

Can Rebates Foster Equity in Congestion Pricing Programs?

James Sallee, Ph.D., Associate Professor, Department of Agricultural and Resource Economics, University of California, Berkeley
Matthew Tarduno, Ph.D. Candidate, Department of Agricultural and Resource Economics, University of California, Berkeley

August 2022

Technical Report Documentation Page

1. Report No. UC-ITS-2020-26		2. Government Accession No. N/A		3. Recipient's Catalog No. N/A	
4. Title and Subtitle Can Rebates Foster Equity in Congestion Pricing Programs?				5. Report Date August 2022	
				6. Performing Organization Code ITS Berkeley	
7. Author(s) James M. Sallee, ORCID: 0000-0001-6590-9897 Matthew A. Tarduno, ORCID: 0000-0001-9682-479X				8. Performing Organization Report No. N/A	
9. Performing Organization Name and Address Institute of Transportation Studies, Berkeley 109 McLaughlin Hall, MC1720 Berkeley, CA 94720-1720				10. Work Unit No. N/A	
				11. Contract or Grant No. UC-ITS-2020-26	
12. Sponsoring Agency Name and Address The University of California Institute of Transportation Studies www.ucits.org				13. Type of Report and Period Covered Final Report (July 2019 – Jan 2022)	
				14. Sponsoring Agency Code UC ITS	
15. Supplementary Notes DOI:10.7922/G2V40SJ4					
16. Abstract Congestion pricing improves economic efficiency, but it may lead to inequitable outcomes. A key policy priority in California is identifying ways to avoid the hardship of congestion pricing on low income or other vulnerable populations. This study uses data from a congestion pricing experiment in the Seattle metro area to examine the feasibility of using revenue from congestion pricing to compensate those harmed by the policy. Results indicate that the initial burden of congestion pricing is highly inequitable, with the lowest income drivers paying an average of 7 percent of their weekly income in congestion charges. There are also considerable differences in burdens within income groups. We show that policymakers face a tradeoff in ameliorating these two types of unequal burdens. Returning an equal fraction of the toll revenue to all drivers can make a policy progressive on average, but doing so leaves many drivers either over-compensated or under-compensated. We then show that while compensation packages based on basic demographic information could improve targeting, many low-income drivers would be left with large proportional burdens because of the fundamental difficulty in predicting individual-level tax burdens. Survey data on travel behavior from Seattle and California metro areas show that the difficulty of designing equitable transfers would be similar in the California metro areas most likely to consider adopting some form of congestion pricing.					
17. Key Words Congestion pricing, low income groups, social equity, vehicle miles of travel, travel behavior, incentives, traffic data				18. Distribution Statement No restrictions.	
19. Security Classification (of this report) Unclassified		20. Security Classification (of this page) Unclassified		21. No. of Pages 47	21. Price N/A

About the UC Institute of Transportation Studies

The University of California Institute of Transportation Studies (UC ITS) is a network of faculty, research and administrative staff, and students dedicated to advancing the state of the art in transportation engineering, planning, and policy for the people of California. Established by the Legislature in 1947, ITS has branches at UC Berkeley, UC Davis, UC Irvine, and UCLA.

Acknowledgments

This study was made possible with funding received by the University of California Institute of Transportation Studies from the State of California through the Public Transportation Account and the Road Repair and Accountability Act of 2017 (Senate Bill 1). The authors would like to thank the State of California for its support of university-based research, and especially for the funding received for this project.

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the State of California in the interest of information exchange. The State of California assumes no liability for the contents or use thereof. Nor does the content necessarily reflect the official views or policies of the State of California. This report does not constitute a standard, specification, or regulation.

Can Rebates Foster Equity in Congestion Pricing Programs?

James Sallee, Ph.D., Associate Professor, Department of Agricultural and Resource Economics, University of California, Berkeley

Matthew Tarduno, Ph.D. Candidate, Department of Agricultural and Resource Economics, University of California, Berkeley

August 2022

Table

of

Contents

Table of Contents

Executive Summary	1
Introduction	4
Congested Pricing Research	9
Congestion Charges Tend to be Regressive	9
Horizontal Equity and Congestion Charges	10
Congestion Pricing in Economic Theory	12
Pigouvian Taxes and Economic Efficiency	12
Transfers Can Mitigate Regressivity Without Sacrificing Efficiency, but Some May Still Be Left Worse Off.	12
Summary of Seattle Experiment	14
Results	17
Direct Burdens of Congestion Pricing are Regressive	17
Performance of Conditioned Transfer Programs	20
Alternatives for Conditioning Transfers	21
Measuring the Regressivity of Congestion Pricing	26
Applicability to California	27
Conclusion	31
References	32
Appendix	34
Congestion Pricing in Economic Theory	34
Pigouvian Taxes and Economic Efficiency	34
One’s Burden from a Pigouvian Tax Depends on the Quantity of the Good or Service Consumed	35
The Challenge of Targeted Compensation	36

List of Tables

- Table 1. Toll Schedule 14
- Table 2. Summary Statistics 16
- Table 3. Costs Rise with Income but are Diffuse 18
- Table 4. Predicting Program Costs 20
- Table 5. Using the Gini to Measure Regressivity..... 27
- Table 6. California – Seattle Comparison 28

List of Figures

- Figure 1. Costs Rise with Income but are Diffuse 17
- Figure 2. Congestion Charges are Regressive and Diffuse 19
- Figure 3. Net Burdens Under Uniform Lump-Sum Transfers by Income Category 22
- Figure 4. Distribution of Net Burdens Under Uniform Lump-Sum Transfers 23
- Figure 5. Distribution of Net Burdens Under Targeted Transfers, by Income Category 24
- Figure 6. Distribution of Net Burdens Under VMT-Targeted Transfers..... 25
- Figure 7. The VMT-Income Relationship in California Cities and Seattle..... 29
- Figure 8. Predicted Proportional Burden, CA 30
- Figure 9. Tolled Roads in Seattle’s Transportation Choices Study. 38

Executive

Summary

Executive Summary

Road pricing offers an attractive tool for addressing urban congestion. A key policy priority in California is identifying ways to design road pricing that avoid creating undue hardship on low income or other vulnerable populations. This study uses a congestion pricing pilot project in Seattle to examine the initial burden of road pricing on households across income levels and the feasibility of using revenue from congestion pricing to compensate those households burdened by the policy.

Results show that while high-income households in Seattle paid more in congestion fees than did low-income households, in proportional terms the initial burden of the congestion pricing project was highly inequitable. The lowest income drivers paid an average of 6.7 percent of their weekly income in congestion charges, whereas the highest-income drivers paid less than one percent of their weekly income. Providing rebates or exemptions to low-income households for the tolls they pay would be one way to reduce the burden on these populations, but doing so would erode potential efficiency gains by weakening the program's built in disincentives for using congested roadways. The aim of this study is to ask whether it is possible to achieve equitable outcomes from congestion pricing schemes without compromising their efficiency through the use of lump-sum transfer payments that return a portion of program revenue to households.

We first show that transfers funded by congestion pricing revenues can be used to make the net impact of the program progressive. In the Seattle experiment, using half of the revenues to fund a uniform transfer to each household that has a vehicle would undo the regressive impact and make the program “equity neutral.” That is, under such a transfer, the net burden as a fraction of income would be roughly equivalent across different income groups. If *all* program revenue was used to fund a uniform lump-sum transfer, the program would be progressive: on average, high-income households would experience net costs equal to 0.5 percent of their income, whereas low-income households would see net *benefits* equal to 3.8 percent of their weekly income. Of course, congestion pricing programs could be made even more progressive if the revenue is returned as a function of income, with lower-income households receiving a larger absolute amount.

While such a transfer could equalize the average impact on households across income groups, data from the Seattle experiment suggest that there would still be substantial variation in the impact of congestion pricing between households *within* an income category. This means that, while the program would be progressive on average, it would also cause many households, including some low income households, to be worse off than they were before the congestion pricing system was introduced. This has both an equity implication and a political economy implication because households made worse off by the congestion pricing scheme may organize politically to oppose the system.

In theory, *targeted* lump-sum transfers, which return more program revenue to households that pay more into the system, could be used to address these unequal burdens and to minimize the number of households made worse off by congestion pricing. Care must be taken to design target transfers to avoid eroding the economic

efficiency of the congestion pricing. As we explain below, it is possible to design targeted transfers that preserve efficiency, but the ability of targeted transfers to achieve various equity goals depends upon the degree to which available household characteristics, such as income, household size or neighborhood, can predict variation in the burdens households face from a congestion pricing system.

Empirically, we find that readily available household characteristics are not strong predictors of the individual costs of congestion pricing programs, explaining just 27 percent of the variation in program burdens. Put differently, most of the variation across households in how much they would pay into a congestion pricing system is unpredictable. A consequence of this is that policy that targets transfers based on this information is unlikely to effectively compensate the largest losers. As such, we conclude that, even with targeted transfers, a congestion pricing scheme will create some households who are made worse off by a congestion pricing system, and who may therefore represent a constituency opposed to such policies.

Transfers based directly on past driving behavior—such as vehicle miles traveled (VMT)—can more closely match burdens thereby reducing the number of people made worse off by a congestion pricing system, but this approach comes with two drawbacks. First, tying transfers to driving behavior may require the agency administering the congestion pricing system to collect additional data on drivers. Second, if an individual's transfer is conditioned on their driving behavior, this undermines the primary goal of congestion pricing because drivers will understand that the more they drive, the bigger their transfer will be.

