Prices Versus Quantities Revisited: What Do Policymakers Need to Know to Set Pigouvian Taxes and Subsidies?

Denise DiPasquale* Edward Glaeser†

Adam M. Guren‡ Paul S. Willen§

June 2026

What information do policymakers need to design Pigouvian taxes or subsidies? Standard logic suggests that it is sufficient to know the size of the externality and unnecessary to know about quantities. Yet this logic is incorrect if interventions have fixed costs, taxes create deadweight losses, or there are distributional concerns. We present a model in which these considerations can make it more valuable for policymakers to learn about equilibrium quantities. We apply the model to congestion pricing, which has high fixed costs, and to a proposed housing subsidy in Boston that features deadweight losses and distributional concerns.

Acknowledgments: We thank Nick Ariniello, Patricia Belden, Deep Biswas, Michael A. Cantalupa, Arion Dako, Brian Golden, Jeffrey Goldstein, Reuben Kantor, Brian Kavoogian, Renee LeFevre, Abe Menzin, Young Park, Jerome L. Rappaport, and Matthew Resseger for providing data and comments. We thank Tim Gray, Stephanie Kestelman, and Jamie Murray for editorial assistance. We are responsible for the analysis and the conclusions provided in the report.

*Harvard University and City Research.

†Harvard University and NBER. Glaeser has given paid talks for real estate investment groups and received research support from Tishman/Speyer.

‡Boston University and NBER.

§Federal Reserve Bank of Boston. Email: Paul.Willen@bos.frb.org. The views expressed herein are solely those of the authors and should not be reported as representing the views of the Federal Reserve Bank of Boston, the principals of the Federal Reserve Board of Governors, or the Federal Reserve System.

1 INTRODUCTION

What information do policymakers need to decide how much to tax or subsidize externality-creating activities? Pigou (1920)'s classic result is that the tax or subsidy should equal the size of the externality, which implies, for instance, that policymakers considering a congestion charge need only know the potential harm caused by added driving, not the impact that such a charge would have on traffic. Yet economists have long known that when taxes are distortionary and distributional consequences are significant, the optimal tax or subsidy will differ from the externality. In these cases, it may be optimal to spend more time researching market conditions instead of the externality. In this paper, we develop a framework that provides conditions under which a policymaker will prefer to analyze the size of the externality or the quantity response to a tax, and we apply it to two real-world applications: congestion pricing and construction subsidies.

In Section 2, we present a simple model of research on prices and quantities that builds on Weitzman (1974)'s classic evaluation of whether government market interventions should use price or quantity instruments. We assume that a policymaker's only policy instrument is a Pigouvian tax or subsidy and that the policymaker is uncertain about both the size of the externality and the intercept of the supply and demand curves. The optimal subsidy is lower—and the optimal tax is higher—than the policymaker's estimate of the externality if there are deadweight losses from taxation or if the policymaker dislikes redistribution from consumers to producers.

We then provide conditions under which the policymaker would prefer to learn either the true externality or the equilibrium quantity in the absence of policy intervention.[1] When distortions from taxes are low and distributional concerns are small, the policymaker will prefer to eliminate uncertainty about the cost imposed by the externality. Conversely, when taxes are socially costly and the distributional consequences are significant, the policymaker

[1] In our framework, the impact of the policy is known, but the equilibrium quantity in the absence of the policy (or, equivalently, under the policy) is unknown. In practice, the policy's effect on quantity may also be unknown and worth learning about; therefore, we occasionally use language about the policy's "quantity impact." However, our model focuses on learning about equilibrium quantities, not equilibrium quantity changes under the policy. Intuitively, the quantity matters because it relates to the number of inframarginal goods that would receive a subsidy.

will prefer to clear up uncertainty about quantities. The benefit per dollar of a tax or subsidy depends on the change in quantity, so when taxes are accompanied by other costs, knowledge of quantities becomes more important. Our results generalize to the case of continuous investments in different types of information.

Our first application, in Section 3, is congestion pricing and its implementation in London and New York City. The 1964 Smeed Report, which initiated public dialogue on road charges in the United Kingdom, is almost perfectly Pigouvian in its focus on the magnitude of traffic externalities. However, congestion pricing departs from the Pigouvian framework because the fixed costs associated with operating these systems are extremely high. Consequently, adopting a road pricing system makes sense only if it significantly changes behavior. Over the decades, the policy discussion pivoted toward the impact of congestion pricing on traffic. We argue that this shift makes sense because the high fixed costs of a congestion pricing system require a focus on the quantity impact of the subsidy rather than the size of the externality. In New York, which began considering congestion pricing much later, the discussion has focused only on quantity.

Our second application, in Section 4, is the subsidization of new construction of market-rate multifamily rental housing in Boston. In this case, the deviations from standard assumptions are both distributional—the city government may value profits earned by developers less than services provided to poorer Bostonians—and efficiency-related—the city faces a fixed budget. We argue in this section that the external benefit from new housing is largely political and difficult to quantify but that the central question for the policymaker is quantitative: how many extra units will actually be built due to the subsidy. A property tax abatement is paid on every qualifying project, but only the *marginal* units—that is, those that would not have been built without the subsidy—deliver the policy's intended benefit. The effective cost of the subsidy per marginal unit is therefore the nominal subsidy multiplied by the ratio of total units to marginal units. Estimating this ratio is the central task of the empirical analysis. Specifically, our analysis examines the potential impact of a property tax abatement aimed at unlocking a backlog of proposed but not yet built properties in Boston in 2023, following rising interest rates and skyrocketing construction costs in the wake of the COVID-19 pandemic. As suggested by our theory, our analysis focuses on the quantitative impact of a subsidy and, in particular, on the ratio of total units to marginal units. We find this ratio to be large, implying a significant present value of the subsidy required to build a marginal unit. We conclude by arguing that subsidies for residential construction are likely to be more favorable from a cost–benefit perspective only if they redevelop blighted neighborhoods, improve local amenities, or generate additional tax receipts from neighboring buildings. In our framework, this reduces the distortion and induces policymakers to consider the size of the externality.

This paper should be viewed as a special case within the broader literature on optimal information acquisition by policymakers. Paciello and Wiederholt (2014) and Zhong (2022) are two excellent examples of papers in this literature. We see the value of our paper not so much in generating theoretically surprising results as in asking a very basic information-acquisition question about a significant class of policy interventions.

2 A FRAMEWORK FOR EVALUATING TAXES AND SUBSIDIES

2.A Model of Information Acquisition and Pigouvian Interventions

In this section, we present a model of information acquisition and Pigouvian interventions that analyzes the tradeoff between the returns to researching the magnitude of an externality and the returns to studying the quantity impact of an intervention. The externality could itself reflect other public policies such as regulations that impose costs on businesses and provide other local benefits. The model is meant to capture settings in which either policymakers or their research teams have limited resources for information gathering and need to focus on a subset of possible questions. While our model focuses on a pure choice between researching two topics, its logic extends seamlessly to a setting with simultaneous, costly investment in both types of research, as shown in Appendix B.

Traditional Pigouvian analysis pushes toward studying the externality alone, although Pigou (1920) himself recognized that his analysis held only if there were non-distortionary sources of funds:

> It follows that, under conditions of simple competition, for every industry in which the value of the marginal social net product is greater than that of the

> marginal private net product, there will be certain rates of bounty, the granting of which by the State would modify output in such a way as to make the value of the marginal social net product there more nearly equal to the value of the marginal social net product of resources in general, thus—provided that the funds for the bounty can be raised by a mere transfer that does not inflict any indirect injury on production—increasing the size of the national dividend and the sum of economic welfare; and there will be one rate of bounty, the granting of which would have the *optimum* effect in this respect.

Pigou argued that an optimal subsidy (or "bounty") will equalize "marginal private net product" and "marginal social net product," with the substantial proviso that "the funds for the bounty can be raised by a mere transfer that does not inflict any indirect injury on production." While Pigou produced no formal treatment of cases without non-distortionary taxation, numerous authors have remedied that lacuna (Sandmo 1975; Kaplow 2012). A literature now discusses whether Pigouvian environmental taxes offer a "double dividend" by both reducing externalities and allowing governments to reduce other distortionary taxes (for example, Goulder (2005)).

In his later *A Study in Public Finance*, Pigou (1928) emphasized that "administrative costs" and distributional concerns could also impact optimal Pigouvian taxes. He noted that "since different commodities are purchased in different proportions by rich and poor persons, no tax-bounty scheme could be worked in practice without modifying distribution" and that "these considerations would need to be taken into account." A substantial body of subsequent literature has explored how distributional concerns can affect optimal Pigouvian taxes (see, for example, Casler and Rafiqui (1993), Kaplow (2012), Pai and Strack (2022)).

In contrast to the bulk of this literature, we focus on which information best informs an optimal tax, not on the optimal tax itself. Our core question is whether it is more valuable to research the size of the externality or the impact of a subsidy on the quantity consumed. This question is in the tradition of the prices-versus-quantities question posed by Weitzman (1974), and we utilize his notation as much as possible. However, unlike in Weitzman's analysis, the policymaker in our setting does not have the option of setting either prices or quantities and can only tax or subsidize; the price-versus-quantity margin is primarily about information collection. Moreover, other sources of revenue are either distortionary or fixed, meaning the cost of a subsidy will exceed its nominal amount. This additional cost may also reflect the program's administrative expenses.

Our model has implications for subsidies and taxes, but there are asymmetries between the two cases. A Pigouvian subsidy may cost more than its nominal amount, due to both the administrative costs and the distortions associated with the taxes needed to fund it. A Pigouvian tax may yield less revenue than its nominal value due to administrative costs. Still, the value of that revenue may exceed its nominal value because it reduces distortionary taxation, as recognized by the "double dividend" literature. A cash-constrained government that cannot change other tax rates must accept the distributional consequences of a Pigouvian subsidy. If a Pigouvian tax generates revenue, that revenue can be spent to offset any adverse distributional consequences of the tax. Consequently, the classical Pigouvian assumptions are more likely to be strongly violated for subsidies than for taxes, which is why our language and applications focus on subsidies.

The distributional concern in our context is not between consumers, as discussed by Pigou and most of the literature, but between consumers and producers. Unlike Kaplow (2012), we assume that the policymaker has no other tools with which to undo the distributional consequences of the Pigouvian tax. Unlike Baumol (1972), we do not allow the decisionmaker to learn by iterating among policies.

2.B Set-up

We consider a static model with the following timing. First, the policymaker chooses what information to acquire. Second, the policymaker receives signals about market conditions and the size of the externality. Third, the policymaker chooses a policy. Fourth and finally, all information is revealed, and the market clears. Given this timing, and following Weitzman (1974), the policymaker chooses what information to acquire to maximize the expectation of social welfare, W:

$$W = B^I(q, \eta_I) + B^P(q, \eta_P) - (p + (1 + \lambda)s)q + \delta((p + s)q - C(q, \theta)) - KI_{s \neq 0}. \qquad (2.1)$$

Social welfare is the sum of the consumer surplus and producer surplus. The consumer surplus is the sum of private and external benefits, net of consumer and administrative costs and deadweight losses. Private benefits $B^I(q, \eta_I)$ are a function of the quantity produced q and a demand parameter η_I. External benefits $B^P(q, \eta_P)$ are a function of quantity and a demand parameter η_P. The cost of the good, including the subsidy, is $q(p + s)$, and $\lambda \geq 0$ reflects the combined administrative costs and deadweight losses of taxation per dollar of subsidy sq. There is a welfare weight δ on the producer surplus relative to the consumer surplus, which represents distributional concerns. The producer surplus is equal to the cost of the good, including the subsidy, net of the cost of production $C(q, \theta)$, where θ is a productivity parameter. Finally, the fixed cost of implementing the system is K, and $I_{s \neq 0}$ is an indicator function that equals zero when $s = 0$ and one otherwise.

The divergence between public and private benefits could reflect other public policies, especially in housing markets. If a city government imposes a requirement on builders to produce local public goods, such as parks or housing for lower-income residents, then these activities presumably generate social benefits. There is no reason to think that the builder will directly internalize those benefits.[2] We assume that these policies cannot themselves be changed but can be offset by a subsidy to the builders.

Using the convention that $F_k(X)$ refers to the derivative of function $F(\cdot)$ with respect to its kth argument in the vector X, the market equilibrium always satisfies

$$B_1^I(q, \eta_I) + s = p + s = C_1(q, \theta). \tag{2.2}$$

We define the elasticity of quantity with respect to the subsidy, which is typically an equilibrium object rather than a fixed parameter, as $\varepsilon_{q,s} = \frac{s}{q} \cdot \frac{dq}{ds}$, which will be positive for a subsidy and negative for a tax. If the policymaker had complete information and had already

[2]For example, assume that in the absence of such policies, the development community earned profits of $pq - C_{Real}(q, \theta)$ but that developers also had to pay ωq in a regulatory tax and that this tax generated $z(\omega q)$ community benefits. In that case, the social planner would choose s to maximize $B^I(q, \eta_I) + z(\omega q) - (p + (1 + \lambda)s)q + \delta((p + s - \omega)q - C_{Real}(q, \theta))$. This problem is mathematically isomorphic to our problem if $B^P(q, \eta_P) = z(\omega q)$ and $C_{Real}(q, \theta) + \omega q = C(q, \theta)$.

paid the fixed cost needed to have a nonzero s, the optimal subsidy, denoted s^*, satisfies

$$s^* = \frac{B_1^P(q, \eta_P)}{1 + \lambda + \frac{1}{\varepsilon_{q,s}} \left(\lambda + (1-\delta)\left(1 + \frac{dp}{ds}\right)\right)}. \tag{2.3}$$

Both distributional concerns ($\delta < 1$) and deadweight losses from taxation ($\lambda \neq 0$) reduce the optimal subsidy, and either implies that the optimal subsidy is smaller if quantity is less elastic ($\varepsilon_{q,s}$ is smaller). Inelastic supply effectively means an increase in the relative size of inframarginal consumption to marginal consumption, and both distributional concerns and deadweight losses make it socially costly to transfer benefits to inframarginal producers and consumers.

