A Non-Parametric Analysis of Welfare Redistribution: The Case of Stockholm's Congestion Pricing Trial

Joel P. Franklin[*]
Department of Urban Design & Planning
University of Washington
Box 355740
Seattle, WA 98195

Submitted to Transportation Research, Part A: Policy and Practice
December 1, 2005

Abstract

This study uses non-parametric distributional comparison tools to evaluate the equity effects of a transportation policy. I consider as a case study a congestion pricing plan for Stockholm, Sweden. This case is particularly relevant because the plan has been criticized for its potentially negative equity effects, and the literature is so-far inconclusive. The welfare effects for a sample of individuals were simulated using a travel modeling system and sample enumeration, then the results examined using a series of non-parametric tools.

The results confirmed earlier findings that the magnitudes of effects are small compared to absolute welfare. The overall progressivity or regressivity is strongly related to the choice of refund scheme, with the most progressive scenario being a lump sum payment subsidy and the most regressive being a reduction in income taxes. Disregarding the refund scheme, the results suggest that the welfare effects of the plan itself are borne largely by those who initially use auto. Since these are predominantly high-income individuals, the plan itself tends to be progressive, but low-income individuals who began using the auto mode also bear a high burden.

Keywords: Congestion pricing; equity; distributional impacts; non-parametric methods
JEL Classification: D63, H23, R41, R48

1 Introduction

When public policies are expected to affect a population in complex ways, the normal evaluation measures, such as mean effect, may fail to capture all of the effects that are important to societal goals. As an example, if a policy is expected to be "regressive", benefiting mainly those who are already in the best position or costing mainly those who are already in the worst position, then we might say that its poor equity overwhelms any advantage in average benefit, depending on how heavily we weigh equity relative to that average benefit. The converse is a policy that is "progressive", playing the Robin Hood role of redistributing welfare from richer to poorer members of society.

[*]Tel: 206-616-4499, fax: 206-616-6625, e-mail: joelpf{at}u{dot}washington{dot}edu.

This study demonstrates the application of a set of non-parametric distributional comparison tools in the evaluation of a transportation policy whose effects are simulated using a complex travel modeling system. As a case study, I consider a planned trial of peak-period roadway tolls for Stockholm, Sweden. This case is particularly relevant because the congestion pricing plan has been criticized for its potential to unduly burden low-income individuals, relative to high-income individuals. Moreover, the equity effects of congestion pricing in general have been discussed extensively in literature, without consensus. While some have argued that congestion pricing burdens lower-income individuals who are least flexible to change travel patterns, others have argued that lower-income individuals are less likely to be tolled because they tend to use public transportation.

The welfare effects of Stockholm's congestion pricing plan were obtained from simulated travel demand modeling data that was produced for an earlier study by Eliasson and Mattsson (2005) on the distributional effects of Stockholm's congestion pricing plan. The results were examined using non-parametric comparison tools, such as locally-weighted regression, Relative Distributions, and Mean Relative Polarization indices, as well as some more traditional measurements such as aggregate welfare change.

This paper continues in Section 2 with some background on both the quantitative analysis of equity and of recent equity discussions in the context of congestion pricing. A systematic sequence of research questions for examining equity is presented in Section 3. In Section 4, we see in more detail the context of the case study for Stockholm's congestion pricing plan. Section 5 presents the transportation modeling that led to estimates of individual welfare, both without and with the congestion pricing plan. In Section 6, we see the results of several non-parametric methods in responding to the research questions. Some additional results pertaining to data needs for inference testing are described in Section 7. Finally, Section 8 presents some conclusions and implications for policy and future research.

2 Background

2.1 The Welfare Effects Congestion Tolls

Congestion tolls have long been proposed as a simple means of improving the economic efficiency of roadways. According to prevailing theories of roadway dynamics, as traffic demand approaches a roadway's capacity, overall traffic speeds decrease. The resulting increased travel times are externalities of individual motorists' decisions to take a particular roadway, since their decisions neglect the travel time costs on other users. Consequently, the user-equilibrium traffic distribution is distinct from the globally optimal traffic distribution. Congestion pricing systems, first proposed by Pigou (1952), charge motorists for the difference between the user cost and the full social cost of the decision to take a particular roadway, thereby bringing the user equilibrium and the global optimum into convergence. The literature (e.g. Vickrey (1968); Small (1983); Arnott et al. (1993)) largely agrees on the efficiency benefits of a theoretical first-best congestion pricing system, as well as of a more feasible second-best system that is subject to various unavoidable distortions in how the transportation system is priced. However, the arguments on the equity impacts of congestion have been varied.

"Congestion Pricing Is Regressive" Most arguments regarding the equity effects of congestion pricing have supported the notion that the effects of congestion pricing favor higher-income individuals. For example, Richardson (1974) noted that higher-income individuals value travel time greater, and therefore receive the greater benefit. Moreover, in several theoretical studies using a queue-based congestion model with commuters that are heterogeneous in their travel time and delay costs, Arnott et al. (1993, 1994, 1998) found that higher-income individuals have the greatest flexibility in departure time, allowing them to more easily absorb the toll. These theoretical works confirm earlier findings of regressive effects by Cohen (1987), Evans (1992), and Layard (1977). More recent work by Teubel (2000) and by Raux and Souche (2004) also confirm these results. Small (1983) found that the gains under several refund schemes could benefit all income levels, although once travel time benefits are accounted for (which may be equal across incomes), higher-income groups will continue to enjoy the greatest benefit, due to their higher value of time.

"Congestion Pricing Is Progressive" In contrast to the above evidence, the notion of progressive equity effects from congestion pricing have been supported by some authors, such as Foster and Shneyerov (1999), who argue that higher-income individuals are those most likely to be affected by a toll, since they choose to drive private cars more often and they tend to have a home-to-work trip originating in the suburbs and ending in the central city. This result was borne out in a study by Santos and Rojey (2004), although they note that there are a small number of individual cases where a low-income payer would experience regressive effects.

In reconciling the above views, Eliasson and Mattsson (2005) argue that the variation in evidence can be explained by differences in both the initial set of travel patterns and in the scheme selected for redistributing collected tolls. This second point is supported by the theoretical arguments set forth both by Mayeres and Proost (1997) and by Levine and Garb (2002). Still, Giuliano (1994) argues that redistribution cannot resolve all equity concerns: individual- and household-level mode and schedule flexibility is so varied, even within income groups, that the complexity of a truly "fair" refund scheme would be too great to implement.

Eliasson and Mattsson (2005) also note that the magnitudes of welfare effects, as observed in past studies, are overwhelmed by the magnitudes of incomes themselves, implying that whatever progressivity or regressivity exists should be quite small. Instead, they approach the question of equity by comparing average net effects on discrete or discretized demographic dimensions, regardless of the exact initial income level, such as gender, employment status, family type, geographic area of residence, and income categories.

2.2 Approaches to Equity Evaluation

The inconclusive findings on equity effects of congestion pricing may stem from the use of different approaches to quantifying equity. To help frame the present approach to equity, it will help to recognize a taxonomy of approaches and review their limitations.

In a broad range of social science questions, the total consumer welfare is not the only issue under consideration; also important is how consumer welfare is distributed among society's members, or how equitable is that distribution. Equity questions are often considered with respect to two types of equity, each of which can be linked to principles of justice from Rawls (1971). *Horizontal* equity refers to the distribution of welfare among individuals who are otherwise identical, drawing from Rawls's "principle of equal opportunity". *Vertical* equity refers to the distribution of welfare among individuals who are unequal in other respects, drawing from Rawls's "principle of difference". Raux and Souche (2004) identified a third form of equity in the context of roadway pricing: *spatial* equity, which implements Rawls's "principle of liberty" as providing the right of access to any location in space. In this study, I focus on vertical equity, where the individuals vary with respect to initial welfare, as indicated by household income. Consistent with Giuliano (1994), I consider as equity the distribution of a policy's effects, both positive and negative, in reference to household income as a baseline.

By far, the most common approach to quantifying equity and the one employed by the majority of literature cited in Section 2.1 is to treat the individuals as members of income categories, and to compare the average effects on each group. This approach is simple to implement and easy to understand. However, it does may not always produce a reliable response to the question of whether a result is progressive or regressive; unless the welfare results are clearly ordered by income category, the conclusion may be ambiguous. Moreover, by treating income as in categorical terms and computing group means, this approach has the potential to ignore equity effects that are present *within* income categories.

The limitations of categorical approaches suggest that a generalized approach may be more desirable. These stem largely from the "measurement of inequality" literature, which is most authoritatively summarized by Cowell (1977). Most common of these is the Lorenz Curve graphical tool and its associated summary statistic, the Gini Coefficient, which has been used in the context of transportation policy evaluation by, for example, Teubel (2000) and Bae (1997). However, the Lorenz Curve only illustrates the progressivity or regressivity for a single population under a single set of circumstances, limiting its ability to explicitly compare multiple scenarios. Moreover, these tools may hide useful information about equity effects that are confined to the tails of the distributions being compared, since Lorenz Curves approach asymptotically toward the 45-degree line at the two ends of the distribution. Lorenz Curves and Gini Coefficients have the desirable property of being *multiplicatively* scale-independent, i.e. they produce the same result regardless of whether all of the welfare values are multiplied by a constant. This allows them to be insensitive to a choice of unit of analysis. However, they are not *additively* scale invariant, so the addition or subtraction of a constant to all welfare values will change the result. Hence, if we have reason to believe that some constant portion of welfare should be disregarded (e.g. to account for a baseline cost-of-living), then the results would be sensitive to that constant.