Our analysis further highlights a tension created by targeting. Targeting transfers to reduce inequality within groups comes at the cost of addressing inequality across groups. Because high-income drivers tend to pay more in congestion tolls, targeted transfers tend to be larger for high-income households, regardless of the variables used to design the transfer. As a result, targeted transfers produce a less progressive distribution of outcomes than do uniform transfers. In Seattle, for example, low-income households would see a net benefit from a uniform transfer program equal to 3.8 percent of their income, on average, whereas basing household transfers on VMT would lead to an average benefit of just 0.7. This presents a tradeoff for policymakers who wish to both protect low-income households from unreasonable burdens and to limit the number of people who are net losers from a policy, and therefore may oppose it, across all income groups. Taken together, these results show that a congestion pricing program can be made progressive on average, but also suggest that it will be difficult to effectively shield households who have the heaviest driving needs from experiencing a substantial burden.

The final portion of this report uses data from the National Highway Transportation Survey to compare driving and demographic patterns in Seattle to California metro areas, to assess the applicability of the results from Seattle to California. The relationship between vehicle use and income appears to be quite similar in California and Seattle. Likewise, the data suggest that there is a similarly high level of variation in driving habits based on observable characteristics across the regions. This suggests that the qualitative results from the analysis of the Seattle pilot study—namely the difficulty in designing targeted transfers, and the tradeoff between progressivity and the number of overall losers—are likely to hold for metro areas in California.

Contents

Introduction

Congestion pricing programs aim to alleviate traffic congestion by charging drivers tolls for using certain busy roads during peak travel times. This study uses economic theory and data from an experiment in the Seattle metropolitan area to assess the fairness of congestion pricing programs.

Traffic congestion is symptomatic of inefficient choices. Because drivers are not forced to bear the costs that they impose on others by slowing down travel, people generally drive too much and on socially inefficient routes. This is an example of what economists call an externality. An externality is a setting where some of the cost of an action is not borne by the actors making the decision.

Economists have a solution to inefficiencies caused by externalities, which is to tax externality-creating behaviors by an amount equal to this extra, external cost (Pigou 1920). Applied to traffic congestion, this solution takes the form of congestion pricing, which forces drivers who choose to drive at congested times or in congested places to pay a toll that represents the costs that they impose on others.

The idea of congestion pricing goes back at least as far as Vickrey (1959). In economic terms, it is theoretically sound. In terms of technology, it is imminently feasible. In terms of social value, it is substantial, as the problem it alleviates is ubiquitous and growing. Yet, congestion pricing remains relatively rare. Why is it not more common?

One barrier to the adoption of congestion pricing programs are concerns about equity and fairness. Equity concerns come in two distinct flavors: vertical equity and horizontal equity. Vertical equity is a fairness principle that suggests that high-income and low-income individuals should not pay equal taxes, and that tax burdens should instead increase with income. Horizontal equity is a fairness principle that suggests that people with similar means and similar circumstances should pay similar amounts in taxes.

Like most programs that directly tax an externality-creating behavior, the direct costs of congestion pricing are generally found to be quite regressive—that is, they create a burden on lower-income drivers that is a higher proportion of their income than is the case for higher-income drivers. This is an example of vertical inequality. Levinson (2009) provides a useful survey of the vertical equity impacts of congestion pricing programs.

Congestion pricing can also create horizontal inequities: among individuals with similar income, congestion charges impose very different burdens on different individuals according to their transportation needs and the availability of alternatives. For example, some drivers may need to pay congestion prices to get to their daily workplace, where others may be able to take an alternative route, use public transit, or work from home. Although public discourse on equity typically focuses on vertical equity, we also analyze these horizontal inequities for political economy reasons. If a sufficiently large subset of the population feels unfairly burdened by a policy, they can often form a blocking coalition even if the burden does not fall disproportionately on low-income groups.

Economists have theoretically elegant solutions to both concerns. Regarding vertical equity, congestion pricing programs raise revenue, which can be used to provide direct *lump-sum transfers*¹ to households. These transfers can lead to progressive overall impacts of the policy without sacrificing the efficiency of congestion pricing. Regarding the second type of equity, horizontal equity can also in theory be addressed by lump-sum transfers if the transfers are *targeted* such that those facing greater costs from the program receive larger payouts.

In sum, the economist's solution to equity concerns is to simply shuffle around government revenue so as to smooth out any inequities. As such, transfers that return some portion of program revenue to households can be used to ensure equity, while preserving the economic efficiency of the congestion pricing program. So long as the transfers are not conditioned on driving behavior, drivers continue to face a congestion price when choosing to drive at peak times or through congested corridors.

This apparently simple solution, however, often runs into practical problems. The first is that the revenue is often dedicated for specified purposes (typically transportation infrastructure), and therefore may not be available to spend on ensuring fairness.

The second is that targeting transfers can compromise economic efficiency if it is not done carefully. To alleviate horizontal inequity, policymakers may wish to give larger transfers to those who pay more in tolls. Conditioning transfers on driving behavior (either tolls paid, or VMT), however, undermines the efficiency of congestion pricing by unraveling the intended incentives. If a driver knows they will get a larger transfer at the end of the year if they drive more miles or use toll roads more often, then this erodes the incentive to reduce driving, which is the entire goal of congestion pricing programs.

To preserve efficiency, then, transfers must be based on characteristics or behaviors that do not distort driving decisions, like household demographics, locational variables, or driving activity measured before a program was announced. In other words, to preserve efficiency, the transfers accompanying an externality tax must be lump-sum; that is, they cannot depend on driving behavior once the policy is in place.

It is against this backdrop that the current report examines strategies for mitigating the unequal costs of congestion pricing in ways that do not sacrifice efficiency. We focus on two common transfer designs: **uniform lump-sum transfers**, where revenue from the program is used to provide every road user with an equal payment, and **targeted lump-sum transfers**, where payments are conditioned on variables that do not undermine the goals for the policy.

We note that these transfer designs are not the only ways to address the equity issues in congestion pricing. Two alternatives that we do not cover in this report bear mentioning. The first alternative is to exempt low-income drivers. Exemptions can improve equity, but they erode the efficiency of congestion pricing because exempt drivers will again generate uncorrected externalities. Our goal in this analysis is to examine the

¹ In public economics, a *lump-sum transfer* is a payment to individuals that does not depend on the subsequent actions of that individual. Lump-sum transfers can be uniform (where every individual receives the same transfer) or targeted (where immutable individual characteristics, or individual actions taken prior to the policy are used to determine transfers).

feasibility of systems that address equity without compromising efficiency. Policymakers may ultimately prefer to use exemptions to address equity, but they should do so only after understanding the potential of policies that preserve efficiency.

A second design that we do not explicitly cover is income-tested transfers, or policies that transfer policy revenue only to low-income drivers. This type of policy will be more progressive than the uniform rebates we cover in this report, but they pose a significant administrative burden, as transportation authorities do not already possess credible income information. Moreover, they will perform more poorly on horizontal equity measures than will uniform transfers. Despite these concerns, income-tested transfers may be an attractive option in certain settings, especially where vertical equity concerns are acute, or income information is already available to the congestion pricing authority.

To understand how uniform or targeted transfers would impact equity, the report utilizes data from a congestion pricing experiment conducted in the Seattle metropolitan area by the Puget Sound Regional Council in 2005. The report documents differences in the direct burden experienced by the households in the sample due to their driving patterns, and then evaluates both the initial regressivity of the program and the potential benefits of targeted transfer programs that distribute revenue based on household characteristics like residential location, homeownership status, vehicles, number of children, age, gender, employment and student status, or number of weekly commute trips.

The report presents two major sets of findings. First, the burden of congestion pricing in the experiment has a very weak relationship to household income. This means that the dollar burdens of the policy are very similar across households with high incomes and households with low incomes. In turn, this means that the pricing program is quite regressive when viewed as a proportion of income. This is consistent with prior studies of the distributional burden of congestion pricing, which we review below.

These regressive effects can be overcome if some of the program revenue were used to reduce some of that burden. A direct payment from program revenues that was equal for all households that own and operate a vehicle would make the program moderately progressive. This echoes findings from the literature on carbon taxation (Metcalf 2007). A precedent for this type of policy is the California Climate Credit, which provides a uniform rebate to California households from a portion of revenues from the state's cap-and-trade program for carbon. We show that roughly half of all program revenue would need to be recycled as a uniform lump-sum transfer in order to make the program equity neutral.

The second set of findings pertains to the challenge of targeting payments in order to minimize the horizontal inequity in the pricing program's impact. Households with similar demographic and geographic characteristics have very different driving patterns, and subsequently face very different initial burdens from congestion pricing. The program costs the median household \$28.50 per week, but the standard deviation in this weekly burden was \$21.80, meaning that many households pay far less or far more than the median amount. Little of this variation (8 percent) is explained by income, and only a modest amount (24 percent) is predicted by an expansive array of other household demographic characteristics including homeownership, number of vehicles,

number of children, age, gender, employment and student status, number of weekly commute trips. Almost none of the variation is predicted by neighborhood (as opposed to household) characteristics. This is significant because it implies that the variables that are commonly used to condition cash transfers are poor indicators of the burden of a congestion pricing program and will thus have a limited ability to prevent those who travel most from being harmed by congestion pricing.

