOTES: The figure plots event study estimates of the average differential overtime hours worked by above-median plan checkers (compared to below-median plan checkers) in each pay period relative to the date of when a new building code takes effect. Since a pay period spans two weeks, the last date of the pay period is used to determine the relative event time. The event study coefficients are β_w estimated from the following regression: $Y_{jt} = \sum_{w=-4}^{5} \beta_w \left(\mathbf{1}[\tau_w = w] \times G_j\right) + \tau_w + \mu_{c(t)} \times G_j + \varepsilon_{jt}$ where τ_w indicates event-time dummy variables for relative pay period w, G_j is an indicator variable equal to 1 for a plan checker j whose $\hat{\alpha}_j$ is above the median $\hat{\alpha}_j$, and $\mu_{c(t)}$ are code year cycle indicator variables. The coefficient reported in the figure is β_1 estimated via the following regression: $Y_{jt} = \beta_1(\text{Post}_{jt} \times G_j) + \beta_2 \text{Post}_{jt} + \beta_3 G_j + \mu_{c(t)} \times G_j + \varepsilon_{it}$ where Post_{jt} is an indicator variable equal to 1 for the 5 pay periods after a code change date. [Back to text]

Figure A9: Weekly plan checker availability by submittal office

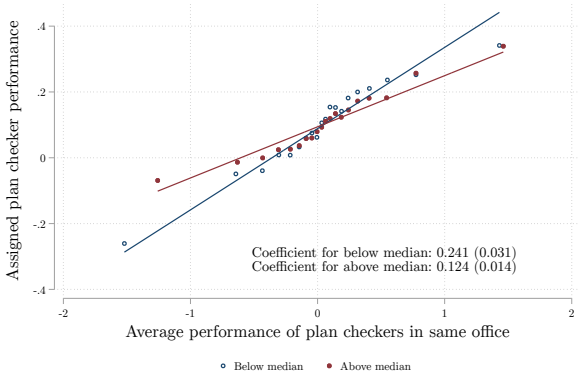


NOTES: This figure plots the number of plan checkers that are assigned to all permits in a given week, separately for the three largest LADBS offices. [Back to text]

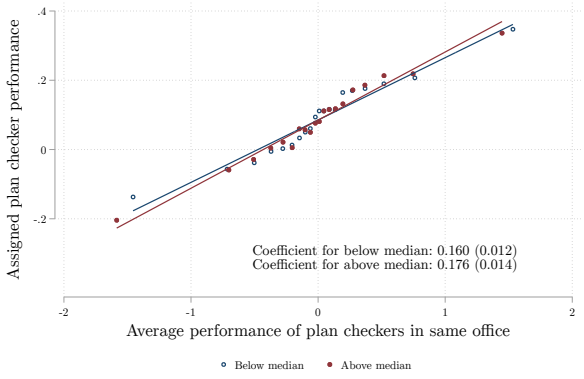
Figure A10: Monotonicity test



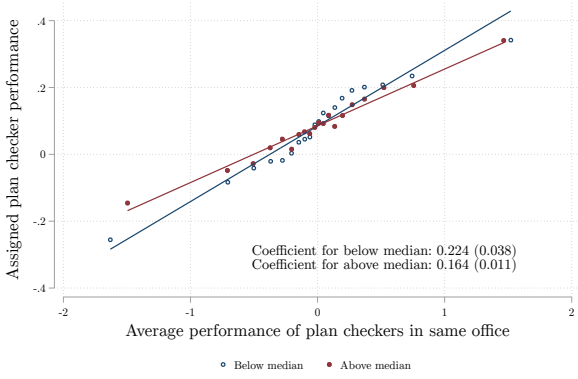
(a) By project value



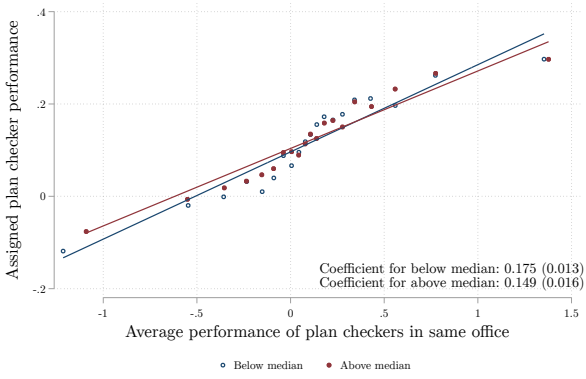
(b) By square footage



(c) By slope



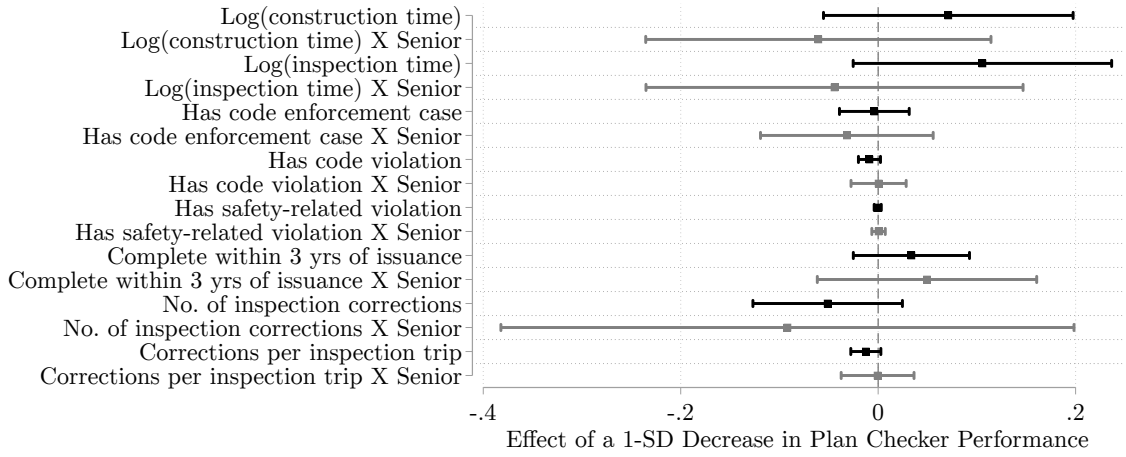
(d) By elevation



(e) By height

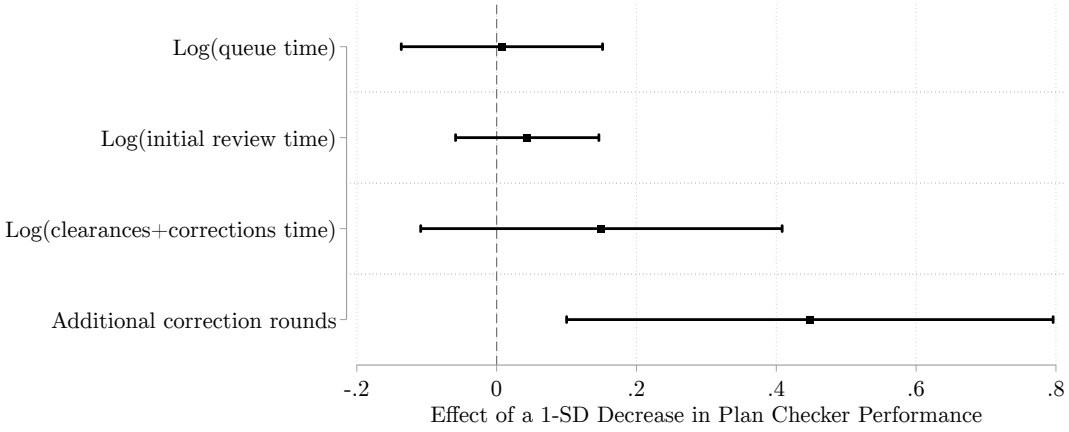
NOTES: This figure plots graphical first stage relationships for various subgroups defined by the median value of the listed permit characteristics. The same binscatter approach used in Figure 9 is used here. The reported coefficients are the estimated slope of the line-of-best-fit for each subsample, and the standard errors are reported in parentheses. Standard errors are clustered at the plan checker and week level. [Back to text]

Figure A11: IV estimates of plan checker performance on permit outcomes: by seniority

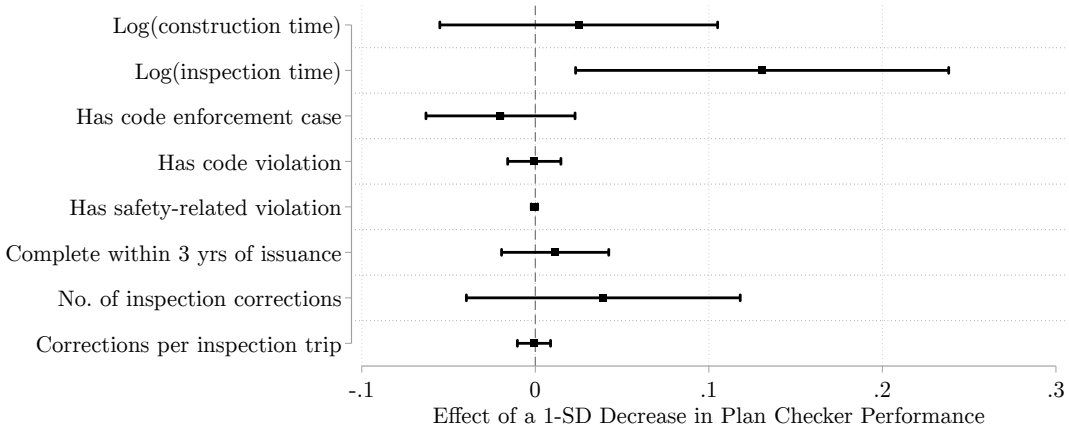


NOTES: This figure plots the second-stage coefficients from estimating Equation (8) for each outcome variable as listed. Positive values on the x-axis represent being assigned a slower plan checker. Standard errors are clustered at the plan checker and week level. [Back to text]

Figure A12: IV estimates of plan checker performance on permit outcomes: 3 largest permit offices



