

NBER WORKING PAPER SERIES

“THE WAR FOR THE FARE”:
HOW DRIVER COMPENSATION AFFECTS
BUS SYSTEM PERFORMANCE

Ryan M. Johnson
David H. Reiley
Juan Carlos Muñoz

Working Paper 11744
<http://www.nber.org/papers/w11744>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2005

Johnson: Department of Economics, University of Arizona, <RyanJ@email.arizona.edu>. Reiley: Department of Economics, University of Arizona, <Reiley@eller.arizona.edu>. Muñoz: Department of Transportation Engineering, Pontificia Universidad Católica (Santiago, Chile), <jcm@ing.puc.cl>. Johnson acknowledges the United Nations Economic Commission for Latin America and the Caribbean for an internship enabling him to work on this project, and Ian Thomson and Sidney Rezende for their individual assistance. He also acknowledges funding from the University of Arizona Honors College for a grant used to fund this study. Reiley gratefully acknowledges support from National Science Foundation grant SES-0094800. Muñoz acknowledges the support of the Chilean Fund for Scientific and Technological Development (FONDECYT) through Project 1040604 and Fundación Andes. We also thank Manuela Angelucci, Oriana Bandiera, Jim Cox, Greg Crawford, Kei Hirano, Tom Hubbard, Mike Urbancic, and seminar participants at the University of Cologne and the European University Institute for their helpful suggestions. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Ryan M. Johnson, David H. Reiley, and Juan Carlos Muñoz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

“The War for the Fare”: How Driver Compensation Affects Bus System Performance
Ryan M. Johnson, David H. Reiley, and Juan Carlos Muñoz
NBER Working Paper No. 11744
November 2005
JEL No. D0, J3, L9

ABSTRACT

Two systems of bus driver compensation exist in Santiago, Chile. Most drivers are paid per passenger transported, while a second system compensates other drivers with a fixed wage. Compared with fixed-wage drivers, per-passenger drivers have incentives to engage in “La Guerra por el Boleto” (“The War for the Fare”), in which drivers change their driving patterns to compete for passengers. This paper takes advantage of a natural experiment provided by the coexistence of these two compensation schemes on similar routes in the same city. Using data on intervals between bus arrivals, we find that the fixed-wage contract leads to more bunching of buses, and hence longer average passenger wait times. The per-passenger drivers are assisted by a group of independent information intermediaries called “sapos” who earn their living by standing at bus stops, recording arrival times, and selling the information to subsequent drivers who drive past. We find that a typical bus passenger in Santiago waits roughly 10% longer for a bus on a fixed-wage route relative to an incentive-contract route. However, the incentives also lead drivers to drive noticeably more aggressively, causing approximately 67% more accidents per kilometer driven. Our results have implications for the design of incentives in public transportation systems.

Ryan M. Johnson
Department of Economics
401 McClelland Hall
University of Arizona
Tucson, AZ 85718
ryanj@email.arizona.edu

Juan Carlos Muñoz
Department of Transportation and Engineering
Pontificia Universidad Católica
Santiago, Chile
jcm@ing.puc.cl

David H. Reiley
Department of Economics
401 McClelland Hall
University of Arizona
Tucson, AZ 85718
and NBER
reiley@eller.arizona.edu

1. Introduction

While riding a public bus in Santiago, Chile, one of us noticed an interesting phenomenon. A man with a notepad got on the bus and yelled some numbers at the driver. The driver then gave the man a coin, and the man got back off the bus. Careful subsequent observation indicated that a number of similar men, called *sapos*,² likewise stood at bus stops in Santiago for hours at a time, recording the arrival times of buses on a notepad, and selling their data to drivers.

This presented a puzzle. Why would bus drivers want to pay for information about the timing of other buses? It turns out that, unlike the typical system in the United States, many Santiago bus drivers receive compensation based on their passenger receipts.³ Such drivers therefore have an incentive to drive in a way that maximizes the number of passengers they transport. This depends significantly on the time interval (called *headway* in the transportation literature) between their bus and the bus immediately ahead on the same route. If the bus in front is far ahead, many passengers will have accumulated since the last bus came by, thus providing high profits for the driver. By contrast, if the bus in front is very close, then the driver can expect to be picking up few passengers and low profits.

For their part, the *sapos* provide valuable headway information to the drivers. Each time a new bus arrives, a *sapo* marks the minute on his notepad, as well as

² Literally, “frogs;” to our knowledge there is no formal documentation of how the name started. The term *sapo* is a derogatory term in Chile. In Peru, they are known, perhaps more appropriately, as *dateros* (data men). so named because they jump on and off buses all day.

³ At the time of the study, the fare cost 310 Chilean pesos, or roughly 50 cents American. Drivers earn 10-20% of this fare under their incentive compensation.

telling the driver (for a fee) his headway, in minutes, with the immediately preceding bus. Given this information, the driver can choose to drive somewhat faster or slower in order to create a more profitable spacing. For example, if the typical headway on a route is 10 minutes, but a driver has gone slowly enough to allow that headway to grow to 20 minutes, more passengers will be waiting and the driver will make more money. However, the bus behind that driver will then have a short headway, thus giving that second driver a strong incentive to change the spacing. Unlike drivers paid a fixed hourly wage, drivers receiving per-passenger compensation play a strategic game with each other, changing their driving in order to maximize profits given other drivers' behavior. In Santiago, this game is commonly known as "*La Guerra por el Boleto*", or "The War for the Fare."

