

CONTRACTIBILITY AND ASSET OWNERSHIP: ON-BOARD COMPUTERS AND GOVERNANCE IN U. S. TRUCKING*

GEORGE P. BAKER AND THOMAS N. HUBBARD

We investigate how contractual incompleteness affects asset ownership in trucking by examining cross-sectional patterns in truck ownership and how truck ownership has changed with the diffusion of on-board computers (OBCs). We find that driver ownership of trucks is greater for long than short hauls, and when hauls require equipment for which demands are unidirectional rather than bidirectional. We then find that driver ownership decreases with OBC adoption, particularly for longer hauls. These results are consistent with the hypothesis that truck ownership reflects trade-offs between driving incentives and bargaining costs, and indicate that improvements in the contracting environment have led to less independent contracting and larger firms.

I. INTRODUCTION

What determines who owns assets in the economy? This question, which goes back at least as far as Coase [1937], is central to understanding firms' boundaries. The theoretical work since Coase has highlighted a number of factors, including asset specificity, noncontractible investments, and multitasking problems, as important in the determination of asset ownership.¹ These theories all share the view that optimal asset ownership hinges on the contracting environment. In this paper we examine the relationship between asset ownership and the contracting environment in the United States trucking industry. Using detailed truck-level data, we investigate what determines whether drivers own the trucks they operate, and how ownership patterns change as the contracting environment changes.

We develop an analytic framework that draws heavily on the property rights theories of Grossman and Hart [1986] and Hart and Moore [1990]. This framework highlights how contractual incompleteness can affect the comparative advantage of using an owner-operator for a haul relative to a company driver. We propose that an important benefit of having the driver own the truck

* We would like to thank Gary Chamberlain, Robert Gibbons, Oliver Hart, Francine Lafontaine, Paul Oyer, Brian Silverman, Margaret Slade, Jerry Zimmerman, anonymous referees, and many seminar participants for comments. We also thank Michael Crum and several dispatchers and drivers for useful discussions. We gratefully acknowledge support from NSF grant SES-9975413 and the Harvard Business School Division of Research.

1. See Klein, Crawford, and Alchian [1978], Williamson [1975, 1985], Grossman and Hart [1986], Holmstrom and Milgrom [1994], and many others.

is that the driver drives in ways that better preserve the truck's value. However, an important drawback is that, when residual rights of control over the truck are allocated to the driver, the individual responsible for planning how trucks should be used—the dispatcher—no longer has critical control rights over the truck. This leads to inefficiencies associated with bargaining over the truck's use; for example, dispatchers may underinvest in finding good “backhauls” (return trips) for trucks, or drivers may engage in inefficient rent-seeking behavior. This analytic framework generates empirical propositions that allow us to examine both sides of the trade-off that we propose.

Our empirical analysis uses truck-level data from the Census' 1987 and 1992 Truck Inventory and Use Surveys. We first show that driver ownership of trucks is greater for longer hauls, and when hauls require equipment for which demands are unidirectional (i.e., backhauls are unlikely) rather than bidirectional. We then develop an empirical strategy that allows us to examine how a new monitoring technology (on-board computers or “OBCs”) that becomes available in the middle of our sample period affects ownership patterns. We show that driver ownership of trucks decreases with OBC adoption, and that this relationship is strongest for long hauls, where the monitoring technology is the most valuable. Finally, we test whether OBCs change how drivers drive, by assessing the fuel economy of trucks driven by company drivers and owner-operators with and without OBCs. We find that while fuel economy is better for trucks with OBCs than without them, this difference is greater for company drivers than owner-operators.

Overall, our evidence supports our analytic framework, and suggests that contractual improvements have led to more integrated asset ownership in trucking, especially in circumstances where allocating control rights to drivers is costly. Contractual improvements have led carriers to subcontract less of their hauls to owner-operators, and thus have led to larger, more integrated firms.

This paper extends several strains of the empirical literature on organizations.² In particular, it is closely related to Baker and

2. Other recent work that investigates organizational issues in trucking includes Chakraborty and Kazarosian [1999], Hubbard [2000, 2001], Lafontaine and Masten [2002], and Nickerson and Silverman [2003]. See Brickley and Dark [1987], Lafontaine [1992], and Shepard [1993] for evidence on contractibility and ownership in franchising, and Brynjolffson and Hitt [1997] and citations for

Hubbard [2003], which examines relationships between OBC adoption and shippers' make-or-buy decision; i.e., whether shippers use a truck from their private fleet for a haul, or outsource their shipping needs to for-hire carriers. In this companion piece, we propose a model in which shipper ownership of trucks is a function of the importance of service quality to a particular haul. We ignore service issues in this paper, because we do not believe them to be relevant to the margin we examine. As we explain in the other paper, the inefficiencies associated with using owner-operators for hauls involving service tasks (such as sorting cargo) are so large that they should not be used on these types of hauls. In practice, owner-operators are rarely used for hauls with significant service requirements.³

An outline of the rest of the paper follows. In Section II we describe the production process and contracting environment, highlighting how the contracting environment affects driver incentives. We then discuss asset ownership, and build the analytic framework that generates the hypotheses to be tested. In Section III we describe the data and present cross-sectional patterns with respect to ownership and OBC use. In Section IV we present and interpret our main results, estimates of relationships between OBC adoption and organizational change. In Section V we present some evidence of OBCs' incentive effects by examining relationships between OBC use and fuel economy. In Section VI we conclude.

II. INCENTIVE PROBLEMS AND ASSET OWNERSHIP IN TRUCKING

Production in trucking involves the movement of goods. Hauls differ along many dimensions, and the type of cargo determines what kind of trailer can be used.⁴ Packaged goods that do not require refrigeration can be hauled in nonrefrigerated vans,

evidence on relationships between information technology adoption and organizational form.

3. Driver ownership is an organizational option for all hauls, including those where shippers choose instead to use trucks from their internal fleet. Although many hauls for which shippers choose to use private fleets are inframarginal "company driver" hauls, excluding them from the empirical analysis would introduce sample selection problems. Like in our companion piece, our empirical analysis therefore uses data from all tractor-trailers in the TIUS (subject to some minor restrictions described later).

4. This is not the only variation in equipment requirements that can affect organizational form. Nickerson and Silverman [2003] argue that asset specificity (in the form of interactions between drive-train configurations and haul charac-

the most common class of trailer, but other goods require trailers that are more specialized to the specific good. For example, logs and vehicles are hauled on trailers that have special features that prevent them from rolling off. Demanders of trucking services are called *shippers*; suppliers are called *carriers*, which include both for-hire trucking firms and trucking divisions of firms that are not trucking specialists, so-called "private fleets." Dispatchers and drivers perform work for carriers, and it is common for a carrier's drivers to be a mix of *owner-operators* (drivers who own their truck) and *company drivers* (drivers who do not).

Dispatchers receive and solicit orders from shippers and assign trucks and drivers to hauls. The dispatcher's job is crucial for maintaining high levels of capacity utilization. One of dispatchers' principal tasks is scheduling "backhauls," or return trips. It is particularly valuable to set up a "backhaul" for trucks when hauls take them outside of their local area, and it is generally possible to do so when hauls use trailers for which demand tends to be bidirectional. This tends to be the case for trailers that are not too specialized to a particular type of cargo.⁵ Because the exact time and place of shippers' demands is usually unknown at the time trucks depart, dispatchers tend not to firm up their plans for the backhaul until trucks are en route, often around the time trucks arrive at their destination. Utilizing capacity efficiently generally implies deferring assignment-setting as much as possible.

Two types of incentive problems exist in the relationship between drivers and carriers. One involves how the truck is driven. It has traditionally been difficult to verify how drivers operate trucks, since they are operated remotely and, other than knowing whether the truck and driver arrived on time at their destination, the carrier has had little information about a truck once it is on the road. Wear and tear on the truck is minimized when drivers drive at a steady and moderate speed, but drivers may prefer to drive fast and then take longer breaks because it allows them to rest longer, visit friends, etc., and still arrive on time. Drivers' scope for this type of nonoptimal driving is

teristics) discourage drivers from owning certain trucks and provide evidence consistent with this.

5. Thus, demand for nonrefrigerated vans or platform trailers is generally bidirectional, but demand for trailers such as logging trailers tends to be unidirectional.

particularly high for longer hauls, because there is more opportunity to make up time.

In recent years, a new technology developed that allows carriers to monitor drivers' behavior much more closely. On-board computers come in two forms: trip recorders and Electronic Vehicle Monitoring Systems (EVMS).⁶ Trip recorders collect information about trucks' operation; one can think of them as trucks' "black boxes." They record when trucks are turned on and off, their speed over time, acceleration and deceleration patterns, fuel use, and variables related to engine performance. Data from trip recorders are collected when drivers return to their base; drivers give dispatchers a chart, floppy disk, or data cartridge with data. These data allow carriers to better know how the truck was driven and give mechanics information that allows them to better maintain the truck's engine. EVMS have several additional features that help dispatchers coordinate the movement of their fleets. For example, they can transmit trucks' real-time location to carriers, and allow dispatchers and drivers to send short text messages to each other. The advent of OBCs has significantly changed the ability of carriers to verify how drivers drive. We analyze the effect of this change below.