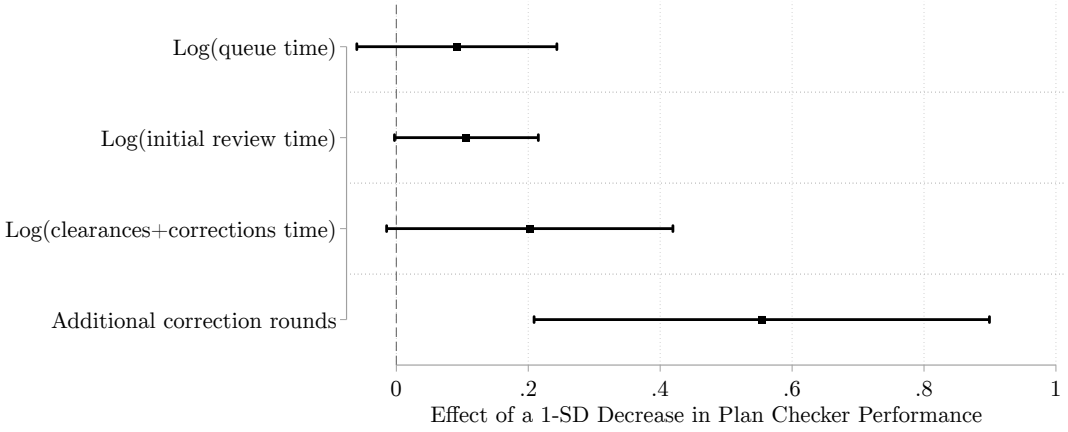
(a) Plan check outcomes



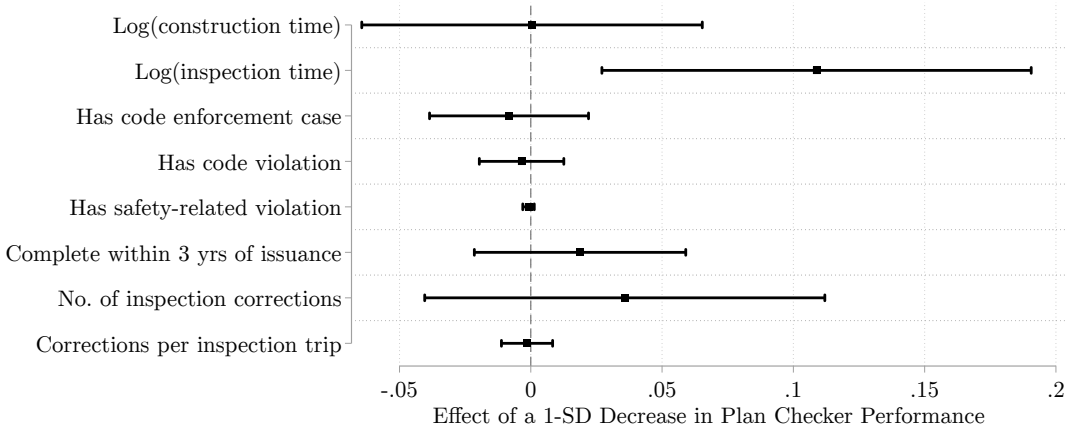
(b) Post-plan check outcomes

NOTES: This figure plots the second-stage coefficients from Equation (6) for each outcome variable as listed. Positive values on the x-axis represent being assigned to a slower plan checker. Standard errors are clustered at the plan checker and week level. The sample is restricted to permits submitted at the 3 largest permit offices. [Back to text]

Figure A13: IV estimates of plan checker performance on permit outcomes: exclude "rushed" permits



(a) Plan check outcomes



(b) Post-plan check outcomes

NOTES: This figure plots the second-stage coefficients from Equation (6) for each outcome variable as listed. Positive values on the x-axis represent being assigned to a slower plan checker. Standard errors are clustered at the plan checker and week level. The sample is restricted to permits not filed in the 4 weeks leading up to a building code change date. [Back to text]

Table A1: Estimates of selection bias

	(1) Pred. Y	(2) Pred. Y of past project	(3) Pred. Y using excluded chars.	(4) Actual Y	(5) Actual Y — X
Plan checker effect	-0.001	-0.001	0.001*	0.940***	0.938***
	(0.011)	(0.022)	(0.000)	(0.034)	(0.035)
Project covariates					X
N	53744	53214	52030	52030	52030
Coef. diff.					0.0021

NOTES: This table reports coefficients from "forecast bias" regressions, as described by Chetty et al. (2014). The dependent variable in column 1 is log issuance time predicted using the covariates X_i described in Equation (1). The dependent variable in column 2 is the predicted log issuance time of the past project assigned to a plan checker. The dependent variable in column 3 is log issuance time predicted using excluded covariates from X_i (specifically, square footage, number of units, and imputed height). The dependent variable in column 4 is log issuance time (residualized using time and administrative fixed effects). The dependent variable in column 5 is residualized log issuance time, including covariates X_i. Standard errors are clustered at the plan checker and week level. [Back to text]

Table A2: Sensitivity of plan checker estimates to control set

	Panel A: Sensitivity to baseline controls					
	Means (1)	Year-month FE (2)	Year-month × office FE (3)	Add admin. effects (4)	Add covariates (5)	Baseline (6)
Plan checker FE: corr. with baseline	0.672	0.813	0.858	0.912	0.940	1.000
Applicant FE: corr. with baseline	0.767	0.835	0.848	0.899	0.965	1.000
Controls						
Plan checker FE	Y	Y	Y	Y	Y	Y
Applicant FE	Y	Y	Y	Y	Y	Y
Year-month FE		Y	Y	Y	Y	Y
Year-month × office FE			Y	Y	Y	Y
Admin. FE				Y	Y	Y
Covariates					Y	Y
Clearances						Y
N	53,744	53,744	53,744	53,744	53,744	53,744
R^2	0.373	0.438	0.475	0.514	0.547	0.568
Adj. R^2	0.343	0.408	0.434	0.476	0.511	0.533
p-value: plan checker FEs = 0				0.000	0.000	0.000

	Panel B: Sensitivity to additional controls and sample					
	Workload (7)	Years of experience (8)	Peers (9)	Excl. rush periods (10)	Drop first 6 months (11)	Include agency fixed effects (12)
Plan checker FE: corr. with baseline	0.984	0.996	0.973	0.968	0.985	0.924
Applicant FE: corr. with baseline	0.994	0.990	0.993	0.982	0.988	0.948
Baseline controls	Y	Y	Y	Y	Y	Y
Additional controls / sample						
Workload (active permits)	Y					
Years of experience		Y				
Peer controls			Y			
Excl. filing-rush periods				Y		
Drop first 6 months of tenure					Y	
Include approval agencies						Y
N	53,555	51,745	53,744	49,731	51,657	45,786
R^2	0.567	0.566	0.568	0.570	0.574	0.619
Adj. R^2	0.532	0.530	0.533	0.533	0.538	0.581

NOTES: The first row displays the correlation coefficients of the plan checker fixed effect parameters $\alpha_{j(i,t)}$ from Equation (1), where each column uses a different set of controls to estimate the plan checker fixed effects. The set of controls is indicated by the "Y" marks and column titles. The reported p-value tests the null hypothesis that plan checker effects are jointly zero. "Peers" refers to including a vector of peer plan checker dummy variables equal to one when the peer plan checker received assignments in a given week and office. "Excl. rush periods" refers to excluding permits submitted during busy periods, as defined in Section 6.2. "Drop first 6 months" refers to dropping permits assigned to a plan checker in their first 6 months after hire/entry. "Include agency fixed effects" refers to including indicator variables equal to one if a permit required at least one clearance from a city or county agency. [Back to text]

Table A3: Balance table around code changes for fast and slow plan checkers

	Non-rushed permits (not treated)		Rushed permits (treated)		Diff. in means
	(1)	(2)	(3)	(4)	(5)
	Below median	Above median	Below median	Above median	Pre-trends test
	Mean/(SD)	Mean/(SD)	Mean/(SD)	Mean/(SD)	Coef./SE
Valuation (2023 $)	2719056.77	1817011.03	1505673.85	1450354.02	902045.74
	(12306759.02)	(10047366.55)	(5254651.88)	(5370462.67)	(1111688.93)
Project sq. ft.	12038.06	7459.60	6676.20	6379.55	4578.46
	(53330.43)	(28265.23)	(19095.55)	(18456.08)	(4421.22)
No. of units	9.54	5.46	4.72	4.52	4.08
	(50.75)	(25.04)	(19.18)	(17.57)	(4.17)
Height (ft.)	29.24	28.44	31.16	29.95	0.79
	(22.18)	(12.05)	(12.46)	(10.16)	(2.04)
Elevation	210.89	145.77	227.23	193.66	65.12
	(147.65)	(117.36)	(156.51)	(127.62)	(38.97)
Slope	6.53	2.95	7.52	3.60	3.58^{+}
	(9.20)	(5.85)	(10.77)	(6.29)	(1.96)
In hillside zone	0.30	0.16	0.35	0.22	0.14^{+}
	(0.46)	(0.36)	(0.48)	(0.41)	(0.08)
In specific plan zoning overlay	0.43	0.39	0.45	0.32	0.04
	(0.50)	(0.49)	(0.50)	(0.47)	(0.08)
Observations	513	435	1455	1398	948

NOTES: This table tabulates permit characteristics of permits submitted in the eight weeks prior to a new code adoption date, separately for fast and slow plan checkers (as defined by the median $\hat{\alpha}_{j(i,t)}$ from estimating Equation (1) while excluding the "rush period" permits). The first four columns report the means and standard deviations of each group, and the last column reports coefficients and standard errors from regressions of each covariate on an indicator variable for a below-median plan checker for permits submitted in the non-rush period. [Back to text]

Table A4: Alternative DID estimates around code changes

	Log(issuance time)	Log(initial review time)	Log(queue time)	Log(clearances+corrections time)	Additional corrections
Above median PC	-0.159*	-0.412**	-0.109	-0.029	-0.293
	(0.076)	(0.135)	(0.096)	(0.065)	(0.271)
Rushed	0.145+	-0.277+	0.530**	0.237**	0.058
	(0.084)	(0.161)	(0.133)	(0.075)	(0.164)
Above median PC	0.002	0.249	0.113	-0.171*	-0.059
× Rushed	(0.091)	(0.159)	(0.132)	(0.076)	(0.236)
N	3346	3320	3329	3316	3346

NOTES: The table reports coefficients from estimating the following equation: $Y_{it} = \beta_1 (\text{Rush}_{it} \times \text{G}_j) + \beta_2 \text{Rush}_{it} + \beta_3 \text{G}_j + \mathbf{X}_i' \gamma + \phi_k + \tau_t + \nu + \varepsilon_i$ in which G_j is an indicator variable equal to one for permits assigned to an above-median plan checker and the remaining variables are the same as those from Equation (3). The fixed effects include indicator variables for year-office, project type, plan check type, and project applicant. The additional covariates include categorical variables for a project's geographical location (council district, community plan area, area planning commission, hillside zones, and special plan areas), clearances, height, project valuation, land use density, elevation, and slope. Standard errors are clustered at the plan checker and week level. [Back to text]

Table A5: Can observables predict future plan checker performance?