What are the effects of this game on passenger welfare? First, let's consider a common complaint about bus systems. Passengers of many bus systems worldwide complain about the spacing of buses. A typical complaint, drawn from an Internet message board for the city of New York, is "I remember many mornings of waiting for the bus only to have 4 or 5 buses arrive right behind each other" (Rider Diaries, 2005). A transit operator for the city of Chicago indicates that "bus bunching is the number one complaint that I hear from our customers" ("*CTA Expands Efforts to Reduce Bus Bunching*," 2000). Passenger advocate groups in Chicago and New York have conducted studies of bus bunching, respectively claiming 40% ("*The Late State of Buses*," 2004) and 60% (Pearson, 2003) of the buses arrived bunched on the routes they examined. Such groups often accuse bus dispatchers of incompetence,

saying that they allow such bunching to occur, but in fact there are mathematical reasons why, even if dispatched precisely evenly, buses should evolve towards a state of bunching (Eastaway, 1998).

Given these common spacing problems, it seems possible that a per-passenger compensation system, by giving drivers incentives to monitor and change their spacing, could actually improve passenger welfare. Drivers, by using their discretion to make longer or shorter stops, or to drive faster or slower, might actually correct a natural instability of the system, and improve the regularity of arrivals of buses for passengers. Whether this occurs in practice is the central question of our paper.⁴

Absent detailed knowledge about the drivers' control variables, we find it impractical to construct a precise mathematical model of the game between the drivers. Instead, we take an empirical approach to the question, because a natural experiment is available to us.

Per-passenger compensation occurs on 96% of Santiago's 8000 buses and 300 bus routes. However, two bus companies⁵, with 332 buses on 25 routes, pay their drivers a fixed wage. The government created these new fixed-wage routes in 2001. Because of complaints about the per-passenger compensation system, the government auctioned off these route contracts with the stipulation that driver

⁴ Krbálek and Seba (2000) have previously noted the existence of *sapos* and of strategic interaction between drivers on a bus route in Cuernavaca, Mexico, for which they recorded 3500 bus arrivals to estimate the empirical distribution of arrival intervals. However, as physicists rather than economists, they did not examine consumer welfare effects, but rather on demonstrating consistency of the data with a quantum chaotic model of a one-dimensional gas.

⁵ The two companies are RedBus, S.A. and Alsa, S.A. Both operate under the Metrobus brand name.

compensation include a fixed-wage component. If we assume that the compensation scheme for these new routes was chosen exogenously by the government (that is, no unobservable driving conditions on the routes made the government think these routes would be particularly good for fixed-wage as opposed to per-passenger compensation), then comparisons between routes with the two different compensation schemes serves as a natural experiment. This experiment unfortunately has no before-and-after comparison, but we can treat the fixed-wage routes as an experimental treatment and the per-passenger routes as an experimental control.

For our empirical project, we managed to locate fixed-wage routes and per-passenger routes that both traveled through the same sector of the city and had similar route characteristics. By comparing bus arrival data from similar routes, we can test whether per-passenger compensation for drivers results in more even bus headways, relative to a fixed-wage compensation scheme.

In order to relate bus arrival data to consumer welfare, we use a theoretical expression for the expected waiting time for a random passenger showing up at a bus stop. Waiting time is especially significant in passenger welfare given the calculation by Mohring *et al* (1987), that the disutility of waiting for a bus is three times as high as that of traveling on a bus. It can be shown that independent of the functional form of the distribution of bus-arrival intervals, the expected waiting time is a simple expression of just the mean and variance of this unknown distribution. We use this result to guide our empirical model.

Using 10,824 observations of bus arrivals collected from multiple points along each of 32 different bus routes, we estimate a regression model. Half of the routes use per-passenger incentive contracts with drivers, while the other half use fixed-wage contracts. We allow the headway variance to change over the course of a single bus route, so that passengers' expected waiting times may be longer towards the end of the route. Since routes vary in their characteristics, we include route fixed effects, and ask whether the deterioration of passenger waiting times increases faster over the course of an incentive-contract route than it does over the course of a fixed-wage route. An interesting technical point is that we are able to use the first four moments of the distribution of bus arrivals at each location in order to compute the variance-covariance matrix for efficient GLS estimation of our regression coefficients. The answer to our question proves to be yes: the per-passenger incentive contract results in more regular bus headways than the fixed-wage contract does.

However, the benefits of per-passenger driver compensation come with some negative side effects for passengers. Initially we conjectured that drivers would improve their spacing by slowing down if they got too close to the bus immediately ahead, but this turned out to be incorrect. In fact, once they get sufficiently close, they attempt to *pass* the bus in front.⁶ This ameliorates the problem of bus bunching: an empty bus proceeds more quickly (making quicker stops) than a bus full of

⁶ In a survey we administered to drivers paid per-passenger (see section 2.3 below), 63% said that they "always" tried to pass the bus in front of them when the two buses were traveling together.

passengers, so putting the empty bus in front of the full bus tends to reduce bunching. However, this technique often involves aggressive driving by the driver attempting to pass, which can result in an uncomfortable passenger ride or an increased probability of accidents. To examine the size of these side effects, we conducted surveys of several hundred Santiago bus passengers and drivers, the results of which we report below. We also managed to collect accident data for all Santiago buses, broken down by bus company, in order to discover whether incentive-contract buses are more prone to traffic accidents than fixed-wage buses. The answer again turns out to be affirmative, and we produce a quantitative estimate of the net effect of the compensation system on deaths and injuries.

We have thus been able to produce estimates of both the benefits and the costs of per-passenger incentive contracts relative to fixed-wage contracts. The results have relevance for policymakers, as government agencies may regulate the contractual arrangements of bus companies. Indeed, the two companies that currently pay fixed wages in Santiago do so because the government required of bidders that a high percentage of drivers' compensation be paid as a fixed wage.

The question of contract design is an important one for bus transportation policy. In Latin American cities, bus travel accounts for more than half of all passenger trips (Wright, 2001), so the safety and service characteristics of urban buses have a significant impact on the quality of life. Bus transportation remains important in many cities outside Latin America as well, and recent developments in global-positioning technology may make it possible to design new types of contracts

that benefit passengers in new ways. Below, we speculate about possibilities for new types of technology-enabled contract design that might provide the benefits of the Santiago incentive system without as many of the costs.