If $\lambda = 0$ (no costs of taxation) and $\delta = 1$ (no distributional concerns), then (2.3) implies that $s^* = B_1^P(q, \eta_P)$, and we return to the classic Pigouvian result that the subsidy (or tax) equals the marginal social costs of the externality. In that case, the formula implies that knowing quantity responses is unnecessary to determine the optimal subsidy level. Knowing the quantity responses is still valuable when deciding whether to pay the fixed cost, K, to implement a positive subsidy.

If $\lambda > 0$ (taxation is costly) and $\delta = 1$ (no distributional concerns), then $s^* = \frac{B_1^P(q,\eta_P)}{1+\lambda+\frac{\lambda}{\varepsilon_{q,s}}}$.[3] If $\lambda = 0$ and $\delta < 1$, then $s^* = \frac{B_1^P(q,\eta_P)}{1+\frac{1-\delta}{\varepsilon_{q,s}}\left(1+\frac{dp}{ds}\right)}$. In both cases, the subsidy is below the Pigouvian benchmark and depends on both the externality and $\varepsilon_{q,s}$. Distributional concerns also make it valuable to understand how much of the subsidy will be embedded in the price.

2.C The Social Returns to Better Information

We now turn to the benefits of improved information on the externality's size and the subsidy's impact on quantity. Like Weitzman (1974), we prove our primary results using second-order approximations. We let $\hat{s}$ represent the subsidy that maximizes the unconditional expectation of welfare and $\hat{q}$ represent the market equilibrium quantity associated with that level of subsidy. Following Weitzman's language, we let the symbol $\cong$ denote an "accurate

[3]If $\lambda > 0$, $\delta = 1$ and $B_1^P(q, \eta_P) < 0$, so that $s < 0$ and there is a tax, then $\frac{B_1^P(q,\eta_P)}{1+\lambda+\frac{\lambda}{\varepsilon_{q,s}}} < B_1^P(q, \eta_P)$ if and only if $\lambda + \frac{\lambda}{\varepsilon_{q,s}} < 0$ or, equivalently, $\varepsilon_{q,s} > -1$. The tax increases with λ if and only if $-sq$ (total tax revenues) increases as s increases.

local approximation," which means that he and we work with quadratic approximations. We assume that

$$\begin{aligned} B^I(q,\eta_I) &\cong b^I(\eta_I) + (B_1^I + \beta^I(\eta_I))(q-\widehat{q}) + 0.5B_2^I(q-\widehat{q})^2 \\ B^P(q,\eta_P) &\cong b^P(\eta_P) + (B_1^P + \beta^P(\eta_P))(q-\widehat{q}) \\ C(q,\theta) &\cong a(\theta) + (C_1 + \alpha(\theta))(q-\widehat{q}) + 0.5C_2(q-\widehat{q})^2, \end{aligned} \tag{2.4}$$

where B_1^I, C_1, B_1^P, B_2^I, and C_2 are all known constants. The other terms are functions of the demand and productivity parameters that the policymaker imperfectly observes.[4]

Again following Weitzman, these constants are arbitrarily close to $E[B_1^I(\hat{q},\eta_I)]$, $E[C_1(\hat{q},\theta)]$, $E[B_1^P(\hat{q},\eta_P)]$, $B_{1,1}^I(\hat{q},\eta_I)$, and $C_{1,1}(\hat{q},\theta)$, respectively. We also assume that $C_2 > 0 > B_2^I$, and the unconditional expected values of $\beta^I(\eta_I)$, $\beta^P(\eta_P)$, and $\alpha(\theta)$ are all zero. Moreover, we assume that the market exists without a subsidy.

Note that while we use quadratic approximations to the private benefit B^I and cost function C, we have only a linear approximation to the external benefit B^P. This asymmetry exists because it is difficult enough to estimate the size of the externality, let alone its concavity. However, concavity in the private benefit and cost functions is necessary for a downward-sloping demand curve and upward-sloping supply curve, respectively.[5]

In Appendix A, we show that while $\beta^I(\eta_I)$ and $\alpha(\theta)$ have independent effects on baseline welfare, these terms only interact with the welfare impact of subsidies (and hence optimal subsidies) through their difference, $\beta^I(\eta_I) - \alpha(\theta)$ (that is, the gap between the intercepts of the demand curve and the supply curve). Hence, we define a new variable $\beta^T(\eta_T) = \beta^I(\eta_I) - \alpha(\theta)$, which can be interpreted as a shock to market quantity as opposed to $\beta^P(\eta_P)$, which is the shock to the externality or the optimal subsidy under canonical Pigouvian conditions.

We assume that the policymaker observes $\widehat{\beta}^T(\eta_T) = \beta^T(\eta_T) + \varepsilon_T$ and $\widehat{\beta}^P(\eta_P) = \beta^P(\eta_P) + \varepsilon_P$, and forms beliefs about $\beta^T(\eta_T)$ and $\beta^P(\eta_P)$ using the Bayesian normal signal extraction

[4]We have changed Weitzman's notation slightly because we find that using apostrophes in fixed parameters is confusing.

[5]We recognize that we are assuming away another reason that knowing the quantity response is important; if the per-unit externality were a function of the quantity of units, then it would be important to know the number of units after the tax or subsidy.

formula. The posterior beliefs about these error terms, denoted $\widetilde{\beta}^j(\eta_j)$ for $j = T, P$, satisfy $\widetilde{\beta}^j(\eta_j) = \frac{Var(\beta^j(\eta_j))}{Var(\beta^j(\eta_j))+Var(\varepsilon_j)}(\beta^j(\eta_j) + \varepsilon_j)$.[6] We use the notation q_{NS} for the quantity without subsidy, Δ_P for the change in price due to the subsidy, and Δ_Q for the change in quantity due to the subsidy. The ratio $\frac{\Delta_P}{\Delta_Q}$ equals $\frac{C_2}{C_2-B_2^I}$ and is a constant.

Proposition 1. *The Determinants of Welfare*

(i) For a subsidy of fixed size s, social welfare is higher with a subsidy than without it if and only if $B_1^P + \beta^P(\eta_P) > 0.5s + \lambda s\left(\frac{q_{NS}+\Delta_Q}{\Delta_Q}\right) + (1-\delta)(s+\Delta_P)\left(\frac{q_{NS}+0.5\Delta_Q}{\Delta_Q}\right) + \frac{K}{\Delta_Q}$.

(ii) If the social planner always implements a subsidy, then the unconditional expectation of social welfare is always declining with $Var(\varepsilon_P)$ *and is independent of* $Var(\varepsilon_T)$ *if* $\lambda = 0$ *and* $\delta = 1$*. If* $\lambda B_2^I \neq (1-\delta+\lambda)C_2$*, then social welfare is always declining in* $Var(\varepsilon_T)$*.*

All propositions are proved in Appendix A.

A subsidy of size s is better than no subsidy at all if and only if the external benefits from more production are greater than one-half of the subsidy plus the loss from the social cost of taxation per unit created, $\lambda s(\frac{q_{NS}+\Delta_Q}{\Delta_Q})$, plus the loss from redistributing from consumers to producers, $(1-\delta)(s+\Delta_P)\left(\frac{q_{NS}+0.5\Delta_Q}{\Delta_Q}\right)$, plus the fixed cost divided by the change in quantity, $\frac{K}{\Delta_Q}$. The effective tax cost per unit of housing created is proportional to $\frac{q_{NS}+\Delta_Q}{\Delta_Q}$ because this is the ratio of the number of units that receive the subsidy to the number of new units created. This ratio plays an important role in our work on housing in Boston.

Proposition 2. *Optimal Information*

If the social planner always implements a subsidy, there exists a value of μ*, denoted* μ^**, at which the decisionmaker is indifferent between learning more about the externality and learning more about market quantity. The decisionmaker prefers learning about the externality if and only if* $\mu > \mu^*$*. The value of* μ^* *is increasing in* λ, $\frac{C_2}{C_2-B_2^I}$ *and* ρ_P *and decreasing with* δ *and* ρ_T*.*

In this setting, the more uncertainty there is about a variable, the more valuable it is to learn its value. Consequently, focusing on the size of the externality is more important

[6] While these are the standard normal signal extraction formulas, we are not assuming that these noise terms are normal, and in particular, we assume that both the cost and private benefit of the good are always positive and that a positive quantity of consumption is optimal and occurs in equilibrium.

when the variability in the size of the externality, which is determined by μ, is larger. There is a maximum value of μ at which it is more desirable to research the market conditions of producer costs and consumer demand than the size of the externality. That maximum, and consequently the range of values at which research on market conditions is more important, is increasing in λ, the waste parameter, and decreasing in δ, since higher values of δ mean that distributional concerns are less important. The key implication of the model is that administrative costs and distributional concerns can increase the returns to research on quantities rather than on the size of the externality.

The composite term $\frac{C_2}{C_2 - B_2^I}$ determines the extent to which the subsidy becomes embedded in higher returns for producers, and that also makes it more important to research quantities. Finally, the parameters ρ_P and ρ_T determine the relative efficacy of researching the externality and researching the market. When it is hard to learn about the externality, it makes more sense to focus on the quantity impact and vice versa.

2.D Model Applications

This framework highlights that the returns to researching prices and quantities depend on the informational environment, administrative costs, and distributional concerns. We now turn to two real-world situations: congestion pricing and local building subsidies. Operational costs are particularly important for congestion pricing. Distributional issues and deadweight losses are central to local building subsidies. Consequently, in both of these settings, policymakers may prefer to learn about the quantity impact of a tax or subsidy rather than the Pigouvian externality.

The classical Pigouvian prescription for addressing externalities—that is, setting taxes or subsidies equal to marginal social damages or benefits—yields a counterintuitive result: The market's quantity response to the intervention is irrelevant for optimal policy design. This is surprising because subsidies are blunt instruments: A property tax abatement, for example, lowers taxes on every qualifying project, but only some of those projects would not have been built without the subsidy. The city pays for both groups, but only the latter delivers the policy's intended effect. The dollars going to the projects that would have been built anyway are simply a transfer from taxpayers to developers, and under the classical welfare assumptions of lump-sum redistribution and equal welfare weights, pure transfers wash out of the social welfare calculation. In other words, the "waste" on units that would have been produced anyway becomes immaterial to the optimal policy calculation. We examine how relaxing these assumptions fundamentally alters the information requirements for optimal policy design, making quantity impacts crucial.

2.D.1 Congestion Pricing

We now simplify the model to focus on the core element of our congestion pricing application: fixed implementation costs. Many of the costs of these systems involve the physical equipment needed to restrict access to fee-paying drivers or to monitor roads and charge drivers for entering. To focus on the K parameter, we assume that $C(q,\theta) = 0$, $B^P(q,\eta_P) < 0$, $s = -t$, and $\lambda = 0$, so social welfare is just $B^I(q,\eta_I) + B^P(q,\eta_P) - KI_{Cong=1}$, where $I_{Cong=1}$ is an indicator function that takes on a value of one if and only if a congestion price has been imposed. We also use the notation $t = -s$ to match the tax setting and assume that the only price is the tax. In this case, a corollary of Proposition 1 applies:

Corollary 1. *A congestion charge of* $t = -[B_1^P + \widetilde{\beta}^P(\eta_P)]$ *optimizes welfare conditional upon imposing a congestion charge, and imposing a congestion charge at this optimal level increases welfare, if and only if one-half multiplied by the expected change in driving multiplied by the expected size of the externality is greater than the administrative cost:* $\frac{(B_1^P + \widetilde{\beta}^P(\eta_P))^2}{-2B_2^I} > K$.

The logic of Corollary 1 is that the fixed cost is weighed against the net benefit from reducing drives. The net benefit is the change in the number of rides multiplied by the size of the externality minus a term that represents the loss in consumer surplus from reducing the number of rides. If the demand curves are linear and there are no income effects (thus Marshallian demand equals Hicksian demand), as in our case, the lost consumer surplus from the price change is one-half the change in price times the change in quantity, which is the familiar Harberger triangle. As the change in price equals the tax, which equals the expected magnitude of the externality, the net benefit is one-half the tax multiplied by the change in the number of riders. This result hinges upon the congestion tax accurately reflecting the size of the externality. If the true externality were much larger than the tax, a congestion tax may still increase welfare, even if implementation costs comprise a large share of the tax revenue.

The Weitzman notation does not capture traffic congestion perfectly because the externality works by reducing the benefits per drive. To show that this logic continues to hold if the externality determines the benefit by impacting the flow welfare, we assume that $B^P(q,\eta_P) = (B_1^P + \beta^P(\eta_P))(q-\hat{q})^2$, which multiplies an externality term $(B_1^P + \beta^P(\eta_P))(q-\widehat{q})$ by a term related to the number of drivers $(q-\widehat{q})$.

