Development Approval Timelines, Approval Uncertainty, and New Housing Supply: Evidence from Los Angeles*

Stuart Gabriel† Edward Kung‡

June 18, 2024

Abstract

We provide credible estimates of the effect of duration and uncertainty in local regulatory approval times on the rate of housing production. The analysis derives from a novel dataset of development timelines for all multifamily housing projects permitted in the City of Los Angeles between 2010 and 2022. As a lower bound, simply by pulling forward in time the completion of already started projects, we estimate that reductions of 25% in approval time duration and uncertainty would increase the rate of housing production by 11.9%. If we also account for the role of approval times in incentivizing new development, we estimate that the 25% reduction in approval time would increase the rate of housing production by a full 33.0%. Both the expected value and the uncertainty in approval times are salient to incentivizing new development. The results provide new evidence that local approval processes are a significant driver of housing supply and reinforce the notion that municipal regulatory reform is an important component of housing reform.

Keywords: local land-use regulation, approval timing and uncertainty, housing production
JEL Classification: R31, R38, R52, R14.

*We thank Mary Leslie, Elaina Houser and seminar and conference participants at UCLA, the University of Wisconsin, and the University of Tokyo for helpful comments. Research funding was provided in part by the Los Angeles Business Council Institute. The information presented herein (including any applicable tables and charts) is based on data provided by the Los Angeles Department of Building and Safety, the Los Angeles Zoning and Map Access System, and the Los Angeles Department of Water and Power. The views expressed in this paper are solely those of the authors and do not necessarily reflect those of the LABC Institute.

†Corresponding Author: UCLA Anderson School of Management, 110 Westwood Plaza, Suite C412, Los Angeles, CA 90095-1481; stuart.gabriel@anderson.ucla.edu

‡California State University, Northridge, David Nazarian College of Business and Economics, 18111 Nordoff St, Northridge, CA 91330; edward.kung@csun.edu

1 Introduction

Substantial academic research and policy debate have focused on the housing market effects of local land-use regulation.[1] Local regulatory regimes often include growth management ordinances or other explicit mechanisms designed to achieve fiscal, environmental, or social goals. Without commenting on the worthiness of the these goals, the use of such regulatory instruments introduces bureaucratic frictions that may have unintended consequences for the efficient provision of housing. Indeed, substantial bureaucratic timelines and related uncertainty in the granting of local regulatory approvals may effectively depress the provision of affordable housing via costly delay and disincentivizing of new investment. Whether consistent with or contrary to local policy goals, the bureaucratic frictions imposed by municipal approval processes may be highly salient to the supply of affordable housing, and of housing overall.

Concerns regarding inadequacy of affordable housing supply are broadly evident in Los Angeles. As one of the least affordable major cities in the U.S., Los Angeles has long failed to address issues of housing inadequacy. A full one-third of renter-households in Los Angeles are severely cost burdened, defined as spending more than one-half of their income on rent.[2] Shortages of affordable housing factor negatively into the shelter prospects of the roughly 47,000 homeless persons living in the city.[3] Further, limited housing supply discourages households and firms from locating in Los Angeles, reduces the economic competitiveness of the city (Gabriel and Painter, 2020), and contributes to a broader spatial misallocation of resources (Hsieh and Moretti, 2019). By all estimates, the city will fall far short of the State of California

[1] See, for example, Quigley and Raphael (2005), Glaeser and Ward (2009), Turner et al. (2014), Hilber and Vermeulen (2016), Jackson (2016), Albouy and Ehrlich (2018), Glaeser and Gyourko (2018), Brueckner and Singh (2020), Gyourko et al. (2021), Gyourko and Krimmel (2021), Ben-Moshe and Genesove (2022). Gyourko and Molloy (2015) provides a further review of the literature.

[2] American Community Survey, 2019 ACS Estimates.

[3] The Greater Los Angeles Homelessness Count, Los Angeles Homelessness Authority, June 2023.

Regional Housing Needs Assessment (RHNA) goal of 457,000 new housing units for the 2021-2029 period, which requires a roughly five-fold increase in production over the 84,000 units produced from 2010-2019.[4]

In this paper, we undertake theoretical modelling and empirical analysis of duration and uncertainty in development regulatory approvals and their effects on new housing supply. We uniquely compile multiple administrative datasets that allow us to track development time for the universe of multi-family projects permitted by the City of Los Angeles between January 2010 and November 2022. We identify specific elements of the approvals and entitlements process and provide new estimates of their incremental effect on approval and construction timelines. We empirically model the distribution of approval and construction times and use these models to simulate housing production outcomes if approval times were reduced. To our knowledge, this is the first paper to use detailed administrative data on actual approval times to estimate robust and credible effects of approval time on housing production. Further, we provide estimates of how both the expected value and the variance of approval time affect project starts by incentivizing new development.

Our research findings provide new estimates of the effect of approval time on the rate of housing production. Over the 2010-2022 sample timeframe, approval times comprised roughly 37 percent (on average) of the nearly 4 years required to complete a multi-family project in the City of Los Angeles. As a lower bound, this paper finds that a 25 percent reduction in the mean and volatility of approval times during the 2010-2022 period would have increased the rate of housing production in Los Angeles by 11.9 percent, simply by pulling forward in time the completion of already started projects. If we further take into account the effect of approval times in incentivizing new development, we estimate that this same 25 percent reduction in the mean and

[4] Los Angeles City Planning 2021-2029 Housing Element

variance of approval times would have increased the rate of housing production by a full 33.0 percent. We show that both the expected value and the uncertainty in approval times matter for incentivizing new development. The analysis shows that reductions in approval time can have a large and economically meaningful effect on housing supply, and therefore reinforces the notion that municipal regulatory reform is an important component of housing reform.

The plan of the paper is as follows. The next section provides a theoretical framework that rigorously describes the economic mechanisms through which reductions in approval time can increase the steady state rate of housing production. Section 3 then describes the data and institutional context of the empirical analysis. Section 4 specifies empirical models of the development approval process and reports on the salience of specific factors to the estimated timelines. Given model estimates, section 5 simulates the effects of reductions in approval times on project completions. Section 6 then estimates the effect of both expected value and volatility in approval times on new construction. In section 7, we provide concluding remarks.

2 Theoretical Framework

To help us understand the role of approval time in housing production, we develop a housing supply model in which development time and the risk thereof is explicitly modeled. In the model, risk-averse developers face volatility in development costs due to uncertainty in development time.[5] A reduction in the mean and volatility of development time increases the expected present value of housing development,

[5] In the model, we abstract from any differences in approval time vs. construction time. In the empirical section of this paper, we focus on approval time because that is where policy can have the most immediate effect. The model, however, is applicable to any factor contributing to long development times, whether during the approval or the construction phase. What matters is that there are costs associated with time in both phases. It would be straightforward to extend the model to allow for different costs during the approval phase and construction phases of a project.

thus incentivizing more housing starts. In addition, the number of developers and the amount of development resources are limited. Increasing the speed of project completion therefore frees up more resources to be used on new development in the steady state.

Development

Present value of housing development

We start by considering the incentives of a housing developer. A housing developer has the opportunity to develop a single housing unit which they will sell for price p at the time of completion. The development cost is ζ per unit of time spent in development and the developer discounts cash flows at a rate r. It takes x units of time to complete the unit. The present value of this opportunity is therefore:

$$PV = -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) e^{-rx} \tag{1}$$

Uncertainty in development time

When developers start development, they know the development cost ζ and the discount rate r. They do not, however, know the amount of time it will take to complete development. Let X be a random variable denoting the time to complete a housing unit. X is distributed according to the p.d.f. $f(x)$, the c.d.f. $F(x)$, and the m.g.f. $m(t)$.[6]

PV is therefore a random variable that depends on X. The expected value and

[6] As a reminder, the moment generating function is defined as $m(t) = E\left[e^{tX}\right] = 1 + E\left[X\right]t + \frac{1}{2}E\left[X^2\right]t^2 + \ldots$ and is obtained from a Taylor expansion around $t \approx 0$.

the variance of PV are:

$$E\left[PV\right] = -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) m\left(-r\right) \tag{2}$$

$$V\left[PV\right] = \left(p + \tfrac{\zeta}{r}\right)^2 \left[m\left(-2r\right) - m\left(-r\right)^2\right] \tag{3}$$

Risk averse developers

Developers are risk averse. Their expected utility over a random present value PV has the following mean-variance form:

$$E\left[u(PV)\right] = E\left[PV\right] - \gamma V\left[PV\right] \tag{4}$$

where γ is a parameter measuring risk aversion.

Substituting equations (2) and (3) into equation (4) gives us an exact formula for the expected present value of a development opportunity:

$$E\left[u(PV)\right] = -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) m\left(-r\right) - \gamma \left(p + \tfrac{\zeta}{r}\right)^2 \left[m\left(-2r\right) - m\left(-r\right)^2\right] \tag{5}$$

Taking a first order approximation of the mean term and a second order approximation of the variance term gives us:

$$E\left[u(PV)\right] \approx -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) \left(1 - rE\left[X\right]\right) - \gamma \left(p + \tfrac{\zeta}{r}\right)^2 r^2 V\left[X\right] \tag{6}$$

which clearly shows how the expected utility of a development opportunity is decreasing in the mean and variance of development time. Equation (6) also shows that the impact of development time is greater when r is higher, which is intuitive because a higher r means a higher time value of money.