	(1) Next period plan checker performance	(2) Next period plan checker performance	(3) Next period plan checker performance
Valuation	-0.003	-0.006	-0.006
	(0.005)	(0.005)	(0.005)
Sq. Ft.	-0.005	-0.004	-0.004
	(0.007)	(0.007)	(0.007)
No. of units	0.007	0.008	0.008
	(0.006)	(0.006)	(0.006)
Height	-0.000	-0.000	-0.000
	(0.003)	(0.003)	(0.003)
Avg. Past Projects	-0.003	-0.001	-0.001
	(0.006)	(0.006)	(0.006)
Elevation	0.006	0.008	0.008
	(0.010)	(0.009)	(0.009)
Slope	0.007	0.009	0.009
	(0.008)	(0.008)	(0.008)
In Hillside Zone	-0.014	-0.019	-0.019
	(0.021)	(0.021)	(0.021)
In Specific Plan Zoning Overlay	-0.001	-0.004	-0.004
	(0.012)	(0.012)	(0.012)
N	46832	46832	46832
F-stat	0.38	0.58	0.58
p-value	0.95	0.81	0.82
Current PC perf.		X	X
Current & future PC count			X

NOTES: The table reports coefficients from regressions of the next period's average plan checker performance on permit characteristics and other covariates. The dependent variable for all regressions is the average plan checker performance measure of all plan checkers receiving assignments in the following week, where plan checker performance is measured using split sample estimates of $\hat{\alpha}_j$ as described in Appendix G. The fixed effects include indicator variables for year-office, project type, plan check type, project applicant, and split sample indicator. Standard errors are clustered at the week level. [Back to text]

Table A6: Plan check outcomes, permit characteristics, and construction outcomes for fast and slow plan checkers

	(1) Slower plan checkers Mean/(SD)	(2) Faster plan checkers Mean/(SD)	(3) Raw diff. in means Coef./SE	(4) Adj. diff. in means Coef./SE
Plan check outcomes				
Total issuance time	321.22	224.45	-96.77**	-72.08**
	(274.47)	(251.99)	(2.47)	(7.84)
Queue time	20.75	17.12	-3.63**	-2.30**
	(17.79)	(30.31)	(0.24)	(0.58)
Initial review time	10.18	9.90	-0.29	-2.06**
	(15.36)	(23.65)	(0.19)	(0.46)
Clearances+corrections time	270.25	227.48	-42.77**	-66.28**
	(258.72)	(263.10)	(2.93)	(7.76)
Additional correction rounds	2.20	1.05	-1.16**	-0.74**
	(3.01)	(2.14)	(0.02)	(0.18)
Mean days between corr. rounds	99.84	95.07	-4.77**	-7.10+
	(112.94)	(129.83)	(1.61)	(4.19)
Total clearances	11.96	10.50	-1.46**	-0.74*
	(5.14)	(5.33)	(0.05)	(0.31)
No. of agencies	4.66	4.07	-0.59**	-0.10+
	(1.55)	(1.72)	(0.02)	(0.06)
Permit characteristics				
Project valuation (2023 $)	1208906.88	1987962.93	779056.05**	282648.46
	(6865714.52)	(11078328.98)	(86503.26)	(248088.91)
Project sq. ft.	5271.83	8286.78	3014.95**	1033.20
	(24382.30)	(34545.38)	(283.78)	(737.00)
No. of units	3.31	5.49	2.18**	0.93
	(19.70)	(28.55)	(0.23)	(0.62)
Project height (feet)	29.04	32.70	3.66**	0.64
	(12.96)	(34.18)	(0.24)	(0.42)
Elevation	182.97	191.29	8.32**	-7.81
	(136.99)	(140.96)	(1.30)	(6.20)
Slope	4.65	5.24	0.59**	-0.25
	(7.99)	(8.12)	(0.08)	(0.37)
In hillside	0.25	0.29	0.04**	-0.00
	(0.43)	(0.45)	(0.00)	(0.02)
Construction outcomes				
No. of inspection corrections	3.75	3.55	-0.20**	-0.04
	(4.87)	(4.62)	(0.05)	(0.16)
Construction time	655.29	744.62	89.34**	8.50
	(558.02)	(757.72)	(6.56)	(19.32)
Has code enforcement case	0.11	0.18	0.08**	0.03**
	(0.31)	(0.38)	(0.00)	(0.01)
Has code violation	0.02	0.04	0.02**	0.01*
	(0.15)	(0.20)	(0.00)	(0.00)
Has safety-related violation	0.00	0.00	0.00*	-0.00
	(0.04)	(0.05)	(0.00)	(0.00)
Complete within 3 yrs of issuance	0.80	0.72	-0.08**	-0.01
	(0.40)	(0.45)	(0.00)	(0.01)
Observations	22624	22904	45528	45431

NOTES: This table tabulates plan check outcomes, permit characteristics, and construction outcomes for fast and slow plan checkers (as defined by the median $\hat{\alpha}_{j(i,t)}$). The second two columns report the average difference in means between the two groups, and the last column reports coefficients and standard errors from regressions that include the baseline control variables described in the main text. [Back to text]

A Data Construction

The main dataset is the set of all building permits submitted to LADBS, which I then merge with additional information. I merge it with the web scraped permit data using the unique permit ID. The web scraped data contains information on the names of staff verifying each update, dates of various updates, clearances assigned, project applicants, and permit inspections.

To identify the assigned plan checker for each permit, I search for the name associated with the permit status update "Assigned to PC Engineer" or "Assigned to Plan Check Engineer" (as shown in Figure A1). Then I use a combination of the texts of the status updates and the dates to identify how long various permit steps take via elapsed days. I also identify the project completion status based on whether the permit was ever finaled or issued a certificate of occupancy. This yields a dataset that has one observation per permit ID.

For the clearance information, I use a combination of string matching and staff names to map each assigned clearance to a city or county agency. The city of Los Angeles has a "Clearance Handbook" that details the set of all possible clearances and their respective agency responsible for issuing the clearance. I first use string matching on the clearance item text to assign clearances to an agency; then for the remaining non-assigned clearances, I use the mapping between staff names and already-assigned clearance items of the assigned group to fill in the agencies of the non-assigned clearances. I finally collapse this to the permit ID level by aggregating how many clearances are assigned per permit ID, separately for each agency.

For the project applicant information, each permit will list all of the recorded engineers, architects, contractors, geologists, and other professionals associated with a project. For each of these professions, the name of the individual and their license number are listed, which I use to merge with the licensing information publicly available from California's Department of Consumer Affairs (for engineers, architects, and other professional jobs) and the Contractors State License Board (for contractors). I assign the project applicant to be the first non-missing individual across this list: engineer, architect, and contractor. For all 3 of these web scraped datasets, the resulting cleaned dataset is unique at the permit ID level for merging to the main dataset of LADBS building permits.

For the permit inspection information, each permit will have a table listing all of the inspection updates. Some status updates are "Corrections issued," indicating the inspector found some aspect that was not built up to code. I count the number of these instances among the table of inspection updates and I designate each unique date with an inspection update as a distinct inspection trip.

The publicly available LADBS code enforcement data contains a case number that I use to web scrape additional information about each code enforcement case and subsequent code violations. Specifically, the web scraped information contains additional dates and the text of the code violation, which I use to classify the violation as being safety-related or not. I merge this data to the building permit data using the parcel identification number (PIN).

To augment the existing permit characteristics in the building permit data, I use the geographical coordinates (or if coordinates are missing, I use the address or PINs to identify the coordinates) to spatially merge the permit data with publicly available datasets

on land-use density, hillside ordinance areas, hillside zones, historical preservation overlay zones, and special plan areas. Controlling for whether a project is in one of these areas is important, as often they are subject to greater regulatory stringency. I also use the USGS' Digital Elevation Model to identify the elevation and slope of each permit's parcel centroid. I also merge in ATTOM property tax deeds to identify the last parcel transaction before the permit was filed (i.e. sales price of the empty/undeveloped lot), yet this only yields matches for approximately 17% of my sample.

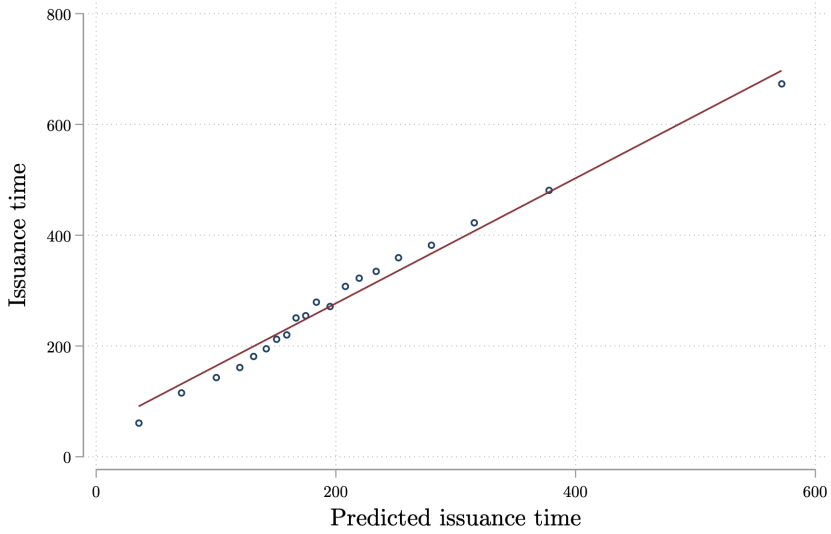
For my main analyses, I restrict these permits to those that are: new residential, not supplemental permits (often submitted due to changes to project design), not express permits (these are issued same day as submittal), submitted by an applicant with at least 2 permits, assigned to a plan checker with at least 2 permits, and in the largest connected set.[24]

B An Alternative Model of Permit Issuance Time

One might be worried that there are unobserved confounders related to project complexity that could bias the estimates from Equation (1). While my preferred approach is to use an exhaustive list of covariates to control for the fact that more complex projects (as measured by clearances, project value, etc) are likely not to be assigned to the most junior plan checker, I also use an alternative model that includes a covariate for expected project complexity as proxied by length of issuance time. I train a simple machine learning model following the approach in Iaria et al. (2024), where I use the text in the work description field of each new residential permit to predict the log issuance time. I follow their steps for processing the text (i.e. filtering out non-alphanumerical characters, removing common words, and stemming the words), then I extract all unigrams and bigrams from the work description to obtain a matrix of permit-by-1,2 grams that is reweighted by the relative importance of n-grams based on frequency (term-frequency inverse-document frequency reweighting). I then train a L1-regularized regression model (lasso regression) using this matrix as an input, along with indicators for the number of words in the work description. I choose the optimal normalization strength λ using 10-fold cross-validation. The algorithm predicts log issuance time for each permit application, and I plot a binscatter of the predicted values against the true values below. I use the predicted value as a control in Equation (1), and I find that the estimates (Table A7) using the KSS estimator are slightly smaller but similar to those presented in the main text. The plan checker fixed effects explain 3-5% of the residual variation in this specification, while applicants explain 14-22% of the residual variation.

[24]While the modal applicant submits once, there is a long right tail for the distribution of number of permits submitted per applicant, such that these restrictions do not exclude a significant share of all new residential permits.

Figure A14: Actual and Predicted Issuance Time



NOTES: This figure plots a binscatter of the actual permit issuance time against the predicted issuance time.

Table A7: Variance Decomposition with Predicted Issuance Time as a Covariate

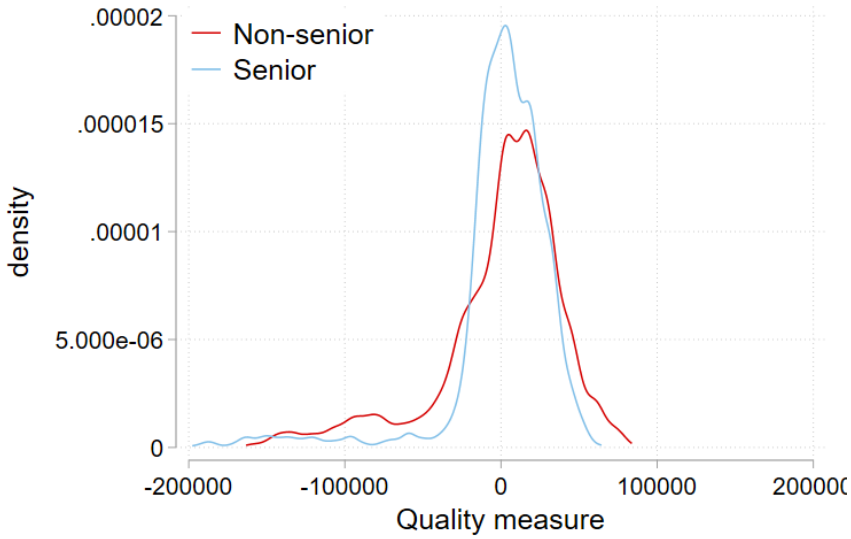
	Plug-in		Shrinkage		KSS estimator	
	Std. Dev.	Variance	Std. Dev.	Variance	Std. Dev.	Variance
Plan Checker	0.173	0.030	0.158	0.025	0.131	0.017
	(0.030)					
Applicant	0.367	0.135	0.353	0.125	0.297	0.088
	(0.009)					
Cov(Plan Checker, Applicant)		0.002				0.004
Log issuance time	1.036	1.073				
Log issuance time| covariates	0.784	0.615				
Number of plan checkers	400					
Number of applicants	2041					
Observations (largest connected set)	53744					
Observations (leave-out connected set)	51453					
Mean no. of applicants/plan checker	52					
Mean issuance time	296					

NOTES: This table reports the variance decomposition from Equation (2). "Plug-in" refers to the baseline plug-in fixed effects. Bootstrap standard errors are reported in parentheses. Shrinkage refers to the the empirical Bayes adjusted estimates. KSS estimator refers to the heteroskedasticity-robust estimation of variance components on the leave-out connected set.

