Remote Work and Optimal Congestion Pricing: Evidence from the COVID-19 Demand Shock

David Bernal*

Abstract

This paper shows how changes in the value of travel time alter the congestion prices implied by a standard marginal-external-cost calculation. I combine hourly freeway traffic data from the California Performance Measurement System with county-level remote-work shares and wages from the American Community Survey for five large California counties. The COVID-19 shift toward remote work provides a plausibly exogenous demand shock: a shift–share instrumental-variables design implies that a 10 percentage point increase in working from home reduces freeway vehicle miles traveled by approximately 1.25 percent. Holding the value of time fixed at its pre-pandemic level, the associated reduction in congestion lowers the implied optimal per-mile tax in every county. Allowing the value of time to evolve with local wages reverses that conclusion: the post-pandemic tax is higher in every county, and the fixed-value-of-time calculation understates the time-varying benchmark by approximately 17–18 percent. An exact decomposition shows that the increase in the monetary value of delay more than offsets the reduction in marginal congestion delay. The results demonstrate that updating traffic conditions without updating the value of time can produce systematically misleading congestion-price recommendations.

Keywords: Freeway congestion, congestion pricing, optimal congestion taxes, vehicle miles traveled, travel time index, deadweight loss, COVID-19, remote work, welfare analysis

JEL Classification: R40, R41, D62, H43

*University of Houston, Houston Texas, United States. E-mail: vdbernalpavas@uh.edu

1 Introduction

Traffic congestion is one of the largest and most visible costs of urban activity. A driver deciding whether and when to use a crowded road considers her own travel time but not the delay she imposes on other road users. Congestion is therefore a canonical externality: beginning with Pigou (1920) and developed for urban transportation by Vickrey (1969), economists have argued that a charge equal to marginal external congestion cost can decentralize the efficient level of road use. Modern empirical work has made that prescription operational by combining observed traffic flows, congestion technologies, demand responses, and a monetary value of travel time (Small, Verhoef, & Lindsey, 2024; Yang, Purevjav, & Li, 2020).

The value of time is central to this calculation. A minute of delay creates a larger monetary externality when the affected drivers place a higher value on time. Yet applied calculations commonly map delay into dollars using a single average value of time that is held fixed within the policy exercise. That convention is convenient, but it can be consequential when the composition or earnings of travelers change. Revealed-preference studies find substantial heterogeneity in motorists' willingness to pay for faster and more reliable travel (Bento, Roth, & Waxman, 2024; Lam & Small, 2001; Small, Winston, & Yan, 2005), and theoretical work shows that heterogeneity can change both the efficiency and distributional incidence of road pricing (Small & Yan, 2001; van den Berg & Verhoef, 2011; Verhoef & Small, 2004). Thus, updating the congestion technology while leaving the monetary value of delay unchanged need not deliver the correct change in the optimal toll.

This paper quantifies that wedge using the COVID-19 transformation of commuting. The pandemic produced a large, sudden, and persistent expansion of remote work, shifting freeway demand without a corresponding expansion of road capacity. I compare pre-pandemic conditions in 2018–2019 with post-pandemic conditions in 2021–2022 in Los Angeles, Orange, San Diego, San Francisco, and Santa Clara counties.[1] The empirical analysis combines

[1] A complete list of included freeways is provided in Table B.1 in the appendix.

hourly measures of vehicle miles traveled (VMT), the Travel Time Index (TTI), and speed from the California Performance Measurement System (PeMS) with county-year work-from-home shares and wages from the American Community Survey (ACS).

I proceed in two steps. First, I estimate how the remote-work shock shifted freeway demand. Because realized remote work may respond to local economic conditions, I instrument the county work-from-home share using pre-pandemic industry composition interacted with national changes in work-from-home propensity. The two-stage least-squares estimate implies that a 10 percentage point increase in working from home reduces freeway VMT by approximately 1.25 percent. The largest reductions in congestion occur during the traditional morning and afternoon peaks, consistent with the removal of commute trips rather than a uniform contraction in travel.

Second, I map the observed change in congestion into marginal external costs and implied Pigouvian taxes. Holding the value of time at its 2018 level, lower post-pandemic congestion reduces the optimal tax in every county. That is the prediction a conventional fixed-value-of-time calculation would deliver. Once the value of time is updated using contemporaneous county wages, however, the result reverses: the implied post-pandemic tax rises in all five counties. Relative to this time-varying benchmark, fixing the value of time understates the post-pandemic tax by approximately 17–18 percent. An exact decomposition shows why. The decline in marginal external delay pushes the tax downward, but the increase in the monetary value assigned to that delay is larger.

The empirical exercise isolates temporal variation in an average value of time; it does not estimate a full distribution of driver-specific valuations. Nevertheless, heterogeneity sharpens the paper's mechanism. Drivers with high values of time are generally more willing to pay to avoid delay and may be less responsive to a toll. If pricing changes who remains on the road, the value of time relevant for the marginal externality can itself become policy dependent. The fixed-average calculation in this paper is therefore best viewed as a transparent benchmark for a broader problem: congestion prices should reflect both the physical delay generated by traffic and the economic value placed on that delay by the drivers who use the road.

The paper contributes to three literatures. First, it contributes to work measuring congestion costs and the welfare consequences of road pricing (Newbery, 1989; Verhoef, 2002; Yang et al., 2020). Second, it connects the empirical literature documenting heterogeneity in values of time and reliability to the calibration of optimal congestion charges (Bento et al., 2024; Small, 2012; Small et al., 2005). Third, it complements research on induced travel, urban speed, and pandemic-era mobility by using remote work as a demand shifter rather than treating the post-pandemic change in congestion as purely descriptive (Akbar, Couture, Duranton, & Storeygard, 2023; Couture, Duranton, & Turner, 2018; Dingel & Neiman, 2020; Duranton & Turner, 2011).

The remainder of the paper is organized as follows. Section 2 develops the pricing framework and the decomposition of the optimal tax. Section 3 describes PeMS and ACS data and documents the demand shock. Section 4 presents the congestion and demand estimates. Section 5 compares optimal taxes under fixed and time-varying values of time. Section 6 concludes.

2 Model

Computing an optimal congestion charge requires three objects: a congestion technology mapping traffic into travel time, a demand curve describing how road use responds to generalized travel cost, and a value of time that converts delay into money. This section combines the standard marginal-external-cost framework of Newbery (1989) and Quinet and Vickerman (2004) with an empirically tractable BPR-type congestion technology and a constant-elasticity demand curve. The framework is designed to express the optimal per-mile tax using quantities available in PeMS and the ACS. Its distinctive feature is that the value of time is indexed by county and year rather than fixed throughout the policy comparison. This makes it possible to separate changes in the optimal tax caused by the physical congestion environment from changes caused by the monetary value assigned to travel delay.

Building on the frameworks of Quinet and Vickerman (2004) and Newbery (1989), the marginal social cost of congestion is defined as

$$MSC(F) = \frac{\partial TSC}{\partial F} = PMC(F) + MECC(F), \tag{1}$$

where MSC denotes marginal social cost, $PMC(F)$ is the private marginal cost faced by a driver, and $MECC(F)$ is the marginal external congestion cost. In a nonatomic traffic model, an individual driver takes prevailing travel time as given, so $PMC(F)$ is also the average user cost. The private term is internalized when making a trip, whereas $MECC(F)$ captures the additional time cost imposed on other road users. Following Yang et al. (2020),

$$MECC(F) = F \cdot VOT \cdot \frac{\partial T}{\partial F}. \tag{2}$$

Appendix A derives Equations (1) and (2) from aggregate travel-time cost.

The $MECC$ framework provides a foundation to evaluate both the welfare implications of congestion and the corresponding adjustment in the optimal congestion price (Pigouvian tax). To operationalize Equations (1) and (2), I specify the functional form of $T(F)$ using a Bureau of Public Roads (BPR) function.

2.1 Travel Time and Value of Time

I model the relationship between traffic flow and delay using a BPR-type equation:

$$TTI_{c,h,t} = 1 + \alpha \left(F_{\text{obs},c,h,t}\right)^{\beta}, \tag{3}$$

where $TTI_{c,h,t}$ denotes the Travel Time Index in county c, hour h, and year t. The term $TTI_{c,h,t} - 1$ represents proportional delay relative to free-flow conditions (for example, $TTI = 1.2$ implies 20% longer travel time). In the conventional BPR expression, delay depends on $(F/C)^{\beta}$, where C is capacity. Because capacity is not observed in the county-level PeMS extract, I use α as the composite scale parameter $\widetilde{\alpha}C^{-\beta}$. The welfare calculation requires this composite congestion curve, not separate identification of engineering capacity and the conventional BPR intercept.

Observed traffic flow $F_{\text{obs},c,h,t}$ is measured as vehicle miles traveled (VMT) per hour and defined as

$$F_{\text{obs},c,h,t} = VMT_{\text{Passenger},c,h,t} + \delta \cdot VMT_{\text{Truck},c,h,t}, \tag{4}$$

where δ is a truck-to-car equivalence factor reflecting the higher congestion contribution of trucks. Following FHWA guidelines, I set $\delta = 1.5$.[2]

Free-flow travel time per mile is given by $1/S_{\text{free-flow}}$, where $S_{\text{free-flow}}$ denotes the free-flow speed (60 mph in the PeMS data). Therefore, travel time per vehicle mile is

$$T_{c,h,t} = T_{\text{free-flow},c,t} \cdot TTI_{c,h,t}. \tag{5}$$

2.2 Computing Costs

The value of time (VOT) is defined as a constant fraction of the local wage:

$$VOT_{c,t} = \theta \cdot w_{c,t}, \tag{6}$$

where $\theta = 0.5$ following Small (2012). This implies that time is valued at half the median hourly wage in county c and year t.

[2] Typical values range from 1.5 to 2 depending on roadway type.