Housing starts

There is a unit mass of housing developers. Developers are either in a state of active development or they are exploring opportunities. If a developer is exploring opportunities, then in each moment in time they draw an opportunity cost u_0 from a distribution of opportunity costs characterized by the c.d.f. $G(u_0)$. If the opportunity cost is less than the expected utility of housing development—that is, if $u_0 \leq E[u(PV)]$—then the developer will begin developing the unit. Let n be the number of developers in active development, so $1 - n$ is the number of developers exploring opportunities. The number of housing starts, s, is therefore:

$$s = (1 - n)G\left(E\left[u(PV)\right]\right) \tag{7}$$

Since $E\left[u(PV)\right]$ is decreasing in the mean and variance of development time, the number of starts is also decreasing in the mean and variance of development time. We call the effect of development time on $G\left(E\left[u(PV)\right]\right)$ the "incentive effect", because it operates through the present value of housing development.

Market equilibrium and steady state

Starts, completions, and active developments

Let us now consider the evolution of this market over time. Let $s(t)$ be the number of starts in time t, let $c(t)$ be the number of completions, and let $n(t)$ be the number of developers/units in active development. $s(t)$, $c(t)$, and $n(t)$ are related by the following equations:

$$c(t) = \int_0^\infty s(t-x)f(x)dx \tag{8}$$

$$n(t) = \int_0^\infty s(t-x)\left[1-F(x)\right]dx \tag{9}$$

Equation (8) says that the number of completions at time t is equal to the number of starts from x periods ago, times the share of them that have a development time of x, integrated over x. Equation (9) says that the number of units in active development at time t is equal to the number of starts from x periods ago, times the share of them that have a development time greater than x, integrated over x.

Steady state equations

In the steady state, the number of starts, completions, and units in active development are the same in each moment of time, so $s(t) = s$, $c(t) = c$, and $n(t) = n$. The relationships between these variables becomes:

$$c = \int_0^\infty sf(x)dx = s \tag{10}$$

$$n = s\int_0^\infty \left[1-F(x)\right]dx = sA \tag{11}$$

where $A = \int_0^\infty \left[1-F(x)\right]dx$ is a constant that depends only the distribution of development times. If development times are faster, then A will be smaller meaning there will be fewer units in active development (and more developers exploring opportunities) for any given rate of project starts.[7]

Combining equations (7), (10), and (11), we obtain an equation for the steady state

[7] To see this, simply note that if F' first order stochastically dominates F (i.e. $F'(x) < F(x)$ for all x), which implies that development times are slower under F' than F, then $A' > A$.

number of completions; in other words, the steady state rate of housing production:

$$c = \frac{G\left(E\left[PV\right]\right)}{1 + AG\left(E\left[PV\right]\right)} \tag{12}$$

This equation is increasing in the present value of development, $E[PV]$, and decreasing in A (which is itself decreasing in development time.)

Incentive effect and pull-forward effect

Equation (12) highlights two channels through which reductions in development time affect the steady state rate of housing production:

1. A reduction in the mean and volatility of development time increases the present value of housing development, $E[PV]$, which in turn increases the steady state rate of housing production, c. We call this the *incentive effect* because it operates through the channel of incentivizing developers to engage in development.

2. A reduction in development time reduces the amount of resources tied up in active development, A, which in turn frees up these resources for new starts, thus increasing the steady state rate of housing production, c. We call this the *pull-forward* effect because it operates through the channel of pulling forward the completion of housing units in time, which increases the overall rate of production.

Housing demand, population growth, and equilibrium price

We close the model by discussing price formation. We assume there is a growing population which necessitates a steady state increase in housing supply. However, the rate of population growth in the market is moderated by the price p. (It can be assumed that those for whom the price is too high choose to live somewhere else.) Let the number of net new housing demanders in each period be $d(p)$, with $d'(p) < 0$ (i.e. a downward sloping demand curve). The net rate of housing production is $c-\delta$, where δ is the rate of housing depreciation. In the steady state equilibrium, the number of new demanders in each period must equal the net rate of housing production:

$$d(p) = c - \delta \tag{13}$$

Since d is downward sloping in p and c is upward sloping in p, this is a standard supply and demand model of house price formation. However, in this model the mean and variance of housing development time will shift the supply curve through the effect on c. A reduction in the mean and variance of development time would therefore shift the supply curve outwards and reduce equilibrium price and increase the equilibrium rate of housing production.

3 Data and Institutional Background

Housing development in Los Angeles

To build new housing in Los Angeles, the developer must first obtain a permit from the Los Angeles Department of Building and Safety (DBS). In addition, if the project requires exceptions to zoning codes or meets other criteria that trigger discretionary review, then the project must also obtain entitlements from the Los Angeles Department of City Planning. Examples of discretionary review include Site Plan Review, which is triggered for projects that add 50 or more dwelling units, and Environmental Impact Review, which is triggered by California CEQA rules. A developer can submit their DBS permit application before all entitlements are approved, but the entitlements generally must be approved before the DBS permit can be issued. Construction can begin once the DBS permit is issued. When construction is complete and the project is fit for habitation and meets all other requirements, DBS issues a Certificate of Occupancy (CofO).

Permits data from Department of Building and Safety

To study the role of approval times in housing development, we compiled a rich dataset containing the universe of all multifamily housing projects permitted by DBS between January 2010 and November 2022. The dataset contains detailed information about each project, including the dates of permit application and issuance, the date of CofO issuance, the application and approval dates of related entitlements, physical characteristics of the project (height, square footage, number of units), and geographic characteristics (address, City Council District).

The primary data source is a publicly available dataset provided by DBS and available on Los Angeles's open data portal.[8] This dataset contains information on *all* new building permits issued by DBS between January 2010 and November 2022, including single-family, multi-family, commercial, and mixed-use buildings. From the dataset, we extracted the universe of multi-family housing projects using a combination of the permit subtype and the project use description.[9] The dataset contains both complete and incomplete projects, where "complete" is defined as having been issued a CofO as of November 28th, 2022. To focus only on projects that were either completed or in active development, we excluded any projects with a closed or expired

[8] A live link to the public dataset can be found here: `https://data.lacity.org/City-Infrastructure-Service-Requests/New-Building-Permits-2010-to-Present/46r2-n9vp`. We accessed and downloaded this data on November 28th, 2022.

[9] The permit subtype can be "1 or 2 Family Dwelling", "Apartment", or "Commercial". We excluded permits with the "1 or 2 Family Dwelling" subtype. We included all permits with the "Apartment" subtype. We included "Commercial" permits only if its use description indicated some residential usage.

permit.

For each project, we determined the number of income-restricted units (i.e. "affordable" units) using the permit work description. An affordable unit is defined as a unit that the project owner has agreed to set aside for residents within an income range, usually determined as a percentage of the area median income. The city provides incentives for developers to set aside units in the affordable range, typically via density bonuses or other exceptions to the zoning code. We classified projects into three groups: "market-rate", "mixed-income", and "100%-affordable". A project was classified as "100%-affordable" if all its units except one or two manager units were affordable units. A project with no affordable units was classified as "market-rate". All other projects were classified as "mixed-income".

The resulting dataset contained 2,677 projects representing 120,213 total dwelling units, of which 102,897 were market-rate and 17,316 were affordable. 1,712 projects were completed (issued a CofO) as of November 28th, 2022, and 965 projects were unfinished as of November 28th, 2022. The full breakdown of counts by project type are reported in Table 1.

[Table 1 around here.]

Entitlements data from City Planning

In that one of our primary goals is to assess the speed of housing production, it is important that we use a meaningful and consistent definition of project start date. The date of DBS permit application is not always a good measure of start date because the process of obtaining entitlements can start multiple years prior to the submission of plans to DBS. This is especially true of projects that require long discretionary reviews, such as projects requiring Environmental Impact Reviews (EIRs). Therefore, a conceptually appealing measure of start date requires that we observe both the permit submission date *and* the entitlement application dates.

The DBS data described above do not contain information on entitlements. To obtain information on entitlements, we took advantage of the Los Angeles Zoning Information and Map Access System (ZIMAS) and the Department of City Planning's Planning Document Information System (PDIS).[10,11] We link DBS permits to entitlement cases using a three-step procedure. First, we take the primary address and associated PINs of each project in the DBS data and use ZIMAS to retrieve all Planning Department cases associated with those parcels. Second, we use PDIS to retrieve information about each of those linked cases. Third, we determine whether each linked case was relevant to the development project. Using this procedure, we linked 1,389 projects to at least one entitlement case. 1,288 projects were not linked to any case. Of the projects that were linked to at least one case, the average number of cases linked was 2.7.[12] With the additional data on entitlements, we were able to measure the start date of a project based on the date of the first seen entitlement or permit application.

Electrical data from Department of Water and Power

In addition to the data on entitlements, we also sought data on new power service installations from the Los Angeles Department of Water and Power (DWP). This was motivated by our conversations with industry professionals who identified long wait

[10]ZIMAS: `https://zimas.lacity.org`. ZIMAS allows users to input an address, assessor parcel number (APN), or parcel identification number (PIN) and retrieve zoning information about the parcel, including relevant Planning Department case numbers.

[11]PDIS: `https://planning.lacity.org/pdiscaseinfo/search`. PDIS, a system maintained by the Planning Department, allows users to retrieve case information by inputting a case number. The information retrieved by PDIS on entitlement cases includes both the application date and the approval date of the entitlement request.

[12]We discuss the rules used to determine relevancy and the quality of the linkage in the appendix.

times for electrical service installation as a major bottleneck.

To obtain this data, we provided DWP with the primary site address for each project in the permits data. DWP then returned, for each project, all new power service work requests associated with that address from 2010 to present.[13] We focus on work requests for the design and installation of new overhead and underground circuits and service voltage. There were 808 overhead work requests and 1,291 underground work requests. In our empirical analysis, we considered the impact of a project requiring *any* new overhead or underground installations on project timelines.[14]

Example project: Florence Towne Apartments

To provide readers with an example of a project in our data, we discuss the Florence Towne Apartments. The Florence Towne Apartments is a 51-unit affordable housing development located at 410 E. Florence Ave. in South L.A. It was entitled under Planning Department cases DIR-2017-4059-TOC and ENV-2017-4060-CE. Both cases were submitted on October 10th, 2017 and approved on February 6th, 2018. The new building permit application was submitted to DBS on March 1st, 2018 and issued on January 18th, 2019. The Certificate of Occupancy was issued on July 1st, 2022.