In this case, we assume that q is chosen treating the externality-related cost, $B^P(q,\eta_P)$, as an exogenous constant, so that

$$q - \widehat{q} = \frac{B_1^I + \beta^I(\eta_I) - t + (B_1^P + \beta^P(\eta_P))(q-\widehat{q})}{-B_2^I} = \frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I - B_1^P - \beta^P(\eta_P)}. \tag{2.5}$$

With these assumptions and the assumptions of Corollary 1 ($C(q,\theta) = 0$, $B^P(q,\eta_P) < 0$, and $\lambda = 0$), Corollary 2 follows:

Corollary 2. *If the congestion charge is chosen optimally, then social welfare with the charge will be greater than social welfare without the charge if and only if the administrative costs of the charge, K, are less than the charge multiplied by the expected change in ridership multiplied by a term bounded by* 0.5 *and one. If $\beta^P(\eta_P)$ is always zero, then the optimal charge is $\frac{B_1^P(B_1^I + E(\beta^I(\eta_I)))}{2B_1^P + B_2^I}$, and then social welfare with the charge will be greater than social welfare without the charge if and only if the administrative costs of the charge (K) are less than the charge multiplied by the expected change in ridership times $\frac{B_1^P + 0.5B_2^I}{B_1^P + B_2^I}$.*

Corollary 2 again shows our basic logic. A congestion price with a flat fee is justified only if the change in ridership, multiplied by the fee and a constant between one-half and one, exceeds the administrative cost. The term $\frac{B_1^P + 0.5B_2^I}{B_1^P + B_2^I}$ provides some intuition about when the term is closer to one and when it is closer to one-half. When B_2^I/B_1^P is close to infinity, which occurs when the externality is small or when the slope of the demand curve is large, then the term is close to one-half. When B_2^I/B_1^P is small, which occurs when the externality is large or when the slope of demand is small, then the term is closer to one.

This algebra suggests that acquiring information about the magnitude of the change in the quantity is preferable when administrative costs are high. In the absence of such costs or distributional concerns, the Pigouvian formula provides precisely the correct answer: Set the tax equal to the externality and do not worry about quantity changes. However, when those costs are substantial, forecasting ridership changes is crucial for determining whether implementing such a scheme is sensible.

Corollary 2 holds only for optimal taxes and linear externalities. If the congestion charge is set far below the actual externality, perhaps for distributional reasons, then this rule is no longer relevant, which may be the case in both New York and London today. To see how nonlinearities can also vitiate the rule, consider a scenario in which driving five miles incurs no externality but driving six miles generates an externality of $10. Assume further that a $1 congestion charge limits everyone's driving to five miles, but all of that money goes to implementation. In that case, the $1 charge is welfare-enhancing because of the externality's nonlinearity, even though it generates no additional revenue.

2.D.2 Distortions and Distributional Concerns in Subsidizing Residential Construction

Our second application concerns subsidies for new housing production. The social benefit of housing may exceed the private benefit of building because local land-use regulations create a significant gap between the cost of buying a home and the cost of building one (Glaeser et al. 2005). Some cities, including Boston, also require builders to deliver community benefits such as affordable housing for poorer residents. We assume that these other policies are immutable. Other possible social benefits of building include distributional benefits from lower housing prices and reduced gentrification. Glaeser et al. (2023) present a model in which gentrification creates social losses by eliminating idiosyncratic stores that yield greater utility to consumers than chain stores. These relatively amorphous social benefits depend strongly on policymakers' preferences, which is likely to make quantification particularly difficult.

Moreover, city governments are often extremely limited in their ability to raise taxes and have a keen interest in distributional concerns. Real estate taxes are among the most important sources of revenue for most city governments; indeed, in Boston, they account for 72 percent of the City's revenue.[7] Providing subsidies large enough to induce developers

[7] See `https://www.boston.gov/departments/budget/fy26-property-taxes`.

to build immediately will dramatically impact a city's budget, and this cannot be offset by non-distortionary lump-sum taxes. In the language of our framework, λ is therefore likely to be large.

Local politicians typically care more about taxpaying residents than about developers. In Boston, this preference is reflected in an Inclusionary Development Policy that requires 17 to 20 percent of all units in market-rate buildings to be set aside as rent-restricted affordable housing. The distributional concerns between consumers and producers correspond to a low δ in our framework.

The combination of distributional concerns, a limited ability to tax, and challenges in quantifying benefits implies that this is a setting in which identifying quantity responses is more valuable than measuring the size of the externality. In this case, the key is to distinguish between marginal units built solely because of the subsidy and inframarginal units that would have been constructed in the absence of the subsidy. The ratio of total units to marginal units produced appears in our formulae in Proposition 1, and its importance there motivates our empirical estimates of construction's response to different levels of tax subsidies in Boston.

3 CONGESTION PRICING

Pigou (1920) himself used congestion as an example of an externality.[8] William Vickrey advocated congestion pricing for subways (Vickrey 1955) and roads (Vickrey 1959, 1963). He was particularly concerned with reducing the congestion pricing system's high administrative costs, which we document in this section. We show that the focus of the public debate about congestion pricing in the United Kingdom shifted, as our model suggests it should, from the size of the externality to the change in traffic associated with the congestion charge.

[8] He discussed two hypothetical roads (B and C) and noted that if "it would be possible, by shifting a few carts from route B to route C, greatly to lessen the trouble of driving those still left on B, while only slightly increasing the trouble of driving along C," then "a rightly chosen measure of differential taxation against road B would create an 'artificial' situation superior to the 'natural' one." Knight (1924) argues that private road ownership should achieve exactly the outcome that Pigou desires.

3.A Administrative Costs for Congestion Pricing

One requirement of the theory of congestion pricing introduced in Section 2.D.1 is that the fixed administrative costs of operating the tax system must be high. This appears to be the case in practice.

Prud'homme and Bocarejo (2005) undertake a preliminary analysis of London congestion pricing and conclude that "the yearly amortization and operation costs of the charge system appear to be significantly higher than the economic benefit produced by the system." There are two critical numbers from their analysis: (1) The congestion charge was associated with a 16 percent reduction in traffic, and (2) administrative "collection costs" were 94 percent of revenues. While these numbers are disputable—and Raux (2005) argues that congestion pricing achieved more than just reducing traffic—our preceding framework suggests that the policy conclusion of Prud'homme and Bocarejo (2005) might have been correct.

Croci (2016) updates the figures on London's costs and compares the congestion pricing systems of London, Stockholm, and Milan. While operating costs are "only" 28 percent of revenues in Stockholm, they are 65 percent of revenues for Milan's Area C program and over 100 percent of revenues for Milan's Ecopass. London's costs initially accounted for 65 percent of revenues, but by 2008, they had fallen to 39 percent. These estimates are large both in absolute terms and relative to the changes in traffic that Croci (2016) associates with the congestion charge.

3.B What Information Was Used to Evaluate Congestion Pricing in London and New York?

London's path toward congestion pricing begins with the 1964 Smeed Report (Smeed 1964), which, as noted, is Pigouvian to its core. The report follows Walters (1961), who measures the congestion-related externality associated with driving by focusing on the engineering relationship between crowding and travel speeds.[9] On its second page, the Smeed Report notes that "a useful general rule upon which to base prices is that the road user should pay a sum equal to the costs he imposes upon others"; Appendix 2 then calculates that the externality per driver per mile ranges from 4.6 pence (at 20 miles per hour) to 212 pence (at 5

[9]Walters (1954) also extended Pigou's thinking on congestion pricing. Strikingly, neither Vickrey nor Walters cites the other.

miles per hour). The report notes that "we would not expect the cost of running the system to be more than 5–10 percent of the measurable benefits," suggesting that the benefits of estimating quantity responses are limited. Indeed, the report's recommendations for future work do not include estimating quantity responses. While the Smeed Report has long been justly admired by academics, its political impact was less fortunate. Goodwin (2024) writes, "As Deputy Director of the Road Research Laboratory, Smeed was tipped to become Director, but on the report's publication two years later, in 1964, the then Conservative Prime Minister, Sir Alec Douglas-Home, hated it." Shortly after, Smeed left public service for academia.

In 1989, Dame Patricia Hewitt wrote *A Cleaner, Faster London: Road Pricing, Transport Policy and the Environment* (Hewitt 1989). Hewitt's more politically astute report breaks sharply with Pigou and implicitly embraces a quantity target. She writes that "precise calculations of the external costs involved is not, however, necessary" because "the price can be fixed at whatever level is needed to persuade enough car-users to switch to another form of transport, or to change the time of their journey." She discusses the fixed administrative costs of electronic road pricing in a study on Hong Kong and extensively examines the quantity impact of congestion pricing in Singapore. She also describes quantity-related studies conducted in Hong Kong, the Netherlands, and the United Kingdom.[10]

Hewitt's study was followed in 1991 by *Transport: The New Realism* (Goodwin et al. 1991), which advocates the "principle" that "users should pay the costs they impose on others" and discusses extensively the total cost of congestion but does not try to measure the marginal costs of congestion. Instead, the primary section on congestion pricing focuses on the authors' "rule of three" plan, which will supposedly lead to "increased speed, especially at congested times, e.g., peak-period speeds to increase by 3–8 km/h." This recommendation follows Hewitt by emphasizing a quantity impact rather than measuring the externality. They also provide survey-based information on drivers' behavioral responses to interventions, such as higher gas prices, which can be used to estimate quantity responses.

In 1995, MVA Consultancy (1995) published *The London Congestion Charging Research Programme*, a quasi-official document commissioned by the Government Office for London.

[10] She also dismisses distributional issues in a section titled "Is Road Pricing Fair?".

The report never discusses the size of the externalities, and one-fifth of the document (10 pages) is allocated to "an assessment of the impact on traffic flows, and the economic benefit to travelers." Another fifth of the report is jointly allocated to two sections labeled "the implementation and operating costs."

In 1998, John Prescott published a government white paper titled "A New Deal for Transport: Better for Everyone" (Prescott 1998). The paper makes no mention of the magnitude of any externality, but it does cite the finding of MVA Consultancy (1995) that "it was estimated that vehicle miles would fall by 15% and CO_2 emissions by 14.5%."

In sum, while the discussion of congestion pricing in London began with an estimate of the optimal Pigouvian tax in 1964, by 1998, it focused solely on the tax's impact on driving. In New York, by contrast, the policy discussion around congestion pricing began much later and has consequently focused primarily on its impact on driving.

The modern discussion of congestion pricing in New York City began in 2007 with PlaNYC (City of New York 2007).[11] Unlike Hewitt (1989) or Goodwin et al. (1991), PlaNYC does not even articulate the core principle of a fee based on social costs. Instead, the document notes that "the main benefit of congestion pricing would be reduced traffic congestion" because "traffic within the zone would decrease 6.3%" and "speeds are projected to increase 7.2%."

In response to PlaNYC, the state legislature commissioned a 2008 "Traffic Congestion Mitigation Commission" report, which also fails to discuss the size of the congestion externalities imposed by drivers. The report instead focuses on operating costs and reductions in vehicle miles traveled, aligning with the predictions of Corollary 2. Nearly a decade later, in October 2017, Governor Andrew Cuomo commissioned a "Fix NYC Advisory Panel" that produced a plan (Fix NYC Advisory Panel 2018) proposing potential charges but not linking those charges to the measurement of any congestion externality. Rather, the advisory panel report predicts significant, highly quantitative behavioral responses: a 13 percent reduction in weekday trips and a 9 percent increase in weekday speeds. In September 2019, the Regional Plan Association (RPA), a venerable transportation planning group in the

[11] While there were some earlier informal discussions, we follow Schaller (2010) and treat PlaNYC as the genesis of the policy discussion.

New York region, released its own report supporting congestion pricing (Regional Plan Association 2019). This report also shows little interest in the size of externalities, pushing instead for two quantity-related targets: "Set the congestion charge high enough to meet both congestion and revenue targets."

The key documents relating to both the failed 2008 and later successful 2025 attempts to implement congestion pricing in New York pay no attention to the externalities that would justify a proper Pigouvian tax. Like the later London documents, the New York reports focus on the change in traffic and speeds. We take this quantity focus as an example of the relevance of our model's prediction that learning about quantity is more valuable than learning about the externality's size when administrative costs are high. Early evidence bears this out: Cook et al. (2025) find that the cordon toll raised speeds in the priced zone by roughly 11 percent, with spillovers onto unpriced trips across the metro area—a quantity response of just the kind that these documents and our model treat as central.

4 TAX ABATEMENT FOR RESIDENTIAL CONSTRUCTION

Our second application concerns local government policies to subsidize housing construction. The rising cost of housing has put pressure on local politicians to find ways to lower costs. Subsidies to stimulate housing construction are a plausible policy intervention. For this application, we provide a more detailed road map for how our framework may guide empirical policy analysis.

We focus on the City of Boston. In 2023, we served on a committee convened by Mayor Michelle Wu to evaluate potential residential property tax-abatement policies to catalyze construction. We provide a relatively high-level overview of our analysis, with more detail in our report, Alejandro et al. (2024).

The mayor convened our committee due to a significant backlog of projects in the development pipeline that had received city planning approval from the Boston Planning and Development Agency (BPDA) but had not begun construction primarily because of high interest rates and elevated construction costs in the wake of the COVID-19 pandemic.[12]

[12]We distinguish between *approved* projects and *permitted* projects that have paid fees and are ready to break ground. Because the permitting fee is approximately 1 percent of the project's value, developers

The scale of the housing backlog was substantial, as indicated in Table 1. According to BPDA officials, approximately 199 approved projects representing 22,000 to 23,000 units had failed to proceed to construction as of 2023. These "stuck" projects represented a significant missed opportunity to address the city's housing shortage, prompting interest in tax-abatement policies as a potential solution to unlock stalled development.

One policy proposal floated to the mayor was the abatement of real estate taxes for multifamily residential buildings that break ground immediately, along the lines of the 421-a policy that existed for 50 years in New York City. Real estate taxes are the City's primary source of revenue, and abating taxes so that developers pay based on the pre-development value of the real estate rather than the post-development value is the mayor's primary lever to subsidize construction. We were asked to consider various permutations of such a policy.

In analyzing this issue, we came to view the quantity impact of a subsidy—and, in particular, the ratio of total units to marginal units—as crucial to understanding its costs and benefits, an observation that informed the development of the framework presented in Section 2. We now describe our empirical analysis and findings as a guide to applying the insights of our framework.

4.A Data and Background

To analyze this phenomenon and evaluate potential policy responses, we obtained a project-level data set from the BPDA spanning the past decade. The data set includes all board approvals with project characteristics, timelines, neighborhood locations, unit counts, and development types. These micro-level data allow us to track individual projects from initial approval through permitting, providing the foundation for our empirical analysis of the factors driving construction delays and the potential effectiveness of tax-abatement policies.