The tools used the present study are generally non-parametric measures of distributional change, in particular Relative Distributions and Mean Relative Polarization indices as the principal methods of examining redistributional and equity effects, which have seen only limited use in applied social science research (Handcock and Morris, 1999; Bernhardt et al., 1999). The Relative Distribution, when its probability density function is plotted, illustrates the change in population density between two scenarios, indexed across quantiles of the welfare indicator (i.e. from the zeroth-percentile welfare level to the 100th-percentile welfare level). The Mean Relative Polarization index measures the degree to which one welfare distribution increases or decreases the polarization of welfares present in a previous welfare distribution. In contrast to the previous methods described, the Relative Distributions and Mean Relative Polarization indices are both additively and multiplicatively scale-invariant, meaning they show the same result regardless of any monotonic transformations of the welfare measure. Also, the Relative Distribution graphical tool and the Mean Relative Polarization index's summary measure both perform well in capturing effects within the tails of the distribution.

3 Research Questions

This study seeks to bring the non-parametric approach to equity analysis into the question of congestion pricing's effects by expanding on the equity analysis conducted by Eliasson and Mattsson (2005). The appropriate choice of non-parametric tools requires a systematic approach to answering questions of equity, moving from the most general detection of an effect and to some specific characterizations of the equity effects that the components of the plan contribute toward those effects. I organize this progression in terms of the eight research questions below.

Question 1. Starting with the most basic question, does Stockholm's congestion pricing plan have an effect of *any* kind on the distribution of welfare levels in Stockholm? To answer this, I use the Kullback-Leibler information-based measure of distributional divergence, which can detect any kind of change in the shape of a distribution.

Question 2. Assuming from Question 1 that the pricing plan has some effect, does it improve or degrade *aggregate* welfare levels across the sample? To answer this, I simply sum the effects across the sample to determine the overall net effect.

Question 3. How popular would the congestion pricing plan be, given the specific effect on each individual? More specifically, does the number of people with a net benefit outnumber the number of people with a net loss, suggesting that a referendum on the plan would pass? I include this question because its results may be informative on what one might expect to occur when the real referendum on the congestion pricing plan is held, as I discuss in Section 4.

Question 4. Assuming from Question 1 that the congestion pricing plan has a detectable effect on welfare, is there any detectable relationship between an individual's income level and the magnitude of the effect on that individual? Here is where we first gauge the potential for disparate effects of the plan at different income levels, regardless of what aggregate result was found in Question 2, although this question does not seek to characterize that relationship expressly as "regressive" or "progressive". I use a locally-weighted linear regression model to determine whether a relationship can be discerned between income level and welfare effect. By using a non-parametric regression method, we can make a relatively weak assumption about such a relationship: that it would be smooth and that along an infinitesimal income range there would exist a linear relationship. This stands in contrast to a conventional linear regression, where the parameters are globally estimated, requiring the strong assumption of a uniform linear functional form for the entire dataset.

Question 5. Assuming from Question 4 that there is a relationship between income and the congestion pricing plan's effect, *how* does the plan change the distribution of individuals' *actual* welfare levels? Here we start to see the congestion pricing plan as an agent of welfare-redistribution. I use a Relative Distribution to graphically characterize the effects of the congestion pricing plan. These graphical tools can detect equity-regressive effects that are both global and local to a particular portion of the welfare spectrum, but they can be inconclusive about overall equity effects.

Question 6. How does the congestion pricing plan redistribute individuals' welfare levels *relative to each other*? Put another way, *after accounting for aggregate effects* (as represented by the distributions' means), will the plan have the same distribution of welfare levels as that which is already exhibited in the distribution of prior incomes? This question, like Question 5, uses a Relative Distribution, but in this case we compare *mean-matched* welfare distributions that are shifted in magnitude to control for aggregate effects, allowing us to see only the redistributive effects among individuals in the sample.

Question 7. When compared to prior welfare levels, will the congestion pricing plan have either a progressive or a regressive effect on welfare levels of greater Stockholm residents? Put another way, how does the plan affect the polarization between high- and low-welfare individuals, relative to each other and regardless of overall improvements or diminishments? I use a final tool here, the Mean Relative Polarization Index, which measures the degree to which the population's welfare levels are either spread into the tails or concentrated in the center of the distribution.

Question 8. Assuming that Question 7 found a polarizing or depolarizing effect of the congestion pricing plan, what components of the plan's effects were responsible for that effect? In other words, to what extent did the tolls paid, the travel time savings, the penalty of having to switch modes, or the benefit of a refund check contribute to a polarizing (i.e. regressive) or depolarizing (i.e. progressive) effect?

The last two questions are where we can finally make generalizations about the "progressivity" or "regressivity" of the proposed plan. Yet the prior questions are still necessary if we are to identify the proper caveats on those generalizations. Clearly there is a difference between, on one hand, a policy that is uniformly beneficial or progressive in its equity effects, and the other hand, a policy whose effects are beneficial or progressive in the aggregate but exhibit undesirable features among some smaller regions of the sample. The systematic progression of questions above and the non-parametric approaches to interpreting the results allow us to make such distinctions in many situations.

4 Case Study: Stockholm's Congestion Pricing Trial

The Municipality of Stockholm and the Swedish national government are planning a seven-month, full-scale trial of congestion pricing for central Stockholm, beginning on January 3rd, 2006 and continuing through July 31st, 2006. The trial will be followed on September 17, 2006, by a popular referendum on whether to continue the toll on a permanent basis. Although the plan will certainly affect travelers from across the greater Stockholm metropolitan area, only residents of the Municipality of Stockholm will be polled.

While the plan has changed as it progressed through planning stages, the plan modeled here consists of time-differentiated toll levels during most of the day, collected along two cordon lines. Since the scheme for how to allocate the collected revenues remained under active discussion when this research began, three different refund schemes were identified for analysis, along with a fourth hypothetical scenario where no refund occurs.

4.1 Toll Collection Scheme

The plan defines two cordon lines, both shown in Figure 1: the first circumscribes the central city area, crossing several major entrances where fees would be collected, and the second bisects the central city area in an east-west direction along the Lake Mälaren/Baltic Sea waterway that separates the north side from the south side[1]. The exception to this is that no tolls would be collected on the Essinge highway, which is the westernmost crossing over Lake Mälaren and constitutes a segment of the partial-ring highway surrounding the central city. The level of tolls would vary between 15 Swedish Crowns (SEK) during peak hours (7-9 a.m. and 4-7 p.m.) and 10 SEK during midday hours (9 a.m. to 4 p.m.). No tolls would be collected during the nighttime hours of 7 p.m. to 7 a.m.

Figure 1: Central City Stockholm and the Toll Cordon Lines

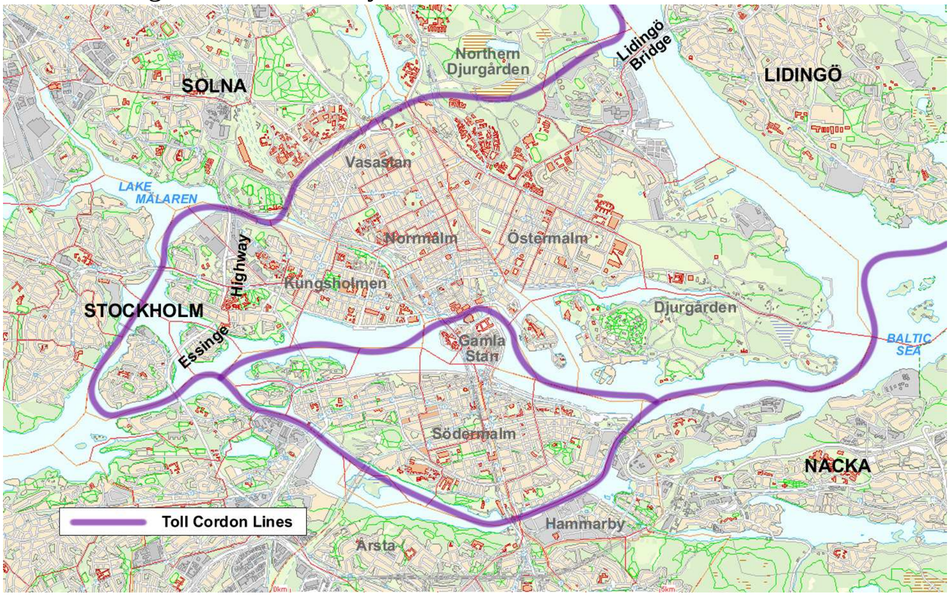

4.2 Refund Scenarios

As suggested by the review in Section 2.1, some of the disagreement in the literature regarding congestion pricing's distributional effects may be attributable to whether, and how, collected revenues are refunded to the population. To investigate this further, the effects of the congestion pricing scheme were analyzed both without any refunds and with one of three plans for allocating refunds: 1) a *Lump Sum* scenario, where all individuals receive the same refund payment; 2) a *Transit Subsidy* scenario, where the amount of benefit is proportional to the number of trips taken by transit; and 3) a *Tax Reduction* scenario, where the revenue is used to finance an across-the-board reduction in the tax rate, resulting in a benefit that is proportional to income. The four total scenarios are summarized in Table 1, along with how they were operationalized and what progressivity or regressivity we should expect to come from the refund itself.