In addition to the policy questions, we feel we have made an important contribution to the economics literature by documenting not just the behavior changes induced by an incentive contract but also an estimate of the resulting welfare costs and benefits.

2. Methodology

In evaluating the difference between the two systems, we used three different criteria: average passenger waiting time, service quality, and number of accidents.

2.1 Average Waiting Time

Ideally, at periods of constant demand, buses should arrive at evenly-spaced intervals. For a frequency of 6 buses per hour, a bus should arrive exactly every 10 minutes at any given point along the route. However, because of various factors, including different driving patterns, uneven passenger arrivals at stops, and the varying incidence of congestion, and different driving patterns, buses do not arrive at set intervals. As noted above, passengers commonly report dissatisfaction with the bunching of buses.

In fact, it turns out that even spacing is mathematically unstable, so that buses tend to bunch, or “platoon” along the route (Strathman, 2003). Buses may start out with even intervals, but a small random shock, such as local traffic congestion or the arrival of a sudden influx of more passengers, causes one bus (say bus A) to be

stopped longer than usual at a stop. This may cause the bus to fall behind schedule. As it falls behind schedule, more and more passengers arrive at stops to wait for its arrival, which slows it down even more. The driver must spend extra time boarding those passengers and collecting their fares and later unboarding them.⁷ Meanwhile, Bus B, immediately following A on the same route, starts collecting fewer passengers than usual because the interval between A and B has reduced. The small initial change thus gets amplified, as Bus A makes longer and longer stops to pick up and drop off more passengers, while Bus B similarly makes shorter and shorter stops. This process continues until Bus B completely catches up to Bus A.

The more uneven the intervals between buses, the more time the average passenger has to wait. To see this intuitively, suppose that the buses on a route have an average spacing of 10 minutes. Suppose that at some point we observe the interval between Bus A and Bus B to be 15 minutes long, with the subsequent interval between Bus B and Bus C only 5 minutes long. The passengers waiting for Bus B will be waiting longer than average, while the passengers waiting for Bus C will be waiting shorter than average. However, more passengers will accumulate in line for Bus B than for Bus C, causing the (weighted) average number of minutes to be higher than if the buses were equally spaced.

This phenomenon is known in the transportation literature as the inspection paradox: a passenger arriving at the bus stop will likely have to wait during a

⁷ Time spent collecting fares is an especially important consideration in Santiago, since drivers are in charge of collecting them, and usually paying the exact fare requires many coins.

longer-than-average bus interval. Let us call h the random variable representing the length of a headway, and let its first and second moments be μ and σ^2 . Then, the expected waiting time, $E(w)$, can be shown to be as follows (see Daganzo, 1997, for a demonstration):

$$E(w) = \frac{E(h)}{2} + \frac{\text{Var}(h)}{2E(h)} = \frac{\mu}{2} + \frac{\sigma^2}{2\mu} \quad (1)$$

This result holds for any bus arrival distribution. The first term corresponds to the average passenger waiting time caused by the average interval, while the second to the one caused by unevenness in the intervals. This shows that waiting time is minimized by minimizing the variance of the headways. In the limit, if headways are totally regular, waiting time reaches its lower bound at half of the average headway. Thus, dividing the second term by the first one we obtain:

$$\frac{\text{Var}(h)}{E(h)^2} = \frac{\sigma^2}{\mu^2} \quad (2)$$

which corresponds to the percentage of wait time attributable to uneven spacing with respect to the minimum possible⁸. We initially hoped that we could simply compare expected passenger wait times across compensation systems. However, we realized that different routes have very different baseline means and variances of bus intervals, depending on traffic patterns. Furthermore, we cannot guarantee that the routes we observe are randomly assigned to a compensation scheme: there may

⁸ For simplicity we refer we hereafter refer to this as the percentage of wait time attributable to uneven spacing.

be correlation between these baseline numbers and the form of the compensation system, which would cause spurious correlation between the compensation system and average passenger wait times. Therefore, instead of relying merely on variation between routes, we choose to rely on variation within routes. In particular, we exploit our natural experiment to examine the differences in rate of deterioration of even spacing.

Routes start off with a low variance of intervals because an inspector regulates departures.⁹ This variance grows as the buses proceed along the route. Our goal was to measure the rate at which the variance increases, separately for the two different compensation systems.

To collect the appropriate data, we isolated a sector in the city that had fixed-wage and per-passenger routes of similar trajectories. We chose 8 fixed-wage routes and 8 per-passenger routes, using both directions for an effective 16 routes of each. We restricted attention to the hours between 6 am and 1 pm. We started our data collection by paying *sapos* for their notebooks of data. However, few *sapos* had notebooks with data of sufficient quality for this analysis, either because they took breaks, didn't work the right hours, or kept inaccurate records.¹⁰ Instead, we hired

⁹ Some routes traveled over parts of the city that we couldn't measure, so we decided to isolate the sector where we had measuring points and treated our first observation point as the dispatch point. For some routes we couldn't measure the variance over the entire length of the city. Instead, we isolated the sector where we had measuring points and treated our first observation point as the dispatch point.

¹⁰ Because of the sheer number of routes that *sapos* often cover, they cannot accurately mark every bus. For example, when three buses of the same route all arrive within a minute of each other, the *sapo* might not write down all three arrival times because the following bus will only need to know the time in minutes since that group passed. Also, *sapos* will sometimes fabricate a number to drivers. For instance, in order to still be paid, a *sapo* will sometimes signal a number to a driver even though he didn't have the interval information. Because of the varying quality of *sapo* data, we only used data for *sapos* whom we were able to audit carefully. In the

measurers that could be supervised and would focus only on the specific routes we were studying. One of us also took a number of measurements himself in order to audit the data. Finally, we also obtained Global Positioning System (GPS) data on bus arrivals on two per-passenger routes, and data from three fixed-wage routes that the operating company collected by hand.¹¹

We aimed for 5 measurement points for each route and measured over three to five days, depending on the route.¹² Descriptive statistics of the number of measurements and points per route are below. In some instances these numbers would be higher, as in the GPS data, while in other instances, such as one company’s 1-day, 3-point study, they would be lower. We also tried to position the measurement points as evenly as possible throughout the route.

