

How Sensitive is Transit Ridership to Gas Prices?

Andrew Foote

April 12, 2010

Abstract

This thesis designs an econometric study to test the extent to which gas prices affect transit ridership. This study is an econometric application of the theory of binary modal choice, which describes the costs of each mode of transportation and how it affects a consumer's choice of transit mode. While the focus of this paper is the elasticity of transit ridership with respect to gas prices, there are a number of other determinants of transit ridership. The most important explanatory variables are wage rate, average commuting time, service frequency of public transit, and transit fare. My data is drawn from ten cities in the United States on a monthly basis from the years 2004 to 2008. In order to estimate the hypothesized relationship, I will use OLS fixed-effects multiple regression in order to control for the differences between urbanized areas that cannot be quantified. Using this econometric technique, this paper constructs an econometric methodology that I argue throughout is statistically superior to the predominant methodology in the literature. In addition, because of the potential for simultaneous causality between service frequency and transit ridership, I propose a number of instrumental variables in order to ensure that simultaneous causality does not present itself as a threat to internal validity in my estimations. I find that the use of instrumental

variables improves my estimations of elasticities, suggesting that the use of instrumental variables is merited.

1 Introduction

The purpose of this thesis is to test the extent to which transit ridership is affected by changes in gas prices. My analysis is based on the theory of binary modal choice, which describes a consumer's choice between two substitutable modes of commuting: driving and public transit. Because these modes are substitutes, as the price of gas rises, the corresponding price of driving also rises, inducing some people who were marginal to substitute from driving to public transit. The theory is most clearly articulated in the seminal paper by McFadden [1974], the framework of which is addressed in detail in section 3, which outlines the theoretical basis for my research. For the purpose of my analysis, I control for a number of other variables, since they affect transit ridership as well. These variables are supply of transit, measured in terms of service frequency, transit fare, average commuting time, working population, unemployment rate, average hourly wage, and cost of living as measured by city-specific CPI. Additionally, there are variables that are included in other studies of transit demand that I do not include because of my use of fixed-effects estimation. These variables are population density, parking costs and percentage of employment in the central business district, all of which do not vary widely in a city over time. Finally, I consider using seasonal dummies to control for exogenous fluctuations in demand of transit throughout the year. My data will be drawn from a number of sources, most notably from the American Public Transit Association (APTA), and the Bureau of Labor Statistics (BLS). An extensive discussion of the variables included in the regression and their expected signs is in the methodology section (section 3), and the description of the dataset is in section 5. Because my data will be longitudinal, I will use an OLS fixed-effects regression as a

way to control for potentially unmeasured variables. This technique and my reasoning for using it will be addressed in section 4. Furthermore, I am going to address the possibility of simultaneous causality between the service frequency variables and transit ridership, as well as possible solutions. These will be fully discussed in sections 4.1 and 6.2, which discuss the rationale for and regression results using instrumental variables.

The results of this econometric study will be useful for three reasons. First, knowing the cross-elasticity between transit ridership and gas prices will help transit authorities to forecast what demand will be in future periods, which will be useful in order to accommodate sudden surges in demand for transit, as we saw in the United States in the summer of 2008 when gas prices suddenly spiked. Second, transit authorities will be able to better understand the underlying structure of demand for public transit, enabling them to encourage transit ridership. Also, if my estimated elasticities are significantly different than other estimates in the literature, it will show that by using cross-sectional data, previous studies have been misspecifying the functional form of their regressions, thereby supporting my contention that the regression specification in this paper is more statistically sound.

While a number of extensive studies of transit demand analysis exist, the predominant methodology in these studies is cross-sectional analysis. However, differences in gas prices across cities do not necessarily affect transit ridership directly. Essentially, I argue that these previous studies have faulty estimates of the elasticity of transit ridership with respect to gas prices, because they have a misspecification of functional form, by misdefining the market for transit. As gas prices rise in a city, people *in that city* substitute on the margin from driving to public transit. However, previous studies have measured gas price levels across cities, which does not determine the demand for transit, since people do not substitute from driving in one city to public transit in the other or vis-versa. Therefore, I argue that using panel data and a OLS fixed-effects estimation is the correction regression specification

in order to estimate the elasticity in question, because we observe gas prices and transit ridership changing through time, controlling for other variables. For this reason, my methodology, which is discussed in depth in section 4, is much stronger for estimating the gas price elasticity.

The remainder of this paper is organized as follows. Section 2 reviews the literature on the subject, and is separated into subsections which summarize background information, similar empirical papers, and the econometric techniques used in the literature. Section 3 discusses the methodology of the study and reasoning for inclusion of the variables that were mentioned earlier. Section 4 discusses the econometric techniques used and potential estimation problems, while section 5 discusses the structure of the dataset and data sources. Section 6 presents my regression results and discusses possible improvements to the analysis, while section 7 concludes the study.

2 Literature Survey and Analysis

There is a wealth of literature on the subject of modal choice theory, and I attempt to review the most important articles here. The structure of this section is as follows: Subsection 2.1 summarizes articles that provide background information about the relationship between transit ridership and other important variables, most notably gas prices. Subsection 2.2 discusses and analyzes empirical papers that have addressed topics that are similar to, or contribute to the framework of, the current question. Finally, subsection 2.3 looks at the econometric techniques that were used in the other empirical papers, analyzing the strength of each technique.

2.1 Background Information

A number of articles and books provide background information relevant for my hypothesis. The most groundbreaking of these literature reviews is by Goodwin [1992], who attempted to unify the research of transit elasticities.

He notes in his introduction that earlier literature reviews on the subject of transit demand did not have any articles in common, because of the stark division between the research done by peer-reviewed journals and government agencies. These two strains of the published literature rarely intersected before Goodwin, who makes a push to unify these two strains of transit demand analysis. His review covers the effects of car costs, principally measured in gas prices, public transportation costs, which is concerned mostly with the fare elasticity, and accessibility and car ownership, which is important in determining the populations of marginal and transit-dependent riders. Particularly important for my analysis is the observation that consumers adjust to changes in gas prices over a fairly long time horizon, sometimes taking up to five years for the adjustments to stabilize. This suggests that I should include lagged gas price variables into my analysis. I discuss this issue further in section 4.1, which addresses the use of a lagged gas price variable.

Following in this strain, Litman [2004] also spans the divide between journal publications and government agency reports. Litman covers the factors that most directly affect the elasticities of transit demand. He breaks down these effects into the following categories: user type, trip type, geography, time horizon and transit type. User type is very important, because transit dependent riders are relatively insensitive to price, while discretionary riders (who have the option of driving themselves to work) are more price sensitive. Applying this observation to the current empirical question, this fact implies that if there are more transit dependent riders, we would expect the cross-elasticity with gas prices to be lower, because less people have the option to switch from driving to transit use. This is because transit dependent riders do not have the option to drive, and so gas prices do not directly affect their commuting mode choice.

Furthermore, geography has an impact on the elasticities of transit demand, because the larger a given city is, the higher the population of transit dependent riders. Additionally, larger cities have more congestion and gener-

ally longer commuting distances, all of which contribute to ridership growth. Moreover, larger cities experience greater economies of scale in supplying public transit.

Finally, transit type contributes to the transit elasticities, because bus and rail generally serve different markets. While buses are designed to primarily serve transit dependent riders, light and heavy rail are designed to serve discretionary riders. For this reason, we would expect light-rail ridership to exhibit a higher elasticity with respect to gas prices, because light-rail riders are more affected by gas prices. Additionally, rail trips are generally longer, which means that the comparative driving costs are higher, meaning gas prices contribute a larger amount to overall trip cost.

Litman also summarizes the findings of a number of different studies on the elasticities of transit. Of particular note, he summarizes the ranges of the elasticity of transit ridership with respect to gas prices. He finds that the elasticity for commuting trips is about 0.20, while it is lower in other cases, because most other trips are discretionary. Finally, he notes that overall, elasticity with respect to gas prices is about 0.34, from Goodwin [1992].

Another article that contributes to background knowledge in the literature is Bento et al. [2003], which discusses the impact of the size and shape of an urban area and its effect on transit ridership. They list a number of different measures of urban form that are important in transit demand analysis, some of which I try to incorporate into my analysis. First, they argue that the road density is important, because if a city has a less dense road network, the cost of driving is higher (because it is hard to directly get to where one works), and therefore more people would be encouraged to take public transit. However, at the same time, this makes transit less accessible and more costly for the transit authority to widely provide, which makes any effect ambiguous from a purely theoretical standpoint; for this reason, I do not consider road density. Also, it is likely constant over time, and is controlled for in fixed effects. Second, they argue that city shape is an important

determinant of the cost of driving, since “theory suggests that trip distances should be longer in long, narrow cities than in circular cities with radial road networks.” (page 467) This observation is an important one, because it is possible to have lower transit costs in a more oblong city, which would encourage transit use, a point that Kain and Liu [1999] argue in their article, which addresses transit in San Diego. Third, they maintain that residential land use patterns are important for addressing transit demand. They note that, “in a circular city the natural measure of the pattern of residential land use is the population density gradient.” (page 468) However, in cities with less circular shapes, this measure does not particularly fit the actual population density gradient. For this reason, they reject any measure of the pattern of residential land use (because all other measures are also imperfect). They instead create a measure of population centrality, “in which the actual population at each distance from the CBD is weighted by distance.” (page 468) This measure would be computationally similar to the density gradient in a circular city, but more accurately reflects cities such as San Diego, which are not circular.

The final variable of interest for my analysis is the distribution of employment measure. This variable is important because if the distribution of employment is uniform, then the likelihood that people are closer to their place of work is higher, leading to lower driving times. However, if employment is concentrated in a few areas, congestion costs will be higher, parking will be expensive and driving times will also be increased, all of which will encourage transit use on the margin. For this reason, the distribution of employment is probably the most important measure that Bento et al. discuss, since it quantifies some costs of driving that are not otherwise captured by other variables. Surprisingly, this variable, constructed as the jobs-housing balance, was not statistically significant in any of their regressions, possibly because it is highly correlated to city shape and population centrality.

