

DISCUSSION PAPER SERIES

IZA DP No. 10835

**Training Contracts, Employee Turnover,
and the Returns from Firm-Sponsored
General Training**

Mitchell Hoffman
Stephen V. Burks

JUNE 2017

DISCUSSION PAPER SERIES

IZA DP No. 10835

Training Contracts, Employee Turnover, and the Returns from Firm-Sponsored General Training

Mitchell Hoffman

University of Toronto and NBER

Stephen V. Burks

University of Minnesota, Morris and IZA

JUNE 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Training Contracts, Employee Turnover, and the Returns from Firm-Sponsored General Training*

Firms may be reluctant to provide general training if workers can quit and use their gained skills elsewhere. “Training contracts” that impose a penalty for premature quitting can help alleviate this inefficiency. Using plausibly exogenous contractual variation from a leading trucking firm, we show that two training contracts significantly reduced post-training quitting, particularly when workers are approaching the end of their contracts. Simulating a structural model, we show that observed worker quit behavior exhibits aspects of optimization (for one of the two contracts), and that the contracts increased firm profits from training and reduced worker welfare relative to no contract.

JEL Classification: J24, M53, J41

Keywords: training contract, firm-sponsored general training, organizations, trucking, truck driver, truckload

Corresponding author:

Stephen V. Burks
University of Minnesota, Morris
Division of Social Sciences
600 East 4th Street
Morris, MN 56267-2134
USA

E-mail: svburks@morris.umn.edu

* A previous paper appeared as “Training Contracts, Worker Overconfidence, and the Return from Firm-Sponsored General Training.” That previous paper has been divided in two, with the present paper focusing on the impact of training contracts. The other paper (Hoffman and Burks, 2017) focuses on worker overconfidence and uses a different dataset (though also based on workers from Firm A). Some small portions of text are similar across the two papers. We are very grateful to David Card, Stefano DellaVigna, John Morgan, and Steve Tadelis for their encouragement and advice. We also thank Ben Handel, Ben Hermalin, Ken Judd, Pat Kline, Botond Koszegi, Matthew Rabin, and Kenneth Train for especially detailed comments, and we also thank many seminar/conference participants. We thank the many trucking managers and drivers who shared their insights with us. We thank managers at Firms A for sharing their data and for facilitating on-site data collection. We thank Mark Gergen, Anthony Kraus, and Gregory Lemmer for guidance in understanding relevant legal issues. We thank Graham Beattie, Christina Chew, Dan Ershov, Sandrena Frischer, Will Kuffel, Amol Lingnurkar, Kristjan Sigurdson, and Irina Titova for outstanding research assistance. Hoffman acknowledges financial support from the National Science Foundation IGERT Fellowship, the Kauffman Foundation, and the Social Science and Humanities Research Council of Canada. Burks and the Truckers & Turnover Project acknowledge financial support from Firm A, the MacArthur Foundation, the Sloan Foundation, the Trucking Industry Program at Georgia Tech, and University of Minnesota, Morris.

NON-TECHNICAL SUMMARY

Firms may hesitate to provide training which is general enough that it can be used in other jobs, if employees are likely to leave shortly after being trained. A “training contract,” which penalizes the employee for quitting quickly after being trained, may help solve this problem. Using data from a large US trucking company, we show that two different training contracts did reduce quitting after training, especially near the end of the contracts. For one contract the observed patterns of turnover closely match the predictions of a theoretical model of turnover. Our analysis shows that the contracts increased the profitability of training for the firm and decreased worker welfare relative to no contract being in place, though not necessarily relative to having no training opportunity at all.

1 Introduction

Firm-sponsored training is common in many countries and is a central means by which workers accumulate human capital (e.g., Acemoglu and Pischke, 1998, 1999a; Autor, 2001). However, it has been recognized since Pigou (1912) that general training is subject to a “hold-up” problem: firms may be reluctant to train if workers are likely to quit after training. Becker’s (1964) canonical solution is for workers to pay for training themselves, but this may fail for a number of reasons, e.g., credit constraints. Understanding what makes training profitable for firms is important for policy, given concerns that US workers receive less training than workers in countries with lower turnover, and for economic theory.¹

To discourage workers from quitting after training, firms often use training contracts: the firm pays for training and workers are fined for premature quitting. Training contracts of this form are used for many jobs (e.g., truckers, policemen, nurses, pilots, and federal employees, to name a few), but have received limited attention from economists.² How do training contracts affect quitting, as well as firm profits from training? We demonstrate that training contracts significantly reduce post-training quitting and increase the profitability of general training.

We show that training contracts reduce quitting using plausibly exogenous contractual variation from the trucking industry. Although it is one specific industry, trucking is ideal for our study because it is large (see Section 3.1 for details), training contracts are widely used, and productivity (miles driven per week) is easy to measure. At a leading US-based trucking firm, which we call Firm A, training was initially provided with no contractual obligation. In an effort to increase retention, the firm introduced a training contract fining workers for quitting in their first 12 months. Different states approved the contract for use at different speeds, leading the contract to be phased into different training schools at different times. Several years later, contracts were changed a second

¹See e.g., Acemoglu and Pischke (1998, 1999a); Autor (2001); Barron et al. (1999); Cappelli (2004) for evidence that firms often pay for general training, both nominally and in terms of incidence, and for general discussion of firm-sponsored general training. US firms appear to spend less on training than firms in other countries (Lynch, 1993; Brunello and Medio, 2001), but firm training expenditures are still substantial, estimated in 1994 at \$50 billion annually for private-sector formal training programs (Baron and Kreps, 1999).

²We discuss related economics work below. For our examples of workers with training contracts, as well as discussion of relevant legal issues, see the law articles by Kraus (1993, 2008). Other workers with training contracts include firefighters, pilots, mechanics, salesmen, paramedics, electricians, accountants, teachers, flight attendants, bank workers, repairmen, firm-sponsored MBAs, and social workers. Courts in the US and abroad have generally ruled that training contracts are legally permissible, arguing they serve the public good by promoting investment in training. See Appendix E.1 for more on legal issues related to training contracts.

time to an 18-month pro-rated contract, also in a staggered fashion.

Exploiting the staggered contractual roll-out in a “diff-in-diff,” we estimate that the two contracts both reduced quitting by about 15%, relative to no contract. The contract effects vary substantially by worker quarter of tenure, with the largest effects occurring closer to when the contracts expire. Separate cohort-based event study analysis shows that the contracts led to sharp changes in quitting in the quarter after the change, as well as no evidence of pre-trends. While these results may seem intuitive, they violate a basic prediction of the Coase Theorem where negotiation through side payments leads to efficient turnover irrespective of the quit penalty in place (Lazear, 1990). The results thereby suggest that some condition violating the Coase Theorem (such as private information and/or costs to contract renegotiation) is at play in this market.

Having estimated the causal impacts of the contracts, we next discuss mechanisms. In general, contracts may produce treatment effects by affecting behavior (“incentives”) and/or by affecting employee selection (“selection”) (Lazear, 2000). In our setting, the evidence is most consistent with incentives being the driver of the impacts we observe, as the workers brought in under the training contracts do not appear superior in terms of productivity or firing rates, and do not appear different in terms of background characteristics. After discussing mechanisms, we argue against several potential confounds / alternative explanations.

Last, we use a structural model of turnover developed in Hoffman and Burks (2017) to examine to what extent observed behavior across the contracts is consistent with a dynamic programming model, as well as to analyze how training contracts affect worker welfare and firm profits. After receiving training, workers make a decision every week whether to quit, gradually learning about their underlying productivity by observing their weekly miles (similar to Jovanovic (1979)). If individuals quit, they are fined according to a training contract, which specifies penalties at different tenure levels. The model is estimated entirely off a subset of about 700 workers facing the 12-month contract for whom we collected very rich data.

Using the model to simulate how training contracts affect quitting, we show that the model delivers similar quantitative estimates to our quasi-experimental estimates, suggesting that observed worker quit reductions were roughly comparable to those predicted by dynamic optimization. For the 12-month contract, we show that the time path of the contract impacts relative to no contract is also consistent with theory. Interestingly, for the 18-month contract, the time path of contract

impacts is quite different from the theoretical prediction, even though the model is roughly accurate in predicting the magnitude of the overall contract impact. Furthermore, we show that the 12-month and 18-month training contracts led to fairly similar increases in profits relative to not having a contract. In addition, both contracts moderately decreased worker welfare.

The central contribution of the paper is to show that training contracts significantly reduce quitting, estimating the effects using plausibly exogenous intra-firm contractual variation. While there is a large literature on why firms pay for general training, we provide the first evidence (to our knowledge) that training contracts play an important role in reducing turnover after training and making training profitable for firms, and we provide evidence that our results are causal.³ Acemoglu and Pischke (1999a) review the literature on firm-sponsored general training. There are many other reasons besides credit constraints that firms may pay for general training including informational asymmetries, search frictions, and labor market institutions that compress the wage distribution. The part of the training literature most related to our paper is that on tuition reimbursement. Employer-provided tuition reimbursement programs are quite common in the US, being offered in as many as 85% of medium and large firms (Cappelli, 2004). In a sample of part-time MBA students, Manchester (2010) found that 87% received tuition assistance, with 42% of those obligated to come back to the firm for 12 or more months after completing the MBA.⁴ We also contribute to the broader empirical contracts literature; Chiappori and Salanie (2003) argue that theory on contracts has outstripped empirics, reflecting that plausibly exogenous variation in contracts (such as in our paper) is rare.⁵

³Economists have studied turnover impacts from contracts other than training contracts, such as non-compete agreements (e.g., Marx et al., 2009; Balasubramanian et al., 2017), as well as tuition reimbursement programs without post-training retention requirements (e.g., Manchester, 2012). Some of such papers are discussed below in Section 4.1, when we discuss the magnitudes of our quit impacts. More general than the training contracts of the type we study, we provide some of the most definitive causal evidence to date that formal contracts can serve a critical role in reducing post-training turnover.