Table 1. Descriptive Statistics of Measurements.

	Mean	Median	Minimum	Maximum
Total km of the observed portion of the route	17.6 km	15 km	12 km	27 km
# of measurement points per route	4.1	5	2	6
# of observations per measurement point	82.6	61	21	376
# of km between measurement points on the same route	5.6 km	4 km	.5 km	25 km

end only about 10% of our usable data came from *sapos*.

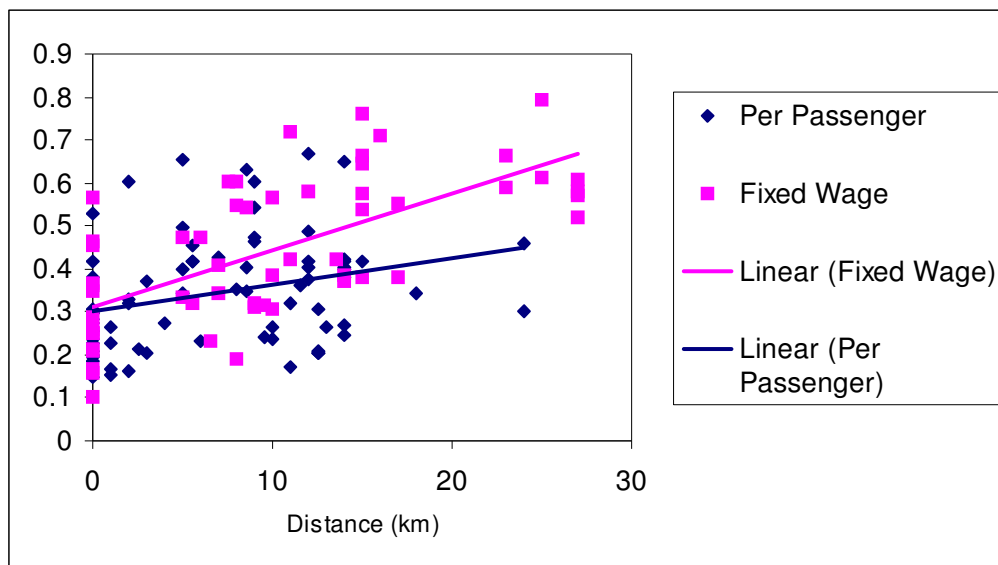
¹¹ The GPS data came from Cantares de Chile, S.A., and Lokal Traffik, S.A. (one route each). They use GPS for driver accountability, not to actively manage spacing. RedBus, S.A. collected one day’s worth of data on three of its routes.

¹² The vast majority of the time, the route-points were measured on the same days. For about a third of the routes, we added additional measuring days when we noticed that we wanted more reliable data.

In all, we obtained 10,824 observations of bus arrival times at 130 route-points on 32 routes. We choose to focus on average passenger wait times, which we showed above to depend both on the mean and the variance of bus time intervals. At each route-point, we aggregate our arrival observations to compute a sample estimate for σ^2/μ^2 , the proportion of wait time attributable to nonzero variance.

Figure 1 shows a scatterplot of the values obtained at the 130 different route-points. We plot our sample estimates of σ^2/μ^2 versus distance along the route. For each route, we have a first measurement point; we arbitrarily define the first point's distance to be zero. For other points along a given route, the distances are measured in number of kilometers from that route's zero point. By comparing the best-fit lines for the two different compensation systems, we can see our first indication that expected passenger wait times grow faster on fixed-wage routes.

Figure 1. Proportion of Wait Time Attributable to Uneven Intervals.



We next develop a more formal regression model in which we assume that

σ^2/μ^2 increases linearly with the distance traveled over the course of the route. Our research question is to measure the extent to which this rate of increase depends on the drivers' contract form.

To do so, we define the following variables:

f_r = value of σ^2/μ^2 at the first place where route r enters the area of measurement;

this acts as a regression fixed effect for route r

d_{mr} = distance, in kilometers, traveled by buses of route r from the first place where they enter the area of measurement to measurement point m .

pp_r = dummy variable corresponding to one if drivers of route r are paid per passenger, and 0 if paid a fixed salary.

The regression model then explains the proportion of average passenger wait time due to uneven spacing as:

$$\left(\frac{\sigma^2}{\mu^2}\right)_{mr} = f_r + \alpha * d_{mr} + \beta * pp_r * d_{mr} + \varepsilon_{mr} \quad (3)$$

Table 2 displays the results of the regression estimates for α and β .

Table 2. Ordinary Least Squares Regression.¹³

Explanatory Variable	Coefficient (SE)
d	.0135773 (.0015166)
$pp * d$	-.0076021 (.0024127)

Since we observed very different numbers of bus arrivals at different route-

¹³ Fixed-effect estimates have been suppressed. Standard errors are in parentheses.

points, we might be concerned that our dependent variable was measured less precisely for some routes than for others, thus causing heteroskedasticity.¹⁴ To correct for this heteroskedasticity, we choose to estimate the uncertainty in our estimates of σ^2/μ^2 and run the regression using generalized least squares. Since σ^2/μ^2 is a function of the first two moments of the bus arrival data, its variance turns out to be a function of the first four moments of the bus arrival data. We look up formulas for the variances and covariance of the first two sample moments of a distribution, and combine them using the Delta Method to obtain a formula for the sample variance of the error term in our regression.