An alternative is to condition transfers not on demographic or neighborhood variables but instead on *past* travel behavior. For example, for a congestion charge that due to start in 2023, a policymaker could design a system that provides bigger rebates to households that drove more in a prior “baseline” year, say 2022. In theory, if the policymaker has access to driving data from a baseline year during which drivers were unaware that their driving behavior would be used to determine a future transfer size, economic efficiency is preserved. In practice, this is difficult to execute because policymakers are unlikely to have travel data before a program begins, and because there will be a need to update the measure over time, which introduces potential inefficiencies. Nevertheless, we study the effect of conditioning transfers on baseline travel as a “best case” scenario for targeting.

The Seattle experiment measured baseline driving behavior by installing measurement devices in the vehicles of volunteers before revealing that these devices would be used to introduce congestion pricing. In our data, we thus have a measure of driving behavior prior to the introduction of pricing that can be used as a baseline measure. Among the experiment’s households, pre-experiment average mileage is a strong predictor of program burdens, and pre-experiment counterfactual costs (what participants *would have paid* if the experiment tolls had been in place during the pre-experiment period) is an even stronger predictor. Even so, there remains substantial variation in program burdens that are not explained even by these driving data. This highlights the challenge of precisely targeting transfers because future programs are unlikely to have this sort of baseline measure available.

The final section of this report comments on the relevance of these results from the Seattle experiment to the California context. We compare data on travel behavior from the National Highway Transportation Survey and use the results from the Seattle experiment to extrapolate the likely regressivity of similar congestion pricing programs should they be implemented in California’s congested metropolitan areas. In brief, the basic features of the driving population in Seattle—namely the high variation in driving behavior, and the VMT-income relationship—are also present in California, which implies similar results would be seen in metropolitan areas in the Golden State.

Existing congestion pricing programs have tended not to include direct payments from toll revenues. Examples include both cordon zones in Milan, London, and Gothenburg and variable tolling on congested roads in Stockholm and Singapore. Despite concern that congestion pricing may be regressive in these cities, revenue from these programs is generally not transferred to low-income road users. Instead, the aforementioned road pricing programs all funnel excess revenue towards transit-related projects.

In California, authorities are actively developing proposals for cordon zones in both Los Angeles and San Francisco. In both cases, there is concern about the equity of such programs. Rebates, as well as exemptions, appear to be an active part of the policy conversation. This report helps offer insights about the potential to achieve progressivity and the potential role targeted payments could play.

Before presenting our empirical results, this report reviews the key lessons from prior research and discusses economic theory that suggests how data on driving behavior can be used to estimate the burdens of a congestion pricing policy.

Congested Pricing Research

Congestion Charges Tend to be Regressive

To understand the distributional impacts of congestion pricing, it is useful to distinguish between *absolute* and *proportional* regressivity. A policy is *absolutely* regressive if it costs more in gross terms to low-income households than it costs high-income households. A policy is *proportionally* regressive if it costs more to low-income households than high income households as a proportion of income. The latter is the conventional criteria for characterizing a policy as regressive and the benchmark that we adopt in this report.

The discussion of the distributional impacts of road pricing in economics dates at least to Foster (1974), who argued that road pricing would be progressive because of the tendency of rich families to drive more, and Richardson (1974), who countered that even if high-income households drive more, road use was unlikely to rise proportionally with income, and road pricing would therefore be proportionally regressive.

Following these theoretical treatments of regressivity in congestion pricing, recent empirical work investigating real-world road pricing policies generally suggests that if drivers are not compensated in some way—be that through rebates or through public investment—congestion charges are proportionally regressive (Levinson 2010, Anas and Lindsey 2011, Cohen D’Agostino, Pellaton and White 2020).² For example, in an analysis of road pricing in Dresden, Germany, Teubel (2000) concludes that without compensation, both the cost and the time savings induced by road pricing would be regressive, confirming the hypothesis presented by Richardson (1974). Schweitzer and Taylor (2008) similarly find that before accounting for the use of toll revenue, the incidence of tolls in Southern California is proportionally regressive. Cain and Jones (2008) report that the proposed tolling system in Edinburgh, Scotland would be regressive for households that drive. Because low-income households are less likely to drive, however, the authors find that average driving costs as a proportion of income are relatively flat across income quintiles. In an analysis of several proposed pricing programs in Paris, Bureau and Glachant (2008) report that while certain tolls may be absolutely progressive, none are proportionally progressive: “In relative terms, tolls are always more detrimental to low-income individuals, meaning that tolling is regressive.” Importantly, the stylized fact that uncompensated tolling is proportionally regressive holds whether time savings are accounted for (as in Teubel (2000) and Bureau and Glachant (2008)), or whether time savings are ignored (as in Cain and Jones (2008)).

Although the majority of empirical studies find that road pricing is regressive before compensation is considered, there is no clear consensus about whether or not investments funded with toll revenue can yield progressive

² In rare cases, research suggests that congestion pricing could be progressive *without* redistributing revenue, as in Santos and Rojey (2004), who simulate several congestion pricing programs in Britain. These findings typically require either large disparities in road use with income, or travel time benefits to other modes (e.g., urban busses) that are heavily used by low-income commuters.

outcomes. Eliasson and Mattsson (2006), for example, find that the fact that toll revenue is devoted largely to public transportation makes Stockholm’s congestion pricing system progressive on average. Schweitzer (2009), however, finds in a review of seven congestion pricing studies that charges are “broadly regressive,” with “the variance in the findings depend[ing] much less on differences in methodology or data than on policy design and the distribution of revenues.”

Existing research does agree, however, that uniform revenue redistribution to road users can generate progressive outcomes (e.g., Eliasson and Mattsson (2006) and Bureau and Glachant (2008)). Because gross toll payments tend to be lower for low-income households, low-income households are better off on average after uniform revenue distribution, and high-income households are worse off on average. To see this, imagine an example with 100 households, half of which are low-income (and pay \$50 in tolls per week) and the other half are high-income (and pay \$150 per week). If all of the toll revenue is redistributed in a uniform fashion, each household will receive a \$100 transfer, making low-income households better off on net, and high-income households worse off on net.

Our study confirms these main findings in the Seattle context, and it also addresses some topics that are largely absent in the prior literature. Missing from the prior literature is a discussion of the amount of revenue that is required to achieve progressivity through a lump-sum transfer program. Although progressivity may be theoretically attainable, in practice congestion pricing programs typically that require significant portions of revenue be used for the construction and maintenance of the tolling system, and planning bodies typically earmark excess funds for other public projects. It is therefore unclear whether progressive road pricing is feasible given these competing revenue uses. This report uses the data-rich Seattle setting to address this gap in the literature—we show that uniform payments can lead to a policy that is on average progressive while using as little as 52 percent of total revenue.

Horizontal Equity and Congestion Charges

There are considerably fewer research papers that discuss horizontal equity than there are papers that consider vertical equity in congestion pricing. Discussions of horizontal equity typically focus on difficulties in achieving Pareto improvements (outcomes where everyone benefits) through congestion pricing. For example, Arnott, de Palma, and Lindsey (1994) use a theoretical model of travel demand to demonstrate how different driver characteristics (schedule flexibility and the value of travel time) imply different benefits under congestion pricing with lump-sum rebates. That paper noted that congestion pricing may leave certain types of drivers worse off, even after toll revenue is redistributed lump-sum.

Several studies have demonstrated the wide variation in the individual costs of congestion pricing programs empirically, which provides suggestive evidence that lump-sum payments may leave some road users worse-off, even after accounting for improvements in traffic conditions. In an analysis of toll lanes in Southern California, Schweitzer and Taylor (2008) find that within income groups, heavy road users pay more than twice the group average in tolls. Eliasson and Mattsson (2006) show that the benefits of congestion pricing in

Stockholm vary significantly by geography, with large travel time benefits and low costs to city dwellers, and larger costs and smaller benefits accruing to drivers who live outside of the city. Of particular concern is the possibility that differences in driving patterns will leave a fraction of low-income drivers worse off even with progressive revenue redistribution, as highlighted by Vickerman and Börjesson (2021).

This report extends the understanding of horizontal equity in congestion pricing in two ways. First, using tools developed by Sallee (2019), this report offers an empirical test of whether or not variations in driving behavior makes it impossible to make everyone better off (achieve a Pareto improvement) with congestion pricing. Second, in exploring the distributional impacts of different programs for redistributing congestion pricing revenue (e.g., uniform versus VMT-based payments) this report highlights a policy tradeoff between programs that perform well in minimizing the number of losers and programs that perform well on progressivity measures.

Congestion Pricing in Economic Theory

This section briefly reviews concepts that underpin the empirical analysis performed later. These concepts are standard fare in economics, but they are frequently misunderstood. We provide a more detailed discussion of the economic theory behind congestion pricing in the Appendix.

Pigouvian Taxes and Economic Efficiency

When a good sold in a market creates an externality, the free market equilibrium will feature a quantity of that good that is inefficient from society's point of view. Pigou (1920) proposed a solution to this inefficiency, which was to add a tax equal to whatever costs were being ignored by the free market (i.e., the size of the externality in dollar terms). Such a tax is often referred to as either a Pigouvian tax or a corrective tax. A corrective tax raises the price to buyers and/or lowers the price received by sellers, which causes a reduction in the equilibrium quantity sold. This generates a social gain by reducing the externality, but the incidence of the tax also causes losses of economic welfare to some agents.