⁴Manchester (2010) also examines the relation between having a post-MBA service requirement and a student's surveyed intended probability of quitting their firm in the next 12 months. She finds that a longer post-MBA service requirement is associated with lower intended turnover (relative to no service requirement), but, interestingly, that a shorter post-MBA service requirement is associated with higher intended turnover (relative to no service requirement). For a recent theoretical analysis of employee bonding and turnover, which includes analysis of training contracts, see Peterson (2010). Dustmann and Schoenberg (2012) analyze a different commitment issue in training: the ability of firms to commit to providing a certain quality of training.

⁵Instead of analyzing variation in contracts, Naidu and Yuchtman (2013) study how exogenous labor demand shocks affect prosecutions for servants breaching employment contracts in 19th century Britain. Naidu et al. (2016) study the consequences of a Middle Eastern reform that made it easier for migrants to switch employers. By studying formal training contracts, our paper differs from that studying more informal contracts, which are surveyed by MacLeod (2007).

A secondary contribution of our paper is to analyze worker behavior under training contracts using a structural model, using the model to estimate impacts on worker welfare and profits. Because of the rich, plausibly exogenous contractual variation, we are able to compare quasi-experimental and structural estimates, something which is rare in the literature (for exceptions, see Card and Hyslop (2005), Todd and Wolpin (2006), and Galiani et al. (2015)), but something that scholars increasingly advocate when possible (Keane and Wolpin, 2007; Todd and Wolpin, 2006). This is valuable in two respects. First, it helps shed light on to what extent observed quasi-experimental impacts are consistent with dynamic optimization. Second, it helps validate the structural model.

The paper proceeds as follows. Section 2 provides theoretical discussion on how training contracts may affect worker behavior. Section 3 provides background information on the trucking industry and describes the data. Section 4 provides reduced-form estimates of the contracts, as well as discusses mechanisms and alternative explanations. Section 5 performs structural simulations. Section 6 concludes.

2 Theoretical Discussion

In Becker’s (1964) model of training, workers finance their own training through reduced wages during the training period. However, if workers have limited credit and training is brief (e.g., a few weeks, as in trucking), workers will not be able to finance training through a reduced training period wage without making their wage highly negative. As we model in Appendix B, this situation can be remedied with a training contract, helping the firm recover training costs after quits and potentially also reducing quitting. The Appendix B model is a stylized one-period version of the dynamic structural model considered in Section 5, and it allows us to derive a proposition analytically.

It is not obvious that a training contract will affect quitting. Suppose that workers and firms have no private information and that bargaining is costless. Then, by the Coase Theorem, turnover will be efficient, occurring when the sum of the worker’s and firm’s outside options exceeds the value of the match. Thus, turnover will be unaffected by a training contract, as a training contract is merely a “property right” held by firms over the quit decision.

To see why, consider the case where it is socially optimal for the worker to quit, but disadvantageous for the firm. Without a training contract, the worker will quit; the firm will try to “bribe” the worker to stay, but the maximum bribe the firm is willing to offer will still not be high enough to retain the worker. With the same situation and a training contract in place, the worker and firm will bargain such that (after negotiation) the worker will still quit. Whether the worker must bribe the firm to let him quit may be affected by the training contract, but the quitting outcome will not be.

In our context, however, it seems unlikely that the conditions of the Coase Theorem will hold. Workers likely have private information (about their taste for the job or their outside option) and renegotiating contracts with thousands of workers may be costly for a large firm like Firm A. Assuming that workers have private information and that there is no renegotiation, allowing for training contracts increases firm profits from training and reduces turnover (Proposition 1 or P1).⁶

In a related application of the Coase Theorem, Lazear (1990) analyzes job security provisions in Europe, where firms are “fined” (e.g., they must pay severance pay) for firing workers. He shows theoretically how the Coase Theorem may fail to hold and shows empirically that job security provisions do indeed affect firm firing.

3 Background on our Setting and Data Description

3.1 Background

Truckdriving in the US. In 2010, there were roughly 1.8 million US workers operating heavy trucks such as those used by Firm A (BLS, 2010). Firm A is in the long-distance truckload segment of the for-hire trucking industry, which is the biggest employment setting in this occupation. There is an important distinction between short-haul and long-haul trucking. Whereas long-haul truckload drivers are typically paid by the mile (a piece rate) (Belzer, 2000) and drive long distances away from home, short-haul truckload drivers are not usually paid by the mile and typically spend fewer nights away from home.⁷

⁶We note that even if the conditions of the Coase Theorem held, changing training contracts could influence quit rates by affecting worker selection—Section 4.2, however, finds fairly limited evidence of selection effects.

⁷We highlight several other institutional details here. In truckload, most drivers do not own their own trucks and the unionization rate is low; in addition, employee turnover is very high, frequently exceeding 100% annually (Burks et al., 2008). Baker and Hubbard (2004) report that in 1992, around 10% of trucks were driven by drivers

For heavy truckdrivers, the main training is that needed to acquire a commercial driver's license (CDL). Most new truckers complete a formal CDL training course, and it is required by law in some states (BLS, 2010). CDL training is provided at various venues, including truck driving schools run by trucking firms, private truck driving schools, and some community colleges. In phone surveys we did with 30 of the largest US-based truckload firms, about half the firms report providing CDL training at some point from 2001-2010, with firms providing CDL training typically using a training contract (see Appendix E.3 for more on the survey). CDL training at Firm A lasted about 2-3 weeks, though the length of training varied somewhat by training school and over time. Firm A CDL training included classroom lectures, simulator driving, and actual behind-the-wheel truck driving. The market price for CDL training varies, but is often several thousand dollars at private training schools.

Production. Truckload drivers transport full loads between various locations. Hours worked are constrained by the federal legal limit of about 60 hours per week. While our data do not contain driver hours, managers informed us that drivers often work up to the limit.

At Firm A, load assignment is done by a central dispatching system. Loads are assigned primarily based on proximity, as well as hours remaining up to the federal limit. Upon finishing a load, a driver may begin a new one.

In long-haul trucking, productivity is measured in miles driven per week. Productivity reflects both significant persistent differences across drivers and substantial idiosyncratic variation within drivers. According to managers at the firm, productivity differences across drivers reflect various factors, including speed, skill at avoiding traffic, route planning (miles are calculated not by distance traveled, but according to a pre-specified distance between two points), not getting lost, and coordinating with other people to unload the truck. Given that miles are calculated based on a pre-specified distance between two points, poor route planning or getting lost means that drivers are wasting time instead of accumulating paid miles. Arriving late might mean that a driver needs to wait around to have someone unload their truck, which can also be quite damaging to accumulating high miles per week. As for the sources of the substantial week-to-week idiosyncratic

who own their own truck (owner-operators), with the remainder driven by drivers driving company-owned trucks (company drivers). All the drivers we study are non-union company drivers. Hubbard (2003) provides an analysis of productivity in trucking.

variation, managers emphasized traffic, weather, variable unloading/loading time, and disadvantageous load assignments. Thus, driver miles reflect both driver performance and effort, as well as factors that drivers do not control and may be hard to predict. For more on measuring productivity, see Appendix E.2 and Hoffman and Burks (2017).

Contract Changes. To examine the effects of training contracts, we analyze two large contract changes at Firm A, a leading trucking firm. At the start of our data period, Firm A provided CDL training to thousands of new drivers per year at no cost at several geographically dispersed training schools. There was no contractual obligation. We omit certain details from our descriptions to preserve Firm A’s anonymity.

Around late 2000, management proposed implementing a training contract.⁸ The primary motivation was to increase retention, with a secondary motivation being to help recover costs. To implement the contracts, contracts had to be certified by the states where the training schools were located. This certification process took different amounts of time in different states. In three training schools, the contract was certified in 2001 and was put in use on a particular date in early spring 2002. At another training school, certification did not occur until late spring 2002 and the contract was not used until fall 2002, and in one school the contract was never used due to certification issues. Managers told us that cross-state differences in time for state certification seemed idiosyncratic and were unlikely to be related to the type of impact the contracts might have. The quit penalty varied slightly by training school and was between \$3,500 and \$4,000. The contract lasted 12 months and the quit penalty was constant throughout. The contract applied for both quits and fires.⁹

After several years of the 12-month contract, in an effort to further increase retention, management decided to switch to an 18-month contract. The initial penalty for quitting was increased to roughly \$5,000, but gradually decreased with worker tenure—the penalty was reduced by about \$65 per week of service.¹⁰ Again, the contract was phased in gradually, being introduced at one

⁸Our impression from talking with managers is that the idea of using a training contract was proposed by a newly promoted manager, and that the firm had not previously considered the possibility of using a training contract. Also, note that our description of the timing of contract changes is for the five training schools in our sample.

⁹According to managers, the contract also covered fires so as to prevent workers who wanted to quit from trying to get fired; however, according to these managers, the firm did not intentionally fire workers to collect training penalties.

¹⁰Of the roughly \$65 per week that was deducted from the worker’s initial quit penalty, the worker had around \$13 each week deducted from his pay check. After two years with the company, the driver would receive a bonus

school before being brought to other schools on a particular date about one year later.¹¹ Adoption of the 12-month and 18-month contracts were made without additional changes to driver pay. Both the changes to the 12-month and 18-month contracts applied only to new drivers; drivers who had already signed on with no contract or the 12-month contract did not have their contracts altered.

Appendix A.1 gives further background on the contract changes. Appendix Table A1 lists the (fuzzed) dates of contract changes, summarizing the information that is presented here in Section 3.1.

Enforcement. While it may seem hard to enforce training contracts with truckers, Firm A made strenuous enforcement efforts. Drivers signed a written contract specifying penalties for early exit. No bond was posted. Upon early exit, drivers were contacted by Firm A to pay the amount due. If drivers did not pay, they were often referred to one of multiple collection agencies. For drivers who remained delinquent, credit agencies were notified. Although comprehensive driver-level collection data are not available, the available data and conversations with managers suggest that approximately 30% of quit penalties were collected (details in Appendix A.2).