We note that the Y variable in our regression is the square of the coefficient of variation, which can be rewritten as a function of the first two noncentral sample moments:

$$\begin{aligned}
 Y &= \frac{\hat{\sigma}^2}{\hat{\mu}^2} \\
 &= \frac{\overline{x^2} - (\bar{x})^2}{(\bar{x})^2} \\
 &= \frac{\overline{(x^2)}}{(\bar{x})^2} - 1
 \end{aligned} \tag{4}$$

For simplicity of notation, we let A and B equal the first two noncentral sample moments of the bus arrival data, and note that our dependent variable Y is a simple nonlinear function of A and B :

¹⁴ This implies that the heteroskedasticity is due to basing our aggregates on different sample sizes. Other sources of heteroskedasticity are possible: for example, one of our workers might record arrival times with greater accuracy than another. They performed a difficult task, picking out specific buses in heavy traffic for up to eight hours at a time, so even though we did our best to audit their work for accuracy, these samples likely contain small, differing amounts of measurement error.

$$\begin{aligned}
A &= \bar{x} \\
B &= \overline{x^2} \\
Y &= \frac{B}{A^2} - 1
\end{aligned} \tag{5}$$

Defining α_k ($k = 1, 2, 3, 4$) to be the population parameter equal to the k^{th} non-central moment of the distribution, we know from van der Vaart (1998) that the variances and covariance of the two sample moments A and B can be written as follows:

$$\begin{aligned}
\text{Var}(A) &= \frac{1}{m}(\alpha_2 - \alpha_1^2) \\
\text{Var}(B) &= \frac{1}{m}(\alpha_3 - \alpha_1\alpha_2) \\
\text{Cov}(A, B) &= \frac{1}{m}(\alpha_4 - \alpha_2^2)
\end{aligned} \tag{6}$$

We combine these expressions using the Delta Method to obtain the desired variance of Y :

$$\begin{aligned}
\text{Var}(Y) &= \left(\frac{dY}{dA}\right)^2 \text{Var}(A) + 2\left(\frac{dY}{dA}\right)\left(\frac{dY}{dB}\right) \text{Cov}(A, B) + \left(\frac{dY}{dB}\right)^2 \text{Var}(B) \\
&= \left(\frac{-2B}{A^3}\right)^2 \left(\frac{1}{m}\right) (\alpha_2 - \alpha_1^2) + 2\left(\frac{-2B}{A^3}\right)\left(\frac{1}{A^2}\right)\left(\frac{1}{m}\right) (\alpha_4 - \alpha_2^2) + \left(\frac{1}{A^2}\right)^2 \left(\frac{1}{m}\right) (\alpha_3 - \alpha_1\alpha_2)
\end{aligned} \tag{7}$$

To estimate the variance of each individual observation Y_i , we replace the population parameters with their sample estimates from the bus arrival data. We then re-estimate the previous regression using generalized least squares, where the regression error variance is assumed to be a diagonal matrix with elements given by our estimates. The result of this regression is shown in table 3:

Table 3. Weighted Least Squares Regression.¹⁰

Explanatory Variable	Coefficient (SE)
<i>D</i>	.01034950132037 (.00130090015205)
<i>pp * d</i>	-.00515280067441 (.00258690018658)

Note that the correcting for heteroskedasticity does not change the results significantly. We will continue to use GLS for the remainder of the paper. relative relationship stayed the same, but the magnitude of both effects was roughly reduced by a third. This is expected, however, because by definition the weighting decreases the influence of the route-points with high variance and increases the influence of the route-points with low variances.

We were also concerned that variance could have been misinterpreting by our combining route data for peak and off peak periods. Because mean arrival rates vary over the course of the day we check for robustness by constructing separate observations for peak versus off-peak periods.¹⁵ Because our initial data didn't contain hour-of-day information (data purchased from *sapos* indicate only the minute of the hour), we ended up with hour-of-day information for only 12 of the 32 routes. We define the peak period to be from 6 a.m. to 10 a.m and off-peak to be from 10 a.m. to 1:00 p.m. We then create two observations (peak, off-peak) for each of the 12 routes and redo the regression with a total of 44 "routes". The GLS results were as follows.

¹⁵ Because our initial data didn't contain hour-of-day information (data purchased from *sapos* indicate only the minute of the hour), this was possible to do for only 12 of the 32 routes. We define the peak period to be from 6 a.m. to 10 a.m and off-peak to be from 10 a.m. to 1:00 p.m. We then create two observations (peak, off-peak) for each of the 12 routes and redo the regression with a total of 44 "routes".

Table 4. Weighted Least Squares for splitting peak versus off-peak periods.¹⁰

Explanatory Variable	Coefficient (SE)
<i>Dd</i>	.0091795013294 (.00128290016198)
<i>pp * d</i>	-.0050845 008629 (.00146140019864)

These results are qualitatively similar to our original estimates, with slightly smaller standard errors, as one might expect if pooling two different periods introduces noise. To check this hypothesis further, we restricted our sample to just those 12 routes we were capable of splitting into peak and off-peak. The regression was as follows:

Table 5. Comparison of Weighted Least Squares Results of splitting versus not splitting peak and off-peak periods. ¹⁰

Explanatory Variable	Coefficient (SE)	Coefficient (SE)
	For 12 unsplit routes	For 12 routes split into 24
<i>dD</i>	.0089043015365 (.0018821002207)	.00756560152513 (.001620022669)
<i>pp * d</i>	-.00819050094381 (.00258690037077)	-.0107642 0182183 (.0023016003458)

SBy splitting the twelve routes into peak versus off-peak periods, we find even stronger results: the magnitude of the effect of the per-passenger compensation scheme nearly doubles, and the standard errors get slightly smaller. However, because we have to throw nearly two-thirds of the data to run this specification, this is not our preferred specification. One troubling feature of this final specification is that the point estimates predict *improvement* of the variance over the course of a per-passenger route (0.015-0.018=-0.003). However, an F test shows that the sum of the two coefficients is not statistically significant at any conventional level of confidence. Notice also that the full sample of route observations yields lower standard errors (Table 3) than the restricted sample (Table 5), indicating that the econometric efficiency gains of retaining the full sample outweigh the efficiency gains of splitting into peak versus off-peak routes. We therefore prefer our original specification (Table 3).