The second important incentive problem that affects the relationship between drivers and carriers results from the incomplete nature of contracting over the use of the truck. Agreements between carriers and drivers generally cover multiple periods, and hence multiple hauls, but they generally do not specify in advance exactly which hauls drivers will complete because flexibility in scheduling can be extremely important for capacity utilization. Conflicts between carriers and drivers arise because hauls vary in their desirability to drivers in ways that are not captured in agreements with carriers. Those that take drivers into congested or dangerous areas are less desirable than those that do not. Hauls that involve layovers or empty miles can be undesirable for "over-the-road" (i.e., nonlocal) drivers, whose compensation is generally output-based.⁷ Dispatchers negotiate with drivers to induce them to accept undesirable hauls, particularly when drivers are far from their base and carriers have no other

6. As of 1992, trip recorders cost about \$500. EVMS hardware cost \$3,000–\$4,000 to buy or about \$150/month to lease.

7. Drivers' compensation for intercity hauls is generally based on either miles, loaded miles, or a fraction of the haul's revenues, regardless of whether they own trucks. See Lafontaine and Masten [2002] for an analysis.

drivers in the area. This negotiation usually involves a combination of moral suasion, promises to assign drivers desirable hauls in the future, and sometimes pecuniary compensation.

II.A. Asset Ownership

Truck ownership implies both residual control over how the truck is used and a residual claim on the truck's value. Residual control rights with respect to how trucks are used exist because, as discussed above, it is rarely optimal to specify exactly how trucks should be used more than a few hours in advance. An important convention in the trucking industry is that truck owners have the ultimate say with respect to how trucks are used.⁸ A common expression of this convention is that there is "no forced dispatch" for owner-operators. Taken literally, "no forced dispatch" refers to trucks rather than drivers, since carriers cannot literally force any driver to accept dispatchers' assignment—company drivers can quit as well. But unlike company drivers, owner-operators can take their truck and use it as they wish.⁹

Note that, in principle, the party with residual control rights need not be the residual claimant on its value. But if the party that held residual control rights over the truck did not also have residual claimancy, this party would not have incentives to utilize these rights in a way that preserves trucks' value. A carrier who held residual control rights, but not residual claims, on a truck would have strong incentives to use the truck for hauls that are hard on the truck's engine, for example. This is an important reason for the "no forced dispatch" convention, and is why arrangements whereby owner-operators sign away their right to refuse backhauls—arrangements that would give carriers residual control rights but not residual claimancy—do not appear in this industry.¹⁰

8. This convention, which links ownership of a physical asset to residual decision rights with respect to the asset's use, parallels Grossman and Hart's [1986] definition of asset ownership.

9. A company driver who quits far from his base has to find his own way home. An owner-operator can much more credibly threaten to walk away (actually drive away) from the bargaining. In interviews with drivers and dispatchers, we learned that whether the driver owns both the tractor and the trailer, or only the tractor, matters little to bargaining costs.

10. As discussed at length in a previous version of the paper [Baker and Hubbard 2000], there exist long-term contracts between owner-operators and carriers whose provisions would appear to restrict how owner-operators can use their trucks. But the formal lease terms are misleading; they exist for regulatory reasons unrelated to our analysis. Carriers do not deny owner-operators access to their trucks, even when drivers unilaterally terminate leases prematurely. The

The incentive benefits of having drivers own their truck are clear: if the driver owns the truck, he has incentives (through his residual claim on the truck's value) to make optimal trade-offs with respect to how the truck is driven. If the driver does not own the truck, then absent some contracting technology, he will make decisions about how to drive that are likely to be inefficient.

The incentives induced by driver ownership with respect to negotiation over the backhaul are more complex. When the carrier owns the truck, then if a profitable backhaul is found, the carrier can mandate that the truck be used for that backhaul. However, if the driver owns the truck, the carrier cannot do so. This can lead to at least three forms of inefficiencies, which we collectively label "bargaining costs." First, the carrier may be less likely to try to arrange a highly time-critical pickup (even though it might be highly profitable). Second, the driver's ability to control how his truck is used may encourage him to engage in costly search for alternative hauls, in order to strengthen his bargaining position with the dispatcher. Finally, even if neither party engages in these sorts of *ex ante* inefficient actions to maintain or improve their bargaining positions, they may engage in costly *ex post* haggling that wastes time and effort. The likelihood of all of these types of behavior increases when the driver owns the truck, and can threaten not to carry a particular backhaul lined up by the dispatcher.

This depiction of the costs and benefits of owner-operators is consistent with characterizations in the literature. Dispatchers often claim that they have more difficulty inducing owner-operators to accept hauls than company drivers, and that this makes it more difficult to plan schedules. In his book Maister [1980] observes that: "Owner-operators' refusal of loads is, by a large margin, the most commonly reported disadvantage in utilizing owner-operators rather than a company-owned fleet. Refusals mean that the carrier can plan less well, and, as we have seen, operational planning is a difficult task for any irregular-route carrier because of the 'real-time' nature of planning required of such carriers" [p. 97].

control right provisions in owner-operator leases are, for our intents and purposes, a legal fiction.

II.B. Empirical Propositions

We propose that the optimal ownership of trucks is influenced by the relative costs of these two organizational structures: the agency costs that can arise when the driver does not own the truck, and the bargaining costs (both *ex ante* and *ex post*) that can result when the driver does own the truck. We develop two cross-sectional propositions about asset ownership in trucking based on this simple trade-off. These propositions require that we are able to differentiate hauls by the magnitude of these incentive problems and these bargaining costs.

As discussed above, longer hauls are likely to induce greater driving problems. This is because on longer hauls, drivers have more scope to drive fast for some period of time, and then use the time saved to engage in other types of activities. Thus, the inefficient driving problem should be greater for long hauls. In contrast, bargaining costs should not systematically differ with distance when looking across hauls that take trucks outside of their local area: situations where backhauls are valuable. The inefficiencies associated with driver ownership are primarily a function of whether backhauls exist and would be profitable, and the bargaining environment varies little depending on whether trucks are, say, 150 or 300 miles from their base. We therefore propose (**P1**) that, among hauls that take trucks outside of their local area, longer hauls are more likely to be completed by owner-operators.

In addition, certain types of hauls use trailers that are more likely to be used for backhauls than others. Hauls carried in general purpose trailers such as nonrefrigerated vans are more likely to suffer from costly backhaul negotiations than hauls that use trailers for which there tends to be little backhaul demand. Since hauls without backhaul problems will not suffer from bargaining costs, they are more likely to be carried by owner-operators. We divide trailers into two groups—those for which aggregate demands are likely to be bidirectional and those for which they are likely to be unidirectional—and propose (**P2**) that (holding the haul length constant) hauls that use unidirectional trailers are more likely to be carried by owner-operators.¹¹

Evidence with respect to P2 is important because it sheds

11. We also break down the “backhaul” group slightly more finely, and show that nonrefrigerated vans (which are the most likely to have bidirectional demand) are still more likely to be company owned.

light on the costs associated with driver ownership. In particular, finding evidence consistent with P2 is inconsistent with the hypothesis that optimal asset ownership reflects a simple trade-off between incentives and risk-sharing, unless any risk-related costs associated with driver ownership were systematically lower for unidirectional than bidirectional trailers. Evidence consistent with P2 would also be inconsistent with a simplistic prediction that asset specialization should lead to greater integration. Unidirectional trailers tend to be those that are specialized to particular types of cargo, and P2 states that hauls that use such trailers should be relatively more likely to be carried by owner-operators.¹²

Our main empirical proposition relates, however, to how ownership patterns change with the introduction of OBCs, and exploits the time dimension of our data. If OBCs reduce the inefficient driving problem by making good driving contractible, they should affect the trade-off between driver and company ownership of the truck for a particular haul. We therefore propose (**P3**) that driver ownership should decline with OBC adoption. The reason for this is simple: OBCs eliminate an important advantage of owner-operators over company drivers. They reduce the agency costs associated with company drivers, but do not change the bargaining costs associated with owner-operators.

Our analysis suggests an additional proposition with respect to the relationship between OBC adoption and ownership. If the agency costs associated with inefficient driving are greater for longer hauls, and OBCs eliminate such costs by making driving contractible, then OBCs reduce the agency costs associated with company drivers more for longer hauls. As a consequence, all else equal, the likelihood that a company driver is used for a haul should increase more for longer hauls. We therefore propose (**P4**) that the relationship between OBC adoption and ownership change should be stronger for long hauls than for short hauls.

We also analyze whether the relationship between OBC adoption and ownership change varies between unidirectional and bidirectional hauls. This provides some additional evidence regarding whether bargaining costs associated with backhauls influence asset ownership. Suppose that bargaining costs do not differ systematically between these classes of hauls. One would

12. It would not necessarily be inconsistent with a prediction that carefully distinguishes between assets specialized to *users* and assets specialized to *uses*.

then not expect the relationship between OBC adoption and ownership to differ. Finding that this relationship varies between these classes of hauls therefore provides additional evidence that the contracting environment varies in this dimension, and would be consistent with the proposition that bargaining problems associated with the backhaul influence ownership patterns.

Finally, we will examine the hypothesis that OBC adoption should lead company drivers to drive better by analyzing how fuel economy, which is correlated with how drivers drive, varies with whether trucks have OBCs installed. We propose (P5) that this relationship should be stronger when comparing across company drivers than across owner-operators. We discuss this test and its empirical implementation in more detail in Section V, after we have presented the data and our results with respect to ownership and adoption.

III. DATA AND CROSS-SECTIONAL PATTERNS

The data are from the 1987 and 1992 Truck Inventory and Use Surveys (TIUS) (See Bureau of the Census [1989, 1995] and Hubbard [2000]). The TIUS is a survey of the nation's trucking fleet that the Census takes every five years. The Census sends forms to the owners of a random sample of trucks, and asks owners questions about the characteristics and use of their truck. The characteristics include trucks' physical characteristics such as make and model year, as well as whether certain aftermarket equipment is installed—including whether and what class of OBCs are installed. Questions about use yield information on the state in which the truck was based, how far from home it was generally operated, the class of trailer to which it was generally attached, and the class of products it primarily hauled.¹³ For trucks that operate outside of their local area, the class of products trucks primarily haul reflects what they carry on "fronthauls," since the cargo individual trucks carry on fronthauls is more consistent than the cargo they carry on "backhauls." The survey also asks whether the truck was driven by an owner-operator or a company driver. This paper uses observations of diesel-powered truck-tractors, the front halves of tractor-trailer

13. Trucks are not always attached to the same trailer. However, it turns out that trailers are detached from tractors less than one might expect, and most tractors end up pulling one type of trailer most of the time.

combinations. We eliminate observations of those that haul goods off-road, haul trash, are driven for less than 500 miles during the year, or have missing values for relevant variables. This leaves 19,308 observations for 1987 and 35,204 for 1992. The sample is larger for 1992 because the Census surveyed more trucks.