3.2 Data

The data from Firm A are ideal for analyzing the impact of training contracts due to the large sample size and high frequency of observation. We focus on new inexperienced drivers who are trained by Firm A. The data contain weekly miles and earnings, as well as basic demographics and other information, for thousands of new drivers for 2002-2009. Drivers are primarily paid by the mile, with small payments for other tasks (e.g., helping unload a truck). The per mile piece rate increases with driver tenure. We restrict our sample to drivers from 5 training schools, excluding several schools where the training provided differed and/or the precise contract change dates were not available. We refer to this sample as the “full sample.” Appendix A.3 provides further details on sample construction.

Several other papers by one or both of the authors have analyzed Firm A data on a subset of roughly 1,000 new drivers trained at one of the firm’s training schools in late 2005-2006.¹² However, payment roughly equal to the total amount deducted from his pay check in the first 18 months.

¹¹Unlike the first contract change, the order of locations for the second contract change appears to have been more actively chosen by the firm. Appendix A.1 discusses why we do not believe it to be a source of bias.

¹²Appendix A.4 describes several unrelated papers using the data subset (e.g., comparing social preferences of

the full 8-year Firm A dataset on all new workers, with hundreds of thousands of worker-weeks, is much larger than the data subset; our paper is the first to analyze the full dataset and the dataset is critical for identifying the impact of training contracts (given the cross-training school contractual variation).¹³

Table 1 presents sample means. The sample size is omitted to preserve firm confidentiality. The sample is primarily male and is majority white. The majority of workers have the 12-month contract, but there are still sizable shares with no contract and the 18-month contract.

In the full sample of drivers, we do not observe driver credit scores. However, in the subset of Firm A drivers studied by Hoffman and Burks (2017), there is information on credit scores: the average credit score is 586 and the median credit score is 564. These are quite low relative to the US median credit score of 723 (median at time of data collection) and indicate that a substantial share of drivers are “subprime.”¹⁴ Such low credit scores are consistent with many drivers facing credit constraints.

Figure 1 compares quit hazards under the three contractual regimes. Quit rates are high (often over 1% per week, though varying substantially based on week of tenure). For drivers with the 12-month contract, there is a spike in quitting at the 52-week mark. There are also smaller bumps at the 52-week mark under the no contract and 18-month regimes. Firm A managers suggested that these bumps in quitting at 52 weeks under the no contract and 18-month contract regimes may result from workers postponing quitting until then to be able to say that they worked for a full year at their last employer when applying for other jobs.

Where do drivers go when they exit? As is common in most personnel datasets, we do not observe where drivers go when they exit the firm. For the subset of drivers studied by Hoffman and Burks (2017), we surveyed drivers by mail to collect this information, and we found that most exiting drivers are not going to jobs in long-haul trucking. Of the drivers responding to our survey,

truckers, students, and non-trucker adults). For details on the Firm A data collection, see Burks et al. (2008).

¹³A later-written paper, Burks et al. (2015), uses data from 9 firms, including the full dataset from Firm A, to study the unrelated question of hiring through employee referrals (for details, see Appendix A.4).

¹⁴The credit score is the FICO-98 and has a range of 300-850. Among the drivers, 53% have a credit score less than 600, compared to a rate of only 15% in the US general population. The credit score definition of a “subprime” borrower varies by lender, though the cutoff is frequently 620 or 640. Drivers are particularly over-represented among those with credit scores below 550 (a very low score), with 43% of drivers being below 550 compared to 7% of the general population. We use the “Report to the Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit” issued by the Federal Reserve Board of Governors (Aug. 2007) for information on credit score statistics. For drivers in the subset studied by Hoffman and Burks (2017), Firm A purchased credit scores.

48% reported moving to a non-trucking job or unemployment, while 25% reported going to a local driving job. Further, 15% of drivers went to a regional trucking job and only 12% report going to a job in long-haul trucking.¹⁵

Why is reducing turnover particularly important for Firm A? As is the case for many firms that train, it is quite costly for Firm A to train a new driver who quits shortly after training, and this is true even after training contracts are introduced (given the imperfect enforcement). Furthermore, there is substantial learning by doing that occurs for new drivers, as driver productivity rises considerably during the first year on the job before flattening off (see Appendix Figure D1). When a worker quits prematurely, the firm has less time to reap the benefits of a worker producing at “peak productivity” and instead must invest more in recruiting and training new workers. Because of this, in conversations with the researchers, Firm A managers expressed a strong general desire to reduce driver turnover.

4 Analysis

4.1 Do Training Contracts Affect Quitting?

Diff-in-Diff for overall effects. Using Cox proportional hazard models, Table 2 shows that training contracts significantly reduced quitting. The regressors of interest are dummies for the 12-month and 18-month contracts, which vary at the school-week of hire level. We also include current quarter-year fixed effects (to control for differences in quit rates over time), month-year of hire fixed effects (to control for differences in quit rate by month of hire), training school dummies, the annual state unemployment rate, a driver’s average productivity to date, and other controls. The training contract dummies are identified by changes in training contracts across workers hired at different dates at a given training school. Throughout the reduced form analysis, standard errors are clustered at the training school-week of hire level (the level of variation for the training contracts). Doing so allows for arbitrary correlation of the error within training school classes (drivers attending the same training school at the same time).

¹⁵The survey had a response rate of only 25%; however, we are not particularly concerned about selection bias, as whether a driver responds to the survey is uncorrelated with most driver characteristics. For further information on the exit survey, see Section 2.2 and Appendix A5 of Hoffman and Burks (2017) for further details.

In column 1, the coefficients on the 12-month and 18-month contract variables are -0.141 and -0.148, respectively, meaning the contracts decreased quitting by about 14% and 15%. The estimates remain similar when controls are added for the driver’s current state unemployment rate, the driver’s average productivity to date (in hundreds of miles driven per week), and demographic controls.¹⁶ The contract coefficients are a bit larger in magnitude once school-specific time trends are included, increasing to between -0.20 and -0.24 . The coefficient on average productivity to date (in hundreds of miles per week) is about -0.09 , reflecting that more productive workers are less likely to quit. Thus, the impact of the training contracts on quitting is roughly the same as an increase of average productivity to date of about 150-200 miles per week, where 200 miles per week is about 10% of weekly average productivity.¹⁷

In assessing magnitudes, it is useful to compare our estimates to those in studies of other types of contracts. Although often used for higher skilled workers than truckers, one contract that is related to a training contract is a non-compete agreement, which prohibit employees from starting competitor firms or moving to a competitor for a given length of time. Marx et al. (2009) study the turnover consequences of allowing enforcement of non-compete agreements for Michigan inventors, finding that enforcing non-competes led to an 8% drop in turnover for inventors outside the auto industry.¹⁸ Another type of contract/benefit that may affect turnover is health insurance eligibility, which may cause “job lock.” Madrian (1994) finds that having spousal health insurance increases job turnover by 25%, implying that job lock reduces turnover by 25% (see Gruber and Madrian (2002) for further discussion of the job lock literature). Based on these estimates, the impact of the training contracts is about twice as large as those from enforcing non-compete agreements, and about 60% as large as those from an employee not having spousal health insurance.

¹⁶We control for average productivity to date (as opposed to some other combination of past productivity realizations) because it is a sufficient statistic for past productivity realizations in a normal learning model. Beyond the overall time effects at the quarter-year of hire level, it is also useful to control for state unemployment, especially given that the later part of the sample period corresponds with a weakening of the labor market.

¹⁷There are also persistent differences across drivers in underlying weekly productivity, with a standard deviation of roughly 275 miles (Hoffman and Burks, 2017). It is not immediately clear that one wants to control for average productivity to date in estimating the overall treatment effects of the contracts on quitting. Changes in average productivity to date may reflect idiosyncratic demand factors beyond a driver’s control, and this would be a reason to control for average productivity. However, changes in underlying productivity may also be reflective of worker selection. We thus believe it is useful to show results with and without controlling for average productivity to date. In practice, this has little impact on the results.

¹⁸That non-competes reduce mobility is also found in a recent study by Balasubramanian et al. (2017), who find that, all else equal, being in the US state with the highest non-compete enforceability increases job tenures of high-tech workers by 7.5% relative to the US state with the lowest non-compete enforceability.

Effects by quarter of tenure. To examine further how the effects of the contracts varied with tenure, we estimate a linear probability model of quitting on interactions of the contracts with worker quarter of tenure. Estimates are shown in Figure 2. Beyond the 16 contract X quarter of tenure terms (12m*Q1, 12m*Q2, ..., 12m*Q8, 18m*Q1, 18m*Q2, ..., 18m*Q8), we also include interactions of the contracts with dummies for tenure being greater than 8 quarters.¹⁹ Under the 12-month contract, relative to no contract, the largest negative impacts on quitting are observed in weeks 27-52, with quitting increasing in weeks 53-65. This is probably workers postponing quitting until their 12-month contract is over and seems consistent with basic worker optimization.

Under the 18-month contract, relative to no contract, negative impacts on quitting are especially strong in quarters 5-6. That is, drivers appear reluctant to quit in quarters 5 and 6, even though (unlike the 12-month contract) the penalty is pro-rated and gradually decreasing. We do not fully understand this, but we speculate this could reflect a (perhaps psychic) value to some drivers of “finishing the contract.” We return to this issue later in Section 5.1, where we compare observed impacts of the contracts to those simulated from a model.

A potential concern is that quit impacts could be driven by economic conditions. Some of the drivers with the 18-month contract were hired toward the end of the sample period (2002-2009) when the labor market began to slow down. We are already addressing national shocks by including current quarter-year fixed effects (as well as month-year of hire fixed effects), and the results in Figure 2 already control for a driver’s annual state unemployment rate. While monthly state unemployment rates may have more noise than annual state unemployment rates, we repeated Figure 2 using monthly state unemployment rates and achieved a similar picture. Thus, we do not believe that changes in economic conditions are biasing the results.

