

THE TRUCKERS AND TURNOVER PROJECT:
CONTEXT, DESIGN, AND A SELECTION OF PILOT RESULTS

Stephen V. Burks, University of Minnesota, Morris¹
Jeffrey Carpenter, Middlebury College
Lorenz Götte, University of Zürich
Kristen Monaco, California State University at Long Beach
Aldo Rustichini, University of Minnesota, Twin Cities
Kay Saager, Cooperating Carrier

This Draft: March 29, 2006

Submitted to the Conference on the Analysis of Firms and Employees (CAFE)

Abstract: The Truckers and Turnover Project is a statistical case study of a single firm and its employees which matches proprietary personnel and operational data to new data collected by the researchers to create a two-year panel study of a large subset of new hires. The project's most distinctive innovation is the data collection process which combines traditional survey instruments with behavioral economics experiments. The survey data include information on demographics, risk and loss aversion, time preference, planning, non-verbal IQ, and the MPQ personality profile. The data collected by behavioral economics experiments include risk and loss aversion, time preferences (discount rates), backward induction, patience, and the preference for cooperation in a social dilemma setting. Subjects will be followed over two years of their work lives. Among the major design goals are to discover the extent to which the survey and experimental measures are correlated, and whether and how much predictive power, with respect to key on-the-job outcome variables, is added by the behavioral measures. The panel study of new hires is being carried out against the backdrop of a second research component, the development of a more conventional in-depth statistical case study of the cooperating firm and its employees. This is a high-turnover service industry setting, and the focus is on the use of survival analysis to model the flow of new employees into and out of employment, and on the correct estimation of the tenure-productivity curve for new hires, accounting for the selection effects of the high turnover.

JEL Categories: C81, C93, L92, J63

Keywords: field experiment, risk aversion, loss aversion, time preference, IQ, MPQ, numeracy, U.S. trucking industry, for-hire carriage, truckload (TL), driver turnover, employment duration, survival model, tenure-productivity curve

¹ Corresponding author; svburks@morris.umn.edu.

1. Introduction²

The Truckers and Turnover Project is a statistical case study of a single large trucking firm and its driver employees. The cooperating firm operates in the largest segment of the for-hire trucking industry segment in the United States, the “full truckload”, or TL segment, in which approximately 800,000 people are employed, according to the 2002 Economic Census. The TL segment has a high-turnover labor market for its main employee group, tractor-trailer drivers, and the project is designed to address a number of academic and business questions that arise in this setting.

One major part of the project matches proprietary personnel and operational data to new data collected by the researchers, to create a two-year panel study of a large subset of new hires. The most distinctive innovation of this project component is the data collection process, which combines traditional survey instruments with behavioral economics experiments. The survey data include information on demographics, risk and loss aversion, time preference, planning, non-verbal IQ, and the MPQ personality profile. The data collected by behavioral economics experiments include risk and loss aversion, time preferences (discount rates), backward induction, patience, and the preference for cooperation in a social dilemma setting. Subjects will be followed over two years of their work lives. Among the major design goals are to discover the extent to which the survey and experimental measures are correlated, and whether and how much predictive power, with respect to key on-the-job outcome variables, is added by the behavioral measures.

The panel study of new hires is being carried out against the backdrop of a second research component, the development of a more conventional in-depth statistical case study of the cooperating firm and its employees. This component involves constructing large historical data sets from fragmented legacy IT sources, and using them to create multivariate models of turnover and productivity. Two main emphases are on the use of survival analysis to model the flow of new employees into and out of employment, and on the correct estimation of the tenure-productivity curve for new hires, accounting for the selection effects of the high turnover.

The project is designed to last three and a half years, with the first half-year for set up, and then a year for the initial intensive data collection in the panel study of new hires, in parallel with the construction of the data sets for the statistical case study, and the initial generation of modeling from these data. Then there will be two years of lower-intensity work while follow-up data is collected from the participant in the panel study of new hires.

The balance of the paper is structured as follows. Section 2 sets the context by describing the U.S. trucking industry, and the role of the TL segment within it. Section 3 discusses the nature of the labor market for TL drivers, and why it has had a high turnover equilibrium for about twenty-five years. Section 4 discusses the nature of the research relationship with the cooperating firm, and how it was constructed. Section 5 discusses the statistical analysis of historical operational and human resource data from

² The authors gratefully acknowledge financial support from the MacArthur Foundation’s Norms and Preferences Research Network, the Sloan Foundation, the Trucking Industry Program at Georgia Tech, the University of Minnesota, Morris, and from the cooperating carrier. We also thank the managers, staff, and employees of the cooperating carrier, without whose involvement and active support the project would be impossible. The experimental designers of the present project would like to thank Catherine Eckel and Kate Johnson for sharing protocol and design details from the large field experimental project in Mexico to which they are key contributors, and for offering very helpful advice on our design issues.

the firm. It has two main subparts, Section 5.1, which exhibits preliminary findings on turnover, and Section 5.2, which does the same for productivity. Section 6 describes the design of the panel study of new hires. It has four main subparts. Section 6.1 describes the process by which new students are trained as tractor-trailer drivers, and Section 6.2 discusses the schedule for the data collection effort at the training school. Section 6.3 lists and describes the five data collection events that take place during the first two-hour session of each data collection event, while Section 6.4 does the same for the second two-hour section of each data collection event. Section 7 concludes.

2. The U.S. Trucking Industry

2.1. Segments within the Industry

To a casual observer one truck looks much like another, but in fact, the operations that provide trucking services in the U.S. are meaningfully differentiated from each other on several dimensions. At the broadest level, trucking operations are broken into private carriage versus for-hire carriage, based on a legal relationship: whether the carrier also owns the freight (private carriage), or is hauling it for another party (for-hire carriage).³ In recent years for-hire carriers (one of which is the focus of the present study), have typically operated about one-third of the heavy trucks⁴ in the overall U.S. fleet, but about three-fifths of the total miles run by such vehicles (Burks, Monaco et al. 2004).

For-hire trucking is itself further broken into a number of distinct segments, separated along three cross-cutting dimensions. Within each segment inter-firm competition is significant, but across segments it may be muted, or in some cases even absent. The 2002 Economic Census, because of its use of the relatively new North American Industrial Classification System (NAICS), which is based on production process characteristics, gives a good overview of the structure of the for-hire trucking industry at this level of segmentation. For-hire truck transportation as a whole, NAICS category 484, generated \$165.56 billion in revenue in 2002, or about 1.56% of that year's GDP.⁵

The first broad scale distinction within for-hire trucking is between firms that use general purpose equipment (i.e. standard enclosed van trailers) to handle general commodities, and those that use specialized equipment to handle special commodities (examples of the latter would be refrigerated vans, flatbeds, tank trailers, and various other types of specialized equipment). According to the Economic Census, in 2002 general freight operations generated \$111.60 billion annual revenue (67.3% of the total) and specialized freight had \$54.01 billion annual revenue (32.3% of the total). A second cross-cutting broad scale distinction is between firms that make long distance inter-city hauls, and those that specialize in operations in and around a particular metropolitan area.

³ Private carriers are firms primarily in non-trucking lines of business who provide trucking services internally as support functions to their primary business operations. Examples might be deliveries of food by a retail grocery chain to its stores in trucks it also owns, or pickups of parts for assembly at an auto plant by the auto manufacturer's freight vehicles.

⁴ Heavy freight vehicles are defined here as having a gross vehicle weight (GVW) of more than 26,000 lbs., the level at which weight alone is sufficient to require the driver to hold a commercial driver's license (CDL).

⁵ Calculation by the authors; total GDP is from the U.S. Department of Commerce, Bureau of Economic Analysis; URL: <http://bea.gov/>.

In 2002 the Economic Census reports \$120.21 billion in annual revenue for long distance trucking (72.6% of the total), and \$45.35 billion for local hauls (27.4%).

A third cross-cutting broad scale distinction is based on the size of the typical shipment hauled, and this dimension on which firms differ is of particular relevance to the present study. It is easiest to understand this distinction by considering the two extremes, full-truckload service, and parcel service. At one end of the spectrum are firms like the one providing data for the current study. The archetypal full truckload (TL) carrier sends a driver with a tractor-trailer to a shipper's dock to fill up the trailer with a load, typically weighing from 10,000 to 48,000 pounds.⁶ The driver takes the loaded trailer wherever in the U.S. the shipment is destined, and unloads at the consignee's dock. The driver is then dispatched empty, possibly after waiting for a while, to the next location where a full load is available for pick up. TL carriers may use specialized equipment for special commodities, but if they haul general commodities they use general purpose equipment to maximize the chance of backhauls.

By contrast, both parcel and LTL firms aggregate large numbers of individual shipments collected at local terminals by local drivers into full trailer loads, and move them between terminal systems on fixed routes. Parcel carriers handle very small shipments (each piece typically being no larger than 150 lbs., with the average nearer to 50 lbs.), and LTL carriers aggregate medium-sized shipments (widely varying, but with average size around 1,000 lbs.). The Economic Census does not group parcel service firms with the for-hire trucking industry, but with air freight carriers. However, it does capture LTL and TL firms within trucking. In 2002 the TL segment dominated the general freight portion of (non-parcel) for-hire trucking, with 67.9% of the total employment, and 83.8% of the total revenue. If the segments of specialized freight that are primarily TL by shipment size are added to the mix⁷, then TL's share of the total employment of 1.137 million jumps to 72.8%, and its share of the total revenue of \$124.50 billion rises to 77.1%.

2.2. Differences in the Type of Competition within Segments

The differences across the segments in the operational routines needed affect the form and intensity of competition within each segment. Specifically, in the parcel and LTL segments, the need for a fixed network of freight re-handling terminals creates an entry barrier.⁸ While competition among parcel and LTL carriers is frequently strong, it generally takes place among incumbents. This is evidenced by the numbers of firms in the long distance parcel and LTL segments. In parcel there are really only four firms with full national coverage (UPS, FedEx, DHL, and the USPS). There are more LTL

⁶ The variation is because some less dense freight exhausts a trailer's volume at low weight levels, while more dense freight hits the weight limit before the volume limit.

⁷ Essentially, this means adding all specialized freight except household goods moving.

⁸ A brand new LTL carrier that wants to serve more than a single metropolitan area must create and operate a network that is of minimum size necessary to attract sufficient traffic from shippers with differing destination demands, relative to the total shipment flow densities in the geographic area it wishes to serve. But such networks exhibit strong economies of density (a combination of both scale and scope)—at low volumes the average costs are high, but they fall rapidly as volume increases. The expenses of running such a network until a large enough market share is obtained to make the new network cost competitive with those of incumbent carriers are non-recoverable (or “sunk”) if the firm exits. And the existence of a sunk cost of entry is the classic definition of an entry barrier.

firms, but the number is still small. The 2002 Economic Census identifies 89 long distance general freight LTL firms with five or more establishments, which is the minimum number of terminals needed to give any kind of geographic scope; there are only 57 firms with ten or more.

But in TL there are essentially no entry barriers. Since TL carriers do not normally re-handle freight once it is loaded, they do not typically require terminals, nor regular route patterns, for cost-competitive operations. So a one-truck carrier can cover the entire nation, and in doing so is competitive, on a load-by-load basis, with most of the services offered by one of the TL-segment's giants. When more complex service coordination is the key factor in market penetration, small firms can subcontract to third party logistics providers.⁹ And in fact, there is a continual flow into, and out of, the TL segment, mostly by firms operating at small to medium scales. In TL, the 2002 Economic Census identified 25,831 long distance general freight firms.¹⁰ The market concentration levels in these two segments also show the differing nature of competition. In LTL, the 2002 Economic Census puts the revenue share of the top 4 long distance general freight LTL firms at 36.3%, while it calculates the share of the top 4 long distance general freight TL firms to be only 14.7%.

The implication of these facts is that most of TL service is what business analysts call a "commodity business", and what economists call "perfectly competitive. As a result, the firms "at the margin", whose choices set prices for the whole market, are in TL not the big players, but instead are the small firms in the competitive fringe of the industry segment. Their pricing is in turn driven significantly by the wages drivers in such firms are willing to accept. Small firms generally face more modest wage expectations from their employees than do large ones, and they also have the benefit of more personal relationships between owners, managers, and drivers. And owner-operators, who make up a significant subset of the small firms, can always choose to pay themselves less in order to get started in the business. Large firms can pay a modest premium above the level set by such firms, because they may have cost efficiencies in other areas, and they may be able to maintain a small price premium due to offering customers a number of different services in an integrated fashion, but if they raise their wages too high they will make their costs uncompetitive. This industry structure sets the context for the derived demand for truck drivers in TL freight, and the consequent nature of the labor market for TL drivers.

3. The Labor Market for TL Drivers

3.1. Segmented Labor Markets Emerge

The American Trucking Associations' (ATA) quarterly turnover report typically shows the average turnover rate at large truckload (TL) motor carriers to be in excess of 100% per year (ATA Economic & Statistics Group 2005). Driver turnover among these carriers is an economically significant phenomenon—truckload carriers make up the

⁹ Since a TL carrier can subcontract actual movements in a spot market to owner-operators, it is possible for a firm to enter TL for-hire carriage initially with zero trucks.

¹⁰ Unlike the case of LTL, since TL firms don't have freight terminal networks, single establishment firms can be of national geographic scope, but in fact 997 of these had more than one establishment, which is still an order of magnitude larger than in LTL.

largest segment of for-hire motor carriage by employment, with approximately six hundred thousand drivers working at any given time (U.S. Census Bureau 2004)¹¹. This segment of the universe of for-hire trucking firms emerged into its present form after the economic deregulation of 1980, which transformed the structure of the trucking industry. Before deregulation the nature of the entry barriers created by government policies resulted in lots of TL output by firms using the LTL-type organization of production, with a fixed network of freight handling terminals (Belzer 1995; Burks 1999). But in the post-deregulation period carriers specialized quite strongly in one or another specific shipment size, from the smallest (parcel), through middle-sized shipments (less-than-truckload, or LTL), to the largest ones (truckload, or TL) (Corsi and Stowers 1991; Belzer 1995; Burks, Monaco et al. 2004).

