

Israel, Mark

Working Paper

Tenure dependence in consumer-firm relationships: An empirical analysis of consumer departures from automobile insurance firms

CSIO Working Paper, No. 0019

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Israel, Mark (2003) : Tenure dependence in consumer-firm relationships: An empirical analysis of consumer departures from automobile insurance firms, CSIO Working Paper, No. 0019, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/38632>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0019

**Tenure Dependence in Consumer-Firm
Relationships:
An Empirical Analysis of Consumer Departures
from Automobile Insurance Firms**

By

Mark Israel*

Northwestern University, Kellogg School of Management

First Draft: July 2001

This Draft: June 2003

* Address for correspondence: 2001 Sheridan Road, 6th Floor, Evanston IL, 60208. m-israel@kellogg.nwu.edu

Abstract

Long term consumer-firm relationships are an important feature of a large and growing number of industries. Consumer demand in many of these markets exhibits tenure dependence. A crucial question for firm strategy in such markets, then, is whether this tenure dependence reflects true relationship effects – so that a particular consumer’s preference for a firm actually increases with tenure. If so, firms face incentives to encourage consumers to remain through the early periods of the relationship, to enter new markets first in order to establish relationships, etc. However, this empirical regularity can also be explained by selection on unobserved heterogeneity. Empirically, the challenge is to use the data generally available in these markets – detailed records of the experiences of the consumers of one firm, with limited information about the broader market – to inform the distinction between these two explanations. This paper relies on individual specific price histories in auto insurance to this end. The principle findings are an economically relevant level of both relationship effects and heterogeneity, but a far more important role for heterogeneity in explaining the observed decline in departure rates.

JEL Classifications: D12, G22, L14, L84

Keywords: Tenure Dependence, State Dependence, Unobserved Heterogeneity, Switching Costs, Consumer-Firm Relationships, Automobile Insurance.

I. Introduction

A large and growing number of markets -- including insurance, telephony, and generally any “continuously provided service” -- are characterized by long-term consumer relationships with particular firms. A key feature of such markets is that consumers may evaluate firms based on the strength of the relationship. If so, a whole new set of strategic issues arises including: the use of low initial prices to develop strong relationships with consumers who will then be willing to pay higher prices and first-mover advantages through relationship development in new markets. The key empirical question behind such issues is one of consumer demand – do consumers in these markets actually prefer firms from which they have purchased repeatedly?

Answering this question raises challenges which are not easily handled by the standard, primarily static, empirical models of consumer demand. The central issue is: how can we measure the impact of a consumer’s tenure with one firm, when that tenure is determined endogenously by previous purchase choices? For example, consider Figure 1.1, which displays the percentage of consumers departing one automobile insurance firm in Georgia by tenure, measured in 6-month policy periods. From its peak at 10.3% in period 2, this percentage falls steadily to 4.8% after five years with the firm. Some would argue that this provides evidence for the importance of relationships, as captured by “true tenure dependence” in demand. But over a panel of consumers, such a pattern does not necessarily imply that the probability of departure falls with tenure for a **given individual**. Instead, this may simply reflect sample selection – often referred to a “spurious tenure dependence” -- generated by the ongoing departure of those “bad matches” who viewed the firm less favorably in the first place.²

In addition to raising a new set of questions, these relationship markets also give rise to datasets which are quite different from those in “traditional product markets.” Data following consumer purchases across firms, or even measuring industry wide market shares, are very hard to find. But, as a side effect of firms’ focus on customer relationship management, there are an increasing number of datasets detailing the complete purchase history of all consumers with one firm. An important, open question is where we can find identification power in such data. In particular, this paper considers how we can exploit the rich information on a consumer’s history with one auto

² A “bad match” is really any consumer who is consistently more likely to depart. This simplest example is a consumer who, for whatever idiosyncratic reason, realizes after initial purchase that her firm is not appealing to her – perhaps due to what she learns about her insurance agent, the company’s billing practices, the details of her policy, etc. But a “bad match” might also be a consumer who is more likely to depart any firm she joins, perhaps because she has lower than average search or switching costs. The key distinction here is between fixed characteristics of the individual which do not change over tenure with one firm vs. changes for a given consumer over tenure.

insurance firm, supplemented by only limited data on the broader market, to separate true and spurious tenure dependence in demand.

The econometric challenge arises because such data generally contain only one “spell” with the firm – a consumer purchases from one insurance company, say, for some period of time, then departs and does not return. Observing only one departure decision per consumer means that we have only one “experiment” per consumer, which makes the standard “fixed effects” identification strategy – include a fixed effect to control for consumer heterogeneity and then identify true tenure effects from the time series of departures given this fixed effect – impossible. So, we must rely on a random effects model, and in particular we must find some source of cross-sectional variation to separate the explanations. This challenge has long been discussed in the labor-econometric literature (see Heckman, 1991 for a summary). And, unfortunately, this literature establishes that, in this context, separating the two explanations is extremely difficult. In the textbook case, where all measured covariates are fixed through time, results are either extremely dependent on functional form assumptions or restricted to highly specialized cases such as the “proportional hazards” framework.

Fortunately, data on consumer-firm relationships allow us to make use of covariates which follow a distinct time-path for each consumer, based on her unique characteristics and purchase history.³ The main new idea in this paper is that we can use the information in these rich, time-varying, individual specific histories – in this case for the Georgia consumers of one automobile insurance firm -- as a powerful source of identifying information in the data.

Three features of the Georgia auto insurance market make these individual specific histories especially powerful. First, as in all states, insurance prices in Georgia vary by a large number of “rating classes” – determined by age, marital status, vehicle usage, zip code, etc. Second, these prices are set through a “prior approval” regulatory system, in which each firm negotiates a price for every rating class with the state insurance commission. As discussed more fully below, this regulatory process is assumed to generate exogenous prices. Finally, the negotiated rates for a given class vary greatly from firm to firm. And a firm’s position in the market price distribution varies greatly from class to class. Combining these features yields rich variation through time in a consumer’s *relative price* (defined throughout the paper as the price at the study firm minus the average price for that class in the overall market). That is, the fully observable changes in a

³ Much previous literature has included covariates which vary though time but follow the same path for all individuals – something like the unemployment rate. This paper relies critically on variables which follow a distinct time path for different individuals.

consumer's rating classes through time generate exogenous, individual specific, relative price histories.⁴

To see the identifying information in these histories, consider the pricing strategy that a firm would pursue given true tenure dependence – low initial prices to encourage relationship development followed by high prices to exploit “locked-in” consumers. Under sample selection, such a policy would fail, as low initial prices would simply encourage more “bad matches” to remain and then depart following a price increase. *Using the exogenous variation in price histories, we can observe exactly this pricing experiment.* That is, we can consider two types of consumers, each of whom has the same tenure with the firm and faces the same price today, while in past periods one type has faced lower prices than the other. If the “low price history” types are significantly more likely to depart than the “high price history” types when faced with today's common, higher price, we have evidence that sample selection has caused more bad matches to remain. If they are just as likely to remain, we have evidence that tenure has a structural effect.

At core, this is a standard sample selection issue. In estimating the departure probability at a given tenure with the firm, we are considering only those consumers who have chosen to remain with the firm to that point. So, the departure decisions in previous periods endogenously select the current sample. To measure the underlying impact of tenure -- and all other characteristics -- we need to control for these selection effects. And as always, to control for sample selection without relying completely on functional form or distributional assumptions requires variables which enter the selection equation but not the current departure probability. In this context – unlike the case of time-invariant covariates – previous prices vary independently of current prices and provide exactly such a variable. *That is, controlling for all current variables, the variation in departure rates associated with different price histories must arise from selection.* By observing these selection effects across many price histories, we are able to trace out the shape of the heterogeneity distribution and thus separate the two explanations.⁵

This argument is helped in two ways by the exogenous variation in prices. First, we do not need to include a model of the firm's dynamic profit maximization problem, which would

⁴ Note that exogeneity requires not only that the price menus are exogenous, but also that consumers do not change rating classes – by moving, choosing cars, experiencing claims, etc. – in response to insurance rates. This assumption is maintained from this point forward. Among other things, this rules out moral hazard. While this may not be completely correct, it seems to be a reasonable first approximation, and one which is consistent with empirical studies of auto insurance. (Chiappori and Salanie, 2000).