[13] We filtered out any work requests that were entered before the project start date or after the project's CofO date.

[14] For the projects that we did not receive DWP data for, it was likely because of an error in finding matching addresses. Addresses are not standardized between the DBS and DWP databases, leading to an imperfect matching process. In our empirical work, we assume that projects without DWP data did not require any overhead or underground circuit installations, but the results are similar if we drop any projects without DWP data. In total, we received data on 70,114 work requests for 2,324 projects.The vast majority of these work requests were for new meter installations. We ignore new meter installations in our analysis because they are required on almost every project.

Figure 1: Florence Towne Apartments: Project Timeline

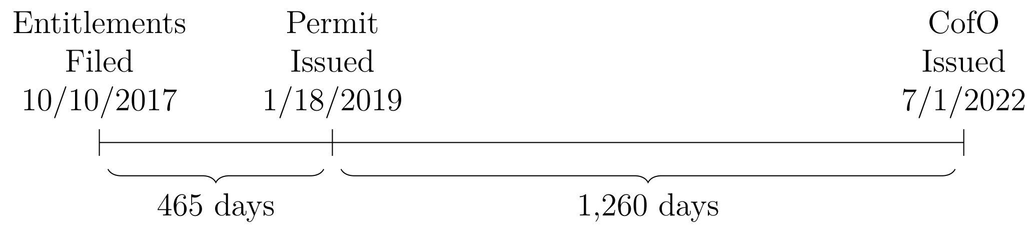

We define a project's "approval" time as the number of days between the first seen entitlement or permit application (i.e. the start date) and the date in which the DBS permit is issued. We define a project's "construction" time as the number of days between the permit issuance and the CofO issuance.[15] Based on these definitions, the Florence Towne Apartments spent 465 days in the approval period and 1,260 days in the construction period. Out of the 465 days in the approval period, 119 days were spent in entitlement and 323 days were spent in permitting, with a few weeks in between.

Summary statistics

Table 2 lists the variable names and definitions from our data. Table 3 shows summary statistics for each of these variables across project types. Of special note are the project timelines. Our data show that the average approval time across all projects was 652 days (1.8 years) and the average construction time was 863 days (2.4 years).[16]

[15]Note that not all the time in "construction" is necessarily spent on physical construction. There may be various approvals and compliance checks that the project must undergo even after the issuance of a permit by DBS.

[16]Approval time is measured for all projects. Completion time is measured only for completed projects.

The average total development time was 1,413 days (3.9 years).[17] Since larger projects take longer to complete, the completion time of the average dwelling unit is even longer than for an average project. The average dwelling unit took 1,784 days (4.9 years) to complete.

[Table 2 around here.]

[Table 3 around here.]

In addition to long average timelines for project development, there is also a significant amount of uncertainty. Figure 2 shows the distribution of development times for completed projects in our data. The 25th percentile for total development time was 946 days (2.6 years) and the 75th percentile was 1,739 days (4.8 years). This means that 1 out of every 4 multifamily housing developments built in Los Angeles between 2010 and 2022 took 4.8 or more years to complete, from first permit or entitlement application to CofO. The variability in total development time is not simply due to predictable factors. Using a simple OLS model of completion time on observable characteristics, we found that the observed characteristics could only explain 25.8% of the variation in total development time.[18] There is thus a significant amount of remaining uncertainty in development time that is not explicated by the variables in our dataset.

[Figure 2 around here]

Figure 3 illustrates the impact of long development times on total housing production. The top blue line in Figure 3 shows the cumulative number of dwelling units

[17] Construction time and total development time are measured only for completed projects. Approval time is measured for all projects.

[18] See Table 2 for a list of the variables used.

started by project start date. The bottom orange line shows the cumulative number of dwelling units completed by CofO date. Over the course of the data period, 120,213 dwelling units were started, but only 71,532 were completed, leaving a gap of 48,681 unfinished units. One of the central questions we address in this paper is the extent to which this gap of 48,681 units could be reduced simply by "pulling up" the orange line through accelerated timelines. We turn to this question next in Sections 4 and 5.

[Figure 3 around here]

4 Statistical Models of Development Time

To assess the role of approval times in housing development, we first develop statistical models of approval and construction times. The models allow us to estimate the mean and variance of approval times and construction times after controlling for project characteristics. Crucially, the models allow us to estimate the distribution of project completion times, which is a latent variable in our data since not all projects were completed. This in turn allows us to simulate how changes in approval time affect project completion rates within a fixed time period.

Approval Time. We statistically model approval time according to the following equation:

$$AT_i = X_i\beta + \epsilon_i \tag{14}$$

AT_i is the approval time, measured in days, for project i. X_i is a vector of project characteristics. This includes structural characteristics, such as the number of units, building height, and building square footage, but also what entitlements the project required and the type of project (market-rate, mixed-income, or 100%-affordable). Table 2 shows the full list of features that we consider in the model. The error term, ϵ_i, is modeled as a logistic distribution with mean 0 and an unknown scale parameter.

Construction Time. Construction time is modeled according to the following equation:

$$CT_i^* = X_i\gamma + \nu_i \tag{15}$$

CT_i^* is the latent construction time, in days, for project i. As above, X_i is a vector of project characteristics. The error term, ν_i, is again modeled as a logistic distribution with mean 0 and an unknown scale parameter.

For simplicity, we assume the independence of ϵ_i, ν_i, and X_i. This may not hold in practice. There may be unobserved factors in ϵ_i and ν_i that are correlated with observed factors in X_i. Moreover, it is likely that ϵ_i and ν_i are correlated, because more complex projects may have both unexpectedly long approval times and unexpectedly long construction times. We therefore cannot interpret (14) and (15) as structural equations. However, our simulation results in Sections 5 and 6 do not require a structural interpretation of equations (14) and (15) since they rely only on the reduced form distributions of AT_i and CT_i^*.

In estimating our model for construction time, it is important to note that CT_i^* is not observed for the unfinished projects. For these projects, we only know that CT_i^* is larger than the number of days between November 28th, 2022 and the permit issuance date. The model for construction time is therefore in the class of accelerated failure time model for which standard estimation techniques have been developed.[19]

[19] See Kalbfleisch and Prentice (2002).

Regression Results. Table 4 reports our coefficient estimates for models (14) and (15). Although we are not primarily interested in the causal effects of X_i on AT_i and CT_i^*, it is both interesting and informative to examine their relationships in the reduced form. Council District fixed effects were included in the model but the results are omitted for space. We comment on some selected results of interest below.

[Table 4 around here.]

Physical Characteristics. Out of the three physical measurements available in our data—number of units, height, and square footage—height was the most statistically significant variable in the model. The results show that every 10 feet of height (about one story) is associated with 4 additional days in approval time and 12.1 additional days in construction time.

Project Types. The results show that there is significant variation in approval and construction times by project type, even after controlling for project characteristics and required entitlements. 100%-affordable projects had shorter approval and construction times than market-rate projects, whereas mixed-income projects had longer approval times than market-rate projects. The faster development times for 100%-affordable projects may be due to city policies that prioritize affordable housing development, whereas the slower developmentment times for mixed-income projects may be due the increased complexity of such projects.

CEQA. There are three entitlement case types related to CEQA (the California Environemntal Quality Act): Categorical Exemption (CE), Mitigated Negative Declaration (MND), and Environmental Impact Report (EIR). According to the State of California:

> "CEQA requires state and local government agencies to inform decision-makers and the public about the potentially significant environmental effects of a proposed project, ways to minimize those effects, and to indicate alternatives to the project. If a project subject to CEQA will not result in potentially significant adverse effects to the environment, the Commission may adopt a document known as a Negative Declaration or a Mitigated Negative Declaration. If the project may cause adverse environmental effects, the Commission will prepare a more detailed informational document called an Environmental Impact Report (EIR). An EIR contains information on potential effects, measures to mitigate those effects, and an analysis of alternatives to the project. A key feature of the CEQA review process is the opportunity for the public to provide input on Negative Declarations, Mitigated Negative Declarations, and EIRs."[20,21]

Our results show how entitlements related to CEQA affect development time. Unsurprisingly, projects requiring an EIR had significantly longer approval times, by 504 days on average. By contrast, projects using a MND did not have significantly longer approval or construction times than projects not subject to CEQA. Projects that were categorically exempt from CEQA had longer construction times but no significant difference in approval times.

Although EIRs add significant delays to a project's approval time, we note that not many projects required an EIR. Only 28 projects in the data required an EIR, affecting a total of 8,311 dwelling units.

[20] `https://www.slc.ca.gov/ceqa/`. Accessed January 2nd, 2023.

[21] In our data, we observed very few Negative Declarations and mostly Mitigated Negative Declarations. In the few instances where a Negative Declaration was observed, it was lumped together with Mitigated Negative Declaration.

Power installation. In our conversations with developers and city officials, the length of time for DWP to complete a circuit installation was often cited as a bottleneck. Our model results are consistent with that sentiment. The results show that requiring a new overhead circuit installation adds 117 days to the construction time, whereas requiring a new underground circuit installation adds 69 days to the approval time and 176 days to the construction time.

Other Entitlements and By-Right. The results show the effect of various other types of entitlements on approval and construction times. For example, the results show that projects requiring approval by the City Planning Commission had 193 days longer approval times and 125 days longer construction time. Projects requiring Site Plan Review had 106 days longer approval time, but no statistically significant difference in construction times.