Overall, the results suggest that there is a strong relationship between per-passenger compensation and passenger waiting time. For example, a passenger waiting for a bus 30 km away from the start of a route can expect to wait 15% longer,

on average, when the driver is paid a fixed wage rather than per passenger.

How large might we expect to be the overall benefit of the per-passenger spacing to passengers in Santiago? Because we model the rate of change of wait time over the course of the route rather than the level of wait time, we can only come up with a rough estimate. One-way route lengths in Santiago range from 30km to 50km. If the average passenger boards halfway through the route, then the bus may have traveled about 20km by the time the passenger boards. The fixed-wage system thus adds $20 \times (0.005) = 10\%$ to the average wait time relative to the per-passenger system. The average Santiago passenger waits 4.5 minutes for each bus,¹⁶ and Santiago government officials estimate the value of a citizen's time at 724 Chilean pesos (US \$1.13) per hour,¹⁷ which likely underestimates the disutility of time spent waiting.¹⁸ So if all routes in Santiago started on a fixed-wage system and converted to a per-passenger system, our rough estimate is that passengers would save 24 million Chilean pesos (US \$38,000) worth of waiting time each day!¹⁹

This model only takes into account welfare effects on passengers waiting at stops. It fails to look at, for example, the value of travel time saved by the generally higher velocity of per-passenger drivers.

It appears that having drivers paid per passenger results in a significant

¹⁶ This is the official figure used by SECTRA (Government Ministry of Transportation Planning), the transportation research branch of the Chilean government. However, we have reason to believe, based on our survey of 200 passengers, that the real number is roughly twice that amount.

¹⁷ This figure is the figure the Chilean government uses for project evaluation (Precios 2004).

¹⁸ Mohring *et al.* (1987) calculated the disutility of time spent waiting at three times the disutility of time spent traveling on a bus.

¹⁹ Based upon 4,457,238 passenger trips per day (Encuesta, 2001).

saving of passenger wait time. The dynamic adjustment of intervals shows a clear improvement over the fixed-salary drivers in the data.

2.2 Service Quality

To understand service quality from the passenger point of view, we administered a survey of 300 passengers. We approached 200 customers of the per-passenger buses, both at bus stops and on the buses themselves, and orally administered a 7-question survey in Spanish. We administered the same survey, at a subway-bus transfer station, to 100 passengers of fixed-wage routes. The two samples reported a distinct difference in service quality. In particular, users of regular (per-passenger) buses were far less likely to feel that the driver waits until they are safely on board before starting, as shown in Table 4. These results back up a larger survey which showed that the bus system is the least popular public service in the city (Adimark 2001).

Table 4. Survey of Passengers on Driver Behavior

When you get in the bus, how often does the driver wait until you're safe to continue?		
	Users of Metrobus (fixed wage)	Users of yellow buses (per passenger)
Always	59%	14%
Almost Always	25%	31%
Half the Time	4%	8%
Rarely	9%	33%
Never	3%	14%

100 users of Metrobuses and 200 users of regular buses surveyed July and August, 2004. $p < .005$.

Another negative side effect of incentive contracts involves buses occasionally

failing to stop to pick up passengers. When a passenger is waiting at a stop alone, sometimes the driver won't stop because the opportunity cost of the time spent picking up that passenger is greater than the income from the fare. In fact, often times a single passenger waiting will have to wait for several buses or until more passengers arrive at the stop (Naudon Dell'Oro, 2004). Once a bus finally does stop, the driver quickly gets the bus moving at full speed, often in complete disregard for the stability or comfort of the passenger. These rapid stops and quick accelerations can occur for the entire duration of the trip.

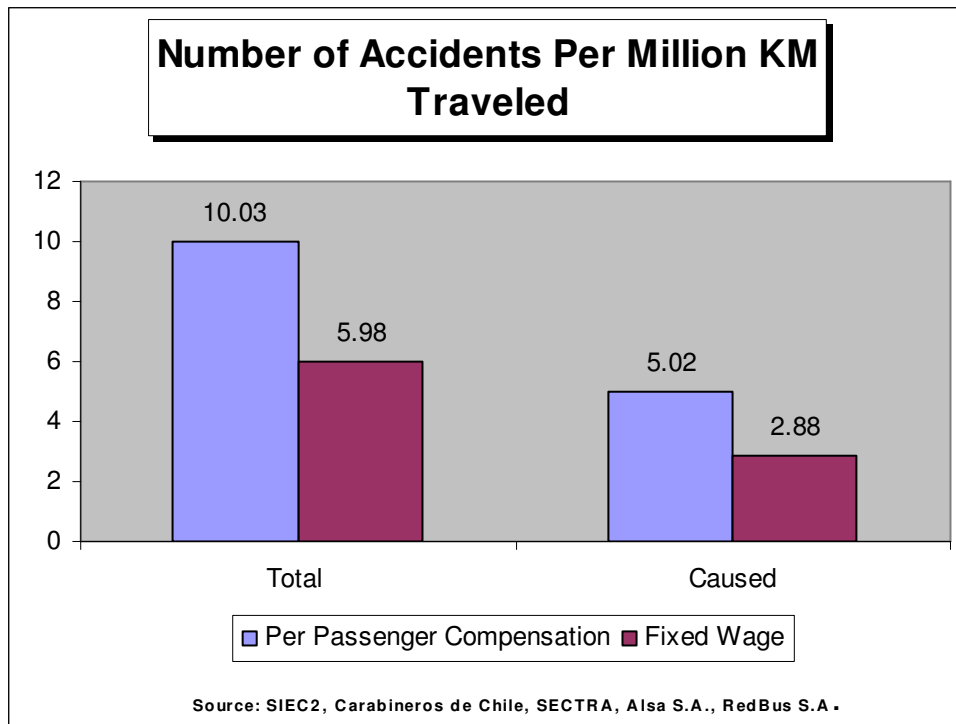
We also administered a 21-question²⁰ survey to 100 per-passenger drivers and 46 fixed-wage drivers at various bus depots. Per-passenger drivers talked about the demands that the per-passenger compensation system puts on them. Sixty-six percent agreed with the statement: "Being paid per passenger makes you drive more aggressively." On the other hand, 83% of fixed-wage drivers agreed with the statement: "Being paid a fixed wage makes you drive more safely."