The derivative of travel time with respect to flow is

$$\frac{dT_{c,h,t}}{dF} = \frac{1}{S_{\text{free-flow}}} \cdot \alpha \cdot \beta \cdot \left(F_{\text{obs},c,h,t}\right)^{\beta-1}. \tag{7}$$

The private marginal cost per vehicle mile is the monetary value of the travel time experienced by the driver:

$$PMC_{c,h,t}(F) = VOT_{c,t}\, T_{c,h,t}(F) = \frac{VOT_{c,t}}{S_{\text{free-flow}}}\left[1 + \alpha F^{\beta}_{\text{obs},c,h,t}\right]. \tag{8}$$

Substituting Equations (3), (7), and the definition of VOT into Equation (1) yields

$$MSC_{c,h,t} = \underbrace{\frac{VOT_{c,t}}{S_{\text{free-flow}}}\left[1 + \alpha \left(F_{\text{obs},c,h,t}\right)^{\beta}\right]}_{\text{Private Marginal Cost (PMC)}} + \underbrace{\frac{VOT_{c,t}}{S_{\text{free-flow}}}\alpha\beta \left(F_{\text{obs},c,h,t}\right)^{\beta}}_{\text{Marginal External Congestion Cost (MECC)}}. \tag{9}$$

The elasticity β and the composite scale of the congestion curve are estimated from Equation (3) using OLS.

2.3 Travel Decisions and Demand Specification

Drivers base their decisions on the private marginal cost $PMC(F_{c,t})$ that they personally bear. Travel demand is therefore expressed as

$$F_{c,t} = A_{ct} \cdot PMC_{ct}^{-\eta}, \tag{10}$$

where $\eta > 0$ is the absolute value of the elasticity of freeway demand with respect to private marginal cost. I set $\eta = 0.1$, retaining the conservative short-run elasticity used in the baseline calibration. The demand shifter A_{ct} captures county-specific and temporal variation and is parameterized in logs:

$$\ln A_{ct} = \mu_c + \sigma Z_{ct} + X'_{ct}\gamma, \tag{11}$$

where μ_c are county fixed effects, X_{ct} are observable controls, and Z_{ct} is the share of people working from home. This log specification ensures that A_{ct} is positive and makes the structural demand equation consistent with the empirical log-linear specification below.

To address potential endogeneity of Z_{ct}, I first define

$$Z_{ct} \equiv \text{WFH share}_{ct}, \tag{12}$$

the share of employed residents in county c and year t whose usual place of work is their home. I then instrument Z_{ct} using a Bartik-type (shift–share) variable,

$$Z_{ct}^{IV} = \sum_j s_{cj}^{2019}\, \tau_{jt}^{WFH}, \tag{13}$$

where s_{cj}^{2019} is the 2019 employment share of industry j in county c, and τ_{jt}^{WFH} is the national change in work-from-home propensity in industry j at time t, constructed using teleworkability indices from Dingel and Neiman (2020).

Intuitively, Z_{ct}^{IV} captures how the work-from-home share in county c would evolve if its pre-pandemic industry structure were exposed only to national shocks to telework feasibility. The identifying assumption is that, conditional on controls and fixed effects, these national industry-level shocks are orthogonal to county-specific unobservables that directly affect congestion, so that Z_{ct}^{IV} shifts freeway travel demand only through its effect on remote work.

2.4 Freeway Optimal Congestion Tax

Let $P(F)$ denote inverse demand, or the marginal willingness to pay for an additional vehicle mile. In the unpriced equilibrium, $P(F_{c,t}) = PMC(F_{c,t})$. A toll τ raises the cost perceived by drivers to $PMC(F) + \tau$; setting $\tau = MECC(F)$ makes drivers face the full marginal social cost. In the social optimum, the flow $F_{c,t}^{opt}$ therefore satisfies

$$P(F_{c,t}^{opt}) = MSC(F_{c,t}^{opt}). \tag{14}$$

Given the demand in Equation (10), the optimal flow can be written as

$$F_{c,t}^{opt} = A_{ct} \left[\frac{1}{MSC(F_{c,t}^{opt})} \right]^{\eta}, \tag{15}$$

and the optimal Pigouvian tax is

$$\tau_{c,t} = MECC(F_{c,t}^{opt}). \tag{16}$$

Under a local linearity assumption for the marginal external congestion cost between $F_{c,t}^{opt}$ and $F_{\text{obs},c,t}$, the deadweight loss (DWL) associated with unpriced congestion can be expressed as

$$DWL_{c,t} = \frac{1}{2} \cdot MECC(F_{\text{obs},c,t}) \cdot (F_{\text{obs},c,t} - F_{c,t}^{opt}). \tag{17}$$

Figure 1. Average and Marginal Social Costs of Congestion

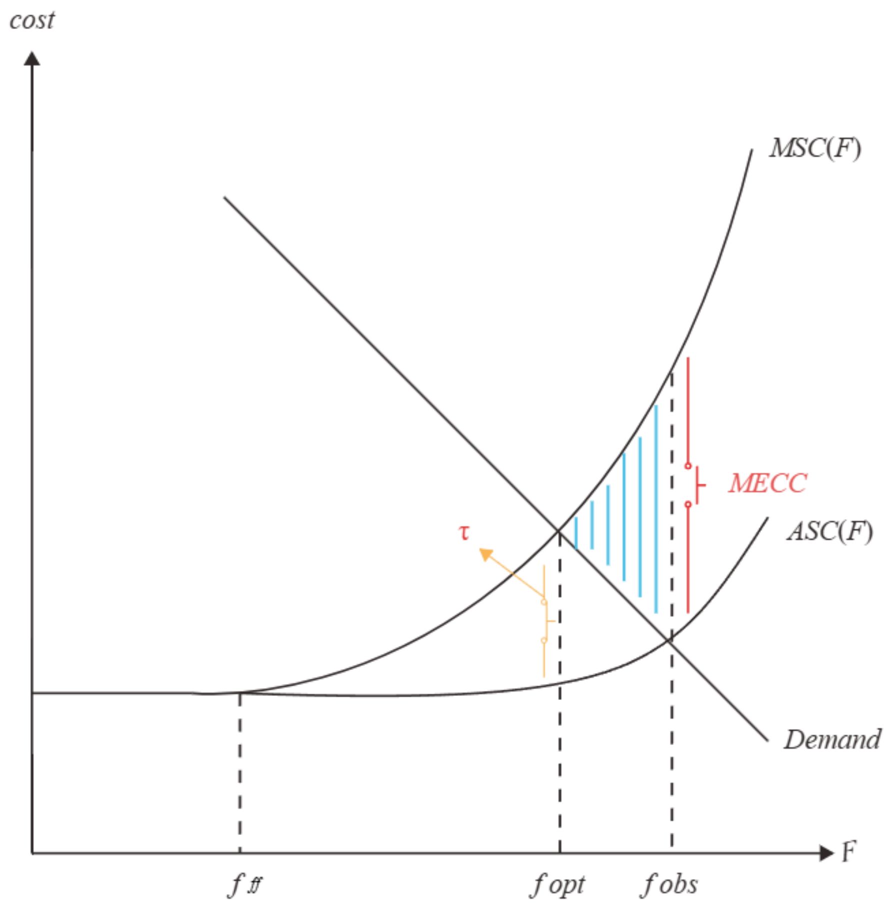

Notes: The curve labeled ASC in the schematic is the average user cost, which is identical to the private marginal cost (PMC) in the nonatomic model. Marginal social cost (MSC) adds the marginal external congestion cost (MECC). The shaded area represents the deadweight loss from unpriced congestion.

Following the pandemic, a negative demand shock shifts the travel demand curve downward, reducing observed flow and congestion. Holding VOT constant, changes in $MECC$ are driven entirely by this demand shift. The pandemic thus provides a natural experiment to identify the change in the optimal congestion tax, τ.

Figure 2. Optimal Tax Diagram with Private and Social Costs

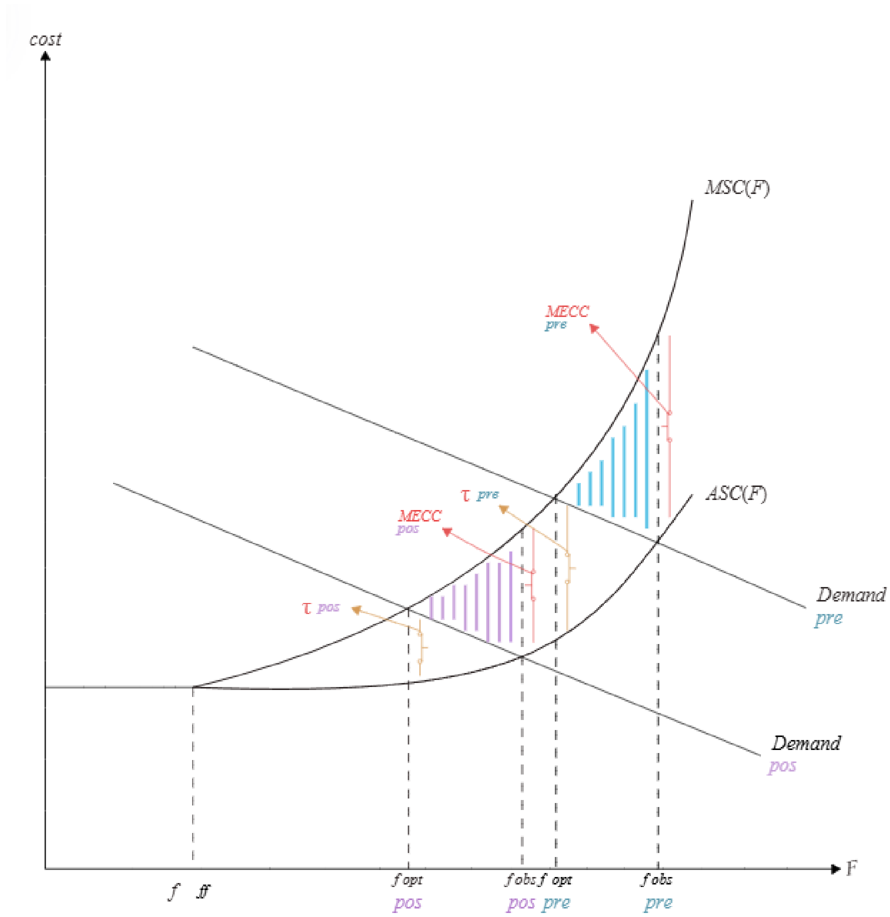

Notes: This figure compares the pre- and post-pandemic equilibria. The difference between $\tau_{\text{pre}} = MECC(F_{\text{opt,pre}})$ and $\tau_{\text{post}} = MECC(F_{\text{opt,post}})$ reflects how optimal congestion pricing responds to the remote work–induced demand shock.