As the truckload industry segment emerged, so did a parallel segmentation of the labor market for truck drivers (Belzer 1995; Burks 1999)¹². Drivers wanting to enter employment at parcel and less-than-truckload carriers generally found job queues¹³, while the labor market for truckload driving jobs began exhibiting high rates of turnover. In fact, the labor market in the truckload segment has essentially been in a high-turnover equilibrium since soon after the end of the recessions of 1981-82.¹⁴

3.2. The TL Driver's Job

To understand this situation, we start with a short description of the human capital investment needed to become a driver, and then discuss the working conditions encountered by the typical driver. Driving a tractor-trailer requires training for, and passing, the state-administered written and driving tests for a commercial drivers' license (CDL). Typically a high-school equivalent level of literacy is required, and training begins with at least two weeks mixed between classroom work and in-truck practice. This is usually followed by between a few days and a few weeks of initial driving experience, which is often obtained with an experienced driver riding in the cab as a coach, while the trainee is still driving on a "learner's permit", before he or she has taken the final test for the CDL. While the CDL test is administered separately by each state, as of 1991 they do so under Federal standards for what must be included. It is comprised of both written and driving portions, and the minimum legal age at which it may be taken is 21. Trucking firms generally considered a driver to be satisfactorily experienced after a year of work, so the level of human capital required places the job somewhere between unskilled and skilled, and it is best labeled as "semi-skilled".

Once a driver is licensed, the key problem in retention is generally perceived to be the working conditions faced by a tractor-trailer operator in the archetypal long-haul, randomly dispatched, 48-state service provided by most TL firms. In addition to the stresses of handling a big rig among swarms of cars, many drivers have very long weekly

¹¹ The calculation is this: in the 2002 Economic Census TL firms have 72.8% of the total employment of 1.137 million workers in (non-parcel) trucking, and the usual rule of thumb is that about 75% of the labor force employed by a TL firm is made up of drivers, the balance being made up of sales, customer service, administrative, and managerial employees.

¹² In fact, the argument of the second cited work is that the labor market segmentation was itself a significant driver of the parallel industry segmentation.

¹³ This was especially at unionized carriers, but was also to some degree at non-union ones.

¹⁴ It is an indication of the institutionalization of the high-turnover secondary labor market equilibrium in TL trucking that the ATA has published its turnover report continuously since 1996.

work hours on an irregular schedule. In one published survey of long haul drivers, 21.9% reported working 70 plus hours each week, and two out of three drivers reported working 60 plus hour weeks (Stephenson and Fox 1996). Other surveys report similar findings (Belman and Monaco 2001). A survey of long haul drivers in the Midwest found the median driver worked 65 hours, with 25% reporting 80 or more hours. In a 24-hour period, the median hours worked was 11, median hours driving 8.5, and median hours in non-driving work 2 (Belman, Monaco et al. 2005). These hours contrast to those in two industries in which there are occupations with similar human capital requirements, manufacturing and construction, which had average work weeks of 40.8 and 38.3 hours in 2004, respectively (Bureau of Labor Statistics 2002).

A related issue is that long-haul drivers are often away from home for multiple weeks at a time, with little predictability about the date of return. In the same survey previously mentioned, only 20.7% of TL drivers reported that they were home almost every day, while 28.7% of drivers in the same study reported being home less often than once every two weeks (Stephenson and Fox 1996). In the survey of drivers from the Midwest, the median long haul driver had last been home four days prior to the interview, though one-quarter had been away from home ten days or longer (Belman et al., 2005). A less tangible issue is that both drivers and firms like to think of CDL holders as professionals, in command of a big rig and responsible for its safe operation. But trucking is a service business, and a primary job function of the driver is to make shippers and receivers happy. The implications vary by customer shipping or receiving location, but this can place drivers somewhat lower than they might expect on the supply chain status hierarchy.

Of course, not every driver in TL operations faces the same conditions. The foregoing description applies to those “running the system”, or being randomly dispatched across the 48 U.S. states. Some TL operations are dedicated to the service of particular large customers, and drivers in these operations have a more restricted set of pickup and delivery locations, more regular schedules on average, and generally enjoy more time at home, as well. And some TL operations move freight between cities via trailer-on-flat-car (TOFC) or container-on-flat-car (COFC) intermodal methods. Drivers in these operations usually have regional or local runs to and from intermodal facilities, and are often home nightly, or nearly so.

Given these facts, a labor economist would expect to observe a “compensating differential” built in to the wages of TL drivers that have the worst conditions. In other words, other things equal, TL firms should offer a higher earnings level than stay-at-home jobs requiring similar human capital, to compensate for their poorer working conditions. But dissatisfaction over wage compensation levels is frequently cited as a leading reason for TL driver turnover (Cox 2004).

3.3. Buying "Effective Labor"

Perhaps a better way to think of the firm’s decision problem, that captures the nature of the driver labor market and the TL driver’s job, is to consider the nature of “effective labor” in this context. For a TL firm this is the application of labor services to move trucks to and from geographically specific customer locations on the particular time schedule desired by the firm. There are three main factors that go into the cost of effective labor in this setting. One is the cost of recruiting and training new drivers to

replace those who leave, to account for the lower productivity of inexperienced drivers, and also to account for any growth in business. A second is the cost of paying compensating differentials to drivers with the worst conditions, to slow driver exits. The third is the operational cost of making driver working conditions better. Frequently the most efficient allocation of equipment, in response to stochastic customer demands, calls for irregular schedules and little time at the driver's home terminal. When this is the case, making schedules more regular and increasing the driver's time at home is costly.

The key point is that these three cost factors can, to a significant degree, be traded off against each other, with higher expenditure in one area lowering the expenditure in another. The firm's goal can then be construed in the standard manner: it is to find the cost-minimizing mix of these factors. Historically, the best thinking among many competing TL firms appears to be that spending more on recruiting and training is a cheaper way to get the needed units of effective labor than paying more to raise compensating wage differentials or improve schedules.¹⁵

A stable equilibrium characterized by high turnover rates defines what labor economists call a "secondary labor market" (Cain 1976; Dickens and Lang 1993).¹⁶ The persistence of the secondary labor market for drivers in TL trucking since sometime in the early nineteen-eighties has occasioned much discussion in the trucking industry trade press over the years, as well as a number of academic studies (examples include (Casey 1987; Griffin, Rodriguez et al. 1992; Stephenson and Fox 1996; Griffin and Kalnbach 2002; Beadle 2004)). Through the American Trucking Associations the industry has commissioned significant analytic efforts to understand the management issues raised by a high turnover business model, and the long term demographic trends affecting the viability of the model (Gallup Organization 1997; ATA Economic & Statistics Group 2005). The major findings suggest that firms are aware of the trade-offs among the components of effective labor, and that within this framework firms adjust to changes in the conditions of the demand for, and supply of, effective labor. It appears that as a result, the labor market as a whole also adjusts, perhaps with some lags, to such changes.

A major study done by consultants at Global Insight for the ATA links the supply of truck drivers to the supply of labor for semi-skilled jobs in construction, since this type of work often represents the next best opportunity for likely truckers. The labor demands in these two industries are driven by significantly different macroeconomic factors. During the 1990's, when the derived demand for drivers was high, there was a modest premium—truckers' earnings were an average of 6-7% above a position demanding similar levels of human capital in construction. The downturn of the economy in 2000-01 created slack in the trucking labor market, but the arrival of low interest rates kept the derived demand in construction relatively stronger. As a result, for a few years the average long haul driver could expect to make less than if employed in the construction industry. By 2004 the gap had narrowed, with long haul drivers 1.5% below that of construction workers (Global Insight Inc. 2005). These facts suggest that wage levels do

¹⁵ There is actually another cost factor in "effective labor" that is non-negligible, the costs of accidents, which inexperienced drivers have at a higher rate than do experienced ones. We do not address that cost in this project.

¹⁶ Correspondingly, the ATA typically reports turnover rates at LTL firms to be in the 10% to 20%, which makes them roughly equivalent in turnover to non-trucking jobs requiring similar amounts of human capital.

adjust over time to changes in the balance of labor supply and labor demand, but the persistence of the high turnover numbers shows that the levels of compensating differential being offered are not sufficient to lower turnover to the levels typical in other blue collar jobs.¹⁷

It is well documented that the flows into and out of industry (as well as related movement of dissatisfied drivers between firms) represent a substantial cost to firms. The Upper Great Plains Transportation Institute found in 1998 that replacing one dry van TL driver conservatively costs \$8,234 and the industry wide cost total was estimated at nearly \$2.8 billion in 1998 dollars (Rodriguez, Kosir et al. 2000). The authors believe that this estimate is conservative. But it gives an idea of the magnitude of the turnover costs that TL firms must balance against the alternative costs of raising wages and/or adjusting operational and dispatching decisions, in order to lower turnover.

One might well ask whether firms have fully explored the possibilities for trade-offs among the three factors behind the cost of effective TL labor. Most firms are operating with high turnover costs, and relatively lower costs for compensating differentials and operational adjustments that improve driver lifestyles. Is it possible that some large discrete shift along the frontier could move a firm out of a “local cost minimum” in this region, to a different local minimum that might be lower in total costs?

In fact, J.B. Hunt, then the second largest firm in the industry, engaged in a highly publicized experiment with switching from a business model with high turnover and modest wage costs to one with higher wage costs but lower turnover in 1996. It took the portion of its workforce facing the worst conditions (long and irregular dispatches) and raised wages by 35%, while at the same time closing down its driver training schools (Cullen 1996; Isidore 1997). The net result was a cut in both turnover and accident rates by approximately one-half (Belzer, Rodriguez et al. 2002). However, the long run net financial benefits were not as clear (Waxler 1997), and most of the other large firms in the industry, including the one providing data for the present study, continue to train many of their new drivers from scratch, and nearly all TL firms use the high-turnover-modest-pay-premium model.

The long run dynamics of driver labor supply and demand are made more complex by the growth of the long-haul TL industry. Between 2004 and 2014 Global Insight projects it will grow at a rate of 2.2% which translates into an additional 320,000 heavy-duty long-haul new jobs. This statistic does not include the number of drivers needed to replace those who will retire during this time; the industry will need to find an estimated 219,000 additional drivers to replace the 1 in 5 drivers who are 55 years old or older and are approaching retirement. Concurrently with an increase in demand for drivers, the overall labor force’s growth rate will slow from 1.4% to .5% between 2005 and 2014 (Global Insight Inc. 2005). Another challenging trend for the industry is that to date Hispanics, who comprise the fastest growing segment in the workforce, represent a lower percentage of drivers than they do of the overall labor supply. It is possible that the conjunction of these factors means that a secular trend towards higher prices for trucking labor has begun. This in turn could shift the nature of the tradeoffs that firms face among

¹⁷ The Global Insight study used government data that does not distinguish TL from LTL among drivers for firms in long distance trucking, but TL drivers make up the predominant share of the categories they analyze.

the components of effective TL labor, and it could also dampen the long run growth prospects of the industry (Reiskin 2006).

4. Working with the Cooperating Firm

The cooperating trucking firm is a large company of national geographic scope, with divisions that operate in several of the segments of TL trucking, including long-haul random dispatch service, dedicated carriage for large customers, and intermodal services. By revenue and employment it is among the top one hundred firms in TL. The firm began as a family-owned enterprise in the regulatory era, although it has grown through multiple acquisitions, as well as internal expansions, and the original family has not been centrally involved in top management for some time.

Under family control the management culture was stable and effective, but was also, by design, relatively inward looking. It was based on long-term employment relationships with managerial and administrative ranks filled with “trucking people”, whose careers tended to be built within this single firm. A significant portion of the management started as front-line driver supervisors, or in some cases as drivers, and then worked their way up. Managers at the firm tended to learn their skills on the job, and did not see much need to look elsewhere, except to service vendors who could provide expertise relevant to particular practical business problems, such as targeted marketing surveys.

During the period between deregulation and the end of the twentieth century the firm made many major and critical strategic moves, some of which were quite daring. But the decisions leading to these moves were primarily based on the vision and judgment calls of the trucking people in top managerial positions. There was little thought of broad strategic planning in the formal sense. Early in the new millennium a new CEO, who had significant formal training in management-related areas, directed the first exercise in formal strategic planning in the firm’s history, following a process recipe provided by a major consultancy. This exercise began to increase the interest within the firm in planning as a useful activity, and also increased interest in establishing the analytic foundations for planning work.

UMM faculty began to work with the firm starting in the fall of 2002, initially on a single pilot project in the form of faculty-guided analysis by an advanced undergraduate student. The project was successful, and laid the foundation for an expanding series of faculty-guided research projects over the next two years on a variety of topics. These projects operated on a gift-exchange basis: the faculty and students contributed their time as teaching and learning functions, and the firm paid out-of-pocket expenses and provided access (under appropriate confidentiality restrictions) to proprietary business data. The core of the process involved selecting topics of both business and academic interest, and for which advanced undergraduates could provide analyses of business use, as well as generating course-level academic output. By the third year of such projects, about twenty students supervised by six different faculty members had done small projects on several continuing topics, from the analysis of exit interviews, to some initial turnover and productivity analyses, to work on the recruitment and retention of Hispanic employees.

Within the firm the linchpin of the process was a senior executive who had joined the firm from the outside, and who had significant prior experience working fruitfully

with academics. He was promoted to responsibility for a number of the aspects of human resources and driver training, and moved into his new role just as the firm as a whole was opening up internally to the importance of strategic analysis. From this initial contact UMM came to work with several other executives, at similar or higher levels of authority and responsibility, on specific projects.