These data are well-suited to studies of organizational form, since theories of organizational form commonly take the transaction as the unit of analysis. Both the analytic framework presented above and the empirical framework we present below take this as a starting point. Because individual trucks tend to be used for similar types of hauls from period to period, observing ownership and OBC use at the truck level is much like observing ownership and OBC use for a sequence of similar transactions. We will therefore think of these observations of trucks as observations of hauls in our analysis, and interpret our empirical results using this perspective.

Table I reports owner-operator shares for different haul categories in 1987 and 1992. In 1987, 14.6 percent of tractor-trailers

TABLE I
SHARE OF TRUCKS DRIVEN BY OWNER-OPERATORS, BY TRAILER AND DISTANCE

	All distances	How far from its base was the truck generally operated?		
		<50 miles	50-200 miles	200+ miles
<i>1987 owner-operator shares</i>				
All trailers	0.146	0.084	0.118	0.211
“No backhaul” trailers	0.181	0.136	0.198	0.286
“Backhaul” trailers	0.140	0.070	0.101	0.208
Platform, refrigerated vans, tank trucks	0.179	0.081	0.121	0.269
Nonrefrigerated vans	0.125	0.083	0.095	0.158
<i>1992 owner-operator shares</i>				
All trailers	0.101	0.045	0.091	0.139
“No backhaul” trailers	0.124	0.062	0.155	0.202
“Backhaul” trailers	0.097	0.039	0.075	0.136
Platform, refrigerated vans, tank trucks	0.114	0.038	0.076	0.163
Nonrefrigerated vans	0.091	0.046	0.084	0.110

Source: 1987 and 1992 Truck Inventory and Use Surveys, authors' calculations. All calculations use Census-provided sampling weights. N = 19,308 for 1987; N = 35,204 for 1992. “No backhaul” trailers include dump, grain, livestock, and logging trailers. “Backhaul” trailers include all others.

were driven by their owners.¹⁴ The share is higher for trucks used for longer hauls; over one-fifth of trucks primarily used for long hauls were owner-operated. Looking across distances for which backhauls are valuable, the owner-operator share is higher for trucks that generally operate more than 200 miles from their base than those that generally operate between 50 and 200 miles from their base. This is consistent with (P1). We split the sample according to whether trucks were generally attached to trailers where demands are usually unidirectional. "No backhaul" trailers include dump, grain body, livestock, and logging trailers. "Backhaul" trailers include all other trailer types; vans, refrigerated vans, platforms, and tank trucks make up most of this category (and are the most prevalent trailers in general). About 18 percent of trucks commonly attached to "no backhaul" trailers were owner-operated, compared with 14 percent of trucks commonly attached to "backhaul" trailers. Moving to the right, a greater share of trucks attached to "no backhaul" trailers are owner-operated in each distance category. We further analyze the importance of the backhaul negotiation problem by examining two subcategories of the "backhaul" trailers. It may be the case that hauls using nonrefrigerated vans (the most general-purpose trailer) are more likely to be bidirectional than platforms, refrigerated vans, and tank trucks. While we are not confident that this distinction is as sharp as that between the "no backhaul" and "backhaul" trailers, the comparison provides similar results. The owner-operator share is smaller for vans than these other trailer types overall, and the difference is largest for long hauls. Our evidence is thus consistent with (P1) and (P2). Owner-operators are used more for longer (nonlocal) hauls, and for hauls that use trailers for which demands tend to be unidirectional rather than bidirectional.

The bottom panel reports analogous figures for 1992. The owner-operator share fell by about 30 percent between 1987 and 1992, from 14.6 percent to 10.1 percent, and declined in each of the distance-trailer cells reported in this table.

Table II reports OBC adoption rates, by organizational form and distance, for 1992. OBC adoption is negligible during 1987, and is treated as zero for that year throughout the paper. Adop-

14. Note that the sample contains trucks within both private and for-hire fleets. About half of the nation's truck-tractors operate within private fleets; all trucks within private fleets are driven by company drivers. Also, the 1992 Survey contains more detailed distance categories than the 1987 Survey. We convert the five 1992 categories to the three 1987 ones when comparing the two years.

TABLE II
1992 ON-BOARD COMPUTER ADOPTION RATES

	How far from its base was the truck generally operated?				
	<50 miles	50–100 miles	100–200 miles	200–500 miles	500+ miles
OBC					
Owner-operator	0.037	0.031	0.040	0.070	0.098
Company driver	0.071	0.126	0.211	0.274	0.348
Trip recorder					
Owner-operator	0.017	0.012	0.009	0.023	0.024
Company driver	0.043	0.078	0.127	0.120	0.084
EVMS					
Owner-operator	0.020	0.020	0.031	0.048	0.074
Company driver	0.028	0.049	0.084	0.154	0.265

Source: 1987 and 1992 Truck Inventory and Use Surveys, authors' calculations. All calculations use Census-provided sampling weights. N = 19,308 for 1987; N = 35,204 for 1992.

tion is higher for trucks driven by company drivers than owner-operators (for whom OBCs are useful only for improving maintenance and coordination), and increases with how far trucks operate from home. Almost 35 percent of trucks used for hauls of 500 or more miles and operated by company drivers had either trip recorders or EVMS installed.

Tables I and II thus indicate that OBC adoption coincided with ownership changes in the aggregate. Hauls in general moved from owner-operators to company drivers at the same time OBCs were beginning to diffuse. Ownership changes and OBC adoption were both greatest for long hauls. These broad trends set the stage for more detailed analysis that investigates whether ownership changes and OBC adoption are related, and thus provides evidence with respect to (P3) and (P4). The rest of the section develops the empirical framework that supports our analysis.

III.A. Cross-Sectional Relationships, Individual Data

Our analytic framework, in keeping with the organizational economics literature, assumes that efficient organizational forms are chosen. We therefore begin by specifying total surplus for a particular haul under the two organizational alternatives. Let S_{iot} represent total surplus of haul i at time t , if a driver owns the truck, and S_{ict} represent total surplus of haul i , if a carrier owns

the truck. Note that when haul i is nonlocal, it has a fronthaul and backhaul component. The characteristics of the former are known at the time organizational form is chosen, but not the latter: for example, the backhaul will necessarily use the same trailer as the fronthaul, but need not involve the same product. Specify these as

$$(1) \quad \begin{aligned} S_{iot} &= X_{it}\beta_o + Z_i\gamma_{ot} + \delta_o d_{it} + \varepsilon_{iot} \\ S_{ict} &= X_{it}\beta_c + Z_i\gamma_{ct} + \delta_c d_{it} + \varepsilon_{ict}, \end{aligned}$$

where X_{it} and Z_i are vectors depicting time-varying and time-invariant haul characteristics and d_{it} is a dummy variable that equals 1 if OBCs are used for the haul. ε_{iot} and ε_{ict} capture how haul characteristics not observed by the econometrician but observed by carriers and drivers affect surplus when using owner-operators and company drivers, respectively.

Assuming that ownership choices are efficient, company drivers will be chosen if and only if $S_{ict} > S_{iot}$. Assuming that ε_{iot} and ε_{ict} are i.i.d. type I extreme value, the probability the carrier owns the truck used for haul i , conditional on X_{it} , is

$$(2) \quad \begin{aligned} P_{it} &= \frac{e^{X_{it}(\beta_c - \beta_o) + Z_i(\gamma_{ct} - \gamma_{ot}) + d_{it}(\delta_c - \delta_o)}}{1 + e^{X_{it}(\beta_c - \beta_o) + Z_i(\gamma_{ct} - \gamma_{ot}) + d_{it}(\delta_c - \delta_o)}} \\ &= \frac{e^{X_{it}\beta + Z_i\gamma_t + d_{it}\delta}}{1 + e^{X_{it}\beta + Z_i\gamma_t + d_{it}\delta}} = \Lambda(X_{it}\beta + Z_i\gamma_t + d_{it}\delta), \end{aligned}$$

where $\Lambda(a) = \exp(a)/(1 + \exp(a))$.

The top panel of Table III contains results from estimating this model using simple logits on the 1992 data. We present estimates for all distances, then for short, medium, and long hauls separately. The dependent variable is a dummy variable that equals one if the truck was driven by a company driver and zero if an owner-operator. The independent variables are a dummy variable that equals one if the truck has an OBC installed and zero otherwise, a vector of dummies that indicate how far from home the truck generally operated, and $\ln(\text{trailer density})$. The latter is the number of trucks based in the same state that are attached to the same trailer type, normalized by the developed land in the state. This is a measure of local fronthaul market thickness; it is high for logging trailers in Oregon and low in Kansas, for example. (See Hubbard [2001] for an extensive discussion.)