Finally, I define the change in the optimal congestion tax as $\Delta\tau_c = \tau_{\text{post}} - \tau_{\text{pre}}$, and the welfare gain from the reduction in unpriced congestion as

$$\Delta W_c = \sum_{t\in\text{pre}} DWL_{c,t} - \sum_{t\in\text{post}} DWL_{c,t}. \tag{18}$$

This framework links observed congestion data with optimal pricing and welfare, highlighting how shifts in travel demand, particularly through remote work, affect the efficiency of freeway usage.

2.5 Toll Pass-Through and the Congestion-Cost Offset

The toll and private marginal cost enter the travel decision together, but they do not move one-for-one in equilibrium. Suppressing county and time subscripts, let the price perceived by a driver be $p(F, \tau) = PMC(F) + \tau$, where τ is a per-mile toll. Equation (10) then becomes

$$F(\tau) = A\left[PMC(F(\tau)) + \tau\right]^{-\eta}. \tag{19}$$

At a given flow, the toll does not enter the private travel-time cost directly:

$$\left.\frac{\partial PMC}{\partial \tau}\right|_{F} = 0.$$

The relevant object for incidence, however, is the total equilibrium derivative. Differentiating Equation (19) and collecting the terms involving $dF/d\tau$ yields

$$\frac{dF}{d\tau} = -\frac{\eta F}{PMC(F) + \tau + \eta F PMC'(F)} < 0. \tag{20}$$

The denominator contains the usual demand response and an endogenous congestion feedback. A higher toll reduces flow, but the resulting improvement in travel conditions lowers PMC and partially offsets the initial price increase. By the chain rule,

$$\frac{dPMC}{d\tau} = PMC'(F)\frac{dF}{d\tau} = -\frac{\eta F PMC'(F)}{PMC(F) + \tau + \eta F PMC'(F)} < 0. \tag{21}$$

Thus, $-dPMC/d\tau$ is the fraction of a marginal toll increase returned to drivers through lower congestion costs. The corresponding pass-through into the full price perceived by drivers is

$$\frac{dp}{d\tau} = 1 + \frac{dPMC}{d\tau} = \frac{PMC(F) + \tau}{PMC(F) + \tau + \eta F PMC'(F)}, \qquad 0 < \frac{dp}{d\tau} \leq 1. \tag{22}$$

Full pass-through occurs under free-flow conditions, when $PMC'(F) = 0$. When congestion is responsive to traffic, some of the toll is offset by a reduction in the time cost of travel.

The BPR specification makes the offset transparent. Define the free-flow private cost as $PMC_0 = VOT/S_{\text{free-flow}}$. Equation (8) implies

$$FPMC'(F) = \beta \left[PMC(F) - PMC_0\right]. \tag{23}$$

Consequently, imposing $\eta = 0.1$ gives

$$\frac{dF}{d\tau} = -\frac{0.1F}{PMC(F) + \tau + 0.1\beta[PMC(F) - PMC_0]}, \tag{24}$$

$$\frac{dPMC}{d\tau} = -\frac{0.1\beta[PMC(F) - PMC_0]}{PMC(F) + \tau + 0.1\beta[PMC(F) - PMC_0]}. \tag{25}$$

The congestion-cost offset is therefore larger when demand is more price responsive, the delay–flow relationship is steeper, or congestion accounts for a larger share of private travel cost. With the conservative elasticity of 0.1, the short-run flow response is limited unless traffic is located on a steep portion of the congestion curve. Finally, the two price components have different welfare interpretations: PMC is a real time cost, whereas toll payments are transfers to the pricing authority. A toll can therefore increase the monetary price faced by drivers while reducing the real resources lost to congestion.

2.6 Wages and the Monetary Value of Delay

From Equation (6), a higher local wage raises $VOT_{c,t}$ and therefore scales both private and external time costs. Holding flow fixed, the effect on marginal social cost is

$$\left.\frac{\partial MSC_{c,h,t}}{\partial w_{c,t}}\right|_F = \frac{\theta}{S_{\text{free-flow}}}\left[1 + \alpha(1+\beta)F_{c,h,t}^{\beta}\right] > 0. \tag{26}$$

Likewise, the direct effect on the congestion charge at a given flow is

$$\left.\frac{\partial MECC_{c,h,t}}{\partial w_{c,t}}\right|_{F} = \frac{\theta}{S_{\text{free-flow}}} \alpha \beta F_{c,h,t}^{\beta} > 0. \tag{27}$$

These are partial effects. A higher VOT also changes efficient flow, which in turn changes congestion and the tax base. The total change in the optimal tax is therefore not signed by the equations above alone. The empirical analysis resolves that ambiguity by solving for optimal flow in each county-period and then decomposing the resulting tax change into a congestion component and a VOT component.

2.7 Decomposition of Changes in the Optimal Congestion Tax

To identify which economic force contributes most to the change in the tax, write the optimal Pigouvian congestion tax per mile at time t as

$$\tau_t = VOT_t \, g_t, \tag{28}$$

where VOT_t is the value of travel time (dollars per hour) and g_t denotes the marginal external travel delay per mile (hours per mile) implied by the congestion technology evaluated at optimal traffic conditions. In the model underlying equation (9), g_t is given by:

$$g_t = \frac{1}{S_{\text{free-flow}}} \alpha \beta \left(F_t^{\text{opt}}\right)^{\beta} \tag{29}$$

Two-period change. Consider two periods, a pre-period and a post-period. Define

$$\Delta\tau \equiv \tau_{\text{post}} - \tau_{\text{pre}}. \tag{30}$$

Exact decomposition (counterfactual approach). Define the counterfactual tax that holds VOT at its pre-period level while allowing congestion to change:

$$\tau^g \equiv VOT_{\text{pre}}\, g_{\text{post}}. \tag{31}$$

Then,

$$\begin{aligned}\Delta\tau &= \tau_{\text{post}} - \tau_{\text{pre}} \\ &= \underbrace{\left(\tau^g - \tau_{\text{pre}}\right)}_{\Delta\tau_{\text{cong}}} + \underbrace{\left(\tau_{\text{post}} - \tau^g\right)}_{\Delta\tau_{\text{vot}}}.\end{aligned} \tag{32}$$

Substituting from (28) and (31) yields

$$\Delta\tau_{\text{cong}} = VOT_{\text{pre}}\left(g_{\text{post}} - g_{\text{pre}}\right), \tag{33}$$

$$\Delta\tau_{\text{vot}} = g_{\text{post}}\left(VOT_{\text{post}} - VOT_{\text{pre}}\right). \tag{34}$$

Equations (33)–(34) provide an exact decomposition:

$$\Delta\tau = \Delta\tau_{\text{cong}} + \Delta\tau_{\text{vot}}. \tag{35}$$

3 Data

This section introduces the data used in the analysis and documents how the COVID-19 pandemic produced a large, sudden, and persistent reduction in congestion across California's main metropolitan counties. The observed patterns are consistent with a negative demand shock to travel behavior that reduced vehicle flows and travel times.

3.1 Data Sources

The main dataset used in this study is the California Performance Measurement System (PeMS), a real-time database managed by Caltrans that records hourly traffic conditions from more than thirty-nine thousand loop detectors across the state. PeMS provides information on Vehicle Miles Traveled (VMT), Vehicle Hours Traveled (VHT), average speed, and the Travel Time Index (TTI), which measures the ratio of actual travel time to free-flow travel time. In PeMS, the free-flow speed is fixed at sixty miles per hour, so TTI values above one indicate congestion.

To complement PeMS, I use data from the American Community Survey (ACS) to obtain median wages, population, and the share of people working from home. These variables are later used to compute the value of time (VOT) and to interpret behavioral changes in travel patterns.

3.2 The Pandemic as a Natural Experiment

The COVID-19 pandemic created a unique setting to study congestion and welfare. In March 2020, California implemented a statewide lockdown that sharply reduced mobility across all urban areas. Because no major freeway expansions, closures, or pricing policies occurred during this period, the pandemic can be interpreted as an exogenous event that altered travel behavior through restrictions on movement and the expansion of remote work, rather than through changes in road supply.

Appendix Figure B.1 shows the evolution of the work-from-home share from 2016 to 2023. Before the pandemic, remote work remained below seven percent in every county. In 2021, it exceeded forty-five percent in San Francisco and reached approximately thirty-five percent in Santa Clara. Although these shares declined after 2021, they remained well above their pre-pandemic levels through 2023. The persistence of this change distinguishes the episode from a transitory traffic interruption and motivates the pre/post comparison used below.

3.3 Changes in Traffic and Congestion

Traffic data reveal a clear and lasting change in freeway conditions following the pandemic. Figure 3 summarizes the three central traffic measures. VMT declined sharply during the lockdown and remained below its pre-pandemic trend, while average speed increased and TTI fell. The three series therefore tell a consistent story: the post-pandemic period combined lower freeway use with shorter travel times.

Figure 3. Freeway Traffic and Congestion

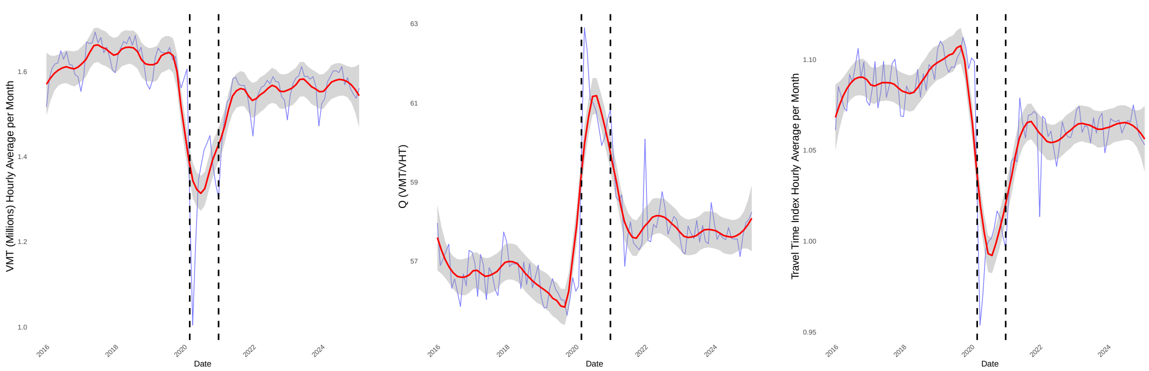

(a) Vehicle Miles Traveled (b) Average Speed (c) Travel Time Index

Notes: Monthly county averages from PeMS for 2016–2025. Red lines show smoothed trends. Lower TTI and higher speed indicate less congestion.