Economists often refer to corrective taxes as "efficient," meaning that the market outcome after the imposition of a corrective tax is associated with higher total welfare when compared to the market without a corrective tax. In other words, the total benefits from reducing externalities (e.g., congestion) outweigh the total reduction in welfare from the costs of the policy to producers or consumers (e.g., the tax burdens experienced by drivers).

In theory, it is possible to make *everyone* better off using corrective taxation. By the above definition of efficiency, the total welfare in the economy is larger than it was before the imposition of the tax, so it should be possible to ensure that everyone benefits relative to the status quo. In practice, however, the costs and benefits are often distributed unevenly among economic agents, and the government collects revenue in the process of taxing, which needs to be returned to individuals in order to reach a state where every individual is better off than they were prior to the corrective tax.

Transfers Can Mitigate Regressivity Without Sacrificing Efficiency, but Some May Still Be Left Worse Off.

Given that some individuals may be made worse off by a corrective tax, Governments may wish to design transfer payments to individuals to make them better off than they were prior to the imposition of the tax. Designing these transfers comes with two significant challenges.

The first challenge is that if policymakers wish to maintain the economic efficiency induced by the tax, these transfers cannot be tied directly to the amount that someone pays in corrective taxes. If, for example, people

knew that they would receive a refund for any corrective taxes they pay, then they would not face any incentive to avoid externality-generating activities (like driving on congested streets), and the goal of the corrective tax (to reduce externalities) would be undone by these transfers.

This challenge may sound familiar to policymakers who have faced equity-efficiency tradeoffs in other settings. An important result from economic theory, however, is that an equity-efficiency tradeoff need not exist in the case of externality-correcting taxes. In this paper, we focus on a class of transfers where there is *not* an equity-efficiency tradeoff. Specifically, any transfer that is not tied to an individual's actions (i.e., a lump-sum transfer) will not undermine economic efficiency. As we show below, such transfers can generally be designed to be progressive on average. Some examples of transfers that would not undermine efficiency are (a) a uniform transfer, where everyone gets an equal share of the tax revenue, or (b) a transfer based on actions taken before the tax was announced (e.g., VMT prior to the passage of a congestion tax).

The second challenge is that transfers designed to mitigate tax burdens are often imprecise. When individuals have underlying differences that lead them to consume externality-generating goods (e.g., vehicle trips) at different rates, but policymakers only observe certain characteristics of individuals that imperfectly capture these preferences, some individuals are inevitably overcompensated, while others are undercompensated. Sallee (2019) shows if there is too much unpredictable variation in the initial consumption of the good—and thus in the initial burdens created by the policy—then there is not enough “extra surplus” to be redistributed through transfers to ensure that no one is harmed on net. In practice, this result demonstrates that it is usually very difficult to introduce an efficiency-enhancing correct tax without making some people worse off.

In what follows, we use data from a congestion pricing experiment in Seattle to understand empirically how these efficient transfers perform along different equity measures, with the goal of helping policymakers evaluate the relative merits of different uses of congestion pricing revenue.

Summary of Seattle Experiment

In 2005, the Puget Sound Regional Council (PSRC) recruited 255 households in the Seattle area to participate in a congestion pricing pilot study. The goal of the study was to test the viability of congestion pricing technologies and to evaluate consumer responses to congestion pricing. We briefly summarize the program here. Further details are available in the Council's report on the program (Puget Sound Regional Council 2008).

Households enrolled in the program completed a survey before the program began. Then each participating household had all their vehicles fitted with a GPS monitor that tracked vehicle location and a dashboard meter that conveyed pricing information to drivers based on which roads they travelled on in the region (see Figure 9 in the appendix). Households drove as normal for several months during a baseline period before pricing began. The length of the baseline period varies across households based on the date on which their devices were installed. In addition, a problem with the technology required the units to undergo a software update during the baseline period. This process took vehicles out of commission for a period of time, and we drop data from this period (February 14, 2005 – March 8, 2005).

All participants began facing prices in July of 2005, ending in April of 2006. We refer to these 10-months as the treatment period. During the treatment period, participants had to pay a cost per mile for use of certain roads at certain times. Roads were divided into three categories: tolled freeways, tolled non-freeways, and unpriced roads.

Pricing differed by time of day and weekday versus weekends. Rates did not vary with congestion conditions in the way that a fully efficient, real-time road pricing program would. An advantage of this program is that it was transparent and consistent, which arguably allowed for maximal driver response because the rates were known with certainty. The rates, which are summarized in Table 1, were designed to approximate the average social cost (externality) of driving per mile in each of the time segments and types of roadways (Puget Sound Regional Council 2008).

Table 1. Toll Schedule

Time Period	Hours/Days	Toll Rate (Freeway)	Toll Rate (Non-freeway)
PM Peak Period	4-7 PM (Mon-Fri)	\$0.50/mi	\$0.25/mi
AM Peak Period	6-9 AM (Mon-Fri)	\$0.40/mi	\$0.20/mi
Midday	9 AM – 4 PM (Mon-Fri)	\$0.15/mi	\$0.075/mi
Weekend Peak	10 AM – 7 PM (Sat-Sun)	\$0.20/mi	\$0.10/mi

Time Period	Hours/Days	Toll Rate (Freeway)	Toll Rate (Non-freeway)
Off Peak & Weekend Off Peak	7-10 PM (Mon-Sun), 7-10 AM (Sat-Sun)	\$0.10/mi	\$0.05/mi
Late Night	10 PM-6 AM (Mon-Sun)	N/C	N/C

Notes: This table shows the fees per mile used in the experiment. They range from no fee during the late-night hours to \$0.50 per mile on freeways during the evening peak. During the treatment period, participants were given information about these rates as well as details on the priced routes. A helpful map is included in Figure 9 which appears in the appendix.

The dashboard meter looks similar to a traditional fare meter in a taxicab. During both the baseline and the treatment period, the GPS device recorded locations and travel time for every trip. During the experiment, the dashboard meter displayed a cumulative total cost of each ride, as well as the price per mile of the current segment, along with the name of the road. During the baseline period, the meter simply displayed the name of the road without any pricing information. During the experiment, participants were given a map of the Seattle area that detailed which roads were in each price tier. They could log into an online account and review the route and cost of any past trip, and they received monthly statements detailing their travel.

The program administrators wished to ensure that program participants would not experience a financial penalty. As a result, each household received an endowment in an online account before the program began. Their account was then drawn down over the course of the experiment, and households were allowed to keep the remaining balance when the program ended. If their account was driven to zero, they would not be required to pay anything. Thus the economic incentive facing drivers is precisely that of the cost per mile, but it should be kept in mind that participants did not actually pay a bill, but rather saw a declining bank account.

Program administrators designed the program hoping to prevent anyone from drawing their account to zero, as this would leave them with no financial incentive to modify their behavior and effectively cause them to drop out of the experiment. As a result, administrators used baseline driving behavior to project total program costs and then provided endowments that would be expected to cover the full cost of driving during the program period with some additional cushion. Importantly, participants were not told that their driving patterns during the baseline would have any financial implications, and in fact the aims of the study were not explained to the participants until pricing began.

Individuals might be expected to respond to congestion pricing in different ways in the long and short run. The fact that the program was explicitly temporary means that some long-run responses, such as moving or forming a carpool, are not observed. The aim of this report is not to evaluate behavioral effects of the program, but instead to evaluate its distributional impacts.

Table 2 presents summary information about the cost of the program among participant households.

Table 2. Summary Household Statistics

Variable	Mean (PSRC)	SD (PSRC)	Mean (NHTS)	SD (NHTS)
Income	64000	47528	106202.21	65152.14
Age (Head of HH)	45.914	15.140	40.49	21.37
Commute Distance	7.847	10.925	3.89	8.57
Number of Drivers	1.548	0.707	2.51	0.91
Weeks Pre-Experiment	13.608	1.179	-	-
Weeks Post-Experiment	32.412	4.636	-	-
Dollars Per Week Pre-Experiment	34.83	21.87	-	-
Dollars Per Week Post-Experiment	33.12	21.84	-	-

Note: The first two columns of this table show summary statistics for participants in the Puget Sound Regional Council Traffic Choices Study. The last two columns show summary statistics for NHTS respondents in the Seattle-Tacoma MSA. *Weeks Pre-Experiment* and *Weeks Post-Experiment* represent the average number of weeks that we observe PSRC participants during the unpriced and priced periods, respectively. *Dollars Per Week Pre-Experiment* represents the average cost to participants during the pre-experiment period if their trips had been tolled. *Dollars Per Week Post-Experiment* is the average weekly cost incurred by participating households during the tolled period.

Participants were recruited from an initial random sample of Seattle area residents who were contacted by telephone. Potential participants were screened on several dimensions. Households that did not own a vehicle or that owned more than three vehicles were ineligible. Households that did not currently have at least one commuter who used roads to be tolled on a regular basis were ineligible, as were commuters who reported that they regularly carpooled to work. The data include statistical survey sampling weights that can be used to ensure that data in the study match the population along specified characteristics. Even so, the study consists of a set of participants interested and willing to participate in the study and therefore may not be representative of the full population.


In the next section we present the results of the study and how they can support targeting payments to make the program more equitable. We also compare the participants in the study to the sample of drivers from the Seattle metro area from the National Household Travel Survey to gauge how much the sample differs from the average. Households in the study had household heads that are slightly older than average. Households in the experiment drive more than average. This is to be expected because the experiment was limited to those who drive a minimum amount.