Event study analysis. A complementary approach to identification is to analyze the quit patterns of new drivers hired before and after the contract changes using an event study. To maximize statistical power, we focus on tenure periods during which the contracts had a big effect, as indicated in Figure 2. For the transition from no contract to the 12-month contract, we analyze

¹⁹The coefficient on 12m*(Tenure > 8 quarters) is -2.0e-3 (se=6.4e-4) and the coefficient on 18m*(Tenure > 8 quarters) is -2.9e-3 (se=9.5e-4). That is, both contracts caused sizable reductions in quitting beyond quarter 8 of tenure, both of which are statistically significant at the 1% level. The coefficient on 12m*(Tenure > 8 quarters) is slightly larger in magnitude than the coefficient on 12m*Q8, whereas the coefficients on 18m*Q8 and 18m*(Tenure > 8 quarters) are quite similar. These coefficients are omitted from the figure for visual simplicity.

quitting in the 4th quarter of tenure (weeks 40-52). Under the 12-month contract, drivers may optimally wait until their year is up before quitting, whereas the same incentive is not present for drivers with no contract. For the transition to the 18-month contract, we analyze quitting in the 5th and 6th quarters of tenure (weeks 53-78). Those under the 12-month contract may have waited for the year to end, whereas these weeks are still under contract for the 18-month contract. For the transition from the 12-month to the 18-month contract, the event study is estimated using:

$$Quit_{isqt} = \sum_{j=\underline{T}}^{\bar{T}} \theta_j D_{sq}^j + X_{isqt} \lambda + \epsilon_{isqt} \quad (1)$$

where $Quit_{isqt}$ is a dummy for quitting in week t by driver i from school s with quarter of hire q and ϵ_{isqt} is an error. D_{sq}^j is a dummy for whether those in quarter of hire q at training school s are j periods from the introduction of the 18-month contract; formally, $D_{sq}^j = \mathbf{1}(q - e_s = j)$, where e_s is the quarter when school s adopted the 18-month contract. \underline{T} and \bar{T} define the range of event time under study. X_{isqt} is the same vector of control variables as in column 2 of Table 2 (and includes school fixed effects and month-year of hire fixed effects). We normalize $\theta_{-1} = 0$.

As seen in Figure 3, for the transition to the 12-month contract, the probability of quitting in a week between weeks 40-52 drops by roughly 0.6 percentage points.²⁰ For the transition to the 18-month contract, there is a substantial decrease in quitting during weeks 53-78 (by over 1 percentage point). These event study impacts are quite sizable relative to the average quit rates shown in Figure 1. The effects are sudden and persist over time.

4.2 Incentives or Selection?

A decrease in quitting from training contracts may result through incentive and/or selection effects. If a worker is penalized for quitting, he may become less likely to quit, i.e., an *incentive* effect. However, adding a training contract could also affect the *selection* of workers who choose to work at the firm, e.g., quit penalties might deter workers with low productivity or low taste for the job from signing up. We perform two tests suggesting that the effect of the training contracts on quitting operated primarily through incentives.

²⁰In panel (a), there is an imbalance in the panel due to the fact that 3 of the training schools only have 1 quarter of data before the contract change. In Appendix Figure D2, we re-make panel (a) while restricting to the training school for which we have the longest pre-period before the 12-month contract change (this school got the 12-month contract in fall 2002) and at which the 12-month contract is eventually introduced. The result is qualitatively similar to panel (a) of Figure 3. In contrast, for panel (b), we have significant pre- and post-periods for all schools.

Our first test is to examine whether selection occurred on various observable characteristics, the most obvious being productivity: Did adding a training contract lead more productive workers to begin working at the firm? As seen in Panel A of Table 3, there is no evidence that the contracts increased productivity, and the results are reasonably precisely estimated. Columns 1-2 present results without trimming, whereas columns 3-4 present results while trimming miles at the 5th and 95th percentiles (trimming after dropping 0 mile weeks) to reduce the influence of productivity outliers. In column 4, the 95% confidence intervals on the 12-month and 18-month contracts are about $[-22,+34]$ and $[-35,+36]$, respectively, meaning we can rule out that the contracts either decreased or increased productivity by more than about 35 miles per week (where 40 miles per week is around 2% of weekly productivity).²¹ We also see no impact of the contracts on weekly driver earnings (Appendix Table D2), which, given the null impact on miles in Table 3, is unsurprising given that drivers are primarily paid by the mile.

One can also test for selection by looking at whether workers with other characteristics (potentially correlated with taste for the job or tendency to quit) are more likely to choose to work for the firm once training contracts are in place. Panel B shows little evidence for this, with 1 out of 10 coefficients are statistically significant at 5% significance. An additional worker characteristic/outcome to examine is whether an employee gets fired. If training contracts deterred low-quality drivers from working at Firm A, one would expect a decrease in the rate of firing. Appendix Table D4 estimates negative coefficients (i.e., drops in firing from the training contracts), but none of the coefficients are statistically significant. We do not emphasize the firing results because the firing estimates are less precise than those for quitting (fires are about 3 times rarer than quits).

Our second test of selection aims at testing whether there was selection on unobserved taste for trucking. Suppose that there are two types of drivers: “Good drivers,” who are productive and who have a high taste for trucking, and “Bad drivers,” who are less productive and have a low taste for trucking. The training contract would induce positive selection if it caused a greater share of new workers at the firm to be “Good drivers.” If the contracts caused positive selection, controlling

²¹We display our test for selection using all weeks of data to maximize power. A concern, however, is that because more productive people are less likely to quit, the impact of the training contracts on quitting could confound their impact on selection. To address this issue, we re-did our analysis restricting to drivers in their first 6 months for drivers who stay with the firm for at least 6 months. In these results, we find no evidence that the contracts induced positive selection.

for productivity should reduce the estimated magnitude of the coefficients on the contract variables in quit hazard models. As can be seen in column 3 of Table 2, the contract dummy coefficients are very similar once productivity is controlled for. Thus, this test provides no evidence to support the idea that the contract induced significant positive selection.²²

These two tests do not find much evidence of selection due training contracts. Given the strong evidence of selection effects of contracts in other personnel settings (e.g., Lazear, 2000), why does positive selection here seem relatively small or absent?²³ One possibility is that for some reason the contract may not have been salient to drivers when they signed up for the job. However, this seems unlikely to be the case, as a discussion of training contracts was a mandatory part of interviews at Firm A, according to managers. What seems a more likely possibility to us is that workers lack private information about their productivity when signing up. Unlike in Lazear (2000), the workers here are new to long-haul trucking. Long-haul trucking is very different from most other jobs and it may be difficult to predict how good one will be at it.

Was quitting reduced among low-productivity or high-productivity workers? Following Lazear (2000), we can also think about selection of employees out of the firm instead of into it. Specifically, did the training contracts reduce quitting among low-productivity or high-productivity workers? Appendix Table D1 shows that there were fairly similar absolute quit impacts of the training contracts among both high- and low-productivity workers. To classify drivers as high- or low-productivity, we regress residualized miles per week (residualized on the controls in column 1 of Table 2) on driver fixed effects. Drivers with above-median fixed effects are defined as high-productivity drivers. For the 12-month contract, the absolute reduction in the weekly quit rate is slightly larger in magnitude for low-productivity drivers, whereas for the 18-month contract, the absolute reduction in the weekly quit rate is slightly larger in magnitude for high-productivity drivers.²⁴

²²This test is inspired by the test for selection in Lazear (2000), who tests for selection by analyzing whether the coefficient on the piece-rate dummy changes once individual fixed effects are added. He finds that the coefficient on the contract dummy falls by half, leading him to conclude that selection explains half the treatment effect of the contract. Our test is significantly more indirect, given that we cannot observe the same individual under multiple contractual regimes. Accounting for selection by adding individual fixed effects is also used in other papers such as Lafontaine and Shaw (2016), who study the phenomenon of serial entrepreneurship.

²³Another recent piece on evidence on incentives vs. selections comes from Manchester (2012). She shows that instituting a tuition reimbursement program at a university decreased staff turnover, and provides evidence of effects operating through selection.

²⁴It is important to note that high productivity drivers have substantially lower quitting risk compared to low

4.3 Threats to Identification

We discuss potential threats to identification and why they are unlikely to affect our results.

Endogenous Contract Changes. Our estimation assumes that implementation of training contracts is orthogonal to unobserved factors affecting quitting. However, if training contracts were implemented in areas expected to have higher quitting, then we will underestimate the effect of training contracts on quitting. Alternatively, if it was easier to implement training contracts where future quitting was expected to be lower, then we will overestimate the effect. As noted earlier, our conversations with managers suggest that contract roll-out seems unlikely to be driven by such factors. In the data, training contract adoption is not predicted by state unemployment rates (which may be correlated with unobserved factors affecting quitting), as seen in Table 4. Also, as noted above, the estimated reductions in quitting are actually a bit larger once training school-specific linear quarter-year of hire trends are included (columns 6-7 of Table 2), suggesting the estimated reductions are not the result of pre-existing trends in quitting.