5 Modeling Methodology

To estimate the welfare effects for a representative sample of Stockholm area residents, this study used the enumerated results of a travel demand model when run both without the congestion pricing plan, and with the plan. The modeling was carried out for an earlier study by Eliasson and Mattsson

[1]Note that since this analysis was conducted, the bisecting cordon line has been eliminated from the planned trial.

Table 1: Refund Scenarios

Scenario	Description	Implementation	Expectations
No Refund	A purely theoretical scenario where no toll revenue is refunded.	Refunds for all are set to zero.	With no refunded revenue, this should reflect the character of the congestion pricing system itself, plus a general worsening of welfare due to the disappearing revenue.
Lump Sum	Toll revenues are divided evenly among all adults living in Greater Stockholm.	Refunds are set to total revenue divided by total population.	Since payments are equal for all individuals, there should be no equity effect from the refund. Hence, this should be just as progressive or regressive as the No Refund scenario.
Public Transit Subsidy	Toll revenues are used to improve public transportation throughout Stockholm.	Refunds are allocated to individuals in proportion to the number of trips they take on public transportation.	As long as lower income individuals take transit more than high income individuals, this plan should be more progressive than the No Refund scenario.
Income Tax Reduction	Toll revenues enter into the general fund and the income tax rate for Greater Stockholmers is reduced accordingly.	Refunds are allocated to individuals in proportion to their income level.	This scenario should be more regressive than the No Refund scenario, since it will send larger refunds to those with greater incomes.

(2005); although the approach was well documented by the authors, I repeat below the aspects of the modeling methodology that are most important to the present study.

5.1 Travel Survey

The precursor study by Eliasson and Mattsson (2005) began by using a travel survey of 2,893 adult travelers residing within Greater Stockholm, where each household's trip-making patterns were recorded along with socio-economic data, such as household income. Of these, 173 individuals, or 6.0%, had a zero household income recorded in the survey data. Since the reason for this is unknown, we cannot say whether these individuals actually have an extremely low initial welfare level. It may be that they actually do have zero income; conversely, it may be that they simply report a zero income for tax reasons, and in fact receive substantial income from unconventional sources. In either case, we cannot consider income to be a reliable indicator of welfare for these individuals, so those individuals were excluded from the analyses that follow. Consequently, the valid sample had an overall median household income per adult of 174,000 Swedish Crowns (SEK) per year and a maximum of 2,732,000 SEK per year. At November, 2005 exchange rates, these translate to approximately US$21,400 and US$335,400 per year, respectively.

5.2 Modeling Travel Behavior Responses

Travel data in both the *No-Toll* and *With-Toll* scenarios were generated using the SAMPERS travel demand model for Greater Stockholm (Beser and Algers, 2001). The model results were used together with a separate system of logit models for decisions on travel mode and destination choice, estimated separately for men and for women, and for work trips and non-work trips, to compute travel times and travel costs for travel between all of the Stockholm area's transportation analysis zones (TAZs).

Sample enumeration was used to produce, for each observed trip in the household survey, a set of estimated choice probabilities, travel costs, and travel times for every pairwise combination among five travel modes and 20 destinations (one destination as observed in the survey, and 19 others randomly sampled from the full set of TAZs). From the travel time and cost information, we can estimate a generalized cost of each mode-destination pair, where the travel time component is converted to monetary units using a categorical marginal-value-of-time for each trip purpose and each gender that is obtained from the estimated SAMPERS travel model. The sample enumeration procedure is explained in more detail by Eliasson and Mattsson (2005).

5.3 Representing Differential Welfare Effects

A "Rule of Half" consumer welfare computation was used to estimate the cost effects of the congestion pricing plan on each sampled individual. This method reduces the data requirements from knowing the entire shape of the demand curve, to knowing only its locations for the two scenarios being compared. The procedure employs an assumption that travel demand changes approximately linearly at small changes in prices and travel times. As such, the welfare change can be estimated by computing the trapezoidal area under the demand curve, which is we consider to be a function of generalized cost. This trapezoidal area is shown in Figure 2 as the hatched area extending to the left of the demand function between points A and B.

Figure 2: The Rule-of-Half

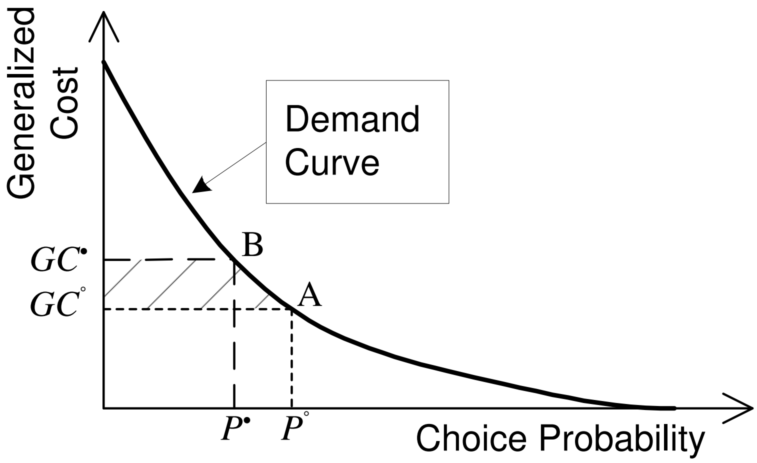

When applied to an individual's travel choice, the demand curve in Figure 2 represents the probability of choosing a particular mode/destination pair, given that a trip will take place, rather than the total quantity of trips demanded. Hence, the change from point A to B represents the change in that probability, P, that would occur if the generalized costs, GC, of a particular mode/destination pair rose between the *No-Toll* scenario, "∘", and the *With-Toll* scenario, "•". We would have similar curves for each other mode/destination pair that spans the range of mode alternatives and the 20 destinations (one observed, 19 sampled) for each individual. Since we take the trip's occurrence as given, these probability curves collectively sum to one, given any vector of generalized costs for all alternatives. Thus, knowing the probability functions and the generalized costs, we can compute the expected value of the change in generalized cost that would be incurred for an individual.

Formally, for each individual i, the Rule-of-Half sums the trapezoidal areas for each mode j's and destination k's demand curve as follows:

$$\mathbb{E}\left[\Delta w_i^{\bullet\circ}\right] = \frac{1}{2}\sum_{j,k}\left(P_{ijk}^{\circ} + P_{ijk}^{\bullet}\right)\left[\left(c_{ijk}^{\circ} + \theta_i t_{ijk}^{\circ}\right) - \left(c_{ijk}^{\bullet} + \theta_i t_{ijk}^{\bullet}\right)\right], \tag{1}$$

where $\Delta w_i^{\bullet\circ}$ is individual i's welfare change from the *No-Toll* scenario to the *With-Toll* scenario and θ_i is the marginal cost of travel time for i, given the trip type and i's gender, and obtained from the travel model. P_{ijk}° and $P_{ijk}^{\bullet}$ are i's probabilities of taking mode j to destination k, c_{ijk}° and $c_{ijk}^{\bullet}$ are i's monetary costs of that choice, and t_{ijk}° and $t_{ijk}^{\bullet}$ are i's travel times of that choice, all for the *No-Toll* and *With-Toll* scenarios, respectively.

Now note that in the congestion pricing plan, the only changes that occur are to the auto mode: the cost increases in the *With-Toll* scenario by a toll of $\tau_{ak}^{\bullet}$, and the travel time may decrease as fewer travelers choose the auto mode. We can therefore restate (1) with only the auto mode. Moreover, following the formulation by Eliasson and Mattsson (2005), we can also state the welfare change as a sum of three terms:

$$\mathbb{E}\left[\Delta w_i^{\bullet\circ}\right] = \underbrace{\frac{1}{2}\theta_i \sum_k \left(P_{\text{i}ak}^{\circ} + P_{\text{i}ak}^{\bullet}\right)\left(t_{ak}^{\circ} - t_{ak}^{\bullet}\right)}_{\textbf{I}} \underbrace{-^{\bullet}\sum_k P_{\text{i}ak}^{\bullet}\tau_{ak}}_{\textbf{II}} \underbrace{-\frac{1}{2}^{\bullet}\sum_k \left(P_{\text{i}ak}^{\circ} - P_{\text{i}ak}^{\bullet}\right)\tau_{ak}}_{\textbf{III}}, \tag{2}$$

where the three terms above represent three distinct effects of the congestion pricing plan:

I. *Value of Travel Time Savings:* Those who remain in the auto mode will likely experience reduced congestion, saving travel time, after the toll is introduced.

II. *Total Paid Tolls:* All travelers who take auto into or out of the central city pay the toll.

III. *Change in Travel Pattern:* The remaining term represents the penalty experienced by those who are forced to adjust travel mode or destination to avoid the toll.

Using the decomposition in (2), we can parse out these three effects, when we identify the sources of any polarizing or depolarizing effects in Question 8.

5.4 Representing Absolute Welfare

The rule-of-half computation gives us only the *change* in welfare from the *No-Toll* scenario to the *With-Toll* scenario. In fact, we would like to know the *absolute* welfare level of each individual in each scenario. This is important in the present analysis because many of the questions stated in Section 3 consider not only the incremental changes in *individual* welfares, where absolute welfare can be safely disregarded, but also changes in the *distribution* of individual welfare, where we cannot disregard the absolute welfare levels of individuals.