On the UMM side the linchpin was an industry studies connection: the initial supervising faculty member (Burks) worked with the Sloan-funded Trucking Industry Program as a doctoral student, and as a post-doctoral fellow.¹⁸ This added academic depth and polish to trucking industry institutional knowledge he had originally begun acquiring in his youth, as a tractor-trailer driver. Burks' background, along with a passion for all things trucking-related, gave him credibility with executives, and allowed him to guide the UMM side of the relationship so that that useful business deliverables always accompanied the academic results of interest to faculty and students.

On the basis of the relationship constructed through the student projects, Burks and a second UMM researcher, biostatistician Jon Anderson, developed a small project contractually sponsored by the firm for the summer of 2004. This project began exploring the historical data retained by the firm for strategic purposes, including the analysis of the determinants of driver productivity and turnover. The larger scale design of the Truckers and Turnover Project was developed from the starting point provided by the results of this project. Burks, who is devoting a sabbatical year to the project, is the principal organizer, and he has been joined in creating and developing the substantive content of the project by the co-authors of the present paper, as well as by a number of other colleagues, who are based in several other institutions.¹⁹

5. Research Component One: Statistical Case Study of Historical Data

Research Component One is a statistical case study of some of the historical personnel and operations data of the cooperating trucking firm. There are three interrelated parts to this component. The first is building the data sets needed for analysis, the second is running some key analyses of turnover, and the third is running some key analyses of driver productivity. The goal of the first part is to take the many different data and report outputs produced by the fragmented legacy IT resources at the firm, and construct from them data sets that permit useful strategic and tactical analyses. Because the firm's IT investments began in the early mainframe era, and those investments were focused primarily on solving succeeding generations of practical business problems, the data storage and reporting functions at the firm do not lend themselves easily to strategic use. Data set assembly, documentation, and validation are consuming, and will continue to consume, a very large part of the project's resources.

The goal of the second part is to use survival analysis to map the differences in turnover by driver group, to use hazard functions to explore the different time paths of exits by driver group, and to use Cox proportional hazard multivariate regression to analyze the interaction between the various factors that can affect exits. The goal of the third part is to use panel data multivariate regression models to map the tenure-productivity curve of new drivers as they gain experience, using a fixed effects variant to

¹⁸ Burks was a doctoral student at the University of Massachusetts at Amherst; TIP was then located at the University of Michigan, and is now hosted by Georgia Tech.

¹⁹ A complete list of co-investigators appears in Appendix A.

make a first-approximation adjustment for the impact of selection on the tenure-productivity relationship. Once the panel data model is sufficiently robust, the estimated fixed effects will then be further dissected statistically.

A key (proprietary) business deliverable from this part of the project will be the assembly of the results of the turnover and productivity models to create an “expected net value of human capital” model for the investment in recruiting and training various types of drivers, who are utilized in various types of operational settings at the firm. Central academic results are expected to be generated from both the turnover and productivity models. Additionally, the analysis of Research Component Two, the panel study of new hires, will be integrated with the results of the analyses from the statistical case study. We next briefly describe the challenges and sketch a few pilot findings from the turnover and productivity analyses.

5.1. Initial Work on Turnover

The proprietary human resource data set used for initial turnover analysis was constructed from three distinct initial data files, which share the feature that each record provides information on one driver during one calendar week. The constituent files covered different calendar periods, so we utilize the calendar window during which all three overlap, September 1, 2001, through March 31, 2005. The first file, Weekly Hires, consists of some of the data elements recorded about a driver during the week he or she is hired. Drivers who are rehired during the calendar window have more than one line in this file. The second file, Weekly Separations, contains information recorded about a driver during the week that he or she separates from the firm. Drivers who are rehired and who, as a result, also separate more than once during our calendar window, have more than one line in this file. The third file, Weekly Employment, consists of one observation in each week for each driver employed during that week. Combining all three data sets gives a complete picture, week by week, of flows in, flows out, and who is currently working, for the firm.

However, there are some important limitations in these data, and a resulting major problem with analyzing them. The Weekly Hire and Weekly Separations data files contain a number of useful variables, including several key breakout variables, such as the driver’s division (e.g. dedicated, intermodal, system) and what kind of prior training or experience the driver had when they joined the firm.²⁰ Unfortunately, the Weekly Employment data file is missing these key variables. This means that at the present initial stage of the analysis we don’t have this information on the drivers who do not experience either a hire or a separation event during our calendar window. And our information is incomplete for drivers who experience only a hire or only a separation event. In particular, the division to which the driver is assigned is known prospectively at the time of the hire event. But it changes later for many drivers, and we only have the updated information in the separation event record for that subset that does depart.

To partially compensate for these problems, we take the following steps. Breakout variables that are of interest in the present study are carried forward to all observations on a given driver, from that driver’s hiring observation. This gives us reasonably accurate information on the previous trucking industry training or experience

²⁰ Not included, on the other hand, are items such as age, gender, level of formal education, or ethnic category.

of each driver (since this is not information that changes with tenure). It also tells us which division of the firm's operations a new driver is expected to be assigned to at the time of hire. Because the data on the type of work assignment is so noisy after this process, and because we would only be able to update it for those who exit, we do not pursue specific findings about the impact of the type of work on retention in the present analysis.²¹

A further implication of the data limitations is that we restrict ourselves in this initial work to the subset of drivers for which we observe a hiring event during our calendar window, because we do not have either hire or separation observations for long-time incumbent employees, and so are missing their key breakout variable values. Given an industry context in which there are large inflows all the time however, this subgroup is of significant independent interest, irrespective of what might be found if a more inclusive group could be analyzed. Also, because we are not confident that we can correctly identify all the characteristics of second or later spells of employment, we here only examine the first spell of employment during our calendar window, for those drivers who have more than one observed hiring event.²² These restrictions still leave us with a lot of data: we analyze a set of more than one half million observations covering more than 5,000 distinct individual drivers, observed during the period from September 1, 2001, through March 31, 2005.²³

Our procedure will be to first examine the survival curve for the entire set of drivers we consider here, along with the associated hazard function, which exhibits the time path of exit risk that gives rise to the survival curve. Then we will separate out the survival curves for discrete subgroups of interest, and test for differences between them, and we will also examine the hazard functions for each subgroup for useful insights. It should be noted that our analysis does not distinguish between the possible different reasons for separation. In particular, of the separation events that we observe, 76.4% are voluntary quits, while 23.6% are discharges for cause, but our survival curves and hazard functions include both.²⁴

²¹ We experimented with the following procedure. We flowed the values from the separation observation backward, to all prior observations of that particular driver, for the variable recording division to which the driver is assigned--for those drivers who have an observed separation only. (This overwrote the forward-flowed divisional assignment data from the time of hire for those separated drivers for whom we observe the hire event.) This gives us improved information on those who separated, but at the cost that noise is differentially left in the observations on those who do not separate. The results were not credible, so we abandoned this part of the analysis until further information can be added to the data set.

²² This does not prevent us from examining rehires, as a significant number of the first spells we observe are of rehired drivers.

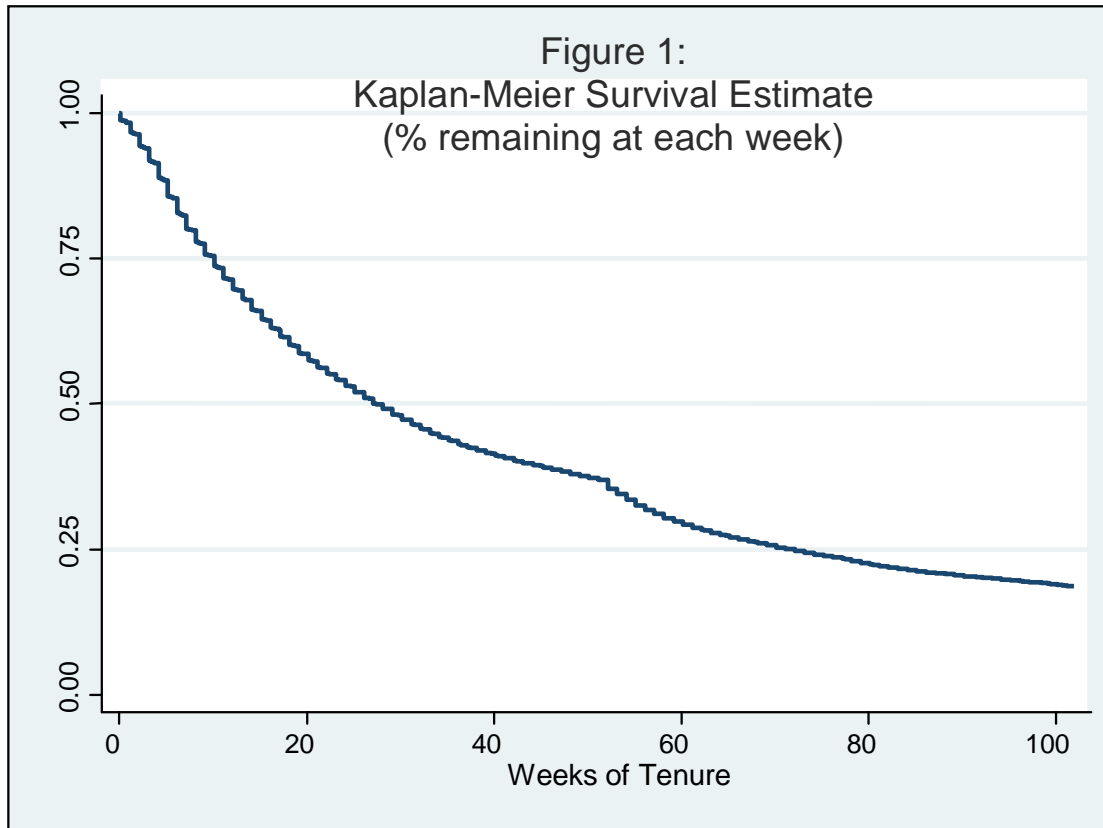
²³ The precise number of drivers and observations is suppressed for confidentiality reasons.

²⁴ The primary statistical methodology is survival analysis. Standard descriptive and analytical methods are problematic when the key dependent variable (here, the length of job tenure) is a time period, since ongoing spells observed at any given point in time are censored: they continue for an unknown further period. Instead, a conditional probability approach is needed, to correctly take into account the statistical information contained in censored observations

Kiefer, N. M. (1988). "Economic Duration Data and Hazard Functions." *Journal of Economic Literature* **26**: 646-679,

Cleves, M. A., W. A. Gould, et al. (2004). *An Introduction to Survival Analysis Using Stata, Revised Edition*. College Station, TX, Stata Press.


5.1.1. Descriptive Results for All First-Hire-Event Employment Spells



We begin by examining the survival pattern for the first observed employment spell of all drivers having a hire-event during our calendar window. Figure 1 displays the central results. The vertical axis indicates the percentage of the population initially entering employment that remains after each amount of time on the job, shown on the horizontal axis in weeks from the start of employment.

Some key qualitative facts emerge from this picture. First, turnover rates do look extremely high. At 10.1 weeks 25% of the population is gone, 50% have left by 29.1 weeks (the median survival time), and 75% have departed by 75 weeks. Second, there is a leveling off of departures in the second six months on the job, followed by an acceleration at the end of the first year. This is consistent with the fact that most of the trainees observed here who undergo the firm's full training program sign a contract to pay back about half the cost of training (several thousand dollars) if they do not complete a year of service after training. Plus, the job options within trucking are more plentiful for drivers with a year of experience. The surprise, in fact, is that so many new drivers leave before the first year is up. Clearly, these departures cause both the firm and the drivers to incur real costs.

Further insights may be obtained by examining the hazard function for this group of drivers. The vertical axis indicates the probability of leaving during any particular week shown on the horizontal axis, given that the driver made it to the beginning of the week. (Or, to be slightly more careful, the vertical axis shows a “departure rate”, because it is the conditional probability just described, divided by the number of analysis-time units contained in each unit on the horizontal axis. In our case the denominator is one, so



the rate is also a simple conditional probability.)²⁵ Here the differences in risk of departure are shown more clearly. Exit risk is highest at about 6 to 8 weeks, which is approximately when new trainees first pull a load by themselves, without the assistance of an instructor-driver in the cab. Once drivers make it past this stage, exit risk declines sharply, until the one-year mark is reached, when separation risk spikes to almost the same level as at the beginning. Drivers who make it to the end of two years are essentially self-selected to have a high likelihood of turning out to be longer-term employees.


5.1.2. Descriptive Results by Level of Previous Experience or Training

Drivers who are hired by the cooperating firm arrive with different levels of prior training and prior experience. In Figures 3 and 4, and Table 1, the differing performance of these subgroups with respect to retention gives rise to separate survival curves and hazard functions. The best retention is exhibited by the small group (4% of the total) of

²⁵ Formally, the hazard function is defined to be the ratio of the density of exit events to the survival function of exits, or $h(x)=f(x)/S(x)$.

rehires. This can be observed from the fact that their survival curve is well above the curves of the other subgroups, and is quantified in Table 1. We can see in the table that rehires have the longest time period of any group at which 75% still remain (almost 4 months), and at which 50% still remain (over 5 years). Rehires also have a retention period for 25% of the starting population that is so long that it cannot be meaningfully calculated in our data. This is not surprising—rehires are the self-selected subset of drivers who are not only experienced drivers, but who have worked at least once already at the cooperating firm. Having explored other opportunities, they now choose to return to this firm as their best current option.