Additionally, the paper by de Jong and Gunn [2001] is extremely impor-

tant, as it summarizes the impacts on the costs of driving. First, they distinguish between stated preference, found when offering surveys, and revealed preference, the actions taken when changes actually occur. For the purpose of transit demand elasticity analysis, I am only concerned with the revealed preferences, not only because the data is easier to collect, but it is less likely to be biased and there will be fewer instances of errors-in-variables. Second, they summarize the past research on fuel price elasticities, which provides me a range that I can expect for my results. However, since de Jong and Gunn are looking at studies done in Europe, where transit is more prevalent and culturally pervasive, and car costs are much higher, these results do not directly apply to my data, which will be collected from cities in the United States. What is most important from this article are the variables they use to measure responsiveness to a rise in the price of driving. First, they look at vehicle miles traveled and average number of trips, both of which are impacted greatly by an increase in gas prices. Furthermore, they note that “for long-distance trips, public transport...has a relatively strong competitive position.” (page 143) For this reason, I attempt to include a variable for average commute length, in order to capture this effect, and test if it is also true in the United States.

For the most comprehensive discussion of determinants of transit demand and supply, Taylor et al. [2009] and Kohn [2000] are central articles to the literature. For Taylor et al., the determinants of demand fall into three broad categories: regional geography, metropolitan economy and population characteristics. While Kohn does not categorize his determinants in this way, his variables are generally consistent with these categories as well, and so they can be discussed concurrently.

First, regional geography includes variables such as community size, employment concentration, and regional location in the United States. The first two have been discussed at length earlier in this literature review, while the third one is slightly less intuitive. Regional location can capture two partic-

ular aspects of an urban area; first, the ethos of the region may encourage or discourage public transportation use, which exogenously affects the utility of a transit trip. For instance, if a region is known for its environmentalism, then taking transit can be a way to fit into that culture, and so endogenously increasing the utility derived from that transit trip. Conversely, if a region generally discourages public transit, then it would be seen as a breach of cultural norms to use public transit, thereby endogenously decreasing the utility of the trip. Secondly, regional weather patterns are particularly important in determining public transit demand, because inclement weather discourages transit use by increasing the cost of access time (fewer people want to ride public transit when it involves walking through two feet of snow to catch a bus). Because it is sometimes difficult to measure inclement weather, regional dummies are generally used to capture each of the above effects. Since both inclement weather and the ethos of a city do not vary widely from year to year, my use of fixed-effects estimation controls for these variables.

Second, the metropolitan economy plays a large role in determining the level of demand for public transportation. Both Kohn and Taylor et al. note that income distribution, economic growth, working population and unemployment rate all play a large role in determining transit demand. I try to incorporate these measures in the variables of unemployment rate, working population and average wage rate.

Finally, population characteristics play a large role in determining transit demand. Taylor et al. note that age distribution, college population and percent of population below the poverty line also affect the level of transit demand. College population is particularly important, since they usually do not fall under the poverty line, but are very transit dependent. Additionally, the population below the poverty line will likely be extremely transit dependent, making transit ridership less sensitive to changes in the price of gasoline. Kohn also adds that the population growth, immigration rate and fertility rates may also affect transit demand, because as population growth

becomes more rapid, the costs of commuting and congestion increase, making transit relatively more attractive. Also, if fertility rates rise, more mothers (or fathers) stay home with their children, decreasing transit demand.

Also contributing to the literature is Thompson and Brown [2006], who bifurcate transit demand characteristics into external and internal factors with a view on advising transit authorities as to what policy decisions should be made to most impact transit ridership. All of their external factors have been addressed in the previous paragraphs, and I do not want to revisit them now, but their internal factors are important for the current analysis. Included in the internal factors, over which transit authorities have control, are service frequency, service reliability, service amenities (such as covered bus stops) and transit fare. Service frequency, reliability and amenities affect the access costs of transit ridership, which is largest component of the cost of using public transit. I include service frequency and transit fares in my analysis, but exclude service reliability and amenities, which are both subjective and not subject to uniform measurement. Overall, while this article is useful in providing background information, their analysis is unclear and unreliable, as they eliminate insignificant variables from the regression, something that is not an acceptable econometric practice.

There are a number of other articles that contribute information to my analysis. First, Mattson [2008] outlines recent trends in gas prices and transit ridership, focusing his analysis to smaller urban areas (which have comparatively lower rates of public transit use). His empirical study is discussed more in subsection 2.2. Second, Maley and Weinberger [2009] note that because each month has a different number of days, there are slight measurement biases if measurements are taken on a monthly basis; instead, they divide ridership by the number of working days in a given month. Finally, Kain and Liu [1999], discussed above, outline certain reasons why transit ridership growth has been so robust in San Diego and Houston, noting that in both cities, strategic expansions as well as the combination of a decrease in fares

and an increase in service frequency contributed strongly to the growth in both cities. Additionally, San Diego benefited from its geographical anomalies, since it is situated in a valley, making an end-to-end commute very time consuming.

2.2 Similar or Related Hypotheses

Probably most useful in informing the present analysis are the articles that have carried out the same or similar empirical tests. The most comprehensive and contemporary study is Taylor et al. [2009], who find that the coefficient on gas prices is positive and statistically significant. Their analysis is unique in the literature, as they used two-stage least squares in addition to standard OLS. They implemented two-stage least squares in an attempt to solve the simultaneous causality between transit demand and transit supply. Their instrumental variables in the analysis were quite novel, as they use total population and percent voting Democrat in the 2000 presidential election. However, because they do not include an F-statistic for the regression, the statistical strength of these instrumental variables is not clear. In fact, Thompson and Brown [2006] argue that the use of two-stage least squares estimation does not improve the adjusted R-squared value, saying that, "Taylor and Miller, who used two-stage least squares to separate the two effects, found that their method offered little insight over the traditional view." (page 175) Nonetheless, an improvement in the adjusted R-squared value is not the only goal of two-stage least squares analysis. I believe that it is valuable to examine this technique as an option, because it may improve the internal validity of the research, even though it has no effect on the value of the adjusted R-squared statistic. Taylor et al. is particularly useful to my thesis, because they explore all of the possible options for regression analysis, supporting each theoretically and practically. Finally, there is a table in their paper that includes each important variable and the location from which the data was collected, which aided my own collection of data for my

analysis. However, their data was cross-sectional, which significantly hinders their analysis of the transit elasticity with respect to gas prices, because gas prices generally only vary over time; this issue is discussed more fully in section 4.1 on page 28. This is important because individual transit authorities are particularly concerned with changes in variables over time. In fact, most papers in the literature, with two notable exceptions [Mattson, 2008, Bresson et al., 2004], use cross-sectional data. This leads to significant omitted variable bias, which could be easily solved by compiling longitudinal data and estimating the same elasticities using a fixed-effects regression. Therefore, while their analysis gives me a strong foundation, it does not speak to the techniques that I am going to use in my empirical analysis.

Another important empirical paper that contributes to the framework of my analysis is Mattson [2008], who looks at smaller urban and rural transit authorities in his analysis. While the data is not particularly useful (as I am looking at ten major cities), his methodology and analysis are important. First, he develops a lagged regression model to estimate the effect of gas prices on transit ridership, suggesting that the effect of gas price changes is distributed over time. However, he does not include any other variables in this regression, which contributes to omitted variable bias in his analysis. Additionally, Mattson never addresses the fact that gas prices fluctuate seasonally, which may lead to biased results if these expectations are not taken into account. In fact, no paper that I read attempted to address this problem. While I attempt to address this, I have to leave this instead for further research, which is discussed in my conclusion in section 7. Additionally, for the second half of his paper, he uses panel data, but does not include a lagged gas price variable, which was confusing and self-contradictory, given his earlier insistence on its importance. However, Mattson did contribute to my thesis in one significant way. He notes that "the data for the previous year [in APTA reports] are revised from the original reports, and the differences are sometimes significant." (page 21) For this reason, I will follow Mattson's

suggestion and only use the revised figures from the APTA. Mattson also distinguishes between short- and long-run effects for gas price elasticity, in order to measure short- and long-run elasticities. In order to estimate these elasticities, I want to incorporate different lengths of lags into my regression. Since it is not theoretically clear what these lag lengths are, this issue is discussed further in subsection 4.1.

There are also two meta-analyses in the literature [Holmgren, 2007, Hensher, 2007], which look to combine a number of different studies and their results in order to find unified elasticity measures. Both of these papers attempt to explain the wide range of elasticity estimates found in recent studies of transit demand elasticities and cross-elasticities. While Hensher focuses his analysis on the United States, Holmgren's analysis is particularly concerned with the differences between elasticity estimates in Europe, America and Australia. Hensher's findings are that the variation in elasticity estimates is almost entirely due to different model specifications, not due to actual changes in the elasticities.

Overall, Holmgren's analysis is much more useful in understanding the differences in elasticity estimates. He addresses the difference between Europe's transit system and the systems in Australia and the United States, finding that transit elasticities differ significantly between the two systems. In particular, he finds that the fare elasticity for Europe is -0.75, compared with -0.59 in America and Australia. Additionally, he finds the short-run elasticity of transit ridership with respect to gas prices to be 0.4 in Europe, contrasted with 0.82 in America and Australia. Finally, he finds that the long-run elasticity of transit ridership with respect to gas prices is 0.73 in Europe, while it is estimated to be 1.15 in America and Australia. From these results, I can conclude that American and Australian transit ridership is more sensitive to changes in gas prices, possibly because driving and public transit are both viable means of commuting, which is not always true in Europe. Holmgren makes two important points in his conclusion, both of which

shape my methodology. First, he notes that "if there is no significant difference between including and excluding a variable it is still possible that that variable has explanatory value and should be included." (page 1033) This is a direct critique of Thompson and Brown [2006], who do exactly this, eliminating a number of variables from the regression after their coefficients were found to be statistically insignificant. Finally, Holmgren notes that there are a number of variables that need to be determined inside of the model, such as vehicle-kilometers of transit authorities. He notes that there a significant bias in many studies "due to the fact that treating a variable as exogenous when it is an endogen." (page 1034) I address this concern by instrumenting for service frequency. This is discussed in further depth in section 6.2.