2.3 Accidents

In the first 6 months of 2004, 3,961 buses were involved in an accident and 1,960 buses caused an accident. After analyzing the accidents by license plate, to find out which driver compensation system applies in each case, the results show that routes with drivers paid per passenger clearly are involved in and cause relatively more accidents, as is shown in Figure 2.

²⁰ We asked more questions of the drivers than of the passengers because, in addition to questions about service quality, we were trying to understand how drivers make use of *sapo* information.

Figure 2. Number of Accidents per Million Km Traveled.



These results were derived by taking the overall dataset of accidents in Santiago, and taking out accidents for the two companies that pay a fixed wage. The two were then compared and indexed for the number of kilometers traveled. Moreover, this result likely understates the difference in accidents between the two compensation systems due to differences in reporting of minor accidents. The organizational structure of the per-passenger companies is relatively informal, and local transportation experts believe that its drivers are less likely to report accidents to the police. By contrast, the fixed-wage Metrobuses belong to more reputable firms that are more likely to report accidents.

The implication of this difference in the number of accidents is significant.

Specifically, if all buses in Santiago had the same number of accidents per km as those with drivers paid a fixed wage, we estimate that it would save 55 lives per year. It would also eliminate 227 serious injuries, 210 less serious injuries, and 1,293 light injuries.²¹

3. Discussion of results

We find our most interesting result to be the fact that per-passenger compensation yields more regular spacing of buses, and hence lower expected waiting time for passengers, than does fixed-wage compensation. However, this clearly does not mean that per-passenger compensation is a superior system. We find significant costs to the per-passenger system as well. In particular, per-passenger compensation exhibits a much higher incidence of accidents and much lower passenger comfort.

But the more we thought about the problem, the more we began to think that the best way to approach the question is not to say that one system is better than the other. Instead, we try to see if there are alternatives that can combine the positives of the two systems while minimizing the negatives.

The principal advantages of per passenger compensation stem from drivers

²¹ Based upon 27,077,232 km traveled by fixed-salary buses per year, 773,668,768 km traveled per year by per passenger buses. This chart compares the two systems over a 6-month (January to July, 2004) sample of accidents, during which time fixed-salary buses caused 39 accidents and per-passenger buses caused 1,941 accidents. Assumes all buses had one accident per 347,144 km, and that accidents would produce injuries at the same rate as before. Data comes from SIEC2, the transportation statistics division of the Carabineros, the Chilean police.

being more motivated.²² Besides shorter waits for passengers, this makes managing drivers much easier for bus owners. While we were undertaking the study, we noticed a sharp difference in the behavior of drivers at their respective bus depots. Drivers paid per passenger were excited to get back on the road. They took quick bathroom breaks, quick meal breaks, and were always ready to depart when the inspector (who regulates departure times) called them. At the depots of fixed-wage bus companies, drivers were not ready when the inspector called them. They took longer meal breaks, spent more time socializing, and used the bathroom and other excuses to delay leaving.

The costs of per-passenger compensation come from drivers caring so much about maximizing their number of passengers that they impose externalities on others. The problems include more than just accidents due to aggressive or reckless driving. For example, per-passenger drivers stop anywhere there is a passenger, not just at officially designated stops. This is illegal and presents high social costs relative to a system in which all passengers wait only at designated stops. Stopping at non-designated stops decreases traffic flow for everyone on the road, and lengthens overall passenger transport time because the buses stop too frequently. Can these negatives be minimized without losing the benefits of a per-passenger system?²³

²² The productivity benefits of incentive based contracts are not unique to bus drivers. A study by Shearer (2005) of two different compensation systems in the timber industry found that workers paid piecewise (per unit of work) were 20% more productive than workers paid a fixed wage.

²³ Note that even if he preferred fixed-wage compensation, it would be imprudent for an individual bus owner to switch from per-passenger to fixed-wage compensation. Santiago has always had a highly fragmented

It would be useful to find an alternative incentive system for drivers to maintain even intervals between buses, without the costs of lower service quality and more accidents. Giving drivers incentives to meet a fixed timetable is not likely to be valuable in keeping regular intervals, because the timetable is static rather than dynamic. For instance, if one bus breaks down and other drivers follow their time goals, they will not adjust for the resulting gap in a way that would minimize average passenger wait time. On the other hand, in a dynamic system, such as one encouraging even intervals, other buses might adjust their intervals in order to compensate for the broken-down bus.

The *sapos* are a major part of the incentive system's ability to improve bus spacing. According to our survey of 100 per-passenger drivers, the average driver pays 6.5 *sapos* a total of 570 Chilean pesos (US \$0.90) each day. Furthermore, *sapos* are not unique to Santiago. Others have observed the equivalent of *sapos* in other Latin American cities, such as Lima, Peru, where they are known, more appropriately, as *dateros* (data collectors).²⁴ *Sapos* act as key information providers, letting drivers know the locations of the other buses. We find it fascinating that the network of *sapos* springs up through a market process, especially given that both

organization of bus owners. For the 8,000 buses, there are as many as 3,000 owners (Transantiago website 2005). These owners group together to form "paper companies" of 20 to 30 owners per route. These paper companies serve several purposes. First of all, they allow individuals owners to use their collective power to buy from the government the right to operate routes (which are auctioned off). They also allow owners to share overhead, such as inspectors, depots, maintenance facilities, and office space. However, these associations neither coordinate route operation nor share revenues. Owners get their revenue directly from the drivers, and drivers are paid directly by the owners. (Revenue sharing would create free-rider problems among owners.) If one owner of a single bus suddenly decides to pay his driver a fixed wage, that driver could be taken advantage of by per-passenger drivers who have incentives to aggressively seek passengers, causing revenue loss for the owner.