⁵ To see this most starkly, suppose that, for some period of time, some consumers faced prices from their firm low enough to ensure that they would not depart. For this set of consumers, there would be no selection effects. Their departure probability upon moving into a higher pricing class, as a function of tenure, would exactly trace out the true tenure effects. In practice, since no consumers are guaranteed to remain, some functional form assumptions are required, but crucially, parameter estimates within these functional forms are driven by information in the data.

substantially complicate matters. Second, the rich variation in rating classes creates correspondingly rich variation in price histories. It is as though firms have conducted the ideal experiment – lowering prices for some consumers up front and raising these prices at various points in tenure to measure the extent to which consumers are truly “locked-in.” Nevertheless, the basic logic of this approach extends to any relationship market in which data are available on distinct consumer price histories, distinct consumer experiences with the firm, distinct exposures to advertising promotions, etc. The challenge will be to find enough independent variation in these histories, and to extend the model to account for any endogeneity in the relevant variables.

To implement this approach requires a demand model tailored to relationship markets, in which both consumers’ focus and available data are primarily within one firm. This paper develops a simple yet flexible specification. Consumers’ evaluate firms based on quality and price, along with the status of the relationship and match quality at their current firm. To draw the key distinction as sharply as possible, relationship effects are measured with a distinct *tenure effect* for each period with the firm, while selection effects are captured with a *fixed* heterogeneity term. While focus is on the internal relationship with one firm, the decision to maintain or break the current relationship is triggered by “offers” randomly drawn from the broader market. This information arrival provides the model’s fundamental source of randomness, as observationally equivalent consumers may make different decisions due to differing choice sets. So the market **distribution** of prices and qualities – supplemented with consumer information parameters, modeled as the number of offers a consumer draws from the distribution – drives the model’s estimating equations.

The remainder of this paper proceeds as follows. Section II briefly reviews related literature. Section III describes the important institutional features of insurance markets and the data used in this study, and clarifies how these data assist in identification. Section IV details the empirical model used in estimation. Section V presents the results. Section VI extends the model to test its robustness. Section VII concludes.

II. Literature Review

This paper builds on three well-established literatures. First is the work in marketing on “brand loyalty,” or tenure dependence in brand choice (Keane, 1998, is perhaps the best example of this work). This literature focuses primarily on traditional product markets, rather than markets with long term consumer-firm relationships, extending demand models to include terms for experience with each brand. As such, it differs in several crucial ways from the current paper. First, the “scanner” data generally used in these studies contain multiple “spells” of purchases for each consumer for many different brands. As discussed above, this allows the papers to rely on variation

in the decisions of a given consumer across multiple trips to achieve identification, often by including a fixed effect.⁶ Specifically, the total number of purchases of each brand identifies the unobserved preference factors, while the *clustering* of those purchases identifies loyalty or switching costs. While a useful insight, this does not help in the current context.

These papers also address issues that are quite distinct from the relationship value vs. sample selection question considered here. For example, Chintagunta (1998) focuses on inertia vs. variety seeking in the characteristics of detergent and soft drinks. That is, do consumers stick to colas, or do they move back and forth between colas and other soft drinks through time? Separately, Shum (2000) considers whether advertising can overcome brand loyalty in breakfast cereal purchases.

Second, the paper builds on the empirical economics literature on switching costs, which consists primarily of reduced form supply side studies designed to determine whether equilibrium prices increase with switching costs. As such, the primary focus of papers like Sharpe (1997), Knittel (1997), and Viard (2001) is to find an exogenous measure of switching costs and then to estimate the comparative static, price response to changes in its level. In contrast, the current paper estimates a structural model allowing for fully general tenure dependence in demand --rather than the 0/1 form usually specified for switching costs -- directly from consumer purchase patterns.⁷

Finally, as noted above, a long literature in labor economics considers the separate identification of true and spurious “state dependence.” Similar to this paper’s focus on time spent with a single firm, the bulk of this literature considers the impact of time spent unemployed on the probability of finding work. However, the data available for such studies frequently include only time-invariant individual characteristics, or “individual-invariant” macro-economic indicators. The central result of this literature is that identification in such a context is either driven entirely by functional form assumptions, or requires restriction to the “proportional hazard” class of functions, in which fully non-parametric identification can be achieved (Heckman, 1991). Substantial progress has been made within this framework, with non-parametric estimators now available (see van den Berg and van Ours, 1996). Nevertheless, proportional hazard models remain quite restrictive and hard to justify economically.⁸ The contribution of this paper, then, is to make use of the identifying

⁶ Technically, a fixed effect only allows the constant term to vary by consumer. Goldfarb (2002) takes this methodological approach to its extreme, estimating separate regressions for each consumer in a panel of internet “clickstream” data, allowing fully general unobserved heterogeneity.

⁷ It seems quite likely that switching costs grow over tenure. As a consumer spends time with one firm, the effort required to “re-enter” the market – collecting the information required to establish new accounts, evaluate options, etc. – certainly grows. And over time with one firm, consumers take steps which make it painful to switch – selecting favorable billing options, informing relevant 3rd parties of their current firm’s identity, etc.

⁸ At core, the proportional hazard model makes two simplifying assumption. First, it limits the interaction between observable covariates and the unobservable error. A similar assumption is required here – in this case, we use the standard assumption that observables and unobservables are additively separable in the preference function.

information in individual specific histories, and thus to allow estimation with flexible functional forms which do not meet the proportional hazard restrictions.

III. Auto Insurance Prices, Data, and Identification

A. Auto Insurance Regulations and Prices:

The Georgia auto insurance market has two helpful features for this study. First, auto insurance is mandatory for those who wish to drive. So, it's reasonable to assume that at all times, each driver is purchasing insurance from some company. This paper models the decision to remain with that "current" firm or switch to another firm in the broader market.

Second, and even more importantly, prices are regulated through a "prior approval" regulatory system. Basically this means that each firm's menu of prices by rating class and claims history must be approved the state insurance commissioner, which requires actuarial cost data in support of all prices.⁹ In some states, such prior approval systems are basically "rubber stamps," but Georgia was selected for this paper because representatives of the study firm indicate that the commissioner "substantially restricts" pricing freedom.

We rely on this regulatory process to make simplifying assumptions about prices. Most importantly, prices are assumed to be exogenous to consumer demand, which in this context means that they are uncorrelated with any unobservables in the consumer's departure function. There are two sources of endogeneity that one might be concerned about. First, in the cross-section, if the firm knows that certain consumers are less likely to depart, it may charge them higher prices. Since we observe all the variables the firm uses to set prices, and have estimated specifications controlling for each, such endogeneity would have to arise from complicated, non-linear functions of the covariates. The core assumption, then, is that the "substantial restriction" on pricing freedom prevents any such subtle, demand-based pricing. Second, if the mix of consumers changes over time, firms may respond with changes in the price menu. But the regulatory process makes it quite costly to change prices -- the study firm changed prices only once over the study period -- limiting such concerns. In addition, we have estimated specifications including year dummies, with effectively no change in the results.

Second, it assumes that the elasticity of the departure rate with any variable does not vary with tenure nor the unobserved match quality. In contrast, one might expect price elasticity, say, to be lower for those who match well with the firm. In any case, the variation of price elasticity with tenure and observable consumer characteristics is observable in the data. A major advantage of this paper is that it maintains flexibility in this elasticity pattern, relying on the observed variation to pin it down, as clarified in Appendix I.