Also, the collected toll revenue is to be returned to the population in some manner, representing a fourth component of welfare change. I assume here that the costs of operating the toll system and the refund system are negligible, so that the entire amount of collected revenue is available to refund. Using the one of the four refund schemes in Section 4.2, we can compute a refund payment $r_i^\bullet$ for each individual i, potentially based on that individual's income or transit usage.

For the *No-Toll* scenario, we obtain the absolute welfare simply from reported income: $w_i^\circ = y_i$, where y_i is individual i's income. For the *With-Toll* scenario, we sum the reported income, the monetized *change* in welfare that results from the rule-of-half computation, and the refund payment, giving us: $w_i^\bullet = y_i + \Delta w_i^{\bullet\circ} + r_i^\bullet$.

6 Distributional Analysis

After modeling all scenarios and computing the welfare levels of each individual, the collective results across all individuals constitute the welfare distribution vectors $\mathbf{w}^\circ$ and $\mathbf{w}^\bullet$ for the *No-Toll* and *With-Toll* scenarios, respectively. In fact, *With-Toll* comprises the four refund scenarios described in Table 1. Some summary results for the changes in welfare from the *No-Toll* scenario to each of the four *With-Toll* refund scenarios are shown in Figure 3.

Figure 3: Summary of Distributional Results

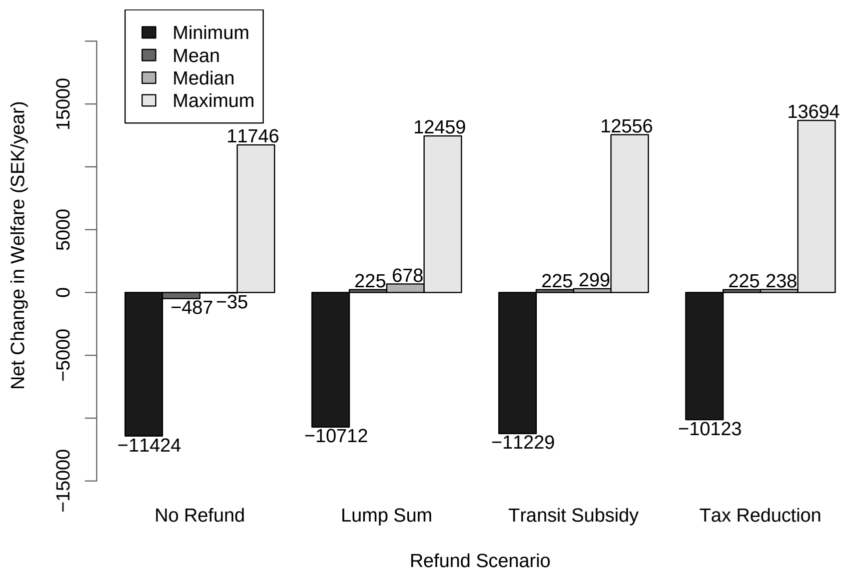

Let us leave aside until Question 2 the mean changes in welfare from Figure 3. Looking instead at the minimums, maximums, and medians, we see that the individual changes in welfare range from positive to negative multiple thousands of Swedish Crowns per year. In the *No Refund* scenario, these statistics suggest that most people do worse than any of the other scenarios, and, in particular, the median person has a net loss. This makes sense, since in this scenario, none of the collected toll revenue is refunded to the population in any way. When we do refund the tolls in some way, the median person does best in the *Lump Scenario*, but the worst in the *Tax Reduction* scenario.

To explore these results further, a variety of methods were used to interpret the welfare vectors $\mathbf{w}^{\circ}$ and $\mathbf{w}^{\bullet}$ with respect to the eight research questions. In many of these methods, we treat the welfare distributions $\mathbf{w}^{\circ}$ and $\mathbf{w}^{\bullet}$ as if they were collections of random draws from the two random variables, W° and $W^{\bullet}$, for the *No-Toll* and *With-Toll* scenarios, respectively. This metaphor allows us to define for each welfare distribution a corresponding Cumulative Distribution Function (CDF), F° and $F^{\bullet}$, and Probability Density Function (PDF), f° and $f^{\bullet}$.

6.1 Question 1: General Effect

We start with the most basic question of whether or not the congestion pricing plan changes the distribution of welfare levels in any of the four refund scenarios. To test this, we use the Kullback-Leibler Divergence.

6.1.1 Kullback-Leibler Divergence Method

The *Kullback-Leibler Divergence* (KL) is an information theoretic measure of whether two distributions are equivalent (Kullback and Leibler, 1951). It can be expressed in terms of the two distributions' PDFs:

$$\mathrm{KL}\left(f^{\bullet}; f^{\circ}\right) = \int_{-\infty}^{\infty} \log\left(\frac{f^{\bullet}(w)}{f^{\circ}(w)}\right) dF^{\bullet}(w), \tag{3}$$

where f° and $f^{\bullet}$ are the PDFs for the *No-Toll* and *With-Toll* welfare distributions, respectively, and $F^{\bullet}$ is the CDF of the *With-Toll* welfare distribution. The KL divergence can be interpreted as measuring the informational distance between two distributions; hence, if the two distributions being compared were identical, the KL would be zero, and positive values indicate increasing divergence (the KL cannot take on negative values). The KL divergence is also known as the "relative entropy" between two distributions because, as we will see later, the KL divergence is equal to the entropy of the *Relative Distribution* between two distributions.

6.1.2 Kullback-Leibler Divergence Results

The results of the KL computations for each of the four *With-Toll* refund scenarios, when compared to the *No-Toll* scenario, are shown in Figure 4. These results suggest that the *Lump Sum* scenario departs most significantly from the *No-Toll* scenario, followed by the *No Refund* and *Transit Subsidy* scenarios. The *Tax Reduction* refund scenario produces a welfare distribution most similar to the *No-Toll* scenario. From these results we can say that none of the *With-Toll* scenarios have an identical welfare distribution to the *No-Toll* scenario.

6.2 Question 2: Aggregate Effect

Second, we examine the aggregate effect of the congestion pricing plan under each refund scenario, using the mean change in welfare across all individuals as an indicator. The means are reported in Figure 3 in the third row of data. The results show exactly what we should expect: that when the toll revenues are not refunded, there is an average net loss (−487 SEK/year/person) to all individuals, and that in all other scenarios, there is a net gain of +225 SEK/year/person, i.e. the same mean

Figure 4: Kullback-Leibler Divergence Results

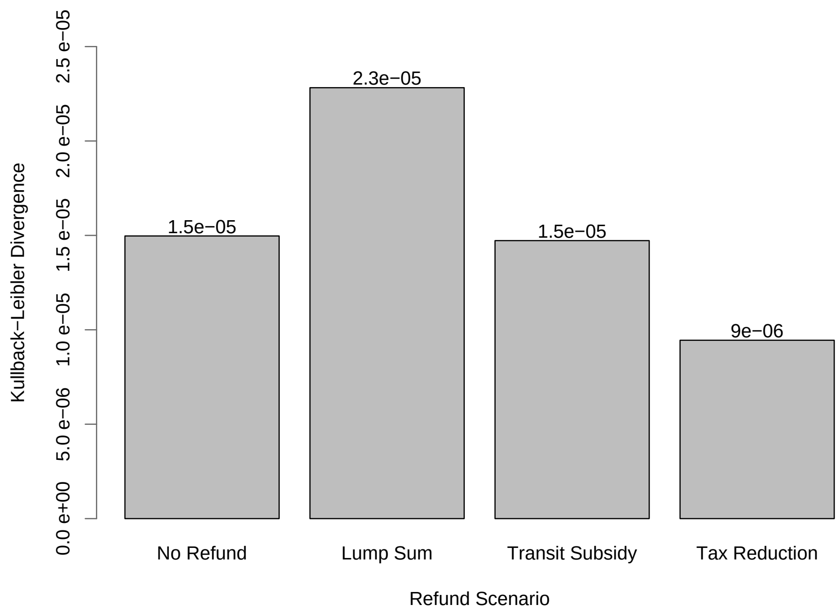

regardless of how the revenues are refunded. These means must be the same since the only component of welfare change that varies across these scenarios is the refund scheme, and in all refund scenarios except the *No Refund* scenario, there are both the same total amount of revenue to spend and the same number of individuals receive it. From a utilitarian welfare perspective, these results mean that all scenarios except the *No Refund* scenario are equally desirable.

6.3 Question 3: Popularity

A second approach to measuring the desirability of the congestion pricing plan is to assess its popularity in a hypothetical referendum, in this case by estimating what proportions of the population would have a net gain in welfare, and thus vote for the plan, and what proportion would have a net loss, voting against the plan. The conclusion of this approach will favor the majority group, regardless of what the aggregate, or average, net change in welfare was in Question 2.

The dark gray and light gray bars in Figure 5 show the proportions of the population with a net-loss and with a net-gain following the congestion pricing plan, respectively, along with a third set of medium-gray bars showing the net popularity of the plan (i.e. winners minus losers), in percentage-points. The results for the *No Refund* scenario show that there are fewer winners (61%) than losers (39%). Since revenue is disappearing in this scenario, this result makes sense; it is only when collected revenue is redistributed in some way that the winners outnumber the losers. The *Lump Sum* scenario makes most Stockholmers winners (56%) instead of losers (16%). The *Transit Subsidy* and *Tax Reduction* scenarios are slightly less popular, but still have a net positive popularity, suggesting that as long as revenues are redistributed in some manner, then the plan would pass a referendum. Note that this analysis indicates the popularity of the plan among *all* residents of Greater Stockholm. In contrast, the referendum that is to be conducted after the conclusion of the trial will only poll residents of the Municipality of Stockholm.