There are a number of other articles in the literature that contribute somewhat to the question at hand. First, Kohn [2000] analyzes gas price elasticity of transit ridership. However, his regressions are in semi-log specification, which leads to problems in interpreting the coefficients. Also, Wang and Skinner [1984] use a Box-Cox transformation to choose between linear and logarithmic specifications, which adds some validity to their conclusions on price elasticity. However, their conclusions fall generally in line with similar studies, calling into question the value-added of performing such complex calculations. A distinctly novel study was performed by Kyte et al. [1988], who looked at the demand determinants of transit ridership in Portland, Oregon. Their time-series analysis was quite expansive, and looked at every bus line that Tri-Met, the transit authority in Portland, ran from the years 1971 to 1982. They find that both service level and market size explain the vast majority of the variation in transit ridership. I incorporate these variables into my analysis as service frequency and working population, which serves as a close proxy for how many people are in the market for public transit in a given month.

Additionally, Bresson et al. [2004] use a Bayesian approach and compare it to the results from the traditional fixed-effects model that is commonly

used with panel data. This paper was particularly useful because of their discussion of time-series estimation, which informs the analysis in this paper. Their conclusion is that the Bayesian approach to empirical estimation is a significant improvement over fixed-effects, because it allows elasticities to differ by region and by city. This is an important and unique paper in the literature, because they use panel data and fixed-effects, something no other paper in the literature does.

Finally, Kuby et al. [2004] look at the factors which influence light-rail ridership, and their conclusions can be applied to all public transit demand analysis. Their analysis is incredibly strong, looking at population around the light-rail boarding stations, while also including climate of the cities in which they study. Climate is an important determinant of transit ridership, and they note that, "[excessively] hot or cold weather discourages transit use because it often requires walking and waiting outdoors." (page 235) Because unfavorable weather increases the access cost of using public transit, many marginal riders will choose to drive, decreasing ridership. The estimate of the coefficient on unfavorable weather in their study was negative and statistically significant, which means that on average, cities with more frequent unfavorable weather will likely observe lower ridership, all other things equal. However, instead of incorporating this variable into my analysis, because the frequency of inclement weather does not vary much from year to year, I solve this measurement problem by using longitudinal data and fixed-effects regression analysis. Incorporating such a variable fully into the analysis is left for future work.

Table 1 on the next page summarizes the elasticities of interest found in the literature that has been reviewed in this paper.

2.3 Econometric Techniques

There are a number of papers that I consulted in constructing the methodology that I use in the current research. First, because I am estimating

Table 1: Summary of Elasticities of Transit Ridership w.r.t. Gas Prices

Study	Location	Type of Elasticity	Elasticity
TRACE, 1999	Europe	—	0.20
Taylor et. al	United States	non-instrumented	0.73
		instrumented	1.45
Bresson et. al	France	Short-Run	0.08
		Long-Run	0.14
Mattson	Barnesville	—	0.042
	Detroit Lakes	—	0.065
Wang and Skinner	Des Moines, IO	—	0.80
	New York	Surface Transit	0.17
Kyte et. al	Portland, OR	—	0.30
de Jong and Gunn	The Netherlands	Short-Run	0.18
		Long-Run	0.16
	Italy	Short-Run	0.22
		Long-Run	0.22
Holmgren	United States	Short-Run	0.82
		Long-Run	1.15

Courtesy: Taylor et al. [2009], Litman [2004], de Jong and Gunn [2001], Mattson [2008], Wang and Skinner [1984], Bresson et al. [2004], Kyte et al. [1988], Holmgren [2007]

elasticities, my regression will be expressed in a double-log specification, as noted by Stock and Watson [2003]. Furthermore, in developing my use of fixed-effects regressions, I consulted Bureau, who outlines the advantages and disadvantages of the estimation technique used in this paper. The discussion of my methodology is found in section 4 on page 27.

One of the main differences in the literature from paper to paper is the regression technique that is used, which Holmgren [2007] notes is significant in determining the elasticity, while the type of data set that was used is also important. Many papers [Maley and Weinberger, 2009, Thompson and Brown, 2006, Mattson, 2008, Kuby et al., 2004] use only OLS multiple regression technique and a cross-sectional data set, as it is relatively straight-forward for an estimation of elasticities. Additionally, Wang and Skinner [1984] use an OLS regression with some adjustments to the general framework. In particular, they factored in a number of dummy variables in order to avoid auto-correlation, and used a Box-Cox transformation to choose between a log or linear specification.

Also, although it is contested, Taylor et al. [2009] use a two-stage least squares regression to estimate elasticities. Their rationale for the use of this technique is that transit demand and transit supply are co-determining, and so to avoid this simultaneous causality, they instrument for transit supply. However, the transit market is not a conventional market, because there is only one agent on the supply side of the market, which means that supply generally changes abruptly, usually in reaction to earlier demand changes. For this reason, it may be the case that there is not simultaneous causality in the current period, but that transit supply in past periods determines current demand levels or the forecasted supply determines demand. For this reason, one of my instrumental variables is annual scheduled vehicle revenue miles. This issue will be more fully addressed in sections 4.1 and 6.2, where I discuss each variable in depth and then present the results using these instruments.

Additionally, because the overarching theory that applies to this analysis is modal choice, Bento et al. [2003] estimates a multinomial logit model of mode choice, where the choices are driving to work, taking rail transit, taking non-rail transit, walking or biking to work. This analysis is particularly useful when looking at the effect that changes in variables have on an individual's choice of transit mode, but it is not particularly useful to estimating aggregate elasticities of transit, which is the question that I am addressing.

Finally, Bresson et al. [2004] use Bayesian analysis with shrinkage estimators to estimate transit demand elasticities. This is in contrast to the fixed-effects model that is used for most time-series regression analysis, and it is particularly strong because it allows the computation of elasticities for each urban area. In order to test the goodness of fit for Bayesian analysis, Bresson et al. test the homogeneity of coefficients (see their paper, pages 272-275 for full details), concluding that there is some heterogeneity in the coefficients for each city. This means that the transit elasticities differ from city to city. However, because the number of regressors is much higher (because it estimates city-specific elasticities), Bayesian estimation requires a much larger dataset in order to have sufficient degrees of freedom. Although this technique is intriguing and quite novel, my dataset is simply too small for Bayesian estimation to be a viable option (with only 600 observations). Furthermore, the purpose of this paper is to estimate a general elasticity of transit ridership with respect to gas prices, not city-specific elasticities, and so this econometric technique is not applicable to my research question. However, the Bayesian approach is a possible extension to this research in the future.

3 Theoretical Framework

3.1 Regression Specification

The purpose of the present paper is to explore the relationship between two variables: transit ridership and gas prices, while including a number of other explanatory variables, consistent with economic theory. The equation to be estimated is the following:

$$\begin{aligned} TR_{it} = & \beta_0 + \beta_1(GP_{it}) + \beta_2(TF_{it}) + \beta_3(SF_{it}) + \beta_4(CT_{it}) \\ & + \beta_6(WP_{it}) + \beta_7(UR_{it}) + \beta_9(\omega_{it}) + \beta_{11}(CPI_{it}) \\ & + \beta_{12}(SeasonalDummies) + \epsilon_{it} \end{aligned} \tag{1}$$

In the above equation, TR_{it} is transit ridership in city i at time t . GP_{it} is gas price in city i at time t , TF_{it} is transit fare in city i at time t , SF_{it} is service frequency of public transit in city i at time t , and CT_{it} is average commuting time in city i at time t . Furthermore, WP_{it} is the working population of city i at time t and UR_{it} is the unemployment rate in city i at time t . Finally, ω_{it} is the average wage rate in city i at time t and CPI_{it} is the consumer price index in city i at time t . *SeasonalDummies* is a set of seasonal dummies controlling for time of year, and are discussed later. ϵ_{it} is the error term. Equation 1 will be in a double-log specification, which means that each variable will be expressed in logarithmic form, with the exception of the *SeasonalDummies* and *UR* (which is already expressed as a percentage), in order to have the coefficients estimate the respective elasticities for transit demand.

I look at three different lengths of lags, based on the literature and general consensus: one month, three months and six months. Essentially, the one-month lag measures short-term response, three-month lag measures medium-term response, and the six-month lag measures long-term response to gas price changes. I also include a regression with current gas prices; the reason-

ing behind this will be discussed later in section 5.2, which explains how I constructed the data set.

The economic theory of binary modal choice explores the conceptual relationship between the dependent and explanatory variables. This theory is explained most comprehensively by McFadden [1974], and the next section looks at this theory in depth.

3.2 Binary Modal Choice Theory

What follows is a brief explanation of the theory of binary modal choice.¹ For a more extensive discussion on the theory, consult McFadden [1974] and Oum et al. [1992]. Suppose that a mode of commuting, mode i , has a vector of attributes, such that $x_i = (C_i, Tv_i, Ta_i, K_i)$, where C_i is the cost of mode i , Tv_i is the on-vehicle time of mode i , Ta_i is the access time of mode i and K_i is the comfort level of mode i . Generally, because there is no robust measure of comfort, empirical analysis focuses on the first three attributes of each mode of transportation. Suppose further that there is some prevailing average wage rate, called ω (notice that this is the same as the ω in equation 1), and that CPI is an index of the cost of living in the city (identical to CPI used in equation 1). Therefore, the consumer will choose mode i instead of mode j when the following condition holds:

$$\theta_4(K_i - K_j) > -\theta_1 \frac{(C_j - C_i)}{\omega} - \theta_2(Tv_j - Tv_i) \left(\frac{\omega}{CPI}\right) - \theta_3(Ta_j - Ta_i) \left(\frac{\omega}{CPI}\right) \quad (2)$$

Also, in the above equation, each θ_i is the weight attributed to each type of cost. For instance, we can say θ_3 is greater than θ_2 because the level of discomfort experienced because of accessing transit is higher than the discomfort of in-vehicle time on public transit. These weights are derived in McFadden [1974]. While they are important, they are not crucial for

¹This description borrows heavily from section 2.3 of the paper by McFadden [1974].

understanding the comparative statics that follow. For further explanation, consult section 2 and following in his paper.