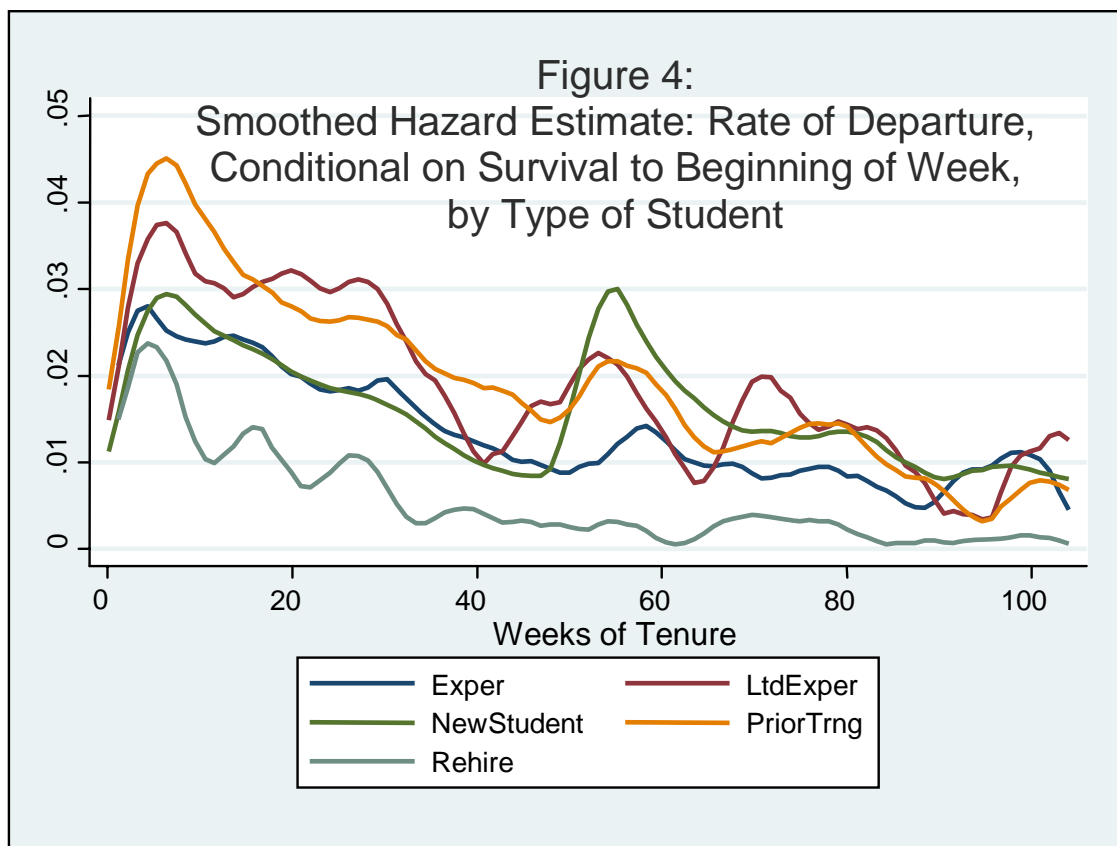
The hazard function for these drivers is distinctive, as well. It shows a modest spike in exit probability early, with falling exit risk thereafter, and also a very distinct periodicity during the first year, which likely reflects the incentive effects of the firm’s quarterly bonus system. Rehires are eligible for the firm’s quarterly bonus immediately upon starting work, and also have experience with the incentive provided by the particular bonus system offered by the firm. The periodicity in the rehire hazard function suggests drivers in this group who may consider leaving during the first year are likely to wait until they have completed a quarter and have qualified for the bonus, before separating. Also noteworthy, and sensible, is that there is no “first-year-effect” spike in



the rehire hazard rate—this effect in the aggregate hazard function is entirely due to the behavior of other subgroups.

Next consider experienced drivers. These are students who have significant levels of over-the-road tractor-trailer experience with other employers, before coming to the cooperating firm. Like rehires, they only have to take a refresher training course that takes a few days, instead of the multiple-week basic training course all other drivers new to the firm are required to pass. Their retention performance is not as good as that of the rehires, but it is still well above that of the lowest groups, with 75th, 50th, and 25th percentile retention periods of 10.4, 29.4, and 98.3 weeks, respectively. Their hazard function shows the usual pattern of an early peak, with later declines, and appears to have a muted version of the periodicity seen in rehires. This would make sense, as experienced drivers are eligible for the bonus system immediately, but don't have as much experience with its incentives as rehires.

The next item to note is akin to Sherlock Holmes' famous observation about the mysterious behavior of the dog in the night. The dog didn't bark when it should have, and correspondingly one would expect new students with no prior background of any kind in trucking to have different (and in particular, poorer) retention performance than



experienced drivers.²⁶ But in these data both new students who are learning the industry from scratch, and experienced drivers who are new to the cooperating firm, have closely similar retention behavior for nearly the first entire year of employment. New students actually do slightly better than experienced drivers, near the end of the first year. At that point their hazard function spikes very sharply, and their performance drops below that of

²⁶ The mysterious behavior (in “The Silver Blaze”) was that the dog did not bark when someone removed a valuable race horse from the barn, which was a clue to the thief’s identity.

experienced drivers. This is likely associated with the facts that their training contracts are completed and they then have enough experience to easily switch trucking firms if they desire. Since new students are by far the largest group (73%) of drivers for whom we observe a hire event, their behavior is very important in determining that of the entire aggregate driver population. Thus the size of their initial aggregate spike in exit risk, as well as that after a year of service, both strongly shape the aggregate survival curve and hazard function.

As it turns out, a Chi-square statistical test of the significance of the difference in overall survival performance between new drivers and those with experience at firms other than the one providing the data shows that experienced drivers do better overall, at the 5% significance level ($p = .018$). But, as Table 1 shows, the effect is all driven by the one-year exits of new drivers, and the magnitude of the effect is much smaller than the difference between either of these groups and rehires.²⁷ For instance, 50% of the rehire group is estimated to still be at work for the cooperating firm 5.48 years after the hire event we observe, while for drivers with experience at other firms it is only 6.8 months, and for new students it is essentially the same, at 6.9 months.²⁸ At longer durations of employment we see a modest difference: 25% of the drivers with experience at other firms still remain at 22.6 months, while it is only 16.9 months for the same proportion of

new drivers.

Last, consider the retention performance of the two final groups: drivers with some prior experience, and those with some prior training. Both these groups are identified by the driver recruiting staff at the

Drivers for whom a "hire event" is observed; N > 5,000	Estimated Job Tenure in Weeks			
	Percent of Drivers	75% of Drivers Remaining	50% of Drivers Remaining	25% of Drivers Remaining
All Drivers	100%	10.1	27.4	72.1
Rehire	4%	16.6	284.7	n/a (long)
Experienced	8%	10.4	29.4	98.3
New Students	73%	11.1	30.1	73.1
Ltd Experience	3%	8.1	21.1	53.1
Prior Training	14%	6.7	18.1	49.1

cooperating firm as having some background in trucking, but not enough to qualify the student to take only the short training course for fully experienced drivers. To extend the previous allusion, here is a dog barking loudly—these two groups do quite badly, by comparison to students wholly new to trucking. The job tenure lengths for the retention of the 75th, 50th, and 25th percentiles of students with limited driving experience is 1.87, 4.94, and 12.25 months, respectively. This tells us that only 25% make it to the completion of their one-year-service-after-training employment contract; the other 75% are incurring a multi-thousand dollar debt in order to leave early.²⁹ Students with only some prior training, but no prior experience, do even worse, with retention periods for the 75th, 50th, and 25th percentiles of only 1.58, 4.18, and 11.33 months, respectively. So less

²⁷ The pairwise differences between rehires and new drivers, and between rehires and experienced drivers, are both significant—the Chi-square p-values for Type 1 error are zero to four decimal places.

²⁸ The base time unit for the statistical analysis is weeks, so months are everywhere calculated as weeks divided by 4.33.

²⁹ Except for those who are hired by a rival firm that is willing to pay off their indebtedness—something which is known to occur in this labor market.

than one quarter of these students complete their training contracts. (The difference between these two groups is significant by the Chi-square test, at the 5% level (p-value of .045.³⁰))

Why should these students be at the bottom of the performance ranking, when normally prior training or experience would be expected to improve retention? A reasonable hypothesis is that it has to do with the distinctive characteristics of a high-turnover, secondary, labor market. In this type of market there is always demand for drivers at some job or other. So someone with prior experience of any kind, as well as the graduates of any of the many commercial driver training schools, can get some job, as long as they have a CDL. It may not be a very desirable job, but it is possible to accumulate experience if one is willing to put up with some of the poorer working conditions available in an industry segment known for having poor conditions on average. In this context, coming to the cooperating firm, and being willing to assume the debt contract that accompanies the full training program, is a bad signal. There may be many specific reasons outside a prospective driver's control that lead to such a decision. For example, the student could have experienced some kind of family event that stopped their prior training before the CDL exam, or caused them to quit a prior job quickly. But on average, students with some prior training or some prior experience are likely either to be job switchers who just couldn't do better for the time being, but who will be looking to leave as soon as possible, or to be job candidates who were unsuccessful at someone else's training course, or were otherwise judged inadequate by other firms. Either of these reasons means the student is more likely to depart.

5.2. Pilot Work on Productivity

The pilot work on productivity utilized a different set of data files from the cooperating firm than did the turnover work described above. We began with two data files, one containing basic information (especially, hire date, and separation date, if any) on all the drivers who had separated during the period of one year (for example, in some of the pilot work we used 2003), and the second, extracted at the end of the that year, containing similar information on all currently employed drivers. Then two separate additional files containing demographic information, and racial and ethnic identity from voluntary EEOC employee disclosure forms³¹, were added.

Merging these using the internal employee number ("driver number") as an identifier immediately caused problems. It turned out that driver numbers are not unique, but are recycled on a regular basis, so we had to delete some duplicate cases that really represented different drivers.³² "Hire date", a key variable for survival analysis also turned out to be problematic. As one might expect in a high-turnover setting, a small but significant number of drivers become re-employed, some having as many as four or five successive employment spells. The problem was that drivers gone for less time than

³⁰ The pairwise differences between either of these groups and any of those with better retention performance is highly significant—the Chi-square p-values for Type 1 error are zero to four decimal places.

³¹ "EEOC" is the Equal Employment Opportunity Commission.

³² For the pilot work we did not want the responsibility of making use of social security numbers, although a secure method for making use of the relevant identification information has been developed for later work.

some threshold (six months at one point, but varied over time) kept their original hire date, while those gone longer were assigned a new one. The latter fact made it impossible to distinguish rehires from new drivers with recycled driver numbers.

To do a productivity analysis the key addition to the records already described was information from the firm's payroll records, which provide a week-by-week compilation of the items added to (or deducted from) each employee's pay, with each such transaction constituting a line of data. The taxes and fringe benefit co-pays were in a separate data source we did not try to access, but even so the initial files had as many as 44 transactions per driver per pay period, with more than one million lines of data per file. We proceeded to document the different variables that contained coded information about the driver's work assignment and pay structure, consulting subject-matter experts at the firm

regularly. Each variable could take on multiple values, the meanings of which to some degree changed over time as operational needs changed. In addition we began to document all the meanings of the values of the key variable specifying




Figure 5: Variance! Miles per week by week of driver tenure.

what type of transaction each line of the payroll file contained. There were several hundred distinct values of this variable, including values denoting several different types of mileage pay, dozens of types of lump sum pay for specific tasks, dozens of types of pay advances and pay deductions, and so on.

After documentation, we next "rolled up" the payroll file. We sorted the file by driver and pay-week, and then accumulated all the transaction-level information we were interested in having on a weekly basis into new variables, so that the last transaction in each driver-pay-week record contained cumulative information for the week. The kinds of information in the resulting records included such key items as the total (paid) miles and the amount paid for them, the total number of dispatches. Also included was information on various kinds of ancillary activities when they generated a pay

transaction, such as paid customer stops, pay supplements for very short runs, paid maintenance delays, and so on. The payroll data thus provides a very rich set of information about what each driver does during each week.

However, the payroll file records what drivers are actually paid for, which is in general a subset what they actually do. So, for instance, the first pickup stop and first delivery stop on each loaded dispatch are not separately compensated. Extra pickup or delivery stops are paid when they occur on long-distance random-dispatch loads, but only some of the time when they are on a scheduled run dedicated to a particular customer, that is engineered to have multiple stops. Most drivers are primarily compensated by the mile, and these drivers are paid miles for all their dispatches, which normally includes loaded miles, plus miles pulling an empty trailer, repositioning for a new load, and also any bobtail miles (i.e. without a trailer). However, drivers generally run more miles than those for which they are paid. Paid miles are based on a least-distance routing algorithm which is historically standard in the industry, but which undercounts the actual miles by several percent (recent guesstimates by managers at our firm for the average undercount range from 4% to 6%).³³ Despite these limitations, the payroll data provide a very useful starting point for the productivity analysis.³⁴

³³ This is in part because the standard algorithms are to and from standard reference points, and given the circuitry of the road network, this undercounts actual miles on average. It is also because drivers are responsible for selecting a practical route for a large loaded tractor-trailer, which is often more circuitous than the least-distance version. In addition, drivers may choose to deviate for other reasons (for example, to run on a turnpike where the salt trucks will be out at night in the Pennsylvania mountains in winter, as opposed to a non-toll highway on which such services are more uncertain), as long as they don't exceed certain percentage standards for excess miles, and meet delivery schedules.

³⁴ For later work it is expected supplemental information will be added from a separate operational events data set also maintained by the firm. It is not the place to start because it has its own limitations, and also because it is about an order of magnitude larger than the payroll data set.

5.2.1. Descriptive Productivity Results for Inexperienced Long-haul Random Dispatch Drivers

We began our pilot work with a subset of drivers for the years 2002 and 2003. The subset is those drivers who were inexperienced at hire (i.e. who had to take the full training course offered by the firm), who were assigned to drove solo (as opposed to in a team) on long-haul random dispatch runs, and who were in their 5th week to 156th week of tenure with the firm.³⁵ This gave us more than 100,000 pay-week observations on more than 2,000 drivers. Examining the key dependent variable, miles per week, we observed very high variance (see Figure 5). In particular, there were negative values and also very high positive values. The former turned out to be due to mistaken pay being charged back against a driver's earnings, and the latter to a small number of drivers from the firm's early days who were permitted to accumulate vacation earnings over several years and were being paid upon retirement. We decided to trim the extremes, and had to choose whether to leave in zero-miles weeks or use only positive-miles ones, and what upper bound to use.