⁹ As noted above, because each firm uses its own actuarial data and methods to justify prices, and for a variety of other institutional reasons, the process leads to prices which vary greatly across firms, with the standard deviation of prices for a given rating class on the order of 25% of the mean.

Relying on this cost to changing prices, we also assume that the *distribution* of prices from other firms in the market is fixed over the study period for all rating classes. And because the regulator encourages all firms to use similar rating variables, we assume that all firms in the market use the same rating classes as the study firm. These assumptions greatly limit the data required to obtain useful measures of market prices.

Finally, the details of auto insurance pricing present two additional complications;

1. What Price is Used to Compare Policies? The choice of an appropriate price to use in comparing firms is complicated by the fact that consumers can choose which coverage to buy at each firm. In addition to the core liability policy required by law, consumers can choose to purchase extra liability coverage, or “first party” coverage for damage to their own person or vehicle. Theoretically, the choice of optional coverages may change if the consumer switches firms, making it difficult to compute the appropriate comparison price. For tractability, this paper assumes that consumers will purchase the same set of coverages at any alternate firm that they purchase at their current firm.¹⁰

2. How do Consumers Account for Future Price Changes? In a study of the impact of relationships, we need to account for the fact that consumers’ prices change over the course of that relationship. Such controls are possible here because we observe the full menu of prices at the study firm. However, lacking data on the full price menu at any other firm, we maintain the assumption that, while they are certainly forward looking, consumers compute expected price changes at any firm using their knowledge of the menu at the current firm.¹¹

These future price changes have two relevant impacts on the analysis. First, under the Georgia regulatory scheme, firms are allowed to offer discounts after consumers accrue enough “claim-free tenure.” Since these discounts provide an additional explanation for the falling departure probability over tenure, all analysis explicitly controls for expected price changes. Second, interpreting the impact of price would be difficult if, say, if a consumer with a low current price – relative to the price he would pay at other firms in the market – is more likely to have a high future price. There is no causal reason to believe this would be the case. Prices change only due to *observable* changes in a consumer’s rating class, and depending on which of the many possible class

¹⁰ All that is really required is that consumers compare their current firm to alternate firms based on the coverages they are buying at the current firm. Since they are unlikely to ask for the full menu of prices at alternate firms when price shopping, this basis for comparison seems natural.

¹¹ Along with the assumption on prices for alternate coverages, this implements the idea that consumers in relationship markets have limited information on firms in the broader market and thus base decision on what they have learned at their current firm.

changes occurs, this *relative price* may go up or down.¹² Still, assuming that consumers use the expected price change menu at the current firm to compute price changes at any firm deals with this in a precise way, as this assumption implies that the current difference between two firms' prices is the consumer's best indicator of the price difference going forward.¹³

B. Data

The data used in the study consist of a panel of 18,595 Georgia consumers joining one large, long-standing auto insurance firm between October 1991 and December 1998. Data of this form raise two immediate concerns. First, since the model's key parameters – for tenure dependence and unobserved heterogeneity – are assumed to be stable over the study period, one might worry if the firm is changing rapidly, say by gaining or losing market share. Fortunately, the firm's market share, both overall and broken down by rating class, county, etc. is extremely stable over the study period (and before and after this period as well). Second, in any study of a single firm, one worries whether it is an outlier. Fortunately, this is a large, mainstream firm, and a demographic comparison of its customers to the full set of Georgia insurance customers (available from the author on request) indicates that, at least on these observable measures, the firm is quite representative.

A consumer is defined throughout the analysis as a single auto insurance policy. Households generally keep the same policy when they switch cars, so this tracks a household through time. Each consumer is observed beginning with her initial purchase from the firm. 31.5% of these consumers voluntarily depart the firm during the sample period, with the remaining consumers either randomly falling off the sample or surviving through 1998.¹⁴ Table 3.1 presents the detail on entry and exit from the sample.

The record for each consumer is divided into 6-month periods, the length of an auto insurance policy. The primary variable to be explained is whether the consumer **completes** the period with the study firm, or departs before the end. Figures 3.1 through 3.3 mimic the graph of Figure 1.1 – percentage departing as a function number of periods with the firm – this time broken down by whether the primary insured is male or female, under or over 55, and a resident of Atlanta

¹² For example, prices change when a consumer turns 25. For some classes, the difference between the firm's price and the market average price is higher for those above 25 than those below 25, while for other classes the difference is smaller.

¹³ If consumers instead assume that the current price difference is temporary, and that following any class change both firms will draw a new price from a common distribution, so that the expected difference will be 0, all results go through. The only difference is the interpretation of the price coefficient, as it would be the impact of a *temporary* price difference.

¹⁴ Consumers randomly fall off the sample if they move to another state, or for a variety of administrative reasons which lead them to be assigned a new policy ID number.

or the rest of the state, respectively. The main point to take away is that a steadily declining departure probability is seen for all observable categories, including dozens which were computed but not reported. So, this phenomenon does not simply reflect selection on *observables*, further motivating the paper’s focus on true tenure dependence vs. selection on *unobservables*.

For each period, available explanatory variables include all those required to determine the consumer’s rating class – age, gender, marital status, vehicle characteristics, zip code, etc – collectively denoted as k_t for policy period t . These are assumed to be exogenous throughout the analysis. In addition, consumers are categorized in one of 4 “limit classes,” y_b , defined according to the selected limits of coverage for bodily injury and property damage liability.¹⁵ Representatives of the study firm indicate that more than 95% of the variation in limits can be explained by variation in income, so this is assumed to be exogenous and is used as a proxy for income in the analysis.

The fully observable claims history consists of a record of all past “chargeable” claims – those “at-fault” *claims* which impact prices. Only these chargeable claims are relevant for this analysis, and their occurrence is explicitly modeled for each consumer-period and estimated jointly with the departure probability. The period by period chargeable claims history for all periods prior to t (used to set prices) is denoted L_{t-1} . The number of consecutive periods at the study firm with no chargeable claims, used by firms to give price discounts, is denoted τ_t^{nc} .

The study firm’s menu of prices as a function of rating class and claims history, denoted $p_c(k_b, L_b, \tau_t^{nc})$, is fully observable. Prices in the broader market are summarized by the average price for each coverage as a function of rating class and claims history, denoted $\mu_p(k_b, L_{t-1})$.¹⁶ As noted above, the full price menu at each firm is assumed to be an exogenous outcome of the regulatory system. All prices are deflated by the CPI from January of the year in which the policy begins. Additional details on computation of this average price are contained in Appendix III.

Table 3.2 presents descriptive statistics for the continuous variables used in analysis, and Table 3.3 contains frequency tabulations for the discrete variables. One note here – while the study firm’s average price is somewhat lower than the market average, it is certainly not considered to be a “discount” firm, and as discussed in Appendix III, some of this difference reflects rating factors omitted from the average price measure (which we control for in estimation).

¹⁵ Bodily injury limits give the total amount payable for injuries to other individuals caused by the insured. Property damage limits give the total amount payable for damage to others’ property. y_t is set equal to 0 if the sum of these limits is less than or equal to \$75,000, to 1 if the sum is greater than \$75,000 and less than or equal to \$150,000, to 2 if the sum is greater than \$150,000 and less than or equal to \$350,000, and to 4 if the sum is greater than \$350,000.

¹⁶ Over the study period, insurance companies in Georgia share all claims data. So, both the current price and the market average price are correctly specified as functions of the consumer’s full claims history.

C. Identification:

The goal of this paper is to use these data to inform the distinction between true tenure dependence and sample selection. To frame this identification issue, note that the departure probability for a given individual in a given period can be written in general as:

$$\Pr(\text{Departure}) = D(\tau_t, x_t | \varepsilon) \quad (3.1)$$

where τ_t is the number of periods with the firm at time t – the study’s key relationship variable – x_t includes all other observables, and ε is any source of unobserved heterogeneity which impacts the departure probability – idiosyncratic match quality, differences in search or switching costs, etc. Crucially, all economic hypotheses about individual preferences – including that one’s preference for a firm grows with tenure – refer to this equation.