Our empirical analysis is based on data covering 598 projects approved by the BPDA for development between January 2013 and June 2023. From that set of projects, we focus on "market-rate rental" projects, defined as rental developments that are not condominiums, have less than 25 percent Inclusionary Development Policy (IDP) units (that is, units restricted to be affordable), are not Low-Income Housing Tax Credit (LIHTC) projects, and

generally do not pay until they intend to begin construction.

are not public housing owned by the Boston Housing Authority. We also exclude one large development, Suffolk Downs, which is subject to unique planning challenges. The hazard model is estimated on all 598 BPDA approvals to improve precision; the policy counterfactuals are then applied to the pipeline of stalled market-rate projects eligible for tax abatement, which contains 77 market-rate rental projects representing 7,795 units and $4.0 billion in planned investment, demonstrating the substantial scale of stalled development that could be unlocked through policy intervention.

While developers initially attributed the construction slowdown primarily to rising interest rates, our analysis reveals a more complex picture driven by multiple market factors, most notably the dramatic increase in construction costs following the pandemic. Understanding these underlying economics is crucial for evaluating whether tax abatements can effectively address the fundamental constraints that prevent approved projects from moving forward or whether such policies would primarily subsidize inframarginal developments, that is, those that were already viable.

Our analysis combines two complementary tools. The forecasting tool is a logit hazard model that relates the historical permitting decision to the financial environment a project faces; we use it to predict how policy interventions that change the post-tax cap rate affect the rate at which approved projects move through permitting. To complement the forecast, we compute a project-level return on capital and plot its evolution alongside the cap rate, providing a transparent picture of how project economics deteriorated over the sample period.

4.B A Discrete-choice Model of Permitting

We model the permitting decision as the optimal exercise of an irreversible-investment option in the spirit of Capozza and Helsley (1990), Williams (1991), and Dixit and Pindyck (1994). The empirical implementation is in the spirit of Bulan et al. (2009), who estimate a discrete-time hazard for residential construction and interpret it as a reduced-form approximation of the optimal exercise threshold from a real-options model; Murphy (2018) provides the corresponding fully structural treatment for residential land development. As is standard in this literature, we follow Rust (1987) and McFadden (1973) by adding i.i.d. Type-1 extreme-value shocks to the developer's payoffs, which yields a logit form for the probability of construction.

Consider a developer who has received BPDA approval and at time $t = 1$ faces a decision between permitting the project today or waiting until $t = 2$. Let NOI_1 denote the net operating income from the built project in the first period, C the construction cost, r the long-term interest rate, g the expected growth rate of NOI, and τ the property tax rate. By the Gordon growth formula, the present value of the built project at the moment of completion is $NOI/(r - g + \tau)$, and the quantity $r - g + \tau$ is the cap rate that prevailing market practice applies to such a property. We assume r is constant across periods and discount with $1/(1 + r)$. There is uncertainty in NOI: The developer learns NOI_2 at the start of $t = 2$ but knows only its distribution at $t = 1$.

The developer compares two payoffs in present-value terms, with i.i.d. Type-1 extreme value shocks of scale σ capturing project-specific factors that the econometrician cannot observe:

$$\Pi^{\text{build}} = \frac{NOI_1}{r - g + \tau} - C + \sigma\,\varepsilon^{\text{build}}, \tag{4.1}$$

$$\Pi^{\text{wait}} = \frac{1}{1 + r}\,\mathbb{E}\left[\max\left\{\frac{NOI_2}{r - g + \tau} - C,\ 0\right\}\right] + \sigma\,\varepsilon^{\text{wait}}. \tag{4.2}$$

The optimal decision rule is to build whenever $\Pi^{\text{build}} > \Pi^{\text{wait}}$. Standard McFadden algebra delivers a logit form for the build probability:

$$\Pr(\text{build at } t = 1) = \frac{1}{1 + \exp\left(-\frac{(NOI_1/(r - g + \tau) - C) - W_2}{\sigma}\right)}, \tag{4.3}$$

where $W_2 = \frac{1}{1+r}\,\mathbb{E}[\max\{NOI_2/(r - g + \tau) - C,\, 0\}]$ is the option value of waiting until the second period.

The logit nests two simpler decision rules as limiting cases. With no uncertainty about the developer's payoff ($\sigma \to 0$), the logit collapses to a step function: Build if and only if $NOI_1/(r - g + \tau) > C + W_2$, which we can rearrange in return-on-capital form as

$$ROC = \frac{NOI_1}{C} > (r - g + \tau)\left(1 + \frac{W_2}{C}\right). \tag{4.4}$$

The relevant threshold is not the cap rate $r - g + \tau$ itself but the cap rate scaled up by an option-value-adjusted multiplier $(1 + W_2/C)$. If we further assume no uncertainty about future NOI, then the option value collapses to zero, and the rule reduces to a standard hurdle-rate calculation:

$$ROC = \frac{NOI_1}{C} > r - g + \tau. \tag{4.5}$$

This is also the textbook condition for residential construction: Build when the return on capital exceeds the cap rate. Because the property tax rate in Boston is roughly 1 percent of assessed value, and assessed value is roughly equal to construction cost, a full property tax abatement can offset at most a roughly 1 percentage point increase in r.

4.C Taking the Model to the Data

Equation (4.3) is a logit hazard with an index, which we call η, as its argument. We estimate η as a project-level variable that is a linear combination of observable covariates:

$$\eta_i = \beta' X_i + \gamma r_t, \tag{4.6}$$

with β and γ estimated freely. The relationship between covariates and the index is not constrained to satisfy any accounting identity. Instead, the data tell us how each covariate maps into the build probability, including the cap rate. Tax abatement enters the policy counterfactual through whatever coefficient the cap rate carries in the regression; we cannot separately identify a "tax channel" from a "cap rate channel" because τ enters the developer's problem only as part of $r - g + \tau$.

The logit error arises because the econometrician cannot observe everything that enters the developer's build-or-wait decision. Two sources of unobserved heterogeneity drive the extreme-value shocks $\varepsilon^{\text{build}}$ and $\varepsilon^{\text{wait}}$. The first is unobservable project-specific factors: the developer's financing terms, partner negotiations, contractor relationships, working capital constraints, and the actual land cost. The second is measurement error in the project-level economic inputs. In our analysis, CoStar rent indices proxy for project-specific rents; fixed soft-cost markups stand in for project-specific building costs; the census's Survey of Construction cost index for New England stands in for Boston-specific construction costs; and assessed land values stand in for actual land costs. Both sources move true project-level costs relative to our empirical proxies, and neither can be entirely eliminated simply by including additional covariates. The appropriate object for each project is therefore the probability that it will be in the money, not a binary indicator of whether it is.

The hazard model is estimated on the project-quarter panel, with the dependent variable equal to one in the quarter in which a project receives its building permit. The sample includes all BPDA-approved projects from 2013 to 2023 to improve precision; policy simulations later restrict predictions to market-rate rentals that were approved but had not begun construction as of July 2023. Each specification controls for time since BPDA approval, neighborhood fixed effects, housing-type fixed effects (condo or not), project size in units, an indicator for high IDP (affordability-restricted) share, and season (quarter).

Figures 1A–C plot the key time-varying covariates entering the regression: Supply chain pressure spiked in late 2020, sale prices peaked in mid-2021, and short-term rates (LIBOR) rose more than 3.5 percentage points by mid-2023 relative to their trough in 2020.

Table 2 presents the regression results, building up our preferred specification one group of controls at a time, with all columns including fixed effects for time since approval, neighborhood, housing type, high IDP share, and season. Column (1) controls for project size bins, with projects with 17 units or fewer as the excluded category, and confirms that larger projects proceed more slowly: A 50-unit project has about a 50 percent lower quarterly permitting probability than a 10-unit project. Column (2) adds an indicator for projects approved after 2019 and confirms that progress slowed sharply in the post-pandemic period. Columns (3)–(5) drop the post-2019 indicator and replace it with groups of market fundamentals that may explain the post-pandemic decline.

Column (3) introduces the level of the six-month LIBOR and the FRBNY Global Supply Chain Pressure Index (GSCPI), a composite published by the Federal Reserve Bank of New York that combines transportation costs and survey-based delivery delays. The GSCPI is meant to capture the sharp post-2020 increase in construction costs that does not show up directly in observed prices but does show up in developers' expected build-out costs and timelines. A one-standard-deviation increase in the GSCPI is associated with a 16 percent reduction in the quarterly permitting probability; a 100-basis-point increase in LIBOR reduces it by 11–14 percent.

Column (4) adds the cumulative change since BPDA approval in the neighborhood capitalization rate ("cap rate") from CoStar and in rent per square foot. The capitalization-rate change captures both the opportunity cost of cash and changes in expectations about future rent growth, and the rent change captures contemporaneous revenue conditions. A 100 basis point increase in the cap rate roughly halves the quarterly permitting probability; a one-standard-deviation increase in neighborhood rent per square foot raises it by 23–25 percent. Column (5), our preferred specification, adds the log of assessed land value as a project-fixed measure of project quality not captured by the size dummies. Each economic channel—revenue, financing cost, and supply chain pressure—enters with the sign and magnitude that the developer's decision problem would predict.

We assess model performance using the area under the receiver operating characteristic curve (AUC). The in-sample AUC for column (5) is 0.68; the out-of-sample AUC—estimated on data through mid-2023 and evaluated on subsequent permits—is 0.63, well above the 0.50 coin-toss benchmark. The model does not precisely predict individual project outcomes, which depend on idiosyncratic factors, but it captures the aggregate relationship between financial conditions and permitting flows.

4.D Understanding the Post-2020 Pipeline Backlog

The estimated model allows us to conduct counterfactual analyses showing what the BPDA pipeline would have looked like absent either the post-2019 break or the post-2021 rate increases. In both counterfactuals, we assume that entry into the pipeline was unaffected—this is justifiable since the BPDA approval process typically begins years before the approval is secured—and that only the probability of permitting responds to the shocks.

Our first counterfactual uses the column (2) specification, which includes only project-size dummies, time since approval, neighborhood effects, and a post-2019 indicator. The red dashed line in Figure 2 shows the simulation with the post-2019 indicator set to zero. The pipeline stabilizes near 48 projects—reasonable if roughly five projects enter and 10 percent of pipeline projects are permitted each quarter, balancing inflow and outflow. The actual pipeline at the end of the sample is 77 projects; the gap of 29 projects, at about 108 units per project, suggests that post-2019 disruptions account for about 3,100 fewer units permitted than pre-pandemic patterns would have predicted.

Our second counterfactual uses the column (5) specification and freezes the cap rate and the six-month LIBOR at their 2021:Q4 levels while letting rents, the supply chain index, and other variables follow their realized paths. The dashed orange line in Figure 2 shows the result. The pipeline grows more slowly after 2021 but does not stabilize; at the end of the sample, it sits at about 62 projects. The gap of 15 projects between the actual line and this counterfactual—about 1,600 units—is the share of the post-2021 backlog attributable specifically to rate increases. The remaining 14 projects (about 1,500 units) of the post-2019 break reflect other forces, most prominently the post-2020 construction-cost surge that we examine more directly in Figure 3.

4.E The Impact of Property Tax Abatement on the BPDA Pipeline

Using the preferred specification from column (5) of Table 2, we simulate two types of tax abatement: proportional reductions (100 and 75 percent) and per-unit credits capped at \$5,000 or \$2,500 per year. Per-unit credits are converted to cap-rate equivalents by dividing by the current per-unit tax payment and multiplying by the tax rate. For example, a \$5,000 credit on a project with a \$7,500 per-unit obligation is equivalent to a 67 basis point cap-rate reduction; projects with tax bills below the cap receive a full abatement. We assume the subsidy is available for any project permitted by the end of 2025 and that market variables follow the baseline CoStar forecasts.

The original proposal advanced to the Mayor Wu's office in early 2023 was a full 100 percent abatement of residential property tax for new market-rate rental construction, modeled on New York's 421-a program and lasting roughly the same 29-year horizon. We use that proposal as the anchor point for our policy menu and consider variations along three dimensions. First, we vary the abatement *duration*—perpetual, 29 years, and 15 years—to ask how much of the cost can be avoided by simply truncating the subsidy stream. Second, we vary the abatement *depth*—100 percent versus 75 percent—to ask whether scaling back the proportional reduction meaningfully changes the marginal-unit count. Third, we replace the proportional reduction entirely with *per-unit dollar caps* of \$5,000 and \$2,500, which limit the subsidy that the most expensive luxury projects can capture and concentrate relief on smaller and more affordable buildings. The menu spans the range of designs the mayor's working group considered, from the most generous variant of the original proposal to the most targeted alternative we judged practical to administer.

Table 3 presents the results. Column (1) reports total units permitted; columns (2)–(4) report the NPV fiscal impact (revenue, service costs, and net cost); columns (5)–(7) express costs per marginal unit. We conservatively assume service costs per unit equal per-unit tax revenue under baseline—that is, property taxes are well calibrated to service costs.

About 3,400 units are permitted absent any policy change. A perpetual full abatement generates about 1,800 marginal units at roughly $549,000 each. The fundamental problem is the ratio of total to marginal units reported in column (6): Even the most generous policy produces only about one marginal unit for every three total units subsidized. The revenue loss per unit may be modest, but the City must extend the abatement to all qualifying projects, and the majority would have been built anyway. The fiscal cost is therefore dominated by foregone revenue on inframarginal units, not the direct cost of inducing new ones.

Shorter durations and per-unit caps improve efficiency by shrinking the per-unit subsidy rather than the ratio of total to marginal units. The 29-year full abatement costs about $427,000 per marginal unit; the $2,500 per-unit cap costs roughly $299,000. Per-unit caps are more efficient because they provide full relief to affordable projects while limiting the subsidy to expensive buildings. A 75 percent abatement generates fewer marginal units than the $5,000 cap but costs more per unit because percentage reductions give the largest dollar subsidies to the most expensive projects.