Figure 5: Summary of Popularity Results

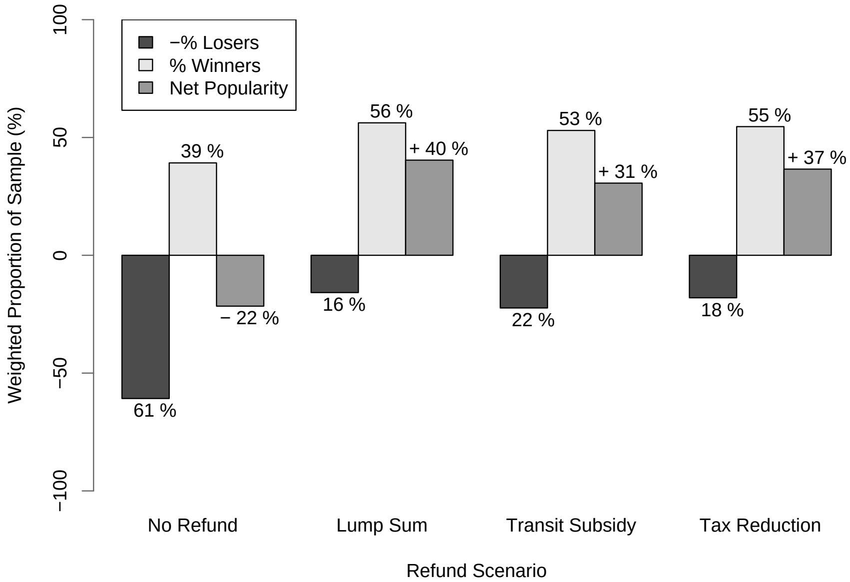

6.4 Question 4: Relationship with Income

We found in Question 1 that each of the *With-Toll* scenarios diverges measurably from the *No-Toll* scenario. In the interest of studying equity, our next natural question would be whether there is any kind of relationship between income level and the kind of net welfare effect the plan has. We can examine this question by making relatively weak assumptions about what relationship might exist: we assume only that the relationship between income and welfare-effect, if it exists, would follow a continuous function, but we do not assume a particular functional form. To test whether this is the case, we can estimate locally-weighted regression functions of welfare-change by income.

6.4.1 Locally-Weighted Regression Method

A particular method of implementing this approach, known as the LOESS method, was introduced by Cleveland (1979) and further refined by Cleveland and Devlin (1988). In this method, we estimate for every point along the income range a continuously varying low-order polynomial regression function, where at each point in the range, the regression parameters are estimated using a weighted subset of the observed welfare-change data. Using a moving Kernel weight function, the data are weighted by their distance from the particular point on the curve where we are estimating the regression function, with farther data being weighted close to zero, and data with income levels close to the point of analysis being strongly weighted. In this case, I use a tri-cube weighting function:

$$K(d)\begin{cases} \left(1 - |d|^3\right)^3, & \text{for } |d| < 1; \\ 0 & \text{for } |d| \geq 1, \end{cases} \tag{4}$$

where $K(\cdot)$ is the Kernel weight function and d is the standardized distance along the income scale between the point where the regression function is being evaluated and the point where a data point is located.

6.4.2 Locally-Weighted Regression Results

The results of the locally-weighted regression estimations for each of the four refund scenarios are shown in Figure 6. Each line in the figure represents the locally-weighted average change in welfare for one of the four refund scenarios, plotted across the range of income quantiles. The most important finding here is that all four of the refund scenarios exhibit an average welfare change that varies across income levels, indicating that equitable effects of the plan may be an warrant further investigation.

Specifically, in both the *No Refund* and *Lump Sum* scenarios, the upper and lower ends of the income distribution fare better than the middle of the distribution, suggesting a mixture of progressivity (low-income individuals doing better than mid-income individuals) and regressivity (mid-income individuals doing worse than high-income individuals). The *Transit Subsidy* scenario is similar, except that a much smaller high-income range exhibits a net gain, and a larger low-income range gains, suggesting that, on balance, this scenario is probably progressive in equity effects. A clearly opposite pattern appears in the *Tax Reduction* scenario: benefits increase monotonically with income level, suggesting a strongly regressive pattern.

Figure 6: Locally-Weighted Regression of Welfare Change by Income

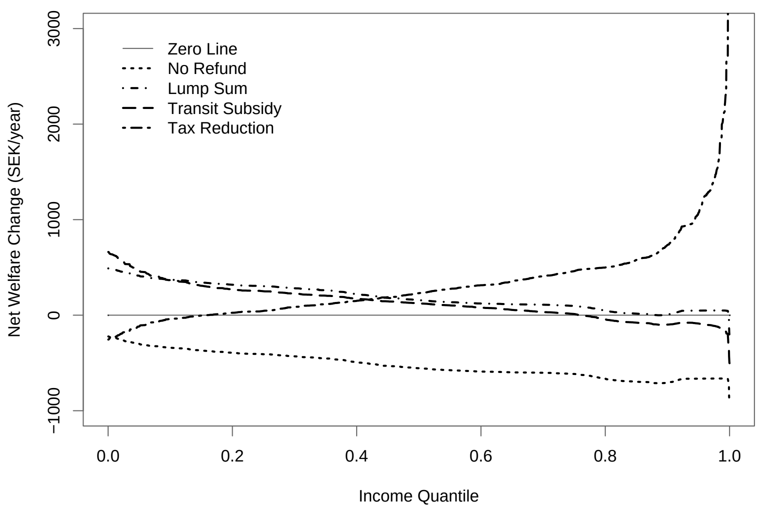

6.5 Question 5: Redistribution of Welfare

We now have some evidence that there is a relationship between income level and the change in welfare that an individual would be expected to experience after the congestion pricing plan is implemented. We might say that by looking at the *distribution of welfare changes*, with respect to income, we found a pattern. We next begin to investigate the character of that relationship further, starting with a graphical tool that helps us recognize the *changes in the distribution of welfare*. Our interest is now in what welfare levels would be achieved by what proportions of the population, if the congestion pricing plan were actually implemented, relative to the previous welfare levels. We can visualize these changes by estimating a Relative Distribution.

6.5.1 Relative Distribution Method

The *Relative Distribution* is treated thoroughly by Handcock and Morris (1999) for application to the social sciences, but has roots as early as the closely related *comparison distribution* from Parzen (1979, 1992). Recall our two random variables, W° and $W^{\bullet}$. The first, W°, represents a *No-Toll* distribution, and acts as a basis for the other comparisons. The second, $W^{\bullet}$, represents the *With-Toll* distribution. Respectively, these two random variables have cumulative distribution functions (CDFs) given by F° and $F^{\bullet}$, and probability density functions (PDFs) given by f° and $f^{\bullet}$.

We can define the Relative Distribution as a random variable generated by a transformation of the *With-Toll* random variable $W^{\bullet}$ through the CDF of the *No-Toll* random variable:

$$R^{\bullet\circ} = F^{\circ}\left(W^{\bullet}\right), \tag{5}$$

where $R^{\bullet\circ}$ is the Relative Distribution's random variable, F° is the CDF of the *No-Toll* distribution, and $W^{\bullet}$ is the *With-Toll* distribution. The Relative Distribution can then be interpreted as producing quantile levels that data drawn from the *With-Toll* distribution would have if they had come from the *No-Toll* distribution. Hence, with R's output being quantiles, they can hold values from 0 to 1.

The Relative Distribution's PDF is useful in a variety of circumstances, not least the graphical depiction of a Relative Distribution's features. The PDF is given by:

$$g^{\bullet\circ}(r) = \frac{f^{\bullet}\left(Q^{\circ}(r)\right)}{f^{\circ}\left(Q^{\circ}(r)\right)}, \tag{6}$$

where $g^{\bullet\circ}(r)$ is the Relative Distribution's PDF, $f^{\bullet}(r)$ is the *With-Toll* distribution's PDF, and $f^{\circ}(r)$ is the *No-Toll* distribution's PDF. The Relative Distribution is also useful in that it links directly to the Kullback-Leibler divergence from Question 1; the KL divergence is simply the entropy of the Relative Distribution's PDF:

$$\mathrm{KL}\left(f^{\bullet}; f^{\circ}\right) = \mathrm{Entropy}(g^{\bullet\circ}) = \int_0^1 \log\left(g^{\bullet\circ}(r)\right) g^{\bullet\circ}(r)\, dr. \tag{7}$$

The PDF of the Relative Distribution will be the focus of our results. As suggested by (6), the shape of g can be interpreted as the density ratio between the probability density functions for the two scenarios being compared. The ratio is evaluated as a function of the quantiles, r, of the *No-Toll* distribution. It spans the range of quantiles of the *No-Toll* distribution from zero to one, and it can take on values ranging from zero to, asymptotically, $+\infty$. Values of zero occur where the *With-Toll* distribution has zero density, in other words at welfare levels that no one in the sample achieves in the *With-Toll* scenario. Very high positive values arise at welfare levels for which the *With-Toll* distribution has positive density but the *No-Toll* distribution has nearly zero density. Finally, the relative density takes on a value of one for welfare levels at which the probability densities of the two distributions are equal. For two probability distributions that are completely identical, the Relative Distribution takes on a value of one for all welfare quantiles.