In the first column the coefficient on the OBC dummy is positive

TABLE III
TRUCK OWNERSHIP AND OBC ADOPTION—LEVELS ESTIMATES, 1992

	All distances	<50 miles	50–200 miles	200+ miles
Individual trucks				
OBC	1.587 (0.071)	0.728 (0.292)	1.717 (0.172)	1.584 (0.081)
N	33283	7998	11429	13856
Cohorts, observed ownership and adoption shares				
OBC	1.560 (0.189)	–2.701 (2.018)	1.204 (0.389)	1.698 (0.228)
N	426	38	123	265
Cohorts, Bayesian estimates of ownership and adoption shares				
OBC	0.681 (0.070)	– 3.552 (0.367)	1.381 (0.153)	1.447 (0.106)
N	3676	1049	1332	1295

This table presents the coefficients from three logit specifications. The top panel uses observations of individual trucks; the dependent variable is a dummy variable that equals one if the truck is driven by a company driver and zero otherwise. The coefficient reported above is that on a dummy that equals one if the truck has an OBC installed and zero otherwise. Observations are weighted using Census sampling weights.

The middle and bottom panels use truck cohorts, where cohorts are defined by state-product-trailer-distance combinations. In these multivariate regressions, the dependent variables are ln(company driver share/owner-operator share) in 1987 and 1992. The coefficient reported above is that on the share of trucks within the cohort with OBCs installed. In the middle panel, the ownership and OBC adoption shares are the actual shares; we include only cohorts with nonzero shares of company drivers and owner-operators in both years. In the bottom panel the ownership and OBC adoption shares are constructed using the “Bayesian” formula as described in the text. Cohort observations are weighted. The weight of cohort r equals $[n(r, 1987) * k(r, 1987) + n(r, 1992) * k(r, 1992)] / 2$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

All specifications include dummy variables that indicate whether the truck was operated <50 miles, 50–200 miles, or >200 miles from its base and ln(trailer density) as controls.

Boldface indicates rejection of the null $\beta = 0$ using a two sided t -test of size 0.05.

and significant: hauls that use trucks with OBCs tend to be completed by company drivers more than hauls that do not use trucks with OBCs. The magnitude of the point estimate implies that, holding the controls at their means, the probability that a haul is completed by a company driver is about 11 percentage points higher—about 0.96 rather than about 0.85—if the truck has an OBC than if it does not.¹⁵ The coefficient on the OBC dummy is positive and significant for short, medium, and long hauls, but the correlation between OBC use and truck ownership is weakest for short hauls. The cross-sectional evidence thus is consistent with P3 and P4, but is also consistent with hypotheses where adoption need not lead to

15. $\Lambda(X\beta + 1.587) - \Lambda(X\beta) = 0.96 - 0.85 = 0.11$, where X is the mean value of the controls and β are the coefficient estimates associated with these controls.

changes in ownership. For example, one would expect OBC use and carrier ownership to be correlated in the cross section if the returns to monitoring are greater when trucks are not driver-owned. The time dimension of our data provides a significant advantage in confronting this issue.

III.B. Cohorts

The data are multiple cross sections rather than panel data; we do not observe exactly the same truck from period to period. To exploit the data's time dimension, we construct "cohorts" of individual observations based on state-product-trailer-distance combinations that are observed in both of our sample periods. An example is "long-distance hauls of food in refrigerated vans by trucks based in California." There are 131,274 possible combinations (51 states*33 products*26 trailers*3 distances); only about 3 percent of these have positive observations in both years, mainly because it is rare for a product class to be hauled in more than a few trailer types (transportation equipment is never hauled in tank trucks, for example). We base cohorts on state-product-trailer-distance combinations because it aggregates the data up to narrowly defined haul segments. Defining cohorts narrowly minimizes within-segment heterogeneity in haul characteristics which would otherwise tend to bias our estimates, as explained below. Our empirical work will relate within-segment changes in OBC use to changes in driver ownership of trucks: does the owner-operator share decrease the most in segments where OBC adoption is greatest, and is there evidence that adoption causes ownership to change? This will provide evidence regarding whether and how changes in contractibility relate to changes in the comparative advantage of using company drivers relative to owner-operators.

The first column of Table IV presents summary statistics for the 3676 cohorts with at least one observation in both years. On average, segments are based on relatively few observations; this is a drawback of defining cohorts narrowly. Because of this, many of our segments, particularly the very smallest ones, have either 0 percent or 100 percent company drivers in one or both years; nearly half have 100 percent company drivers in both. This is not surprising, given that most hauls are completed by company drivers, especially short hauls. But it creates some empirical problems because our empirical specifications below are logit-based regressions that use log-odds ratios of the ownership

TABLE IV
COHORT SUMMARY STATISTICS

	All cohorts	Cohorts with positive owner-operator and company driver shares in both years
Cohorts	3676	426
Observations/cohort, 1987	4.13	10.61
Observations/cohort, 1992	6.42	17.80
Owner-operator share, 1987	0.14	0.27
Owner-operator share, 1992	0.10	0.18
Change in owner-operator share	-0.04	-0.09
OBC adoption, 1992	0.19	0.24
Trip recorder adoption, 1992	0.09	0.10
EVMS adoption, 1992	0.10	0.14

Averages are computed using weights. The weight of cohort r in time t equals $n(r, t) * k(r, t)$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

shares as the dependent variable. While specifying the model in a regression framework allows us to difference out cohort-specific fixed effects, the log-odds ratios are only well-defined when cohorts have nonzero company driver and owner-operator shares in both years.

We have addressed the problem of 0 percent or 100 percent owner-operator shares in several ways. One is simply to use only cohorts with nonzero company driver and owner-operator shares in both years. This allows our empirical specifications to be connected to the framework and estimates discussed above, but leads the analysis to be based on a relatively small part of our data; only 426 of the 3676 cohorts satisfy this criterion. As reported in Table IV, these 426 cohorts tend to have many more observations per cohort than those with a zero company driver or owner-operator share in at least one of the years; they are only 12 percent of the cohorts, but contain over 30 percent of the observations in each year. The average owner-operator share tends to be larger for these cohorts, reflecting that populations in which owner-operators are rare are more likely to have zero owner-operator shares than those where they are common. In both columns the owner-operator share declined by about 30 percent. Below, we will show that the cross-sectional relationships between OBC adoption and ownership for this subsample are also similar to those in the broader population. Combined, this pro-

vides evidence that the relationship between ownership and OBC use within this subsample resembles that in the broader population, and provides some assurance that estimates based on this subsample do not misrepresent relationships between OBC adoption and the comparative advantage of driver ownership in the population as a whole. This approach provides our main empirical results.

We have also estimated linear probability specifications of the model. This is a potentially attractive alternative because the dependent variable is well-defined for both the “zero” and the “nonzero” cohorts. This approach has several drawbacks, however. One is related to a general problem with linear probability models: it treats changes in the owner-operator share from 0.05 to 0.10 the same as those from 0.50 to 0.55, even though the former may indicate a greater underlying change in the comparative advantage of driver ownership. Since many of our observations have owner-operator shares that are less than 0.2, this affects our results more than it would if the owner-operator and company driver shares were more equal. Another, possibly more important, drawback arises in our first-difference specifications, and it arises precisely from using observations where the share of owner-operators is zero or one in both years. The problem is that changes in the owner-operator share do not fully reflect the underlying change in the comparative advantage of owner-operators for these observations.¹⁶ This problem is particularly relevant for us because nearly half of the 3676 cohorts with at least one observation in both years have no owner-operators in either year. If OBC adoption increased the comparative advantage of company drivers relative to owner-operators within these cohorts, first difference estimates would not pick this up (as it is impossible for the owner-operator share to decrease further), and this effect would bias our estimates toward zero. Other truncation-related problems arise for cohorts that have no owner-operators

16. The econometric issues that arise here are similar to those in Chamberlain's [1980] analysis of the fixed effect logit model. The conditional maximum likelihood estimator he proposes does not apply directly to our case, where the observations are grouped rather than individual data. Because the dependent variable can take any value between zero and one, the sets upon which one would condition would be very small. To our knowledge, the econometrics literature has not addressed the issue of first-difference estimation of qualitative response models with grouped data. See Maddala [1987] for a discussion of limited dependent variable models using individual-level panel data.

in one of the two years, or that have all owner-operators in either or both years.

Our third approach attempts to use the information in the cohorts with 0 percent or 100 percent owner-operators while retaining the logit-based specification. This approach treats the observed owner-operator shares as informative, but not fully informative, of the true shares across the population of trucks within each cohort. Consistent with the theoretical specification outlined above, we assume that cohorts with all or no owner-operators are observed not because one of the organizational forms has an insurmountable comparative advantage, but because we observe a finite, often small, number of observations within each cohort. To implement this approach, we estimate the true owner-operator share within each cohort with a weighted average that puts some weight on the mean share across all cohorts in the same distance category; the weight on the observed shares increases with the number of observations in the specific cohort. This results in estimates of the owner-operator shares that are bounded away from zero and one, and mitigates the truncation-related problems discussed above. For consistency, we create estimates of the true OBC adoption shares using an analogous procedure. These estimates of the shares can be thought of as Bayesian, using our data to update priors about the true shares of driver ownership and OBC use within each cohort. We use the resulting posteriors, which can never be zero or one, in logit-based regressions analogous to those discussed above. The formulas for these “Bayesian” estimates are given in Appendix 1.¹⁷

In the results section below, we present two sets of estimates that we will use for each of our tests. One set uses the observed ownership and adoption shares from the 426 cohorts described above. The other set uses the Bayesian estimates of the ownership and adoption shares rather than the observed shares.¹⁸ All calculations and estimates involving cohorts use weights that

17. It is natural to base the initial priors for these Bayesian estimates on distance-specific means because ownership and adoption vary systematically with distance, and because the cells within each of the distance categories are made up of many truck-level observations. Because of the latter, there is little sampling error associated with these distance-specific means. This would not be the case if we were to base these means on more narrowly defined categories.

18. Note that while our dependent and independent variables are constructed using a Bayesian procedure, the regression coefficients themselves are not Bayesian estimates.

reflect differences in the number of observations within cohorts and in the rate in which the Census sampled trucks.¹⁹ Thus, while the latter set of estimates incorporate information from many small cohorts that are omitted in estimates that use the observed shares and collectively make up much of the industry, most of the cohorts that are added receive little weight individually.