Worker Sorting Between Schools. Another potential confound to identifying the impact of training contracts on quitting would be worker sorting between schools. (Note that workers sorting between schools is different from the above-mentioned possibility of workers selecting into the firm.) For example, a worker who believed he had a high chance of quitting might prefer to attend a training school that did not have a training contract. This is unlikely to be an issue at Firm A because drivers generally attend a training school based on their state of home residence. Specifically, 93% of drivers in the full sample attend the modal training school in their state.²⁵

Concurrent Firm Policy Changes. As mentioned above, the contracts did not affect driver earnings (Appendix Table D2), nor did they affect pay per mile (Appendix Table D3).

productivity drivers, as indicated also in Table 2. If we perform the analyses in Appendix Table D1, but using a Cox model instead of OLS, we see higher percentage reductions in quitting for high productivity than low productivity workers for both the 12-month and 18-month contracts. However, we prefer to present OLS results here because we believe that the absolute level of quitting reduction is more relevant than the percentage reduction in quitting for the issues of selection and whether the contracts could affect productivity by differentially affecting the quitting of different types of workers. We also repeated our OLS analysis while splitting productivity into 4 categories instead of 2 categories. Quit impacts were negative for both contracts for all 4 categories (though not always statistically significantly so), but as for 2 categories, there was no clear evidence that absolute quit impacts were larger as underlying productivity increased or decreased.

²⁵This is generally the training school closest to their state. The percentage attending the modal training school in their state is 89% if we look at all training schools instead of the 5 training schools in the full sample. Drivers who don't attend their state's modal training school often live in states roughly equidistant from two training schools.

Furthermore, according to managers, contract changes at training schools were not accompanied by other changes in Firm A policy, such as changes in worker benefits or applicant screening.

Tenure-Varying Contract Enforcement. Under the 12-month contract, quitting after 51 weeks leads to the same penalty as quitting a few weeks after training. Could it have been that training contracts were enforced differently depending on worker tenure? Such a possibility does not threaten identification of the overall impact of the training contracts on quitting, but may affect the interpretation of impacts by tenure. Although disaggregated data on contract enforcement are not available, managers said that worker tenure did not affect contract enforcement.²⁶

5 Structural Simulations of the Different Training Contracts

The analysis so far has considered reduced-form impacts of the contracts. In this section, we use a structural model developed and estimated in Hoffman and Burks (2017) to simulate retention, as well as calculate worker welfare and firm profits, under the different contracts. A main advantage of the structural approach is that it enables one to perform simulations while accounting for key features of the work environment studied, including learning about ability, unobserved heterogeneity in taste for the job, and idiosyncratic taste shocks. The welfare consequences of “locking in” a worker via a training contract may seem to depend importantly on these features; for example, locking in a worker who is roughly indifferent between his inside and outside options may have little worker welfare consequences, but may be much more important when there are large differences.

For brevity, we lay out the model in detail in Appendix C. Focusing on the main elements here, workers solve an “optimal stopping problem” of when (if ever) to quit the firm. Workers have an underlying productivity, which is initially unknown to them, but drawn from a known distribution. Every week, workers observe their miles, which is a noisy realization of their underlying productivity, and this enables them to learn about their productivity. The learning process is a generalization of a purely rational version of Bayes’ Rule, thereby accommodating a broader range of updating behavior than would be imposed by strict rationality. Workers are compensated by a piece rate,

²⁶We note also, if for some reason, there was some workers who were “let off the hook” somewhat for a quitting event coming shortly before 52 weeks, and workers understood this, we would expect that this would make the 12-month contract less impactful than it would be otherwise. That is, such time-varying contract enforcement would work against us finding strong results in weeks 40-52, where we are currently finding strong impacts, and correcting it would likely make the results stronger than they already are.

w_t , that depends on their tenure. Workers are risk-neutral and make stay/quit decisions in order to maximize their perceived expected utility, $V_t(\mathbf{x}_t)$:

$$V_t(\mathbf{x}_t) = \max_{d_t, d_{t+1}, \dots} E_t \left(\sum_{s=t}^{\infty} \delta^{s-t} u_s(d_s, \mathbf{x}_s) \mid d_t, \mathbf{x}_t \right). \quad (2)$$

where d_t is a decision each week t about whether to stay or quit; δ is the discount factor; and \mathbf{x}_t is the vector of state variables, which includes the realizations of past miles, y_1, \dots, y_{t-1} . If a worker stays, each week they receive their earnings, plus their persistent non-pecuniary taste for the job, plus an idiosyncratic error. That is, $u_t(1, \mathbf{x}_t) = \alpha + w_t y_t + \varepsilon_t^S$, where α is the mass-point-distributed non-pecuniary taste for the job and ε_t^S is the idiosyncratic error that is i.i.d distributed according to an Extreme Value Type 1 distribution. In contrast, if the worker quits, they pay their training contract fine (if any), and receive their outside option thereafter, and in the week of quitting they also receive an idiosyncratic error. That is, $u_t(0, \mathbf{x}_t) = -k_t + \frac{r_t}{1-\delta} + \varepsilon_t^Q$, where k_t is the training contract penalty for quitting in week t , r_t is the outside option per period, and ε_t^Q is the idiosyncratic error (drawn from the same distribution as ε_t^S).

The model in Hoffman and Burks (2017) is estimated on a subset of 699 workers facing the 12-month contract for whom we collected very rich data, most importantly, weekly data on workers' subjective beliefs about their productivity. These workers were hired in late 2005 and 2006 at one of the firm's training schools. The subjective belief data are quite valuable because they allow us to relax the rational expectations assumption in worker productivity forecasting, something which Hoffman and Burks (2017) show is quantitatively important. Burks et al. (2008) and Hoffman and Burks (2017) provide details on data collection for this subset of workers.

Critically, we are performing our out-of-sample prediction below using data that are not used in any fashion in Hoffman and Burks (2017).

5.1 Out-of-sample prediction of contract impacts on quitting

Of course, the extent to which a structural model is useful for examining welfare and profits depends critically on whether it provides a reasonable account of the data and of the impact of the training contracts. To do so, we simulate the full data generating process under the three different contractual regimes.²⁷

²⁷In simulating the 18-month contract, we account for the details of the 18-month contract described in footnote 10, including importantly that there was a bonus paid to drivers after two years that returned the roughly \$13 that

Figure 4 shows simulated worker survival under the three different regimes. Under the 12-month contract, many workers optimally postpone quitting until after one year, echoing our earlier reduced form results. To assess the similarity of these results to the quasi-experimental estimates from Table 2, we ran a Cox proportional hazard model using the simulated data. Table 5 shows coefficient estimates of -0.182 and -0.185, meaning that the 12-month and 18-month contracts reduce quitting by around 18-19% relative to no contract in the simulated data, estimates which lie within the range of estimates of quit reductions presented in Table 2. We simulated 40,000 workers for each of the three contractual conditions—because the sample size here is chosen by the researcher (and is thus arbitrary), we do not present standard errors in the table.²⁸

It is worth emphasizing that there is no mechanical reason why the structurally estimated quit reductions would be similar to the quasi-experimental estimates, as the model is estimated entirely based off of a subset of workers facing the 12-month contract. This suggests that the quit reductions we estimate are broadly consistent with those from a worker who is “optimizing.”²⁹

In addition to examining whether the overall simulated impacts of the contract are consistent with our quasi-experimental estimates, we can examine whether this is also true for the time path of estimates. To do this, Figure 5 replicates Figure 2, but using the simulated data. It shows how the 12-month contract affects quitting relative to no contract. Panel (a) of Figure 5 is similar to panel (a) of Figure 2, with significant quit reductions, particularly close to the end of the first year of tenure, and with an increase after that.

In contrast, for analyzing how the 18-month contract affects quitting relative to no contract, panel (b) of Appendix Figure 5 is very different relative to panel (b) of Figure 2. In Figure 5, the largest impacts of the 18-month contract take place early in a driver’s tenure when the contract penalty is still the highest. This is because the contract is pro-rated. For drivers who realize they are less productive as truckers or who have low taste for the job, it may be optimal to postpone quitting until quarter 4, 5, or 6 of tenure when the quit penalty is lower.³⁰

was deducted from their paycheck each week under the 18-month contract.

²⁸Simulating 40,000 workers per contract seems to be plenty of workers for the simulation to deliver accurate results. Performing our Cox model on the simulated data and clustering by driver, the standard error on the contract coefficients is 0.008 for both contracts.

²⁹We use the term “optimizing” with quotation marks for two reasons. First, as we discuss below, some of the simulated impacts over time differ substantially from the quasi-experimental estimates. Second, the structural model assumed is not a fully rational model. In particular, it allows workers to have productivity beliefs that may differ from those predicted by Bayes’ rule, with differences both in means and variances.

³⁰Besides seeing whether the model is consistent with the reduced form impacts of the training contracts, an

Even though the model is estimated on workers facing the 12-month contract, there is no mechanical expectation that it be able to replicate the 12-month contract impacts relative to the no contract regime. Behavior under the no contract regime represents an out-of-sample prediction for the model.

Discussion. For the 18-month contract relative to no contract, why might the time-path of simulated contract impacts differ so much from the quasi-experimental contract impacts? One possibility alluded to earlier is that there is a particular non-monetary benefit for workers of finishing the contract. For example, drivers may believe that there would be reputational consequences for their future labor market activities if they did not finish the contract or may feel that it is the “right thing to do” to finish their contract with the firm that trained them (e.g., out of reciprocity or gift exchange (Akerlof, 1982)).

Driver behavior could also potentially reflect less than fully rational or “behavioral” beliefs. After taking on an 18-month contract, drivers may mentally “anchor” (e.g., Ariely et al., 2003; Maniadis et al., 2014) on 18 months as the minimum amount they intend to stay even if they are not very happy with some aspects of the job. Another related possibility is that the contract highlighted 18 months of tenure as a goal for workers to achieve (e.g., Corgnet et al., 2015). A further possibility is that some drivers may not easily understand the difference between a cliff format contract (like the 12-month contract) and the 18-month pro-rated contract, and might have responded to the contract as if the penalty did not vary with tenure.³¹

Ultimately, however, it is difficult for us to distinguish the validity of these different explanations. Our contribution here is to document that the quit impacts of the 18-month contract differ from those predicted by a simple dynamic model.

alternative approach is to see whether the model can closely match retention curves out of sample (e.g., Card and Hyslop, 2005). This is not our preferred approach because although the variation across the training contracts is plausibly exogenous conditional on control variables, we do not have a randomized experiment. When we have attempted this approach, we found that the model estimated under 699 workers with the 12-month contract could do a reasonable job predicting no contract worker behavior out of sample, but struggled somewhat in predicting behavior under the 18-month contract.