4.F A Back-of-the-envelope Check on Project Economics

To complement the regression-based forecast, we compute the project-level return-on-capital gap—that is, the difference between the two sides of equation (4.5), $NOI_i/C_i - (r - g + \tau)$—and compare its evolution with the cap rate.[13] Because we do not directly observe project-level revenues, costs, or financing terms, we estimate them from observable project

[13] In our report to Mayor Wu, we used these return-on-capital calculations together with an assumed variance in hurdle cap rates to represent measurement error as part of the policy analysis itself; here, we focus on the logit specification and report the gap purely as a descriptive picture of how project economics evolved over time. Details of the construction of our project-level return on capital are in Appendix C.

characteristics—units, parking spaces, height, square footage, and neighborhood—combined with BPDA filings and market indexes. Net operating income is derived from regressions of rent and expense ratios on observables, fitted to BPDA Tax Increment Program filings, and scaled over time by CoStar's neighborhood rent index. Construction costs are estimated using a hedonic regression of cost per square foot (derived from the BPDA permitting fee) on project characteristics, with a 12 percent post-2020 inflation adjustment. Land costs are based on the assessor's value, and construction-period financing is added at a 60 percent loan-to-cost ratio at the two-year LIBOR rate plus 200 basis points over a 22-month construction timetable. The project-level return on capital is therefore a back-of-the-envelope summary rather than a precise measure of viability; Appendix C gives the full details.

Figure 3 shows the mean gap over time, with the cap rate plotted alongside for reference. Two features stand out. First, the gap was comfortably positive in 2019, climbed through the low-rate environment of 2020–2021, and then fell sharply over 2022 and 2023; by 2023, even projects one standard deviation above the mean were below the cap rate. The collapse coincides with the rise in interest rates and the surge in construction costs that followed the pandemic. Second, the dashed line—a counterfactual in which construction costs grow at their pre-pandemic trend rather than at their post-2020 surge—tells a useful story about the relative roles of each shock in explaining the decline in construction. Under the cost counterfactual, even with the actual post-2022 rise in interest rates, the gap would have stayed close to its 2019 level. The implication is that the apparent "normalcy" of 2019–2021 was driven by abnormally low interest rates more than by any underlying favorable trend and that the post-2022 deterioration is primarily due to rising construction costs rather than rising interest rates. This matches the regression picture: The cost-shock channel is the largest single driver of the post-2020 construction freeze.

Overall, our finding is that the ratio of total units to marginal units is high, which our framework indicates is the key margin for analysis due to distributional concerns and deadweight losses associated with taxation. A policymaker would have to weigh these quantity impacts and implicit subsidy cost per unit built against the perceived benefits of immediate construction.

4.G When Does Tax Abatement for Residential Construction Make Sense?

Subsidies for the construction of market-rate residential buildings are sometimes proposed when rising costs of living create political imperatives to spur development despite challenging economic conditions, as was the case in Boston. Our analysis of Boston suggests that quantity impacts are relatively modest, and per-unit costs are high due to a high ratio of total to marginal units built; however, this need not be the case in all circumstances.

Subsidies can also occur when a city uses tax abatement to rejuvenate neighborhoods by encouraging development in abandoned or blighted areas, which was not a goal of the Boston policy proposal. Such policies can create substantial amenities that catalyze further local improvements, increasing neighboring property values and property tax receipts for the City in the long run. Indeed, Rossi-Hansberg et al. (2010) find that Richmond's "Neighborhood-in-Bloom" program raised tax revenues by $2 to $6 for every dollar spent on subsidies, and Schwartz et al. (2006) find that New York City's "Ten Year Plan" generated increased revenues that more than offset the cost of the subsidies. These urban renewal projects would have been far less costly per unit subsidized if nothing had been built in the neighborhood without the subsidy. If the non-subsidy quantity is essentially zero, measuring it is far less relevant, and the project's externalities on neighborhood amenities are far more important, consistent with our theory.

5 CONCLUSION

Knowledge is scarce, and so is the effort available to generate new knowledge. Frequently, economists must decide what topic is worth their time. Fifty years ago, Weitzman (1974) taught us when it made sense to intervene through prices—perhaps by imposing a Pigouvian tax—and when to intervene through quantity regulation.

In this paper, we present a model that analyzes when it makes sense to learn about optimal prices or about how quantities will respond to price changes. Our model implies that when implementation costs are low or when distributional issues are negligible, knowledge of "price" is sufficient. There is no need to understand market responses to price changes. That result has, of course, been known for over a century. However, when there are fixed policy costs, such as with congestion pricing, or when the taxes needed to fund subsidies generate substantial deadweight losses, adding precision about quantity responses can be more valuable than learning the size of an externality.

We apply this framework to the discussion of congestion pricing in London and New York City. The policy discussion of congestion pricing in the United Kingdom began with the Smeed Report, which focuses almost entirely on the size of the externalities associated with driving. Over the next 40 years, writers gave increasing attention to how much a congestion fee would reduce traffic, and they were right to do so. New York City followed suit. Congestion pricing involves extremely high administrative fixed costs, which are justifiable only if it has a significant impact on traffic.

We then turn to our analysis of subsidizing housing in Boston. In this setting, Mayor Wu had a sense of the social value of new housing and wanted to know the costs. The key question is how many incremental units could be produced for each dollar of subsidy. We use a regression-based approach to estimate the relationship between financial conditions and new housing production, treating a subsidy as a change in the effective capitalization rate. We find that building subsidies in Boston would be costly because they subsidize many units that would otherwise have been built—exactly the total-to-marginal ratio that our framework implies a policymaker should consider when taxes create deadweight losses or there are distributional concerns.

Economists have often shown far more enthusiasm for Pigouvian taxes than the general public, perhaps because we have assumed away the messy costs that sometimes accompany them. The fact that costs and distributional consequences arise implies that good policymaking is based on the size of any impact and the quantity changes implied by any intervention. We hope that future research provides more information about both of these objects. Understanding both prices and quantities can strengthen the case for future Pigouvian interventions.

References

Alejandro, Patricia, Mary Ellen Carter, Denise DiPasquale, Edward Glaeser, Adam M. Guren, and Paul S. Willen. 2024. "Report on the Potential Impacts of Property Tax Abatement on Rental Housing Construction in Boston." Federal Reserve Bank of Boston Working Paper 24-1.

Baumol, William J. 1972. "On Taxation and the Control of Externalities." *The American Economic Review* 62 (3): 307–322.

Bulan, Laarni, Christopher Mayer, and C. Tsuriel Somerville. 2009. "Irreversible Investment, Real Options, and Competition: Evidence from Real Estate Development." *Journal of Urban Economics* 65 (3): 237–251.

Capozza, Dennis R., and Robert W. Helsley. 1990. "The Stochastic City." *Journal of Urban Economics* 28 (2): 187–203.

Casler, Stephen D, and Aisha Rafiqui. 1993. "Evaluating Fuel Tax Equity: Direct and Indirect Distributional Effects." *National Tax Journal* 46 (2): 197–205.

City of New York. 2007. "PlaNYC: A Greener, Greater New York." City of New York Report.

Cook, Cody, Aboudy Kreidieh, Neha Arora, Hunt Allcott, Freek van Sambeek, Andrew Tomkins, Eray Turkel, and Shoshana Vasserman. 2025. "The Short-Run Effects of Congestion Pricing in New York City." NBER Working Paper 33584, National Bureau of Economic Research, `https://www.nber.org/papers/w33584`.

Croci, Edoardo. 2016. "Urban Road Pricing: A Comparative Study on the Experiences of London, Stockholm and Milan." *Transportation Research Procedia* 14 253–262.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty.* Princeton University Press.

Fix NYC Advisory Panel. 2018. "Fix NYC Advisory Panel Report." State of New York Report.

Glaeser, Edward L., Joseph Gyourko, and Raven Saks. 2005. "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices." *The Journal of Law and Economics* 48 (2): 331–369.

Glaeser, Edward L., Michael Luca, and Erica Moszkowski. 2023. "Gentrification and Retail Churn: Theory and Evidence." *Regional Science and Urban Economics* 100 103879.

Goodwin, Phil. 2024. "The Smeed Report on Road Pricing: Still Influential After 60 Years, But the Difficult Issues Still Remain." Available online at https://tapas.network/64/goodwin.php.

Goodwin, Phil, Sharon Hallett, Francesca Kenny, and Gordon Stokes. 1991. "Transport: The New Realism." Transport Studies Unit, University of Oxford, Report 624.

Goulder, Lawrence H. 2005. "Environmental Taxation and the Double Dividend: A Reader's Guide." In *Environmental Policy Making in Economies With Prior Tax Distortions*, edited by Goulder, Lawrence H. Chap. 3, Edward Elgar Publishing.

Hewitt, Patricia. 1989. *A Cleaner, Faster London*. London: Institute for Public Policy Research.

Kaplow, Louis. 2012. "Optimal Control of Externalities in the Presence of Income Taxation." *International Economic Review* 53 (2): 487–509.

Knight, F. H. 1924. "Some Fallacies in the Interpretation of Social Cost." *The Quarterly Journal of Economics* 38 (4): 582–606.

McFadden, Daniel. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Zarembka, P. 105–142, Academic Press.

Murphy, Alvin. 2018. "A Dynamic Model of Housing Supply." *American Economic Journal: Economic Policy* 10 (4): 243–267.

MVA Consultancy. 1995. "The London Congestion Charging Research Programme: Principal Findings."

Paciello, Luigi, and Mirko Wiederholt. 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *The Review of Economic Studies* 81 (1): 356–388.

Pai, Mallesh, and Philipp Strack. 2022. "Taxing Externalities Without Hurting the Poor." Yale University Working Paper, Available at SSRN 4180522.

Pigou, Arthur C. 1920. *The Economics of Welfare*. London: Macmillan.

Pigou, Arthur C. 1928. *A Study In Public Finance*. London: Macmillan.

Prescott, John. 1998. "A New Deal for Transport: Better for Everyone." The Government's White Paper on the Future of Transport, London.

Prud'homme, Remy, and Juan Pablo Bocarejo. 2005. "The London Congestion Charge: A Tentative Economic Appraisal." *Transport Policy* 12 (3): 279–287.

Raux, Charles. 2005. "Comments on 'The London Congestion Charge: A Tentative Economic Appraisal'." *Transport Policy* 12 (4): 368–371.

Regional Plan Association. 2019. "Congestion Pricing in NYC: Getting It Right." Available online at https://rpa.org/work/reports/congestion-pricing-in-nyc.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens, III. 2010. "Housing Externalities." *Journal of Political Economy* 118 (3): 485–535.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033.

Sandmo, Agnar. 1975. "Optimal Taxation in the Presence of Externalities." *The Swedish Journal of Economics* 77 (1): 86–98.

Schaller, Bruce. 2010. "New York City's Congestion Pricing Experience and Implications for Road Pricing Acceptance in the United States." *Transport Policy* 17 (4): 266–273.

Schwartz, Amy Ellen, Ingrid Gould Ellen, Ioan Voicu, and Michael H. Schill. 2006. "The External Effects of Place-Based Subsidized Housing." *Regional Science and Urban Economics* 36 (6): 679–707.

Smeed, R. J. 1964. "Road Pricing: the Economic and Technical Possibilities."Technical report, HMSO, `https://archive.org/details/op1265810-1001`.

Vickrey, William S. 1955. "A Proposal for Revising New York's Subway Fare Structure." *Journal of the Operations Research Society of America* 3 (1): 38–68.

Vickrey, William S. 1959. "Statement on the Pricing of Urban Street Use." In *Hearing of US Joint Committee on Metropolitan Washington Problems*, Volume 11. 42–65.

Vickrey, William S. 1963. "Pricing in Urban and Suburban Transport." *The American Economic Review* 53 (2): 452–465.

Walters, Alan A. 1954. "Track Costs and Motor Taxation." *The Journal of Industrial Economics* 2 (2): 135–146.

Walters, Alan A. 1961. "The Theory and Measurement of Private and Social Cost of Highway Congestion." *Econometrica* 676–699.

Weitzman, Martin L. 1974. "Prices vs. Quantities." *The Review of Economic Studies* 41 (4): 477–491.

Williams, Joseph T. 1991. "Real Estate Development as an Option." *Journal of Real Estate Finance and Economics* 4 (2): 191–208.

Zhong, Weijie. 2022. "Optimal Dynamic Information Acquisition." *Econometrica* 90 (4): 1537–1582.

Table 1. This table describes the sample used in the analysis. **Panel A (top)** documents the construction of the policy-simulation sample from the universe of BPDA-approved projects: The left block shows the full pipeline of 199 projects comprising 22,872 units and $13.9 billion in investment, and reports the units excluded (owner-occupied projects, the Suffolk Downs mega-development, Boston Housing Authority projects, and projects with inclusionary-development percentages above 25 percent—all of which face permitting incentives that differ substantially from those of market-rate rental developments), leaving a policy-simulation sample of 77 market-rate rental projects with 7,795 units and $4.0 billion in investment; the right block shows the corresponding figures for all BPDA-approved projects. (The hazard model itself is estimated on all BPDA approvals; the 77-project pipeline is the sample to which the policy counterfactuals are applied.) **Panel B (bottom)** reports, for each approval-year cohort, the number of projects approved, the mean project size (in units), the prevailing cap rate and six-month LIBOR rate at approval, and the share of projects that received a building permit (overall and within two years); the right block restricts attention to projects still in the pipeline as of July 1, 2023. The sharp drop in permitting rates for the 2022 and 2023 cohorts, together with the rising pipeline count, motivates the hazard-model analysis.