Another feature of the Relative Distribution's PDF that makes it useful for interpretation is the fact that, because it is a PDF of a random variable, we know that the area underneath the curve will always equal one. This allows us to interpret the shape of the curve as a true *redistribution* of the population density across the range of welfare. Hence, if we see the curve dipping below the one line at welfare level w_1, representing a reduction in density, then we know there must be a complementary increase in density at some other welfare level w_3. Such a situation might suggest that some individuals move directly from welfare level w_1 to welfare level w_3, but the interpretation should actually be broader than this: there may instead be some individuals moving from welfare w_1 to an intermediate welfare w_2, and an equal number moving from w_2 to w_3. The Relative Distribution method treats these two situations the same; it does not matter whether we match specific individuals between the two states being compared, only that the overall distribution of anonymous individuals changes density in certain ways.

6.5.2 Relative Distribution Results

Relative Distributions were estimated from the simulated sample data using a weighted local likelihood density estimation with kernel weighting of the observed data (see Handcock and Janssen (2002) for more details on this estimation method). To begin examining the redistributive effects of the congestion pricing plan on actual welfare levels, we graph the estimated Relative Distribution's PDFs for each scenario, as shown in Figure 7. These curves show us how each scenario's probability density compares to the probability density without congestion pricing—which parts of the original distribution's range has been made more dense with individuals, and which parts less dense.

Figure 7: Relative Distributions of "With-Toll" vs. "No-Toll" Welfare

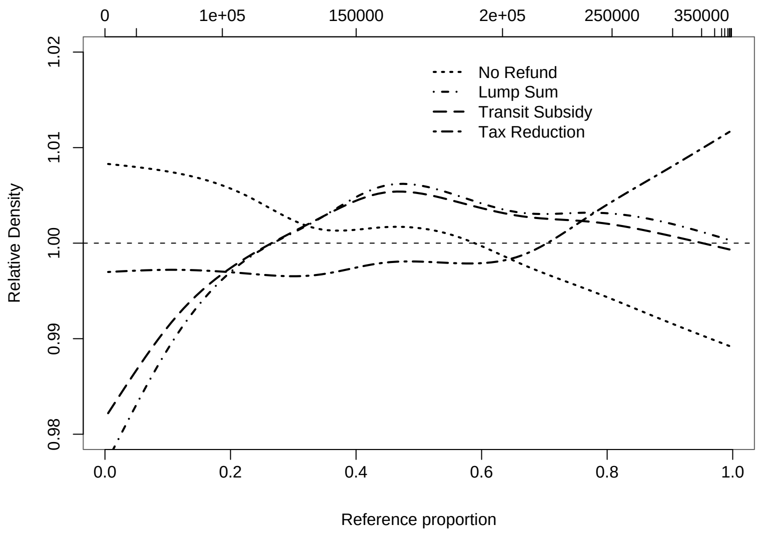

Note that the *x*-axis in Figure 7 represents quantiles of welfare levels in the *No-Toll* scenario, ranging from 0 (i.e. the zeroth percentile) to 1.0 (i.e. the 100th percentile). A second horizontal scale is given by the irregularly spaced tick marks along the top edge of the plot, representing the absolute welfare levels (in SEK/year) that correspond to the quantile scale. Next, note that the *y*-axis, which represents the ratio of probability densities, includes only a very narrow range of values, all in close vicinity to one. This is because the welfare distributions are nearly equal across the four scenarios. However, even at this narrow range the results have interesting properties.

The *No Refund* scenario has a steadily decreasing function over nearly its entire length, indicating that the overall distribution of total welfare after introducing the congestion pricing plan is more concentrated at the low end and less concentrated at the high end, when compared to the total welfare values without congestion pricing. This result corresponds to our expectation that, because tolls are being collected but not returned in any way to the population, welfare values decrease across all income levels.

The *Lump Sum* and *Transit Subsidy* scenarios both show concave functions with a slightly increasing tilt. This suggests a combination of an increase in overall welfare and a concentration of welfare, or a progressive equity effect. The *Tax Reduction* scenario, on the other hand, exhibits a slightly upward-tilting convex curve, which suggests an overall increase in welfare but also a slightly regressive equity effect.

6.6 Question 6: Mean-Adjusted Redistribution of Welfare

The interpretations from Figure 7 are somewhat ambiguous in that the shapes of the estimated curves are being used to comment on both the aggregate effects and the equity effects. To focus on the equity results of each refund scenario, it would be useful to control for overall changes in welfare magnitude across the entire sample. In this way, we can focus on redistributions of welfare around some central tendency, since in both circumstances being compared, the same total amount of welfare would then exist; it would only be distributed differently among the individuals. To do this, we produce a Relative Distribution that compares the *With-Toll* distribution with a new *mean-adjusted No-Toll* distribution, which is identical to the *No-Toll* distribution in *shape* but has its values adjusted in magnitude such that its mean matches the *With-Toll* distribution.

6.6.1 Mean-Adjusted Relative Distribution Method

First, we denote an intermediate random variable, $W^{\hat{\circ}}$, which produces values identical to the *No-Toll* distribution, except that the values are all shifted upward or downward by some constant value such that the mean of its results matches the *With-Toll* distribution's mean. The random variable, its CDF, and its PDF, respectively, are given by:

$$W^{\hat{\circ}} = W^{\circ} + \rho^{\bullet\circ}, \tag{8}$$

$$F^{\hat{\circ}}(w) = F^{\circ}(w - \rho^{\bullet\circ}), \text{ and} \tag{9}$$

$$f^{\hat{\circ}}(w) = f^{\circ}(w - \rho^{\bullet\circ}), \tag{10}$$

where $\rho^{\bullet\circ} = \overline{w}^{\bullet} - \overline{w}^{\circ}$ is the estimated difference in means between the two distributions.

Using $W^{\bullet}$, W°, and $W^{\hat{\circ}}$, we can define two intermediate Relative Distributions, the first representing differences in location and the second representing residual changes in shape:

$$R^{\hat{\circ}\circ} = F^{\circ}\left(W^{\hat{\circ}}\right) = F^{\circ}(W^{\circ} + \rho^{\bullet\circ}), \text{ and} \tag{11}$$

$$R^{\bullet\hat{\circ}} = F^{\hat{\circ}}(W^{\bullet}) = F^{\circ}(W^{\bullet} - \rho^{\bullet\circ}). \tag{12}$$

The first of these, $R^{\hat{\circ}\circ}$, represents the Relative Distribution between the original *No-Toll* distribution and the mean-adjusted *No-Toll* distribution. In other words, it captures the changes in density due only to an overall shift in magnitude. The second of these, $R^{\bullet\hat{\circ}}$, represents the remaining differences in shape; this is the *mean-adjusted Relative Distribution.*

The three Relative Distributions we have now defined ($R^{\bullet\circ}$, $R^{\hat{\circ}\circ}$, and $R^{\bullet\hat{\circ}}$) can be related to each other in terms of probability density functions, as defined in (6):

$$g^{\bullet\circ}(r) = g^{\hat{\circ}\circ}(r) \times g^{\bullet\hat{\circ}}(p) \tag{13a}$$

$$\Rightarrow \frac{f^{\bullet}(w_r)}{f^{\circ}(w_r)} = \frac{f^{\hat{\circ}}(w_r)}{f^{\circ}(w_r)} \times \frac{f^{\bullet}(w_r)}{f^{\hat{\circ}}(w_r)}, \tag{13b}$$

where r is the quantile in the *No-Toll* distribution for a given value w_r, and p is the quantile in the location-adjusted *No-Toll* distribution for the same value w_r.

6.6.2 Mean-Adjusted Relative Distribution Results

The mean-adjusted Relative Distributions for each refund scenario are shown in Figure 8. The first two scenarios, *No Refund* and *Lump Sum*, show exactly the same results: a generally concave curve, suggesting a progressive equity effect. It is in fact appropriate that these two scenarios would be identical to each other, since the only difference between them is that the *Lump Sum* scenario pays a fixed amount to each individual. Since these results are adjusted by such fixed effects, that payment's effect disappears. Another feature of these curves is that the upper tail of the distributions flatten near the one-line, while the lower tail reaches far below the one-line. This finding indicates that the progressivity of the congestion pricing plan mostly affects the lower and middle portions, moving more people into the center of the distribution, but that the upper end of the welfare range remains relatively unchanged.

Figure 8: Mean-Matched Relative Distributions of “With-Toll” vs. “No-Toll” Welfare

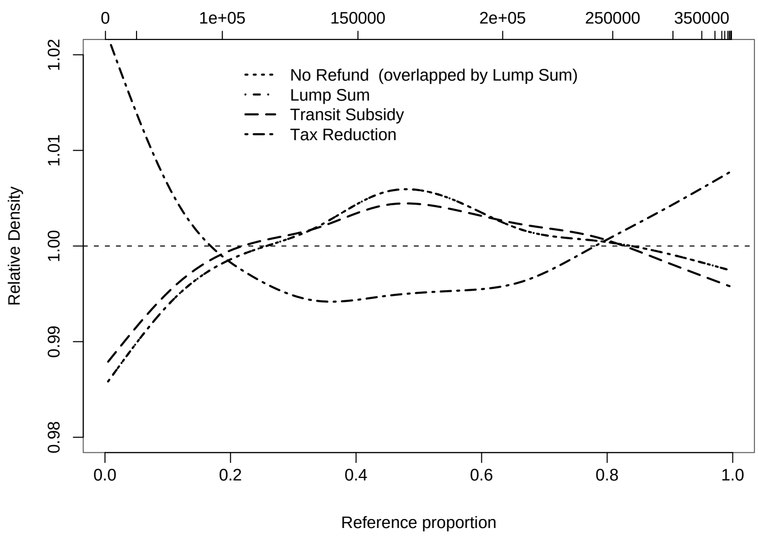

The *Transit Subsidy* also shows a concave, equity-progressive result, but in this case, both the upper and lower tails exhibit decreasing density, suggesting that a greater number of high-income individuals experience a reduction in welfare, moving them toward the center of the distribution.