However, without observing ε , 3.1 can not be measured directly. If we had multiple spells per consumer, we could employ fixed-effect techniques to measure the other parameters of the equation, holding ε fixed. But with a single spell per consumer, we are forced to rely on a mixture of terms like 3.1 across different consumers with different values of ε . The mixture probabilities are determined by the selection out of the sample which has occurred since period 1, and thus are functions of the full history of x variables with the firm (including price) up to t , denoted as x^t . For example, suppose ε can take on J discrete value, $j = 1 \dots J$. Then, the observable mixture of departure probabilities, D_m is given by:

$$D_m(\tau_t, x_t | x^t) = \sum_{j=1}^J D(\tau_t, x_t | \varepsilon_j) \rho_j^t(x^t) \quad (3.2)$$

where $\rho_j^t(x^t)$ is the probability that $\varepsilon = \varepsilon_j$ at t conditional on the history of observables, x^t . Given an initial set of probabilities, effective in the first period with the firm, ρ_j^1 , the mixture probabilities are:

$$\rho_j^t(x^t) = \frac{\rho_j^1 \prod_{s=1}^{t-1} (1 - D(s, x_s | \varepsilon_j))}{\sum_{i=1}^J \rho_i^1 \prod_{s=1}^{t-1} (1 - D(s, x_s | \varepsilon_i))} \quad (3.3)$$

The problem is that the observable decline in $D_m(\bullet)$ with tenure can be explained either by true tenure effects in $D(\bullet)$ or by the shift in the $\rho_j^t(x^t)$ terms toward the “good matches.” So, we can not rely on this tenure path alone to separate the two explanations.

As always, to achieve this separation using information in the data requires a variable in the selection term, $\rho_j^t(x^t)$, that is not in the underlying departure equation. This is exactly where the distinct relative price histories come in. The key point is that the past prices do not affect the current departure probability for a given consumer, 3.1. But they do affect the mixture probabilities, $\rho_j^t(x^t)$. So, controlling for current prices, the impact of past prices on current departure

probabilities identifies the selection effects, leaving the variation in observed departure rates with tenure free to identify the true tenure effects.

Clearly, this argument rests critically on the amount of variation in a consumer’s relative price history, controlling for the current value. Fortunately, consumers in the panel change rating classes frequently, leading to relative price histories which vary greatly even after controlling for the current rating class and price. In 20.4% of all successive periods, consumers change relative prices, and only 23.6% of all observations of 6 consecutive periods are at a constant relative price. As a result, the standard deviation of the residual from a regression of relative price at $t - 3$ on relative price at t is \$36.12, 19.8% of the mean policy price of \$182.38. For $t - 6$ this increases to \$47.66.

Table 3.4 relies on this identification logic to provide initial evidence for the importance of sample selection, by showing that, even after controlling for current prices, tenure, and characteristics, past prices impact the current departure probability. The table reports the key coefficients from a Probit of the following form:

$$\Pr(\text{Departure}) = \Phi(a_0 p_t + a_1 p_{t-1} + a_2 p_{t-2} + a_3 p_{t-3} + a_5 (\tau_t = 4) + \dots + a_{16} (\tau_t \geq 12) + f(x_t)) \quad (3.4)$$

where p_t is the current relative price, p_{t-s} represents a lagged relative prices, $(\tau_t = x)$ is an indicator for tenure equal to x , and $f(x_t)$ is a set of demographic controls.¹⁷ The main point to take from this while the current price has the expected positive sign (higher prices make departure more likely), all three of the lagged price terms are significantly negative. Following the logic of the paper this means that consumers who have paid higher relative prices in prior periods – and thus have “revealed more preference” by remaining with the study firm – are more likely to remain today, suggesting the importance of persistent unobserved heterogeneity, which these prior prices are capturing.

While informative, results such as these are somewhat awkward and hard to interpret. So, the remainder of the paper uses a structural model of consumer choice, which makes use of such results to directly *measure* the importance of true tenure dependence vs. sample selection. The goal in developing the model is that it be simple and economically reasonable, yet flexible enough to allow the data to choose the appropriate specification. Section IV develops the details of the model, and with these in hand, Appendix I provides some additional notes on identification.

¹⁷ Obviously with 3 lagged prices, we can only include consumers in their 4th or later period with the firm. So, we do not include additional lags. Additionally, as discussed in Appendix III, the measure of market average price leaves out some rating factors, which we adjust for in the full model. So, the measure used here is quite rough.

IV. Empirical Model

A. Chargeable Claim Occurrence

The occurrence of a chargeable claim, denoted by $l_t = 1$, is explicitly modeled for each period with the study firm, conditional on the claims history at the time of initial purchase, L_0 .¹⁸

$$P(l_t = 1) = \lambda(k_t; L_0 | \eta) \quad (4.1)$$

η is an individual specific, time-invariant claims risk error, unobservable to both the firm and the econometrician. Its presence motivates the firm to use prior claims history in setting prices. For simplicity, η is implemented as a discrete random variable with 2 points of support, and assumed to be independent of all other variables in the model.¹⁹

B. Consumer Preferences

Because the model contains several elements which are necessarily forward looking—relationship development, expected price changes, etc. – we directly specify a consumer’s expected value function from remaining with the firm in period t .²⁰ This value function depends on the standard variables impacting current period utility – price, firm quality, and idiosyncratic match quality. It also depends on the state of the consumer’s relationship with her current firm, summarized by two variables: τ_t , the number of consecutive periods the consumers has spent with the firm, and τ_t^{nc} , the number of consecutive periods with the firm with no chargeable claims, included to control for the price discounts are given after 6 and 12 such periods.

It is specified as for the current firm, f_c , as:

$$\begin{aligned} V(f_c) = & \xi_c - \alpha_1(y_t)p_c(k_t, L_{t-1}, \tau_t^{nc}) \\ & + \alpha_2(\tau_t) + \varepsilon \\ & - \alpha_3(y_t)\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta) \end{aligned} \quad (4.2)$$

The first line captures the standard characteristics-based utility, a function of firm quality (ξ_c) and price, with price sensitivity determined by the limit class.²¹ $\alpha_2(\tau_t)$ measures the value placed on

¹⁸ The exact parameterization of this claim probability function is contained in Appendix II. Note that for simplicity, we model it as 0 or 1+ claims in each period and estimate a probit probability function. In less than 0.1% of all observed periods do consumers have 2 or more chargeable claims, so this simplification is natural.

¹⁹ What we actually see in the data is the distribution of η , *conditional* on joining the firm. For tractability, we directly specify this conditional distribution.

²⁰ Directly specifying V as a function of current period characteristics and state variables is a simplification. A full model could specify a period specific utility function as a primitive, then derive the value function. Future period utilities would depend on changes in characteristics as well as changes in the firm selected. Incorporating the induced probability of switching firms makes this problem intractable. So, this approach directly specifies the value function as a primitive. This generates a comparison rule between firms as a function of price and quality, augmented by the status of the relationship with the current firm, as desired.

²¹ Technically, price should enter as a function of remaining income, $y-p$. We can not be that precise since we do not have an exact measure of income, so this linear form is used for tractability.

the current consumer-firm relationship. The only restriction placed on this function is that the relationship is fully formed after 12 periods with the firm, so that:

$$\alpha_2(\tau_t) = \alpha_2(12) \quad \forall \tau_t \geq 12 \quad (4.3)$$

ε represents any source of unobserved heterogeneity in this preference function.²² To maintain focus on sample selection over the entire course of the consumer-firm relationship, it is assumed to be constant. And to simplify analysis, it is assumed to be unobserved until *immediately after* joining the firm. To limit distributional restrictions, a discrete distribution with five points of support – $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5$ and corresponding probabilities $\rho_1^\varepsilon, \rho_2^\varepsilon, \rho_3^\varepsilon, \rho_4^\varepsilon, \rho_5^\varepsilon$ – is estimated. The standard restrictions that probabilities sum to 1 and that the expected value is 0 leaves 8 free parameters, with the middle cases, ε_3 and ρ_3^ε , defined to satisfy the restrictions. As usual, ε is assumed to be independent of all other variables.