³¹To explore this possibility, we simulated the structural model assuming that the 18-month contract was flat at its initial level instead of pro-rated. As seen in Appendix Figure D3, this does a better job matching the impacts of the 18-month contract relative to no contract, but it is still quite imperfect.

5.2 Impacts on profits and welfare

Profits are computed as production profits, plus penalties from the training contracts, minus firm costs from training. For a worker staying T periods at the firm before quitting, profits are given by:

$$\pi = \sum_{t=1}^T \delta^{t-1} ((P - mc - w_t)y_t - FC) + \sum_{t=1}^T \delta^{t-1} \theta k_t q_t - TC \quad (3)$$

Here, δ is the discount factor, P is the price charged by the firm for a mile of shipping, mc denotes the non-wage marginal cost per mile (such as fuel costs and truck wear), y_t is the driver’s productivity (miles in a week), FC denotes fixed costs per week (e.g., back office support for drivers), q_t is a dummy variable for quitting in a given week, θ is the share of the training contract penalty collected by the firm, and TC represents the training cost per worker. Worker welfare is calculated by adding up earnings in trucking, taste for trucking, idiosyncratic shocks, and realizations of the fixed outside option using equation (6). For calculating firm profits and worker welfare, we use the same parameters and use the same simulation procedure as in Hoffman and Burks (2017) (also detailed in this paper in Appendix C.2).³²

Table 5 shows that profits are higher with the two training contracts (compared to no contract), but that worker welfare is lower. In addition, profits are slightly higher and worker welfare slightly lower under the 18-month contract relative to the 12-month contract. For the firm, workers stay longer under the contracts, so they have more of an opportunity to reap profits after having made a training investment. For the worker, the contracts limit worker ability to costlessly leave the job if they find it to be non-lucrative or unsatisfying.³³

Why may the firm not have been optimizing to start with? For the firm we study, profits were increased by adopting a training contract versus having no contract. However, ad-

³²As in Hoffman and Burks (2017), we focus separately on profits and worker welfare, as opposed to analyzing total welfare. We have found our conclusions on profits and worker welfare (separately) to be quite robust to different assumptions. However, across a range of counterfactuals that we have considered using our model (including for work outside of this paper), we found total welfare to depend more closely on particular assumptions made. That some counterfactuals may not lead to unambiguous changes in total welfare is not surprising, given our model allows for multiple market failures, including private information about taste for the job and shocks; potentially biased beliefs; quitting externalities; and monopsony power in the training market.

³³One important feature of the structural model and counterfactual simulations is that the quality of new workers is not affected by which training contracts are used. This is consistent with our empirical results in Section 4.2 that the observed contracts do not appear to have significantly affected worker quality. It could be the case that a sufficiently punitive training contract would affect worker quality, but that the observed contracts do not rise to this level.

ditional attempts at optimization toward picking a better contract produced more modest gains. Some readers may be tempted to ask the question of why might the firm may not have been optimizing in the first place. A simple explanation is that the use of training contracts may be thought of as a management practice, for which technological adoption is not immediate, as explored in Bloom et al. (2016). That firms may gradually learn about successful human resource management practices is a common theme in the personnel economics literature. For a recent example, see Friebel et al. (2017), where adoption of a particular management practice (namely, a shop-level bonus) led to significant improvements in worker performance.

6 Conclusion

Given significant concern about under-investment, understanding what makes training profitable for firms is critical both for economic theory and for policy. This paper explores the role of commonly-used contracts that fine workers for quitting after training.

Using plausibly exogenous contractual variation created by the staggered introduction of training contracts across training schools within a firm, we show that implementing a training contract reduced quitting by around 15%. The impacts on quitting are evident across different research designs, including difference-in-difference and event study analysis, and are strongest when workers are nearing the end of their contract. These impacts appear to be driven by making employees less likely to quit, as opposed to selecting better employees, though we caveat that our analysis faces the challenge that individuals are observed only under one contract; thus, we cannot control for individual fixed effects to employ a common test of selection vs. incentive effects of contracts. Simulating the contract changes in a structural model, we show that observed quit impacts are in line with simulated predictions; that the time path of quits exhibits aspects of optimization (for the 12-month contract but not the 18-month contract); and that the contracts significantly increase firm profits from training and decrease worker welfare.

While we find that the two contracts reduced worker welfare relative to not having a contract at this firm, this does not imply that the existence of training contracts reduces worker welfare (e.g., that worker welfare would increase if training contracts were banned). If training contracts were not available, it is possible that firms would not be willing to make as much investment in worker

training, possibly limiting the number of drivers that could go into trucking. Furthermore, our analysis of profits and welfare is a “partial equilibrium” analysis for one particular firm, and does not analyze what happens if all firms went from having no contract to having training contracts. Future research with data on multiple firms is needed to address such issues.³⁴ Still, our analysis shows that training contracts can play a significant role for profits and welfare within one firm.

Although we focus on a single industry, training contracts are used in many jobs, both other blue-collar jobs (e.g., mechanics and electricians) and high-skill jobs (e.g., pilots, accountants, and stockbrokers). Future work should examine whether the impacts of training contracts are similar. For example, one might imagine that workers in high-skill jobs may be more adept at renegotiating contracts compared to workers in lower-skill jobs. As such, training contracts may be thought to have less of an impact in those settings.

Finally, future work may wish to examine how training contracts may interact with other contracts used by firms, as well as heterogeneity in the types of workers that are most affected by training contracts. Firms already use a variety of contractual devices to reduce employee turnover, such as employee stock options (Oyer and Schaefer, 2005) and deferred compensation schemes. Would training contracts serve as complements or substitutes with these sorts of contracts? Turning to types of workers, our study found interesting results that worker responses for one of the two training contracts seemed to differ from those predicted by dynamic optimization. Could it be that worker responsiveness to training contracts varies with cognitive ability (or could even affect selection along cognitive ability)? While there is growing research in public finance that labor supply (e.g., Liebman and Zeckhauser, 2004) and consumer demand (e.g., Ito, 2014) exhibit aspects of limited rationality in responding to contracts, such insights may also be fruitfully applied to personnel economics. In addition, the worker overconfidence documented by Hoffman and Burks (2017) could have important interactions with the effectiveness of training contracts. We hope that such issues are explored in future research.

³⁴There are other questions that might be explored with data on multiple firms. For example, with data on multiple firms (along with plausibly exogenous contractual variation), researchers could study if there is a relation between training contract usage/structure and the quality of training provided. Still, by studying employee turnover and firm profits from training, our paper provides an important step in understanding how training contracts affect the incentive of firms to make training investments.


References

- Acemoglu, Daron and Jorn-Steffen Pischke**, “Why Do Firms Train? Theory And Evidence,” *Quarterly Journal of Economics*, 1998, *113* (1), 78–118.
- and – , “Beyond Becker: Training in Imperfect Labour Markets,” *Economic Journal*, 1999a, *109* (453), F112–42.
- Akerlof, George A.**, “Labor Contracts as Partial Gift Exchange,” *Quarterly Journal of Economics*, 1982, *97* (4), 543–569.
- Ariely, Dan, George Loewenstein, and Drazen Prelec**, ““Coherent arbitrariness”: Stable demand curves without stable preferences,” *Quarterly Journal of Economics*, 2003, *118* (1), 73–106.
- Autor, David H.**, “Why Do Temporary Help Firms Provide Free General Skills Training?,” *Quarterly Journal of Economics*, 2001, *116* (4), 1409–1448.
- Baker, George P. and Thomas N. Hubbard**, “Contractibility And Asset Ownership: On-Board Computers and Governance In U. S. Trucking,” *Quarterly Journal of Economics*, 2004, *119* (4), 1443–1479.
- Balasubramanian, Natarajan, Jin Woo Chang, Mariko Sakakibara, Jagadeesh Sivadasan, and Evan P. Starr**, “Locked In? The Enforceability of Covenants Not to Compete and the Careers of High-Tech Workers,” 2017. Working paper, University of Michigan.
- Baron, James and David Kreps**, *Strategic Human Resources: Frameworks for General Managers*, New York, NY, 1999.
- Barron, John M., Mark C. Berger, and Dan A. Black**, “Do Workers Pay for On-The-Job Training?,” *Journal of Human Resources*, 1999, *34* (2), 235–252.
- Becker, Gary**, *Human Capital: A Theoretical and Empirical Analysis*, Chicago, IL: University of Chicago Press, 1964.
- Belzer, Michael**, *Sweatshops on Wheels: Winners and Losers in Trucking Deregulation*, Oxford University Press, 2000.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen**, “Management as a Technology?,” Working Paper 22327, National Bureau of Economic Research June 2016.
- BLS**, “Bureau of Labor Statistics Occupational Outlook Handbook, 2010-11: Truck Drivers and Driver/Sales Workers,” 2010.
- Brunello, Giorgio and Alfredo Medio**, “An Explanation of International Differences in Education and Workplace Training,” *European Economic Review*, 2001, *45* (2), 307–322.
- Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman**, “The Value of Hiring through Employee Referrals,” *Quarterly Journal of Economics*, 2015, *130* (2), 805–839.
- , **Jeffrey Carpenter, Lorenz Goette, Kristen Monaco, Kay Porter, and Aldo Rustichini**, “Using Behavioral Economic Field Experiments at a Firm: The Context and Design of the Truckers and Turnover Project,” in “The Analysis of Firms and Employees: Quantitative and Qualitative Approaches” 2008.