Table 1 presents a before-and-after comparison for VMT and average speed using hourly data for the years 2018 to 2019 and 2021 to 2022. The results show a significant decline in VMT and an increase in speed, confirming the descriptive evidence. The changes are strongest during rush-hour periods, when congestion had previously been most severe.

Appendix Figures B.2 and B.3 show how the post-pandemic changes vary over the day. Congestion relief is concentrated in the morning and afternoon commute peaks: TTI falls and speed rises most strongly precisely when the pre-pandemic network was most congested.

Table 1: Differences in Traffic Before and After the Pandemic

	VMT	Q	VMT	Q	VMT	Q
	All Data		(3pm–6pm)		Weekends	
	(1)	(2)	(3)	(4)	(5)	(6)
Pandemic After	-0.067^{***}	0.028^{***}	-0.040^{***}	0.066^{***}	-0.054^{***}	0.006^{***}
	(0.006)	(0.001)	(0.013)	(0.002)	(0.012)	(0.001)
Constant	13.642^{***}	4.080^{***}	14.200^{***}	3.902^{***}	13.569^{***}	4.134^{***}
	(0.004)	(0.0004)	(0.009)	(0.001)	(0.008)	(0.001)
Observations	181,655	181,655	30,280	30,280	51,695	51,695
R^2	0.001	0.011	0.0003	0.051	0.0004	0.001
Adjusted R^2	0.001	0.011	0.0003	0.051	0.0004	0.001
Note:						*p<0.1; **p<0.05; ***p<0.01

Notes: Each pair of columns reports differences in means of $\log(VMT)$ and $\log(Q)$, where Q denotes average speed. Columns 1–2 use all hourly observations, columns 3–4 use the 3–6 p.m. peak, and columns 5–6 use weekends. The comparison excludes 2020 and contrasts 2018–2019 with 2021–2022.

VMT also falls during commuting hours but changes little late at night. This timing matters economically. Because the delay-flow relationship is nonlinear, removing a vehicle during the peak produces a larger reduction in the delay imposed on others than removing the same vehicle when the road is uncongested.

3.4 Differences Across Counties

The reduction in congestion was not uniform across counties. Figure 4 shows that San Francisco and Santa Clara experienced the largest improvements, followed by Los Angeles, San Diego, and Orange. These differences closely match the intensity of remote work adoption, with counties that had the highest increase in work-from-home rates also showing the greatest decline in congestion.

Figure 4. Change in Congestion by County

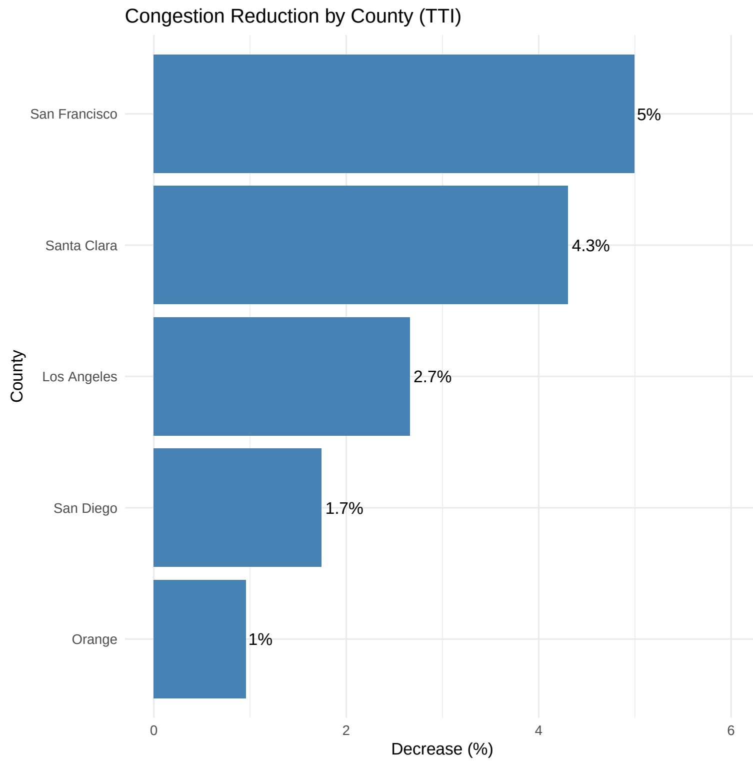

Notes: Estimated changes in the Travel Time Index by county. Counties with larger increases in remote work saw greater reductions in congestion.

The cross-county pattern is consistent with the time-series evidence. San Francisco and Santa Clara combine the largest increases in remote work with the largest reductions in TTI. Appendix Figure B.4 presents the corresponding county-year scatterplots. These correlations are descriptive; the shift–share design in Section 4 is used to isolate the component of remote work driven by predetermined local exposure to national telework shocks.

Alternative explanations for the decline in congestion are less consistent with the data. Freight activity, measured by truck VMT, accounts for less than five percent of total VMT and declined only slightly after the pandemic. Figure B.5 shows the reduction in VMT separately for trucks and non-truck vehicles, with patterns that closely mirror the aggregate results presented above. Population levels fell modestly between 2020 and 2021 but recovered in subsequent years, and average commuting distances remained stable. Taken together, these facts point to a primarily behavioral rather than structural source of the observed congestion relief.

4 Estimation

4.1 Supply side

I start by estimating the BPR equation in (3) using OLS, as specified below:

$$\ln(TTI_{c,h,t} - 1) = \rho + \beta \cdot \ln(F_{\text{obs},c,h,t}) + \gamma_c + \epsilon_{c,h,t}, \tag{36}$$

where γ_c denotes county fixed effects. These fixed effects absorb time-invariant cross-county differences in baseline travel conditions, including persistent differences in highway infrastructure, network design, and other structural features of the road system. Because major changes in aggregate freeway capacity are limited over the relatively short analysis window, treating these county-specific factors as stable is a reasonable identifying assumption.

The estimated congestion elasticity is $\beta = 3.6$, within the range commonly used in congestion-pricing applications (approximately 2.5 to 5) (Small et al., 2024). The regression intercept and county fixed effect jointly recover a county-specific composite scale, $\alpha_c = \exp(\rho + \gamma_c)$. Because this term combines the conventional BPR intercept with unobserved effective capacity, I do not interpret α_c as a transferable engineering parameter; it is used to trace each county's fitted delay-flow relationship over the observed range.

For the baseline calibration, I set $\theta = 0.5$ and use the 2018 county median wage. Section 5 first holds this value fixed to isolate the effect of changing traffic conditions and then updates the value of time using contemporaneous wages. Presenting the two counterfactuals together makes transparent which part of the tax change comes from congestion and which part comes from valuing a given delay differently.

4.2 Demand side

To estimate the demand for freeway travel, I use two-stage least squares (2SLS), instrumenting the potentially endogenous remote work variable Z_{ct} with the shift–share instrument defined in Equation (13).

Demand equation (second stage). The demand for freeway travel is specified as

$$\ln F_{ct} = \alpha_c - 0.1 \ln(PMC_{ct}) + \beta_Z Z_{ct} + u_{ct}, \tag{37}$$

where α_c are county fixed effects, PMC_{ct} is the flow-weighted mean private marginal cost per vehicle mile defined in Equation (8), Z_{ct} = share_Home$_{ct}$ is the share of workers who work from home in county c and year t, and u_{ct} is an error term. The coefficient β_Z captures the causal effect of remote work on the demand shifter. I impose the private-cost elasticity at -0.1, as in the baseline calibration, rather than estimate it from the same equilibrium variation used to construct PMC_{ct}. Thus, the estimating equation can equivalently be written as $\ln F_{ct} + 0.1 \ln(PMC_{ct}) = \alpha_c + \beta_Z Z_{ct} + u_{ct}$.[3]

First stage (2SLS). Because Z_{ct} may be endogenous, I estimate Equation (37) using 2SLS. In the first stage, the work-from-home share is projected on the shift–share instrument:

$$Z_{ct} = \delta_c + \pi\, Z_{ct}^{IV} + \eta_{ct}, \tag{38}$$

where δ_c are county fixed effects, Z_{ct}^{IV} is the Bartik (shift–share) instrument constructed from pre-pandemic industry shares (2018 shares) and national shocks to telework feasibility, and η_{ct} is the first-stage error term.

[3] The model-based PMC captures the private travel-time cost per vehicle mile. A broader generalized cost could add fuel and other operating costs per mile. Gasoline price alone is not the relevant price of a vehicle mile because it omits travel time and does not convert fuel prices into per-mile operating costs.

Substituting the fitted values $\widehat{Z}_{ct}$ from Equation (38) into Equation (37) yields the 2SLS estimating equation:

$$\ln F_{ct} = \alpha_c - 0.1 \ln(PMC_{ct}) + \beta_Z \widehat{Z}_{ct} + \varepsilon_{ct}, \tag{39}$$

where ε_{ct} denotes the composite error term in the second stage.

The 2SLS results are summarized in Table 2.

Table 2: Effect of Working from Home on Freeway Flow (2SLS, PMC-Adjusted Demand)

Dependent Variable:	log(F)
Model:	(1)
Variables	
WFH share (2SLS)	-0.00125**
	(0.00034)
PMC elasticity	-0.1 (imposed)
Fixed-effects	
county	Yes
Fit statistics	
First-stage F-statistic	22.823
Wu-Hausman, p-value	0.40774
Observations	20
R^2	0.99977

Clustered (county) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The two-stage least squares estimates reported in Table 2 show a negative effect of working from home on freeway traffic volumes. The coefficient on the instrumented work-from-home share is -0.00125 with a county-clustered standard error of 0.00034, statistically significant at the 5 percent level. Because the work-from-home share is measured in percentage points, a 10 percentage point increase in remote work reduces freeway traffic by approximately 1.25 percent, holding private marginal cost and county fixed effects constant.