Panel A: Sample Construction

	Pipeline			All Projects		
	# Projects	Units	Investment in $bn	# Projects	Units	Investment in $bn
Total	199	22,872	13.9	598	49,118	29.3
− Owner-Occupied	79	5,836	4.4	259	11,842	9.7
− Suffolk Downs	5	5,720	2.9	5	5,720	2.9
− BHA	5	1,931	1.6	17	2,498	2.0
− >25% IDP	33	1,590	1.0	89	4,544	2.6
= Market Rental	77	7,795	4.0	228	24,514	12.1

Panel B: Summary Statistics by Approval Year

	sample (as of approval)						Pipeline (as of 7/1/2023)				
Year Approved	# of Projects	Mean # of units	Cap Rate	6-mo LIBOR	% permitted	% perm. in 2 yrs	# of Projects	Mean # of units	Δ Cap Rate	Δ rent in %	Cumulative Total
2013	20	182	4.61	0.28	100	65	NA	NA	NA	NA	NA
2014	20	137	4.48	0.24	90	60	2	158	-0.35	24.2	77
2015	11	51	4.23	0.38	91	73	1	15	0.14	21.4	75
2016	33	107	4.14	0.77	91	67	3	74	0.05	19.7	74
2017	23	114	4.05	1.36	78	48	5	126	0.16	17.4	71
2018	19	61	4.00	2.45	79	47	4	47	0.23	16.6	66
2019	29	96	3.95	2.28	72	31	8	137	0.24	11.5	62
2020	15	97	3.78	0.61	53	20	7	88	0.39	15.7	54
2021	23	106	3.52	0.17	39	30	14	126	0.72	9.5	47
2022	23	90	3.54	2.65	4	4	22	78	0.62	4.0	33
2023	11	112	3.91	5.18	0	0	11	112	0.23	2.2	11
Total	227	107	4.01	1.39	66	42	77	101	0.42	9.5	

Table 2. Logistic regressions in which the dependent variable equals one if a project receives a building permit in a given quarter. Column (1) includes only project-size indicators along with fixed effects for time since BPDA approval, neighborhood, housing type, a dummy for high IDP share, and season. Column (2) adds a post-2019 indicator. Column (3) adds the GSCPI and the six-month LIBOR rate. Column (4) adds the cumulative change in neighborhood cap rates and in rent per square foot. Column (5), the preferred specification, adds log assessed land value. Standard errors, clustered by quarter and neighborhood, are in parentheses. Significance: *** 0.01, ** 0.05, * 0.10.

Dependent Variable:	Prob(Permit)				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
(18,35] Units	-0.3423***	-0.2173**	-0.2807**	-0.3339**	-0.4120**
	(0.1232)	(0.1029)	(0.1101)	(0.1322)	(0.1766)
(35,72] Units	-0.4905***	-0.4536***	-0.4732***	-0.4958***	-0.5761***
	(0.1299)	(0.1290)	(0.1166)	(0.1260)	(0.1395)
(72,832] Units	-0.5956***	-0.5664***	-0.5985***	-0.6672***	-0.7303***
	(0.2201)	(0.1870)	(0.1906)	(0.2034)	(0.2366)
Year>2019		-0.7109***			
		(0.1015)			
FRBNY Supply Chain Pressure Index			-0.1564***	-0.1969***	-0.1593***
			(0.0457)	(0.0366)	(0.0442)
6-month LIBOR in %			-0.1537***	-0.1373**	-0.1118*
			(0.0499)	(0.0610)	(0.0664)
Cumulative Δ in cap rate in %				-0.8622*	-0.8712*
				(0.4507)	(0.4895)
Cumulative Δ in rent per square foot (scaled)				0.2514**	0.2320**
				(0.1041)	(0.0916)
Log assessed land value					-0.0417***
					(0.0120)
Fixed-effects					
Time since approval	Yes	Yes	Yes	Yes	Yes
Neighborhood	Yes	Yes	Yes	Yes	Yes
Condo	Yes	Yes	Yes	Yes	Yes
High IDP percentage	Yes	Yes	Yes	Yes	Yes
Seasonal	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	4,120	4,120	4,120	4,107	3,968
Pseudo R^2	0.03488	0.04835	0.04396	0.05106	0.05839
In-sample AUC	0.642	0.667	0.659	0.670	0.679
Out-of-sample AUC	0.502	0.562	0.553	0.604	0.629

Table 3. Tax-abatement policy simulations using the hazard model from column (5) of Table 2. Each row corresponds to a different policy scenario. Column (1) reports total housing units produced. Column (2) reports the NPV of residential tax revenue. Column (3) reports the cost of providing city services (assuming per-unit cost equals per-unit revenue in the no-policy-change scenario). Column (4) reports the net fiscal cost. Column (5) reports marginal units induced relative to the no-policy-change scenario. Column (6) reports the ratio of total to marginal units. Column (7) reports the net cost per marginal unit. Marginal units in all columns are measured relative to the no-policy-change scenario. Market variables follow CoStar's forecast. All figures are NPV in millions (or thousands per unit) discounted at 3 percent.

			Budget Impact			*Cost Per New Unit*		
Policy	Horizon	(1) # of Units	(2) Revenue NPV in $ millions	(3) Costs NPV in $ millions	(4)=(3)-(2) Net Cost NPV in $ millions	(5) Marginal Units	(6)=(1)/(5) Total/ Marginal	(7)=(4)/(5) Cost per unit NPV in $ thousands
Assuming Costar forecast cap rates								
No policy change	∞	3362	643.6	643.6	0.0			
100% of revenue	∞	5147	0.0	980.2	980.2	1786	2.9	549.0
100% of revenue	29	4540	362.9	866.2	503.3	1178	3.9	427.1
100% of revenue	15	4134	503.5	789.7	286.2	772	5.4	370.7
75% of revenue	29	4206	453.3	803.4	350.1	845	5.0	414.6
$5000/unit	29	4363	457.5	803.1	345.5	1002	4.4	344.9
$2500/unit	29	3969	551.4	732.7	181.4	607	6.5	298.6

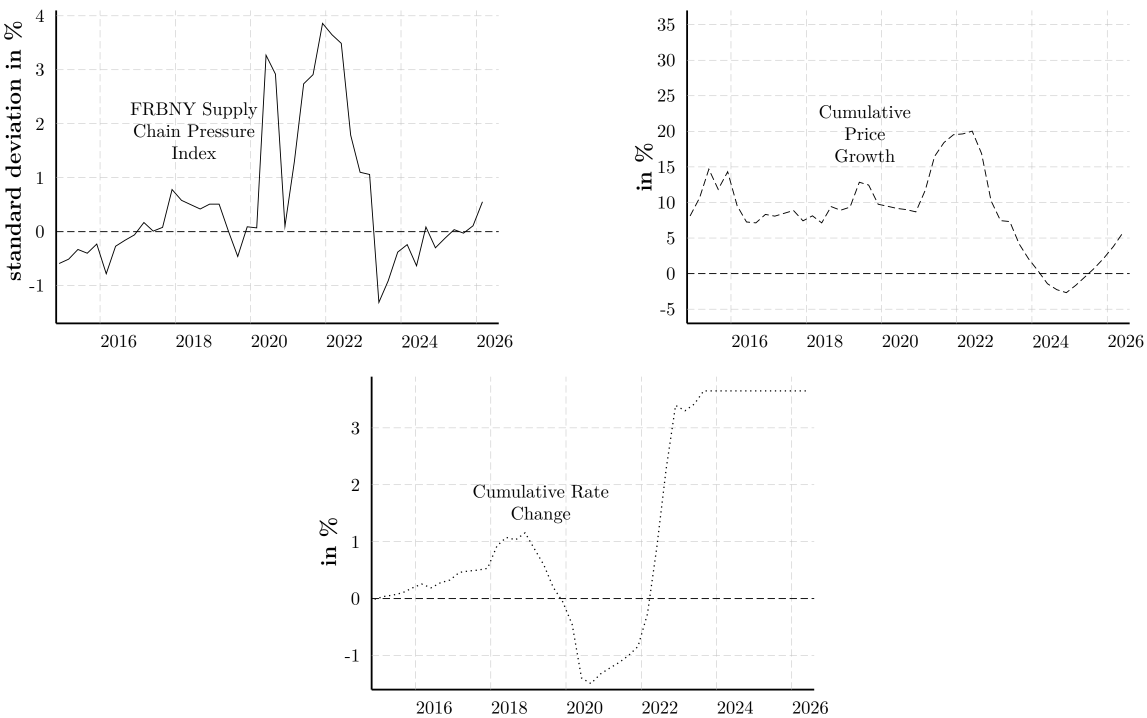

Figure 1. Time-varying covariates used in the hazard models. Panel A (top left) shows the FRBNY Global Supply Chain Pressure Index (GSCPI) in standard deviations from its historical mean. Panel B (top right) shows cumulative growth in the multifamily sale price per unit relative to each project's BPDA-approval date. Panel C (bottom) shows the cumulative change in the six-month LIBOR rate from each project's approval date, in percentage points.

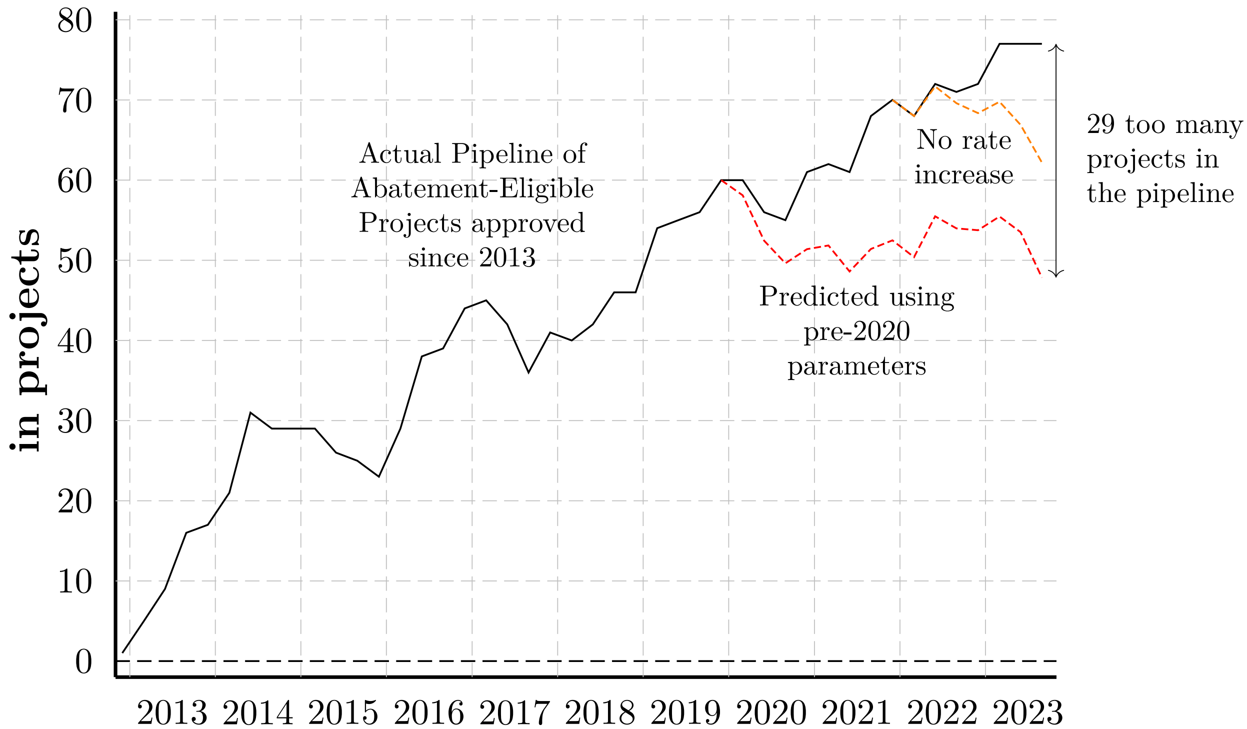

Figure 2. Actual pipeline of abatement-eligible projects approved since 2013 (solid black) alongside two counterfactual scenarios. The dashed red line uses the specification from column (2) of Table 2, which includes only project characteristics and a post-2019 indicator, and sets that indicator to zero—showing what the pipeline would have looked like absent all post-2019 disruptions. The gap at the end of the sample—about 29 projects—represents excess projects beyond what pre-pandemic patterns would predict. The dashed orange line uses the full specification from column (5) of Table 2 but freezes interest rates and cap rates at their 2021:Q4 levels; the difference between the orange and black lines isolates the effect of the post-2021 rate increases.