Even after controlling for a number of different entitlement types, the results show that projects requiring *no entitlements* (i.e. "by-right" projects) had lower approval times by an average of 197 days.[22] In other words, requiring any entitlements added, on average, 197 days to the approval time, even after controlling for the types of entitlements we observe in the data.

Scale Parameters. In addition to the coefficient estimates, Table 4 also shows our estimated scale parameters for ϵ_i and ν_i. The estimated scale parameter for ϵ_i (the residual of approval time) is 208.6, which implies a standard deviation of 378 days.[23] The estimated scale parameter of ν_i (the residual of construction time) is 331.7, which implies a standard deviation of 602 days.

[22] By-right projects still require approval from DBS, so their approval times are not zero.

[23] The standard deviation of a logistic distribution is $s\pi/\sqrt{3}$, where s is the scale parameter.

Salience of Factors. Although some factors have a large effect on approval or construction time, such as EIR on approval time and CPIOC on construction time, these factors may not apply to many projects. To measure the overall salience of a factor on development time, we computed its combined effect on approval time and construction time and multiplied that by the total number of dwelling units in projects impacted by that factor. The salience of each factor is reported in Figure 4. Salience is measured in unit-years. A salience of x means that if the marginal effect of that factor were eliminated, an additional x units would have been completed one year sooner. Figure 4 shows that if all projects requiring entitlement were instead made by-right, then 48,085 unit-years in development time would have been saved. If all projects requiring new underground circuit installations instead did not require it, then 46,957 unit-years of development time would have been saved. By-right and underground power installation are the two most salient factors, followed by City Planning Commission and Site Plan Review.

[Figure 4 around here.]

5 The Impact of Approval Time on Project Completions

To assess the impact of approval times on project completions, we consider the following thought experiment. Suppose the expected approval time or the uncertainty in approval times was reduced by X%. Taking as given the set of projects started between January 2010 and November 2022, how many additional units would have been completed by November 28th, 2022 if approval times had been reduced in this way? Note that in this counterfactual, the newly completed units must come from projects that in the baseline *did not* complete by November 28th, 2022. Thus, to conduct this exercise, we need to simulate the distribution of latent completion times for the unfinished projects. We describe the simulation procedure below.

First, let $\hat{\beta}$ be the estimated coefficients for the approval time model given by equation (14), let $\hat{\epsilon}_i$ be the estimated residual, and let $\hat{s}_A$ be the estimated scale parameter. By definition, $AT_i = X_i\hat{\beta} + \hat{\epsilon}_i$. Similarly, let $\hat{\gamma}$ be the estimated coefficients for the construction time model given by equation (15) and let $\hat{s}_C$ be the estimated scale parameter. Because CT_i^* is not observed, we cannot directly estimate the construction-time residual, ν_i, for each project; instead, we must simulate it by drawing from its estimated distribution.

Let T_i^0 be the start date of project i, which we take to be exogenous, and let $\bar{T}$ be November 28th, 2022, the censoring date. Finally, let c_i be an indicator equal to 1 if project i was completed by November 28th, 2022 and 0 otherwise. By definition,

$$c_i = \begin{cases} 1 \text{ if } T_i^0 + AT_i + CT_i^* \leq \bar{T} \\ 0 \text{ otherwise} \end{cases}$$

We now describe a procedure for simulating ν_i, the latent residual in the construction time model. When $c_i = 1$, CT_i^* is observed and we can estimate ν_i using the equation $\nu_i = CT_i^* - X_i\gamma$.

When $c_i = 0$, CT_i^* is not observed. Instead, it is only known that $C_T^* > \bar{T} - T_i^0 - AT_i$. Thus, when $c_i = 0$, we only know that $\nu_i \geq \bar{T} - T_i^0 - AT_i - X_i\gamma$. We can therefore estimate a distribution for the latent residual, $\hat{\nu}_i$, using the following formula:

$$\hat{\nu}_i = \begin{cases} CT_i^* - X_i\hat{\gamma} & \text{if } c_i = 1 \\ R \mid R > \bar{T} - T_i^0 - AT_i - X_i\hat{\gamma} & \text{if } c_i = 0 \end{cases}$$

where R is a random variable distributed according to the logistic distribution with mean 0 and scale $\hat{s}_C$.

The intuition for $\hat{\nu}_i$ is quite straightforward. For projects completed by November 28th, 2022, $\hat{\nu}_i$ is simply the difference between the observed construction time and the model's predicted construction time. For projects that were unfinished as of November 28th, 2022, $\hat{\nu}_i$ is drawn from the estimated distribution of ν_i, conditional on being large enough so that the latent completion time is past the censoring date, as is consistent with the data.

We now describe how to calculate the number of units added in the counterfactual. Suppose approval times are changed in such a way that every project receives a new approval time, AT_i'. We assume that in the counterfactual, latent construction times remain unchanged. Let $\hat{c}_i$ be an indicator for whether project i would have been completed by November 28, 2022 in the counterfactual:

$$\hat{c}_i' = \begin{cases} 1 \text{ if } T_i^0 + AT_i' + X_i\hat{\gamma} + \hat{\nu}_i \leq \bar{T} \\ 0 \text{ otherwise} \end{cases}$$

We can therefore calculate the number of units added in the counterfactual by adding up the number of units in projects that were not completed in the baseline but completed in the counterfactual:

$$\# \text{ Units Gained} = \sum_i w_i(1 - c_i)\hat{c}_i \tag{16}$$

where w_i is the number of dwelling units in project i.

Simulating a 25% reduction in expected approval times. We now turn to a simulation in which the mean approval time is reduced by 25%, but the amount of uncertainty stays the same. The counterfactual approval times are defined by:

$$AT_i' = 0.75X_i\hat{\beta} + \hat{\epsilon}_i$$

Under this counterfactual, we find that a 25% reduction in expected approval time increases the number of units produced by 7,375, a 10.3% gain over the baseline of 71,532.

Simulating a 25% reduction in the uncertainty of approval times. To simulate a reduction in the uncertainty in approval times, we define the counterfactual approval times by:

$$AT_i' = X_i\hat{\beta} + 0.75\hat{\epsilon}_i$$

In this counterfactual, the expected approval times conditional on project characteristics remain the same, but the standard deviation is reduced by 25%. This has the effect of compressing the distribution of approval times, making some projects take longer to approve but also reducing the size of the long tail. We find that reducing the uncertainty in approval times by 25% leads to an increase in the number of units produced by 1,545, a 2.2% gain over the baseline.

Simulating a reduction of 25% in total approval times. Lastly, we simulate the effect of reducing total approval time by 25%. This reduces both the mean and the uncertainty of approval times. In other words, we let:

$$AT_i' = 0.75AT_i$$

Under this counterfactual, we find that a 25% reduction in total approval time increases the number of units produced by 8,478, a 11.9% gain over the baseline.

Discussion. The exercises above show how reductions in approval time translate to increased housing production through a single channel: the pulling forward in time of project completions. We therefore call this effect the "pull-forward" effect. Figure 5 illustrates the pull-forward effect visually. The 8,478 new units are gained by "pulling" forward the completions in time.

[Figure 5 around here.]

The estimate of the pull-forward effect relies on few assumptions other than the distribution of the construction time residuals.[24] The estimated pull-forward effects can therefore be considered both robust and conservative estimates of what is possible through the reduction in approval times. The estimates are robust because the pull-forward effect relies only on a distributional assumption and is robust to different choices of said distribution. The estimates are conservative because they allow reductions in approval time to operate through only one channel. In reality, a reduction in

[24] In unreported results, we verified that the estimated pull-forward effects are robust to different choices of the distribution of ν_i.

approval time would also incentivize new development, a channel we turn to in the next section.

Pull-forward effect response spectrum. The above exercises can be repeated for a range of approval time reductions, producing a spectrum of effects. This spectrum is shown in Figure 6 and gives a sense of the bounds of possibility through just pull-forward effect. If all approval times were reduced by 50% (clearly an ambitious goal in real life), then our model suggests that housing production would increase by 25.9% due to the pull-forward effect. If approval times were reduced by a more modest 10%, housing production would have increased by 5.3% due to the pull-forward effect. The spectrum gives policymakers a sense of what reductions in approval time would be needed to achieve different rates of increased housing production. Overall, the results suggest that the pull-forward effect by itself is economically significant in magnitude and thus it is worthwhile for policymakers to target approval time reductions.

[Figure 6 around here.]

6 The Impact of Approval Time on Project Starts

We now turn to estimating the effect of approval times on project starts. Risk-averse developers experience high holding costs and opportunity costs due to delayed development. We therefore expect that reductions in both the expected value and the volatility of approval times would be associated with increased project starts.

To investigate this effect, we exploit variation in the mean and variance of approval times over time and across the 15 Los Angeles City Council Districts. We run panel regressions of the number of projects started in Council District j in a window surrounding time period t on the backward looking mean and volatility of approval times for projects in that district.

To operationalize this idea, let N_{jt} be the number of projects started in council district j between the dates of $t-h$ and $t+h$. Let μ_{jt} be the mean of log approval times for projects permitted in Council District j between the dates of $t-2h$ and t. Let σ_{jt} be the standard deviation of log approval times for projects permitted in Council District j between the dates of $t-2h$ and t.

N_{jt} is modeled as a Poisson distribution in which:

$$E[\ln N_{jt}] = \theta_0 + \theta_1 \mu_{jt} + \theta_2 \sigma_{jt} \tag{17}$$

That is, the expected log number of project starts between $t-h$ and $t+h$ in Council District j depends linearly on the mean and standard deviation of log approval time for projects permitted in Council District j between times $t-2h$ and t. The behavioral model we have in mind is that developers look at recently permitted projects to form expectations about how long their own project will take to approve, and this impacts their decision-making as to whether or not to start a project.