The first-stage F-statistic for the excluded instrument is 22.8, indicating that the shift–share instrument remains strongly correlated with the endogenous work-from-home share. The Wu–Hausman p-value is 0.408, so the null that the work-from-home share is exogenous is not rejected in this small sample. I nevertheless retain the IV specification because local remote-work adoption can respond to unobserved economic conditions that also affect freeway travel. The estimate should be interpreted with the limitations of the 20-observation county-year sample in mind.

5 Results

5.1 Congestion Costs with a Fixed Value of Time

Given the demand and congestion estimates, I compute optimal flow using Equation (15) and deadweight loss using Equation (17). Figure 5 combines the two aggregate cost measures. Holding the value of time fixed at its 2018 level, both total social cost and deadweight loss decline after the pandemic in every county. The ranking across counties reflects exposure as well as congestion: Los Angeles has the largest aggregate freeway cost because substantially more travel occurs on its freeway network, whereas the San Francisco calculation covers a smaller freeway system.

Figure 5. Congestion Costs with the Value of Time Fixed at 2018 Levels

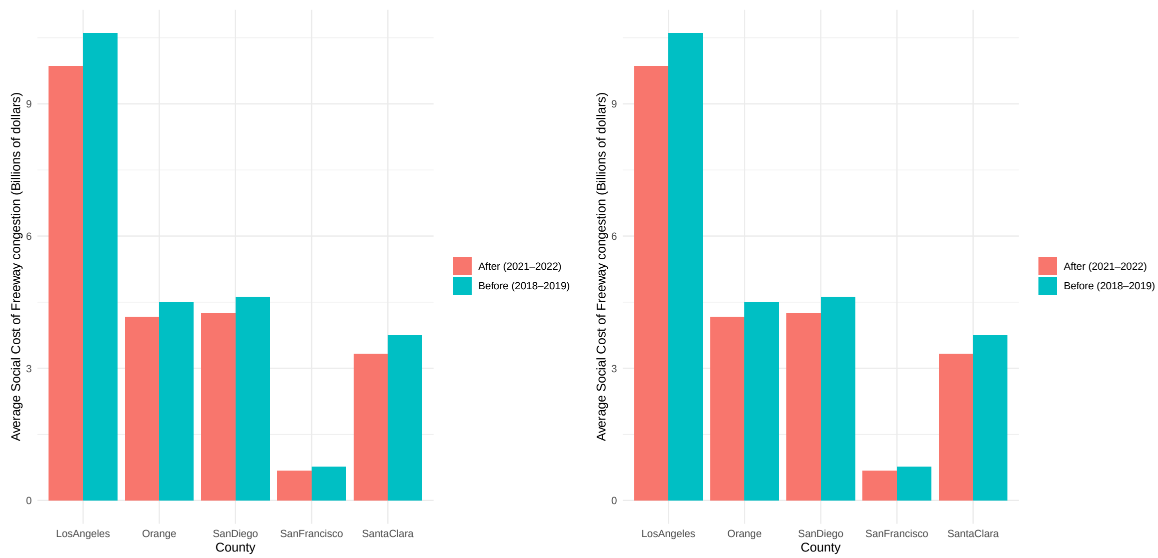

(a) Total Social Cost

(b) Deadweight Loss

Notes: Panel (a) reports the total social cost implied by Equation (9); panel (b) reports the local-linear deadweight-loss approximation in Equation (17). Both panels use 2018 county wages to value travel time.

The decline is economically largest where the remote-work shock removed travel from the steep portion of the congestion curve. This is why the timing documented in Appendix Figures B.2 and B.3 matters: reductions concentrated in commute peaks generate more external-cost relief than equal reductions spread uniformly across the day. These fixed-VOT calculations isolate that physical congestion channel. They do not yet answer the paper's central comparative-static question, because the monetary value assigned to each minute of delay also changed.

5.2 The Value-of-Time Wedge in Optimal Taxes

Figure 6 shows the evolution of the value-of-time proxy, defined as $VOT_{ct} = 0.5\,w_{ct}$. The increase is present in every county and is particularly pronounced in San Diego and San Francisco. Because $MECC(F) = VOT \times FT'(F)$, this change scales the dollar value of the marginal delay externality even when the physical congestion term $FT'(F)$ falls.

Figure 6. Value of Time by County

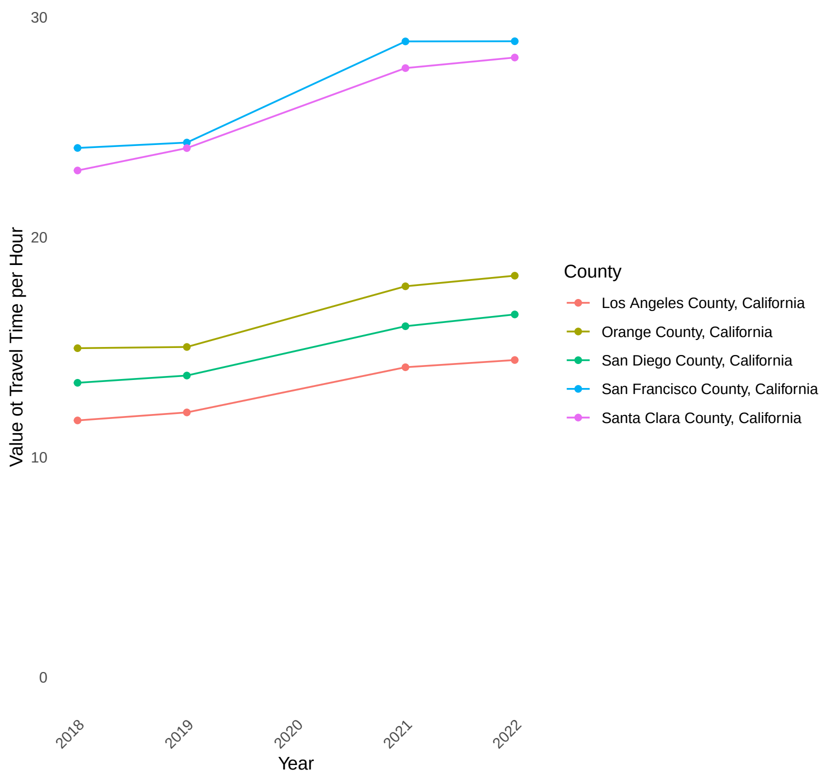

Notes: Value of time as a fixed proportion of the median wage in each county. The increase over time is particularly pronounced in San Diego and San Francisco.

Figure 7 is the paper's central comparison. Panel (a) holds VOT fixed and therefore answers what a pricing authority would conclude if it updated traffic conditions but continued to use the pre-pandemic value of time. Under that assumption, the implied optimal tax falls in every county. Panel (b) updates VOT with contemporaneous county wages. The conclusion reverses: the post-pandemic tax rises in every county. For example, the fixed-VOT calculation places Los Angeles at approximately $\$0.093$ per mile after the pandemic, while the time-varying calculation gives approximately $\$0.114$. The corresponding comparisons are $\$0.071$ versus $\$0.086$ in Orange, $\$0.083$ versus $\$0.101$ in San Diego, $\$0.183$ versus $\$0.220$ in San Francisco, and $\$0.146$ versus $\$0.177$ in Santa Clara.

Figure 7. Optimal Congestion Taxes Under Alternative Value-of-Time Assumptions

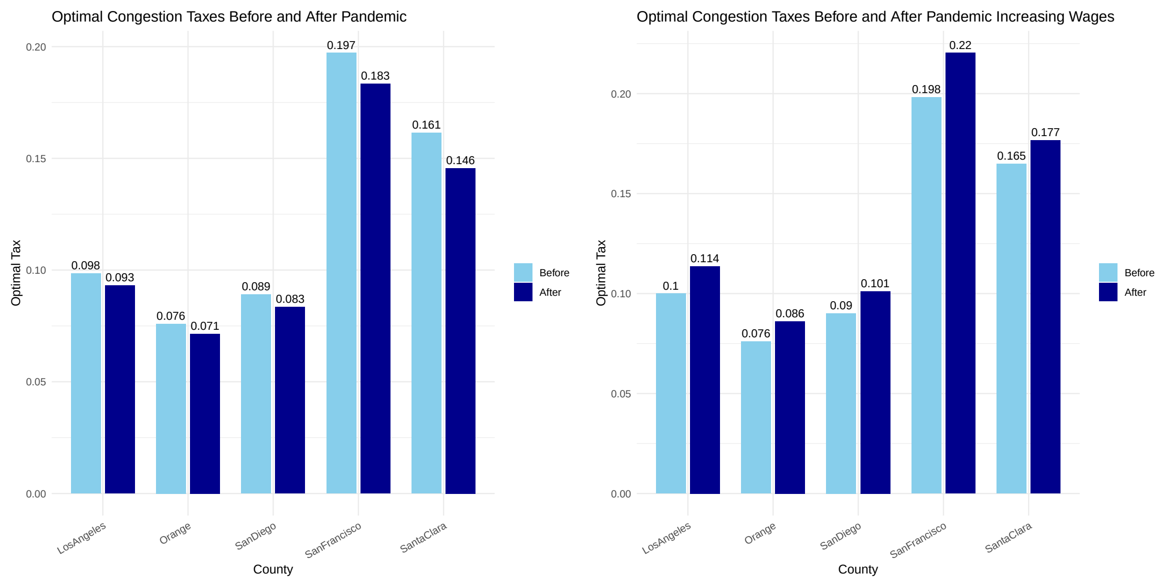

(a) Value of Time Fixed at 2018 Level (b) Value of Time Updated Annually

Notes: Implied optimal congestion taxes in dollars per vehicle mile. Light bars report the 2018–2019 mean and dark bars report the 2021–2022 mean. Panel (a) uses 2018 county wages in both periods; panel (b) uses contemporaneous county wages. Taxes equal marginal external congestion costs evaluated at optimal flow.

Relative to the time-varying benchmark, the fixed-VOT calculation understates the post-pandemic tax by approximately 17–18 percent in each county. The stability of this percentage across locations follows from the multiplicative role of VOT: the error is not primarily a statement about which county is most congested, but about how an outdated conversion from minutes to dollars scales the marginal external cost. In levels, however, the wedge is largest in San Francisco and Santa Clara because their baseline values of time and optimal taxes are higher.