Finally, the *Tax Reduction* scenario is convex, suggesting a regressive effect. These findings are consistent with the expectations set forth in Table 1. The result is particularly alarming because of the height of the lower tail: for the most part, density from the middle of the mean-adjusted welfare scale is moving to the low end of the scale, while a relatively small amount is moving to the upper end of the scale.

6.7 Question 7: Polarization Effects

While the graphical estimates of Relative Distributions and mean-adjusted Relative Distributions are instructive on a policy's redistribution patterns across the welfare scale, it remains useful to summarize regressive or progressive tendencies using a single metric. To do this, we use indices of relative polarization, which are described in Handcock and Morris (1999). "Polarization" in this case means that we are interested in the degree to which the congestion pricing plan moves the density of individuals either from the margins of the distribution into the center, or from the center out to the tails. A *polarizing* effect would cause some greater disparity in welfare levels, and would therefore be interpreted as *regressive*, while a *depolarizing* effect, where individuals move toward the center, would equalize welfares and be be interpreted as *progressive*. The Mean Relative Polarization indices are computed based on the mean-adjusted Relative Distribution. Consequently, they indicate the relative polarization that results only from changes in the distribution's shape, regardless of any aggregate changes in welfare.

6.7.1 Mean Relative Polarization Method

We start with the Mean Relative Polarization (MRP) Index of a distribution $W^{\bullet}$, relative to a baseline distribution W°, which indicates the polarization across the entire distribution. It can be expressed either using expected values of the distribution's distances from one-half, or using an integral area calculation for the distribution's PDF:

$$\mathrm{MRP}\left(F^{\bullet};F^{\circ}\right)=4\mathbb{E}\left[\left|R^{\bullet\hat{\circ}}-\tfrac{1}{2}\right|\right]-1=4\int_{0}^{1}\left|r-\tfrac{1}{2}\right|g^{\bullet\hat{\circ}}\left(r\right)dr-1, \tag{14}$$

where $\mathbb{E}\left[\cdot\right]$ is the expected value function, $R^{\bullet\hat{\circ}}=F^{\hat{\circ}}\left(W^{\bullet}\right)=F^{\circ}\left(W^{\bullet}-\rho^{\bullet\circ}\right)$ is the mean-adjusted Relative Distribution of the *With-Toll* distribution, $W^{\bullet}$, relative to the *No-Toll* distribution, W°, and $g^{\bullet\hat{\circ}}\left(r\right)$ is the PDF of the mean-adjusted Relative Distribution.

As suggested by the integral form in (14), the MRP is essentially a weighted measure of the area underneath the Relative Distribution's PDF curve. The measurement is weighted such that areas near the center of the curve are weighted the least, and areas near the tails are weighted the greatest. As a result, the MRP will be high when the tails are high but the center is low, indicating a probability density that is dispersing from the mean in the *With-Toll* distribution, when compared to the *No-Toll* distribution. The MRP will be low when the tails are low but the center is high, indicating convergence of values in the population, or increasing equity. An MRP value between 0 and 1 indicates greater polarization, while a value between -1 and 0 indicates less polarization.

6.7.2 Mean Relative Polarization Results

The MRPs were estimated using a set of quasi-relative data generated for a comparison between the *No-Toll* distribution and the estimated-mean-adjusted *With-Toll* distribution (for more details on this estimation method, please see Handcock and Janssen (2002)). The results are shown as the large black dots in Figure 9. These results are based on the same comparisons shown in Figure 8.

As with the mean-matched Relative Distribution curve results, the *No Refund* and *Lump Sum* scenarios also have identical MRPs; again, we expect this result, since these scenarios are only different with respect to a fixed refund payment, which is absorbed in the mean-matching process. The MRPs also agree with respect to the shape of the mean-matched curves. The overall index's negative value agrees with the observation that the mean-matched curves do not exhibit strong concavity of any kind.

The third scenario, *Transit Subsidy*, has a negative Mean Relative Polarization index relative to the first two scenarios, suggesting that this scenario is progressive but not to the same extent. The last scenario, *Tax Reduction*, has a positive Mean Polarization Index, suggesting increased polarization of high and low welfare individuals. This is exactly what we expected of a refund plan that benefits individuals in proportion to their incomes, thereby intensifying the existing dispersions of income.

Figure 9: Mean Relative Polarization Indices & Components

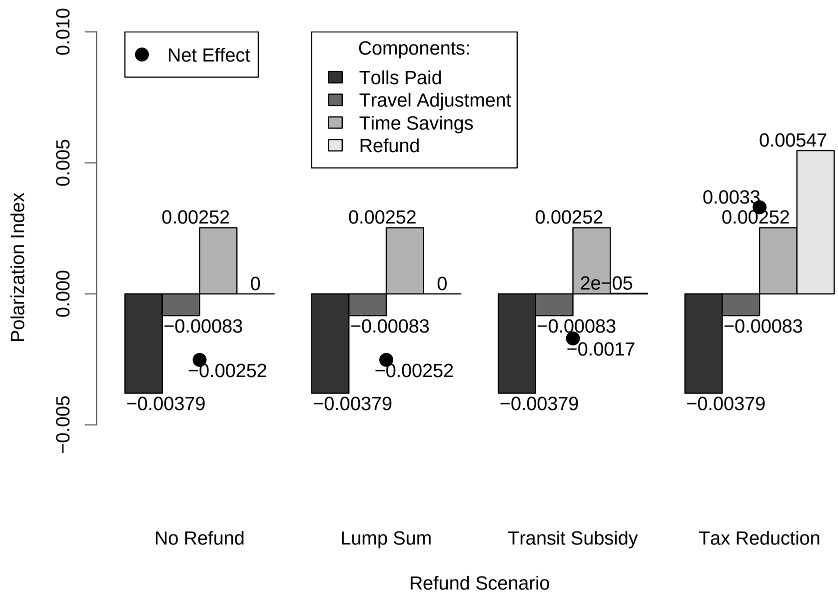

6.8 Question 8: Decomposition of Polarization Effects

Finally, we examine the specific contributions to polarization or depolarization of the four components of the congestion pricing plan's welfare effects from (2): **(I)** Value of Travel Time Savings, **(II)** Total Paid Tolls, **(III)** Change in Travel Pattern, and **(IV)** Toll Revenue Refund Payment. To measure these effects individually, the components were compounded in the following order: 1) Cost of Tolls Paid, 2) Cost of Travel Adjustments, 3) Travel Time Savings, and 4) Benefit from a Refund Payment. MRP indices were then computed for each of these four welfare effect components, relative to the net welfare after the previous welfare component was applied. The results are shown as the bars in Figure 9. Note that the MRP is *not* additive by welfare effect components, so the MRPs for each component will not necessarily sum to the MRP of the overall plan. However, the MRPs all use the same scale, so they can be compared.

First, notice that the first three bars are the same for all four refund scenarios. This is because the only welfare component that changes across scenarios is the fourth component, the size of the refund for each individual. Among the first three components, it appears that the tolls paid and the travel adjustments have depolarizing effects on the welfare distribution, while travel time savings have a polarizing effect on welfare. This can be explained by the prevailing initial trend in mode choice for Stockholm: the auto mode is more strongly favored at higher income levels than at lower income levels. Since each of the first three welfare effects are borne by those who initially used auto, they also tend to be borne most strongly by the upper income levels. Hence, any negative effects (tolls paid and travel adjustments) will be depolarizing, and any positive effects (travel time savings) will be polarizing.

The polarization effect of the refund payment varies by method of allocating the refunds. In the first two scenarios, there is no polarization effect. Of course, in the *No Refund* scenario, there should be no effect since there is no refund, but the same result appears in the *Lump Sum* scenario. This is because the payment is constant across all income levels, and therefore the welfare distributions are identical except for a constant increase across all individuals. Therefore, when computing the *mean-adjusted* Relative Distribution, the effect of the refund payment disappears.

In the *Transit Subsidy* scenario, the refund is very slightly polarizing. This is counterintuitive, since those who use transit, and therefore those who receive the benefit, tend to have lower incomes compared to those who do not take transit. However, here we are measuring the effect of the refund *after* the other components have had their effect. Therefore, this result suggests that the refund may actually be *overcompensating* low-income individuals, since they have already benefited from the direct effects of the congestion pricing plan (relative to higher income individuals). Still, this polarizing effect of the refund only marginally decreases this scenario's overall depolarizing effect.

Finally, a strong polarizing effect is found in the *Tax Reduction* scenario. Because the refund payment increases proportionally with income, there is a clear and expected polarizing effect, and its magnitude is greater than any of the other effects measured. This causes the overall effect of this scenario be highly polarizing compared to the others.