The actual maximum number of miles that a solo driver could legally run during this period, given state speed limits and Federal Hours of Service Regulations for operators of commercial vehicles was about 4,000 per week. But during at least part of this period, until the practice was ended, some drivers at the firm were paid for their runs only after they submitted completed paperwork for each dispatch. This meant that if they held their paperwork they could have one (or even two) weeks in a row with zero paid miles, and then a week with very high miles. We decided to trim only the negative values, leaving zero-miles weeks in, and also trimmed values over 6,500, after looking at the distribution of the upper tail.

³⁵ Drivers begin receiving mileage pay when they first pull a load on their own, without a trainer in the truck with them, and the earliest this occurs is about the 5 the week.

Further examination showed that almost 20% of our observations were of zero miles pay weeks. So we first trimmed out all the pay-week observations that were associated with any payroll transaction that could exogenously cause the driver to either miss work, or be paid on a non-mileage basis. This included pay weeks with disability pay, vacation or holiday pay, salary (sometimes paid to driver-trainers), lump-sum training pay, and the like. Then we discovered that the payroll system was generating dummy paychecks for drivers who had separated from the firm, for several weeks after separation when the driver left owing money, e.g. for things like cash pay advances or purchases at a store at a company terminal. When all of these cases were trimmed out, we reduced the number of weeks with zero miles substantially, but 6% of our observations remained with zero miles per week.

Figure 6 exhibits a simple descriptive version of the tenure-productivity curve for this subset of drivers. Even after all the trimming, the remaining weeks with zero miles affect the mean values quite significantly. Without zero-miles weeks, the initial increase to full productivity is achieved at about nine months, whereas with zero-miles weeks it is nearer to a year. There is a sharp drop in the curve at one year of tenure with zero-miles included.

This is undoubtedly related in some way to the fact that drivers with one year of experience can more easily switch firms, and also to the fact that most of the drivers in this subset, all of whom have taken the firm's training, assume a debt of several thousand dollars for




Figure 6: Average miles per week by week of tenure, with and without 0-miles

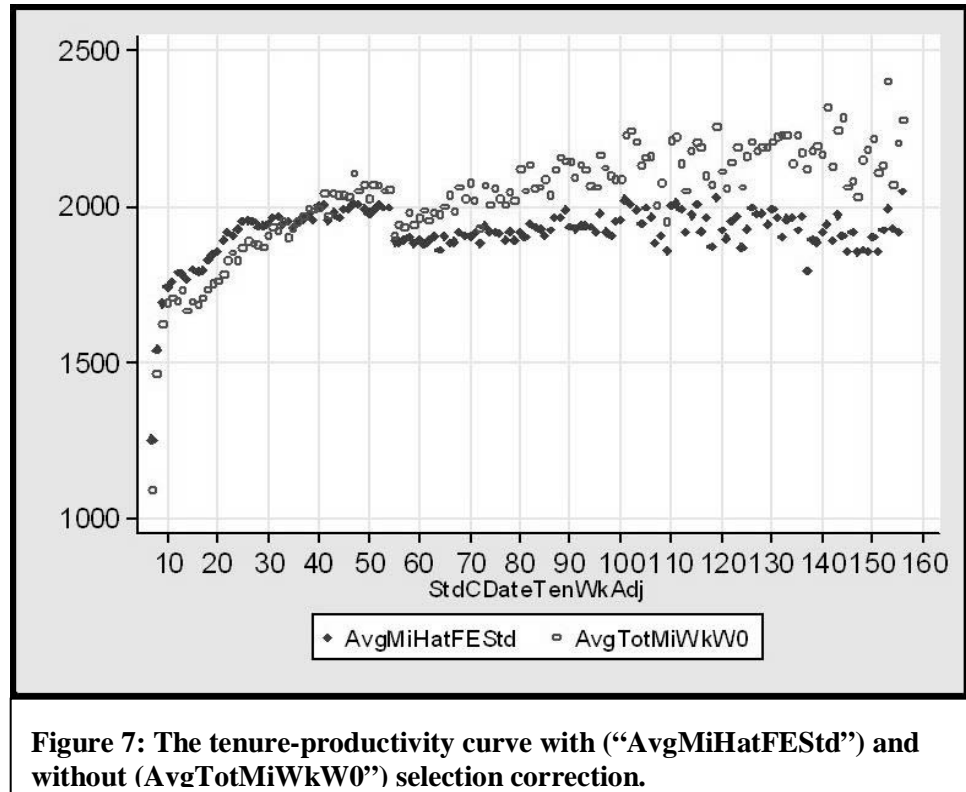
its cost, which is forgiven at the end of a year of service after training is completed.³⁶ The balance of the pilot analysis keeps all the zero-miles weeks in the data set, but a goal

³⁶ New drivers also earn their first week of paid vacation at this point, but that cannot be the reason for the drop in the averages, as those weeks have been removed from the data.

of the full statistical case study is to dig deeper into this phenomenon and develop better evidence on whether they all should be included when analyzing productivity, or not.

5.2.2. The Impact of Selection on the Tenure-Productivity Curve

A major goal of the statistical investigation of productivity is to analyze the true causal effect of increasing tenure on expected miles per week (the “treatment effect”), while accounting for any impact the high turnover rate might have (the “selection effect”). *A priori*, a reasonable hypothesis would be that drivers with lower productive capacity would be more likely to leave at any given level of tenure.



In order to test this hypothesis we start by running a “fixed effects” panel data OLS regression model, with total miles per week as the dependent variable. We use all the independent variables from the payroll data set that plausibly measure exogenous factors that affect productivity. These include driver tenure (a linear term and as many higher-order terms as prove significant), the number of dispatches (linear term plus those higher-order terms that prove significant). We also include variables such as the number of short-haul pay supplements, the number of paid maintenance delays, as well as dummy variables for each week of the calendar year (to capture any time-period effects or time trends), and a dummy variable for each terminal at which drivers are based (to capture any geographic effects of the home base). We use robust standard errors. The pilot version of this model has an adjusted R^2 of .66.

The fixed effects model constrains the coefficients on all the independent variables to be the same across all drivers, but permits each driver to have his or her own regression plane intercept, or constant. This constant, or “fixed effect”, which is estimated by the regression, may be thought of as a measure, specific to the model and the data, of the degree of “job match” between the driver and his employment at the firm. In the context of the model, it is the number of miles which the driver “brings to the job each week” (which can be positive or negative), according to the model estimate.

Allowing this specific flexibility in the regression model provides a statistical adjustment for the relative speed with which drivers of high and low job match turn over, in predicting the miles each driver will operate per week. So, we take the predicted values from the model, and compute the average of these values over each week of tenure on the job. These averages, when graphed, produce a “selection-corrected” tenure-productivity curve. In Figure 7 we compare this new tenure-productivity curve with the simple descriptive version of the same curve we exhibited in Figure 6.³⁷

When the selection-corrected tenure-productivity curve is below the old curve, it says that the true effect of tenure on productivity is smaller than it appears in the simple descriptive case. This implies that drivers with poor job matches are leaving differentially faster than those with good ones, which is in accord with our hypothesis. By contrast, when the selection-corrected tenure-productivity curve is above the old curve, it says that the true effect of tenure on productivity is larger than it appears in the simple descriptive case. This implies that drivers with good job matches are leaving differentially faster than those with bad ones. It is a bit of a surprise to see that the pilot results suggest that highly productive drivers are leaving faster than less productive ones early in their tenure, when the firm has its highest rates of separations. According to Figure 7, our initial hypothesis is only true from about the ninth or tenth month on, and not before. A major goal of the full statistical case study is to verify or refute this tentative finding.

6. Research Component Two: Panel Study of New Hires

Research Component Two is a study of approximately 1,000 new driver-trainees who have been recruited by the firm to start their education at a specific training school operated by the firm. The basic design of the panel study is quite straightforward in conception, although it is extremely labor intensive and costly to carry out. A large amount of data is being collected on each driver trainee, starting with an initial contact while each is in the first phase of training, and then data will continued to be collected over two years of the trainee’s work life at the firm, or until exit.

The follow-up data collections include these elements: (1) a follow-up paper survey for the driver mailed to his or her home every six months, for two years, as long as they stayed employed by the firm; (2) an exit survey for the driver mailed to the driver’s home soon after their separation, if it occurs during the first two years; (3) a weekly survey with two questions sent over the satellite unit to the driver’s truck (this is planned, but IT issues mean it may not be implemented soon, or even at all)³⁸; (4) an initial survey (and consent form) for the driver’s spouse or significant other asking about family issues mailed to the driver’s family soon after the driver enters the study, (5) a follow-up survey to the driver’s family mailed every six months up to two years, and (6) an exit survey for the driver’s family mailed upon driver separation, if it occurs during the first two years. As subjects are informed as they enter the study, a cash gift of \$5 is included in each survey mailing, with the goal of increasing the response rate. Finally, the driver’s on-the-

³⁷ Both curves are for the case in which zero-miles weeks that cannot be specifically explained away are retained in the data set

³⁸ The two questions are: “How happy are you with your job right now?” (Likert scale response), and “How many miles do you expect to run next week?”

job performance data will be collected as part of future updating of the master data files for the turnover and productivity studies of Research Component Two.

A central project design goal is to perform a multivariate statistical analysis of the relationship between all the factors that are being measured, and the success on the job of the trainees, where employee success is measured first by the length of time they are retained by the firm, and second, by their productivity on the job.³⁹ While the researchers (and the firm's managers) start with a number of hypotheses about what might matter in predicting each of these outcome variables, this research component is essentially a large scale empirical investigation, designed to let the data tell us which factors matter statistically, and which do not. In this regard, it is particularly exciting that we have the opportunity to employ both traditional measurement instruments, and a selected set of economic experiments.

One methodological point should be noted. We plan to look at the relationships among the various measures we are collecting on each subject, as there is little evidence in the literature on this topic for many of our measures, let alone evidence using the subject population from which we are drawing our participants. If this were our only goal it would be critical vary the order in which the different measures are implemented during the initial data collection, as it is quite possible that order effects could be important for some of these relationships. Such variations are a standard feature of many studies in experimental economics.

However, since the central design goal is to examine the predictive power of the various measures with respect to on-the-job outcomes, a countervailing methodological need is to present the same exact set of stimuli to each subject, so that the relationship to the outcome variables is always the same. In addition, the complexity of administering the data collection process at the field site makes it very costly to do vary the order, once we have a sequence that fits the schedule. For this reason the sequence of measurements in the initial data collection has not been varied, and the only variation we contemplate introducing later in the data collection process is to alternate the order of Session 1 and Session 2; a final decision has not been reached on this.

6.1. The Driver Training Process

The initial data collection process takes place on a Saturday in the middle of the two-week training process for new driver trainees. We next describe the training process, to provide context for understanding where the initial data collection fits in the new driver's work life.

The two-week training includes a large dose of classroom work in which students are introduced to the firm and learn essential facts about the equipment they'll be operating, the regulations governing commercial vehicle operators and operations, map reading and course-plotting, and the safety rules and procedures specific to the firm. It also alternates time in truck simulators with time driving in an actual Class 8 tractor-trailer.⁴⁰ Trainees first drive bobtail, and then with a empty and a loaded trailer, on the

³⁹ The analysis has the potential to be extended to include safety performance, but that is not part of the present project.

⁴⁰ By definition, Class 8 vehicles have a maximum gross vehicle weight (GVW) of between 33,000 lbs. and 80,000 lbs. The units utilized by most truckload carriers, including the cooperating firm, are at the top of this range, and have a maximum GVW of 80,000. lbs.

school's property. Once the trainee has qualified for a commercial driver's license learning permit, and is judged ready by the instructors, they quickly begin driving on the actual roads and highways surrounding the school, with a driver-trainer in the right seat of the tractor.

The simulators have two purposes. One is to speed familiarization with the basic features and operational characteristics of the Class 8 tractor-trailer. An example is how to correctly shift a ten-speed transmission that requires frequent double-clutching, and that also requires attention to engine RPM's in comparison to road speed, when down-shifting. Another basic feature drivers are first exposed to in the simulator is how to maneuver a rig that is about 65 feet (19.8 meters) long, and that includes a trailer that is 53-feet (16.2 meters) long, around corners on city streets. A second function of the simulators is to give drivers practice at maneuvering in dangerous settings which could never be practiced in real life, such as how to respond to a blow-out on a steering axle tire while on glare ice on a freeway.

Trainees who complete the initial two-week training process are officially hired by the firm on their date of completion, after passing a basic skills qualification test. Once hired the driver goes back to his or her home base, at one of the firm's terminal locations. While the content of the Commercial Driver's License exam is federally regulated, the exact rules governing driver training vary from state to state. So, depending on their proficiency level and on the state in which they will be based, trainees may also take the CDL exam at the training school, or they may do so later, at or nearby their home base. Whether they have the CDL yet or not, all new drivers have a second training phase, during which the trainee works for between one and three more weeks, actually hauling freight from their home base, but with a certified driver trainer in the right seat of the tractor. When the trainer judges the trainee ready (and he or she has acquired a CDL), the new driver is assigned a tractor of his or her own, and goes to work on their own.

As is mentioned above in previous sections, it is typical for many of the TL carriers that train all or most of their new drivers to ask trainees to sign a contract for their training in the form of the extension of credit to the driver-trainee for a substantial portion of the cost of the training (several thousand dollars, an amount that is about half the full average cost per trainee), and this firm is no exception. The contract provides both phases of the training, including access to the classroom, instructors, and trucks, plus needed classroom supplies. For the first phase of training, also covered is transportation to and from the training school, a hotel room while at training, and lunches during the training day. The trainee owes no payments on the credit extended as long as they stay employed by the firm, and the debt is fully forgiven if the trainee completes a specific period of service after the point they begin working on their own (either one year or eighteen months). Trainees who do not complete the period of service, however, become legally liable for repayment of the amount of credit extended.