Figure 8 applies the exact decomposition in Equation (35). The congestion component is negative in every county: at the pre-pandemic VOT, lower marginal external delay would reduce the tax. The VOT component is positive and larger in absolute value, producing the net increase shown by the black marker. The exercise therefore distinguishes two statements that are sometimes conflated. Post-pandemic California freeways are less congested in physical terms, but the remaining delay is more costly when converted into dollars using the contemporaneous wage-based VOT.

Figure 8. Decomposition of Changes in the Optimal Tax

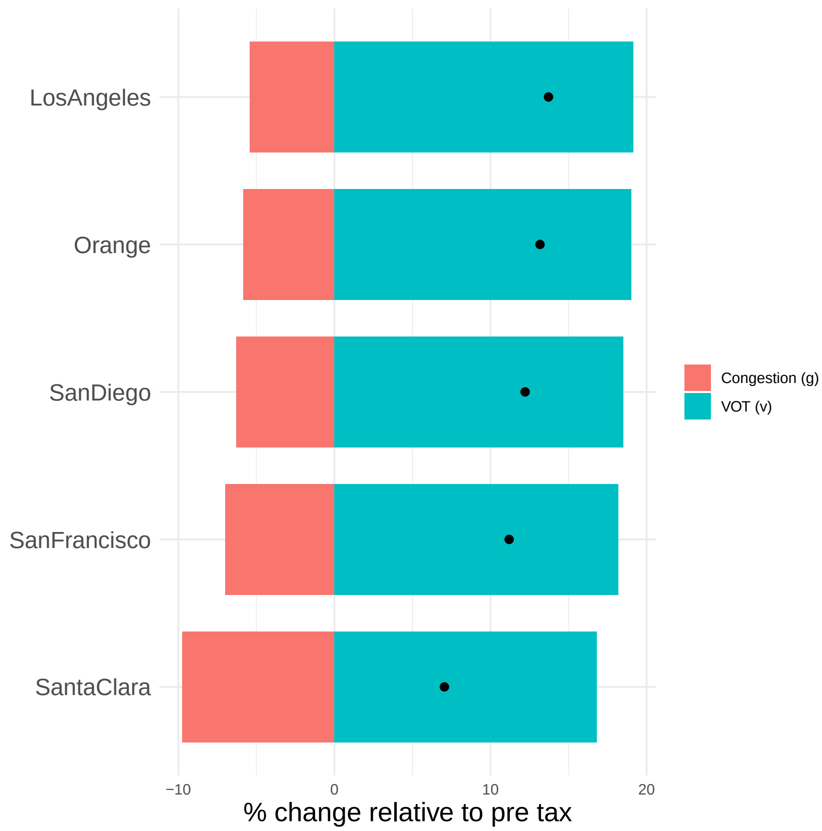

Notes: The figure decomposes changes in the implied optimal tax into a congestion component and a value-of-time component as in equation (35).

The calculation uses a county median wage and therefore remains deliberately conservative about heterogeneity. Express-lane evidence shows wide dispersion in willingness to pay for time savings and reliability (Bento et al., 2024; Lam & Small, 2001; Small et al., 2005). If high-VOT travelers are less elastic and consequently become a larger share of drivers as a toll rises, the relevant VOT at the margin may exceed the county median. The optimal tax and the composition of road users would then be jointly determined. The present results do not estimate that feedback; instead, they demonstrate that even the simpler temporal adjustment to an average VOT is quantitatively large enough to reverse the fixed-VOT conclusion.

6 Conclusion

This paper uses the persistent post-COVID expansion of remote work to study how a large demand shock changes the marginal external cost of freeway use. Combining hourly PeMS traffic measures with ACS remote-work shares and wages, I estimate the demand response and calculate implied congestion prices in five large California counties. The shift–share estimate indicates that a 10 percentage point increase in working from home reduces freeway VMT by approximately 1.25 percent, with the largest congestion improvements occurring during commute peaks.

The principal result is a comparison between two pricing exercises. When the value of time is fixed at its 2018 level, the reduction in marginal delay lowers the implied optimal tax in every county. When the value of time is updated with contemporaneous county wages, the post-pandemic tax instead rises in every county. Using the fixed value understates the time-varying benchmark by approximately 17–18 percent. The exact decomposition makes the economics transparent: less congestion reduces the physical delay externality, but the increase in the monetary value of that delay is larger.

This result has a direct policy implication. Congestion charges should not be updated from traffic conditions alone. A pricing authority that measures current delay accurately but converts it into dollars using an outdated VOT can still systematically understate marginal external cost. The point is especially relevant following structural changes in commuting, when both the quantity of peak travel and the characteristics of road users may change at the same time.

The paper's wage-based VOT is an average benchmark, not a complete model of heterogeneous travelers. Express-lane evidence indicates substantial dispersion in willingness to pay for time and reliability, and high-VOT drivers may be less responsive to tolls. If pricing changes the composition of road users, the VOT relevant at the margin becomes endogenous to the toll. Incorporating that selection mechanism would connect the temporal wedge documented here to the distributional consequences of congestion pricing.

Several limitations define the interpretation of the estimates. The analysis covers freeways in five counties rather than the full multimodal transportation network; it abstracts from route and departure-time substitution, transit crowding, environmental externalities, and the broader costs and benefits of remote work. Accordingly, the welfare calculations should be read as partial-equilibrium measures of the freeway congestion externality. Within that scope, the central lesson is clear: efficient road prices must adapt both to how much delay traffic creates and to how users value the time lost in that delay.

References

Akbar, P., Couture, V., Duranton, G., & Storeygard, A. (2023). Mobility and congestion in urban india. *American Economic Review*, *113*(4), 1083–1111.

Bento, A., Roth, K., & Waxman, A. R. (2024). *The value of urgency: Evidence from real-time congestion pricing* (Working Paper No. 26956). Cambridge, MA: National Bureau of Economic Research.

Couture, V., Duranton, G., & Turner, M. A. (2018). Speed. *Review of Economics and Statistics*, *100*(4), 725–739.

Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home? *Journal of public economics*, *189*, 104235.

Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from us cities. *American Economic Review*, *101*(6), 2616–2652.

Lam, T. C., & Small, K. A. (2001). The value of time and reliability: Measurement from a value pricing experiment. *Transportation Research Part E: Logistics and Transportation Review*, *37*(2–3), 231–251.

Newbery, D. M. (1989). Cost recovery from optimally designed roads. *Economica*, 165–185.

Pigou, A. C. (1920). *The economics of welfare*. London: Macmillan.

Quinet, E., & Vickerman, R. W. (2004). *Principles of transport economics*. Northampton, MA.

Small, K. A. (2012). Valuation of travel time. *Economics of transportation*, *1*(1-2), 2–14.

Small, K. A., Verhoef, E. T., & Lindsey, R. (2024). *The economics of urban transportation*. Routledge.

Small, K. A., Winston, C., & Yan, J. (2005). Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica*, *73*(4), 1367–1382.

Small, K. A., & Yan, J. (2001). The value of "value pricing" of roads: Second-best pricing and product differentiation. *Journal of Urban Economics*, *49*(2), 310–336.

van den Berg, V. A. C., & Verhoef, E. T. (2011). Congestion tolling in the bottleneck model with heterogeneous values of time. *Transportation Research Part B: Methodological*, *45*(1), 60–78.

Verhoef, E. T. (2002). Second-best congestion pricing in general networks. heuristic algorithms for finding second-best optimal toll levels and toll points. *Transportation Research Part B: Methodological*, *36*(8), 707–729.

Verhoef, E. T., & Small, K. A. (2004). Product differentiation on roads: Constrained congestion pricing with heterogeneous users. *Journal of Transport Economics and Policy*, *38*(1), 127–156.

Vickrey, W. S. (1969). Congestion theory and transport investment. *American Economic Review*, *59*(2), 251–260.

Yang, J., Purevjav, A.-O., & Li, S. (2020). The marginal cost of traffic congestion and road pricing: Evidence from a natural experiment in beijing. *American Economic Journal: Economic Policy*, *12*(1), 418–453.

Appendix

A Derivation of Private, External, and Social Costs

This appendix explains the economic logic behind Equations (1) and (2) and then derives them formally. The central distinction is between the cost perceived by an individual driver and the change in total cost caused by an increase in aggregate traffic. These objects coincide when travel time is independent of traffic. They differ under congestion because an additional vehicle mile changes the travel time experienced by other users.

A.1 Primitives and units

Consider one of the hourly freeway markets used in the empirical analysis. Let F denote total vehicle miles traveled during the hour and let $T(F)$ denote travel time per vehicle mile, measured in hours per mile. The congestion technology satisfies $T'(F) \geq 0$: a larger volume of traffic weakly increases the time needed to travel a mile. Let VOT denote the value of travel time in dollars per hour. Multiplying these objects gives the time cost of one vehicle mile, $VOT\,T(F)$, measured in dollars per vehicle mile. Multiplying once more by the F vehicle miles traveled gives aggregate travel-time cost during the hour:

$$TSC(F) = F\,VOT\,T(F). \tag{40}$$

Equation (40) contains two margins. Increasing F adds another vehicle mile that must itself be traveled, and it may also increase the travel time required for every vehicle mile on the road. Equations (1) and (2) come from separating these two effects.

A.2 The driver's problem and private marginal cost

Freeway traffic is modeled as nonatomic: each driver is small relative to the market and treats prevailing traffic conditions as given. A driver deciding whether to travel one additional mile compares the benefit of that mile with the cost personally borne. In the baseline model, this private cost is the monetary value of the time required for the trip:

$$PMC(F) = VOT\,T(F). \tag{41}$$

The word "marginal" in PMC refers to the driver's individual decision. It does not mean differentiating aggregate cost with respect to aggregate flow. From the driver's perspective, one more mile costs $VOT\,T(F)$ because the driver takes $T(F)$ as fixed. For this reason, private marginal cost is also the average user cost per vehicle mile in the nonatomic model.

Without a toll, the decentralized equilibrium occurs where inverse demand, $P(F)$, equals this private cost:

$$P(F) = PMC(F). \tag{42}$$

The driver internalizes the time spent on the trip but not the small contribution the trip makes to congestion faced by everyone else. That omitted effect is the source of the externality.