C An Alternative Measure of Bureaucrat Performance

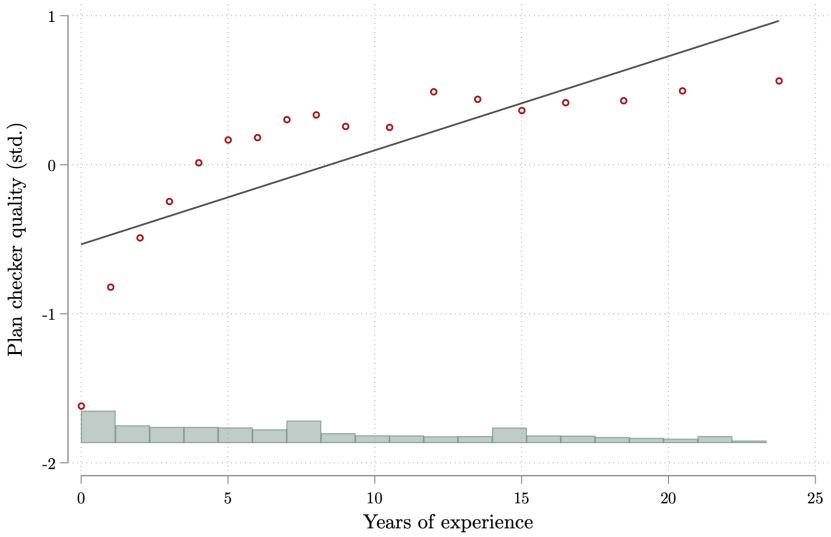
Instead of using the estimated plan checker fixed effects from Equation (1) as my main measure of performance, I also create an alternative measure from plan checkers' salary data. I use the salary data to identify LADBS staff based on name and job titles, then I construct a "quality" proxy similar to Liscow et al. (2024) that uses the residuals from regressions of salary on year dummies, job title dummies, an indicator if there is missing benefits information, and whether the job is a project manager role. The density of this residual quality measure is shown in Figure A15. This quality measure likely reflects a combination of experience and skills, as shown in Figure A16, which correlates the quality measure with years of experience (derived from the permit data). Nonetheless, this is not my preferred performance measure, since pay scales are set by union bargaining units, unlike in the private sector. Notably, this salary residual is not correlated with the estimated plan checker performance measure ($\alpha_{j(i,t)}$).

Figure A15: Distribution of LADBS staff "quality" by seniority



NOTES: This figure plots the density of LADBS staff quality, separately by their seniority (based on job title keywords).

Figure A16: Binscatter of LADBS plan checker "quality" and years of experience



NOTES: This figure plots a binscatter of LADBS plan checkers' quality measure against their years of experience, which is calculated as the current year minus the first year in which a plan checker shows up in the permit data.

D Plan Checker Performance and Overtime

This section explores empirical correlations between the estimated plan checker performance measure ($\hat{\alpha}_{j(i,t)}$) and overtime to expand upon the relationships shown in Figure 7. Figure 7 indicates that plan checkers who work more overtime are also higher performing. The main dataset used in this section is personnel data merged with the estimated plan checker performance measure, $\hat{\alpha}_{j(i,t)}$. To provide more context for the role of overtime, I plot the average annual overtime pay, regular base pay, and total pay for managers and non-managers in Figure A17. This figure highlights how overtime pay is the main way to bridge the regular pay gap between plan checkers who become managers and those that do not. Recall that managers are not allowed to work overtime, which explains the gap between the non-manager and manager series in panel (a) of Figure A17.

To further understand who chooses to work overtime, I regress overtime hours on various plan checker characteristics. Table A8 reiterates previous key relationships. Slower plan checkers work less overtime on average, while those who eventually get promoted to manager roles tend to work less overtime. The positive coefficient on the regular base pay indicates plan checkers do respond to higher pay by working more. This can also be seen visually in Figure A18, where there are jumps in average overtime hours for non-managers as plan checkers move up the job ladder. The interaction term between $\hat{\alpha}_{j(i,t)}$ and manager highlights that low-performing plan checkers who do eventually get promoted to manager are more likely to work overtime, yet the low-performing plan checkers who reach senior plan checker roles (but not managerial roles) tend to work less overtime. This may suggest that overtime can also be a way to make up for being a low performer for those who aim to reach the managerial roles. Overall, this table is providing suggestive evidence that plan checkers' overtime choices seem to respond to incentives related to promotions and pay.

Figure A17: Overtime, regular, and total pay: managers vs non-managers



(a) Overtime pay



(b) Regular base pay



(c) Total pay (regular+overtime)

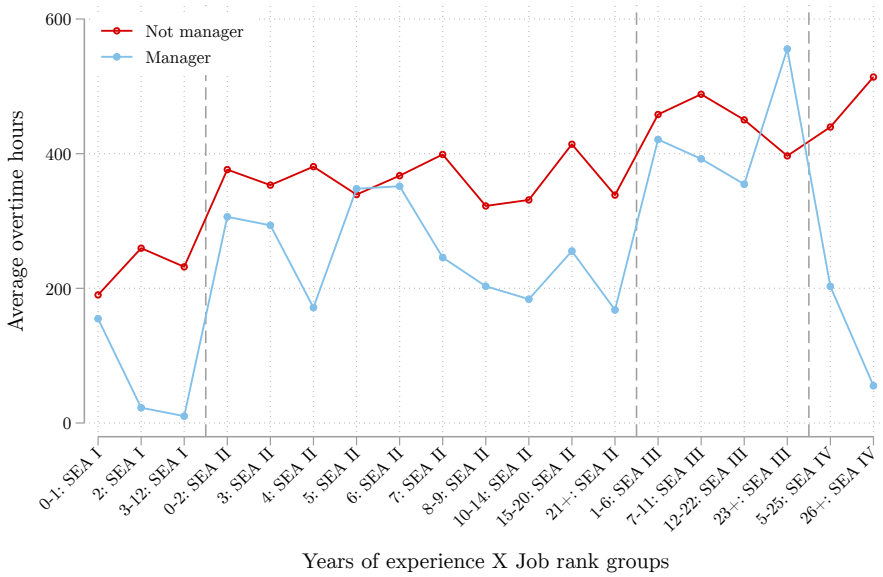
NOTES: This figure plots plan checkers' average overtime, regular, and total pay received by years of experience, separately for those who are ever promoted to managers and those who do not. The dashed vertical line represents the average years of experience of plan checkers when they are promoted to manager.

Table A8: Correlates of overtime hours

	Plan checkers: all		Plan checkers: entry-level start	
	(1)	(2)	(3)	(4)
	OT hours	OT hours	OT hours	OT hours
Ever promoted manager	-3.894^{+}	-3.614	-6.609**	-6.697**
	(2.210)	(2.230)	(2.521)	(2.549)
Ever promoted to senior PC	1.768	1.748	3.593	3.689^{+}
	(1.964)	(1.868)	(2.271)	(2.074)
Plan checker FE (std.)	-1.977*	-1.819^{+}	-1.520	-1.478
	(0.958)	(0.960)	(0.961)	(0.950)
Ever promoted manager=1 × Plan checker FE (std.)	8.553**	7.687*	5.185	5.307
	(3.222)	(3.087)	(3.553)	(3.547)
Ever promoted to senior PC=1 × Plan checker FE (std.)	-5.085*	-5.139*	-7.865*	-7.945*
	(2.437)	(2.388)	(3.245)	(3.269)
Hourly base pay	0.340**	0.220**	0.218**	0.147**
	(0.069)	(0.065)	(0.062)	(0.055)
Years of experience	-0.040	-0.056	0.449	0.322
	(0.098)	(0.105)	(0.281)	(0.310)
N	55305	55305	36385	36385
Year+Month FE	yes	yes	yes	yes
Job title FE	no	yes	no	yes

NOTES: This table displays a regression of biweekly (pay-period) overtime hours on various plan checker characteristics. The first two columns include all plan checkers, while the last two columns are restricted to plan checkers who are observed to start at the entry-level plan checker job title in my sample. Recall that positive values of the plan checker fixed effects correspond to lower-performing bureaucrats. Standard errors are clustered at the plan checker level. [Back to text]

Figure A18: Overtime hours by years of experience and job title



NOTES: This figure plots plan checkers' average overtime by binned years of experience grouped by the 4 plan checker job titles, separately for those who are eventually promoted to managers and those who are not.

E Conceptual Framework

This section presents a simple framework to motivate and inform interpretation of the empirical results. In this setting, plan checkers differ in exogenous skill and choose how many overtime hours to work. The framework clarifies when the cross-sectional correlation between skill and overtime is positive (skill and overtime are complements) or negative (substitutes).

E.1 Setup

In a static setting, suppose all permit plans are homogeneous and uniform (i.e. only require one unit of time for review). Plan checker j has exogenous productivity/skill $\theta_j > 0$ that equals the number of plan reviews completed per unit of time (hour). Regular hours $\bar{H}$ are fixed by the civil service contract. In each pay period, the plan checker chooses overtime hours $h \geq 0$, so total hours are $H = \bar{H} + h$ and output is $y = \theta H$. Compensation is $w\bar{H} + \omega w h$, where $\omega = 1.5$ is the overtime premium and wages w are set by city-wage pay scales. Pay does not depend on output, so skill affects the return to overtime only through its non-pecuniary value, reflected in positive benefit weight $b > 0$ and the non-wage value of output $g(y)$. This could reflect other benefits from completing plan checks, such as intrinsic motivation, career concerns, or the desire to keep the queue moving. There is a cost to working additional overtime hours captured by the function $c(h)$. Plan checker utility is:

$$U(h;\theta) = \max_h w\bar{H} + \omega w h + b\, g\big(\theta(\bar{H} + h)\big) - c(h), \tag{9}$$

Assumption 1 *(i) g is twice differentiable with $g' > 0$ and $g'' \leq 0$. (ii) c is twice differentiable with $c(0) = 0$, $c' \geq 0$ and $c'' > 0$. (iii) Preference parameters (b, c) are distributed independently of θ across checkers.*

Define the elasticity of the marginal value of output,

$$\rho(y) \equiv -\frac{y\, g''(y)}{g'(y)} \geq 0. \tag{10}$$

Under Assumption 1, $U_{hh} = b\theta^2 g''(y) - c''(h) < 0$, so the optimum is unique. An interior optimum $h^* > 0$ satisfies this first-order condition:

$$\omega w + b\,\theta\, g'\big(\theta(\bar{H} + h^*)\big) = c'(h^*). \tag{11}$$

A plan checker will choose their optimal overtime hours such that the marginal benefit of an additional overtime hour (from both additional overtime pay and the value of completing an additional plan review, which depends on skill) equals the marginal cost.