Results

Direct Burdens of Congestion Pricing are Regressive

The initial burden of the congestion pricing program was relatively flat over the income distribution. Figure 1 shows the median (the bar), interquartile range (the box), and 5th and 95th percentiles (the whiskers) for several different income ranges. Richer households do have higher program costs on average, but the difference across income groups is modest, and it is dwarfed by the differences *across* households *within* each bin. A bivariate regression of weekly program cost on household income yields a coefficient that shows every \$10,000 increase in household income is associated with an increase of \$1.58 per week in program costs (standard error of \$0.25).

Figure 1. Costs Rise with Income but are Diffuse



Note: This figure shows the median (solid bar), interquartile range (box height) and 95th percentile range (vertical line) of household average congestion charges per week for each category of self-reported income. The data are from the 2004 Puget Sound Regional Council's Transportation Choices Study.

Table 3. Costs Rise with Income but are Diffuse

Income	0-20k	20k-40k	40k-60k	60k-80k	80k-100k	100k-200k	200k-1M
75 th percentile of cost	41.33	28.26	43.76	40.79	53.51	58.17	83.15
Median cost	26.85	12.281	26.00	32.72	39.45	46.63	47.59
25 th percentile cost	17.74	9.58	18.30	23.65	25.88	29.04	23.23

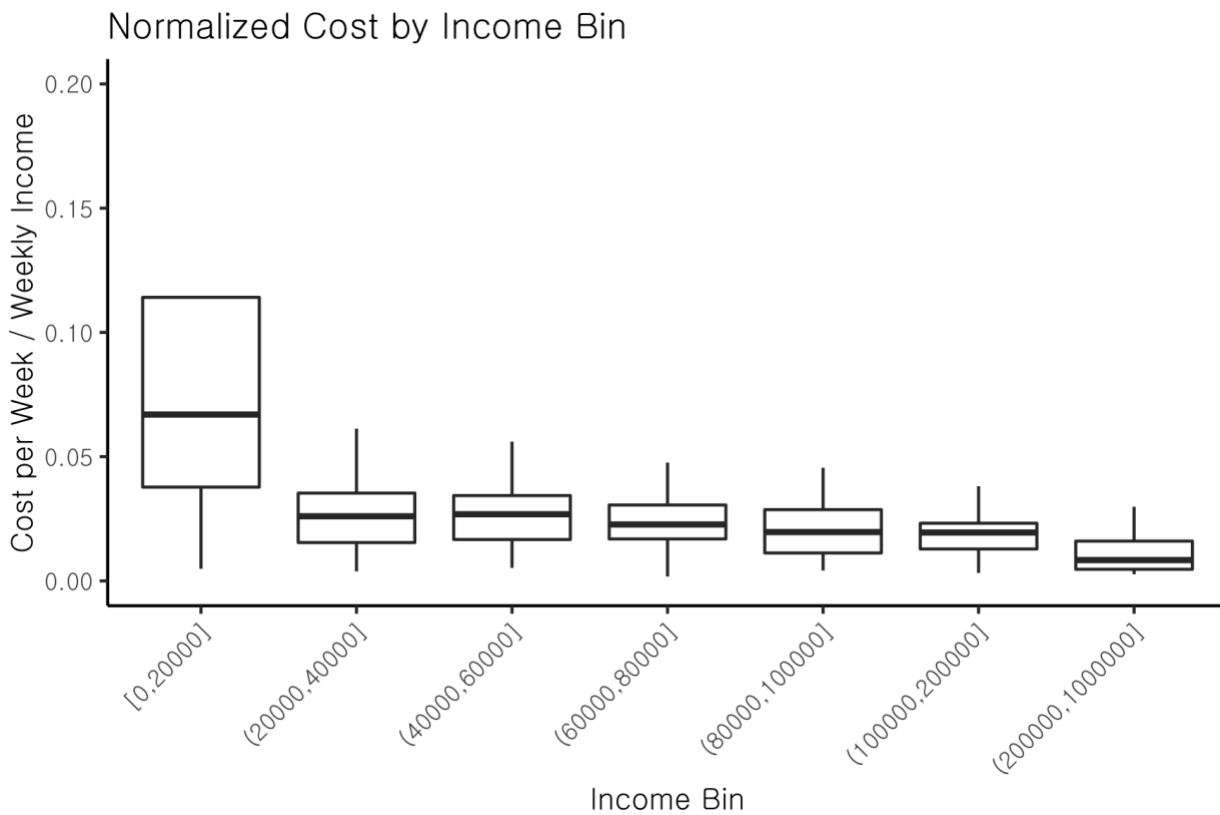
Note: This table accompanies Figure 1. It shows the average weekly cost of the 2004 Puget Sound Regional Council's Transportation Choices Study by income bin.

As discussed in the next section, the large dispersion in direct burdens across households with similar income implies that uniform transfers to all households equally, or even a variable transfers based on a sliding scale of income, will result in gross overcompensation of some households while being insufficient to offset program burdens for many others. The challenge then is to ask whether better targeting of the rebates is possible so as to reduce the dispersion in costs across households.

Table 3 represents a subset of this same information in table form. It shows that the median cost nearly doubles across income categories, but the increase in income goes up by a factor of 10. The relatively flat relationship between weekly burdens and income necessarily implies that the program's burdens will be regressive when measured as a percentage of household income. The standard definition of regressivity in economics is a burden or cost that is a higher proportion of income for low-income households.

Figure 2 demonstrates visually that, measured as a proportion of income, the direct burdens from the congestion pricing program are quite regressive. Low-income drivers in the sample paid an average of 6.7 percent of their weekly income in congestion charges, whereas high income drivers paid an average of 0.8 percent of their weekly income. As was the case with the raw costs, the variation in proportional cost within income groups is quite substantial.

Figure 2. Congestion Charges are Regressive and Diffuse



Note: This figure shows the median (solid bar) interquartile range (box height) and 95th percentile range (vertical line) of household average congestion charges per week as a fraction of weekly income, for each category of self-reported income. The data are from the 2004 Puget Sound Regional Council's Transportation Choices Study.

It is worth emphasizing here that the ultimate progressivity or regressivity of the policy depends also on the distribution of benefits from congestion alleviation, as well as on the uses of the revenue raised. In this case, the program was a small pilot, so it is implausible that there were congestion benefits, but the distribution of benefits will be important for any program deployed at scale.

Revenue can be used either to fund some expenditure that would not otherwise occur, in which case the distribution of benefits from that expenditure is directly relevant, or to offset the need for raising revenue through some other means, in which case the relevant consideration is how these distributional impacts compare to the alternative funding mechanism. Or revenue could be directly redistributed. The latter possibility is the focus of the analysis in the next section.

Performance of Conditioned Transfer Programs

The goal of an externality-correcting tax is not to punish economic actors, but rather to create the right marginal incentives.

The above figures show that the burdens of congestion pricing are highly varied. The central question of this report is whether these program costs can be predicted by variables that can be used to set targeted payments that mitigate the harm to households that face the greatest impact from congestion in a way that does not compromise the benefits of the program for alleviating congestion.

Table 4. Predicting Program Costs

Model	(1) VMT	(2) Demographics	(3) Toll Spending	(4) All Variables
Adjusted R-squared	0.40	0.24	0.57	0.58
Mean Absolute Error (\$/Week)	10.47	13.13	7.19	6.96
Vehicle Miles Traveled	✓			✓
Demographic Information		✓		✓
Pre-Period Dollars Per Week			✓	✓

Note: This table displays the R-squared values from regressions of weekly program costs on observable variables using data from the 2004 Puget Sound Regional Council’s Transportation Choices Study. Checkmarks denote which variables are used to predict program costs in each column.

Each row in Table 4 represents the R-squared from regressions predicting weekly program costs. Column (1) shows that the number of vehicle miles traveled during the pre-pricing period alone explains 40 percent of the variation in cost. Column (2) shows that demographic information—income, number of vehicles, number of children, age, gender, employment and student status, homeownership status, and the number of weekly commute trips—together explain only 24 percent of the variation in program costs. Pre-period dollars per week, which measures how extensively program participants were using tolled roads before the pricing period began, is the strongest predictor of program costs (57 percent).

As discussed in the Introduction, directly conditioning transfers on vehicle miles traveled or baseline period costs may undermine the policy’s goals of reducing traffic congestion. One cannot use these variables during a period when prices are in effect because that would undo the intended incentive of road pricing (if drivers

know they will receive higher payments in the form of a future rebate, the program's efficiency would be compromised). In theory, if data can be gathered from a baseline period when drivers do not know that their behavior is being used for these purposes (so that they cannot manipulate their driving strategically during the baseline period), then the incentives need not be misaligned. But, these types of data do not readily exist, and there will inevitably be challenges in incorporating new households that form or move to an area over time. Thus, it may or may not be practical to use these variables to condition transfers. We analyze the performance of this type of targeted lump-sum transfer in this report to determine whether gathering this type of information would allow for the design of transfers that perform well in reducing unequitable policy burdens.

The results from this table illustrate that that the information commonly used to condition other government transfers does little to explain the differences in program costs across household. Given information on the intensive margin of driving habits, however, a policymaker could much more accurately target payments.