In addition, we present two sets of results from linear probability specifications in Appendix 2, one for the 426 “nonzero” cohorts, and one for all cohorts. As would be expected, the results are similar to our logit-based specifications when using only the nonzero cohorts, but the estimates are small and not statistically different from zero when including all of the “zero” cohorts.

III.C. Aggregation-Related Biases

There is a potential aggregation-related bias introduced by using cohorts as the unit of analysis. This bias works against finding the relationships between OBC adoption and organizational change that we predict. The issue arises if, as in our framework, OBCs and driver ownership are incentive substitutes and if hauls differ within cohorts. In cohorts where good driving is particularly important, either OBCs or driver ownership will be used to provide drivers’ incentives. If hauls within these cohorts are identical, the same solution to the driving incentive problem should be chosen for each, but if there is within-cohort heterogeneity, OBCs will be used for some hauls and driver ownership will be used for others. In cohorts where good driving is unimportant, one should observe neither OBCs nor owner-operators. One could therefore observe a negative correlation at the cohort level between OBCs and carrier ownership, even if the haul-level correlation is positive. Thus, aggregating the observations into cohorts biases us against finding a positive correlation between OBC adoption and company ownership.²⁰ While we define cohorts narrowly to make within-cohort differences as small as possible, this does not necessarily eliminate this problem.

We examine this problem’s empirical relevance below by

19. The formula is $(n_{r,1987} * k_{r,1987} + n_{r,1992} * k_{r,1992})/2$, where $n_{r,t}$ is the number of observations in cohort r and $k_{r,t}$ is the average Census weighting factor in cohort r in year t . Census sampling rates, and thus $k_{r,t}$, differ primarily across states, not across trucks within states during a particular year. The results in Section V are robust to variations in weighting.

20. See Deaton [1985] for a general depiction of this problem.

comparing estimates of cross-sectional relationships between OBC use and ownership from the individual and cohort data. Areas where the cohort-based cross-sectional estimates differ from the individual-based ones indicate situations where the bias described above is likely to affect our first-difference estimates, which necessarily rely only on the cohort data.

III.D. Cross-Sectional Relationships, Cohort Data

The cohort analog to equation (2) is

$$(3) \quad s_{rt} = \Lambda(X_{rt}\beta + Z_r\gamma_t + d_{rt}\delta),$$

where s_{rt} is the share of hauls in cohort r at time t for which company drivers are used, X_{rt} is a vector of average haul characteristics for cohort r in time t , Z_r represents time-invariant haul characteristics (such as distance), and d_{rt} is the OBC adoption rate within cohort r . One can estimate the parameters of this equation by estimating the linear regression,

$$(4) \quad \ln(s_{rt}/(1 - s_{rt})) = X_{rt}\beta + Z_r\gamma_t + d_{rt}\delta + \varphi_{rt}.$$

Because we have two years of data, we estimate the system of equations:

$$(5) \quad \begin{aligned} \ln(s_{r,1987}/(1 - s_{r,1987})) &= X_{r,1987}\beta + Z_r\gamma_{1987} \\ &+ d_{r,1987}\delta + \psi_r + \eta_{r,1987} \\ \ln(s_{r,1992}/(1 - s_{r,1992})) &= X_{r,1992}\beta + Z_r\gamma_{1992} \\ &+ d_{r,1992}\delta + \psi_r + \eta_{r,1992}. \end{aligned}$$

For purposes of the discussion below, we have decomposed the error term into time-varying (η_{rt}) and time-invariant (ψ_r) components.

Returning to Table III, the bottom two panels contain the “levels” estimates of δ from multivariate regressions using the cohort data.²¹ Note that $d_{r,1987} = 0$, since OBCs were not installed on trucks at this time. δ thus reflects cross-sectional relationships between OBC use and truck ownership during 1992. The dependent and independent variables are analogous to

21. We use multivariate regressions using both years of data here so that the specifications and sample are comparable to those in the first difference results reported below. We have also run univariate regressions that use only the 1992 data, thus estimating only the second equation in (5); the results are similar.

those in the top panel. In the middle panel we use observed ownership shares in calculating $\ln(s_{rt}/(1 - s_{rt}))$ and d_{rt} , and therefore can use only the 426 cohorts that have nonzero owner-operator and company driver shares in both years. In the bottom panel we use the Bayesian estimates of the ownership and adoption shares. Since none of these estimates generate zero owner-operator or company driver shares, this panel uses all 3676 cohorts with observations in both years. Note that the estimates in the middle and bottom panels are similar, indicating that our Bayesian estimates do not greatly distort the cross-sectional relationship between OBC use and ownership.

The estimates for medium and long haul cohorts are similar to those when we use the individual-level data. However, the estimates from the short haul cohorts are not: the coefficient on OBC is strongly negative and, in the bottom panel, statistically significant. Combined, the results suggest that the aggregation-related bias described above does not much affect the medium or long haul estimates, but strongly affects the short haul estimates. One explanation for this is that there is little within-cohort heterogeneity in the degree of the incentive problem for medium and long hauls, but significant within-cohort heterogeneity for short hauls. This would be the case if incentive problems were driven more by idiosyncratic factors for short hauls than medium or long hauls.

Below we will find that this negative relationship for short hauls appears in first difference estimates as well, but we will not focus on this result because the cross-sectional evidence strongly suggests that it reflects a negative bias in the estimates.

IV. OBC ADOPTION AND OWNERSHIP CHANGES

This section contains the main empirical evidence in this paper, which concerns relationships between OBC adoption and ownership changes. Before discussing the results, we describe the conditions under which one can and cannot interpret our estimates as reflecting causal relationships.

The central issue regarding causality is that OBCs are not adopted at random. In the levels estimates in Table III, OBC use is econometrically endogenous if it is not independent of unobserved factors that affect ownership trade-offs; that is, if $E(d_{rt}|\psi_r + \eta_{rt}) \neq 0$. This would be the case if, for example, there are unobserved differences in market conditions across segments.

Suppose, for example, that the thickness of the backhaul market differs across segments in unobserved ways; in some cohorts a larger fraction of hauls are “backhaul” versus “no backhaul” hauls. This would lead to unobserved differences in the comparative advantage of company drivers. This may, in turn, affect the returns to OBC adoption within the segment, especially if motivating good driving is more valuable when trucks are hauling cargo than empty. OBC use and driver ownership of trucks would be negatively correlated even if OBCs did not directly affect ownership patterns.

Taking the difference between the equations in (5), and recalling that $d_{r,1987} = 0$, yields

$$\begin{aligned}
 (6) \quad & \ln(s_{r,1992}/(1 - s_{r,1992})) - \ln(s_{r,1987}/(1 - s_{r,1987})) \\
 & = (X_{r,1992} - X_{r,1987})\beta + Z_r(\gamma_{1992} - \gamma_{1987}) \\
 & \quad + d_{r,1992}\delta + (\eta_{r,1992} - \eta_{r,1987}).
 \end{aligned}$$

In first-difference estimates, OBC adoption is econometrically exogenous if $E(d_{r,1992}|\eta_{r,1992} - \eta_{r,1987}) = 0$; that is, if OBC adoption is independent of unobserved changes in organizational form. This condition is much weaker than the corresponding condition when estimating the model in levels because ψ_r , which represents unobserved time invariant factors that affect the comparative advantage of driver ownership, has been differenced out. The condition allows for unobserved differences in incentive problems in the cross section (driver ownership may create greater rent-seeking problems in some parts of the country than in others), but requires such differences to be constant over time. Relationships between OBC adoption and changes in driver ownership therefore depict causal relationships if, within market segments defined by state-distance-trailer-product combinations, incentive problems with drivers are stable over time. There is some reason to believe that such problems are stable: since the composition of demand evolves very slowly in this industry (e.g., shipping patterns in 1987 are similar to those in 1992), the characteristics of hauls in a segment are probably fairly similar from one period to the next.

While our main results will come from simple first-difference specifications, we will also present and discuss results from instrumental variables specifications. We do this to provide some additional evidence with respect to causality: while unobserved

factors affecting ownership probably vary more cross-sectionally than over time, they might not be completely stable. For example, suppose that some backhaul markets become unobservedly thicker, moving some segments to move from mostly “no backhaul” to mostly “backhaul.” This would lead the comparative advantage of using company drivers to increase, and may independently encourage the adoption of OBCs within these segments. If so, OBC adoption would be correlated with unobserved changes in the comparative advantage of driver ownership, and the simple first difference estimates need not represent how much OBC adoption led to changes in truck ownership.

In these additional specifications we use product class dummies as instruments for OBC adoption. Using product dummies, which in our data reflect fronthaul cargo, as instruments is attractive for two reasons. First, as described above, OBC adoption offers benefits other than improving contracts with drivers. These benefits, which include verifying trucks’ operation to third parties such as insurers and customers with lean inventories, vary systematically with the cargo. Hubbard [2000] tests this proposition empirically, and finds evidence in favor: for example, conditional on who owns the truck and haul length, OBCs are used more when trucks haul dangerous cargo such as petroleum or chemicals or haul products for which sales/inventory ratios are high. Thus, product characteristics are shifters of OBC adoption. Second, unobserved changes in the comparative advantage of using an owner-operator relative to a company driver should not vary across products, given a haul’s location, distance, and trailer requirements. To see this, consider two hauls with the same origin and destination that use the same trailer but transport different products. Bargaining costs should not differ between these two hauls, because they are identical from the perspective of the backhaul. These costs are manifested after the truck reaches its destination at a time when its trailer is empty. On the other side of the trade-off, the benefits of good driving may differ between these two hauls, for example if the cost of an accident varies with the product being hauled. But this difference should not change from year to year. As a consequence, absent OBCs, the ownership changes we examine in our first-difference specifications should not systematically differ across products. We assume this to be true in our instrumental variables specifications below.