E.2 Skill and Overtime: Complements or Substitutes

Proposition 1 *At an interior optimum, with $y^* = \theta(\bar{H} + h^*)$,*

$$\frac{dh^*}{d\theta} = \frac{b\, g'(y^*)\left[1 - \rho(y^*)\right]}{c''(h^*) - b\theta^2 g''(y^*)}, \qquad \textit{so} \qquad \operatorname{sign}\frac{dh^*}{d\theta} = \operatorname{sign}\big(1 - \rho(y^*)\big). \tag{12}$$

In particular, if $g(y) = y^{1-\rho}/(1-\rho)$ with constant ρ (and $g = \log y$ for $\rho = 1$), then skill and overtime are complements if $\rho < 1$, independent if $\rho = 1$, and substitutes if $\rho > 1$.

By the implicit function theorem applied to (11), $dh^*/d\theta = -U_{h\theta}/U_{hh}$. Differentiating, $U_{h\theta} = b\big[g'(y) + \theta(\bar{H} + h)g''(y)\big] = b\,g'(y)\big[1 - \rho(y)\big]$, and $-U_{hh} = c''(h) - b\theta^2 g''(y) > 0$. The comparative statics highlight the relationship between skill/productivity and overtime hours. A more skilled plan checker reviews more plans per overtime hour, which raises the return to overtime. But a more skilled checker also reaches a satisfactory level of output within regular hours, which lowers the marginal value of further output. The overtime wage ωw does not vary with θ, so it shifts the level of overtime but not the sign of its relationship with skill.

E.3 Two examples

Example 1 (Linear output: complements) *Let $g(y) = y$ and $c(h) = \kappa h^2/2$. Then $\rho = 0$ and*

$$h^* = \frac{\omega w + b\theta}{\kappa},$$

which is increasing in θ, so skill and overtime are complements in this example.

Example 2 (Workload target: substitutes) *Let the plan checker value clearing an assigned workload N, with $g(y) = -\frac{1}{2}(N - y)^2$ for $y \leq N$, and output beyond this target has no value. Let $c(h) = \kappa h^2/2$. Then*

$$h^* = \frac{\omega w + b\theta(N - \theta\bar{H})}{\kappa + b\theta^2}, \qquad \rho(y) = \frac{y}{N - y},$$

so skill and overtime are substitutes whenever $y^* > N/2$, that is, when the plan checker is not too far behind on their workload target.

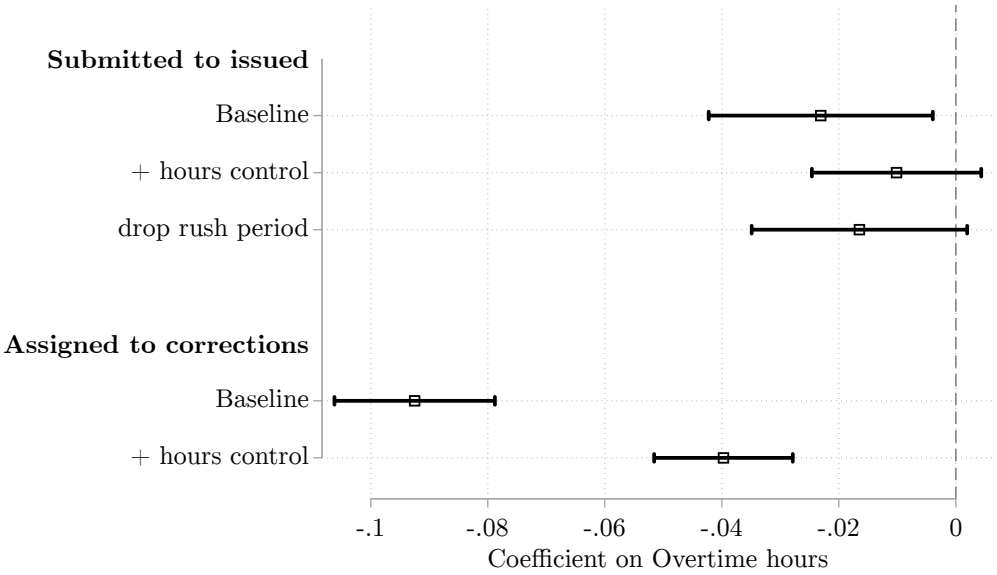
E.4 Empirical Implications

Let s_i denote checker i's log speed. If issuance time is inversely proportional to $\theta_i H_i$, then $s_i = \log\theta_i + \log H_i$ up to office-period terms, where H_i is checker-average total hours. Let $s_i^{\text{raw}} = -\alpha_i$ be the checker effect from Equation (1) without an hours control, and s_i^{net} the effect with a control for log hours, so that $s_i^{\text{net}} = \log\theta_i$. Assume constant ρ, so that each component of H_i is monotone in θ_i in the direction given by $\text{sign}(1 - \rho)$.

Prediction 1 (Net-of-hours speed and overtime) $\text{sign}\,\text{Cov}(s_i^{net}, \log H_i) = \text{sign}(1 - \rho)$.

The plan checker effect will be positively correlated with overtime if skill and overtime are complementary. This is what is shown in Figure A19, which correlates overtime hours with various estimates of plan checker speed that adjust for hours worked. Across various different ways of estimating $\alpha_{j(i,t)}$, all of the coefficients between overtime and plan checker performance are negative, highlighting that higher-performing (i.e. faster) plan checkers tend to work more overtime. This evidence points to overtime and skill as complementary in this setting.

Figure A19: Relationship between plan checker performance and overtime hours

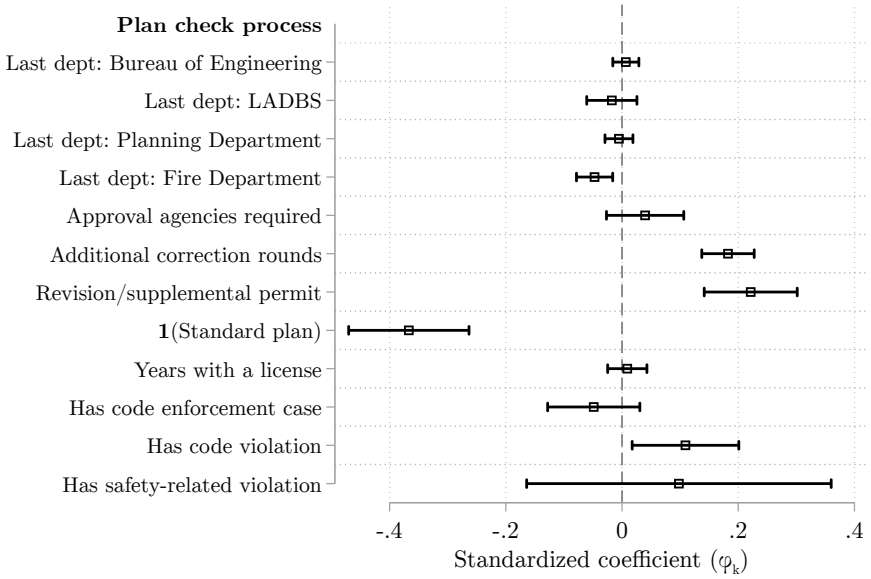


NOTES: The figure plots the coefficients of bivariate regressions between the estimated plan checker fixed effect ($\alpha_{j(i,t)}$) and annual overtime hours. The y-axis labels indicate the sample or additional controls, while the headings indicate which time outcome variable is used to estimate $\alpha_{j(i,t)}$. $\alpha_{j(i,t)}$ and overtime hours are standardized to have mean zero and unit variance. The regressions are permit-weighted. For the permit-level regressions, each covariate is first residualized on the control vector (administrative and time fixed effects from Equation (1)). Confidence intervals are computed using the heteroskedasticity-robust standard errors proposed by Kline et al. (2020).

F Project Applicant Performance

Given that project applicants do explain a non-trivial share of the variation in permit issuance times, I also examine what predicts applicant performance. I correlate the shrinkage estimates of ϕ_k with project and applicant attributes. For the applicants, the takeaway from Figure A20 is that faster applicants propose simpler projects, which will require less time. This is also supported by the correlation showing that slower applicants have more revisions to their permits and a large negative correlation on the dummy variable for whether the project description mentions "Standard Plan" (which are generally pre-approved plans for a section of a home or for ADUs).[25] Interestingly, projects submitted by slower applicants are positively correlated with receiving a future code violation, suggesting that quality impacts are more likely to be driven by applicants than the plan checkers.

Figure A20: Correlates of project applicant performance



NOTES: The figure plots the coefficients of bivariate regressions between the estimated applicant fixed effect (ϕ_k) and the listed covariate. ϕ_k and the covariates are standardized to have mean zero and unit variance. Except for the regressions using applicant covariates, all other regressions include the vector of time and administrative fixed effects. Standard errors are clustered at the plan checker level.

[25] Standard plans can refer to various features of a home that follow a pre-approved plan, such as stairs, pools, etc. In recent years, LA has piloted having pre-approved plans for ADUs.

G A Split Sample Approach to Identifying Quality Impacts

As a robustness check, I describe another approach for estimating how plan checker performance affects future project outcomes using a split-sample instrumental variables design. The split-sample IV design addresses concerns that there might be permit-level shocks across adjacent weeks that could lead to noisy estimates of α_j. First, I form two partitions of my data (denoted A and B) by randomly sampling weeks associated with each plan checker-applicant combination. I estimate Equation (1) on both partitions, resulting in two estimates of plan checker performance $\hat{\alpha}_{j(i,t)}$ which I denote as $\hat{\alpha}^A_{j(i,t)}$ and $\hat{\alpha}^B_{j(i,t)}$. Then I use these estimates of the plan checker fixed effects from each partition as an out-of-sample estimate of plan checker performance for each permit i to instrument for own plan checker performance.