In our preferred specification, we calculate N_{jt}, μ_{jt}, and σ_{jt} on a monthly basis, using the first of each month as the base time period. To ensure large enough sample sizes to calculate mean and volatility, we use a bandwidth of $h = 2$ years. Because the variables used in the regression are aggregates calculated from underlying microdata, we estimate standard errors using 500 bootstrapped samples formed by resampling the underlying microdata 500 times.

Table 5 shows the results of regression (17). We show five specifications. Column 1 shows the baseline result when no variables other than μ_{jt} and σ_{jt} are included. Column 2 adds Council District and year fixed effects to control for unobserved heterogeneity between Council Districts and across time periods. Column 3 adds a variable called "Council Member Real Estate Score", which is equal to the current Council Member's share of campaign contributions from donors associated with the real estate industry in the last election cycle.[25] The real estate score is meant to control for contemporaneous political and economic factors that could influence a developer's decision to start a multi-family housing project in a particular Council District. Column 4, instead of controlling for the real estate score, controls for Council Member fixed effects. Column 5 additionally controls for contemporaneous economic factors measured annually for each Council District: the median household income in the Council District, the unemployment rate, and the percent of adults in the Council District with a bachelor's degree or higher.[26]

[Table 5 around here.]

The results show that there is a statistically significant and robust negative association between project starts and the backward-looking mean and volatility of approval time. Using specification 5, our preferred specification, we estimate that a 25% reduction in the mean of approval times increases project starts by 11.4%, and a 25% reduction in the volatility of approval times increases project starts by 8.5%.[27] If both the mean and the standard deviation of approval times are reduced by 25%, the model suggests that project starts would increase by 20.9%. We call this the incentive effect because it operates through developers' decisions to start new projects

[25] A donor associated with the real estate industry is a donor whose name or occupation contains any of the words "architect", "developer", "building", or "construction". Raw data for campaign contributions comes from the Los Angeles City Ethics Commission website, `ethics.lacity.org` (accessed March 14th, 2024).

[26] These variables were constructed from the ACS 5-year estimates at the Census Tract level, then aggregated to the level of a Council District using geographic crosswalks.

[27] Our estimated effect of a 25% reduction in mean approval time is on the same order of magnitude as those presented in Casey et al. (2022). Casey et al. used a financial model to estimate that reducing approval times in Los Angeles by 25% would result in 9.8% additional housing units becoming financially feasible for development.

in response to lower expected approval times and reduced uncertainty in approval times.

Combined incentive and pull-forward effects. To combine the incentive effect with the pull-forward effect, we apply the incentive effect to the number of units started at any given point in time, assuming the distribution of project start times remains the same as in the data. With this assumption, the combined incentive and pull-forward effect can be calculated as $z = (1 + x)(1 + y) - 1$, where x is the percent gain in project starts due to the incentive effect and y is the percent gain in project completions due to the pull-forward effect. z is the percent gain in the number of units that would have been completed between January 2010 and November 2022.

Our estimated incentive effect of a 25% reduction in approval time (both mean and volatility) was 20.9% and our estimated pull-forward effect was 11.9%. We estimate a combined effect of 33.0%. In other words, we estimate that if the mean and volatility of approval times were both reduced by 25%, then 33.0% additional units would have been completed between January 2010 and November 2022, relative to the baseline number of 71,532, or 23,634 additional units. Figure 7 illustrates the combined incentive and pull-forward effects by plotting their effects on starts and completions over the data period.

[Figure 7 around here.]

Figure 8 shows the response spectrum for the combined pull-forward and incentive effects. If approval times were reduced by an ambitious 50%, housing production would have increased by 73.6% taking into account both the effect on project starts and the effect on project completions. If approval times were reduced by a more modest 10%, we estimate that housing production would have increased by 13.3%.

Discussion. Our estimates show that reductions in approval time can have a large effect on housing production. As a lower bound, we estimate that a 25% reduction in approval time would increase the rate of housing production by 11.9%, simply by pulling forward in time the completion of already started projects. When we additionally take into account the incentive effect, we estimate that a 25% reduction in approval time would increase the total rate of housing production by 33.0%.

We are agnostic as to whether or not a 25% reduction in approval time is a realistic objective for the City of Los Angeles. We leave that question to city officials and agency heads to figure out. We do note, however, that our estimated combined effect of 33.0% is a partial equilibrium estimate.

Among limitations of our counterfactual exercise, note that we do not consider the behavioral response of opponents to housing development. For example, even if the city was able to reform its processes so that approval times were reduced by 25%, opponents to development may find other means to extend, delay, or even cancel projects during the construction phase. Also, the counterfactual exercise does not consider second order general equilibrium effects operating through prices and resource allocations in the broader economy. Finally, the simulations use data on projects permitted between January 2010 and November 2022. These projects were developed under specific macroeconomic conditions that may not be reflected going forward.

7 Conclusion

Our analysis is among the first to theoretically model and empirically evaluate the role of local land-use regulatory approval times and uncertainty in determination of housing production. Using a unique dataset on the development timelines for all multifamily housing projects permitted in Los Angeles from 2010 to 2022, we estimate that if mean duration and volatility in approval times were reduced by 25%, the rate of housing production would increase by 11.9%, due simply to the pulling forward in time of projects that were already started. When we additionally account for the effect of incentivizing new development, we find that same 25% reduction in approval time would increase the rate of housing production by a full 33.0%.

Our paper provides robust and credible quantitative evidence that approval policy is a significant driver of the rate of housing production. The effect of approval policy was shown to be quite large, even when using a conservative estimate of just the pull-forward channel. Our paper adds novel evidence that local approval processes are a significant driver of housing supply and reinforces the notion that municipal regulatory reform is an important component of housing reform.

References

Albouy, David and Gabriel Ehrlich, "Housing productivity and the social cost of land-use restrictions," *Journal of Urban Economics*, 2018, *107.*

Ben-Moshe, Dan and David Genesove, "Regulation and Frontier Housing Supply," Working Paper 2022.

Brueckner, Jan K. and Ruchi Singh, "Stringency of land-use regulation: Building heights in U.S. cities," *Journal of Urban Economics*, 2020, *116.*

Casey, Alexander, Ben Metcalf, David Garcia, Haley Rebecca Tiu, and Issi Romem, "Modeling New Housing Supply in Los Angeles: Simulations from the Terner Housing Policy Dashboard," Policy Brief, Terner Center for Housing Innovation at UC Berkeley 2022.

Gabriel, Stuart and Gary Painter, "Why affordability matters," *Regional Science and Urban Economics*, 2020, *80.*

Glaeser, Edward and Bryce A. Ward, "The causes and consequences of land use regulation: Evidence from Greater Boston," *Journal of Urban Economics*, 2009, *65* (3), 265–278.

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

Gyourko, Joseph 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*, 2021, *126.*

_ **and Raven Molloy**, "Regulation and Housing Supply," in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, 2015, pp. 1289–1337.

_ **, 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*, 2021, *124.*

Hilber, Christian A. and Wouter Vermeulen, "The Impact of Supply Constraints on House Prices in England," *The Economic Journal*, 2016, *126* (591), 358–405.

Hsieh, Chang-Tai and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 1–39.

Jackson, Kristoffer, "Do land use regulations stifle residential development? Evidence from California cities," *Journal of Urban Economics*, 2016, *91.*

Kalbfleisch, John D. and Ross L. Prentice, *The Statistical Analysis of Failure Time Data*, Hoboken, NJ: John Wiley & Sons, Inc., 2002.

Quigley, John M. and Steven Raphael, "Regulation and the High Cost of Housing in California," *American Economic Review*, 2005, *95* (2), 323–328.

Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw, "Land Use Regulation and Welfare," *Econometrica*, 2014, *85* (4), 1341–1403.

Table 1: Summary of the Data - Project and Unit Counts

		All Projects	Market-Rate	Mixed-Income	100%-Affordable
# Projects	Total	2,677	1,681	701	295
	Completed	1,712	1,192	351	169
	Not Completed	965	489	350	126
# Dwelling Units	Total	120,213	70,272	36,269	13,672
	Completed	71,532	47,904	15,929	7,699
	Not Completed	48,681	22,368	20,340	5,973
# Market-Rate Units	Total	102,897	70,272	32,422	203
	Completed	62,493	47,904	14,449	140
	Not Completed	40,404	22,368	17,973	63
# Affordable Units	Total	17,316	0	3,847	13,469
	Completed	9,039	0	1,480	7,559
	Not Completed	8,277	0	2,367	5,910

Notes: Multi-family housing development projects issued a new building permit by DBS between January 2010 and November 2022. "Mixed-Income" refers to projects that include both market-rate and income-restricted units. "Complete" indicates that the project was issued a Certificate of Occupancy by November 28th 2022.