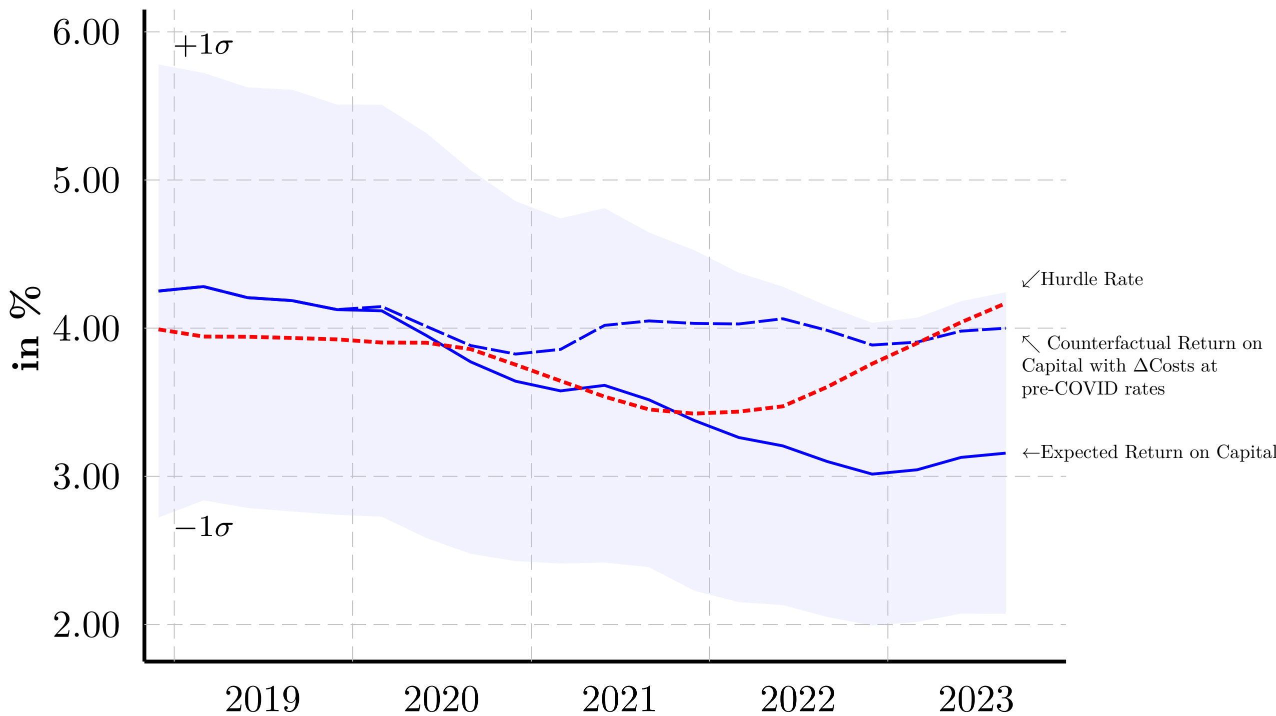

Figure 3. Mean return on capital (ROC) for BPDA pipeline projects (solid blue) alongside the mean market capitalization rate from CoStar (dashed red). The shaded region shows ±1 standard deviation of the cross-sectional ROC distribution. Construction is profitable when the ROC exceeds the cap rate. The dashed blue line shows a counterfactual ROC assuming construction costs followed their pre-pandemic growth trend.

A PROOFS OF PROPOSITIONS

A.A Proof of Proposition 1

With subsidy s, the market equilibrium is $B_1^I + \beta^I(\eta_I) + B_2^I(q-\widehat{q}) = p = C_1 + \alpha(\theta) + C_2(q-\widehat{q}) - s$.

Putting these together, $q - \widehat{q} = \frac{s+B_1^I+\widetilde{\beta}^T(\eta_T)-C_1}{C_2-B_2^I} = q_0 + \frac{s+\widetilde{\beta}^T(\eta_T)}{C_2-B_2^I}$, where $q_0 = \frac{B_1^I-C_1}{C_2-B_2^I}$ and $p = \frac{C_2(B_1^I+\beta^I(\eta_I))-B_2^I(C_1+\alpha(\theta))+B_2^I s}{C_2-B_2^I} = p_0 + \frac{C_2\beta^I(\eta_I)-B_2^I\alpha(\theta)-B_2^I s}{C_2-B_2^I}$, where $p_0 = \frac{C_2B_1^I-B_2^IC_1}{C_2-B_2^I}$. Consequently, $q_{NS} = \widehat{q} + \frac{B_1^I+\widetilde{\beta}^T(\eta_T)-C_1}{C_2-B_2^I}$, $\Delta_Q = \frac{s}{C_2-B_2^I}$ and $\Delta_P = \frac{B_2^I s}{C_2-B_2^I}$.

Social welfare is $B^I(q,\eta_I) + B^P(q,\eta_P) - (p+(1+\lambda)s)q + \delta((p+s)q - C(q,\theta)) - KI_{s\neq 0}$, which, using the supply and demand curves, can be written as

$$\widehat{q}\left(\delta\left(C_1+\alpha(\theta)\right) - B_1^I - \beta^I(\eta_I)\right) + b^P(\eta_P) + b^I(\eta_I) - \delta a(\theta) + 0.5\left(\delta C_2 - B_2^I\right)(q-\widehat{q})^2 + \\ \left(B_1^P + \beta^P(\eta_P) + \delta C_2\widehat{q} - \widehat{q}B_2^I\right)(q-\widehat{q}) - (1+\lambda)sq - KI_{s\neq 0}.$$

Simplifying this, welfare equals

$$V_0 + \frac{s}{C_2-B_2^I}\left(B_1^P + \beta^P(\eta_P) - \frac{(1-\delta+\lambda)C_2 - \lambda B_2^I}{C_2-B_2^I}\left(B_1^I + \beta^I(\eta_I) - C_1 - \alpha(\theta) + \widehat{q}(C_2-B_2^I)\right)\right) \\ -\frac{(1+\lambda-0.5\delta)C_2 - (0.5+\lambda)B_2^I}{\left(C_2-B_2^I\right)^2}s^2 - KI_{s\neq 0},$$

where[14]

$$V_0 = \widehat{q}\left(\delta\left(C_1+\alpha(\theta)\right) - B_1^I - \beta^I(\eta_I)\right) + b^P(\eta_P) + b^I(\eta_I) - \delta a(\theta) + \\ \frac{\left(B_1^P + \beta^P(\eta_P) + \delta C_2\widehat{q} - \widehat{q}B_2^I\right)\left(B_1^I + \beta^I(\eta_I) - C_1 - \alpha(\theta)\right)}{C_2-B_2^I} + \\ 0.5\left(\delta C_2 - B_2^I\right)\left(\frac{B_1^I + \beta^I(\eta_I) - C_1 - \alpha(\theta)}{C_2-B_2^I}\right)^2.$$

[14] An earlier version of this proof included a spurious $-\frac{1}{2}(B_2^I + C_2)(q-\widehat{q})^2$ term in the definition of V_0; symbolic verification (setting $s = 0$ and comparing with the preceding welfare expression) confirms that term does not belong and has thus been removed.

A subsidy is better than nothing if and only if

$$\frac{s}{C_2 - B_2^I}\left(B_1^P + \beta^P\left(\eta_P\right) - \frac{(1-\delta+\lambda)C_2 - \lambda B_2^I}{C_2 - B_2^I}\left(B_1^I + \beta^I\left(\eta_I\right) - C_1 - \alpha(\theta) + \widehat{q}(C_2 - B_2^I)\right)\right)$$
$$> \frac{(1+\lambda-0.5\delta)C_2 - (0.5+\lambda)B_2^I}{\left(C_2 - B_2^I\right)^2}s^2 + K.$$

Substitution gives us that this inequality is equivalent to

$$B_1^P + \beta^P\left(\eta_P\right) > 0.5s + \lambda s\left(\frac{q_{NS}}{\Delta_Q} + 1\right) + (1-\delta)\left(s + \Delta_P\right)\left(\frac{q_{NS}}{\Delta_Q} + 0.5\right) + \frac{K}{\Delta_Q}.$$

If the social planner gets to choose s to maximize expected welfare, their first-order condition is

$$s = \frac{B_1^P + \widetilde{\beta}^P\left(\eta_P\right)}{1 + 2\lambda + \frac{(1-\delta)C_2}{C_2 - B_2^I}} - \frac{\lambda + \frac{(1-\delta)C_2}{C_2 - B_2^I}}{1 + 2\lambda + \frac{(1-\delta)C_2}{C_2 - B_2^I}}\frac{q_{NS}}{\frac{1}{C_2 - B_2^I}}.$$

The planner's expected social welfare is $E(V_0)$ plus the conditional expectation

$$E\left(V_0\right) + \frac{\left(B_1^P + \widetilde{\beta}^P\left(\eta_P\right) - \left((1-\delta+\lambda)C_2 - \lambda B_2^I\right)\left(\widehat{q} + \frac{B_1^I + \widetilde{\beta}^I(\eta_I) - C_1 - \widetilde{\alpha}(\theta)}{C_2 - B_2^I}\right)\right)^2}{2\left((1+2\lambda)(C_2 - B_2^I) + (1-\delta)C_2\right)}.$$

The unconditional expectation of this quantity is $E\left(V_0\right) + \frac{\left(B_1^P - \left((1-\delta+\lambda)C_2 - \lambda B_2^I\right)(q_0 + \widehat{q})\right)^2}{2\left((1+2\lambda)(C_2 - B_2^I) + (1-\delta)C_2\right)}$, which is independent of the variance of the error terms, plus

$$\frac{\frac{\left(Var\left(\beta^P(\eta_P)\right)\right)^2}{Var(\beta^P(\eta_P)) + Var(\varepsilon_P)} + \left(\frac{(1-\delta+\lambda)C_2 - \lambda B_2^I}{C_2 - B_2^I}\right)^2 \frac{\left(Var(\beta^T(\eta_T))\right)^2}{Var(\beta^T(\eta_T)) + Var(\varepsilon_T)}}{4\left((1+2\lambda)(C_2 - B_2^I) + (1-\delta)C_2\right)^2}.$$

This term is always decreasing with $Var\left(\varepsilon_P\right)$ and independent of $Var\left(\varepsilon_T\right)$ as long as $(1-\delta+\lambda)C_2 = \lambda B_2^I$, which is always true if $\lambda = 0$ and $\delta = 1$. As long as $(1-\delta+\lambda)C_2 \neq \lambda B_2^I$, then this expression is always declining in $Var\left(\varepsilon_T\right)$.

A.B Proof of Proposition 2

Using the notation $Var(\varepsilon_j) = \varphi_j Var(\beta^j(\eta_j))$ for $j = T, P$, the increase in welfare if $Var(\varepsilon_T)$ is reduced to $\rho_T Var(\varepsilon_T)$ is equal to the increase in welfare if $Var(\varepsilon_P)$ is reduced to $\rho_P Var(\varepsilon_P)$ if and only if

$$\mu = \frac{(1+\varphi_P)(1+\rho_P\varphi_P)(1-\rho_T)\varphi_T}{(1+\varphi_T)(1+\rho_T\varphi_T)(1-\rho_P)\varphi_P}\left(\lambda + \frac{(1-\delta)C_2}{C_2 - B_2^I}\right)^2 = \mu^*.$$

If $\mu > \frac{(1+\varphi_P)(1+\rho_P\varphi_P)(1-\rho_T)\varphi_T}{(1+\varphi_T)(1+\rho_T\varphi_T)(1-\rho_P)\varphi_P}\left(\lambda + \frac{(1-\delta)C_2}{C_2-B_2^I}\right)^2$, then receiving the signal about $\beta^P(\eta_P)$ yields greater gains, and if $\mu < \frac{(1+\varphi_P)(1+\rho_P\varphi_P)(1-\rho_T)\varphi_T}{(1+\varphi_T)(1+\rho_T\varphi_T)(1-\rho_P)\varphi_P}\left(\lambda + \frac{(1-\delta)C_2}{C_2-B_2^I}\right)^2$, then receiving the signal about $\beta^T(\eta_T)$ yields higher gains. The value of μ^* is increasing in λ, $\frac{C_2}{C_2-B_2^I}$, and ρ_P and decreasing in δ, and ρ_T.

A.C Proof of Corollary 1

With a charge t, the market equilibrium is $B_1^I + \beta^I(\eta_I) + B_2^I(q - \widehat{q}) = t$, or $q - \widehat{q} = \frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}$ and $q_{NS} = \widehat{q} + \frac{B_1^I + \beta^I(\eta_I)}{-B_2^I}$, and $\Delta_Q = \frac{-t}{-B_2^I}$.

Social welfare is $B^I(q, \eta_I) + B^P(q, \eta_P)$ if there is no congestion charge, or

$$\begin{aligned}
&b^I(\eta_I) + \left(B_1^I + \beta^I(\eta_I)\right)\left(\frac{B_1^I + \beta^I(\eta_I)}{-B_2^I}\right) + 0.5B_2^I\left(\frac{B_1^I + \beta^I(\eta_I)}{-B_2^I}\right)^2 \\
&= b^I(\eta_I) + b^P(\eta_P) + \left(B_1^P + \beta^P(\eta_P) + 0.5B_1^I + 0.5\beta^I(\eta_I)\right)\left(\frac{B_1^I + \beta^I(\eta_I)}{-B_2^I}\right) \\
&= b^I(\eta_I) + b^P(\eta_P) + \left(B_1^P + \widetilde{\beta}^P(\eta_P) + 0.5B_1^I + 0.5\widetilde{\beta}^I(\eta_I)\right)\left(\frac{B_1^I + \widetilde{\beta}^I(\eta_I)}{-B_2^I}\right) \\
&\quad + \frac{E\left(\beta^I(\eta_I)\left(0.5\beta^I(\eta_I) + \beta^P(\eta_P)\right)\right) - \widetilde{\beta}^I(\eta_I)\left(0.5\widetilde{\beta}^I(\eta_I) + \widetilde{\beta}^P(\eta_P)\right)}{-B_2^I}.
\end{aligned}$$

If a congestion charge is imposed, social welfare is $B^I(q, \eta_I) + B^P(q, \eta_P) - K$, and the optimal congestion charge maximizes the expectation of

$$b^I(\eta_I) + \left(B_1^I + \beta^I(\eta_I)\right)\left(\frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}\right) + 0.5B_2^I\left(\frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}\right)^2$$
$$+b^P(\eta_P) + \left(B_1^P + \beta^P(\eta_P)\right)\left(\frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}\right).$$

The first-order condition is

$$t = -B_1^P - E\left[\beta^P(\eta_P)\right] = -B_1^P - \widetilde{\beta}^P(\eta_P).$$

At the optimum, welfare is the expected value of

$$\left(B_1^I + \beta^I(\eta_I) + B_1^P + \beta^P(\eta_P)\right)\left(\frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}\right) + 0.5B_2^I\left(\frac{B_1^I + \beta^I(\eta_I) - t}{-B_2^I}\right)^2 + b^P(\eta_P)$$

or, equivalently,

$$b^I(\eta_I) + b^P(\eta_P) + \frac{\left(B_1^I + \widetilde{\beta}^I(\eta_I) + B_1^P + \widetilde{\beta}^P(\eta_P)\right)^2}{-2B_2^I} +$$
$$\frac{E\left[\beta^I(\eta_I)\left(0.5\beta^I(\eta_I) + \beta^P(\eta_P)\right)\right] - \widetilde{\beta}^I(\eta_I)\left(0.5\widetilde{\beta}^I(\eta_I) + \widetilde{\beta}^P(\eta_P)\right)}{-B_2^I} - K.$$

This expected welfare is larger than

$$b^I(\eta_I) + b^P(\eta_P) + \left(B_1^P + \widetilde{\beta}^P(\eta_P) + 0.5B_1^I + 0.5\widetilde{\beta}^I(\eta_I)\right)\left(\frac{B_1^I + \widetilde{\beta}^I(\eta_I)}{-B_2^I}\right) +$$
$$\frac{E\left[\beta^I(\eta_I)\left(0.5\beta^I(\eta_I) + \beta^P(\eta_P)\right)\right] - \widetilde{\beta}^I(\eta_I)\left(0.5\widetilde{\beta}^I(\eta_I) + \widetilde{\beta}^P(\eta_P)\right)}{-B_2^I}$$

if and only if $\frac{\left(B_1^P + \widetilde{\beta}^P(\eta_P)\right)^2}{-2B_2^I} > K$, or 0.5 multiplied by the expected change in ridership multiplied by the size of the externality must be greater than the administrative costs.