Finally, the third line controls for the impact of expected price changes. The function $\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta)$ is the average of the discounted, expected price changes for all future periods, based on changes in the rating class and claims experienced.²³ The valuation placed on these price changes is also allowed to vary with limits, as $\alpha_3(y_t)$. Appendix II contains details on the computation of this function for the current firm vs. alternatives, the relevant construct for the departure probability. Here, simply note that the key parameters (estimated with the model and reported in Section V) are δ_1 , consumers' expected price discount after 6 periods; δ_2 , consumers' expected discount after 12 periods; and β , the discount factor applied to future prices.

To switch to an alternate firm, then, consumers require that the expected value from the alternate firm, less any switching costs, be greater than the expected value at the current firm. The expected value at an alternate firm, $V(f_a)$, can be found by normalizing $\alpha_2(0)$ to 0 and replacing the unknown value of ε with 0, its expected value.²⁴

$$V(f_a) = \xi_a - \alpha_1(y_t) p_a(k_t, L_{t-1}, 0) - \alpha_3(y_t) \Delta_c(k_t, L_{t-1}, 0, \eta) \quad (4.4)$$

²² Other than the possibilities laid out above – idiosyncratic preferences over agent characteristics or firm practices, differences in search or switching costs, etc., we remain agnostic on the source of this heterogeneity, as well as the source of true tenure effects. This study seeks to separate these two basic classes of explanations for observed tenured dependence, with the aim of directing future work toward the more fruitful class to consider.

²³ η is included because it impacts the probability of future claims.

²⁴ Note the implicit assumption that consumers do not maintain relationships with firms they have previously purchased from, so that τ and τ^{nc} are 0 at all other firms. Also note that we use $\Delta_c(k_t, L_{t-1}, 0, \eta)$ because, as discussed above, consumers are assumed to base expectations on the menu at the current firm.

Switching costs are included to account for those conditions which make changing insurance companies particularly difficult. These measures are used. First, households with more than one insured vehicle, indicated by the variable $MCAR_t$, may find it more complex to switch insurers, as may households buying first party coverages for damage to their own vehicles, rather than just the mandated liability insurance, indicated by FP_t . Further, because claims generally take more than 1 period to resolve, it may be difficult to leave a firm in the period following a claim. So, an indicator for claims in the previous period, $ICLM_{t-1}$ is also included, yielding switching costs of:

$$S_t = \alpha_4 + \alpha_5 MCAR_t + \alpha_6 FP_t + \alpha_7 ICLM_{t-1} \quad (4.5)$$

Hence, a consumer's decision rule when comparing the current firm to a potential alternate is, switch if:

$$\begin{aligned} & \xi_a - \alpha_1(y_t)p_a(k_t, L_{t-1}, 0) > \\ & \xi_c - \alpha_1(y_t)p_c(k_t, L_{t-1}, \tau_t^{nc}) + \\ & \alpha_2(\tau_t) + \varepsilon - \alpha_3(y_t)(\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta) - \Delta_c(k_t, L_{t-1}, 0, \eta)) + S_t \end{aligned} \quad (4.6)$$

C. Offers as Triggers for Departure

This decision rule is applied when consumers receive offers from firms in the broader market. Each period, a consumer exogenously receives a number of offers, each of which is an independently drawn price and quality quote, (p_a, ξ_a) , from the market distribution. The consumer chooses whether to accept or reject the offer using 4.6. Rejected offers can not be recalled in future periods.²⁵ These random offers provide a simple empirical framework that incorporates consumers' limited information about the broader market, and utilizes it to provide the model's underlying source of econometric randomness.²⁶

The number of offers a consumer receives, denoted Ω_t , is used to capture the amount of information that a consumer has about the market distribution. This information level is assumed to depend on a consumer's ease of obtaining quotes, measured with a categorical variable representing the number of insurance agents per person in the consumer's zip code, AGT_t .²⁷ It is also allowed to vary with measures of the consumer's opportunity cost of time, including whether she is under 25,

²⁵ This may be because the quotes are only good for one period, because consumers do not recall them from period to period, or because consumers' characteristics change so that the quotes are no longer valid.

²⁶ This is basically a simplified search model, in which the offers arrive exogenously, rather than following explicit consumer search. The simplification is made to maintain focus on the key questions of the paper, while incorporating the fact that consumers do not have full information about market alternatives. And nothing in the data would permit identification of endogenous search intensity.

²⁷ AGT_t is equal to 0 if there are fewer than 0.25 agents per 10,000 people, 1 if this ratio is between 0.25 and 0.65, 2 if the ratio is between 0.65 and 1, and 3 if it's greater than 1.

indicated by $YNG_t = 1$, along with good student status, GS_t for those under 25, and whether the vehicle is used to commute to work $WK_t = 1$, for those over 25.²⁸ So, we have:²⁹

$$\Omega_t = \beta_0 + \beta_1(AGT_t) + \beta_2YNG_t + \beta_3WK_t + \beta_4GS_t \quad (4.7)$$

The relevant market distribution depends on k_t and L_{t-1} since these determine prices. For all values of k_t and L_{t-1} , the distribution is assumed to be bivariate normal:

$$(p, \xi | k, L) \sim N(\mu_p(k, L), \mu_\xi, \sigma_p^2, \sigma_\xi^2, \sigma_{p,\xi}) \quad (4.8)$$

A key advance of this study, relative to previous studies of auto insurance or other relationship markets, is the use of data on $\mu_p(k, L)$ to measure variation in the location of this distribution across consumers. All other moments of this distribution, however, are left as parameters, using the restriction that the variance terms do not depend on k or L to ease estimation.³⁰

D. Departure Probability

Using 4.8, note that the RHS of 4.6 is distributed as:

$$\xi_a - \alpha_1(y_t)p_a(k_t, L_{t-1}) \equiv V^{p,\xi} \sim N(\mu_\xi - \alpha_1(y_t)\mu_p(k_t, L_{t-1}), \sigma^2(y_t)) \quad (4.9)$$

where we've simplified the variance term since y_t is the only term which varies across consumers.³¹

So, the probability that a consumer turns down a given, randomly drawn offer is:

$$\Pr(\text{Turn Down Offer}) = \Phi\left(\frac{\tilde{V}_t}{\sigma(y_t)}\right) \quad (4.10)$$

where:³²

$$\begin{aligned} \tilde{V}_t &= (\xi_c - \mu_\xi) - \alpha_1(y_t)(p_c(k_t, L_{t-1}, \tau_t^{nc}) - \mu_p(k_t, L_{t-1})) \\ &+ \alpha_2(\tau_t) + \varepsilon_t - \alpha_3(y_t)(\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta) - \Delta_c(k_t, L_{t-1}, 0, \eta)) \\ &+ \alpha_4 + \alpha_5MCAR_t + \alpha_6FP_t + \alpha_7ICLM_{t-1} \end{aligned} \quad (4.11)$$

²⁸ While rough, these measures are better than those available for most demand studies. Presumably, young consumers have lower opportunity costs of time. And the availability of student and work status does provide some direct information on time demands.

²⁹ We normalize $\beta_1(0) = 0$. Note that there is no restriction that Ω_t be an integer. While, non-integer values strain the interpretation of Ω_t as a number of offers, they are natural in estimation, and the basic logic that consumers with higher values of Ω_t are more sensitive to the position of the current firm in the market distribution holds. In estimation, we did restrict Ω_t to be positive, although this restriction was never binding.