Table 2: Variable Names and Definitions

Variable Name	Definition
AT	Approval time in days
CT	Construction time in days
UNITS100	Number of units in project divided by 100
HEIGHT10	Project height in feet divided by 10
SQFT10K	Square footage of project divied by 10,000
MIXEDINCOME	Project is a mixed-income project
AFFORDABLE	Project is a 100%-affordable project
BY_RIGHT	Project did not require any entitlements
CPC	Project required review by City Planning Commission
CE	Project had a Categorical Exemption to CEQA requirements
MND	Project adopted a Negative Declaration or Mitigated Negative Declaration for CEQA
EIR	Project required Environmental Impact Report for CEQA
SPR	Project required Site Plan Review
SPP	Project required Specific Plan Permit Compliance
ZAA	Project required Area/Height/Yard/Bldg line adjustments
ZV	Project required a Zone Variance
CPIOC	Project required Community Plan Implementation Overlay Clearance
OVR	Project required Overlay Review
DB	Project requested Density Bonus
POWER_OH	Project required new overhead circuit installation
POWER_UG	Project required new underground circuit installation
CDX	Whether the project is in Council District X

Table 3: Selected Summary Statistics for Variables

	All Projects	Market-Rate	Mixed-Income	100%-Affordable
Approval Time (days)				
25th pctile	308	274	449	250
Median	524	469	735	491
Mean	652	588	838	577
75th pctile	872	769	1,100	818
Construction Time (days)				
25th pctile	574	534	659	622
Median	791	804	864	692
Mean	863	862	917	757
75th pctile	1,068	1,068	1,124	839
# Units, mean	44.9	41.8	51.7	46.3
Height (ft), mean	53.2	51.9	57.1	50.7
Square Footage, mean	49,811	49,820	53,326	41,405
By Right, mean	0.481	0.609	0.211	0.397
CPC, mean	0.066	0.036	0.098	0.156
CE, mean	0.272	0.199	0.459	0.241
MND, mean	0.168	0.124	0.264	0.193
EIR, mean	0.010	0.008	0.014	0.017
SPR, mean	0.112	0.079	0.163	0.186
SPP, mean	0.103	0.108	0.096	0.095
ZAA, mean	0.043	0.049	0.029	0.041
ZV, mean	0.037	0.046	0.014	0.034
CPIOC, mean	0.027	0.012	0.026	0.108
OVR, mean	0.020	0.024	0.013	0.010
DB, mean	0.163	0.036	0.422	0.271
POWER_OH, mean	0.271	0.231	0.358	0.292
POWER_UG, mean	0.365	0.286	0.536	0.407

Notes: Approval time is measured as the number of days from the date of the first seen entitlement or permit application to the date of permit issuance. Construction time is measured as the number of days from the date of permit issuance to the date of CofO issuance. Construction time is only measured for completed projects. Approval time is measured for all projects. All other variables are measured for all projects.

Table 4: Approval and Construction Time Model Coefficients

	AT	CT
UNITS100	-29.333	-84.513^{*}
	(24.496)	(49.798)
HEIGHT10	3.978^{*}	12.137^{***}
	(2.248)	(4.391)
SQFT10K	1.301	7.981^{*}
	(1.950)	(4.240)
AFFORDABLE	-106.183^{***}	-197.394^{***}
	(24.370)	(43.349)
MIXEDINCOME	101.500^{***}	42.589
	(19.790)	(35.920)
BY_RIGHT	-197.187^{***}	12.072
	(27.850)	(48.479)
CPC	192.589^{***}	124.885^{**}
	(36.169)	(61.694)
CE	-17.300	116.733^{***}
	(25.421)	(44.197)
MND	11.461	-72.193
	(29.350)	(48.040)
EIR	504.227^{***}	-75.983
	(85.306)	(141.370)
SPR	105.619^{***}	35.944
	(30.562)	(53.060)
SPP	125.525^{***}	-9.338
	(26.856)	(46.681)
ZAA	222.313^{***}	81.616
	(43.121)	(69.488)
ZV	84.277^{*}	157.718^{**}
	(46.821)	(75.233)
CPIOC	-77.401	530.381^{***}
	(48.929)	(138.028)
OVR	-5.876	393.497^{***}
	(54.528)	(102.124)
DB	63.089^{**}	-34.955
	(26.313)	(44.245)
POWER_OH	22.749	116.687^{***}
	(16.509)	(29.617)
POWER_UG	69.005^{***}	175.623^{***}
	(16.317)	(28.692)
Constant	676.691^{***}	966.533^{***}
	(41.982)	(73.038)
Council District FE	Y	Y
Scale Parameter	208.595	331.675
Observations	2,677	2,677
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 5: Regression of Project Starts on Mean and Volatility of Approval Time

	Dependent variable:				
	log(Project Starts)				
	(1)	(2)	(3)	(4)	(5)
Mean of ln(Approval Time)	-0.167	-0.474^{***}	-0.396^{**}	-0.441^{**}	-0.431^{**}
	(0.141)	(0.184)	(0.192)	(0.191)	(0.179)
S.D. of ln(Approval Time)	-0.178	-0.445^{***}	-0.447^{***}	-0.475^{***}	-0.327^{***}
	(0.189)	(0.0879)	(0.0876)	(0.0928)	(0.0833)
Council Member Real Estate Score			0.138^{**}		0.0748
			(0.0699)		(0.0662)
ln(Median Household Income)					-1.19^{**}
					(0.559)
Umemployment Rate					-0.11^{***}
					(0.0334)
% Bachelors or Higher					0.0759^{***}
					(0.0188)
Council District FE		X	X	X	X
Year FE		X	X	X	X
Council Member FE				X	
Observations	1,260	1,260	1,260	1,260	1,260

Note: *p<0.1; **p<0.05; ***p<0.01

Figure 2: Histogram of Development Time for Completed Projects (Days)

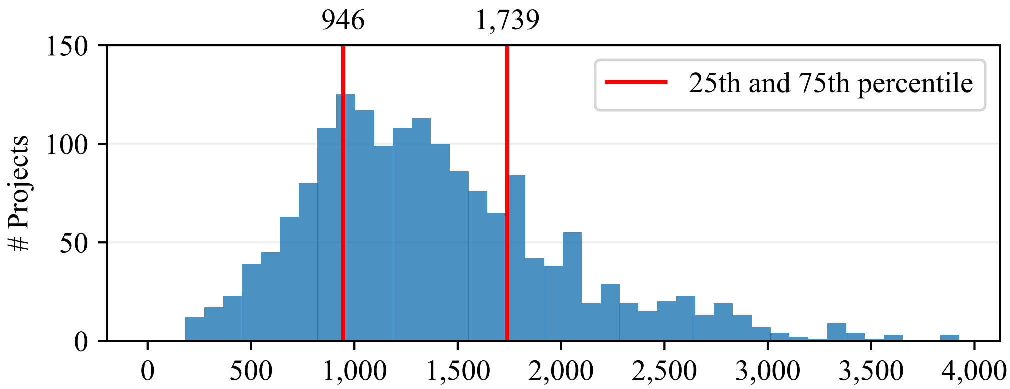

Note: Development time is measured as the number of days from the date of the first seen entitlement or permit application to the date the Certificate of Occupancy was issued.

Figure 3: Units Started vs. Units Completed

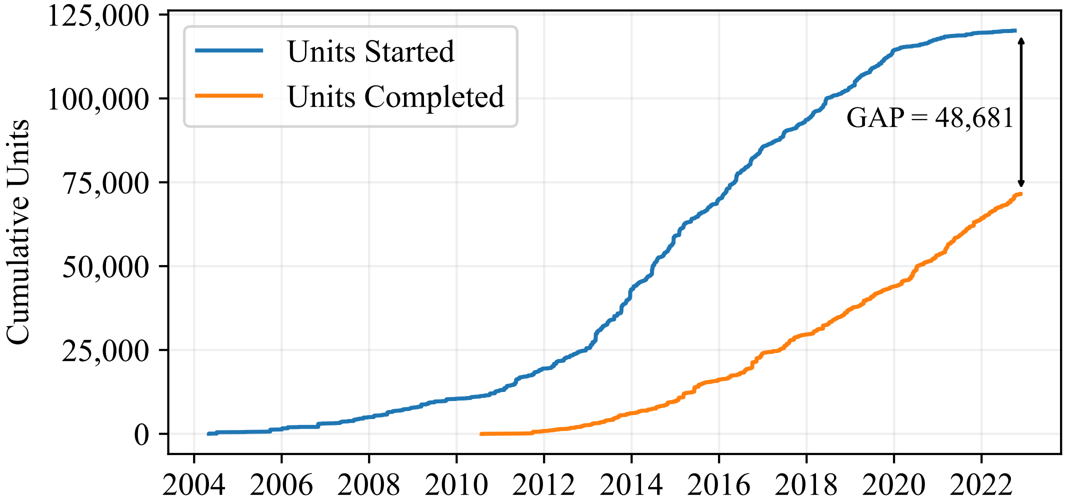

Note: Start date is measured as the date of the first observed entitlement or permit application. Completion date is the issuance date of the Certificate of Occupancy. Not all projects started in the dataset were completed as of November 28th, 2022.

Figure 4: Salience of Development-Related Factors

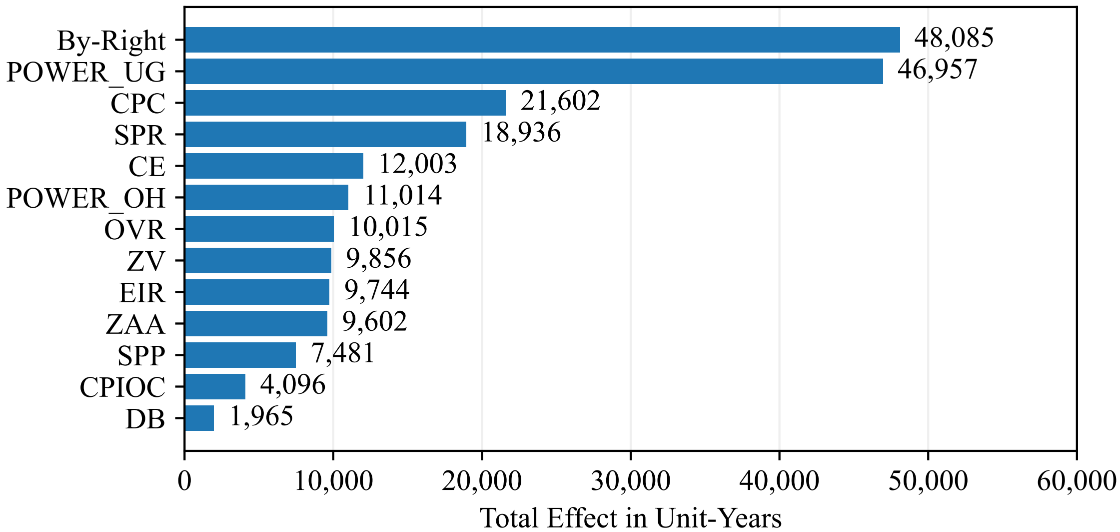

Note: Salience is measured in unit-years. It represents the number of dwelling units that would have been developed one year sooner if the marginal effect of the factor had been eliminated.