6.2.The Initial Data Collection Process

The first phase of driver training at the firm's school at which the initial data collection is done begins on a Saturday, and concludes with graduation on Friday, two weeks later. When they first arrive at the school, trainees receive a one-page flier, advertising that there will be an opportunity coming up the next weekend to take part in

the University of Minnesota study, and that they can expect to earn cash by participating. On the Saturday of data collection Burks conducts an informed consent process, in which he introduces himself as a former driver turned university researcher, and explains the goals and procedures of the study to the potential subjects. The study is governed by the standards of the Human Subjects Protection Committee of the University of Minnesota's Institutional Review Board (IRB), so a formal consent document is used which spells out the risks and benefits of being in the study. This document reinforces the following key facts: that subjects will receive an initial \$20 in cash "thank you gifts" for taking part, and have the opportunity to earn substantially more, that the money paid to them is coming not from the firm but from two non-profit foundations, that the data collected will be kept confidential from their managers at the trucking firm under University rules for protecting research participants, and that participation is voluntary.

The data collection at the training school is set up as two two-hour blocks of time, plus a short break in between, spent doing tasks with the researchers, either on computers or with paper and pencil. Classes range from 35 to 70 students, and the largest group that can be accommodated at one time for data collection is 32 subjects, so each class is broken into two groups. The first group works with the researchers from early to late morning, and the second group from late morning to mid-afternoon. The class day is structured so that those potential subjects who choose not to take part do not have extra training available, but instead can spend the extra time in the break room. The buses that transport students to and from the hotel bring everyone at one time in the early morning, and take everyone back at one time in the afternoon. Given the monetary compensation being offered, the relatively low opportunity cost of taking part, and the credible guarantee of confidentiality from the University, between 85% and 90% of students have been choosing to join the study.

Except for one instrument (the Educational Testing Service's test of Quantitative Literacy, labeled "Numeracy" in our schedule), all the instruments are administered on a wireless network of older Dell notebook computers, with a newer Dell notebook as the master controlling computer. The software is z-Tree (Fischbacher 1999), the toolbox for constructing computer-administered economic experiments and surveys developed at the Institute for Empirical Economics at the University of Zürich.⁴¹ The data collection takes place in a temporary computerized experimental economics laboratory that the research team sets up in one of the classrooms at the training school on Friday, before the data collection on Saturday, and then takes down at the end of the event. The team hand-built a set of cloth dividers hung from frames adapted from portable garment racks, which are installed between all the computer stations in order to separate subjects during the data collection. Because the data collection exceeds the endurance of notebook PC batteries, extra power outlets were installed in the walls and a set of portable extension cords are deployed as part of the laboratory set up. The wireless network is part of a VLAN (virtual local area network) set up by the firm's IT support personnel under a data security agreement approved by the University of Minnesota Human Subjects Committee. The VLAN separates the computers on which the confidential data is being

⁴¹ Because we are running a licensed adaptation of a non-verbal IQ instrument that requires the display of a large number of graphics files, each of which is a scanned image of a page with a pattern-matching task, we are running a beta version of zTree that has been extended to handle graphics files. Our special thanks to developer Urs Fischbacher for providing this new version in time for our project.

collected from the rest of the firm's network, and provides secure data storage until the initial data collection is complete, and the project moves back to the UMM campus.

Table 2 shows the time sequence and schedule for the actual data collection event of February 25, 2006. This class of trainees was relatively large, at 59 students. 51 out of 56, or 91%, chose to take part.

Table 2: Time Allocation 2/25/2006

Start Time	Budgeted Time	Total Time	Informed Consent Process
7:13	7:10	x	Start Informed Consent Process
59	x	X	How many people in the room?
7:29	7:30	0:16	Finish Informed Consent Process

Session 1a			
Start Time	Budgeted Time	Total Time	Activity
7:30	7:30	0:09	Check-In
7:39	7:40	0:04	Information
7:43	7:42	0:19	Activity 1: Prisoner's Dilemma Computer generated Pay on beliefs for what percentage of people will send \$5 as person 1, how much person 2 sends if person 1 sends \$0, and how much person 2 sends if person 1 sends \$5, \$1 each
8:02	8:09	0:39	Activity 2: MPQ No payment
8:41	8:47	0:12	Activity 3: Risk/Loss Aversion 1 Question, #23 Green by drawing out of a bowl Everyone paid for their selection
8:53	8:58	0:27	Activity 4: Demographics No payment
X	9:16	X	Activity 5: Big Red Button -Computer generated
X	X	X	Release Participants
9:09	X	X	First person left at this time
9:20	9:26	X	Last person left at this time

Session 2a			
Start Time	Budgeted Time	Total Time	Activity
9:37	9:40	0:08	Check-In
9:45	9:45	0:02	Information

9:47	9:47	0:09	Activity 1: Time Preference 2 Subjects, #12 and #19 1 Question, #18
9:56	9:57	0:47	Activity 2: IQ Start Time of Test 10:02 2 Subjects, #8 and #21 Pay \$1 for each correct answer Pay on "pre" and "post" beliefs, \$2 each
10:43	10:30	0:25	Activity 3: Numeracy Start Time of Test 10:44 2 Subjects, #8 and #18 Pay \$2 for each correct answer Pay on "pre" and "post" beliefs, \$2 each
11:08	10:55	0:08	Activity 4: Ambiguity 1 Question, #7 Blue by drawing out of a bowl Everyone paid for their selection
11:16	11:05	0:12	Activity 5: Hit 15 Points Computer generated
11:28	11:28	0:12	Activity 6: Risk, Cooperation, Impatience Questions No payment
X	X	X	Release Participants
11:35	X	X	First person left at this time
11:40	11:40	X	Last person left at this time

Session 1b			
Start Time	Budgeted Time	Total Time	Activity
12:15	12:30	0:10	Check-In
12:25	12:40	0:03	Information
12:28	12:42	0:23	Activity 1: Prisoner's Dilemma Computer generated Pay on beliefs for what percentage of people will send \$5 as person 1, how much person 2 sends if person 1 sends \$0, and how much person 2 sends if person 1 sends \$5, \$1 each
12:51	1:09	0:35	Activity 2: MPQ No payment
1:34	1:47	0:13	Activity 3: Risk/Loss Aversion 1 Question, #21 Blue by drawing out of a bowl Everyone paid for their selection
1:47	1:58	0:32	Activity 4: Demographics No payment
X	2:16	X	Activity 5: Big Red Button -Computer generated
X	X	X	Release Participants

2:04	X	X	First person left at this time
2:19	2:26	X	Last person left at this time

Session 2b			
Start Time	Budgeted Time	Total Time	Activity
2:25	2:40	0:09	Check-In
2:34	2:45	0:02	Information
2:36	2:47	0:11	Activity 1: Time Preference 2 Subjects, #17 and #16 1 Question, #7
2:47	2:57	0:39	Activity 2: IQ Start Time of Test 2:53 2 Subjects, #17 and #22 Pay \$1 for each correct answer Pay on "pre" and "post" beliefs, \$2 each
3:26	3:30	0:28	Activity 3: Numeracy Start Time of Test 3:30 2 Subjects, #17 and #16 Pay \$2 for each correct answer Pay on "pre" and "post" beliefs, \$2 each
3:54	3:55	0:11	Activity 4: Ambiguity 1 Question, #22 Blue by drawing out of a bowl Everyone paid for their selection
4:05	4:05	0:16	Activity 5: Hit 15 Points Computer generated
4:21	4:28	0:16	Activity 6: Risk, Cooperation, Impatience Questions No payment
X	X	X	Release Participants
4:29	X	X	First person left at this time
4:37	4:40	X	Last person left at this time

While it is theoretically possible for a subject to end up with only the initial \$20 “thank you gifts”, in fact this has never happened. The actual payouts depend in part on the use of stochastic devices (plastic bowls filled with colored and/or numbered poker chips), as described in the account of each data collection activity, below. So far the average total earnings for the four hour period of the initial data collection has ranged from about \$45 to \$55, with the median usually slightly less. By design we want most subjects to be pleased about their interaction with the researchers, given the follow-up contacts we expect to make with them.

We next turn to a brief description of each item in the sequence of data collection events.

6.3. Session 1 Data Collection Events

6.3.1. Prisoner’s Dilemma

Our version of this experiment is a sequential and strategic one. Person 1 (the first mover) and Person 2 (the second mover) each are allocated \$5. Person 1 can send either \$0 or \$5 to Person 2, and Person 2 can respond by sending \$0, \$1, \$2, \$3, \$4, or \$5 back. All funds sent are doubled by the researchers. Each subject makes both an unconditional decision for the first mover role, and a conditional one for the second mover role (first how to respond to being sent \$0, and second how to respond to being sent \$5, doubled to \$10.) Subjects are randomly matched and their role selected by the computer, after their decisions. This is a variant of the task used in Burks, Carpenter, and Götte (2006).

Before each decision screen, subjects are also asked how they think other participants in the room will act in this experiment. The first question is “What percent of the participants do you think will send their \$5 as Person 1?” and pays \$1 if the subject is correct within plus or minus 5%. The second and third questions are “If Person 1 does not send/does send, what is the average that participants in this room will send back?” and pays \$1 for each question if the subject is within plus or minus \$0.25 of the actual average.

There are no results tabulated yet for this experiment, but its potential application is in predicting on-the-job cooperation. The TL driver’s job is relatively individualized, and there are few metrics for on-the-job cooperation. But there are two places in which the data store from Research Component One may permit the construction of an on-the-job measure of cooperation; the results are not yet in. One has to do with whether a driver takes the time to fill out a form on the in-truck satellite communication unit identifying mechanical faults, when he or she is dropping a trailer that has such a fault, because the driver’s next dispatch is with a different trailer. It is a rule that such a form must always be filled out, but the rule is widely violated, with little consequence, and this has public-goods-game type spillover effects on other drivers. The second is the degree of responsiveness of drivers to particular messages sent over the satellite communications link to his or her truck from a subset of supervisors who plan dispatches, but to whom each driver is essentially anonymous. It is a rule that such messages should always be replied to, but this rule is also widely violated, also with little consequence, and the willingness to respond is potentially an index of cooperation with the company as a whole.

6.3.2. Multidimensional Personality Questionnaire (MPQ)

The Multidimensional Personality Questionnaire (Patrick, Curtin et al. 2002), otherwise known as the MPQ is the second measurement instrument. This is standard personality profile, and consists of 11 different scales that represent primary trait dimensions: wellbeing, social potency, achievement, social closeness, stress reaction, alienation, aggression, control, harm avoidance, traditionalism, and absorption. The short version used in the study has 154 multiple choice questions. Almost all of the 154 questions have the same four possible answers: “Always True”, “Mostly True”, “Mostly False”, and “Always False”. There is no separate payment for this survey, and no results have been tabulated yet.

A principal use of this scale will be to generate a set of control variables in regression modeling of on-the-job outcomes. We also expect to look at how the traits defined by the MPQ relate to answers to the other survey instruments, and to behavior in the experiments.

6.3.3. Risk/Loss Aversion

In the risk/loss aversion experiment there are 24 questions which are divided into four blocks of six questions (see table 3). There are two possible choices for each question, an amount of money received with certainty, and a 50/50 gamble that pays a higher dollar amount if “your color” is chosen and a lower dollar amount if the “other color” is chosen.

Each of the four blocks of six questions follows the same format. That is in each block, the amount for “your color” and the “other color” do not change and the amount of the “for sure” option increases by \$0.50 per question. The design is intended to identify where the subject crosses over from the “for sure” option to the “gamble” option, relative to where a risk-neutral person would do so. After all subjects have made their decisions, one poker chip is drawn by a participant which identifies which question (1-24) everyone will be paid on, and then a second participant draws a colored chip to determine whether the winning color for subjects who chose the gamble is green or blue.

Table 3: Risk/Loss Aversion

Question	For Sure	Your Color	Other Color
1	Win \$2.00	Win \$10.00	Win \$2.00
2	Win \$3.00	Win \$10.00	Win \$2.00
3	Win \$4.00	Win \$10.00	Win \$2.00
4	Win \$5.00	Win \$10.00	Win \$2.00
5	Win \$6.00	Win \$10.00	Win \$2.00
6	Win \$7.00	Win \$10.00	Win \$2.00
7	Win \$0.00	Win \$5.00	Lose \$1.00
8	Win \$0.50	Win \$5.00	Lose \$1.00
9	Win \$1.00	Win \$5.00	Lose \$1.00
10	Win \$1.50	Win \$5.00	Lose \$1.00
11	Win \$2.00	Win \$5.00	Lose \$1.00
12	Win \$2.50	Win \$5.00	Lose \$1.00
13	Lose \$2.50	Lose \$5.00	Win \$1.00
14	Lose \$2.00	Lose \$5.00	Win \$1.00
15	Lose \$1.50	Lose \$5.00	Win \$1.00
16	Lose \$1.00	Lose \$5.00	Win \$1.00
17	Lose \$0.50	Lose \$5.00	Win \$1.00
18	Lose \$0.00	Lose \$5.00	Win \$1.00
19	Win \$1.00	Win \$5.00	Win \$1.00
20	Win \$1.50	Win \$5.00	Win \$1.00
21	Win \$2.00	Win \$5.00	Win \$1.00
22	Win \$2.50	Win \$5.00	Win \$1.00
23	Win \$3.00	Win \$5.00	Win \$1.00
24	Win \$3.50	Win \$5.00	Win \$1.00

Risk aversion and/or loss aversion are potentially key explanatory factors in TL firm turnover, since nearly all of each driver's pay is based on piece rates (cents per mile for a fixed number of miles for each dispatch). So paychecks can vary quite substantially from one week to the next, and thus there is a certain level of financial risk that goes with the job. The question for the study is to what

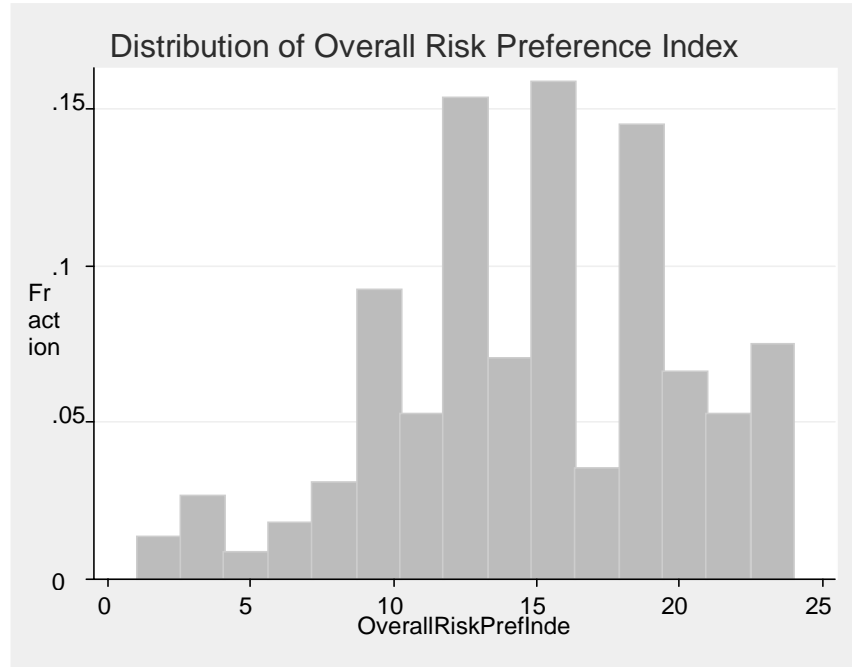


Figure 8: Distribution of the Total Count of Risky Choices

extent this form of experimental measurement of financial risk is predictive of the impact of pay fluctuations on driver turnover, as compared to things such as survey responses and credit score.