³⁰ A nice feature of this framework is that it can easily incorporate as much data as is available on the market distribution. So, future work should seek data on the variance of prices across consumer classes in addition to the mean. Or, if the full distribution is available, it could be used directly in estimation using simulation methods.

³¹ In practice, recall that y_t is a categorical variable. So, we estimate 3 parameters here, $\sigma(1)$, $\sigma(2)$, and $\sigma(3)$, with the normalization that $\sigma(0)=1$.

³² Note that with data on the consumers of one firm, we can not separately identify $(\xi_c - \mu_\xi)$, α_4 , and $\alpha_3(1)$. So, in estimation, we normalize $(\xi_c - \mu_\xi)$ and α_4 to 0. In practice, this means that the base tenure effect also includes the difference between the quality of the study firm and the market average quality and the base level of switching costs.

The probability of departing, then, is 1 minus the probability of turning down all Ω_t offers:

$$\Pr(\text{Departure}) = D(\tau_t, \tau_t^{nc}; y_t, k_t, L_{t-1} | \varepsilon, \eta) \equiv D_t = 1 - [\Phi(\frac{\tilde{V}}{\sigma(y_t)})]^{\Omega_t} \quad (4.12)$$

E. Likelihood Function

The likelihood contribution for a given individual is the joint probability of his departure decision in each period and his chargeable claims for each fully observed period.³³ There are two distinct cases to evaluate – an individual who is observed to depart the firm and an individual who is right censored, so that no departure is observed.

1. Observed Departure: Consider an individual who joins the firm in period t , remains s periods, and then voluntarily departs in $t+s$. His likelihood contribution, conditional on η and ε is given by:

$$\begin{aligned} L(\varepsilon, \eta) &= D(\tau_{t+s}, \tau_{t+s}^{nc}; y_{t+s}, k_{t+s}, L_{t+s-1} | \varepsilon, \eta) \\ & (1 - D(\tau_{t+s-1}, \tau_{t+s-1}^{nc}; y_{t+s-1}, k_{t+s-1}, L_{t+s-2} | \varepsilon, \eta)) \bullet \dots \bullet (1 - D(\tau_t, \tau_t^{nc}; y_t, k_t, L_t | \varepsilon, \eta)) \quad (4.13) \\ & [\lambda(k_{t+s-1}; L_0 | \eta)]^{l_{t+s-1}} [1 - \lambda(k_{t+s-1}; L_0 | \eta)]^{1-l_{t+s-1}} \bullet \dots \bullet [\lambda(k_t; L_0 | \eta)]^{l_t} [1 - \lambda(k_t; L_0 | \eta)]^{1-l_t} \end{aligned}$$

2. Right Censored: The only change in this case is that the first term is removed, since there is no observed departure. So, again with no loss of generality, consider an individual who joins the firm in period t , whose last fully observed period is $t+s-l$, and who is right censored in $t+s$. Her likelihood contribution, conditional on η and ε is given by:

$$\begin{aligned} L(\varepsilon, \eta) &= \\ & (1 - D(\tau_{t+s-1}, \tau_{t+s-1}^{nc}; y_{t+s-1}, k_{t+s-1}, L_{t+s-2} | \varepsilon, \eta)) \bullet \dots \bullet (1 - D(\tau_t, \tau_t^{nc}; y_t, k_t, L_t | \varepsilon, \eta)) \quad (4.14) \\ & [\lambda(k_{t+s-1}; L_0 | \eta)]^{l_{t+s-1}} [1 - \lambda(k_{t+s-1}; L_0 | \eta)]^{1-l_{t+s-1}} \bullet \dots \bullet [\lambda(k_t; L_0 | \eta)]^{l_t} [1 - \lambda(k_t; L_0 | \eta)]^{1-l_t} \end{aligned}$$

In both cases, the unconditional likelihood contribution is given by summing over the possible value of ε and η :³⁴

$$L = \sum_{j=1}^5 \sum_{k=1}^2 \rho_j^\varepsilon \rho_k^\eta L(\varepsilon_j, \eta_k) \quad (4.15)$$

Maximum likelihood estimation proceeds by maximizing the sum of the logs of the unconditional likelihood contributions for all 18,595 consumers in the sample.

³³ Fully observed periods are those for which all 6 months are observed. This excludes periods in which the consumer departs, randomly falls off the sample, or is right censored by the end of the panel, as not all claims can be observed for such periods.

³⁴ Note the implicit assumption that ε and η are independent of one another and all other variables in the model.

F. Flexible Functional Forms

While the model has a large number of moving parts, it gives rise to estimating equations which are actually quite simple and flexible, allowing the data to drive identification. Because the departure and chargeable claim probabilities are based on Normal CDFs, the likelihood function in 4.13 – 4.15 is, at core, a bivariate probit, where the dependent variables are independent of one another, save for the cross equation restriction reflecting the impact of claims risk on expected price changes. The use of external offers in the departure probability further generalizes this specification -- by changing the base number of offers, β_0 , the probability can take the form of a Normal CDF raised to any power.³⁵ And the key terms inside the departure probability are also quite flexible. Relationship effects are specified as a general function of tenure. And the distribution of ε is approximated by a 5-point discrete distribution, which admits a wide range of shapes. Appendix I provides some additional notes on the benefit to and identification of these flexible forms.

V. Results

A. Non-Relationship Parameter Estimates

Before turning to the paper's core questions, consider the "non-relationship" parameter estimates in Table 5.1. These follow the expected patterns. First, consumers with higher liability limits (and presumably higher income) are generally less sensitive to current and future prices. Second, the switching cost estimates suggest more difficulty switching for consumers with more than one insured car, for consumers purchasing first party coverage, and for consumers who have recently experienced a claim. Third, the number of offers increases by roughly 0.1 per period in zip codes with more than 0.25 insurance agents per 10,000 people, relative to those with fewer agents. However, there is no significant difference between agent classes 1 – 3, suggesting that while it helps to have access to some insurance agents, information does not increase continuously with more agents. Fourth, consumers under 25 receive roughly 0.1 more offers per period, perhaps reflecting a lower opportunity cost on the time required to process offers, although this estimate is not significantly different from 0. This impact is almost exactly offset for those under 25 who are good students, perhaps because they face more demands on their time. Finally, adults who commute to work receive roughly 0.2 fewer offers, again consistent with a time-cost interpretation.

³⁵ In addition, robustness tests were conducted by specifying the market distribution of prices and qualities as independent Chi-Square or Logit draws. The key results were quite robust to these changes.

B. True Tenure Effects

Table 5.2 contains estimates of the paper’s key parameters – those for true tenure effects, the heterogeneity distribution, and expectations for tenure based price discounts. The first section of the table reports the true tenure effects, $\alpha_2(\tau)$. To simplify discussion, the marginal impact of each unit of tenure, $\tilde{\alpha}_2(\tau)$, is reported, where:

$$\alpha_2(\tau) = \sum_{t=1}^{\tau} \tilde{\alpha}_2(t) \quad (5.1)$$

The first notable feature of these estimates is the large negative coefficient at tenure 2, corresponding to the period 2 spike in the departure probability. At least some of this is likely an artifact of this dataset, in which we only observe consumers who have already entered the firm. While the discrete time model from Section IV posits separate decisions to join at $t-1$ and then to remain at t , a consumer may actually wait to reconsider her decision to enter the firm until after her first 6-month policy has ended. So conditioning on joining the firm implicitly conditions on remaining through period 1, at least for some consumers.

However, the fact that tenure effects continue to fall slightly in period 3 suggests that there may also be something deeper going on here. Officials of the study firm suggest that this may reflect a “honeymoon period” – relationship effects are strong just after the initial meeting with an agent, but this goodwill depreciates quickly – meaning that tenure effects actually fall briefly, before the long term process of learning and relationship building sets in. In this light, future research should consider data on interactions with the firm as potential direct measures of relationship strength.