To make the consumer decision more clear, we can rewrite this equation below:

$$\theta_4 K_i - \theta_1 \frac{C_i}{\omega} - \theta_2 T v_i \left(\frac{\omega}{CPI} \right) - \theta_3 T a_i \left(\frac{\omega}{CPI} \right) > \theta_4 K_j - \theta_1 \frac{C_j}{\omega} - \theta_2 T v_j \left(\frac{\omega}{CPI} \right) - \theta_3 T a_j \left(\frac{\omega}{CPI} \right) \quad (3)$$

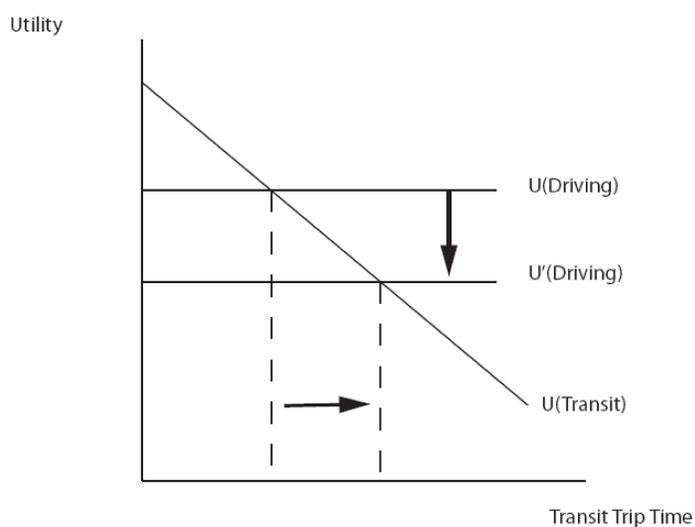
If equation 3 holds, a consumer will choose mode i ; otherwise, they will choose mode j , because the above equation expresses comparative net benefits. Of course, for each consumer, Ta , Tv , ω and K will differ for each mode, and therefore some consumers will choose public transit and others will choose to drive.

3.3 Expected Signs for Variables

Now suppose that mode i is public transit and mode j is driving. Consider the variable, TR , which measures the aggregate number of people for which equation 3 is true. Note that the variable TR is the measure of transit ridership taken from equation 1. In explaining the inclusion of the variables listed in equation 1, I will continually reference equation 3.

Consider first the relationship between the dependent variable, TR_{it} , and the explanatory variable in my hypothesis, GP_{it} . Since gas prices are a portion of the cost of driving, from equation 3 we know that $\frac{\delta TR}{\delta C_j} > 0$, which implies that $\frac{\delta TR}{\delta GP} > 0$, which leads me to expect that the coefficient on gas prices will be positive. Furthermore, this effect can be illustrated graphically, as in figure 1. Notice that in the figure, as gas prices rise, because the cost of driving also rises, the utility of driving decreases at every point. This will lead to a higher collection area for public transit, which will increase transit ridership at the margin. This is intuitive, as economic theory suggests that because driving and public transit are substitutes, when the cost of driving

Figure 1: Transit Ridership Given a Change in Gas Prices

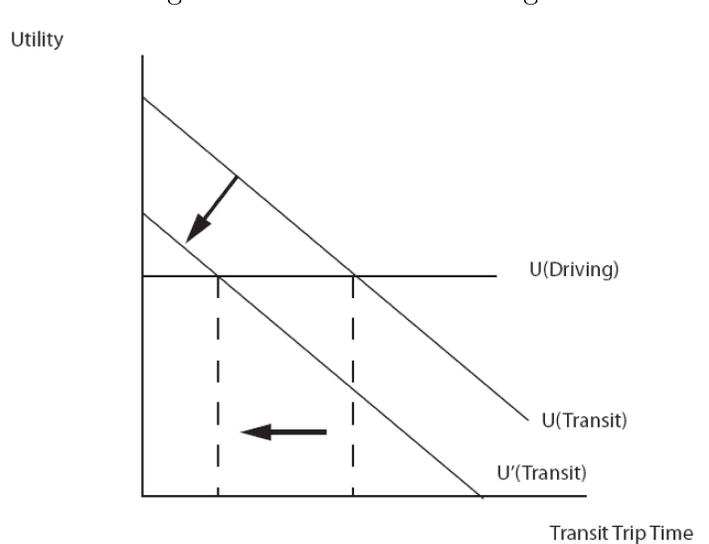


increases, people will substitute away from driving and increase their public transit use.

Next, consider the relationship between the dependent variable TR_{it} and the second explanatory variable TF_{it} . First, note that TF_{it} is a component of the cost of transit, and so it contributes to C_i as expressed in equation 3. For this reason, it is clear from equation 3 that $\frac{\delta TR}{\delta C_i} < 0$, and so we can infer that $\frac{\delta TR}{\delta TF} < 0$ as well. From this result, I expect the coefficient on transit fare, as estimated by equation 1, to be negative. This is also logical, as β_2 is a price elasticity of demand, which should always be negative (I assume, not outlandishly, that public transit is not a Giffen good). As illustrated in Figure 2, we can see graphically that we should predict a negative relationship between these two variables. As the transit fare increases, the utility of using public transit falls at every point. This leads to a smaller collection area for public transit, all else equal. This conclusion is consistent with the partial derivatives I discussed earlier.

Third, consider the impact of service frequency on the dependent variable, TR_{it} . Notice that service frequency, denoted SF_{it} in equation 1, affects the

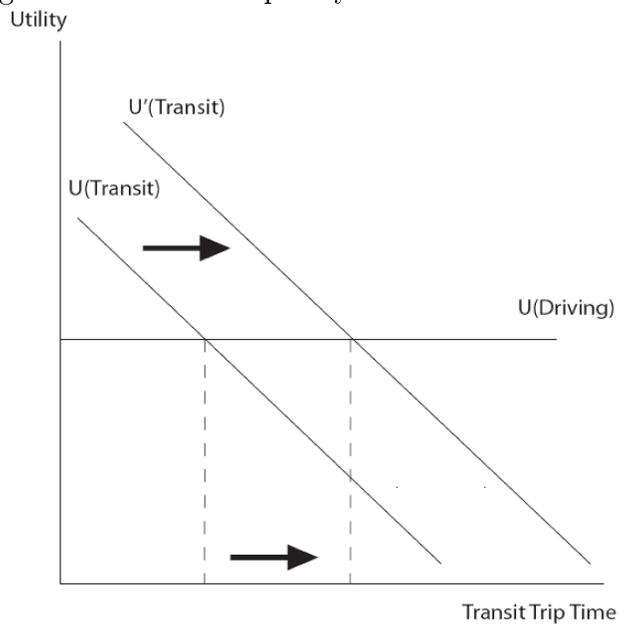
Figure 2: Transit Fare Changes



access time for using public transit, which is denoted Ta_i in equation 3. From this equation, we know that $\frac{\delta TR}{\delta Ta_i} < 0$, which means that as access time for transit increases, transit ridership decreases. Because service frequency and access time are inversely correlated, it can be shown that $\frac{\delta TR}{\delta SF} > 0$. From this result, I expect the coefficient on SF_{it} to be positive, reflecting the fact that increased service frequency decreases the costs of using public transit. This is particularly true in places with more frequent inclement weather, where access time is much more costly (because the disutility of standing in the snow in Buffalo is greater than the disutility of standing in the sunshine of Florida). This concept can also be illustrated graphically, as in figure 3. As shown, when service frequency increases, access time decreases. This causes the utility of transit to increase for every trip time, which will cause marginal riders to take public transit, thereby increasing transit ridership.

Additionally, consider the impact of average commuting time, denoted CT_{it} in the regression equation, on the dependent variable of transit ridership, denoted as TR_{it} . Because average commuting time is a measure of Tv_j in

Figure 3: Service Frequency and Transit Ridership

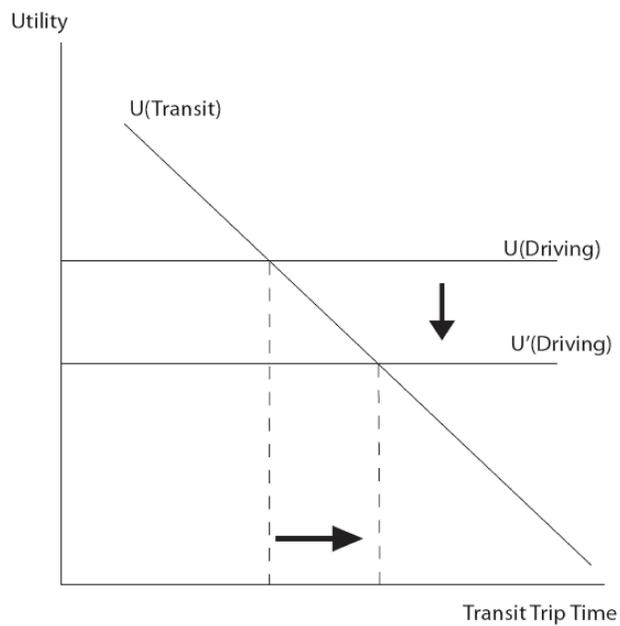


equation 3, we would expect that because $\frac{\delta TR}{\delta T v_j} > 0$, this would imply that $\frac{\delta TR}{\delta CT} > 0$ as well. This is because as in-vehicle time for driving rises, holding everything else constant, the marginal consumer will substitute away from driving and to public transit. This is similar to the effect that a rise in gas prices have on transit ridership, because each are a component to the cost for driving.

The effect of a rise in the average commuting time on transit ridership is illustrated in figure 4. As shown, when average commuting time increases, the utility of driving decreases, which serves to increase the collection area for public transit, which has been discussed above. Between the figure and the partial derivatives above, it is clear that we would expect the coefficient in equation 1 on average commuting time to be positive.