Sallee (2019) derives an empirical test that relates the average absolute error from a prediction regression, like those in Table 3, to the size of the average welfare gain from reductions in the externality. When the average absolute error is greater than two times the average welfare gain, Sallee (2019) demonstrates that it is impossible for an externality-correcting tax with targeted transfers based on the covariates in the prediction regression to yield a Pareto improvement. Instead, even though the tax achieves a Kaldor-Hicks efficiency improvement, it is necessarily the case that some people will be made worse off.

To test this condition in the current context, we need an estimate of the size of the benefits of reduced congestion that one might expect if the road pricing program was implemented at scale. This can then be compared to the average absolute errors from the regression. How big are potential gains? A full-cost (this includes the value of travel time and value of reliability) own-price elasticity of -0.36 from Small et al. (2006), implies a reduction in VMT in Seattle of 12 percent, or 53 kilometers per week for the average vehicle. Applying an estimate of the marginal external cost of congestion of 3.5 cents per vehicle-kilometer from Couture et al. (2016) to this change in VMT yields an average per-person externality of \$1.86 per vehicle per week.

Even before taking into account program costs, the magnitude of this externality compared to the best-case mean absolute prediction error (\$10.47 per vehicle per week when using VMT-based transfers) shows that a true Pareto improvement (everyone made to benefit from the implementation of an efficiency-enhancing congestion pricing program) is impossible. The basic idea is that there is simply too little precision in targeting transfers to possibly allocate revenue from the program sufficiently accurately to make everyone better off.

Alternatives for Conditioning Transfers


We next compare the performance of a uniform transfer to a VMT-based transfer. Each of these alternatives have attractive properties: Uniform rebates are simple, transparent, and are proportionally progressive.

Uniform Lump-Sum Transfers

Figure 3 shows the distribution of net burdens under a program that recycles all revenue back to all households in a uniform lump-sum. Negative costs in the figure mean that the rebate exceeds the cost paid. Overall, this program is quite progressive, with lower-income households experiences larger gains, measured as a proportion of their income. Note, however, that there is still a large dispersion in the net gains within each income category. This means that, within each income category, there are some people bearing a much higher burden than others. This is because driving behavior differs a great deal across otherwise similar households.

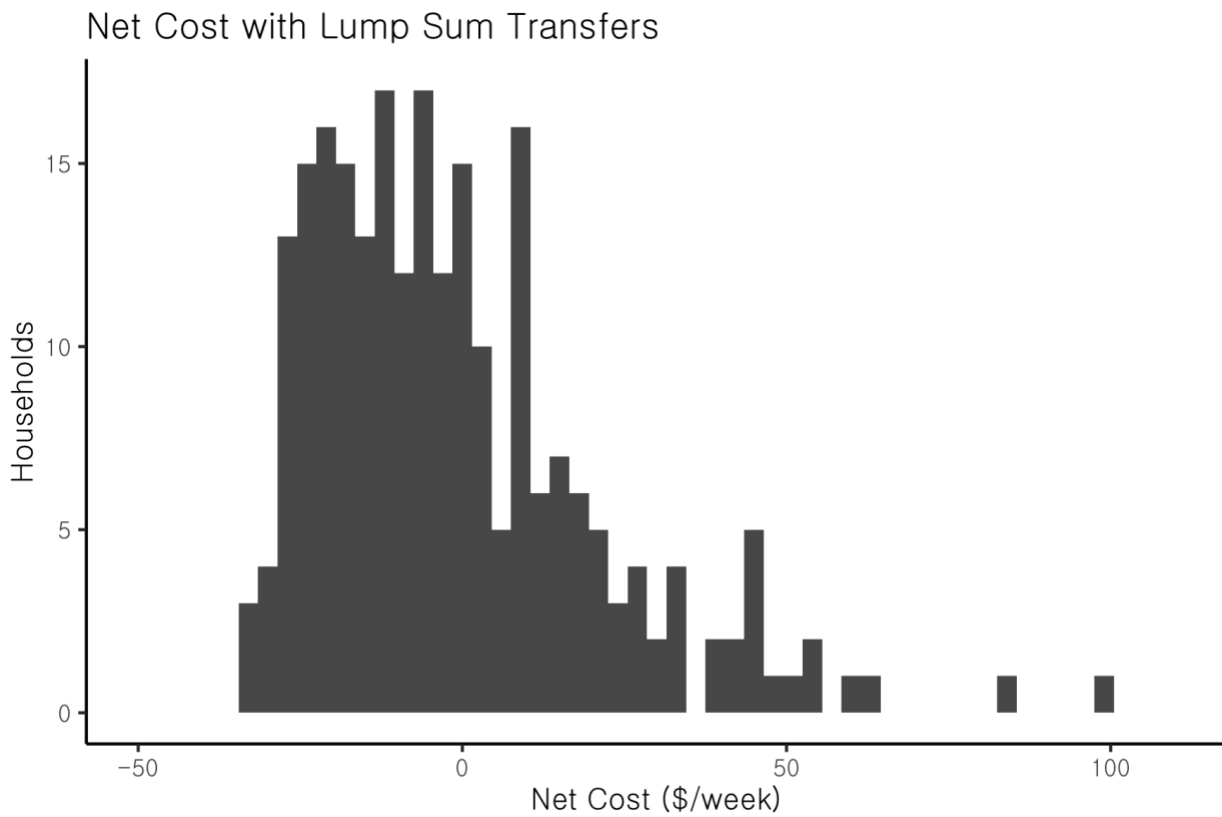
Note that the dispersion of burdens is larger for low-income households largely because the burdens are displayed as a fraction of income. High and low-income groups have similar variation in *gross* burdens within the income group, but in the case of high-income households, even the largest burdens represent just a small fraction of total income. For low-income households, on the other hand, small changes in burden represent a relatively large share of overall income, resulting in a mechanically higher dispersion of proportional burdens for this group.

Figure 3. Net Burdens Under Uniform Lump-Sum Transfers by Income Category



Note: This figure shows the median (solid bar) interquartile range (box height) and 95th percentile range (vertical line) of household average congestion charges per week as a fraction of weekly income, after recycling revenue using a uniform rebate. The data are from the 2004 Puget Sound Regional Council’s Transportation Choices Study.


Figure 4. Distribution of Net Burdens Under Uniform Lump-Sum Transfers



Note: This figure displays the distribution of program costs after recycling revenue using a uniform rebate. The data are from the 2004 Puget Sound Regional Council's Transportation Choices Study.

Figure 4 shows the distribution of these same net burdens in level dollars for the entire sample. Households with a positive cost are net losers (before considering potential congestion reduction benefits). Because there is substantial variation in net burdens, a uniform lump-sum payment leaves many people as net losers from the program. These are the households with higher than average travel demands, who are being undercompensated for their cost by the lump-sum payment.

Figure 5. Distribution of Net Burdens Under Targeted Transfers, by Income Category



Note: This figure shows the median (solid bar) interquartile range (box height) and 95th percentile range (vertical line) of household average congestion charges per week as a fraction of weekly income, after recycling revenue using a targeted payment based on VMT. The data are from the 2004 Puget Sound Regional Council’s Transportation Choices Study.


VMT-Based Transfers

A VMT-based transfer most accurately predicts program costs, and will therefore produce results that are less likely to overcompensate uninjured households at the expense of large losers. But, this approach is less progressive; richer households drive more on average, so a VMT-based transfer program directs more money to higher income households. On average, a VMT-based transfer program is proportionally progressive, but much less so than the uniform policy, which can be seen by comparing the modest slope of average costs against income in Figure 5 (which shows a targeted payment) to the steeper slope in Figure 3 (which shows a uniform payment).

In particular, households with incomes less than \$60,000 a year are still net winners, on average, in the VMT-targeted program. But the positive correlation between VMT and income implies that, in practice, a VMT-targeted system gives larger payments to richer households. This erodes the progressivity of the policy and leaves many lower-income households worse off under the targeted program than under a uniform policy. Because low-income households tend to drive less and incur fewer tolls, a uniform rebate “overcompensates”

that group on average. While blunt, this feature of a uniform rebate helps address the wider variation in program costs among low-income households.

Figure 6. Distribution of Net Burdens Under VMT-Targeted Transfers



Note: This figure displays the distribution of program costs after recycling revenue using a VMT-targeted payment. The data are from the 2004 Puget Sound Regional Council’s Transportation Choices Study.

Targeting based on VMT does reduce the overall dispersion of program costs, reducing the largest losses and erasing some of the largest cases of overcompensation. This can be seen by looking at the compression of the distribution in Figure 6 (which shows a targeted transfer program) versus Figure 4 (which shows the uniform program).

Thus, in practice, a uniform payment results in a wider range of program outcomes, but it is significantly more progressive. This presents a tradeoff for policymakers. One might wish to use better targeting to reduce the overall fraction of households who are made worse off by the congestion pricing scheme (to reduce political opposition) and to reduce the burden on the heaviest drivers, but better targeting tends to erode the progressivity of the policy when wealthier households have higher program costs (in dollar terms, not as a proportion of their income. This suggests that, at least in some contexts, there will be a tension between horizontal and vertical equity.

It is important to note that the policies we are exploring do not represent the limits on progressivity. Conceptually, it is straightforward to make transfers more progressive. For example, one could provide transfers only to households below some income threshold. These transfers could be targeted or uniform. In either case, this would be more progressive than the baselines we study.