A.3 The planner's problem and marginal social cost

A planner asks a different question: by how much does total travel-time cost change when aggregate freeway use increases by one vehicle mile? The answer can be seen without calculus. Starting from flow F, consider a small increase ΔF. To a first-order approximation, the change in total cost has two components:

$$TSC(F+\Delta F) - TSC(F) \simeq \underbrace{VOT\,T(F)\Delta F}_{\text{time cost borne by the new traffic}} + \underbrace{F\,VOT\,T'(F)\Delta F}_{\text{additional delay imposed on existing traffic}} . \tag{43}$$

The first component is familiar: the additional traffic must spend time traveling. The second component is easy to overlook. The increase ΔF raises travel time per mile by approximately $T'(F)\Delta F$. This delay is experienced across the F vehicle miles already on the road, so it creates approximately $FT'(F)\Delta F$ additional vehicle-hours. Multiplication by VOT converts those hours into dollars.

Dividing Equation (43) by ΔF and taking the limit as $\Delta F \to 0$ gives the marginal social cost. Equivalently, applying the product rule directly to Equation (40) yields

$$\begin{aligned} MSC(F) \equiv \frac{\partial TSC(F)}{\partial F} &= VOT\,T(F) + F\,VOT\,\frac{\partial T(F)}{\partial F} \\ &= PMC(F) + MECC(F). \end{aligned} \tag{44}$$

This is Equation (1) in the main text. It is an accounting identity rather than an additional behavioral assumption. The social cost of one more vehicle mile equals the private time cost borne by the driver plus the cost shifted onto other freeway users.

The second term is the marginal external congestion cost:

$$MECC(F) = F\,VOT\,\frac{\partial T(F)}{\partial F}. \tag{45}$$

This is Equation (2). Its three factors have a direct interpretation. The slope $\partial T(F)/\partial F$ measures the increase in hours per mile caused by one more unit of traffic; F measures how many vehicle miles are exposed to that increase; and VOT converts the resulting delay into dollars. The factor F appears because congestion is not merely a cost to the entering driver. The marginal increase in travel time is applied to the entire traffic volume.

The distinction between PMC and MSC is therefore a distinction between viewpoints. The individual treats congestion as given and sees only PMC. The planner recognizes that aggregate flow determines congestion and sees MSC. When $T'(F) = 0$, there is no congestion margin, $MECC(F) = 0$, and the two viewpoints coincide. When $T'(F) > 0$, marginal social cost lies above private marginal cost, so the unpriced equilibrium has more freeway travel than the social optimum.

A.4 Connection to the BPR congestion technology

The paper implements this logic using the estimated BPR-type travel-time function. Let $t_0 = 1/S_{\text{free-flow}}$ denote free-flow travel time per mile. Then

$$T(F) = t_0 \left(1 + \alpha F^{\beta}\right), \qquad T'(F) = t_0 \alpha \beta F^{\beta - 1}. \tag{46}$$

Substitution into Equations (41) and (45) gives

$$PMC(F) = VOT\, t_0 \left(1 + \alpha F^{\beta}\right), \tag{47}$$

$$MECC(F) = VOT\, t_0 \alpha \beta F^{\beta}, \tag{48}$$

$$MSC(F) = VOT\, t_0 \left[1 + \alpha(1 + \beta) F^{\beta}\right]. \tag{49}$$

Define $PMC_0 = VOT\, t_0$, the private time cost under free-flow conditions. The congestion component of private cost is $PMC(F) - PMC_0 = VOT\, t_0 \alpha F^{\beta}$. Therefore,

$$MECC(F) = \beta \left[PMC(F) - PMC_0\right]. \tag{50}$$

Equation (50) provides a useful interpretation of the BPR parameter β. The free-flow part of travel time is a real private cost, but it is not a congestion externality. Only the portion of private cost attributable to congestion generates the external wedge, and the size of that wedge depends on how steeply delay rises with flow. A larger β makes travel time more sensitive to traffic and therefore increases the external cost associated with a given congestion component.

In the conventional BPR formula, congestion depends on the volume-capacity ratio. Because capacity is not separately observed in the county-level PeMS data, the paper's α absorbs the capacity term. Separate identification of capacity is not needed for Equations (47)–(49): the welfare calculation uses the level and slope of the fitted relationship between observed flow and travel time. The interpretation is consequently local to the range of traffic conditions observed in the data, rather than an engineering estimate of physical capacity.

A.5 Why the external cost is the optimal toll

The efficient flow satisfies

$$P(F^{opt}) = MSC(F^{opt}) = PMC(F^{opt}) + MECC(F^{opt}). \tag{51}$$

Drivers, however, respond to the price they personally face. If the pricing authority levies a per-mile toll τ, the perceived price becomes $PMC(F) + \tau$. Choosing

$$\tau^{opt} = MECC(F^{opt}) \tag{52}$$

makes the decentralized condition $P(F) = PMC(F) + \tau$ coincide with the planner's condition in Equation (51). The toll does not represent an additional resource cost: conditional on how the revenue is used, it is a transfer from drivers to the pricing authority. Its role is to make drivers take into account the real delay cost their trips impose on others. Section 2.5 then allows flow and PMC to adjust endogenously to this toll.

The baseline PMC includes private travel-time cost because this is the component measured using PeMS travel times and county-specific values of time. A broader generalized-cost measure could add a flow-independent operating cost per mile, r. Aggregate resource cost would then be $F[r + VOT\,T(F)]$, private marginal cost would be $r + VOT\,T(F)$, and marginal social cost would rise by the same r. The congestion externality in Equation (45) would be unchanged as long as r does not vary with contemporaneous freeway flow. This is also why gasoline price alone is not the appropriate demand price: it omits the time cost and, without fuel economy, is not expressed as a cost per vehicle mile.

B Robustness Tables and Figures

Table B.1: Major Freeways by County in California

County	**Freeways**
Los Angeles	I-5, I-10, I-105, I-110, I-210, I-405, I-605, I-710, SR-1, SR-2, SR-14, SR-60, SR-91
San Francisco	I-80, I-280, US 101, SR-1
Orange	I-5, I-405, SR-22, SR-55, SR-57, SR-73, SR-91, SR-133, SR-241
Santa Clara	I-280, I-680, US 101, SR-17, SR-85, SR-87, SR-237, SR-152
San Diego	I-5, I-8, I-15, I-805, SR-52, SR-54, SR-56, SR-75, SR-94, SR-125, SR-163

Note: The freeway data listed here is derived from Caltrans' Performance Measurement System (PeMS). While PeMS includes data for the entire California freeway network, data quality may vary depending on sensor coverage and performance. However, given the urban nature of the counties analyzed, sensor density and data reliability are generally high in these areas.

Figure B.1. Share of People Working from Home by County

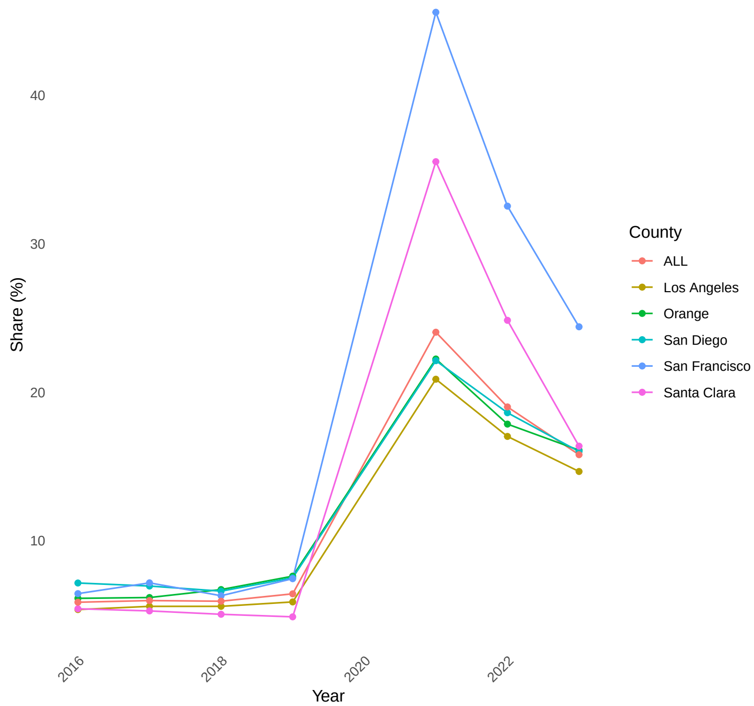

Notes: American Community Survey data. Remote work rose sharply after 2020 and remained above its pre-pandemic level through 2023, particularly in San Francisco and Santa Clara.

Figure B.2. Changes in Congestion by Hour

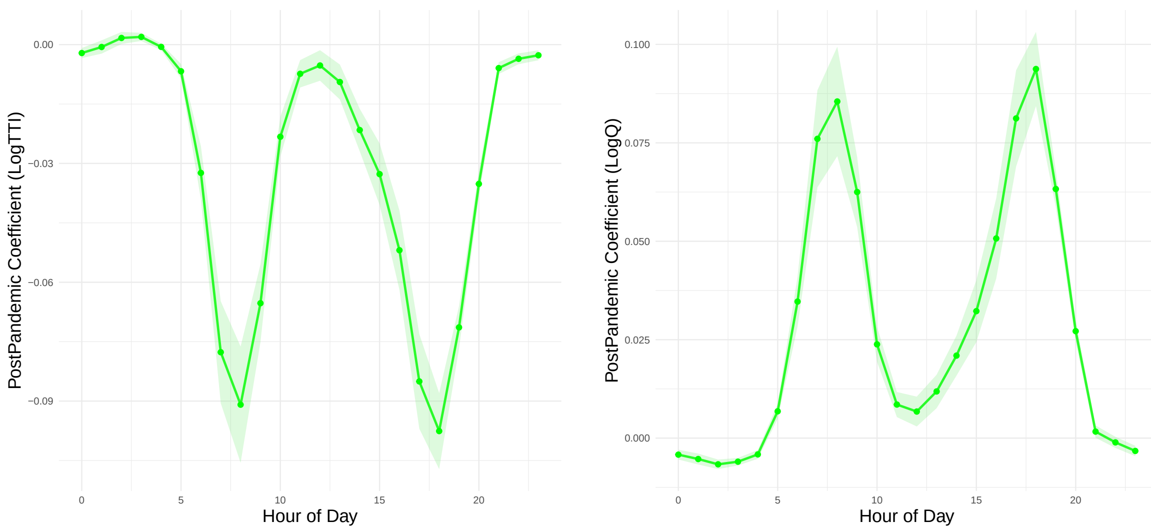

(a) Change in TTI by Hour

(b) Change in Average Speed by Hour

Notes: Hour-specific regression estimates comparing 2018–2019 with 2021–2022. The largest congestion improvements occur during traditional commuting peaks.