The choice blocs overlap the zero dollar point in order to allow inference about the likelihood of different slopes to the implied utility function in the positive and negative domains. But there is no set of choices only in the negative domain because the project is relying on the long-term responsiveness of participants, and pilot work showed that a choice block fully in the negative domain was perceived as behavior that was too unfriendly on the part of the researchers.

6.3.4. Demographics

This instrument is a compilation of 37 demographic questions. The types of questions vary, but a list of some of the areas that are covered are education level, languages, marital status, previous job experience, family lifestyle, racial or ethnic identity, country of birth, how many people are in the household, household income, and networking done while at training. These questions are mostly multiple choice, with a few numerical answer questions as well. There is no separate payment for completing this demographic questionnaire. Figures 9 and 10 show the distribution of ages and years of education for the driver-trainee participant pool so far (N = 227).

The minimum age to acquire the CDL is 21, and there are indeed some new trainees that are the minimum age. But a large fraction is much older: 50% are older than 36, and the mean age is 37. It is also interesting that a bit more than half the population so far has at least some post-secondary education, with about 15% having a four-year college degree or better.




Figure 9: Distribution of Participant Age

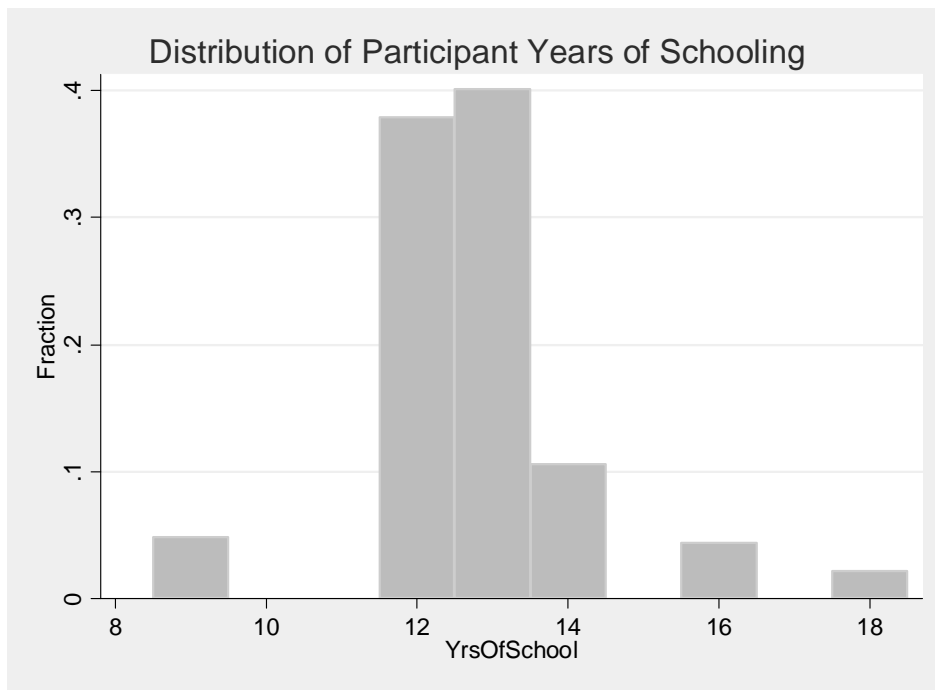


Figure 10: Distribution of Participant Years of Schooling

6.3.5. Big Red Button

This experiment is designed to measure short term impatience. It is the last data collection event in the first of the two two-hour sessions, and it is followed by 10 to 20 minutes of break time, during which participants can use the restroom, stretch their legs

and get something from the vending machines if they wish. It is directly preceded by the demographic questionnaire, to which it is linked. This means that, unlike most of the other instruments, subjects do not have to wait until everyone else completes the prior questionnaire, but instead each participant can start the red button task as soon as they have individually completed the demographic questionnaire. This setup is designed to minimize peer effects in the red button task, since each subject will start the task at a different time.

The red button task takes exactly ten minutes, and the task is to wait quietly until a timer on the task screen counts down from 600 seconds to zero. During each session subjects are requested in general to either work on the task at hand, or wait quietly, not talking to their neighbors, and not pulling out something to read. That message is reinforced at the beginning of the red button task. Participants are paid \$5 if they complete the task by waiting the full ten minutes. However, subjects may choose to end this experiment earlier if they wish. Each click of a button below the countdown clock on the computer screen, which is labeled “Reduce Waiting Time,” will cut their waiting time--but each click costs them \$1. The first click reduces their waiting time by five minutes; the second click reduces the waiting time by three minutes, and the third click reduces the waiting time by an additional two minutes. Once the timer says zero, no matter how long the subject has waited, the he or she is free to get up and leave the room for break time.

Patience is relevant to quit decisions, and being impatient can lead to poor job performance when a significant part of the job involves being able to wait on customer docks, and then hurry up when customers are ready. Being patient in traffic is also a relevant job skill. So far, 72% have waited the entire ten minutes, but 20% have clicked once, 2% twice, and 6% three times. We eventually expect to investigate statistically the extent to which red button choices are explained by prior earnings and prior waiting time during the data collection session. Once this is accomplished, a project goal is to examine whether or not this measure of impatience, as adjusted, will predict impatience on the job, as measured by quit decisions.

6.4. Session 2 Data Collection Events

6.4.1. Time Preferences

In this experiment there are 28 questions which are divided into four blocks of seven questions each (see Table 4). There are two possible choices for each question, a smaller amount of money paid sooner, and a larger amount of money paid later. Each of the four blocks of seven questions follows the same format. The amount for the higher payoff at a later date is always \$80 and the amount for the lower payoff at an earlier time begins at \$75 and decreases by \$5 intervals to \$45. The point at which a subject switches from the later payoff to the earlier one (if they do) provides an implicit point estimate of their discount rate over that time horizon.

The two matched pairs of time frames are designed to allow subjects that have different discount rates for choices with and without a front-end delay, i.e. with and without a “quasi-hyperbolic” discounting pattern. After all subjects have made their decisions, a participant draws one poker chip to select which of the 28 questions will be activated, and then two poker chips are chosen the same way which identify two of the subjects in the session who will be paid for their choices on that question. Payments are,

of course, made on the dates requested, either in person, or with an official bank check mailed in a University of Minnesota envelope. No results have yet been tabulated for this experiment.

Table 4

Question	Today	Tomorrow	Today	Thursday	Monday	1 Week from Monday	Monday	4 Weeks from Monday
1	\$75	\$80						
2	\$70	\$80						
3	\$65	\$80						
4	\$60	\$80						
5	\$55	\$80						
6	\$50	\$80						
7	\$45	\$80						
8			\$75	\$80				
9			\$70	\$80				
10			\$65	\$80				
11			\$60	\$80				
12			\$55	\$80				
13			\$50	\$80				
14			\$45	\$80				
15					\$75	\$80		
16					\$70	\$80		
17					\$65	\$80		
18					\$60	\$80		
19					\$55	\$80		
20					\$50	\$80		
21					\$45	\$80		
22							\$75	\$80
23							\$70	\$80
24							\$65	\$80
25							\$60	\$80
26							\$55	\$80
27							\$50	\$80
28							\$45	\$80

We expect to examine the relationship between this measure and such other measures as non-verbal IQ, risk/loss aversion, and impatience. And we expect to investigate the extent to which it adds predictive power to our statistical models of quits.

6.4.2. Non-Verbal IQ

The IQ instrument used is a computerized adaptation of the Standard Progressive Matrices by J.C. Raven (Raven, Raven et al. 2000); the authors created the adaptation under license from The Psychological Corporation.⁴² It consists of five sections, each containing 12 questions. Our version has only limited comparability to the standard results for the instrument, because we are actively administering only the last four

⁴² Pilot work was done with two other instruments, but one had irresolvable licensing issues, and the other proved unsuitable for our setting. The first several administrations of the Raven's instrument were done with paper and pencil, while we arranged the license for the computerized adaptation.

sections, due to time constraints. The original is primarily benchmarked in an un-timed format, and we cut our subjects off after 31 minutes, having given a warning about 3 minutes earlier (there is one published benchmark from a French subject pool using a 30-minute time limit).

Each question is presented as a graphic image. On top is a large rectangular box containing some kind of a pattern with a piece missing out of the lower right hand corner. On the bottom are six (or eight) possible pieces that could be used to complete the image on top. Each section starts with easy images, and gets progressively more difficult.




Figure 11: Distribution of Adapted Non-Verbal IQ

After both verbal and written instructions and two practice questions, subjects fill out a “confidence question” that asks them how they think they will do as compared to other subjects in the room, e.g. top 20%, bottom 20%, etc. When the Raven’s task has been completed, the same confidence question is asked again. Subjects are paid an additional \$2 for placing themselves in the correct quintile. In addition, two subjects are randomly chosen to be paid \$1 per correct answer, for total possible earnings of \$48 each for their question answers. Starting recently we also began asking whether subjects want to find out their own score and the group average, when they receive their payout.⁴³

The median raw score for the 175 of our subjects who have so far taken the Raven’s adaptation is 48, with a mean of 46.6.⁴⁴ The mode is about 51, and the distribution looks reasonable compared to groups of workers in other countries against which the regular version of the instrument has been benchmarked. The primary use of this measure in the project analysis will be as a control variable in predicting on-the-job outcomes. But we also expect to look at the relationships between this indicator if non-verbal IQ and the other participant characteristics that we measure. We will also examine how IQ, confidence about one’s performance both before and after the task, and the desire for full information about one’s performance, are related.

⁴³ The two selected to be paid for correct answers will learn their score from their payoff, so they will just get the average as new information, but these two subjects are selected after this question is asked.

⁴⁴ This score includes an projected number of correct answers for the initial section, which we are not administering.

6.4.3. Numeracy (Quantitative Literacy)

This instrument is part of the test of adult quantitative literacy from the Educational Testing Service. The full instrument consists of two sections, of which only the first section is being used here, due to time constraints. The section is made up of 12 questions and subjects are given exactly 20 minutes to complete the test. The test requires subjects to be able to add, subtract, compare numbers, fill out a form, and to be able to read and understand a short problem, among other things.

As with the non-verbal IQ, after instructions and a brief practice question, subjects fill out a “confidence question” that asks them how they think they will do as compared to other subjects in the room, by quintiles. When the numeracy task has been completed, the same confidence question is asked again. Subjects are

paid an additional \$2 for placing themselves in the correct quintile. Two subjects are randomly chosen to be paid \$2 per correct answer, for total possible earnings of \$24 each for their question answers. In addition, starting recently we began asking whether subjects want to find out their own score and the group average, when they receive their payout.⁴⁵

Both the mean and the median raw score for those of our subjects who have so far done the numeracy instrument are 8. The distribution is not directly comparable to the distribution for the full instrument (i.e. when both sections are administered), but we appear to be getting a good dispersion in performance (see Figure 12). The primary use of this in the project analysis will be as a control variable in predicting on-the-job outcomes. The job of TL driver requires continual application of numeracy skills, e.g. in map reading and route planning, or in calculating hours remaining and hours coming available at specific future times, under the hours of service regulations for commercial motor vehicle operators. But we also expect to look at the relationships between this indicator of numeracy and the other participant characteristics that we measure. We will

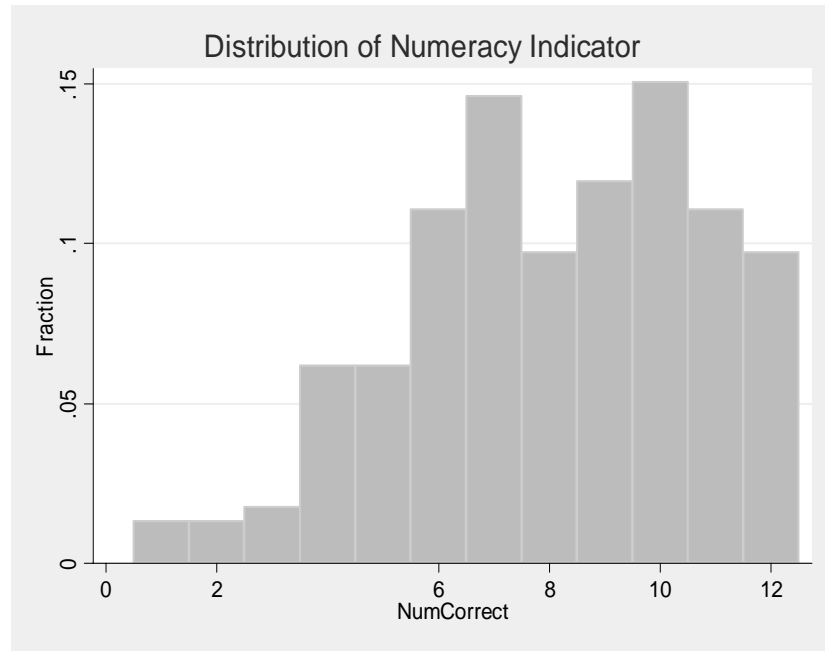


Figure 12: Distribution of Numeracy Indicator Scores

⁴⁵ As with the non-verbal IQ indicator, the two selected to be paid for correct answers will learn their score from their payoff, so they will just get the average as new information, but these two subjects are selected after this question is asked.

also examine how numeracy, confidence about one's performance both before and after the task, and the desire for full information about one's performance are related.