Figure 5: Pull-Forward Effect of a 25% Reduction in Approval Time

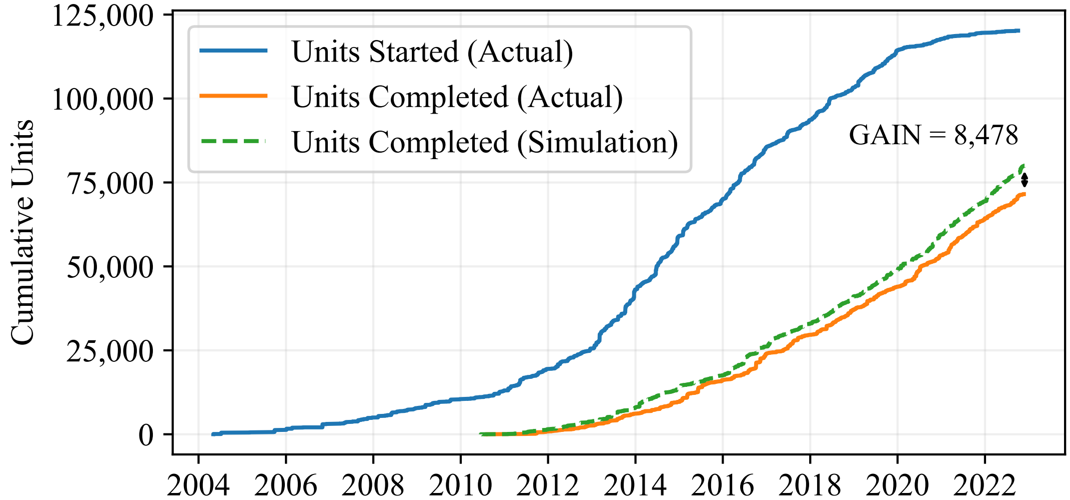

Note: Start date is measured as the date of the first observed entitlement or permit application. Completion date is the issuance date of the Certificate of Occupancy. Not all projects started in the dataset were completed as of November 28th, 2022.

Figure 6: Pull-Forward Effects Due to Reductions in Approval Time

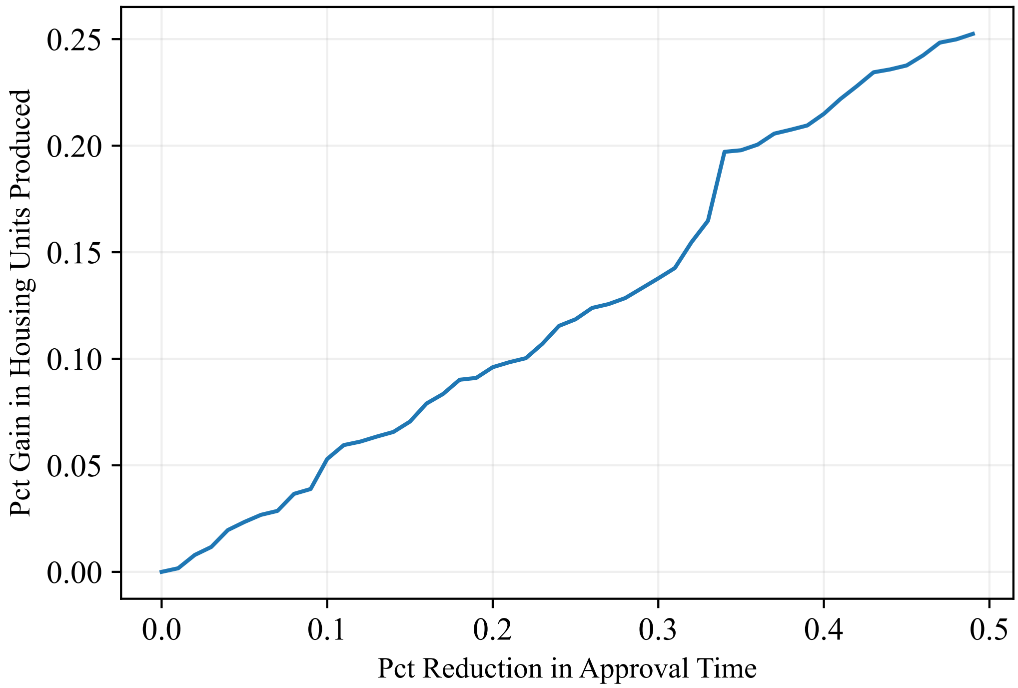

Note: This chart shows the number of additional units that would have been produced between January 2010 and November 2022, taking as given the projects that were started, if approval times had been reduced by X percent. The counterfactual only takes into account the effect of approval time reductions on project completions, i.e. the "pull-forward" effect.

Figure 7: Pull-Forward and Incentive Effects of a 25% Reduction in Approval Time

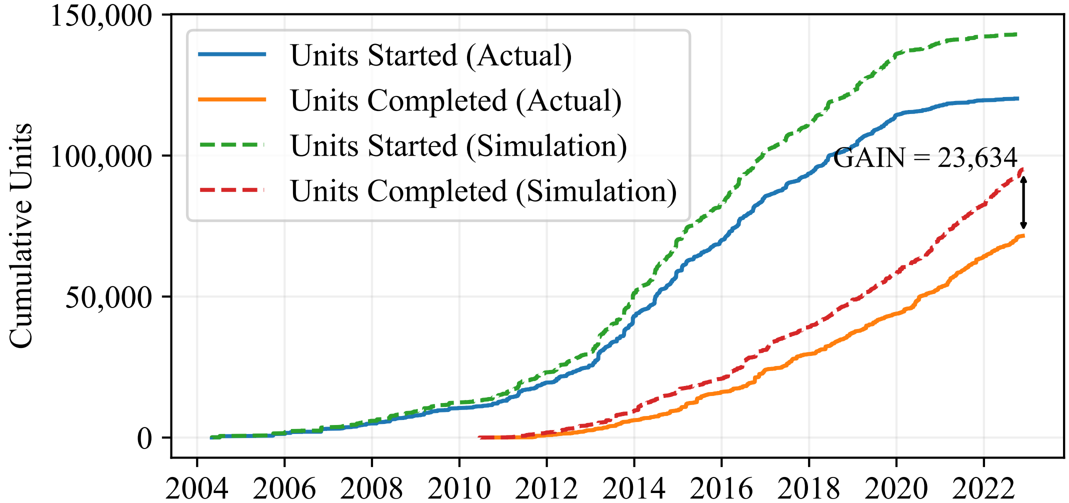

Note: Start date is measured as the date of the first observed entitlement or permit application. Completion date is the issuance date of the Certificate of Occupancy. Not all projects started in the dataset were completed as of November 28th, 2022.

Figure 8: Combined Incentive and Pull-Forward Effects Due to Reductions in Approval Time

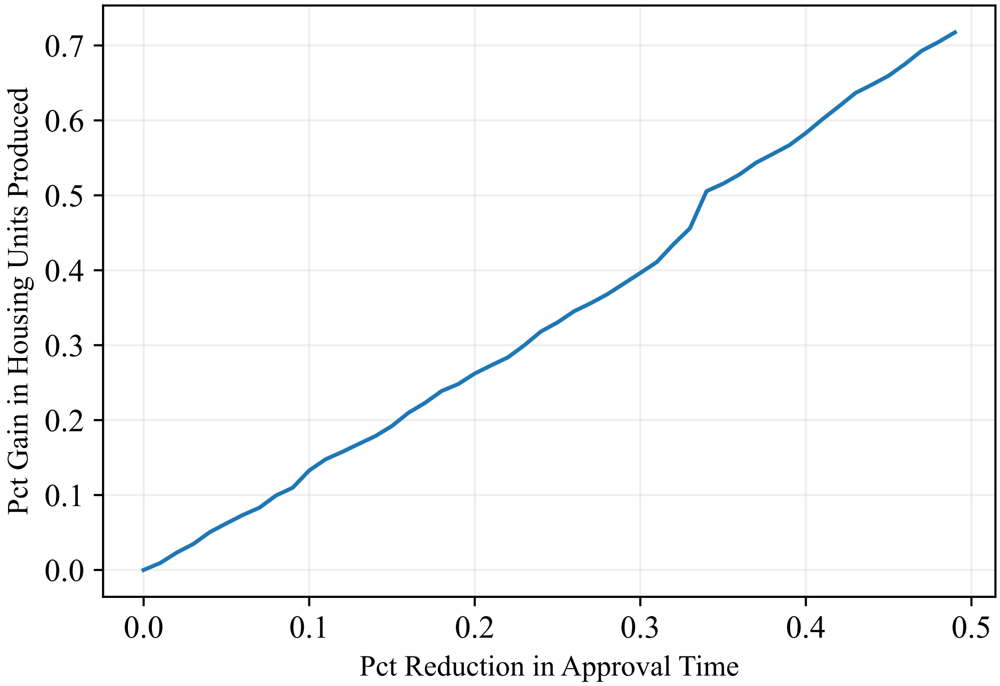

Note: This chart shows the number of additional units that would have been produced between January 2010 and November 2022, if approval times had been reduced by X percent. The counterfactual takes into account both the effect on project completions (the "pull-forward" effect) and the effect on project starts (the "incentive" effect).