More generally, transfers could be some sliding scale of income, and they could potentially be complicated and nonlinear. There is little limitation on how progressive such a system could be, in theory. At this point, however, the transfers begin to function directly as a form of income taxation. As a practical and administrative matter, it is highly unlikely that a transportation authority would have the means to collect and verify income information for a large number of households. So, a system that relied heavily on income for designing the rebates would practically have to operate through an existing income tax or in coordination with government agencies that already have reliable income information. But, such a system then becomes indistinguishable from a system where the tax authority keeps the revenue from the congestion pricing scheme and uses that to lower the income tax. This is an interesting and important possibility, though usually the income tax is not administered at the same level of government as any congestion pricing scheme, but we consider such a system outside the scope of our analysis.

Our focus in this report is on understanding the equity implications of uniform transfers, which are commonly invoked in policy design discussions; and targeted schemes, which can be designed to minimize the number of people made worse off in the hopes of making the program appealing to as many as possible. We interpret the uniform transfers as a realistic policy. In this report, we explore unrealistically complicated targeting schemes, as well as relatively simple, feasible ones, in order to understand what the upper bound possibility is for designing a system that minimizes the number of people who are net financial losers, with an eye to the political implications.

Measuring the Regressivity of Congestion Pricing

The above analyses suggest that (a) without any compensation, congestion pricing would be highly regressive, and (b) uniformly redistributing all the toll revenue would make the Seattle congestion pricing program strongly progressive on average. Because of various practical and political realities, however, a complete redistribution of revenues is unlikely. Can congestion pricing still be made progressive if some of the program revenues are used to cover costs or fund infrastructure projects?

To explore this question, we use the concept of the Gini coefficient, which measures the level of income inequality in an economy; a higher Gini coefficient means an income distribution is less equally distributed. Table 5 (below) uses the Gini coefficient to determine the minimum fraction of the revenue that must be redistributed to result in an income distribution that is less concentrated than the status quo. In other words, we solve for the fraction of revenue required to fund a uniform lump-sum payment program so that the distribution of income (after accounting for the burdens imposed by the congestion pricing program) is no more unequal than the distribution of income before implementing a program.

Rows 1 and 2 of Table 5 show that without redistributing any money, a congestion pricing policy in Seattle would increase the Gini coefficient. That is, the program is regressive before recycling the revenue. Row 3 shows that if 100 percent of the revenue is redistributed uniformly to households in the study, then after-toll incomes would be more equal than the status quo (a lower Gini). That is, the program is progressive after recycling all revenue. The main new result is in Row 4: if policymakers adopt uniform payments, they must distribute at least 52 percent of the revenue to make the policy result in a more equal distribution of after-toll incomes than the status quo. This calculation can help policymakers determine the minimum amount of revenue that needs to be devoted to transfers in order to undo the regressive impacts of congestion pricing programs.

Table 5. Using the Gini to Measure Regressivity

Program	Gini Coefficient
Status Quo (No Congestion Pricing)	0.2947
Congestion Pricing without Rebate	0.2984
Congestion Pricing with 100 percent of Revenues Returned Lump-Sum	0.2914
Congestion Pricing with 52 percent of Revenues Returned Lump-Sum	0.2947

Note: This table displays the Gini coefficient calculated under various congestion pricing and rebate programs. The data are from the 2004 Puget Sound Regional Council’s Transportation Choices Study; the Gini coefficient is calculated using annual income, plus annual toll rebates, less expected annual congestion tolls.

Applicability to California

The data from the Seattle study provide two takeaways for policymakers. First, designing well-targeted transfers to offset losses from congestion charging is difficult. This results from the fact that there is still significant variation in the use of toll facilities that cannot be explained by observable characteristics like VMT or household demographics. Second, because of the weak correlation between VMT and income, there is a tradeoff between programs that perform well in minimizing the number of losers (e.g., VMT based transfers) and programs that perform well on progressivity measures (e.g., uniform transfers).

A comparison between NHTS data in Seattle to NHTS data for California metro areas suggests that these challenges will also exist in the California setting. Table 6 summarizes income, VMT, vehicle use and an

indicator of congestion on daily commutes from Seattle and the major metro areas in California. Broadly, average indicators are similar across the locations.


Most importantly, the relationship between VMT and income looks to be quite similar across urban areas in California and Seattle, both in terms of the mean and in the degree of dispersion. This is shown in Figure 7. These general descriptive statistics do not provide compelling evidence that travel habits are substantively different for residents of Seattle as compared to residents of Californian cities.

As a result, a simulated congestion pricing program in metropolitan areas of California exhibits similar characteristics. High income people drive modestly more than do low-income people, but predicted tolls constitute a much smaller share of their total income. A uniform payment from toll revenue would therefore be progressive on average (see Figure 8), but, as in the case of Seattle, there are still losers and winners within each income group due to the wide variation in driving behavior within groups.

Table 6. California – Seattle Comparison


MSA	Average Income	VMT	Uses Car Daily	Congestion
Seattle-Tacoma	123109	11993	0.904	0.715
Los Angeles	116019	12252	0.945	0.703
Riverside	84947	13242	0.955	0.857
Sacramento	105869	12865	0.928	0.757
San Diego	111083	13072	0.937	0.748
San Francisco	139979	12670	0.883	0.753
San Jose	148014	12310	0.911	0.772

Figure 7. The VMT-Income Relationship in California Cities and Seattle



Note: This figure shows the relationship between income and VMT in Seattle (teal) and California metro areas (pink). The data used to produce this graph are from the National Household Transportation Survey.

Figure 8. Predicted Proportional Burden, CA



Note: This figure shows the median (solid bar) interquartile range (box height) and 95th percentile range (vertical line) of household average congestion charges per week after, recycling revenue recycling, of a hypothetical Seattle-style congestion pricing program in California metro areas. The data used to produce this graph are from the National Household Transportation Survey.

Conclusion

This report analyzes the equity and distributional implications of road congestion pricing by examining a pilot experiment from the Seattle metro area that implemented road pricing. Road pricing is a promising approach to dealing with inefficiencies in road use, but it may lead to inequitable outcomes.

Our analysis both confirms several conventional views and adds some new insights for policy design. Our data align with prior research in finding that the initial burden of road pricing is regressive when burdens are measured proportional to income, but that pricing can be made progressive by recycling revenue from the program back to households in a lump-sum.

Our first finding is that roughly half of the toll revenue must be recycled to make the program equity-neutral, that is, to undo the regressive impact. Our second finding is that most of the variation in burdens from road pricing cannot be predicted by household or neighborhood characteristics that would plausibly form the basis of targeted transfer programs. This suggests that there is limited scope for minimizing harm to those with the highest driving needs and demonstrates that road pricing programs will inevitably create net harm to some households. Our analysis also highlights potential tradeoffs between average progressivity and targeted reforms, as more precise targeting tends to lead to less progressive transfers. These findings help identify some of the key trade-offs that must be assessed in the design of road pricing programs.

References

- Arnott, Richard, André de Palma, and Robin Lindsey. 1994. The welfare effects of congestion tolls with heterogeneous commuters. *Journal of Transport Economics and Policy* 28(1): 139-161.
- Bureau, Benjamin, and Matthieu Glachant. 2008. Distributional effects of road pricing: Assessment of nine scenarios for Paris. *Transportation Research Part A: Policy and Practice* 42(7): 994-1007.
- Cain, Alasdair, and Peter M. Jones. 2003. Using public consultation in developing Edinburgh's congestion-charging-based transport strategy. *Transportation Research Record* 1839(1): 89-97.
- Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2018. Speed. *Review of Economics and Statistics*, 100(4), 725-739.
- Eliasson, Jonas and Lars-Göran Mattsson. 2006. Equity effects of congestion pricing: quantitative methodology and a case study for Stockholm. *Transportation Research Part A: Policy and Practice* 40(7):602-620.
- Foster, Christopher D. 1974. The regressiveness of road pricing. *International Journal of Transport Economics* 1(2): 133-141.
- Richardson, Harry W. 1974. A note on the distributional effects of road pricing. *Journal of Transport Economics and Policy* 111(1): 82-85.
- Fowlie, Meredith and Jeffrey M. Perloff. 2013. Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocations? *Review of Economics and Statistics* 95(5): 1640-1652.
- Kaplow, Louis. 1989. Horizontal equity: measures in search of a principle. *National Tax Journal* 42(2): 139-154.
- Levinson, David. 2010. Equity effects of road pricing: A review. *Transport Reviews* 30 (1):33-57.
- Metcalf, Gilbert E. 2007. A proposal for a U.S. carbon tax swap. Brookings Institution Discussion Paper.
- Montgomery, David W. 1972. Markets in licenses and efficient pollution control programs. *Journal of Economic Theory* 5: 395-418.
- Pigou, Arthur. 1920. *The economics of welfare*. London: MacMillian & Co.
- Puget Sound Regional Council. Traffic Choices Study: Final Report. Online Report, 2008. www.psrc.org/sites/default/files/traffic-choices.pdf
- Sallee, James M. 2019. Pigou creates losers: On the implausibility of achieving Pareto improvements from efficiency-enhancing policies. National Bureau of Economic Research Working Paper. <https://www.nber.org/papers/w25831>

- Santos, Georgina, and Laurent Rojey. 2004. Distributional impacts of road pricing: The truth behind the myth. *Transportation*. 31(1), 21-42.
- Schweitzer, Lisa. 2009. The empirical research on the social equity of gas taxes, emissions fees, and congestion charges. *Transportation Research Board Special Report* 303:1–26.
- Schweitzer, Lisa, and Brian D. Taylor. 2008. Just pricing: the distributional effects of congestion pricing and sales taxes. *Transportation* 35(6): 797-812.
- Small, Kenneth A., Clifford Winston, Jia Yan, Nathaniel Baum-Snow, and José A. Gómez-Ibáñez. 2006. Differentiated road pricing, express lanes, and carpools: Exploiting heterogeneous preferences in policy design [with comments]. *Brookings-Wharton Reports on Urban Affairs*, 53-96.
- Kristoffersson, Ida and Maria Börjesson. 2021. Urban congestion charging in transport planning practice. In Roger Vickerman (Ed.), *International Encyclopedia of Transportation*, Vol. 6: 206-213. Amsterdam: Elsevier.
- Vickrey, William. 1959. The optimum trend of prices. *Southern Economic Journal* 25(3): 315-326.