²⁴ Interview with Ian Thomson, United Nations Economic Commission for Latin America and the Caribbean.

drivers and *sapos* tend to be relatively uneducated. The *sapos* provide the drivers with “knowledge of the particular circumstances of time and place” (Hayek, 1945), enabling them to make decentralized decisions that are in some ways superior to the decisions that can be made by a centralized dispatcher.

Modern technology, however, offers a potential improvement over a network of *sapos*. With a full implementation of GPS technology, drivers could have information on the location of other buses at all times, not just on corners where *sapos* work. Drivers could have a real-time display showing the locations of other buses both in front of and behind them, enabling them to respond with adjustments to the spacing. GPS technology opens up a whole new realm of contractual possibilities. For example, one might pay drivers a bonus based on the continuous-time average spacing between their bus and other buses on the route, thus providing drivers with appropriate incentives to minimize passenger waiting time. Such systems could improve quality of passenger service in cities like Santiago, and could be potentially even more useful in the cities of developed nations.²⁵ Appropriately designed contracts might even be able to provide improved bus arrivals without incurring the costs of aggressive driving. We believe additional research on the combined effects of information technology and contractual incentives could prove immensely valuable to city planning agencies in designing systems that provide better service to passengers.

²⁵ In Santiago, there would be a number of practical and legal concerns, such as whether companies are allowed to track competitor’s buses. In U.S. cities, by contrast, bus service is usually provided by a regulated monopolist, so we would expect fewer competitive concerns and more potential benefits to be had.

References

- Adimark (2001). Presentación de Resultados Índice Nacional de Satisfacción de Consumidores Segundo Semestre .
<http://www.adimark.cl/calidad22001/index.htm> October 1, 2002.
- “CTA Expands Program to Reduce Bus Bunching.” Chicago Transit Authority Press Release. October 4, 2000.
<http://www.transitchicago.com/news/archpress.wu?action=displayarticle_detail&articleid=125309>.
- Daganzo, C. F. (1997) *Fundamentals of Transportation and Traffic Operations*, Pergamon-Elsevier, Oxford, U.K.
- Eastaway, R., and J. Wyndham (1998). *Why Do Buses Come in Threes? The Hidden Mathematics of Everyday Life*. New York: Wiley.
- Encuesta Origen Destino de Viajes: Informe Ejecutivo*. 2001. SECTRA (Chilean Ministry of Transportation Planning).
- Gatica Gutiérrez, Ramón (2003). *Actitud Hacia La Modernización de Los Choferes de la Locomoción Colectiva: Características de su Trabajo y Predisposición Hacia la Profesionalización de su Oficio*. Tesis para optar al Título Profesional de Sociólogo. Universidad de Chile. Facultad de Ciencias Sociales. Departamento de Sociología. December 30, 2003.
- Hayek, F.A (1945). “The Use of Knowledge in Society”. *American Economic Review*. XXXV, No.4. September, 1945. pp. 519-30.
- Krbálek, Milan and Seba, Petr (2000). “The statistical properties of the city transport in Cuernavaca (Mexico) and random matrix ensembles.” *Journal of Physics A: Math. Gen.* 33 No 26 (7 July 2000) L229-L234
- Mohring, H., J. Schroeter and P. Wiboonchutikula (1987). “The values of waiting time, travel time, and a seat on the bus.” *RAND Journal of Economics*. 1987. 18 (1), 40-56.
- Naudon Dell’Oro, María de Los Ángeles (2004). “¿Y a mí qué me importa?” *El Mercurio*. August 28, 2004. Santiago, Chile.
- Precios Sociales Para La Evaluación Social de Proyectos (2004). MIDEPLAN.
- Pearson, Erica (2003). “Bus Bunching”. *The Gotham Gazette*. May 5, 2003. New York. <<http://www.gothamgazette.com/article/20030505/200/372>>.
- Presentación de Resultados Índice Nacional de Satisfacción de Consumidores Segundo Semestre 2001. <http://www.adimark.cl/calidad22001/index.htm> October 1, 2002.
- Rider Diaries. Straphangers Campaign. Internet message board. January 11, 2005.

- <http://www.straphangers.org/ubb/ultimatebb.php?ubb=get_topic;f=25;t=001156>
- Shearer, Bruce (2004). "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment." *The Review of Economic Studies*, Vol. 71, No. 247 (April 2004), pp 513-534.
- Strathman, James G, Kimpel, Thomas J (2003). Headway Deviation Effects on Bus Passenger Loads: Analysis of Tri-Met's Archived AVL-APC Data. Final Research Report for Tri-Met. January 2003.
- The Late State of the Buses: A Survey of Reliability of Service and On Time Performance of CTA Bus Routes*. The Neighborhood Capital Budget Group's Campaign for Better Transit. Chicago, Illinois. February 2004.
<http://www.bettertransit.com/bus_study.pdf>.
- The impacts of subsidies, regulation and different forms of ownership on the service quality and operational efficiency of urban bus systems in Latin America*, CEPAL (Transport and Communications Division), with GTZ, report LC/L.675, 7th August, 1992.
- Transantiago Website (2005).
http://www.transantiago.cl/descargas/transantiago_english.zip. 2005.
- van der Vaart, A.W (1998). Asymptotic Statistics. Pp 25-27. Cambridge University Press. 1998.
- Wright, Lloyd (2001). "Latin American Busways: Moving People Rather than Cars." *Natural Resources Forum*, JNRF 25:2, May 2001.