A Mathematical Appendix

If the time it takes to develop a housing unit is x, then the present value of the development is:

$$\begin{aligned} PV &= -\int_0^x \zeta e^{-rt} dt + pe^{-rx} \\ &= -\tfrac{\zeta}{r}\left[1 - e^{-rx}\right] + pe^{-rx} \\ &= -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) e^{-rx} \end{aligned}$$

The expected value of PV is:

$$\begin{aligned} E[PV] &= -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) E\left[e^{-rx}\right] \\ &= -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) m(-r) \\ &= -\tfrac{\zeta}{r} + \left(p + \tfrac{\zeta}{r}\right) \left[1 - E[X]r + o(r^2)\right] \end{aligned}$$

The variance of PV is:

$$\begin{aligned} V[PV] &= \left(p + \tfrac{\zeta}{r}\right)^2 V\left[e^{-rx}\right] \\ &= \left(p + \tfrac{\zeta}{r}\right)^2 \left(E\left[\left(e^{-rx}\right)^2\right] - E\left[e^{-rx}\right]^2\right) \\ &= \left(p + \tfrac{\zeta}{r}\right)^2 \left(E\left[e^{-2rx}\right] - m(-r)^2\right) \\ &= \left(p + \tfrac{\zeta}{r}\right)^2 \left(E\left[e^{-2rx}\right] - m(-r)^2\right) \\ &= \left(p + \tfrac{\zeta}{r}\right)^2 \left(m(-2r) - m(-r)^2\right) \end{aligned}$$

Note also that:

$$m(-2r) = 1 - 2E[X]r + 2E[X^2]r^2 + o(r^3)$$

and:

$$
\begin{aligned}
m(-r)^2 &= \left(1 - E[X]r + \tfrac{1}{2}E[X^2]r^2 + o(r^3)\right)^2 \\
&= 1 - 2E[X]r + E[X^2]r^2 + E[X]^2r^2 + o(r^3)
\end{aligned}
$$

and thus:

$$
\begin{aligned}
m(-2r) - m(-r)^2 &= E[X^2]r^2 - E[X]^2r^2 + o(r^3) \\
&= r^2V[X] + o(r^3)
\end{aligned}
$$

B Data Appendix

Our goal was to develop a comprehensive dataset with detailed information on the timing, cost, and approvals for a large representative sample of housing development projects in the City of Los Angeles. After investigating publicly available sources and consulting with developers and city officials, we were unable to identify a sufficiently representative sample with reliable cost information inclusive of both affordable and market-rate projects.

Fortunately, we were able to develop a dataset with rich information on development times. We started with a list of all multi-family projects issued a new building permit by the L.A. Department of Building and Safety (DBS) between January 2010 and November 2022.[28] This data is publicly available on the city's open data portal.[29] We exclude closed and expired permits, so the final dataset includes only the projects that are still actively in development or have been completed. We measure project

[28] A project was determined to be multi-family if its permit subtype was not "1 or 2 family dwelling" and if the use description was residential (apartments, senior housing, etc.)

[29] `https://data.lacity.org/City-Infrastructure-Service-Requests/New-Building-Permits-2010-to-Present/46r2-n9vp`.

completion by the issuance of a Certificate of Occupancy (CofO). For each project, we determined the number of income-restricted units (i.e. affordable units) and the number of market-rate units by reading the permit work description and associated entitlement requests. A project with all affordable units except a few manager units were categorized as "100%-Affordable" projects. Projects with no affordable units were categorized as "Market-Rate" projects. All other projects were categorized as "Mixed-Income" projects.

The resulting dataset contained 2,677 projects representing 120,213 total units, of which 102,897 are market-rate and 17,316 are affordable. 1,712 projects were completed (issued a CofO) as of November 28th, 2022, and 965 projects were unfinished as of November 28th, 2022. The full breakdown of counts by project type are reported in Table 1 of the main paper.

Linking DBS permit data to entitlements

In that one of our primary goals was to assess the speed of housing production, it was important that we used a meaningful and consistent definition of project start date. The date of DBS permit submittal is not always a good measure of the start date because the entitlement process (the process of getting formal approval from city planners to commence development of a project) can start many years prior to submission of plans to the DBS. This is especially true of projects that seek to build outside the specifications of existing zoning code.

The DBS dataset does not include information on entitlements. In order to link the DBS data to related project entitlements, we took advantage of the L.A. Zone Information and Map Access System (ZIMAS) and Planning Document Information

System (PDIS).[30,31] ZIMAS allows users to input an address, assessor parcel number (APN), or parcel identification number (PIN) and retrieve zoning information about the parcel, including all relevant planning department cases. PDIS is a system maintained by the L.A. Planning Department that allows users to retrieve information about specific planning department cases based on case number.

To facilitate data collection, we built a tool that performs the following functions:

- First, the tool takes the primary address and the associated PINs of a project from the DBS data. It retrieves and stores all the data returned by ZIMAS for that address and the associated PINs. All planning department case numbers returned by ZIMAS are then linked to the project.

- For each planning case associated with the project, the tool retrieves and stores the data about that case from PDIS. The crucial fields that we make use of are the case filing date, the case completion date, and free text fields describing the project and the requested entitlement. The case number itself also encodes valuable information, such as the level of decision-making required (Administrative Review, Director of Planning, City Planning Commission, etc.) and the requested entitlements (Denstity Bonus, Site Plan Review, etc.)

Filtering relevant cases. Using this tool, we were able to link each project to a list of planning cases relevant to the parcels under development. However, not every such case is relevant to the project development timeline. Some cases, for example, are community plan updates that affect a large number of parcels but were not specifically requested by the developer of the project. Other planning cases may have been filed subsequent to completion of the project. Still others may have been filed many years

[30]ZIMAS: `https://zimas.lacity.org`.
[31]PDIS: `https://planning.lacity.org/pdiscaseinfo/search`.

before the project started and are unrelated to project development. To deal with such cases, we mark a case as not relevant to the project if any of the following hold:

- The case was filed or completed after the project's CofO date. This only applies to completed projects where a CofO was issued.
- The case applicant name was "City of Los Angeles" or if the case is linked to 5 or more separate projects.
- The case was filed more than 5 years before the DBS permit submittal date or completed more than 2 years before the DBS permit submittal date. In addition, we only kept cases filed before the DBS permit submittal date if the case's requested entitlement field contained keywords indicating it was a housing development project.

Using these rules, we linked 1,389 projects to at least one related entitlement case. 1,288 projects were not linked to any case. Of the projects that were linked to at least one case, the average number of cases linked was 2.7.

Discussion. Two types of errors may emerge from our linking procedure. First, cases not relevant to a project may be erroneously marked as relevant. Second, some cases relevant to the project may erroneously be marked as not relevant.

To assess the prevalence of the first type of error, we randomly sampled 30 projects linked to at least one planning case. This resulted in a sample of 30 projects linked to 75 cases. We then investigated each case by hand and assessed our level of confidence that it is indeed relevant to the project. We have a very high degree of confidence in 57 of the 75 linked cases. In these 57 cases, the project description in the entitlement case matches the work description in the DBS permit inclusive of the number of dwelling units proposed. In 16 of the 75 cases we had a medium degree of confidence. In these cases, most of the project features match, such as the number of building stories, but there may be some small discrepancies such as the proposed number of units. A discrepancy does not necessarily mean that the planning case and the permit don't refer to the same project. Discrepancies can arise because project plans may change from the time the entitlement was first requested and the new building permit was submitted. In most cases the discrepancies are quite minor. We had low confidence in only 2 of the 75 cases, both linked to the same project. The error may have been due to an erroneous address in the system.

To assess the prevalence of the second type of error, we randomly sampled another 30 projects with or without any linked planning cases. We then investigated each project by hand to see if any relevant cases were missed. There was one project that our procedure failed to link to any case because the relevant entitlements were filed 6 years before the permit submittal date. For every other project which our procedure failed to link to a case, we did not find any obviously relevant cases.

It is still possible that some cases our procedure marked as not relevant actually are pertinent to the study. For example, developers can push for general or specific plan amendments that affect a large number of parcels at once as a prerequisite to development of a specific project. Yet our procedure would not link the case to any one particular project. To the extent that we are missing such linkages, our procedure can be thought of as a conservative estimate of entitlement and development timelines, because we are measuring fewer entitlements for each project than were actually necessary.

Supplemental data from Department of Water and Power

In addition to data on entitlements, we also sought data on new power service installations from the Los Angeles Department of Water and Power (DWP). This was motivated by conversations with industry professionals, where long delays for electrification were identified as a major bottleneck.

To obtain this data, we provided DWP with the primary site address for each of our 2,677 projects. DWP then returned, for each address, all new power service work requests associated with that address from 2010 to present. We filtered out any work requests that were entered before the project start date and we filtered out any work requests that were completed after the project's CofO date.

In total, we received data on 70,114 work requests for 2,324 of our projects. For the projects that we did not receive data for, it is likely because of an error in finding matching addresses. Addresses are not standardized between the DBS and DWP databases, leading to an imperfect matching process. (This also suggests that improvements to the city's data systems could help accelerate housing development.)

The vast majority of the work requests are new meter installations. We ignore new meter installations in our analysis because they are required on almost every project. Instead, we focus on work requests for the design and installation of new overhead and underground circuits and service voltage. There were 808 overhead work requests and 1,291 underground work requests. In our empirical models, we consider the impact of a project requiring *any* new overhead or underground installations on project timelines.