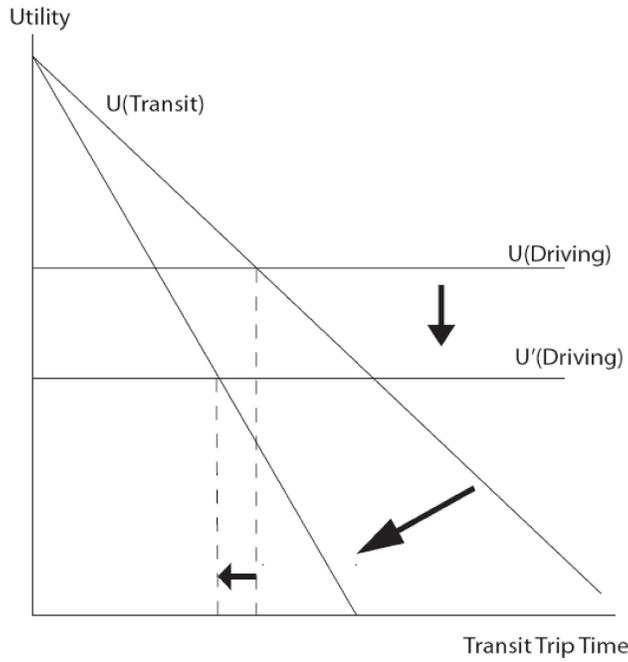
Furthermore, consider the impact of the average wage rate, denoted ω_{it} in the regression equation, on the dependent variable of transit ridership. From equation 3, it is unclear what the impact will be, because ω is on both sides

Figure 4: Average Commuting Time and Transit Ridership



of the inequality. However, we know that as ω increases, the time costs will increase considerably. Because there is no clear comparative static, figure 5 serves as an illustration to foster a discussion of the impact of a change in this variable. As shown, an increase in the wage rate increases the costs of transit and the costs of driving, because of the increase in the opportunity cost of time, which decreases the utility of both modes of commuting. Because the curve for transit use measures decreasing utility as trip time increases, when the wage rate increases, time becomes more costly, making $U(\text{Transit})$ more steep. Also, $U(\text{Driving})$ decreases because time is more costly. The composition of this two effects is ambiguous (because it matters how large the shifts are), but from previous studies, we know empirically that transit is an inferior good [Bresson et al., 2004]. This would imply that an increase in the wage rate would decrease transit ridership, all else equal. We will see if this is empirically confirmed in section 6, which presents the results of data analysis.

Figure 5: Wage Rate and Transit Ridership



Additionally, there are variables that do not directly affect equation 3, but instead affect the number of consumers in the market for public transit, which will affect transit ridership levels. These will be discussed briefly here with the help of comparative statics. First, consider the variable of WP_{it} in equation 1. As the working population increases, the number of consumers in the market for public transit increases (because more people are commuting). This relationship can be expressed in a comparative static, such that $\frac{\delta TR}{\delta WP} > 0$.

Similarly, the unemployment rate, expressed in equation 1 as UR_{it} , is likely to have two conflicting effects, one initial and one long-term, if a higher unemployment rate persists. Initially, a rise in the unemployment rate decreases the demand for public transit, because less people are in the market for commuting. This can be expressed as $\frac{\delta TR}{\delta UR} < 0$. However, if unemployment persists, people might substitute driving for public transit because it

Table 2: Expected Signs of Variables

Variable	Significance	Expected Sign
<i>GP</i>	Gas prices in current month	+
<i>SF</i>	Service Frequency	+
<i>TF</i>	Transit Fare	-
<i>CT</i>	Average Commuting Time	+
<i>WP</i>	Working Population	+
<i>UR</i>	Unemployment Rate	-
ω	Average hourly wage	-
<i>CPI</i>	Consumer Price Index	?

is an inferior good, thereby increasing transit ridership. Finally, there is no clear comparative static for *CPI*, and so we leave the conclusions on this variable to the regression results.

Now that I have established the expected results for each of these variables (summarized in table 2), I will proceed to discuss the techniques used for empirical estimation.

4 Econometric Technique and Methodology

To test my hypothesis, I will employ ordinarily least squares (OLS) fixed-effects multiple regression. I will use a double-log specification because I want the coefficients of the regression to estimate the elasticities of public transit. This technique of fixed-effects is used in order to control for unmeasurable attributes that vary between cities biasing econometric results if they are not addressed. This is the appropriate econometric technique because it directly addresses the empirical question at hand, while avoiding the potential threats to internal validity (these threats are discussed at length in sections 4.1 and 4.2 on page 32).

The use of fixed-effects estimation has particular strengths in estimating elasticities with panel data. Kennedy notes that “the fixed effects estimator is

more robust to selection bias problems...because if the intercepts incorporate selection characteristics they are controlled for in the fixed effects estimation.” (page 290) Additionally, the fixed-effects regression technique is particularly useful when unmeasurable variables differ between the individual entities under study, because the dummy variables that are included in a fixed-effects regression control for those individual-specific variations. This is particularly important in a study on transit ridership, because cities differ in many ways that cannot be measured using conventional data collection; for this reason, I contend that fixed-effects estimation is the appropriate econometric technique for my analysis.

Furthermore, I instrument for service frequency, because of potential simultaneous causality, as transit demand is affected by transit supply (usually measured in service frequency), while at the same time transit authorities might select service frequency based on perceived demand in the next period. Using instrumental variables to predict service frequency would solve the problem of simultaneous causality, if it is present in my data. I propose three possible instruments for service frequency: operators employed by the transit agency, annual scheduled vehicle revenue miles and vehicles available for maximum service. I will discuss the strength of these three instrumental variables in section 6.2. The first variable, operators employed by the transit agency, is an indicator of the maximum number of buses running at any given time. The second possible instrument represents the service frequency as determined before demand is expressed, which neatly unravels the potential for simultaneous causality. Finally, the third variable, vehicles available for maximum service, is a way to measure the extent to which transit authorities can respond to changes in demand.

4.1 Potential Estimation Problems

The empirical literature on the topic of transit demand analysis reveals some potential estimation problems for my thesis question. For each of the poten-

tial problems identified by the literature, I discuss the techniques used in the past to address the problem in addition to explaining whether or not this is a problem in my study and how it could be addressed.

The first problem identified in the literature is deciding on the time span for the gas price elasticity. Because the consumer response to changes in the gas price is not immediate, it is important to decide how long to lag the gas price variable. However, almost no one regresses a lagged gas price variable, because very few studies use longitudinal data. In the literature that I surveyed, only Mattson [2008] addressed any sort of a lagged gas price variable. Using the Akaike Information Criterion, Mattson uses two lags of the gas price variable, implying that a two-year lag is appropriate. However, because his data is yearly, he does not capture the immediate response to a rise in gas prices that occurs on a monthly basis, because consumers have some flexibility to switch between transit and driving. Also, the regressions in which he included a lagged gas price variable did not include any other variables, while the regressions in which he included the other important variables, he did not include a lagged gas price variable. The reasoning behind this decision was incredibly unclear, and because no other papers specifically addressed including a lagged variable, I have no theoretical background on which to base my decisions. As mentioned in section 3.1, I will include the following gas price lags: current, one-month, three-month and six-month lags. This measures current, short-term, medium-term and long-term changes in the reaction to gas prices changes.

These lags are more troublesome than they first appear, for a number of reasons. First, gas prices are highly cyclical, and usually rise throughout the summer and then fall again in the winter. At the same time, as gas prices are falling, more people are working, driving ridership up. Also, lagging the gas price from three months earlier may not be the best measure of gas prices, since it has been shown that as gas prices drop, people resume their normal transportation habits. After the surge of ridership in the summer

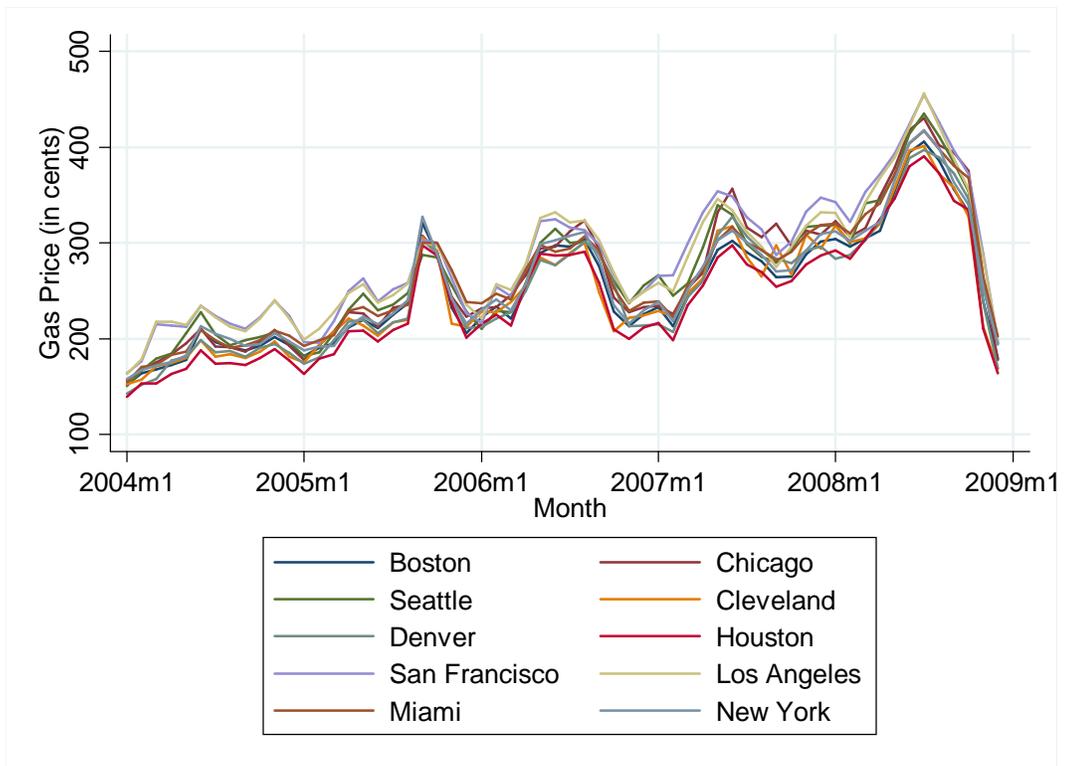
of 2008, ridership numbers fell back to average levels in the fall and winter. To illustrate the cyclical nature of gas prices, I have added figure 6, which describes the phenomenon I am addressing. As is clear from the graph, starting in May of each year, gas prices begin to climb, and then they drop off in September. This is true in every city of my sample. For this reason, while I include lags on gas prices, I regard the coefficients on lagged gas prices, at least in the short- and medium-term with a degree of caution, since gas prices are so volatile from month-to-month. The long-term lag is more informative, because it captures the overall trend of gas prices, not just short-term fluctuations. However, there is a potential advantage in including lagged gas prices in the regressions, because it allows us to control for past gas prices. This enables me to isolate the effect of current prices on transit ridership, potentially addressing the cyclical nature of gas prices.