- Cappelli, Peter**, “Why Do Employers Pay for College?,” *Journal of Econometrics*, 2004, *121* (1-2), 213–241.
- Card, David and Dean R. Hyslop**, “Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers,” *Econometrica*, 2005, *73* (6), 1723–1770.
- Chiappori, Pierre-Andre and Bernard Salanie**, “Testing Contract Theory: A Survey of Some Recent Work,” in L. Hansen M. Dewatripont and S. Turnovsky, eds., *Advances in Economics and Econometrics*, 2003.
- Corngnet, Brice, Joaquín Gómez-Miñambres, and Roberto Hernán-Gonzalez**, “Goal Setting and Monetary Incentives: When Large Stakes are Not Enough,” *Management Science*, 2015, *61* (12), 2926–2944.
- Dustmann, Christian and Uta Schoenberg**, “What Makes Firm-Based Vocational Training Schemes Successful? The Role of Commitment,” *American Economic Journal: Applied Economics*, 2012, *4* (2), 36–61.
- Friebel, Guido, Matthias Heinz, Miriam Krüger, and Nick Zubanov**, “Team incentives and performance: Evidence from a retail chain,” *American Economic Review*, 2017, *Forthcoming*.
- Galiani, Sebastian, Alvin Murphy, and Juan Pantano**, “Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment,” *American Economic Review*, 2015, *105* (11), 3385–3415.
- Gruber, Jonathan and Brigitte C. Madrian**, “Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature,” 2002. NBER Working Paper 8817.
- Hoffman, Mitchell and Stephen V. Burks**, “Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Returns from Training,” 2017. NBER Working Paper 23240.
- Hubbard, Thomas N.**, “Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking,” *American Economic Review*, 2003, *93* (4), 1328–1353.
- Ito, Koichiro**, “Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing,” *American Economic Review*, 2014, *104* (2), 537–563.
- Jovanovic, Boyan**, “Job Matching and the Theory of Turnover,” *Journal of Political Economy*, 1979, *87* (5), 972–90.
- **and Yaw Nyarko**, “Learning by Doing and the Choice of Technology,” *Econometrica*, 1996, *64* (6), 1299–1310.
- Keane, Michael P and Kenneth I. Wolpin**, “Exploring the usefulness of a nonrandom holdout sample for model validation: Welfare effects on female behavior,” *International Economic Review*, 2007, *48* (4), 1351–1378.
- Kraus, Anthony W.**, “Repayment Agreements for Employee Training Costs,” *Labor Law Journal*, 1993, *44* (1), 49–55.
- , “Employee Agreements for Repayment of Training Costs: The Emerging Case Law,” *Labor Law Journal*, 2008, pp. 213–226.

- Lafontaine, Francine and Kathryn Shaw**, “Serial Entrepreneurship: Learning by Doing?,” *Journal of Labor Economics*, 2016, *34* (S2), S217–S254.
- Lazear, Edward P.**, “Job Security Provisions and Employment,” *Quarterly Journal of Economics*, 1990, *105* (3), 699–726.
- , “Performance Pay and Productivity,” *American Economic Review*, 2000, *90* (5), 1346–1361.
- Liebman, Jeffrey B and Richard J Zeckhauser**, “Schmeduling,” 2004. Working paper, Harvard University.
- Lynch, Lisa M.**, “The Economics of Youth Training in the United States,” *Economic Journal*, 1993, *103* (420), 1292–302.
- MacLeod, Bentley W**, “Reputations, relationships, and contract enforcement,” *Journal of economic literature*, 2007, *45* (3), 595–628.
- Madrian, Brigitte C.**, “Employment-Based Health Insurance and Job Mobility: Is There Evidence of Job-Lock?,” *Quarterly Journal of Economics*, 1994, *109* (1), 27–54.
- Manchester, Colleen**, “Investment in General Human Capital and Turnover Intention,” *American Economic Review Papers & Proceedings*, 2010, *100* (2), 209–13.
- , “How Does General Training Increase Retention? Examination Using Tuition Reimbursement Programs,” *Industrial and Labor Relations Review*, 2012, *65* (4), 951–974.
- Maniadis, Zacharias, Fabio Tufano, and John A List**, “One swallow doesn’t make a summer: New evidence on anchoring effects,” *American Economic Review*, 2014, *104* (1), 277–290.
- Marx, Matt, Deborah Strumsky, and Lee Fleming**, “Mobility, skills, and the Michigan non-compete experiment,” *Management Science*, 2009, *55* (6), 875–889.
- Naidu, Suresh and Noam Yuchtman**, “Coercive Contract Enforcement: Law and the Labor Market in Nineteenth Century Industrial Britain,” *American Economic Review*, 2013, *103* (1), 107–144.
- , **Yaw Nyarko, and Shing-Yi Wang**, “Monopsony Power in Migrant Labor Markets: Evidence from the United Arab Emirates,” *Journal of Political Economy*, 2016, *124* (6), 1735 – 1792.
- Oyer, Paul and Scott Schaefer**, “Why Do Some Firms Give Stock Options to All Employees?: An Empirical Examination of Alternative Theories,” *Journal of Financial Economics*, 2005, *76* (1), 99–133.
- Peterson, Jonathan**, “Employee Bonding and Turnover Efficiency,” 2010. Mimeo, Cornell University.
- Pigou, Arthur C.**, *Wealth and Welfare*, London: Macmillan, 1912.
- Todd, Petra E. and Kenneth I. Wolpin**, “Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility,” *American Economic Review*, 2006, *96* (5), 1384–1417.


Figure 1: Training Contracts and the Hazard of Quitting



(a) No Contract




(b) 12-month Contract



(c) 18-month Contract

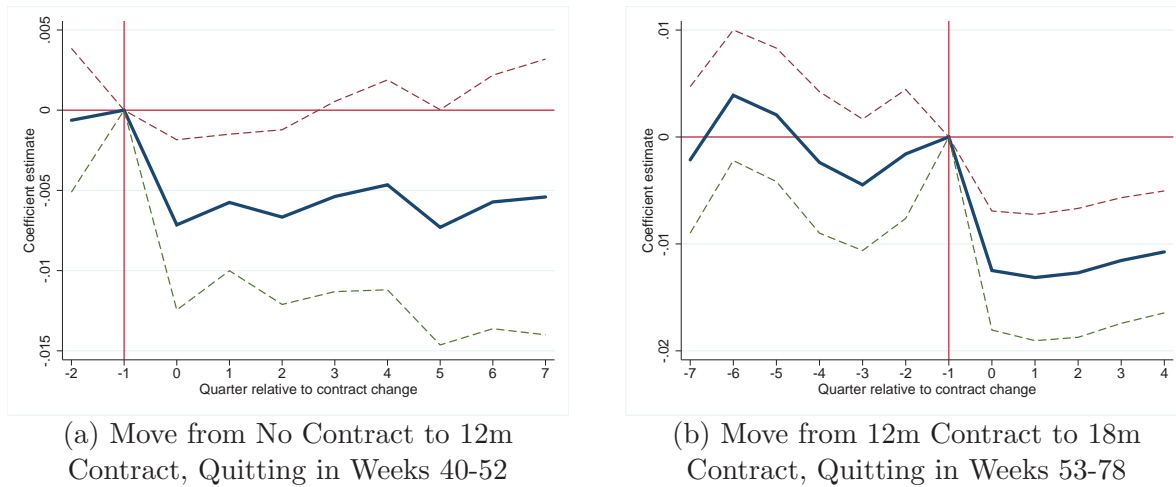
Notes: These figures plot the quitting hazard under the 3 contractual regimes, using drivers in the Firm A full sample. It focuses only on quits (fires are ignored). An Epanechnikov kernel is used. The bandwidth is 4 weeks for the no contract regime, 2 weeks for the 12-month contract regime, and 3 weeks for the 18 month contract regime. In each panel, the x-axis is driver tenure in weeks. The sample size is withheld to protect firm confidentiality, but is much larger than 5,000 drivers.

Figure 2: The Impact of Training Contracts on Quitting by Quarter of Tenure (with 95% Confidence Intervals)




Notes: This figure plots the estimated effect of the two training contracts on quitting at different tenure levels. The solid line denotes the coefficient estimate, with the dotted lines denoting the 95% confidence interval. The coefficients are from an OLS regression of quitting (0 or 1) for a driver in a given week on training contract-quarter of tenure interactions and controls. The controls are the same as in column 2 of Table 2 except we also include week of tenure dummies (in place of the baseline hazard function included in the Cox model in Table 2). Standard errors are clustered at the school-week of hire level. The two figures are based on one regression, with panel (a) plotting interactions of the 12-month contract and different quarters of tenure, and with panel (b) doing the same with respect to the 18-month contract. Beyond the 16 contract X quarter of tenure terms ($12m*Q1$, $12m*Q2$, ..., $12m*Q8$, $18m*Q1$, $18m*Q2$, ..., $18m*Q8$), we also include interactions of the contracts with dummies for tenure being greater than 8 quarters. The coefficient on $12m*(Tenure > 8 \text{ quarters})$ is $-2.0e-3$ ($se=6.4e-4$) and the coefficient on $18m*(Tenure > 8 \text{ quarters})$ is $-2.9e-3$ ($se=9.5e-4$). That is, both contracts caused sizable reductions in quitting beyond quarter 8 of tenure, both of which are statistically significant at the 1% level. The coefficient on $12m*(Tenure > 8 \text{ quarters})$ is slightly larger in magnitude than the coefficient on $12m*Q8$, whereas the coefficients on $18m*Q8$ and $18m*(Tenure > 8 \text{ quarters})$ are quite similar. These coefficients are omitted from the figure for visual simplicity.