The two-stage least squares (2SLS) model is estimated in the following form:

$$\text{First Stage:} \quad \tilde{\alpha}_{j(i,t)} = \gamma Z_i + \mathbf{X}_i\pi + \phi_k + \tau_t + \boldsymbol{\nu} + u_{it} \tag{13}$$

$$\text{Second Stage:} \quad y_{it} = \rho\tilde{\alpha}_{j(i,t)} + \mathbf{X}_i\beta + \phi_k + \tau_t + \boldsymbol{\nu} + \eta_{it} \tag{14}$$

where y_{it} is the outcome of interest, $\tilde{\alpha}_{j(i,t)}$ takes the following form:

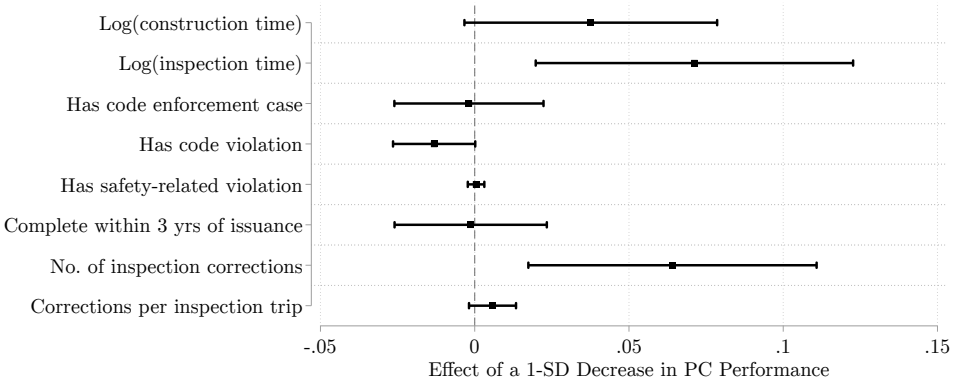
$$\tilde{\alpha}_{j(i,t)} = \begin{cases} \hat{\alpha}^A_{j(i,t)} \text{ if } i \in A \\ \hat{\alpha}^B_{j(i,t)} \text{ if } i \in B \end{cases} \tag{15}$$

and the instrument Z_i is defined as:

$$Z_i = \begin{cases} \hat{\alpha}^A_{j(i,t)} \text{ if } i \in B \\ \hat{\alpha}^B_{j(i,t)} \text{ if } i \in A \end{cases} \tag{16}$$

Estimates from the second stage model are presented below. Again, I do not find strong evidence of negative quality impacts from being assigned a faster plan checker on average, as measured by future probability of a code enforcement case or code violation. However, I do find that being assigned a slower plan checker leads to a statistically significant increase in inspection time and inspection corrections.

Figure A21: Split sample instrumental variables estimates



NOTES: This figure plots the second-stage coefficients from Equation (14) for each outcome variable as listed. Positive values on the x-axis represent being assigned to a slower plan checker. Standard errors are clustered at the plan checker and week level.

H Calculating Net Costs of Avoided Delay Time

This section details the calculations reported in Section 8. For each of the three counterfactual scenarios, I first compute or simulate the number of days saved from implementing the described counterfactual personnel policy. For the two counterfactual scenarios that are improving the slowest plan checkers and increasing plan checker pay, I use the following equation to calculate avoided delay time if one were to replace the observed permit i's plan checker performance $\hat{\alpha}_{j(i)}$ with $\tilde{\alpha}_{j(i)}$:

$$\Delta d_i \;=\; D_i\Big[1 - \exp\left(\tilde{\alpha}_{j(i)} - \hat{\alpha}_{j(i)}\right)\Big]$$

Note that the previous estimate implicitly includes avoided delay time due to certain plan checkers choosing to work more overtime (since a lower $\hat{\alpha}_{j(i)}$, or speed, is correlated with more overtime hours). To separate out the effect of overtime, I also estimate days saved using a different estimate of plan checker performance that controls for hours worked, as described in Appendix E, so that this alternative measure of plan checker performance arguably captures only productivity/skill.

To convert these avoided delay days to a monetary value, I assume that all builders are homogeneous, such that they face the same borrowing and financing rates and make the same timing choices regarding their pre-construction loan[26]. I assume that a delayed project costs the builder an annual flow δ as a share of the project's construction value V_i reported in the permit:

$$\Delta\pi_i \;=\; \delta\, V_i \frac{\Delta d_i}{365}, \qquad \delta \;=\; \underbrace{\ell r}_{\text{financing}} + \underbrace{sh}_{\text{carrying}} + \underbrace{c}_{\text{foregone rent}} \;=\; 8.45\%$$

in which ℓ is a loan-to-value ratio of 0.7, r is an interest rate of 0.06, s is the share of project value attributable to land costs set to 0.25, h is the property tax rate on land equal 0.01, and c is the capitalization rate of 0.04.

The total savings to builders from avoided delay time is the sum of these terms $\Delta\pi_i$ over all affected projects $\mathcal{A}$:

$$S = \frac{\delta}{365} \sum_{i \in \mathcal{A}} V_i\, \Delta d_i$$

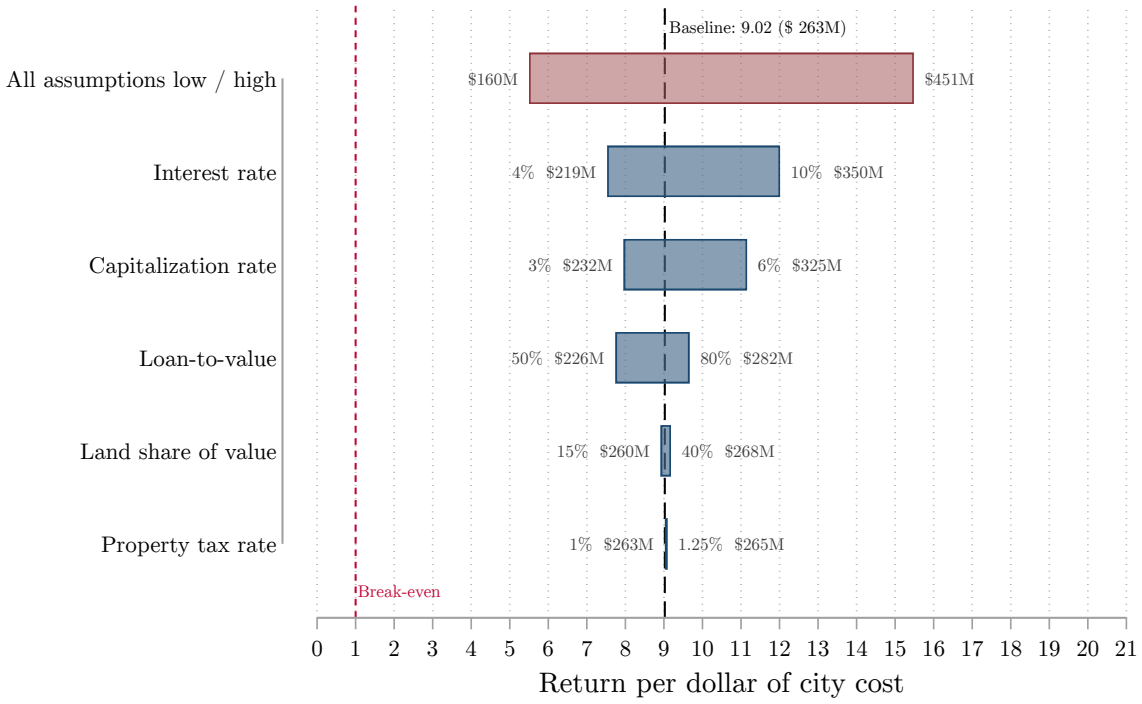
To compare how these private savings would compare to the city's costs of hiring additional personnel, I compute the salary costs using annual observed salaries w_{jt} across plan checker-years $\mathcal{R}$, while assuming there is a hiring friction ϕ of 0.25:

$$C = \Big(N w^{\text{rep}} - \sum_{(j,t) \in \mathcal{R}} w_{jt}\Big) + \phi\, N w^{\text{rep}}$$

[26]The calculations shown here could be under-estimating or over-estimating avoided delay costs, since hassle costs are omitted but the permit issuance is treated as the last hurdle before a pre-construction loan gets converted to a construction loan (which assumes there are no additional delays due to labor or material shortages, for example). Hence, I treat these estimates as illustrative of potential saved delay costs.

where the new replacement wage is w^{rep} for N distinct plan checker-years. The replacement wage varies depending on the counterfactual scenario. The replacement wage includes base pay, overtime pay, and benefit pay. Benefit pay is only observed in 2024-2025, so I use the average share of benefit pay relative to base pay in the observed years to calculate the amount of benefit pay in unobserved years (the share is 12%). Finally, the return from each dollar spent on additional plan checker capacity is computed as the ratio S/C. Figure A22 plots the ratio S/C under different assumptions of the interest rate, capitalization rate, loan-to-value ratio, land share, and property tax rate for the counterfactual scenario of replacing the lowest 25% performing plan checkers with the median.

Figure A22: Sensitivity of counterfactual estimates



Each bar varies one assumption over the range shown, holding the others at baseline (loan-to-value .7, interest 6%, cap rate 4%, land share 25%, property tax 1%).

NOTES: This figure plots the estimates of the ratio of private savings to city costs from the counterfactual scenario of replacing the lowest 25% performing plan checkers with the median plan checker. The dashed black line represent the main baseline estimates shown in Table 10, while the bars represent the range of estimates from varying each of the components labeled on the y-axis, while holding the others fixed at baseline.

I A Queueing Simulation of the Plan Check Process

This section describes the queueing model that I use to simulate how additional staffing affects permitting delays. This multi-server queueing model is adapted from a large operations and management literature (for example, on the theory of queues: Kingman (1962); Little (1961); Whitt (1993)).

I.1 Setup

The queueing environment simulates the processing of permit application i submitted to office o in calendar year t. Standard closed-form queueing approximations (such as Kingman's formula) assume single-server steady-state conditions, homogeneous service times, and stationary arrival rates. Given that there are time-varying trends in permit submissions and heterogeneity in plan checker performance, I treat this setting as a multi-server system with parallel "service stations" (i.e. plan checkers), and I employ a discrete-event simulation framework to model the queueing process.

Each permit application i is submitted on date A_i in office o in year y. Each office has "service stations" staffed by $c_{o,y}$ plan checkers that varies by office and year. Applications are assigned on a first-come-first-serve basis. A plan checker can be assigned up to s applications at once, so a given office o has $c_{o,y} \cdot s$ service capacity in year y; when all plan checkers' service slots are occupied, arriving applications wait in the queue. I read $c_{o,y}$ from the data as the monthly average of the number of distinct plan checkers with an active assignment in office o and year y.

Total permit processing duration (W_i) is decomposed into three distinct phases:

$$W_i = W_{q,i} + W_{r,i} + W_{c,i} \tag{17}$$

where $W_{q,i}$ is the queue time (the time between submittal date A_i and initial assignment to a plan checker), $W_{r,i}$ is the initial review time (the time during which the assigned plan checker conducts initial plan review), and $W_{c,i}$ is the clearances and corrections phase (the time between initial corrections being issued and plan check approval, during which the project may need approval from other agencies). I simulate the new residential permits in the analysis sample of Section 5.1 whose submittal, assignment, initial corrections, and issuance dates are observed and occur in that order. When the plan check approval date is missing or precedes initial corrections, I use the issuance date. The simulation takes the arrival dates $\{A_i\}$ and clearance times $\{W_{c,i}\}$ from the data, draws review times from an estimated model, and generates the queue times $W_{q,i}$.

As described in the main text, plan checkers have heterogeneous processing speeds (or service times). The initial review duration for permit i assigned to checker j is simulated as:

$$\log W_{r,i,j} = \hat{\mu}_i + \tilde{\alpha}^r_j + \varepsilon^*_i, \tag{18}$$

where $\tilde{\alpha}^r_j$ is plan checker j's initial review-time effect and $\hat{\mu}_i$ is the fitted value from the re-estimating Equation (1) but replacing issuance time with initial review time and excluding

the plan checker effect:

$$\log W_{r,i} = X_i\beta^r + \alpha^r_{j(i)} + \phi^r_k + \tau^r_t + \nu^r + \varepsilon^r_i. \tag{19}$$

and ε^*_i is drawn with replacement from the estimated residuals of permits submitted to the same office in the same year.[27]

Motivated by the empirical evidence in the main text, plan checkers remain partially occupied during the clearances/corrections phase due to inter-departmental coordination and additional revisions. The total amount of time that plan checker j is occupied by permit i is $T_{\text{busy},i,j}$:

$$T_{\text{busy},i,j} = W_{r,i,j} + \theta \cdot W_{c,i} \tag{20}$$

where $\theta \in [0, 1]$ parameterizes the administrative load factor imposed on plan checker bandwidth during the clearances phase. Attributing the entire clearances/corrections time period as busy for the plan checker will overstate congestion, but ignoring it is also unrealistic since plan checkers are involved in additional corrections rounds (which occurs for half of the permits in my sample).