For further work, which I discuss in my conclusion, I suggest incorporating a moving average of gas prices over the last three months, or including a standard deviation of gas prices over the past year, in order to measure volatility, since this might affect people's decision to take public transit.

The second problem that I have to address is the measurement of transit fare, because there are many different types of transit fare for a given transit authority. For this reason, I have to look at each transit authority, deciding which fare is closest to the specification of an adult single-trip bus fare. While I realize that many riders use monthly bus passes, these prices are closely correlated with the price level of a single-trip bus fare; as the fare goes up, so does the cost of a monthly bus pass. For this reason, it is likely that the price of riding transit (for the average consumer) might be slightly overstated in the analysis, leading to an underestimation of the price elasticity of transit demand in my regression results.

The final problem that I encounter in my analysis is the inclusion of seasonal dummies. While there is clearly a correlation between transit ridership and the seasons, because of the costs of using public transit to commute are

Figure 6: Gas Prices by City over Time



directly affected by the weather, there is a stronger correlation between the seasonal dummies and gas prices. However, because the level of inclement weather is different in each city, it is likely that fixed-effects estimation actually controls for this seasonal cyclicalilty already. For this reason, I do not include seasonal dummies in my analysis.

4.2 Other Threats to Internal Validity

In this section, I address the common threats to internal validity for an econometric study, and discuss the solutions that I have crafted.

First, many empirical analyses are threatened by omitted variable bias. While almost every empirical analysis is prone to some amount of omitted variable bias (because of the costs of collecting data), I believe that because I use longitudinal data, my empirical analysis suffers from the problem of omitted variable bias to a lesser degree. Also, I attempt to include the variables that other transit demand studies have incorporated; therefore at the very least, I suffer no more omitted variable bias than similar studies.

Another problem that always threatens the internal validity of a study is the misspecification of functional form. I argue that because I use OLS fixed-effects, my methodology is better crafted for estimating transit elasticities, because I look at how variables change in a particular city, not across cities. While this is not the dominant econometric methodology in the literature, I contend throughout this paper that the predominant methodology in the literature simply measures the wrong variations in variables, because most studies use cross-sectional analysis.

Additionally, errors-in-variables threatens the internal validity of an empirical study. However, as I noted on page 34, my data set is actually stronger, because I use city-specific data for gas prices, as well as revised transit ridership numbers. However, as Thompson and Brown [2006] note, these ridership numbers are skewed because they measure unlinked transit passenger trips,²

²This language is taken directly from the APTA ridership reports [APTA, 2008]

as opposed to linked transit passenger trips, a much more accurate measure of transit ridership levels. Because the transit authorities do not have a way to easily measure linked transit passenger trips, these numbers do have errors-in-measurement. Although it is not clear how large this effect is, we know that because of this measurement error, transit ridership numbers are biased upwards; this means that any elasticity measure found might overstate the actual gain in transit ridership, which should lead us to interpret the coefficients from our regression results with caution. Alternatively, we can interpret the coefficient as the percentage change in unlinked passenger trips given a percentage change in gas prices.

Thirdly, because the ten cities I selected as my sample of urbanized areas are all located in large metropolitan areas, my empirical estimations suffer from some sample selection bias. However, the purpose of this study is to provide an estimation of gas price elasticities for transit authorities to use in decision-making. With this aim in mind, I have selected cities that are a cross-section of cities in the United States, as they include almost every mega-region (this is discussed further in section 5.1). Therefore, while sample selection bias is clearly present, I do not believe that it will threaten the validity of the conclusions of my study, because as Bureau notes, using fixed-effects estimation eliminates much of the sample selection bias present (this is discussed in more depth in the previous section).

Finally, many studies of consumer demand suffer from simultaneous causality, because of the identification problem that persists in empirical analysis. Particularly for transit demand, where service levels play a large role in determining the costs of using public transit, simultaneous causality might play a large role. To address this problem, in addition to using OLS fixed-effects regressions, I am going to attempt to instrument for service frequency, and see if this significantly changes my estimations of elasticities, or increases the magnitude of my adjusted R-squared or F-statistic values. This will be discussed more in depth in section 6.2.

Table 3: Data Collection

Variable	Significance	Data Source
<i>TR</i>	Transit Ridership	American Public Transit Authority Ridership Reports
<i>GP</i>	Gas Prices (monthly)	U.S. Department of Energy
<i>TF</i>	Transit Fare	Individual Transit Authorities
<i>SF</i>	Service Frequency	National Transit Database
<i>CT</i>	Average Commuting Time	Census Bureau American Community Survey
<i>WP</i>	Working Population	Bureau of Labor Statistics
<i>UR</i>	Unemployment Rate	Bureau of Labor Statistics
ω	Average Hourly Wage	Bureau of Labor Statistics
<i>CPI</i>	Core Price Index (at the city level)	Bureau of Labor Statistics

5 Discussion of the Data Set

5.1 Variables and Data Sources

The data for the empirical study comes from a variety of sources, and Table 3 on page 34 summarizes the source used to collect each variable. One of the particular strengths of my data, as opposed to similar empirical studies, is that the gas price variable is much more accurate. Most studies use regional gas price numbers; however, for each of the ten cities in my study, I have collected city-level gas price data. Because my data is more accurate than past studies, my study suffers from less errors-in-variables bias, adding to the internal validity of my study compared with other similar ones. Also, taking the advice of Mattson [2008] (noted earlier on page 12), I use revised ridership numbers in my analysis when they are available, since the revisions published by the APTA sometimes differ considerably from originally published ridership figures. This also enhances the internal validity of my empirical analysis.

Table 4: City Corresponding Transit Authorities

City	Transit authority
Boston	Massachusetts Bay Transportation Authority
Chicago	Chicago Transit Authority
Cleveland	The Greater Cleveland Regional Transit Authority
Denver	Denver Regional Transportation District
Houston	Metropolitan Transit Authority of Harris County, Texas
Los Angeles	Los Angeles County Metropolitan Transportation Authority
Miami	Miami-Dade Transit
New York	MTA New York City Transit
San Francisco	San Francisco Municipal Railway
Seattle	King County Department of Transportation - Metro Transit Division

5.2 Structure of the Dataset

In this section, I will discuss the structure of my dataset, as well as certain peculiarities that must be pointed out to understand the data more fully.

My data span the years from 2004 to 2008, containing monthly observations. The data is collected on a monthly basis when available, and if not, are collected on a yearly basis. My data contain ten cities, representing almost every mega-region in the United States (only the Piedmont Atlantic and Arizona Sun Corridor mega-regions are omitted). The ten cities are Boston, Chicago, Cleveland, Denver, Houston, Los Angeles, Miami, New York, San Francisco and Seattle. These cities were selected because the most comprehensive gas price data was available for them through the U.S. Department of Energy, making measurement of gas prices much more accurate. Table 4 presents each city and the transit authority that I use for this research.

During my data collection, I found that seven data points for transit ridership were missing: Boston in December 2004 and the first quarter of 2006, and in Miami in the third quarter of 2004. For whatever reason, this data simply does not exist in any database in the APTA. For this reason, these months are not included in the regressions.

Also, I constructed the gas price variable in the following manner. For each month, I chose the first week of reporting for that month. For instance, if the month is October 2004, then the gas price for that month is the gas price reported for October 4, 2004. In this way, the gas price variable captures the decision that each potential transit rider must make at the beginning of every period, comparing the costs of transit to the costs of driving, according to the theory of binary modal choice, outlined in section 3.2 on page 20. Also, the start of the month is when a potential transit rider is more likely to buy a monthly transit pass, and so gas prices in the first week of a month are most to impact transit ridership.

Finally, I measure service frequency in the following manner. First, the National Transit Database has a yearly statistic for each transit authority that is called Vehicles Operated in Maximum Service, and is a measure of the capacity that a transit authority has to operate. Second, the NTD has another yearly statistic called Vehicle Revenue Miles, which is a measure of how much coverage a transit authority has. If these variables increase year-on-year, it can be implied that more buses were running, or alternatively that buses were running more, which increases service frequency and decreases access time, the effects of which are discussed in sections 3.2 and 3.3. Instrumenting for these service frequency variables is discussed in section 6.2.

6 Data Analysis

In this section, I use the regression specification that I introduced in section 3.1 on page 19. Section 6.1 presents the results of the regressions using both service frequency variables as they are, while in section 6.2, I present the results of my regressions having instrumented for both service frequency variables.

6.1 Regressions without Instrumental Variables

In this section, I present and discuss my regression results that do not include instrumental variables. The tables are presented below. Table 5 presents the regression results that use Vehicles Operated in Maximum Service as the service frequency variable, while Table 6 presents the regression results using Vehicle Revenue Miles as the service frequency variable.

Note that in every regression, using Vehicle Revenue Miles yields a higher adjusted R-squared value. Because using VRM as the service frequency variable consistently leads to better estimation, the remainder of the discussion will pertain to Table 6. Notice that as past gas prices are controlled for, the current gas price variable increases in significance and magnitude. What I believe has occurred in controlling for past gas prices is that I have addressed the problem of cyclicity in gas prices, discussed in section 4.1 on page 28. This is because the coefficient in the fourth regression in Table 6 is the elasticity of transit ridership with respect to current period gas prices, *holding past gas prices constant*.

Table 6 also shows that past gas prices are not significant in determining current transit ridership, as the coefficient decreases in magnitude and significance. This is consistent with economic theory, because as the time horizon extends, more substitutes are available, causing the cross-elasticity between substitutes to tend towards zero. This is the case in the current scenario for a number of reasons. First, as gas prices increase, people may initially take public transit more, but then after a while either organize carpools, substitute to more gas-efficient vehicles or locate closer to their jobs. If this is the case, we can talk about public transit as a short-term substitute for driving, but not a substitute that is seen as a viable long-term option.