IV.A. First-Difference Estimates

Table V presents results from first-difference estimates. The dependent variable is the change in the log-odds ratio from 1987 to 1992 above; the independent variable of interest is the change in OBC use.²² The left panel contains estimates using the observed ownership and adoption shares. In the first column the OBC coefficient is positive and significant: cohorts with high OBC adoption moved the most toward company drivers during this time, consistent with our main theoretical proposition P3. The second column includes the EVMS adoption share separately; thus, the coefficient on OBC picks up the organizational implications of OBCs' incentive-improving capabilities and that on EVMS picks up the effects of their coordination-improving capabilities. The point estimates are both positive, but neither are statistically significantly different from zero. The third and fourth columns are analogous to the first two, but allow the coefficients to differ depending on haul length. In both, the OBC*long coefficient is positive and significant, and statistically significantly larger than either the OBC*medium and OBC*short coefficients, although the latter may reflect the impact of aggregation-related biases on the OBC*short coefficient. None of the EVMS coefficients are statistically significantly different from zero. The point estimates in the first and third columns imply that an increase in the OBC adoption rate from 0 to 0.2 is associated with an 8 percent overall decline in the owner-operator share and a 15 percent decline within the long-haul segment.²³ The first column estimate implies that the 20 percentage point overall increase in OBC use between 1987 and 1992 was related to slightly more than one-fourth of the 30 percent decline in driver ownership of trucks during this time.

The evidence from this panel thus indicates that OBC adoption is correlated with movements toward company drivers, and

22. The specifications also include a constant, the change in trailer density, and distance dummies. We have also estimated specifications that include a full set of state dummies in the vector Z_{it} , thus accounting for possible changes in state-specific economic conditions that affect whether owner-operators or company drivers are used; none of our results change.

23. $\Lambda(X\beta + 0.2*0.532) - \Lambda(X\beta) = 0.011$, where X is the mean value of the controls and β are the coefficient estimates associated with these controls, which is about 8 percent of 0.146, the overall owner-operator share in 1987. The estimate reported for the long-haul segment is computed analogously, as are those reported below that use coefficients from other specifications.

TABLE V
TRUCK OWNERSHIP AND OBC ADOPTION—FIRST DIFFERENCES

Dependent variable: $\ln(1992 \text{ company driver share}/1992 \text{ owner-operator share}) - \ln(1987 \text{ company driver share}/1987 \text{ owner-operator share})$					
Observations are product-trailer-state-distance cohorts.					
	Observed ownership and adoption shares		Bayesian estimates of ownership and adoption shares		
OBC	0.532 (0.269)	0.238 (0.381)	0.759 (0.100)	0.379 (0.145)	
EVMS		0.550 (0.506)		0.678 (0.190)	
OBC*long		0.943 (0.306)	1.021 (0.470)		1.054 (0.119) 0.741 (0.194)
OBC*medium		-0.549 (0.532)	-0.865 (0.670)		0.617 (0.192) 0.504 (0.240)
OBC*short		-2.655 (2.136)	-4.466 (2.374)		-2.611 (0.466) -2.699 (0.558)
EVMS*long			-0.104 (0.567)		0.458 (0.224)
EVMS*medium			1.153 (1.476)		0.397 (0.503)
EVMS*short			10.780 (6.184)		0.298 (1.137)
N		426		3676	

Bayesian estimates use initial priors with distance-specific means for ownership and adoption shares. Cohort observations are weighted. The weight of cohort r equals $[n(r, 1987)k(r, 1987) + n(r, 1992)k(r, 1992)]/2$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

All specifications include dummy variables that indicate whether the truck was operated <50 miles, 50–200 miles, or >200 miles from its base and $\ln(1992 \text{ trailer density}) - \ln(1987 \text{ trailer density})$ as controls.

Boldface indicates rejection of the null $\beta = 0$ using a two-sided t -test of size 0.05.

that the relationship between adoption and organizational change was greater for long than medium hauls. Furthermore, the organizational change appears to be related to OBCs' incentive-improving capabilities. The evidence thus is consistent with P3 and P4 above.

The right panel repeats the analysis using the Bayesian estimates of the ownership and adoption shares. These estimates produce similar evidence: improvements in the contractibility of good driving are correlated with a decrease in driver ownership of trucks. In the first column the OBC coefficient is positive and significant. In the second the OBC and EVMS coefficients are now both positive and significant, as the standard errors are lower than in the left panel because of the larger sample

size.²⁴ This suggests that OBCs' incentive- and coordination-improving capabilities are both correlated with movements toward company drivers. As before, the OBC*long coefficient is positive and significant. Unlike in the left panel, the OBC*medium coefficient is as well: there is some evidence that OBC adoption is correlated with organizational change for medium hauls. The point estimates of the OBC*long coefficient are greater than those of the OBC*medium coefficient. The difference is statistically significant using a *t*-test of size 0.05 in the third column but not the fourth. The OBC*short coefficients are negative and significant, which likely reflects aggregation-related biases. The EVMS*distance interactions are all positive, with EVMS*long being statistically significant. Except for the medium haul interactions, the magnitudes of the point estimates are similar to those in the left panel. For example, the estimates in the first and third column indicate that increasing the OBC adoption rate from 0 to 0.2 is associated with a 13 percent decline in the overall owner-operator share and a 17 percent decline within the long haul segment.

In sum, both panels provide strong evidence in favor of P3, our main proposition: cohorts where adoption was high also moved the most away from driver ownership of trucks. They also provide some evidence in favor of P4, as the relationship between adoption and organizational change is stronger for long-haul than medium-haul trucks. The evidence regarding P4 is not quite so strong, however, as the difference between the OBC*long and OBC*medium coefficients is not statistically significant in all of the specifications.

Table VI reports estimates from analogous specifications that contain interactions between OBC adoption and a dummy variable that equals one if the cohort is a "no backhaul" cohort: one with dump, grain body, livestock, or logging trailers. We estimate these using only the medium and long haul cohorts, both because we suspect the short haul estimates are negatively biased and because firms generally do not try to fill backhauls when trucks operate close to home. From the second column in the left panel, the OBC coefficient is positive and significant while the OBC*"no backhaul" interaction is negative and significant. In the fourth

24. The standard errors reported here do not account for the fact that our dependent and independent variables are estimates, and thus likely overstate the precision of our point estimates. Discussions of statistical significance should be taken in this light.

TABLE VI
TRUCK OWNERSHIP AND OBC ADOPTION—FIRST DIFFERENCES
Medium and long haul cohorts only

Dependent variable: $\ln(1992 \text{ company driver share}/1992 \text{ owner-operator share}) - \ln(1987 \text{ company driver share}/1987 \text{ owner-operator share})$ Observations are product-trailer-state-distance cohorts.								
	Observed ownership and adoption shares				Bayesian estimates of ownership and adoption shares			
OBC	0.576 (0.284)	0.683 (0.284)	0.333 (0.404)	0.424 (0.404)	0.930 (0.104)	0.971 (0.105)	0.604 (0.154)	0.664 (0.156)
EVMS			0.434 (0.513)	0.461 (0.514)			0.557 (0.195)	0.521 (0.196)
OBC*“no backhaul”		-6.993 (2.491)		-6.073 (3.177)		-2.128 (0.579)		-1.952 (1.043)
EVMS*“no backhaul”				-1.764 (3.808)				-0.121 (1.282)
N			388				2627	

Bayesian estimates use initial priors with distance-specific means for ownership and adoption shares. Cohort observations are weighted. The weight of cohort r equals $[n(r, 1987)k(r, 1987) + n(r, 1992)k(r, 1992)]/2$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

All specifications include dummy variables that indicate whether the truck was operated <50 miles, 50–200 miles, or >200 miles from its base and $\ln(1992 \text{ trailer density}) - \ln(1987 \text{ trailer density})$ as controls.

No backhaul includes dump, grain, livestock, and logging trailers. Specifications with “no backhaul” interactions include this variable as a control.

Boldface indicates rejection of the null $\beta = 0$ using a two-sided t -test of size 0.05.

column we allow the effect of adoption to differ for EVMS. The point estimate of the OBC*“no backhaul” coefficient is almost the same as in the second column. It is not statistically significantly different from zero using a two-sided t -test of size 0.05, but is when using one of size 0.10. The right panel repeats the exercise, using the Bayesian estimates of the adoption and ownership shares. The evidence is similar. Driver ownership declines more with OBC adoption when trucks use trailers for which demands tend to be bidirectional rather than unidirectional. This result suggests that the organizational impact of this change in the contracting environment differs across hauls, and is greater for hauls when driver ownership would invite backhaul-related bargaining problems. Combined with the cross-sectional differences in driver ownership, it provides additional evidence that bargaining problems associated with the backhaul affect whether drivers own trucks.

IV.B. Instrumental Variables Estimates

As discussed above, interpreting the first-difference estimates as causal relationships requires the assumption that OBC adoption is independent of unobserved changes in organizational form. In order to provide some evidence with respect to this interpretation, we rerun our analysis using instrumental variables. Following the logic described at the beginning of this section, we use a vector of nineteen product class dummies as instruments for OBC adoption.²⁵ Running a first-stage regression of the OBC adoption share on our controls and this vector, one can strongly reject the null hypothesis that the coefficients on the product class dummies are jointly zero, using a likelihood ratio test of size 0.05.²⁶ As expected, given the results in Hubbard [2000], OBC adoption varies significantly across product classes.