The simulation processes permits sequentially based on submittal date A_i, and permits are assigned to the earliest available plan checker time slot, denoted S_{m^*}. Review begins when the assigned slot becomes available or when the permit arrives, whichever is later:

$$T_{0,i} = \max(A_i, S_{m^*})$$

The baseline queue delay is evaluated as $W^B_{q,i} = T_{0,i} - A_i$. Once a plan checker j' is assigned a permit, that permit occupies one of the s service slots for that plan checker j' for the entire length of $T_{\text{busy},i,j}$.

The model has two unknown parameters, $\psi = (s, \theta)$. The model leaves out several ways in which actual offices respond to congestion. For example, plan checkers work more overtime when workloads rise (Section 6.2). In the model, capacity within an office-year is fixed, so when more applications arrive than an office's slots can review, the backlog grows until arrivals slow down. This feature affects the model's fit, as discussed below.

I.2 Estimation

Moments. I estimate ψ via simulated method of moments, i.e. I choose the values of (s, θ) so that the simulated queue times look as similar as possible to the observed queue times along several dimensions. The targeted moments are four features of the observed queue times $W_{q,i}$:

1. the mean of $W_{q,i}$ among applications filed outside the rush periods (the four weeks before each building code change);

2. the rush effect: the mean of $W_{q,i}$ for applications filed in the four weeks before a

[27] When an office-year has fewer than 30 residuals, I draw from all residuals in the office, and if the office has fewer than 30, from the full sample. Permits without a residual, such as singletons dropped in estimating Equation 19, keep their observed review time, rescaled to the assigned plan checker: $W_{r,i,j} = W_{r,i} \exp(\tilde{\alpha}^r_j - \tilde{\alpha}^r_{j(i)})$.

building code change minus the mean for those filed in the four weeks before that, in code change years;[28]

3. the standard deviation of $W_{q,i}$ outside the rush periods; and

4. the 90th percentile of $W_{q,i}$ outside the rush periods.

Let $\hat{m}$ denote these moments in the data and $m(\psi)$ denote the same moments computed from simulated queue times. The simulation is replicated $R = 30$ times with independent draws of the residuals ε_i^* and plan checker effects $\tilde{\alpha}_j^r$, and $m(\psi)$ is the average of the moments across replications. I use the same 30 sets of draws at every parameter value and in both the baseline and counterfactual simulations (i.e. common random numbers). This ensures that differences in the simulated moments across parameter values are due to the parameters themselves rather than to different random draws, so that most of the simulation noise cancels out when calculating days saved between the baseline and counterfactual simulations. The gap between the simulated and data moments, $m(\psi) - \hat{m}$, reflects both sampling error in the data moments and simulation error in the simulated moments, so I inflate the covariance of this gap by $(1 + 1/R)$ to account for the simulation error. The estimator is:

$$\hat{\psi} = \arg\min_{\psi} \ \big(m(\psi) - \hat{m}\big)' \, \widehat{W} \, \big(m(\psi) - \hat{m}\big), \qquad \widehat{W} = \operatorname{diag}\big(\widehat{V}\big)^{-1}, \tag{21}$$

where $\widehat{V}$ is the covariance of $\hat{m}$, estimated by a bootstrap that resamples office-weeks of submittal ($B = 500$). Applications filed in the same office and week face the same queue, so resampling individual permits would understate sampling error. I use the diagonal weight matrix which weights each moment by the inverse of its own variance, so that the gap in each moment is measured in terms of standard errors.[29] Because the number of slots is an integer, the objective is a step function of s, so I minimize it over a grid of (s, θ) and refine the grid around the minimum.[30]

Identification. Both plan checker capacity s and the clearance load θ both determine how closely each office operates to capacity, and therefore, the level of queue times. The two parameters differ in how they shape the distribution of queue times. Under Equation 20, raising θ lengthens the amount of time each permit occupies a plan checker's slot in proportion to its own clearance time, which is long and highly dispersed. This makes slot occupancy more variable, and slots are tied up most in offices and years with long clearance times, which increased with total issuance time over the sample period. For a given mean, queue times then become more dispersed, both across permits and across offices and years. In contrast,

[28] Unlike the difference-in-differences estimates in Section 6.2, this moment does not difference out the same weeks of the previous year. Since the simulation uses the observed arrivals, seasonal patterns in submittals enter both the data and the simulated moment.

[29] I use the diagonal weight matrix rather than $\widehat{V}^{-1}$, since the four moments are all computed from the same queue times and are highly correlated, which makes $\widehat{V}^{-1}$ unstable and puts large weights on combinations of moments rather than on each moment itself

[30] The initial grid is $s \in \{0.5, 1, \ldots, 10\} \cup \{11, 12, \ldots, 30\}$ and $\theta \in \{0, 0.02, \ldots, 0.30\}$. I then add a finer 15×10 grid that spans the parameter values whose objective (using $\widehat{V}^{-1}$ as the weight matrix) is within the 95th percentile of the χ^2_2 distribution of its minimum, extended by one step of the initial grid in each direction. I repeat this refinement once if the minimum lies on the edge of the finer grid.

raising s expands plan checker capacity proportionally in every office and year and does not have this effect. The rush effect measures how close to capacity offices operate, since an office near capacity cannot absorb a surge in filings, so queue times would rise sharply during the rush. The two parameters still partially substitute for each other, as more slots can be offset by a higher clearance load.

Inference. With four moments and two parameters, the model is overidentified, which allows me to test if the model can match all four moments at once. I report Hansen's (1982) J statistic, $J = \min_\psi J(\psi)$ with $J(\psi) = (m(\psi) - \hat{m})' \widehat{V}^{-1}(m(\psi) - \hat{m})$, which is distributed χ^2_2 if the model is correctly specified.

I.3 Counterfactuals

To evaluate counterfactual staffing levels, I scale office-year capacity by δ: $\tilde{c}_{o,y} = c_{o,y} \cdot (1+\delta)$. The additional plan checkers are assigned the median fixed effect of active checkers in that office-year. Taking the empirical arrival sequence $\{A_i\}$ and review lengths $\{W_{r,i}, W_{c,i}\}$ as given, I re-simulate the permitting process under $\tilde{c}_{o,y}$ to generate counterfactual queue times $\tilde{W}_{q,i}$. The time saved for permit i under the counterfactual policy is calculated as:

$$\Delta W_{q,i} = W^B_{q,i} - \tilde{W}_{q,i}$$

To value the time saved, I apply the accounting framework of Appendix H to each permit's project value and days saved. The city's cost is the number of plan checker-years added, $\delta\, c_{o,y}$ summed over offices and years in 2005–2024, times an annual cost of \$176,450 per plan checker (median plan checker total pay).

At the estimated parameters, a 5% increase in capacity saves 80,956 days of queue time (2 days per permit). This adds 51.6 plan checker-years at a cost of \$9.1 million, and it saves builders \$19.7 million, a return of \$2.16 per dollar. Across the 30 simulation replications, days saved range from 74,117 to 87,411 (5th to 95th percentiles). This range reflects only simulation noise for a given set pf parameters, not uncertainty about the parameters or the model itself.

I.4 Model fit

Table A9 reports the estimated parameters and compares the data and model moments. The estimated clearance load is close to zero ($\hat{\theta} = 0.018$), meaning that a permit occupies its slot for its review time plus about 2% of its own clearances and corrections time (or about five days on average). The model is rejected by the overidentification test ($J = 476.4$, $p < 0.001$).

The model comes close to matching the upper tail of queue times, as the simulated 90th percentile is 35.7 days compared to 35.0 in the data (a gap of 1.4 standard errors). It slightly understates the mean, at 15.3 days compared to 17.3 days in the data. However, the model fits two other features poorly. First, it overstates the dispersion of queue times, as the simulated standard deviation is 38.3 days, which is more than twice the observed standard deviation of 16.1 days. Second, it overstates the rush effect, at 26.2 days compared to 13.0 days in the data.

Figures A23 and A24 show where these gaps arise. Observed mean queue times

are between about 10 and 40 days. In a few office-years, simulated mean queue times are much higher, because arrivals exceed the office's simulated capacity and the backlog grows throughout the year. Matching the overall mean then requires short simulated queues elsewhere, which produces the excess dispersion. The rush effect works in a similar way, as simulated queue times rise during the four weeks before a code change by about twice as much as observed queue times. This suggests that actual offices absorb surges in filings (i.e. through overtime and other adjustments the model lacks), so observed queue times would respond less to a workload shock than a queue with a fixed capacity.

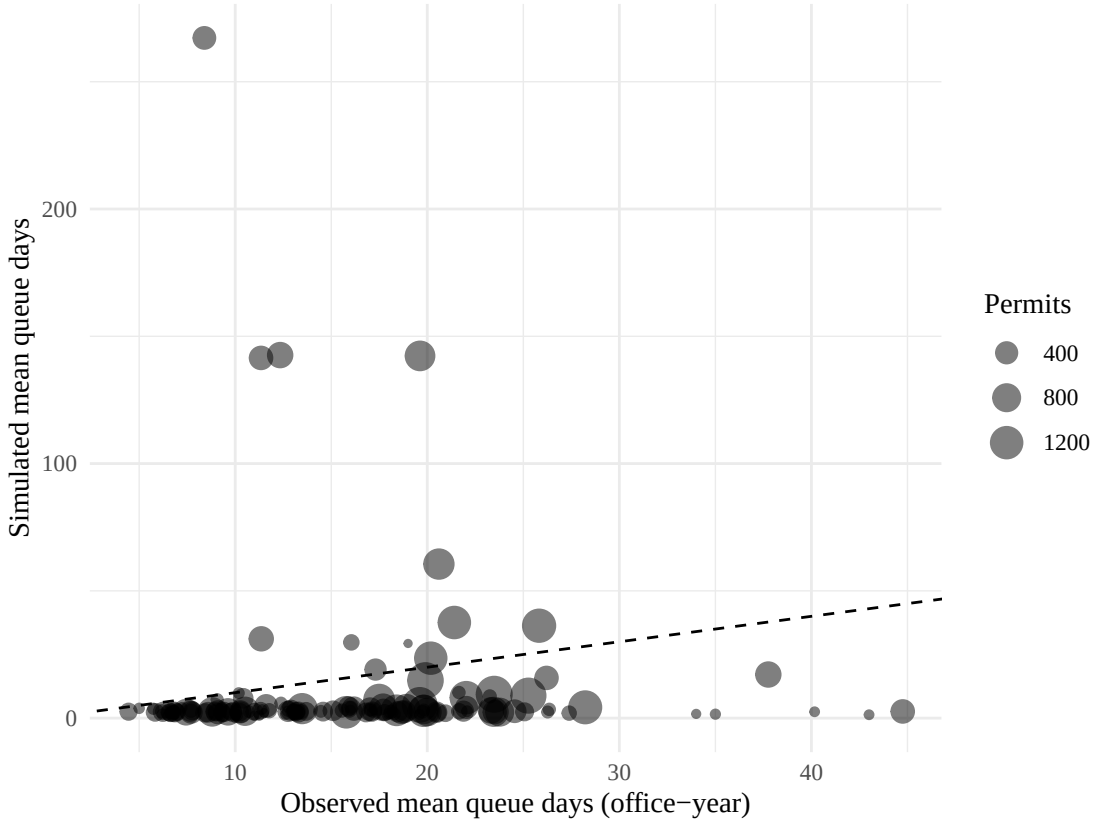
Implications for the counterfactual. The model's misfit also affects estimates of days saved. In the model, queue times are very sensitive to capacity in the office-years where arrivals exceed capacity and insensitive elsewhere. Adding 5% capacity therefore saves many days in the few office-years with long simulated backlogs but saves fewer days in the rest. If actual offices absorb excess workload in ways that the model lacks, the model likely overstates the savings in the most congested office-years, while understating them in office-years where simulated queues are too short. The net direction of these errors is unclear. Moreover, the days saved also exclude any effect of added staff on clearance and corrections times, which understate the time savings. I therefore treat the staffing counterfactual as an illustrative calculation rather than a precise estimate.