Appendix

Congestion Pricing in Economic Theory

Pigouvian Taxes and Economic Efficiency

Since Pigou (1920), economists have favored the idea of correcting market failures from externalities with a tax that imposes the unpriced externality on those responsible for producing it. Such a tax is often referred to as either a Pigouvian tax or a corrective tax.

A corrective tax raises the price to buyers and/or lowers the price received by sellers, which causes the market to reduce the equilibrium quantity sold. This generates a social gain by reducing the externality, but it also causes losses of economic welfare (reduced consumer surplus) in the same way as any tax would.³

Economists often refer to corrective taxes as “efficient” meaning that the market outcome after the imposition of a corrective tax (assuming no other market failures) is itself Pareto efficient—no reallocation of resources can make someone better off without causing harm to another. But, the move from the status quo with an uncorrected externality to the corrected market is not a Pareto *improvement*. It creates welfare gains from a reduction in the externality; it shifts welfare from the market to the government in the form of tax revenue; and it causes the aforementioned reductions in consumer surplus.

The move from the uncorrected market to the corrected one is what economists call a Kaldor-Hicks efficiency gain, or a *potential* Pareto improvement. This means that there is more welfare overall in the economy after the corrective tax is imposed: the gains from the externality reduction and the revenue transferred to the government are greater than the lost producer and consumer surplus in the market. But, unless some portion of the revenue is returned to the market participants, they will be worse off when the policy is implemented.⁴

To transform a *potential* Pareto improvement into an *actual* Pareto improvement, one needs to do something to compensate the losers, such as lump-sum transfers of wealth. When economists say that a corrective tax is efficient, they do not mean that introducing it makes everyone better off, but rather that the new outcome is itself Pareto efficient.

Standard models of social welfare maximization in economics give no special regard to the status quo allocation of welfare in the world of the uncorrected externality, and thus make no requirement that the

³ For clarity, this conversation is framed with respect to consumers only, which is consistent with a model of perfectly competitive supply with constant returns to scale (and thus no economic profit). The basic ideas apply more broadly.

⁴ Moreover, if the benefits of the externality accrue outside the market to a third party, even the full return of the revenue will not be sufficient to compensate all the market participants.

compensation of losers actually be pursued. And many externality correcting taxes have been introduced that do not aim to compensate those who lose through higher prices.

Seeking to compensate those harmed by the corrective tax does, however, accord with the notion of horizontal equity in public economics, which suggests it is desirable for people of similar income to pay similar taxes.⁵ Moreover, Sallee (2019) motivates such compensation as a way to improve the political prospects of corrective taxation.

The core concern of the empirical portion of the report is an assessment of the feasibility of targeting monetary transfers as a way of compensating those harmed by congestion pricing. To pursue that empirical analysis, it is necessary to also briefly comment on the standard way that economists relate welfare burdens, which are not observable phenomenon, to quantities of a good consumed, which are.

One's Burden from a Pigouvian Tax Depends on the Quantity of the Good or Service Consumed

The aim of a corrective tax is not to punish buyers or sellers of the good, but rather to give them the right incentives for consumption and production on the margin. Corrective taxes create a tension, however, because the burden or cost of the tax, which comes in the form of lost consumer or producer surplus, depends on the *total* quantity of the good consumed, not just the marginal unit.

Consider the following setup, which follows Sallee (2019). A group of heterogeneous consumers indexed $i=1, \dots, N$ have demand curves for a good Q that creates a negative externality with marginal external damages equal to ϕ . Q is sold by perfectly competitive sellers that have a constant marginal cost equal to p . Before the introduction of a tax, each individual demands $q_i(p)$ units of the good. For simplicity, assume that the rest of each individual's exogenous income is spent on a some generic other good that does not generate externalities.

The introduction of a corrective tax causes a direct welfare loss for person i equal to the change in their consumer surplus before and after the tax is introduced. Denote this cost as c_i , which we refer to equivalently as the *direct* (i.e., before revenue is reallocated and before accounting for the benefits from the reduced externality) cost or burden of the tax.

The economic theory tells us that, for a small tax increase, the loss in consumer surplus for person i is equal to exactly $q_i(p)$. This means that if one has a measure of the *quantity* of the good that a consumer buys before the

⁵ There is a contentious debate regarding the intellectual merit of horizontal equity (see Kaplow 1989). The invocation here of the concept is meant not as an endorsement of the concept, but rather to point out the connection to an existing intellectual tradition.

hypothesized tax change, then this quantity measure is a direct measure of the burden of the policy for a small change in the tax.

Note that this does not assume there is no behavioral response. On the contrary, it is assumed that the consumer will reduce demand for the good, but this behavioral change does not have a first-order impact on consumer welfare for small changes. Rather, the welfare gain from the introduction of a small tax is equal to the change in the aggregated demand for the good (the sum of the slopes of the demand curves at price p) times the marginal externality (ϕ). This motivates the use of initial consumption levels of a good that experiences a price increase as a proxy for welfare burdens.

For a non-marginal change in the tax, the change in consumer surplus (burden of the policy) will depend on both the baseline consumption $q_i(p)$ and on the size of the consumer's behavioral response (i.e., the slope of their demand curve), but initial consumption will still be a useful approximation.

The Challenge of Targeted Compensation

Because the goal of a corrective tax is not to punish but rather to create the right incentives to achieve market efficiency on the margin, it makes sense to compensate those who are harmed by a tax because they consume the good, so long as this is done without changing the incentive for consumption at the margin (i.e., the tax-inclusive price). The traditional approach to this in economic theory is to consider uniform transfers.⁶

It is convenient to consider using the revenue raised by the tax to fund such compensating transfers, but a greater or lesser amount may also be used. In a setting with perfect information, a social planner or government agent would have the ability to allocate transfers that precisely compensate each actor according to the harm done by a corrective tax, without needing to base the transfers on observed consumption of the good itself. This latter point is essential because if transfers are a function of the consumption of the good itself, a sophisticated market actor will recognize an underpriced good and purchase more than a socially-efficient quantity. With perfect information, transfers could instead be based on underlying primitive (exogenous) variables, like consumer taste, rather than observed (endogenous) choices.

But there are limits to the precision of targeted transfers when consumers differ in their underlying tastes or circumstances, which causes differences in their propensity to consume the good, and the policymaker has information only about choices, not about the underlying preferences. Sallee (2019) formalizes this idea by

⁶ A closely related phenomenon in many environmental regulations is the free allocation of permits to polluting firms in the context of a cap-and-trade program. As long as the permits are given away based on historical or industry-wide pollution, so that going forward firm behavior does not impact one's allocation, then these types of programs can alter the distributional burden of the pollution policy without distorting efficiency. This idea goes back to Montgomery (1972). Fowlie and Perloff (2013) provide empirical support.


deriving an empirically testable condition that must be met in order for complete compensation for the burden of corrective taxation to be possible.

Specifically, Sallee (2019) shows that for a group of exogenous characteristics that can be used to determine a transfer program, if the average absolute error of a prediction equation exceeds twice the average change in the externality, then it is not possible to generate a true Pareto improvement from an externality-correcting tax through targeted transfers, regardless of the distribution of the efficiency gains from the reduction in the externality.

The basic idea is that if there is too much unpredictable variation in the initial consumption of the good—and thus in the initial burdens created by the policy—then there is not enough “extra surplus” to be redistributed through transfers to ensure that no one is harmed on net.

Below we test this condition for the case of congestion pricing. This takes the form of regressions of the observed program cost paid by each household on the set of characteristics that could plausibly be used to target a transfer program, such as income, residential location, homeownership status, number of vehicles, number of children, age, gender, employment and student status, number of weekly commute trips, and pre-experiment driving behavior (VMT and use of toll roads prior to tolling). The more general point of the Sallee (2019) framework is the notion that a policymaker has greater control over the final distribution of net burdens from an externality-correcting policy when the burdens from the policy are more closely correlated with variables that can be used to design a set of targeted transfers. The upshot of the empirical analysis is that demographic information including income explains relatively little of the cost variation in the driving experiment, and so better targeting requires the use of information about pre-experiment driving behavior, which may be more difficult to acquire in real world settings.

Figure 9. Tolled Roads in Seattle's Transportation Choices Study.



Note: This figure shows the tolled roads in the 2004 Transportation Choices Study. Note that freeway (pink) and non-freeway roads (red) were associated with different toll charges.