Table VII reports our instrumental variables estimates. In general, the patterns in the point estimates are similar to those in the simple first-difference specifications. In the first column of each panel, the OBC coefficient is positive and significant, and the coefficients are almost the same. These coefficients are greater than their counterparts in Table V, which indicates that the simple first-difference estimates might understate relationships between OBC adoption and ownership changes. Assuming that unobserved changes in the incentive problem with drivers are independent across products, the Table VII point estimates indicate that increasing adoption rates from 0 to 0.2 decreases the share of owner-operators by 2.3 percentage points. This suggests that absent OBC diffusion, the owner-operator share would have fallen by only about 15 percent between 1987 and 1992 rather than decreasing by 30 percent. The results thus imply that substantial share of the decline in driver ownership during this time is related to OBC-related changes in the contractibility of good driving.

The second column in each panel contains the distance interactions. As in Table V, the OBC*long coefficient is positive and significant in both panels; this provides evidence that OBC adop-

25. We report above that there are 33 product classes in our truck-level data, but some of these are uncommon or are not found in the cohort sample we use here. The instrument vector here includes dummy variables for the eighteen most common product classes (e.g., “food,” “lumber or wood products,” “petroleum products”), plus a dummy that indicates “other product class.”

26. The LR statistic for this test equals 160 when using the 426 nonzero cohorts, and is 312 when using all 3676 cohorts. In both cases, the statistic far exceeds the critical value for a chi-squared distribution with 19 degrees of freedom, which is approximately 30 for a size 0.05 test.

TABLE VII
TRUCK OWNERSHIP AND OBC ADOPTION—INSTRUMENTAL VARIABLES ESTIMATES

Dependent variable: $\ln(1992 \text{ company driver share}/1992 \text{ owner-operator share}) - \ln(1987 \text{ company driver share}/1987 \text{ owner-operator share})$						
Observations are product-trailer-state-distance cohorts.						
	Observed ownership and adoption shares			Bayesian estimates of ownership and adoption shares		
OBC	0.973 (0.426)		1.093 (0.601)	0.990 (0.376)		1.187 (0.493)
OBC*long		1.959 (0.440)			1.104 (0.585)	
OBC*medium		-0.358 (0.595)			1.192 (0.530)	
OBC*short		-6.019 (2.158)			-0.855 (1.280)	
OBC*"no backhaul"			7.615 (5.574)			-2.805 (1.879)
Sample	All	All	Medium, long	All	All	Medium, long
N	426	426	388	3676	3676	2627

Nineteen product class dummies used as instruments for OBC adoption.

Bayesian estimates use initial priors with distance-specific means for ownership and adoption shares.

Cohort observations are weighted. The weight of cohort r equals $[n(r, 1987)k(r, 1987) + n(r, 1992)k(r, 1992)]/2$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

All specifications include dummy variables that indicate whether the truck was operated <50 miles, 50–200 miles, or >200 miles from its base and $\ln(1992 \text{ trailer density}) - \ln(1987 \text{ trailer density})$ as controls.

Boldface indicates rejection of the null $\beta = 0$ using a two-sided t -test of size 0.05.

tion led to organizational changes for long hauls. The OBC*medium coefficient is negative and insignificant in the left panel and positive and significant in the right panel. It is significantly smaller than the OBC*long coefficient in the left panel, but not in the right one. Like the simple first-difference estimates, the instrumental variables estimates provide some evidence that relationships between OBC adoption and ownership changes are greater for longer hauls, but the evidence regarding P4 is not as strong as that regarding P3.

In general, the estimates in Table VII are consistent with the simple first-difference estimates. The point estimates are quite similar to those in the other tables. The strength of the evidence from the instrumental variables estimates varies with the standard errors. It is greatest for our main proposition P3, which concerns the general relationship between adoption and ownership, and provides some additional evidence that this relationship

is causal. It is weaker for P4, which concerns differences in OBCs' effect on ownership. We suspect that this partly reflects that we are running up against limits of the power of our data when we try to estimate interaction effects in first-difference specifications using instrumental variables.

V. ADOPTION AND DRIVING PATTERNS

In this section we present some evidence on whether OBC use affects how company drivers drive. Although our data do not contain any direct information on drivers' driving patterns, the TIUS does ask truck owners to report individual trucks' average fuel economy. Because individual trucks' fuel economy reflects many factors other than how drivers operate them—for example, how they are maintained and the terrain over which they are driven—fuel economy data are not useful for evaluating individual drivers' performance. But if OBC use causes company drivers to drive better on average, systematic differences in driving patterns might show up in fuel economy data from thousands of trucks.

We investigate this by presenting results from some OLS regressions using the truck-level data from 1992. The dependent variable is the truck's reported fuel economy, in miles per gallon. The main independent variables are interactions between dummies that indicate whether drivers own their trucks (one if driver ownership, zero otherwise) and whether the different classes of OBCs are installed (one if installed, zero otherwise). Coefficients on these variables indicate whether trucks with OBCs are more fuel-efficient than those without them. Relationships between OBC use and fuel economy may reflect things other than contracting improvements with drivers. As noted earlier, OBCs supply information that can help mechanics maintain trucks better. To distinguish between the effects of maintenance and incentive improvements, we compare the relationship between OBC use and fuel economy for company drivers and owner-operators. Assuming that the maintenance value of OBCs is the same for company drivers and owner-operators, finding that this relationship is stronger for company drivers than owner-operators is evidence of their incentive-improving effect, because it suggests that the average fuel economy benefits among company driver adopters are greater than that among owner-operators.

Relationships between OBC use and fuel efficiency may also reflect adoption patterns. In general, selection issues work

against finding relationships between OBC use and fuel economy: if OBCs tend to be adopted where agency costs are otherwise high, nonadopting company drivers probably drive better than the adopting ones would absent monitoring. The difference in the relationship between OBC use and fuel economy between company drivers and owner-operators would then understate OBCs' average incentive effect among company driver adopters. Finding that the fuel economy difference between trucks with and without OBCs is greater when comparing company drivers than owner-operators is thus evidence that OBCs induced fuel economy improvements.

Selection also might affect patterns across different types of OBCs. If EVMS are adopted more relative to trip recorders when monitoring's benefits are primarily coordination-related, one would expect the average fuel economy effect among EVMS adopters to be lower than among trip recorder adopters even if they can be used to improve drivers' incentives in the same way. We therefore allow the relationship between OBC use and fuel economy to differ for trip recorders and EVMS.

We include many additional variables as controls. One set contains an extensive set of truck characteristics that affect fuel economy: dummy variables that indicate the truck's make, model year, engine size, the number of driving axles, and whether it has aerodynamic features, as well as the log of the truck's odometer reading, which captures the effects of depreciation.²⁷ They also include variables that capture how the truck is used: how far from home it operates, whether it hauls single, double, or triple trailers, the average weight of the truck plus cargo, and whether it is attached to a refrigerated or specialized trailer. Finally, they include a set of dummy variables that indicate who maintains the truck: the driver, a garage, a trucking company, an equipment leasing firm, etc.

Table VIII reports results from four regressions. The owner-operator coefficient is negative and significant for short hauls, and statistically zero for medium and long hauls. There is no evidence that company drivers without OBCs drive less efficiently than owner-operators for medium and long hauls, and some evidence that fuel economy is higher for company drivers for

27. We do not include these variables in our analysis of ownership because we assume that, unlike on-board computers, these variables' effect on the cost of a haul is the same regardless of whether a company driver or owner-operator is used.

TABLE VIII
1992: FUEL ECONOMY, VEHICLE OWNERSHIP, AND DISTANCE

Dependent variable: miles per gallon				
	All distances	<50 miles	50–200 miles	200+ miles
Owner-operator	-0.042 (0.017)	-0.159 (0.064)	-0.008 (0.033)	-0.015 (0.019)
Trip recorder*owner-operator	0.063 (0.096)	-0.514 (0.330)	0.149 (0.265)	0.127 (0.091)
Trip recorder*company driver	0.186 (0.019)	-0.011 (0.067)	0.108 (0.033)	0.289 (0.021)
EVMS*owner-operator	0.184 (0.064)	0.299 (0.343)	0.346 (0.165)	0.146 (0.059)
EVMS*company driver	0.115 (0.019)	-0.060 (0.084)	0.165 (0.042)	0.126 (0.019)
R^2	0.210	0.154	0.241	0.252
N	35203	8002	11647	15552

This table reports coefficients from linear regressions of a truck's miles per gallon on ownership and OBC adoption variables. These regressions also include controls for distance from home, who maintains truck, refrigerated/specialized trailer, driving axles, vehicle make and model year, equipment dummies (such as for aerodynamic features), average weight, lifetime miles, and engine size.

Observations are weighted using Census sampling weights.

Boldface indicates rejection of the null $\beta = 0$ using a two-sided t -test of size 0.05.

short hauls. The trip recorder and EVMS interactions indicate that medium- and long-haul trucks with OBCs get better fuel economy than those without them. Among long-haul trucks, the point estimate on the trip recorder coefficient for company drivers is more than twice as high as that for owner-operators. The difference is statistically significant when using a t -test of size 0.10. (The owner-operator estimate is noisy because so few owner-operators drive trucks with trip recorders.) The point estimates indicate that, on the average across long-haul trucks for which they were adopted, trip recorders' incentive effect improved fuel economy by at least 0.16 miles per gallon, assuming that selection biases the parameter estimates downward.²⁸ Our estimates imply that this is about equal to aerodynamic hoods' effect on fuel economy. The EVMS coefficients tend to be lower than the trip recorder coefficients, as expected. There is no significant difference in the coefficients on the EVMS interactions.

In sum, these results on fuel economy are consistent with P5. The difference between the long-haul trip recorder coefficients in

28. For a truck that travels 100,000 miles/year, a 0.16 improvement in MPG translates to a \$620 savings per year, assuming that fuel costs \$1/gallon.

the fourth column of Table VIII provides some evidence of OBCs' incentive-improving effects: the difference in fuel economy between trucks with and without trip recorders is greater when comparing company drivers than owner-operators. This evidence does not appear when comparing the EVMS coefficients, possibly reflecting that they are adopted in many circumstances for their coordination- rather than their incentive-improving capabilities.