7 Data Needs for Inference Testing

When applying these methods to empirical data, it is appropriate to quantify the uncertainty associated with their estimates, giving an indication of what the probability is that the variation observed is due not to the theorized relationship, but rather to the effects of unobserved factors or random disturbances. In this application, however, we are simulating data using a known, albeit complex, modeling system. While we may not know what the full set of results will be, we can say that those results will be influenced by factors under our observation.

The results in this study were influenced by three sources of uncertainty: 1) uncertainty in the estimated travel behavioral parameters, 2) internal stochasticity in the simulation of travel data, and 3) uncertainty in fitting non-parametric assessments to the simulated welfare data.

Due to the complexity of the travel model, confidence intervals for the first two sources of uncertainty are not directly obtainable through analytical means, and computational means have not yet been used due to processing time demands. While we can make useful observations about the shapes of these Relative Distributions and the values of the Mean Relative Polarization indices, we cannot yet reject null hypotheses with any specific level of certainty. We can, however, still interpret the results even in the presence of these first two sources of uncertainty, since the two samples in each of those cases have identical members, and it is only the scenario that is different between the two distributions.

With regard to the third source of uncertainty, we can make some initial estimates of what sample size might be required to make these inferences with a given level of confidence. To make these estimates, the sample was replicated and some of the analyses above repeated until the confidence intervals narrowed enough that the result could be said to be significant. Some limited experimentation with replicating the sample found that a new sample 200-times the present sample size (i.e. a valid sample of approximately 544,000 individuals) could potentially demonstrate with 95% confidence that the *Tax Reduction* scenario is regressive. While this sample size is out of range for most household travel surveys, it is likely possible to carry about the above analyses using synthesized demographic data for a 100% sample of a metropolitan area's population, similar to the full datasets commonly used as land use data for travel demand models.

8 Conclusions

This study set out to examine a series of questions concerning the redistributive effects of Stockholm's congestion pricing plan. Foremost among the results is the fact that all of the measured welfare effects are small in magnitude. That result, combined with the inability to compute confidence levels for the results, must qualify all of the findings as inconclusive. The small magnitudes, however, are in fact consistent with earlier findings by Eliasson and Mattsson (2005).

Despite this limitation, the results strongly suggest that the progressive and regressive tendencies of the results related to the choice of method for reallocating collected tolls to the general public, with the most progressive scenario being the lump sum payment and the least progressive scenario being the reduction in the income tax rate. These conclusions are also consistent with the conclusions by Eliasson and Mattsson (2005).

An additional finding was that when we disregard the distributional effects of the refund, the congestion pricing system itself has a progressive equity effect on welfare levels. The decomposition exercise in Question 8 showed that most of this progressive effect stems from the tolls paid by autos; this finding supports the argument in earlier literature that because higher-income individuals are more likely to drive, they are also more likely to pay the toll.

The question remains, however, what burden is placed on low-income individuals who initially took auto. However small this subgroup may be, it persists as an important distributional consideration when implementing a congestion pricing system. Moreover, even though this subgroup was indeed small in the present study on Stockholm, low-income drivers are likely to constitute a much larger subset of the population in an American city, magnifying the importance of this issue. Further research will be needed to examine the distributional effects of congestion pricing, *given* the initial choice of mode, so that the distributional effects on such subgroups can be discerned.

Two additional limitations of the present study are noteworthy. First, it is important to note that the SAMPERS travel modeling system, from which travel data was obtained for this study, is unable to predict scheduling decisions in the face of congestion pricing. This limitation means that individuals are modeled as having fewer options for how to respond to the toll than they would have in reality. With additional options, such as taking an auto trip at a late hour to avoid a toll, individuals would in general be better off, although the distribution of this flexibility is fairly unknown. It is particularly disconcerting to exclude this effect, since one of the primary criticisms of congestion pricing is that wealthier individuals have greater travel flexibility (Arnott et al., 1993).

This limitation suggests that the above results are conservative in their conclusions of regressivity, and liberal in their conclusions of progressivity.

Second, the specifics of the congestion pricing program, as modeled, are different in important ways from the plan that is now being applied for trial implementation. Specifically, the final plan has no bisection cordon line and it treats trips to and from the island community of Lidingö (whose only land connection is through the central city) quite the opposite as they were modeled. Instead of treating Lidingö as part of the central city, such that its trips are free to the central city but charged elsewhere, in the final plan those trips would be charged for entering the central city but given a refund if they exit the central city in some other way within a 30-minute period.

While the limitations above inhibit the validity of the results in judging Stockholm's actual congestion pricing plan, the results remain informative on what kinds of conclusions about equity can be drawn using non-parametric methods.

Acknowledgements

The author thanks Lars-Göran Mattsson of the Royal Institute of Technology, Stockholm, for his invaluable guidance and feedback; Jonas Eliasson of Transek AB for his generosity in providing travel model simulation data and his assistance with its interpretation; Paul Waddell and Mark Handcock of the University of Washington, Seattle, for guidance on the theoretical and methodological issues preconditioning this study; and the University of Washington's Valle Fellowship and International Exchange Program for making possible this research. A substantial portion of this research was carried out while the author was a visiting Ph.D. Candidate at the Royal Institute of Technology.

References

Arnott, R., de Palma, A., and Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, 83(1):167–179.

Arnott, R., de Palma, A., and Lindsey, R. (1994). The welfare effects of congestion tolls with heterogeneous commuters. *Journal of Transport Economics and Policy*, 28(2).

Arnott, R., de Palma, A., and Lindsey, R. (1998). Recent developments in the bottleneck model. In Button, K. J. and Verhoef, E. T., editors, *Road Pricing, Traffic Congestion, and the Environment*, pages 79–110. Edward Elgar, Cheltenham, UK.

Bae, C.-H. C. (1997). The equity impacts of Los Angeles' air quality policies. *Environment & Planning A*, 29(9):1563–84.

Bernhardt, A., Morris, M., Handcock, M. S., and Scott, M. A. (1999). Trends in jobs instability and wages for young adult men. *Journal of Labor Economics*, 17(4-2):S65–S90.

Beser, M. and Algers, S. (2001). SAMPERS - the new Swedish national travel demand forecasting tool. In Lundqvist, L. and Mattsson, L.-G., editors, *National Transport Models*, pages 101–118. Springer, Heidelberg.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74:829–836.

Cleveland, W. S. and Devlin, S. J. (1988). Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83:596–610.

Cohen, Y. (1987). Commuter welfare under peak-period congestion tolls: Who gains and who loses? *International Journal of Transport Economics*, 14:239–66.

Cowell, F. A. (1977). *Measuring Inequality*. Philip Allan, Oxford.

Eliasson, J. and Mattsson, L.-G. (2005). Equity effects of congestion pricing: quantitative methodology and a case study for Stockholm. *Transportation Research Part A: Policy and Practice*, In Press.

Evans, A. (1992). Road congestion pricing: When is it a good policy? *Journal of Transport Economics and Policy*, 26(3):213–243.

Foster, J. E. and Shneyerov, A. A. (1999). A general class of additively decomposable inequality measures. *Economic Theory*.

Giuliano, G. (1994). Equity and fairness considerations of congestion pricing. In *Special Report 242: Curbing Gridlock: Peak-Period Fees to Relieve Traffic Congestion, Volume 2*. Transportation Research Board, National Research Council, Washington, DC.

Handcock, M. S. and Janssen, P. L. (2002). Statistical inference for the relative density. *Sociological Methods & Research*, 30(3):394–424.

Handcock, M. S. and Morris, M. (1999). *Relative Distribution Methods in the Social Sciences*. Springer-Verlag, New York.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 27:986–1005.

Layard, R. (1977). The distributional effects of congestion taxes. *Economica*, 44:297–304.

Levine, J. and Garb, Y. (2002). Congestion pricing's conditional promise: Promotion of accessibility or mobility? *Transport Policy*, 9(3):179–188.

Mayeres, I. and Proost, S. (1997). Optimal tax and public investment rules for congestion type externalities. *Scandinavian Journal of Economics*, 99(2):261–279.

Parzen, E. (1979). Nonparametric statistical data modeling. *Journal of the American Statistical Association*, 74(105-131).

Parzen, E. (1992). Comparison change analysis. In Saleh, A., editor, *Nonparametric Statistics and Related Topics*, pages 3–15. Elsevier, Netherlands.

Pigou, A. C. (1952). *The Economics of Welfare*. Macmillan, London, 4th edition edition.

Raux, C. and Souche, S. (2004). The acceptability of urban road pricing: A theoretical analysis applied to experience in Lyon. *Journal of Transport Economics and Policy*, 38(2):191–216.

Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA, revised edition.

Richardson, H. W. (1974). A note on the redistributional effects of road pricing. *Journal of Transport Economics and Policy*, 8:82–85.

Santos, G. and Rojey, L. (2004). Distributional impacts of road pricing: The truth behind the myth. *Transportation*, 31(1):21–42.

Small, K. A. (1983). The incidence of congestion tolls on urban highways. *Journal of Urban Economics*, 13:90–111.

Teubel, U. (2000). The welfare effects and distributional impacts of road user charges on commuters: An empirical analysis of Dresden. *International Journal of Transport Economics*, 27(2):231–255.

Vickrey, W. S. (1968). Congestion charges and welfare. *Journal of Transport Economics and Policy*, 2:107–118.