6.4.4. Ambiguity Aversion

Our ambiguity aversion experiment is identical to the risk/loss aversion experiment described above in section 6.3.3, except for one detail: it is no longer known to be a 50/50 gamble if the subject chooses the uncertain option. Two blue poker chips and two green poker chips are placed in the bowl from which the winning color is drawn, and then out of sight of the subjects, six more chips are added which can be all green, all blue, or any mixture thereof. As a result, subjects only know that there is at least a 20% chance that green will be drawn and at least a 20% chance that blue will be drawn. All other features of the experiment are unchanged. No results are tabulated yet for this experiment, but we will be looking to see whether subjects choose more or fewer risky options, and whether their shift, if any, is related to other experimental or survey measures, and to on-the-job success.

6.4.5. Hit Fifteen Points

This is a backward induction task, in the form of a small game between subject and computer. The computer and the subject take turns adding points to the "points basket" and during each turn the subject or the computer must add either one, two, or three points to the points basket. The goal is to be the player to add exactly the 15th point. The number of points in the points basket at the beginning of the round varies, and the computer and participant take turns going first. The first round is set so as to give the subjects an example of how the first stage of backward induction works, although this is done by ensuring that the computer wins on this round. Before each decision is made by the subject as to what number of points they want to add to the points basket, the subject is asked who they believe will win that round, the subject or the computer. The subjects are paid \$1 for each round that they win.

Together with numeracy, we anticipate that this measure will be predictive of on-the-job success, and especially of high versus low productivity. TL truck drivers have to do numerical backward induction every day, to calculate back from routing and delivery goals that are from a few hours to a few days ahead, in order to decide on the optimal course of action in the present. One regularly-used application is in figuring out how many hours they have to get to their destination and how many hours they are available to drive under the hours of service regulations.

An error in the programming of this experiment was discovered only recently, after the initial pilots, so little data has accumulated with the corrected version.

6.4.6. Risk, Impatience, and Cooperation Survey

The last instrument used during the Saturday initial data collection events with the panel study participants is a list of attitude questions about risk preferences, patience and impatience, and cooperation. It was assembled from a selection of papers in the literature that propose various survey-response measures for these characteristics of subjects. We intend these to also be tried out as control variables, in order to see if they are useful, and if so, whether they are substitutes or complements to the experimental measures, in predicting behavior on other measures, and on-the-job success.

6.5. Other Measures of Interest

The applicant information collected by the cooperating firm on trainees provides a few other data items of interest. There is an indicator of whether or not the participant is a smoker (whether they requested a hotel room in which smoking is permitted). This may turn out to be related to risk attitudes. In addition, since the trainees apply for credit for the training contract, the firm is able to supply us with a credit score. The one we will have available is the FICO Classic 98 score for credit risk.⁴⁶ We want to ask what relationship there is among our measures of risk/loss aversion, ambiguity aversion, time preference, and impatience, on the one hand, and the credit score, on the other. We also are interested in determining whether the credit score useful in predicting on-the-job success, and if so, whether it is a complement or substitute for our other measures.

7. Conclusion

The Truckers and Turnover Project is a statistical case study of a single firm and its employees in the TL segment of the U.S. trucking industry, which involves both proprietary personnel and operational data, and new data collected by the researchers. The project has two main components. The first component is a conventional in-depth statistical case study of the cooperating firm and its employees. This is a high-turnover service industry setting, and the focus is on the use of survival analysis to model the flow of new employees into and out of employment, and on the correct estimation of the tenure-productivity curve for new hires, accounting for the selection effects of the high turnover.

Descriptive pilot work on turnover shows that less than 40% of new hires remain after a year, and less than 20% after two years, even though most new hires sign a credit contract to repay a substantial amount of the cost of their training if they don't complete a year of service after training. We also find that the survival rates and time paths of exits vary significantly according to the type of prior background that new drivers bring to the firm. Further work will include regression modeling of turnover.

Most drivers in this industry segment, and at this firm, are paid on piece rates, and specifically, by the mile. Pilot work on productivity shows that full miles-per-week productivity is reached by new inexperienced drivers at sometime between nine and twelve months of tenure. There is some uncertainty because 6% of the driver-pay-week observations in the first generation of data have zero miles for reasons that cannot be identified from the payroll-based data set. Further data will be added in later generations, to address this issue.

Under piece rates the standard economist's prediction is that lower-productivity workers should exit at a higher rate than higher-productivity ones, so that the tenure-productivity curve based on raw data will be likely to be contaminated by selection effects. Irrespective of the data issue, it appears there is some evidence from a panel data regression model that this is correct after drivers reach full productivity. But there is also suggestive evidence that higher-productivity workers are leaving faster up to that point, which would contradict the initial expectation. A major goal of the project is to develop clearer evidence for or against this pilot finding.

⁴⁶ FICO 98 is a product of the Fair Isaac Corporation, the originator of the credit score.

The second component of the project is a panel study of approximately 1,000 newly hired driver trainees. The project's most distinctive innovation is a data collection process for the panel study which combines traditional survey instruments with behavioral economics experiments. The survey data include information on demographics, risk and loss aversion, time preference, planning, non-verbal IQ, and the MPQ personality profile. The data collected by behavioral economics experiments include risk and loss aversion, time preferences (discount rates), backward induction, patience, and the preference for cooperation in a social dilemma setting.

Panel study subjects will be followed over two years of their work lives, and the initial data collected will be supplemented by follow-up surveys, and by matching the panel study subjects with the operational and human resource data on their first two years of work, as provided by the firm. Among the major design goals are to discover the extent to which the survey and experimental measures are correlated, and whether and how much predictive power, with respect to key on-the-job outcome variables, is added by the behavioral measures. After extensive development and pilot work, the panel study is currently actively collecting data, and will conclude the intake of new subjects sometime in the summer of 2006, with follow-ups due to be complete in late 2008.

8. Appendix A. The Truckers and Turnover Project Team

Principal Investigator and Organizer:

- 1) Stephen Burks, Associate Professor of Economics and Management, University of Minnesota, Morris, IZA, and Associate of the Sloan Trucking Industry Program

Co-Investigators (each colleague is collaborating on at least one component of the project, some on multiple components):

- 2) Jeffrey Carpenter, Associate Professor of Economics, Middlebury College, IZA, and Core Member, Norms and Preferences Network
- 3) Aldo Rustichini, Professor of Economics, University of Minnesota, Twin Cities
- 4) Lorenz Götte, Assistant Professor of Economics, Institute for Empirical Economic Research, University of Zürich (Switzerland), and IZA
- 5) Kristen Monaco, Professor of Economics, California State University at Long Beach, and Associate of the Sloan Trucking Industry Program
- 6) Andrew Clark, Research Professor, Centre National de la Recherche Scientifique, PSE (France), and IZA
- 7) Francine Lafontaine, Professor of Business Economics, Ross School of Business, University of Michigan, and Associate of the Sloan Trucking Industry Program, and NBER
- 8) Jon Anderson, Associate Professor of Statistics, University of Minnesota, Morris

On-site Team Members:

- 9) Kay Saager, Business Research Manager, Cooperating Firm
- 10) Adam Durand, Research Intern (fall, 2005), University of Minnesota, Morris
- 11) William Leuthner, Research Intern (spring, 2006), University of Minnesota, Morris
- 12) Erin Christenson, Research Intern (summer, 2006), University of Minnesota, Morris

9. References

- ATA Economic & Statistics Group (2005). Employee Turnover and Workforce Changes: Q4: 2004. Trucking Activity Report. **13**: 7.
- Beadle, A. D. (2004). "Building Driver Shortage." Traffic World **268**(42): 26-27.
- Belman, D. L., K. Monaco, et al. (2005). Sailors of the Concrete Sea: A Portrait of Truck Drivers' Work and Lives. Lansing, MI, Michigan State University Press.
- Belman, D. L. and K. A. Monaco (2001). "The Effects of Deregulation, De-Unionization, Technology, and Human Capital on the Work and Work Lives of Truck Drivers." Industrial and Labor Relations Review **Volume 54**(No. 2A): 502-524.
- Belzer, M. H. (1995). "Collective Bargaining After Deregulation: Do the Teamsters Still Count?" Industrial and Labor Relations Review **48**(4): 636-655.
- Belzer, M. H., D. Rodriguez, et al. (2002). Paying for Safety: An Economic Analysis of the Effect of Compensation on Truck Driver Safety, Federal Motor Carrier Safety Administration.
- Bureau of Labor Statistics (2002). Historical hours and earnings--Average hours and earnings of production or non-supervisory workers on private non-farm payrolls by major industry sector, 1964 to date, U.S. Department of Labor.
- Burks, S. V. (1999). The Origins of Parallel Segmented Labor and Product Markets: A Reciprocity-Based Agency Model with an Application to Motor Freight. Amherst, MA, Graduate School of the University of Massachusetts, Amherst: 1-181.
- Burks, S. V., J. Carpenter, et al. (2006). "Performance Pay and the Erosion of Worker Cooperation: Field Experimental Evidence." Institute for the Study of Labor Discussion Papers(#2013).
- Burks, S. V., K. Monaco, et al. (2004). "The Balance between Private and For-Hire Carriage and Trends in the Use of Large Trucks (1977 to 1997)." Journal of the Transportation Research Forum **43**(2): 159-71.
- Cain, G. G. (1976). "The Challenge of Segmented Labor Market Theories to Orthodox Theory: A Survey." Journal of Economic Literature **11**(4): 1215-1257.
- Casey, J. F. (1987). An Assessment of the Truck Driver Shortage, American Trucking Associations: 12.
- Cleves, M. A., W. A. Gould, et al. (2004). An Introduction to Survival Analysis Using Stata, Revised Edition. College Station, TX, Stata Press.
- Corsi, T. and J. R. Stowers (1991). "Effects of a Deregulated Environment on Motor Carriers: A Systematic Multi-Segment Analysis." Transportation Journal **30**(3): 4-28.
- Cox, K. (2004). TL Driver Turnover Rate Reaches New Record High. Transport Topics: p. 5 & 35.
- Cullen, D. (1996). "You get what you pay for: J.B. Hunt Transport Services Inc.'s driver compensation package." Fleet Owner **91**(11): 82-3.
- Dickens, W. T. and K. Lang (1993). Labor Market Segmentation Theory: Reconsidering the Evidence. Labor Economics: Problems in Analyzing Labor Markets. W. Darity, Jr. Norwell, MA, Kluwer Academic: 141-180.
- Fischbacher, U. (1999). "z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual." Institute for Empirical Research in Economics Working Papers(#21).

- Gallup Organization, S. (1997). Empty Seats and Musical Chairs: Critical Success Factors in Truck Driver Retention, ATA Foundation.
- Global Insight Inc. (2005). The U.S. Truck Driver Shortage: Analysis and Forecasts.
- Griffin, G. and L. Kalnbach (2002). Predicting and Classifying Voluntary Turnover Decisions for Truckload Drivers. Fargo, ND, Upper Great Plains Transportation Institute, North Dakota State University: 40.
- Griffin, G. C., J. M. Rodriguez, et al. (1992). Job Satisfaction in U.S. Commercial Drivers. Fargo, North Dakota, The Upper Great Plains Transportation Institute, North Dakota State University: 1-20.
- Isidore, C. (1997). "Hunt gambles on pay raise to keep drivers." Journal of Commerce **411**(28917): 1A.
- Kiefer, N. M. (1988). "Economic Duration Data and Hazard Functions." Journal of Economic Literature **26**: 646-679.
- Patrick, C. J., J. J. Curtin, et al. (2002). "Development and Validation of a Brief Form of the Multidimensional Personality Questionnaire." Psychological Assessment **14**(2): 150-163.
- Raven, J., J. C. Raven, et al. (2000). Section 3: The Standard Progress Matrices. Manual for Raven's Progressive Matrices and Vocabulary Scales. San Antonio, TX, The Psychological Corporation.
- Reiskin, J. (2006). "Shipper Changes Cut Distances That Truckload Freight Travels." Transport Topics(March 13): 1 and 49.
- Rodriguez, J., M. Kosir, et al. (2000). The Costs of Truckload Driver Turnover. Fargo, Upper Great Plains Transportation Institute: 1-18.
- Stephenson, F. J. and R. J. Fox (1996). "Driver retention solutions: strategies for for-hire truckload (TL) employee drivers." Transportation Journal **35**(4).
- U.S. Census Bureau (2004). Truck Transportation: 2002. Economics and Statistics Administration, U.S. Department of Commerce: 1-40.
- Waxler, C. (1997). "J.B. Hunt Transport shares may drop further." Forbes **159**(8): 382-3.