The coefficient on gas prices in my regression results is generally consistent with the literature, falling in the range of past studies, as summarized in table 1 on page 16. While my estimates do not differ much from the literature, I believe that my estimations have more internal validity because I have a

Table 5: Regression Results Using VOMS for Service Frequency

	(1)	(2)	(3)	(4)
LN(GAS PRICE)	0.122** (0.0447)	0.133*** (0.0332)	0.152** (0.0500)	0.145** (0.0544)
LN(GAS PRICE_1)		-0.0191 (0.0408)	-0.0435 (0.0263)	-0.0440 (0.0259)
LN(GAS PRICE_3)			0.0319 (0.0538)	0.0356 (0.0494)
LN(GAS PRICE_6)				-0.0184 (0.0319)
LN(TRANSIT FARE)	-0.0000473 (0.135)	0.00201 (0.135)	-0.00306 (0.137)	0.00186 (0.137)
LN(COMMUTE TIME)	0.234 (0.566)	0.241 (0.565)	0.222 (0.574)	0.237 (0.561)
LN(EMPLOYMENT)	0.849 (1.036)	0.871 (1.020)	0.845 (1.003)	0.848 (0.999)
LN(WAGE)	-0.445 (0.259)	-0.453 (0.265)	-0.438 (0.279)	-0.440 (0.279)
LN(CPI)	0.563 (0.563)	0.596 (0.610)	0.491 (0.748)	0.578 (0.813)
LN(SERVICE FREQUENCY)	-0.00790 (0.0264)	-0.00845 (0.0256)	-0.00775 (0.0262)	-0.00821 (0.0257)
UNEMPLOYMENT RATE	-0.00374 (0.0140)	-0.00319 (0.0136)	-0.00314 (0.0135)	-0.00381 (0.0132)
Observations	593	593	593	593
Adjusted R-squared	0.366	0.365	0.365	0.364

Note: *** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table 6: Regression Results Using VRM for Service Frequency

	(1)	(2)	(3)	(4)
LN(GAS PRICE)	0.139** (0.0447)	0.135*** (0.0319)	0.170*** (0.0518)	0.172** (0.0604)
LN(GAS PRICE_1)		0.00808 (0.0456)	-0.0357 (0.0253)	-0.0355 (0.0250)
LN(GAS PRICE_3)			0.0594 (0.0612)	0.0583 (0.0557)
LN(GAS PRICE_6)				0.00641 (0.0433)
LN(TRANSIT FARE)	0.00975 (0.122)	0.00891 (0.122)	-0.000272 (0.122)	-0.00199 (0.122)
LN(COMMUTE TIME)	0.425 (0.614)	0.423 (0.612)	0.400 (0.616)	0.396 (0.603)
LN(EMPLOYMENT)	0.568 (0.801)	0.558 (0.774)	0.492 (0.744)	0.489 (0.743)
LN(WAGE)	-0.397 (0.225)	-0.394 (0.230)	-0.363 (0.246)	-0.362 (0.247)
LN(CPI)	0.356 (0.599)	0.341 (0.656)	0.128 (0.825)	0.0954 (0.944)
LN(SERVICE FREQUENCY)	0.291 (0.207)	0.293 (0.210)	0.312 (0.214)	0.315 (0.220)
UNEMPLOYMENT RATE	-0.0112 (0.00956)	-0.0114 (0.00915)	-0.0118 (0.00890)	-0.0116 (0.00835)
Observations	593	593	593	593
Adjusted R-squared	0.388	0.387	0.389	0.388

Note: *** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

more appropriate regression specification. Since I have chosen to perform an OLS fixed-effects regression on panel data, in contrast to almost the entire rest of the literature that looks at cross-sectional data, I have controlled for a number of things that were previously captured in the gas price variable. For instance, much of the variation in gas prices that was seen in cross-sectional analyses is explained by supply-chain differences, market structure in a city, as well as city demand for gasoline. However, as is seen in figure 6 on page 31, gas prices do not vary widely from city to city at a given point in time. Therefore, it is more informative to see how many people substitute into and out of the public transit market in a certain city as gas prices vary *in that city* through time. In fact, Taylor et al. [2009] note that this was a problem in their estimation of gas price elasticity, saying “[The marginal significance of gas prices] is likely due to the relatively low levels of variation of average fuel prices between one urbanized area and another.” (page 69) Not only is my regression specification more useful to estimate the true elasticity of transit ridership with respect to gas prices, but it is also more valuable to individual transit authorities, since they are not as interested in the impact that inter-city gas price differences have on transit ridership, but instead the effects of intra-city gas price changes. Therefore, my choice of functional form greatly increases the internal and external validity of my econometric analysis.

Also in Table 6, notice that the elasticity of transit ridership with respect to wage rates is quite large in magnitude and negative, implying that public transit is an inferior good. This is in line with the values that other studies on the determinants of transit demand have found, and so gives me more confidence in the validity of my regression results.

Furthermore, note how large the elasticity of transit ridership with respect to commuting time is. Because this variable is a measurement of the cost of driving, clearly as the commute time increases, people substitute to use public transit. This effect is actually larger in magnitude than the effect of

gas prices increasing because time costs are a larger proportion of the costs of driving.

On a final note, it seems that transit fare is not a significant determinant of transit ridership; not only is it not statistically significant, but the magnitude of the coefficient is incredibly small. This is not particularly surprising, since transit fare is a very small proportion of the total cost of riding public transit (although it is the only explicit cost). Additionally, there is little variation of transit fare over time. In my dataset, over five years, only three cities changed their base transit fare twice (Cleveland, Denver and Miami), while one (New York) did not change their transit fare at all. The remaining transit authorities only changed their fare once. In doing the research for this project, I noticed that for a number of cities, most notably Chicago, fare increases would be announced for a future date and then later rescinded, sometimes just days before the change was to take place, as the transit authority received a one-time grant from the state or federal government. In the case of Chicago, this happened on at least four different occasions, but overall their fares only increased once over the five years included in my data set. This persistence of fare levels means that there is not much variation from which to draw statistical inference. It would be more useful to study ridership levels just before and just after fare changes, holding all else equal; however, that is outside of the scope of the current research.

6.2 Regressions with Instrumental Variables

In previous sections, I have discussed using instrumental variables to improve the internal validity of my research, in order to untangle the possible simultaneous causality that is present between service frequency and transit ridership. In this section, I go into more depth as to the variables being used, the rationale for each of them, and the strength of these variables.

The first instrumental variable that I use is vehicles available for maximum service, hereafter referred to as VAMS. I use this variable to predict

vehicles operated on maximum service, since VAMS is an upper constraint on how much transit authorities can expand service in a given year. Also, while it is likely correlated with transit demand in past periods, it is not likely to change through a year in response to current demand because of the time it takes to secure the financing for additional buses or light-rail cars. For this reason, VAMS is not directly correlated with current transit ridership at the city level, making it a good potential instrument. Also, as summarized in Table 7, the F-statistic for VAMS was 295.42, showing it to be a strong instrument for VOMS.

The second instrumental variable that I use is annual scheduled vehicle revenue miles, hereafter referred to as ASVRM. This is a particularly strong instrumental variable because it is most definitely not correlated with unexpected changes in demand, because ASVRM measures the revenue miles that would be incurred if the schedule were not changed from the beginning of the year through the end. Since transit authorities generally adjust their schedules to accommodate demand, if vehicle revenue miles are predicted using ASVRM, instrumenting for this variable ensures that there is not simultaneous causality present in my regression model. When my regressions use VRM instead of predicted VRM, there was a possibility for simultaneous causality, since VRM measures the responses that transit authorities make throughout the year to changes in transit demand. My regression predicting VRM confirmed that ASVRM is a strong instrument, with a F-statistic of 26.87, which is seen in Table 8.

My third instrumental variable is operators on payroll in each year. The rationale behind this IV is that the number of operators on payroll is a constraint on the transit authority to adjust quickly to changes. However, this measure includes both full- and part-time operators, which means it is not as accurate as desired. My regressions predicting VOMS and VRM using operators on payroll have F-statistics of 9.73 and 0.00; for this reason, I do not include the results of regressions including the predicted service frequency

utilizing operators on payrolls, as it is not a strong enough instrument.

6.2.1 Regression Results using Instrumental Variables

The results of the regressions using VAMS and ASVRM are in Tables 7 and 8, respectively.

Notice particularly, in comparing the regression results using VOMS without an instrument (Table 5 on page 38) and Table 7, which shows results of VOMS being instrumented for with VAMS, that the adjusted R-squared does not improve at all. For this reason, I believe that using VAMS as an instrument does not improve my estimations of transit elasticities.

However, compare the results between Tables 6 and 8, which show regressions including VRM without and with an instrument. The estimations using ASVRM as an instrumental variable for VRM improve the adjusted R-squared by 0.019 in each regression. Also, notice that the coefficients on service frequency have almost tripled, since I have addressed the simultaneous causality present in the regressions. Also, the coefficients on gas prices have decreased slightly. Now my results estimate the elasticity of transit ridership with respect to gas prices to be approximately 0.162 when controlling for past months of gas prices. This implies that if current gas prices increase ten percent, transit ridership, measured by unlinked passenger trips, will increase approximately 1.6%.