Figure 3: Event Studies: The Impact of Training Contracts on Quitting, Comparing Before and After the Contract Changes (with 95% Confidence Intervals)



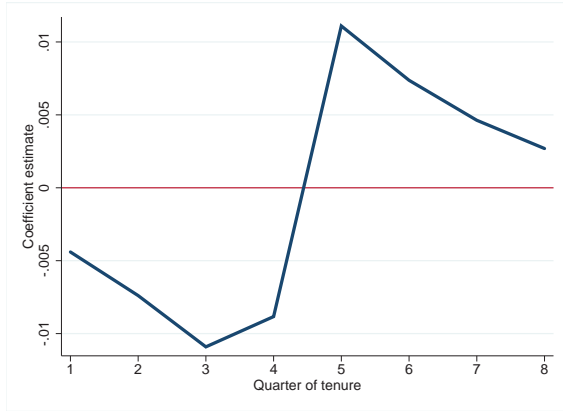
Notes: The solid line denotes the coefficient estimate, with the dotted lines denoting the 95% confidence interval. Panel (a) analyzes quitting in weeks 40-52 before and after the change to the 12-month contract whereas panel (b) analyzes quitting in weeks 53-78 before and after the change to the 18-month contract. The x-axis denotes “event time,” reflecting the contracts being changed at different training schools at different times. Each “quarter” refers to the workers hired in a 3-month block. Quarter 0 is the first quarter after the introduction of each training contract. Each of the two panels reflects a different regression (see equation 1 in the main text), where the control variables (separate from the event time dummies) are the same as in column 2 of Table 2 except we also include week of tenure dummies (in place of the baseline hazard function included in the Cox model in Table 2). Standard errors are clustered at the school-week of hire level. For panel (a), the plotted coefficient for “-2” is an indicator for event time equal to “-3” or to “-2.” We combine them together to increase power. Beyond the event time coefficients plotted, we also include a dummy for event time greater than or equal to 8. Panel (a) excludes one training school that never gets the 12-month contract. For panel (b), beyond the event time dummies plotted, we also include a dummy for event time -8 or less, as well as a dummy for event time of 5 or greater. The number of quarters before and after the event varies between the two contracts due to limits on the number of quarters of data available before or after contract changes. Panel (b) includes one training school that transitioned directly from no contract to the 18-month contract, but the figure is very similar if that training school is removed. For that school, the “event” is the change from no contract to the 18-month contract.

Figure 4: Simulated Retention Under the Three Contracts using the Structural Model




Notes: This figure analyzes the ability of the structural model from Hoffman and Burks (2017) to make out-of-sample predictions. The model is estimated off of 699 drivers under the 12-month contract. We simulate 40,000 drivers in each simulation.

Figure 5: The Impact of Training Contracts on Quitting by Quarter of Tenure using the Simulated Data



(a) 12 Month Contract



(b) 18 Month Contract

Notes: This figure is similar to Figure 2. The difference is that we use the simulated data (also shown in Figure 4) instead of the actual data. Because the sample size is arbitrary here, no standard errors are shown. Because the data are simulated, there are no additional control variables used. As in Figure 2, beyond the 16 contract X quarter of tenure terms (12m*Q1, 12m*Q2, ..., 12m*Q8, 18m*Q1, 18m*Q2, ..., 18m*Q8), we also include interactions of the contracts with dummies for tenure being greater than 8 quarters. The coefficient on 12m*(Tenure > 8 quarters) is 0.0014 and the coefficient on 18m*(Tenure > 8 quarters) is 0.0021. As in Figure 2, these coefficients are omitted from the figure for visual simplicity.

Table 1: Summary Statistics

Variable	Mean
Female	0.09
Black	0.19
Hispanic	0.04
Age	37
Married	0.38
No contract	0.10
12-month contract	0.71
18-month contract	0.19
Number of workers	N

Notes: The sample is drawn from trained drivers at Firm A from 2002 to 2009. The exact number of drivers, N , is withheld to protect the confidentiality of the firm, $N \gg 5,000$. See Appendix A.3 for more details on data and sample construction.

Table 2: Impact of the Training Contracts on Quitting – Cox Model, Diff-in-Diff

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
12m contract	-0.141*** (0.046)	-0.140*** (0.046)	-0.143*** (0.049)	-0.146*** (0.050)	-0.146*** (0.047)	-0.202*** (0.050)	-0.207*** (0.054)
18m contract	-0.148*** (0.064)	-0.150*** (0.065)	-0.151*** (0.068)	-0.148*** (0.068)	-0.151*** (0.065)	-0.234*** (0.075)	-0.242*** (0.079)
Average miles to date (in hundreds of miles)			-0.090*** (0.002)	-0.090*** (0.002)			-0.090*** (0.002)
Quarter-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year of hire FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Training School FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control for state unemp rate	No	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	No	No	No	Yes	Yes	Yes	Yes
School-specific time trends	No	No	No	No	No	Yes	Yes

Notes: An observation is a driver-week. The regressions are Cox proportional hazard models, with coefficients reported and standard errors clustered by school-week of hire in parentheses. It focuses only on quits (fires are ignored). Driver tenure is controlled for non-parametrically. The “state unemp rate” refers to the annual unemployment rate in a driver’s state of residence in the current year. Beyond the controls listed, all regressions include work type controls, as well as dummies for the roughly 2-3 week “transition periods” in between contract phase-in at a school and the first group of trainees graduating CDL training with the new training contract at the school. Demographic controls are controls for gender, race, marital status, and age. Average miles to date is a driver’s average weekly productivity to date over positive mile weeks and is given in terms of hundreds of miles driven per week; its value is set to 0 if the driver has not had any positive mile weeks to that date (and we also include a dummy variable to indicate average miles to date being missing in those circumstances). The “school-specific time trends” are training school-specific linear time trends in month-year of hire. The coefficients can be interpreted as approximate percentage changes in quitting, e.g., column (1) indicates the contracts reduced quitting by about 14 and 15 percent. The sample size is withheld to protect firm confidentiality, but is much larger than 100,000 driver-weeks. The sample is consistent across columns and does not restrict based on driver tenure (it does not restrict to driver-weeks in the first 8 quarters of tenure). For further details on the data, see Appendix A.3. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 3: Training Contracts have Limited Selection Effects

Panel A: Selection on Productivity					
Sample:	All Weeks		Trim 5/95%		
	(1)	(2)	(3)	(4)	
12m contract	-11.55 (21.69)	-9.21 (21.88)	5.39 (14.51)	6.16 (14.42)	
18m contract	-23.45 (26.60)	-22.71 (26.93)	1.16 (17.86)	0.56 (17.90)	
Demographic Controls	No	Yes	No	Yes	
Panel B: Selection on Characteristics					
Dep. Var:	(1)	(2)	(3)	(4)	(5)
	Black	Hispanic	Female	Married	Age
12m contract	-0.010 (0.016)	0.028*** (0.010)	0.001 (0.012)	-0.018 (0.020)	-0.153 (0.452)
18m contract	-0.004 (0.023)	0.015 (0.012)	0.013 (0.015)	0.009 (0.025)	0.387 (0.566)

Notes: Standard errors clustered by school-week of hire in parentheses. Panel A reports OLS regressions of productivity (miles per week) on training contract dummies and controls. In Panel A, an observation is a driver-week. “Trim 5/95%” refers to trimming the lowest 5% and highest 5% of the miles observations (ignoring all 0 mile weeks). In columns 1 and 3 of Panel A, the controls are the same as in column 2 of Table 2, except that we control for a 5th order polynomial in week of tenure (in place of the baseline hazard function). Demographic controls are the same as in Table 2. Panel B reports OLS regressions of driver characteristics on training contract dummies and controls. In Panel B, an observation is a driver. The regressions include quarter-year of hire fixed effects, work type controls, school controls, and the annual state unemployment rate at time of hire. The sample size is withheld to protect firm confidentiality, but is much larger than 100,000 driver-weeks in panel (a) and is much larger than 5,000 drivers in panel (b). * significant at 10%; ** significant at 5%; *** significant at 1%

Table 4: Exogeneity Check: State Unemployment Rates Do Not Predict Training Contract Changes

Dep var:	Has 12m contract (1)	Has 18m contract (2)
State unemployment rate	0.002 (0.005)	-0.010 (0.009)
R-squared	0.850	0.865

Notes: This table examines whether state unemployment rates predict whether drivers have training contracts or not, in an effort to examine whether training contract changes may have been correlated with labor market conditions. An observation is a driver. Each column is an OLS linear probability model where the dependent variable is whether a driver has a 12-month contract (versus no contract) and whether a driver has an 18-month contract (versus the 12-month contract). Column 1 analyzes drivers with either no contract or the 12-month contract. Column 2 analyzes drivers with either the 12-month contract or the 18-month contract. Standard errors in parentheses are clustered by driver’s state of residence. Both regressions include quarter-year of hire fixed effects, training school fixed effects, and work type controls. The sample size is withheld to protect firm confidentiality, but is much larger than 5,000 drivers. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 5: Simulated Retention, Profits, and Welfare Under Three Contractual Regimes

Contractual regime:	No contract	12 month	18 month
Drop in quitting relative to no contract: Cox model coefficients estimated on simulated data		-0.182	-0.185
Retention at 20wks	0.72	0.79	0.82
Retention at 40wks	0.41	0.56	0.53
Retention at 60wks	0.30	0.43	0.40
Profits per worker	\$3,123	\$4,907	\$5,210
Welfare per worker	\$57,933	\$54,840	\$54,231

Notes: This table presents profits per worker and welfare per worker under the three different training contracts used by Firm A. Profits and welfare are described in Section 5 of the main text, and are detailed further in Hoffman and Burks (2017). We use the structural parameters presented in Appendix Table D5. “Drop in quitting relative to no contract” is based on a Cox proportional hazard model estimated using simulated data. We simulate 40,000 workers for each of the 3 contractual regimes, simulating each worker for up to 130 weeks each. The Cox coefficients reported are analogous to those in Table 2. The other numbers (on retention, profits per worker, and welfare per worker) in the table are based on simulating 3,000 workers for up to 1,300 weeks each. “Retention at 20 weeks” represents the share of workers remaining with the firm after 20 weeks. Profits per worker and welfare per worker are both calculated after 110 weeks.