Figure B.3. Changes in Vehicle Miles Traveled by Hour

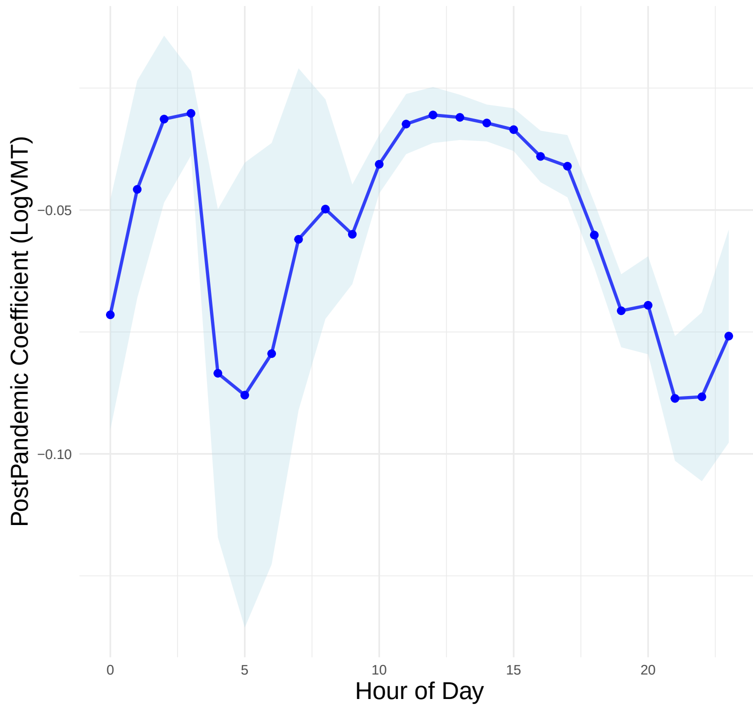

Notes: Hour-specific regression estimates comparing 2018–2019 with 2021–2022. VMT falls most during commuting hours and changes relatively little late at night.

Figure B.4. Changes in Traffic and Congestion versus Remote Work

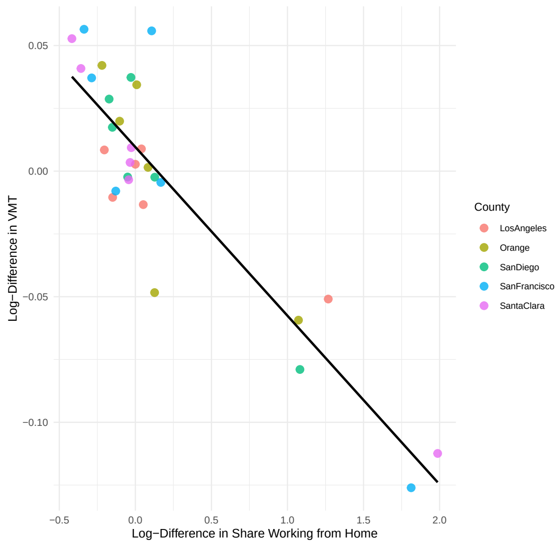

(a) Change in VMT and Remote Work

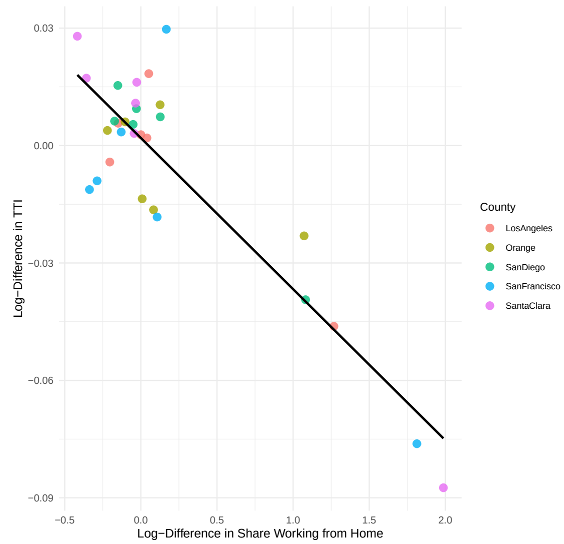

(b) Change in TTI and Remote Work

Notes: Each point is a county-year observation. PeMS traffic outcomes are plotted against ACS work-from-home shares. The correlations are descriptive; causal estimates use the shift–share instrument.

Figure B.5. VMT for Passenger Vehicles and Trucks from January 2016 to February 2025

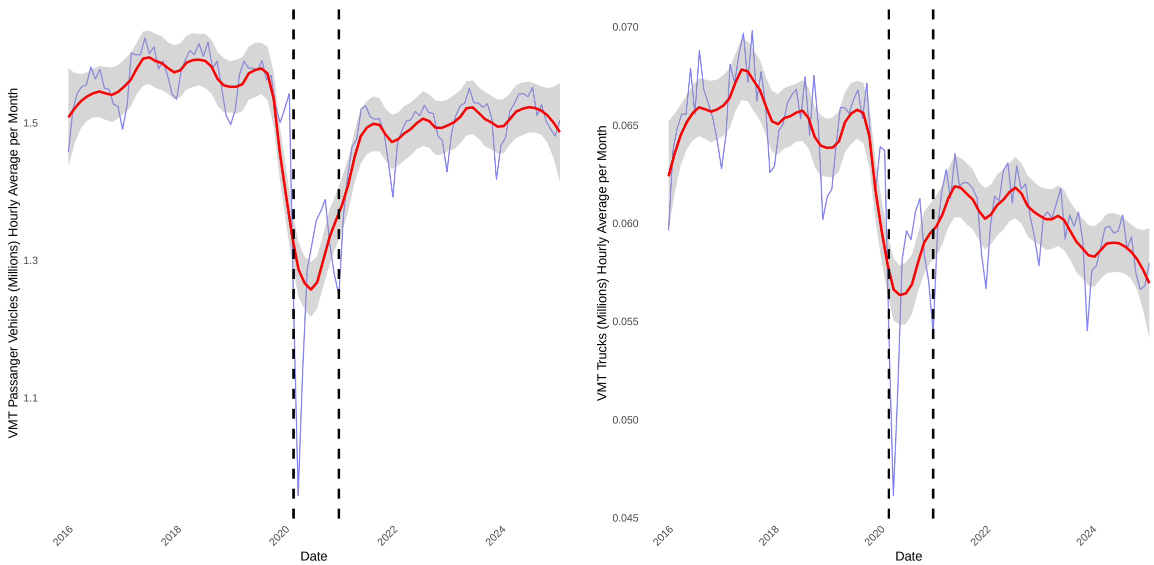

(a) Vehicle Miles Traveled for Passenger Vehicles (b) Vehicle Miles Traveled for Trucks

Notes: This figure shows the evolution of Vehicle Miles Traveled for passenger vehicles and trucks from January 2016 to February 2025, measured as the monthly average per hour. The blue line represents the raw data. The red line is a LOESS smoother with a 95% confidence interval. Vertical dashed black lines indicate the start and end of the lockdown period.

Figure B.6. Yearly Change in Traffic and Congestion vs. Change in Population

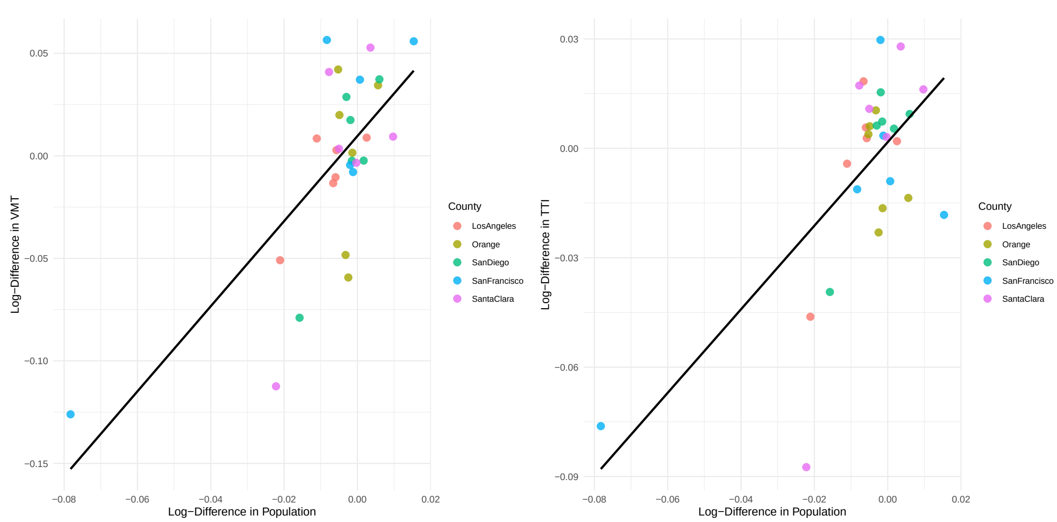

(a) Yearly Change in VMT (b) Yearly Change in TTI

Notes: Yearly approximated percentage change in VMT and TTI plotted against the yearly approximated percentage change in population for each county. VMT and TTI data are from PeMS; population estimates are from the ACS. Each dot represents a county-year observation.

Table B.2: Differences in Traffic Before and After the Pandemic

	TTI		
	All Data	(3pm–6pm)	Weekends
	(1)	(2)	(3)
Pandemic After	-0.029^{***}	-0.065^{***}	-0.007^{***}
	(0.001)	(0.002)	(0.001)
Constant	0.018^{***}	0.197^{***}	-0.036^{***}
	(0.0004)	(0.001)	(0.001)
Observations	181,655	30,280	51,695
R^2	0.011	0.050	0.001
Adjusted R^2	0.011	0.050	0.001
Note:	*p<0.1; **p<0.05; ***p<0.01		

Notes: The table reports differences in mean $\log(TTI)$. Column 1 uses all hourly observations, column 2 uses the 3–6 p.m. peak, and column 3 uses weekends. The comparison excludes 2020 and contrasts 2018–2019 with 2021–2022.