Furthermore, it seems that public transit is still estimated to be an inferior good, as the coefficient on wage rates has remained negative, though decreasing in magnitude. Finally, as in the non-instrumented regression, transit fare is never significant, because it is not a significant percentage of the cost of riding public transit. However, the magnitude of the coefficients on transit fare have increased by an order of magnitude, suggesting that instrumenting for service frequency also affects the estimation of transit fare elasticity. Also, as noted above in section 6.1 on page 37, fares do not vary much over time, which does not give us enough variation to accurately estimate the fare

Table 7: Regression Results Using VAMS as IV

	(1)	(2)	(3)	(4)
LN(GAS PRICE)	0.123** (0.0439)	0.133*** (0.0326)	0.152** (0.0495)	0.145** (0.0541)
LN(GAS PRICE_1)		-0.0178 (0.0405)	-0.0426 (0.0257)	-0.0431 (0.0253)
LN(GAS PRICE_3)			0.0325 (0.0541)	0.0361 (0.0497)
LN(GAS PRICE_6)				-0.0179 (0.0322)
LN(TRANSIT FARE)	0.00183 (0.134)	0.00375 (0.135)	-0.00144 (0.137)	0.00335 (0.137)
LN(COMMUTE TIME)	0.229 (0.569)	0.236 (0.569)	0.217 (0.578)	0.231 (0.566)
LN(EMPLOYMENT)	0.820 (1.025)	0.840 (1.010)	0.814 (0.995)	0.817 (0.991)
LN(WAGE)	-0.438 (0.257)	-0.445 (0.263)	-0.430 (0.277)	-0.432 (0.277)
LN(CPI)	0.545 (0.560)	0.577 (0.606)	0.469 (0.745)	0.554 (0.810)
LN(SERVICE FREQUENCY) (Predicted)	0.00720 (0.0324)	0.00672 (0.0318)	0.00714 (0.0321)	0.00689 (0.0319)
UNEMPLOYMENT RATE	-0.00476 (0.0132)	-0.00425 (0.0128)	-0.00418 (0.0128)	-0.00484 (0.0125)
Observations	593	593	593	593
Adjusted R-squared	0.366	0.365	0.365	0.364

Note: *** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table 8: Regression Results Using ASVRM as IV

	(1)	(2)	(3)	(4)
LN(GAS PRICE)	0.128** (0.0463)	0.133*** (0.0334)	0.163** (0.0519)	0.162** (0.0566)
LN(GAS PRICE_1)		-0.00740 (0.0420)	-0.0469 (0.0256)	-0.0469* (0.0252)
LN(GAS PRICE_3)			0.0518 (0.0554)	0.0522 (0.0512)
LN(GAS PRICE_6)				-0.00221 (0.0339)
LN(TRANSIT FARE)	-0.0534 (0.107)	-0.0525 (0.107)	-0.0622 (0.107)	-0.0615 (0.107)
LN(COMMUTE TIME)	0.584 (0.566)	0.587 (0.564)	0.564 (0.567)	0.566 (0.556)
LN(EMPLOYMENT)	0.727 (0.822)	0.735 (0.800)	0.693 (0.771)	0.694 (0.770)
LN(WAGE)	-0.171 (0.237)	-0.174 (0.247)	-0.144 (0.272)	-0.144 (0.275)
LN(CPI)	0.509 (0.620)	0.522 (0.672)	0.349 (0.815)	0.360 (0.877)
LN(SERVICE FREQUENCY) (Predicted)	0.760 (0.424)	0.759 (0.426)	0.777 (0.435)	0.776 (0.442)
UNEMPLOYMENT RATE	-0.00678 (0.0107)	-0.00658 (0.0102)	-0.00648 (0.00985)	-0.00656 (0.00975)
Observations	593	593	593	593
Adjusted R-squared	0.407	0.406	0.408	0.407

Note: *** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

elasticity for public transit.

7 Conclusion

This paper provides an econometric framework in order to test the extent to which changes in gas prices affect transit ridership levels. Also included in the regressions are transit fare, average hourly wage, service frequency and average commuting time, among others. To test the hypothesis, I employed an OLS fixed-effects regression, in order to control for potentially unmeasurable variations between cities. My use of a OLS fixed-effects regression to estimate transit elasticities is relatively groundbreaking, as only two other papers in the literature have done a study on transit demand that incorporates panel data. Additionally, in order to ensure that I addressed the threat of simultaneous causality, I instrumented for service frequency using vehicles available for maximum service as well as annual scheduled vehicle revenue miles.

I find that the elasticity of transit ridership with respect to gas prices is about 0.162, which accords with the literature. However, my estimate is more statistically sound, because my methodology is intentionally crafted to estimate changes in gas prices correctly, by controlling for city characteristics, allowing me to look at gas price changes at the city level instead of across cities. I argue that the estimation of elasticities for gas prices using panel data and fixed-effects regression accurately captures the response of transit ridership to changes in gas prices through time, controlling for city attributes.

I also find that vehicle revenue miles are a more accurate measure for service frequency, which accords with earlier research on the subject, as the majority of papers use vehicle revenue miles as their transit supply variable. Furthermore, I find that if vehicle revenue miles are instrumented for, by using annual scheduled vehicle revenue miles, my results have considerably more explanatory power, accounting for two percent more of the variance in

the data. I contend that the use of instrumental variables improves the estimations of elasticities, as measured by an increased adjusted R-squared value when using instrumental variables. Annual scheduled vehicle revenue miles is a particularly strong instrument, because the transit authority decides on this value before the year begins, and so it is not correlated with current changes in ridership, addressing the possibility of simultaneous causality.

This paper leaves many topics for future transit demand research. I propose a number of extensions to this research. First, it would be useful to see if the estimates obtained in the course of this research hold if the dataset were expanded to include all of the urbanized areas in the United States, or at least a larger sample of UZAs. Second, incorporating a moving average of gas prices over time, possibly using the past three months, would probably depict the decision-making of individuals in choosing to ride public transit more accurately. Third, incorporating a measure of the recent variation in gas prices would enhance the analysis, since as gas prices become more volatile, risk-averse agents will increasingly opt to take public transportation. Finally, I hope to find some way to control for the seasonality of public transit, while avoiding correlation with the cyclical nature of gas prices. This final extension to the research proves to be the toughest, and solving it would allow an even more accurate estimation of the elasticity of transit ridership with respect to gas prices.

References

San Francisco Municipal Transportation Agency. Transit fare data, Received March 25, 2010.

American Public Transportation Association. Ridership reports, 2004-2008. Technical report.

- American Public Transportation Association. Transit ridership report. Technical report, 2008.
- Chicago Transit Authority. Reports and notices - transit fare data. Technical report, 2008a.
- Greater Cleveland Regional Transit Authority. Rta history - fare data, 2008b.
- Los Angeles County Metropolitan Transportation Authority. Transit fare data, Received March 24, 2010.
- Massachusetts Bay Transportation Authority. Transit fare data, Received March 23, 2010.
- Metropolitan Transportation Authority. Transit fare data, Received on March 18, 2010.
- Antonio Bento, Maureen L. Cropper, Ahmed Mushfiq Mobarak, and Katja Vinha. *The Impact of Urban Spatial Structure on Travel Demand in the United States*. The World Bank, Policy Research Working Paper Series: 3007, 2003.
- Georges Bresson, Joyce Dargay, Jean-Loup Madre, and Alain Pirotte. Economic and structural determinants of the demand for public transport: an analysis on a panel of french urban areas using shrinkage estimators. *Transportation Research Part A: Policy and Practice*, 38(4):269–285, 5 2004.
- U.S. Census Bureau. American community survey, 2004-2008. Technical report.
- National Transit Database. Annual transit databases, 2004-2008. Technical report.

- Gerard de Jong and Hugh Gunn. Recent evidence on car cost and time elasticities of travel demand in europe. *Journal of Transport Economics and Policy*, 35(2):137–160, 05 2001.
- Regional Transportation District. Transit fare data, Received March 23, 2010.
- P. B. Goodwin. A review of new demand elasticities with special reference to short and long run effects of price changes. *Journal of Transport Economics and Policy*, 26(2):155–169, 05 1992.
- David A. Hensher. *Bus Transport: Economics, Policy and Planning*. Research in Transportation Economics series, vol. 18; Amsterdam and Boston.; Elsevier, JAI Press, 2007.
- Johan Holmgren. Meta-analysis of public transport demand. *Transportation Research Part A: Policy and Practice*, 41(10):1021–1035, 12 2007.
- John F. Kain and Zvi Liu. Secrets of success: assessing the large increases in transit ridership achieved by houston and san diego transit providers. *Transportation Research Part A: Policy and Practice*, 33(7-8):601–624, 11 1999.
- Harold Kohn. Factors affecting urban transit ridership. *Statistics Canada Catalogue*, From the Bridging the Gaps Conference, Canadian Transportation Forum, Charlottetown, Prince Edward Island, June 6, 2000.
- Michael Kuby, Anthony Barranda, and Christopher Upchurch. Factors influencing light-rail station boardings in the united states. *Transportation Research Part A: Policy and Practice*, 38(3):223–247, 3 2004.
- Michael Kyte, James Stoner, and Jonathan Cryer. A time-series analysis of public transit ridership in portland, oregon, 1971-1982. *Transportation Research Part A: General*, 22(5):345–359, 9 1988.

- Todd Litman. Transit price elasticities and cross-elasticities. *Journal of Public Transportation*, 7(2):37–58, 2004.
- Donnie Maley and Rachel Weinberger. Does gas price fuel transit ridership? *Panorama*, 4:17–21, 2009.
- Jeremy Mattson. Effects of rising gas prices on bus ridership for small urban and rural transit systems. *Small Urban and Rural Transit Center, Upper Great Plains Transportation Institute, North Dakota State University*, 2008.
- Daniel McFadden. The measurement of urban travel demand. *Journal of Public Economics*, 3(4):303–328, 11 1974.
- Texas Metropolitan Transit Authority of Harris County, Houston. Metro facts & history: Fare data, 2008.
- U.S. Department of Energy. U.s. retail gasoline historical prices, 2004-2008. Technical report.
- U.S. Bureau of Labor Statistics. Local area consumer price index, 2004-2008. Technical report.
- U.S. Bureau of Labor Statistics. National compensation survey. Technical report, 2004-2009.
- U.S. Bureau of Labor Statistics. Local area unemployment statistics, 2004-2008. Technical report, Data extracted February 10, 2010.
- Tae Hoon Oum, II Waters, W.G., and Jong-Say Yong. Concepts of price elasticities of transport demand and recent empirical estimates. *Journal of Transport Economics and Policy*, 26(2):139–154, 05 1992.
- James H. Stock and Mark W. Watson. *Introduction to econometrics*. Addison Wesley, Boston, MA, 2003.

Brian D. Taylor, Douglas Miller, Hiroyuki Iseki, and Camille Fink. Nature and/or nurture? analyzing the determinants of transit ridership across us urbanized areas. *Transportation Research: Part A: Policy and Practice*, 43(1):60–77, 01 2009.

Gregory L. Thompson and Jeffrey R. Brown. Explaining variation in transit ridership in us metropolitan areas between 1990 and 2000: Multivariate analysis. *Transportation Research Record: Journal of the Transportation Research Board*, 1986(1):172–181, 2006.

King County Metro Transit. Annual management reports, 2004-2008. Technical report.

Miami-Dade Transit. Press releases - transit fare data, 2008.

George H. K. Wang and David Skinner. The impact of fare and gasoline price changes on monthly transit ridership: Empirical evidence from seven u.s. transit authorities. *Transportation Research Part B: Methodological*, 18(1):29–41, 2 1984.