Alternatively, previous work by Jovanovic has suggested that such spikes may be a more subtle form of sample selection -- if consumers don’t learn the value of ε immediately, they may wait to accrue enough information before departing (Jovanovic, 1982). In the current context, this may suggest that consumers will wait until after their first claim to consider switching firms. However, Israel (2003) closely examines consumer departures following claims and finds that, while consumers who have experienced claims are more likely to depart, these effects are much smaller than the spike observed here.³⁶ And, even after controlling for claims occurrence (or any other observables) the basic pattern of a decline in the period 2 tenure effect, followed by a steady increase is robust, suggesting that this reflects either a data artifact or a behavioral pattern for all consumers, rather than selection on consumer heterogeneity.

³⁶ That paper finds robust evidence that claims experience *increases* departures, which it interprets as evidence that consumers enter the firm optimistic about service quality and are then disappointed by what they learn from actual service. In any case, this means that impact of learning on repeat purchases – which has been studied by Akerberg (2002) among others – is quite distinct from the tenure effects studied here in which departures *fall* with tenure.

Still, because the period 2 drop in tenure effects is hard to interpret and because our goal is to explain the long term decline in departure rates from the period 2 peak, the rest of this section focuses on the growth in relationship effects from period 2 on, as graphed in Figure 5.1. Even excluding period 1, this growth is far from monotonic, with a large decline at tenure 6, and smaller declines at tenures 3, 9, and 12.³⁷ One lesson from this is the apparently complex interaction between learning/relationship building and the depreciation of the knowledge and relationships, and thus the need for flexible functional forms in studies of these phenomena. In particular, the minimal growth from tenure 10 to 11 and the subsequent decline strongly supports the paper’s assumption that relationship building is essentially complete – and in fact may be depreciating -- after 6 years. For the earlier periods, however, if we assume that the spikes at tenures 5 and 10 are somewhat overstated, relationship effects appear to grow roughly linearly, with an average increase of 0.02 per period. Also, note the three tenures at which relationship effects grow without an offsetting decline in the subsequent period – 4, 7, and 10. Using these natural breaks, the results are conveniently summarized in four intervals – the low at tenures 2 and 3, the increase to periods 4 – 6, the additional increase to periods 7 – 9, and the high in periods 10 – 12. For the remainder of this paper, then, we focus on the situation at the end of each of these intervals, summarizing results with periods 3, 6, 9, and 12.

To what extent do we have *statistical* evidence for the importance of relationship effects? At least somewhat, as the period by period increases are significant for periods 4, 5, and 10, although so is the decrease at period 6. And we can easily reject the joint hypothesis that $\tilde{\alpha}_2(3)$ through $\tilde{\alpha}_2(12)$ are equal to 0 with a likelihood ratio test. However, we can not reject the hypothesis that $\tilde{\alpha}_2(6)$ through $\tilde{\alpha}_2(12)$ are equal to 0, even at the 10% level, so there is limited statistical evidence for relationship effects after 2.5 years with the firm.

More importantly, though, how substantial are these effects *economically*? The first two columns of Table 5.4 consider this, by measuring the cumulative growth in tenure effects, from period 2 on, in dollars. And there is clear evidence for economically relevant relationship effects – worth \$5.52, or 3.0% of the average policy price by period 6; \$13.73, or 7.5% of the average price by period 9, and \$27.12, or 14.9% of the average policy price by period 12. For some additional perspective, note that \$27.12 is 59.7% of the increase in switching costs associated with insuring

³⁷ The declines at periods 6 and 12 are potentially interesting. Recall that these are the tenures just prior to price discounts for consumers with no chargeable claims. So these declines may represent a “Jovanovic effect.” Consumers who are reaching these potential discounts for the first time are unsure exactly what they’ll get, and thus “withhold judgment” on the firm until they observe the outcome. Israel (2002) contains a much more detailed look at consumers’ reaction to expected price changes.

multiple cars. So, relationship effects can not be ignored, but they are not particularly large compared to the impact of other policy characteristics.

C. Unobserved Heterogeneity

Similar to the results for true tenure effects, the estimates in Table 5.2 suggest a complex shape for the heterogeneity distribution, not matching any standard distribution. In particular, we see a marked bi-modal pattern, with the largest number of consumers in class 2, with a slightly worse than average view of the firm, and a second peak in class 5, with a very positive view of the firm. Together these two classes contain 57% of consumers, with the remaining 43% divided quite evenly between the other three classes. Overall, this distribution suggests a large amount of heterogeneity, with a range of 1.761, equivalent to \$294.97, and a standard deviation of 0.525, equivalent to \$87.94. And we can easily reject the hypothesis of no heterogeneity at any reasonable level of significance.

How *economically* substantial is the induced sample selection? The last 2 columns of Table 5.4 consider this by computing the change in the mean value of ϵ by period for a “base case” consumer – with price equal to the market average, no accrued claims free tenure, and all other variables at their sample means. By period 6, this expected value has increased by \$20.76, or 11.4% of the average policy price, meaning that the average remaining consumer values the firm by \$20.76 more than the average consumer at period 2. By tenure 9, this increase is \$30.94, 17.0% of the average policy price, and by tenure 12 it’s \$44.17, 24.2% of the average price. So, selection induces a notable increase in the average preference for the firm over tenure, even holding all observable characteristics fixed.

D. Expected Price Changes

Finally, recall the third way that preferences vary over time spent with one firm, tenure-based price discounts. The final section of Table 5.2 indicates that consumers expect a price break of \$17.10 after 6 periods with no chargeable claims and an additional break of \$11.10 after 12 periods. These are fairly consistent with the average changes actually observed. However, consumers heavily discount future periods, with a 6-month discount factor of 0.564, which implies an annual discount factor of 0.318.³⁸ Table 5.3 lists the average, discounted future price changes associated with these expectations for each accrued period of chargeable-claims free tenure, relative to the value at $\tau^{nc} = 0$. And we see that these averages are fairly small, peaking at -\$9.53 at tenure 6, only 5.2% of the

³⁸ Unlike the discount factors in most dynamic models, this is estimated quite precisely.. This is because we have a definite, measurable, future price change, with which to pin down the discount factor.

average policy price. The impact on the departure probability – computed for a consumer with tenure equal to 1 and all other variables at their sample averages – is correspondingly small, peaking at a 0.6 point reduction from $\tau^{nc} = 0$ to $\tau^{nc} = 6$. So, the remainder of this section fixes τ^{nc} at 0, and focuses on the impact of true tenure effects vs. selection.

E. True Tenure Effects vs. Selection

To this point, we have evidence that the declining departure rate reflects economically relevant levels of both relationship effects and selection. But which is more important? Comparing the figures from Table 5.4 provides a first answer – the true change in preference for a *given* base case consumer vs. the change in the *mean* preference induced by the departure of poor match types. This evidence points toward the importance of selection. At period 6, for example, the increase in the mean value of ϵ is nearly 4 times as large as the increase in the relationship value. As relationship effects continue to accrue, this gap declines somewhat, but the selection effects are still more than twice as large at tenure 9 and more than 1.5 times as large at tenure 12.

However, this comparison does not tell the whole story. Selection actually has two impacts on departure rates, both illustrated in Table 5.5, which shows the departure probability for a base case consumer of each match type and the proportion of base case consumers in each type over tenure. The first impact, which is captured in Table 5.4, is the direct effect of a shift toward consumers with a higher match quality and the associated fall in departure rates. But, the second, more subtle effect is that high match quality consumers are much less impacted by true tenure effects. From period 2 through 12, the lowest match quality consumer's departure probability falls by 4.8 percentage points, while the fall for the highest match quality consumer is only 0.4 points. So, selection both decreases the departure rate directly, and shifts the balance toward those consumers who are less affected by tenure effects. Similarly, the increase in tenure effects implies that, in later periods, the increasing mean of ϵ has less impact on departure rates. So to allow for these interactions, we need to directly measure the relative importance of each explanation in explaining the falling departure rate.