VI. CONCLUSION

This paper investigates factors affecting asset ownership in trucking; in particular, how the contracting environment affects whether drivers own the trucks they drive. Our evidence suggests that improved contracting (through the use of on-board computers) leads to more integrated asset ownership. Owner-operators are used for hauls where noncontractible decisions that affect trucks' value are important but are used less once these decisions become more contractible. The share of trucks that were owner-operated declined by about 30 percent between 1987 and 1992; our evidence suggests that the OBC-related contractual improvements accounted for somewhere between one-fourth and one-half of this decrease.²⁹ We also provide evidence on truck operating performance (in the form of miles per gallon outcomes) that is consistent with the ownership results. Differences in average fuel economy between long-haul trucks with trip recorders and without OBCs are greater among company-owned than driver-owned trucks, reflecting the improved incentives that the company drivers have after the adoption of OBCs.

The analysis in this paper may explain relationships between contractibility and firms' boundaries in other contexts, especially those in which the care of valuable assets is important. Presumably the prevalence of independent contractors in the construction trades is importantly influenced by the requirement to provide incentives for proper operation and maintenance of equipment. The results in this paper suggest that changes in monitoring technology could change the industry structure in this

29. Preliminary analysis using data from 1997 indicates a similar relationship between OBC adoption and ownership changes during 1992–1997. However, because the strongest relationships between OBC adoption and ownership changes are associated with trip recorder adoption, and trip recorder adoption was relatively unimportant during this period, the overall impact of OBCs on driver ownership was smaller than during 1987–1992.

sector. Such changes could similarly affect the professions. The prevalence of “owner-operators” in law and medicine may be driven to a large degree by the need to vest in professionals the value of their reputational assets. It appears that changes in the ability of insurance companies and HMOs to monitor the actions of physicians is causing higher rates of integration in medicine, leading doctors to become employees rather than independent contractors.

Innovations in information technology have led economists, technologists, and business people to theorize about how new informational capabilities will affect the boundaries of the firm. We test a theory concerning one of its capabilities: expanding the set of contractible variables. Our evidence suggests that this capability leads to less subcontracting. But changing information technology offers many other new capabilities, some of which improve resource allocation (“coordination”) along with incentives. In other research [Baker and Hubbard 2003] we examine the organizational impact of some of these other capabilities, in particular how OBCs’ coordination-enhancing capabilities affect shippers’ make-or-buy decision. Combined with this paper, this other work furthers our understanding of how information affects the organization of firms and markets.

APPENDIX 1

Our Bayesian estimates of the shares take the form

$$(11) \quad s_{rt}^b = \frac{a + n_{rt}}{b + N_{rt}} = \frac{b}{b + N_{rt}} s^* + \frac{N_{rt}}{b + N_{rt}} s_{rt},$$

where N_{rt} is the number of observations in cohort r at time t , n_{rt} is the number of positive observations, and s_{rt} is the share of positive observations. This expression is equal to the expectation of a random variable distributed $B(a + n, b - a + N - n)$, where B denotes the Beta distribution. A result from conjugate distribution theory is that $B(a + n, b - a + N - n)$ is the posterior distribution obtained by starting with initial priors $B(a, b - a)$ regarding the unknown mean of a binomial distribution, and Bayesian updating using N independent draws from the distribution, n of which are ones. [Degroot 1970] a and b are parameters that reflect the mean and variance of the distribution of initial priors. $a/b = s^*$ is the mean.

Intuitively, our Bayesian estimates are weighted averages of s^* and s_{rt} , where the weight on the latter increases with the size of the cohort. We set s^* to equal the mean ownership (or adoption) share for hauls in the same distance class in that year. For example, s^* for our Bayesian estimates of the owner-operator share for each long haul cohort is 0.211 in 1987 and 0.139 in 1992 (See Table I). We set $b = 10$; this implies that the observed shares and initial priors receive equal weight when cohorts contain ten observations. We have also estimated our models varying b from 2 to 20, and have found no substantial differences in any of our results.

APPENDIX 2: TRUCK OWNERSHIP AND OBC ADOPTION—FIRST DIFFERENCES
LINEAR PROBABILITY SPECIFICATIONS

Dependent variable: 1992 company driver share – 1987 company driver share
Observations are product-trailer-state-distance cohorts.

	Cohorts with positive owner-operator and company driver shares in both years		All cohorts		
OBC	0.043 (0.043)	0.048 (0.061)	0.025 (0.017)	0.019 (0.025)	
EVMS		-0.010 (0.081)		0.011 (0.034)	
OBC*long		0.084 (0.049)	0.151 (0.076)	0.010 (0.024)	0.010 (0.038)
OBC*medium		-0.067 (0.086)	-0.115 (0.108)	0.071 (0.031)	0.041 (0.039)
OBC*short		-0.275 (0.345)	-0.457 (0.384)	-0.028 (0.050)	-0.015 (0.059)
EVMS*long			-0.103 (0.092)		0.001 (0.046)
EVMS*medium			0.171 (0.239)		0.081 (0.066)
EVMS*short			1.085 (1.001)		-0.046 (0.112)
N		426		3676	

Both sets of estimates use observed ownership and adoption shares.

Cohort observations are weighted. The weight of cohort r equals $[n(r, 1987)k(r, 1987) + n(r, 1992)k(r, 1992)]/2$, where $n(r, t)$ is the number of observations in cohort r in time t and $k(r, t)$ is the average Census sampling weight among trucks in cohort r in time t .

All specifications include dummy variables that indicate whether the truck was operated <50 miles, 50–200 miles, or >200 miles from its base and $\ln(1992 \text{ trailer density}) - \ln(1987 \text{ trailer density})$ as controls.

Boldface indicates rejection of the null $\beta = 0$ using a two-sided t -test of size 0.05.

HARVARD BUSINESS SCHOOL AND NBER

UNIVERSITY OF CHICAGO GRADUATE SCHOOL OF BUSINESS AND NBER

REFERENCES

- Baker, George P., and Thomas N. Hubbard, "Contractibility and Asset Ownership: On-Board Computers and Governance in U. S. Trucking," NBER Working Paper No. 7634, April 2000.
- Baker, George P., and Thomas N. Hubbard, "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information," *American Economic Review*, XCIII (2003), 551–572.
- Brickley, James A., and Fredrick H. Dark, "The Choice of Organizational Form: The Case of Franchising," *Journal of Financial Economics*, XVIII (1987), 401–420.
- Brynjolfsson, Erik, and Lorin Hitt, "Information Technology and Organizational Design: Some Evidence from Micro Data," Massachusetts Institute of Technology, October 1997.
- Bureau of the Census, 1987 Census of Transportation: Truck Inventory and Use Survey (Washington, DC: Government Printing Office, 1989).
- Bureau of the Census, 1992 Census of Transportation: Truck Inventory and Use Survey (Washington, DC: Government Printing Office, 1995).
- Chakraborty, Atreya, and Mark Kazarosian, "Product Differentiation and the Use of Information Technology: New Evidence from the Trucking Industry," Brandeis University, 1999.
- Chamberlain, Gary, "Analysis of Covariance with Qualitative Data," *Review of Economic Studies*, XLVIII (1980), 225–238.
- Coase, Ronald H., "The Nature of the Firm," *Economica*, IV (1937), 386–405.
- Deaton, Angus, "Panel Data from Time Series of Cross-Sections," *Journal of Econometrics*, XXX (1985), 109–126.
- Degroot, Morris H., *Optimal Statistical Decisions* (New York: McGraw-Hill, 1970).
- Grossman, Sanford, and Oliver Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, XCVI (1986), 691–719.
- Hart, Oliver, and John Moore, "Property Rights and the Nature of the Firm," *Journal of Political Economy*, XCVIII (1990), 1119–1158.
- Holmstrom, Bengt, and Paul Milgrom, "The Firm as an Incentive System," *American Economic Review*, LXXXIV (1994), 972–991.
- Hubbard, Thomas N., "The Demand for Monitoring Technologies: The Case of Trucking," *Quarterly Journal of Economics*, CXV (2000), 533–576.
- , "Contractual Form and Market Thickness in Trucking," *Rand Journal of Economics*, XXXII (2001), 369–386.
- Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian, "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *Journal of Law and Economics*, XXI (1978), 297–326.
- Lafontaine, Francine, "Agency Theory and Franchising: Some Empirical Results," *Rand Journal of Economics*, XXIII (1992), 263–283.
- Lafontaine, Francine, and Scott Masten, "Contracting in the Absence of Specific Investments and Moral Hazard: Understanding Carrier-Driver Relations in U. S. Trucking," University of Michigan Business School, 2002.
- Maddala, G. S., "Limited Dependent Variable Models Using Panel Data," *Journal of Human Resources*, XXII (1987), 305–338.
- Maister, David H., *Management of Owner-Operator Fleets* (Lexington, MA: Lexington, 1980).
- Nickerson, Jackson, and Brian S. Silverman, "Why Aren't All Truck Drivers Owner-Operators? Asset Ownership and the Employment Relationship in Interstate For-Hire Trucking," *Journal of Economics and Management Strategy*, XII (2003), 91–118.
- Shepard, Andrea, "Contractual Form, Retail Price, and Asset Characteristics in Gasoline Retailing," *Rand Journal of Economics*, XXIV (1993), 58–77.
- Williamson, Oliver E., *Markets and Hierarchies: Analysis and Antitrust Implications* (New York: Free Press, 1975).
- , *The Economic Institutions of Capitalism* (New York: Free Press, 1985).