A.D Proof of Corollary 2

If there is no tax, social welfare will be $E\left[b^I(\eta_I) - 0.5B_2^I\left(\frac{B_1^I+\beta^I(\eta_I)}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right)^2\right]$. If there is tax, realized welfare is $b^I(\eta_I) + \left(-0.5B_2^I\right)\left(\frac{B_1^I+\beta^I(\eta_I)}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right)^2 - t\frac{(B_1^P+\beta^P(\eta_P))(B_1^I+\beta^I(\eta_I))}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2} + \frac{\left(B_1^P+\beta^P(\eta_P)+0.5B_2^I\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}t^2$.

The optimal tax will satisfy

$$E\left[\frac{\left(B_1^P+\beta^P(\eta_P)\right)\left(B_1^I+\beta^I(\eta_I)\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}\right] / E\left[\frac{2\left(B_1^P+\beta^P(\eta_P)+0.5B_2^I\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}\right] = t.$$

If $\beta^P(\eta_P) = 0$, then this tax equals $\frac{B_1^P E\left(B_1^I+\beta^I(\eta_I)\right)}{2B_1^P+B_2^I} = t$.

We use the notation $\Delta_Q(\eta_P) = \frac{t}{-B_2^I-B_1^P-\beta^P(\eta_P)}$ to denote the magnitude of the change in driving due to the charge. Then welfare can be written as

$$b^I(\eta_I) + \left(-0.5B_2^I\right)\left(\frac{B_1^I+\beta^I(\eta_I)}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right)^2 - t\frac{\left(\left(B_1^P+\beta^P(\eta_P)\right)\left(B_1^I+\beta^I(\eta_I)\right)\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2} + \frac{\left(B_1^P+\beta^P(\eta_P)+0.5B_2^I\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}t^2.$$

Consequently, expected welfare will equal

$$E\left[b^I(\eta_I)\right] + \left(-0.5B_2^I\right)E\left[\frac{B_1^I+\beta^I(\eta_I)}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right]^2 - 0.5tE\left[\frac{\left(B_1^P+\beta^P(\eta_P)\right)\left(B_1^I+\beta^I(\eta_I)\right)}{\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}\right] - K.$$

Welfare will be higher with the congestion charge if and only if

$$t\left(tE\left[\frac{1}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right] + tE\left[\frac{B_2^I}{2\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)^2}\right]\right) > K,$$

and thus K must be less than $E\left[\Delta_Q(\eta_P)\right]$ multiplied by $\frac{E\left[\frac{1}{-B_2^I-B_1^P-\beta^P(\eta_P)}\left(1+\frac{B_2^I}{2\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)}\right)\right]}{E\left[\frac{1}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right]}$.

When $B_2^I < 0$ and $B_1^P + \beta^P(\eta_P) < 0$, then $\left(1+\frac{1}{2\left(-1-\frac{B_1^P+\beta^P(\eta_P)}{B_2^I}\right)}\right)$ is monotonically increasing with B_2^I. This equals one when $B_2^I = 0$ and goes to one-half when B_2^I

goes to minus infinity, and so it must always lie between one-half and one. Consequently $\frac{E\left[\frac{1}{-B_2^I-B_1^P-\beta^P(\eta_P)}\left(1+\frac{B_2^I}{2\left(-B_2^I-B_1^P-\beta^P(\eta_P)\right)}\right)\right]}{E\left[\frac{1}{-B_2^I-B_1^P-\beta^P(\eta_P)}\right]}$ is bounded between one-half and one. If $\beta^P(\eta_P)=0$, this condition becomes[15] $\frac{-t^2(2B_1^P+B_2^I)}{2(B_1^P+B_2^I)^2}>K$, and at the optimal tax $t=B_1^P E\left(B_1^I+\beta^I(\eta_I)\right)/(2B_1^P+B_2^I)$, this simplifies to $\frac{-(B_1^P)^2\left(E\left(B_1^I+\beta^I(\eta_I)\right)\right)^2}{2(B_1^P+B_2^I)^2(2B_1^P+B_2^I)}>K$.

B CONTINUOUS INVESTMENT IN RESEARCH

In many cases, investments will be incremental and, in that case, it is likely to be optimal to invest in learning about all random variables. We now derive more quantitative results by postulating a research technology. We assume that $Var(\varepsilon_j)=Var(\beta^j(\eta_j))\frac{1}{e^{z_j i_j}-1}$ for $j=P,T$, where i_P and i_T represent total research spending on the size of the externality and market conditions, respectively. Each form of spending has the same cost of $1+\lambda$, and we assume that parameter values are such that there is a positive amount of research on all variables. We define estimation accuracy as the R^2 from a linear regression where the true value is regressed on the estimate, which will equal $1-e^{-z_j i_j}$, and inaccuracy as one minus that amount. We can then prove the following proposition:

Proposition 3. *There exists a value of μ, denoted μ^*, at which the decisionmaker invests equally in learning about the externality or learning more about market quantity. The decisionmaker invests more in learning about the externality if and only if $\mu>\mu^*$, and the value of μ^* is increasing in λ and $\frac{C_2}{C_2-B_2^I}$ and decreasing with δ.*

This proposition almost perfectly mirrors Proposition 2. The only major difference is that improvements in investment technology do not automatically lead to increased investment in that technology. That is because the function $-e^{-z_j i_j}$ is concave, and a higher level of z_j can either increase or reduce the marginal return to i_j depending on whether $1>z_j i_j$.

Proof of Proposition: The unconditional expected social welfare is now $E[V_0]-(1+$

[15] An earlier version of this expression compared the left-hand side to zero rather than K. Since the left-hand side is positive under the sign assumptions $B_2^I<0$ and $B_1^P<0$, that comparison was trivially satisfied; symbolic verification confirms the correct right-hand side is K.

$\lambda)(i_P + i_T)$:

$$\frac{Var\left(\beta^P\left(\eta_P\right)\right)\left(1-e^{-z_P i_P}\right)+\left(\frac{(1-\delta+\lambda)C_2-\lambda B_2^I}{C_2-B_2^I}\right)^2\left(Var\left(\beta^T\left(\eta_T\right)\right)\left(1-e^{-z_T i_T}\right)\right)}{4\left((1+2\lambda)(C_2-B_2^I)+(1-\delta)C_2\right)^2}.$$

This yields a first-order condition:

$$z_P \mu Var\left(\beta^P\left(\eta_P\right)\right)e^{-z_P i_P} = z_T\left(\frac{(1-\delta)C_2}{C_2-B_2^I}+\lambda\right)^2 Var\left(\beta^T\left(\eta_T\right)\right)e^{-z_T i_T} \tag{B.1}$$

$$= (1+\lambda)4\left((1+2\lambda)(C_2-B_2^I)+(1-\delta)C_2\right)^2, \tag{B.2}$$

or letting $i_P = i_T + x$,

$$\frac{z_P \mu}{z_T\left(\frac{(1-\delta)C_2}{C_2-B_2^I}+\lambda\right)^2}e^{(z_T-z_P)i_T} = e^{z_P x}. \tag{B.3}$$

If $\frac{z_P \mu}{z_T\left(\frac{(1-\delta)C_2}{C_2-B_2^I}+\lambda\right)^2}e^{(z_T-z_P)i_T} = 1$, then $e^{z_P x} = 1$ and $x = 0$, so the two forms of investment are equal. For all values of μ greater than $\frac{z_T\left(\frac{(1-\delta)C_2}{C_2-B_2^I}+\lambda\right)^2}{z_P}e^{(z_P-z_T)i_T}$, which we define as μ^{**}, we have $e^{z_P x} > 1$ and $x > 0$, and for all values of μ less than $\frac{z_T\left(\frac{(1-\delta)C_2}{C_2-B_2^I}+\lambda\right)^2}{z_P}e^{(z_P-z_T)i_T}$, $e^{z_P x} < 1$ and $x < 0$. The value of μ^{**} is increasing with λ and $\frac{C_2}{C_2-B_2^I}$ and decreasing with δ.

C CONSTRUCTION OF THE PROJECT-LEVEL RETURN-ON-CAPITAL GAP

This appendix details the construction of the project-level return on capital and return-on-capital gap shown in Figure 3 and discussed in Section 4.F. Project-level revenues, costs, and financing terms for projects in the pipeline are not directly observed; we estimate them using observable project characteristics, market data, and BPDA filings.

Net operating income. For each project i, NOI_i is built up from four components: apartment NOI, parking NOI, commercial NOI (where applicable), and an expense ratio applied to gross rent. Each is estimated by a regression on observables:

$$\begin{aligned}
\text{rent}_i &= \alpha_R + \beta_R^u \cdot \text{units}_i + \beta_R^b \cdot \text{bedroom mix}_i + \xi_R^{\text{nbhd}(i)}, \\
\text{expense_pct}_i &= \alpha_E + \beta_E^b \cdot \text{bedrooms}_i + \beta_E^r \cdot \text{rent}_i + \beta_E^u \cdot \text{units}_i + \beta_E^{\text{idp}} \cdot \text{idp}_i + \xi_E^{\text{nbhd}(i)}, \\
\text{park_noi_per_space}_i &= \alpha_P + \beta_P \cdot \text{units}_i + \xi_P^{\text{nbhd}(i)}, \\
\text{comm_noi_per_sf}_i &= \alpha_C + \beta_C \cdot \text{units}_i + \xi_C^{\text{nbhd}(i)}.
\end{aligned}$$

Each regression is fitted on the corresponding income groups in BPDA Tax Increment Program filings (RT2 for apartments, RPK for parking, RT1 for commercial), with neighborhood-region fixed effects ξ. Predicted values are then aggregated to the project-level NOI:

$$NOI_i = 12 \cdot \widehat{\text{rent}}_i \cdot (1 - \widehat{\text{expense\_pct}}_i) \cdot \text{rent\_units}_i + 12 \cdot \widehat{\text{park\_noi\_per\_space}}_i \cdot \text{parking\_spaces}_i + \widehat{\text{comm\_noi\_p}}$$

For four large projects whose pro forma filings were available directly (220 Huntington Avenue, 323–365 Dorchester Avenue, 1234 Soldiers Field Road, and 1170 Soldiers Field Road), we replace the estimated NOI_i with the filed value. NOI is scaled into each calendar quarter by the ratio of the contemporaneous CoStar neighborhood rent index to its level at the project's BPDA approval date.

Construction costs. For projects that have already been built, the total project cost can be backed out of the BPDA permitting fee, which is approximately 1 percent of project value. Using this constructed cost per square foot as a left-hand variable, we fit a hedonic regression on project characteristics:

$$\text{cost\_per\_sqft}_i = \alpha_C + \gamma_1 \text{units}_i + \gamma_2 \text{idp\_pct}_i + \gamma_3 \text{rental}_i + \sum_h \gamma_h^H \nVdash\{x_h\}_i + \gamma_4 \text{parking}_i + \sum_n \gamma_n^N \nVdash\{n\}_i + \sum_t \gamma_t^Y \nVdash$$

where x_h is a height quartile (with separate height-missing and parking-missing flags), $\nVdash\{n\}$ is a coarse-neighborhood dummy, and $\nVdash\{t\}$ is a permit-year dummy. Predicted cost per square foot is then multiplied by the project's leasable square footage to yield estimated hard costs, with the permit year fixed at 2021 and a 12 percent inflation adjustment for projects in the post-2020 pipeline (based on the FRED residential investment chain price index).

Total cost and the moneyness gap. Soft costs are added at 12 percent of hard costs; land costs are the assessor's land value lotvalue_i. Construction-period financing is added at 60 percent loan-to-cost at the two-year LIBOR rate plus 200 basis points over a 22-month build, giving a total cost

$$C_{i,t} = \frac{1.12 \cdot \widehat{\text{hard}}_{i,t} + \text{lotvalue}_i \cdot \frac{\text{sale\_price\_index}_t}{\text{sale\_price\_index}_0}}{1 - 0.6 \cdot (\text{libor}_t^{2y}/100 + 0.02) \cdot 22/12}.$$

The project-level return on capital and gap are then

$$ROC_{i,t} = \frac{NOI_{i,t}}{C_{i,t}}, \qquad \text{ROC gap}_{i,t} = ROC_{i,t} - (r_t - g_t + \tau_t),$$

where $r_t - g_t + \tau_t$ is the CoStar neighborhood market cap rate. Because every component is estimated rather than directly observed, the return-on-capital gap is best read as a back-of-the-envelope summary of how project economics evolved over time, not as a precise project-level measure of viability.