Table A9: Queueing model: parameter estimates and fit

	Data	Model	Difference (s.e.)
Targeted moments (days)			
Mean queue time, non-rush	17.3	15.3	−1.9 (0.26)
Rush effect	13.0	26.2	13.3 (1.61)
Standard deviation, non-rush	16.1	38.3	22.2 (1.71)
90th percentile, non-rush	35.0	35.7	0.7 (0.47)
Parameters			
Slots per plan checker, s		3.79	
Clearance load, θ		0.018	
J statistic (p-value)		476.4 (<0.001)	
Counterfactual: +5% capacity, 2005–2024			
Total queue days saved		80,956	
Simulation draws, 5th–95th pct.		74,117–87,411	
Days saved per permit		2.0	
Share of permits with longer queues		0.0%	
Realized capacity increase		5.1%	
Plan checker-years added		51.6	
Builders' savings ($ million)		19.7	
City cost ($ million)		9.1	
Return per dollar		2.16	

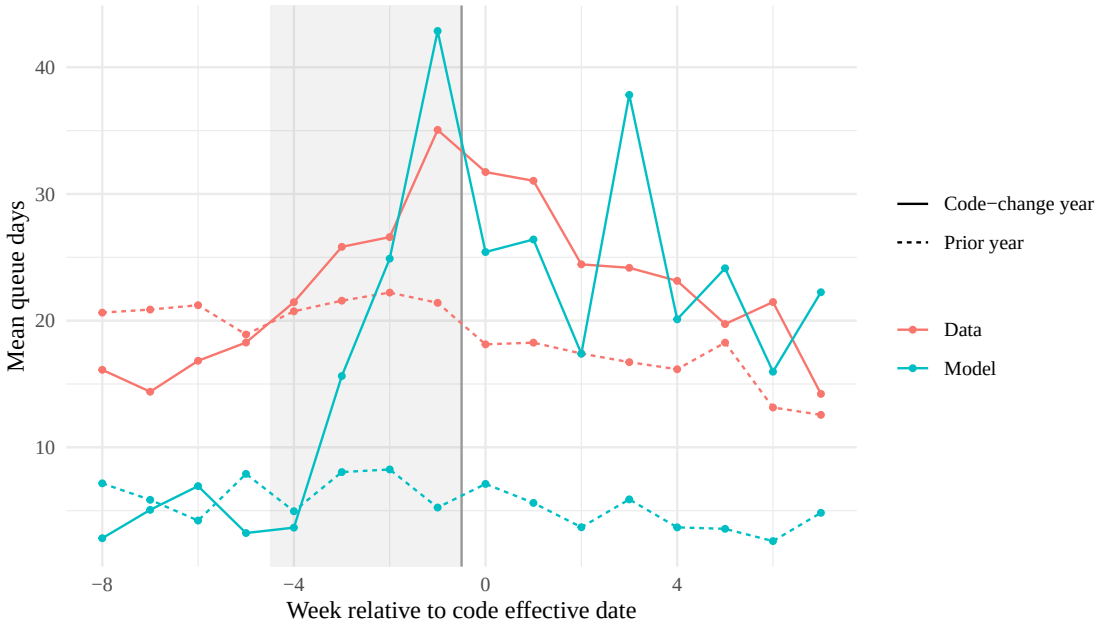
NOTES: This table reports the estimated parameters of the queueing model and compares the targeted moments in the data and the model. Data moments are computed from observed queue times (submittal to assignment) for new residential permits. Model moments are computed with the same code from queue times simulated at the estimated parameters and averaged across $R = 30$ replications. The difference is the model moment minus the data moment, and the standard errors of the data moments are from a bootstrap over office-weeks ($B = 500$). The rush effect is the mean queue time of applications filed in the four weeks before a building code change minus that of applications filed in the preceding four weeks, in code change years. Days saved are savings in queue time only. The realized capacity increase is weighted by permits. The city cost is $176,450 per added plan checker-year (in 2023 $).

Figure A23: Observed and simulated mean queue time by office and year



NOTES: This figure plots the observed and simulated mean queue time for each office-year, where the marker size is proportional to the number of permits. Simulated queue times are computed at the estimated parameters and averaged across replications. The dashed line is the 45-degree line.

Figure A24: Observed and simulated queue time around building code changes



NOTES: This figure plots the mean queue time by week of submittal, from eight weeks before to eight weeks after the effective date of a new building code, separately for code change years and the same weeks one year earlier. The shaded area is the rush period. Only the rush effect in code change years (the difference between the rush weeks and the four weeks before them) is targeted in the estimation.

J Simulating the "Pull-Forward" Effect of Shorter Issuance Times

This section details how I translate the days saved under each personnel policy in Section 8 into the number of additional projects (and units) that would have been completed by the end of my data collection on October 23, 2025 (denoted as $\bar{T}$). I follow the simulation procedure in Section 5 of Gabriel and Kung (2026) and refer readers there for details. Below, I describe my implementation and how it differs from theirs.

Construction time. For permit i submitted on date A_i with issuance time D_i, let CT_i^* be the latent number of days between permit issuance and project completion. I define the completion date as the earliest date on which the permit is finaled or issued a certificate of occupancy. Projects that were issued but not completed by $\bar{T}$ form the set of outstanding projects $\mathcal{O}$, for which construction is unobserved. I exclude any permits that have no updates after more than 10 years after issuance or that have been canceled or revoked.[31]

As in Gabriel and Kung (2026), I model $CT_i^* = X_i\beta^{CT} + \varepsilon_i^{CT}$ with logistic errors, where X_i contains the project characteristics from Equation (1) that are fixed at submission, and I estimate it as an accelerated failure time model that accounts for the right-censoring of construction time among outstanding projects.

Counterfactual completions. My implementation differs from Gabriel and Kung (2026) in two ways. First, rather than reducing the predicted component of approval time by a fixed percentage, I take the days saved for each permit, Δd_i^p, directly from each personnel policy p as discussed in the main text and Appendices H and I.[32] Second, their approval time includes the time spent obtaining entitlements and averages 652 days for multifamily projects, whereas D_i covers only the building permit process and averages 296 days in my sample. Holding fixed submittal dates and latent construction time, an outstanding project is completed by $\bar{T}$ under policy p if $A_i + D_i - \Delta d_i^p + CT_i^* \leq \bar{T}$. Following Gabriel and Kung (2026), I simulate CT_i^* for outstanding projects by drawing the residual from its estimated distribution, conditional on the project not being complete by $\bar{T}$. I average the completion indicator over $M = 100$ draws, using the same draws for every policy, to obtain $\widehat{\text{Complete}}_i^p$. The number of additional completions is computed as $\text{Gain}^p = \sum_{i \in \mathcal{O}} n_i \, \widehat{\text{Complete}}_i^p$, where n_i equals one when counting projects and the number of dwelling units when counting units.

I construct 95% confidence intervals from $B = 500$ bootstrap samples of permits, re-estimating the construction time model and redrawing the residuals in each sample while holding Δd_i^p fixed. Like Gabriel and Kung (2026), the simulation assumes that construction time does not respond to the policy changes, which is consistent with my results in Section 7 that being assigned a faster plan checker does not affect construction time or the likelihood of completion within 3 years (Table 9).

Results. Table A10 reports the results. Replacing the lowest-performing 25% of plan checkers with the median plan checker would have completed an additional 899 units (235 projects) by October 23, 2025, which is 2.7% of the units in outstanding projects. Increasing

[31] The outstanding set consists of 5,461 projects (33,147 units). I restrict the sample to permits submitted between 2005 and 2024, the same years used in Appendices H and I.

[32] Permits that are not included in the queueing simulation are assigned zero days saved, and Δd_i^p is capped at D_i.

plan checker pay by 10% would have completed an additional 790 units (109 projects), and increasing staffing by 5% an additional 26 units (10 projects). The implied elasticities are 0.14 and 0.31 for the first two policies, compared to approximately 0.5 in Gabriel and Kung (2026). Since their approval time is more than twice as long as issuance time in my sample, a 1% reduction saves fewer days in my setting, and scaling their elasticity by the ratio of the two averages yields approximately 0.23, which lies between my two estimates. The staffing policy has a much smaller elasticity (0.03), because its days saved come only from queue time which is a smaller share of total issuance time (Appendix I).

These estimates are lower bounds for two reasons. First, they exclude the "incentive" effect of shorter issuance times on new project starts, if shorter timelines lower uncertainty and incentivize projects that otherwise would not get proposed. Second, they exclude projects still in the plan check phase on October 23, 2025.

Table A10: Projects and units completed by October 23, 2025 under alternative personnel policies

	Replace lowest 25% with median	Increase pay by 10%	Increase staffing by 5%
Avg. reduction in issuance time	4.4%	1.8%	0.6%
Panel A: Projects			
Additional projects completed	235	109	10.4
% of outstanding projects	4.31%	1.99%	0.19%
	[3.97%, 4.68%]	[1.88%, 2.02%]	[0.16%, 0.20%]
% of baseline completions	0.71%	0.33%	0.03%
Elasticity	0.16	0.19	0.05
Panel B: Units			
Additional units completed	899	790	25.5
% of outstanding units	2.71%	2.38%	0.08%
	[1.87%, 3.46%]	[1.90%, 2.82%]	[0.06%, 0.11%]
% of baseline completions	0.62%	0.54%	0.02%
Elasticity	0.14	0.31	0.03
Outstanding and completed as of Oct. 23, 2025			
Outstanding projects (units)		5,461 (33,147)	
Completed projects (units)		32,975 (145,525)	

NOTES: This table reports the additional projects and units that would have been completed by October 23, 2025 under each personnel policy, using the simulation described in Appendix J. Outstanding projects are those issued but not completed as of October 23, 2025, excluding expired, canceled, and abandoned permits. The average reduction in issuance time is total days saved divided by total issuance time. The elasticity is the percent increase in baseline completions divided by the average percent reduction in issuance time. 95% confidence intervals from 500 bootstrapped samples are reported in brackets.