Table 5.6 takes up this challenge, decomposing the decline in departure rates for a base case consumer into the portion which can be explained by true tenure effects alone and the portion which can be explained by heterogeneity alone.³⁹ The impact of tenure effects alone is computed by setting ϵ to -0.1275 -- to match the departure probability predicted by the full model at period 2 -- leaving all other parameters unchanged, and computing the predicted departure probabilities. This still yields

³⁹ Note that the sum of the 2 decomposed effects is greater than the overall decline. This results from the interaction described above – in the absence of selection the impact of increasing tenure effects is greater, and vice versa.

notable declines in departure probabilities, but much smaller than in the full model – roughly 1/5 of the overall decline through period 6, 1/3 through period 9, and 1/2 of the full decline through period 12. The impact of heterogeneity alone is computed by setting $\tilde{\alpha}_2(\tau) = 0$ for all $\tau > 2$, leaving all other parameters unchanged, and computing the predicted departure probabilities. Here the predicted declines are much closer to those from the full model – more than 4/5 of the full decline through periods 6 and 9, and nearly 3/4 of the full decline through period 12. Figure 5.2 graphs each path of predicted departure probabilities. And while the flexibility of the tenure effects allows them to match the period-by-period shape of the full model, the selection model clearly comes much closer to the more important objective of matching the overall decline.

One implication of the importance of selection is that departure probabilities should decline most steeply for those who have faced high prices in previous periods. Table 5.7 demonstrates using the full estimated model. Each column contains the decline in the departure rate from tenure 1 through the specified tenure for a consumer currently in the base case (with price equal to the market average), but who has faced the specified relative price in all previous periods.⁴⁰ Consumers who have faced a price \$100 below the market average in periods 1 -5 see their observed, base case departure probability *rise* by 0.08 percentage points from period 1 to 6, as the limited departure of bad matches is overwhelmed by the drop in tenure effects from period 1 to 2. In contrast, those who have faced a price \$100 above the market average see their observed departure probability fall by 2.44 points, as many bad matches exit. Through period 12, the “low price history” consumers see the departure probability fall by only 2.11 points vs. 4.70 points for those with high price histories. This strong impact of historical prices clearly has implications for pricing strategy – most notably, an incentive to set high early prices to screen out bad matches rather than low prices to “lock consumer in” – which need to be explored in future work.

F. Implications

Finally, we consider the implications of these findings in two areas. First, *econometrically*, what would happen to the estimated parameters if one naively ignored either unobserved heterogeneity or true tenure effects in estimation? Second, *strategically*, what do the estimates say about a firm’s incentive to build relationships?

Table 5.8 considers the econometric issue. Column 1 reproduces the results from the full model, focusing on two summary measures – the full increase in *tenure effects* from period 2 – 12,

⁴⁰ We start in period 1 so that there is no price history, meaning that all three columns begin at the base case departure rate for period 1.

0.162, and the standard deviation of the *heterogeneity* distribution, 0.525. Column 2 contains these results for a model estimated with no heterogeneity term (defined by dropping ε from 4.11). In this case, the estimated tenure effects increase to 0.406, a striking increase of more than 150%. So, interpreting the full decline in departure rates as true relationship effects leads to grossly overstated values, yielding an estimated increase in relationship value from period 2 to 12 of \$68.00, rather than the \$27.12 from the full model.

In contrast, Column 3 contains the results for a model with no tenure effects beyond period 2, estimated by setting $\tilde{\alpha}_2(\tau) = 0$ for all $\tau > 2$. This increases the estimated standard error of ε from 0.525 to 0.598, an increase of only 13.9%. So, ignoring true tenure effects leads us to conclude that the standard deviation in match values is equivalent to \$100.16, rather than \$87.94 in the base model, a fairly small change. Intuitively, this is because the heterogeneity distribution is estimated from the impact of price history on the current departure rate, rather than the tenure path in departures, so eliminating true tenure effects has limited impact.

A more subtle econometric issue is that, absent controls for heterogeneity, estimates of the other coefficients in the model will be biased. Consider the price coefficient. To the extent that prices are positively correlated, a high price today implies a high price history, and thus individuals facing such a high price are more likely to be high ε types, biasing the price coefficient toward 0. So, consistent price estimates require controls for selection effects. Table 5.9 illustrates this econometric impact – in the model with no heterogeneity the estimated price coefficients are substantially lower than in the full model.

What about the strategic implications of these findings? The basic strategic issue raised by consumer-firm relationships has been carefully documented in the switching costs literature – the tradeoff between the desire to encourage consumers to remain to strengthen the relationship and the desire to profit from loyal consumers through higher prices today (Klemperer, 1995). In contrast, unobserved heterogeneity raises no such tradeoff as a firm can not change an individual’s type by inducing him to remain. So, the key strategic question for this paper is: how strong is the incentive to induce consumers to remain to later periods? Put simply, how much would a firm be willing to pay to “change” a tenure 2 consumer into a tenure 3 or higher consumer?

To simplify this calculation, suppose that price regulation generates a markup of M for every consumer in every period. Then, if we denote the departure probability of a particular consumer at tenure τ as $D(\tau)$, the firm’s profit from that consumer from that point forward is given by:

$$\Pi(\tau) = M \sum_{s=0}^{\infty} \gamma^s (1 - D(\tau + s)) \quad (5.2)$$

where we are implicitly assuming that if a consumer departs, there is 0 probability that he'll ever return. Setting the firm's six-month discount factor, γ , to 0.97 and recalling that tenure effects stop accruing after period 12, we can solve this as a function of tenure for any set of consumer characteristics (including a value for ε). The difference between the value at tenure 2 and the value at higher tenures, then, gives the firm's willingness to pay to "move" a consumer to a higher tenure.

Table 5.10 lists these willingness to pay figures in markup units for a base case consumer with $\varepsilon = 0$. The first column indicates that, given the estimated tenure effects, the firm would be willing to pay just over $\frac{1}{2}$ of a single-period markup to move the consumer from tenure 2 to tenure 6 and just less than 1 markup to move her to tenure 12. In contrast, the second column computes these figures if true tenure effects accounted for all of the observed decline in departure rates, that is using the "no heterogeneity" model from Table 5.8. In this case, the firm would be willing to pay nearly 1.5 markup units to move a tenure 2 consumer to the relationship level at tenure 6 and more than 2.3 markups to move her to tenure 12. Overall, then, there is some incentive for firms to encourage relationship development, but they wouldn't sacrifice a full, single-period markup to move a consumer from the base of tenure 2 all the way to tenure 12, and their willingness to pay is roughly 40% of what it would be if the observed decline in departure rates entirely represented true relationship effects.⁴¹

VI. Model Extensions and Robustness

Arguments against this paper's identification logic start from the premise that history impacts current departure probabilities through some mechanism other than selection. This section considers three alternate mechanisms, both as robustness checks and as interesting extensions in their own right.

A. Heterogeneity in Relationship Development

The analysis thus far has assumed that relationships develop at the same rate for all consumers. However, if the firm treats certain consumers differently, relationships may accrue at

⁴¹ In a full, dynamic profit maximization model, we would have to solve for the price schedule as a function of tenure for each consumer type. This price schedule would endogenously determine a firm's willingness to pay for stronger relationships. And the schedule itself would reflect the strength of true tenure effects through the "discount" given in early periods to encourage consumers to remain to higher tenures. Israel (2001) contains a preliminary version of such a model, and uses the observed price path in the unregulated Illinois insurance market to help identify the importance of true tenure effects vs. selection.

different rates. And these treatment differences may be correlated with prices, creating an alternate mechanism for historical prices to impact current decisions.⁴²

The firm actually has an explicit structure in place to charge higher prices to certain consumers whom they consider to be especially high risk. These consumers are placed in the “standard company”, while the vast majority of consumers are placed in the “preferred company.”⁴³ If differential treatment is important it should certainly show up for this distinction. We implement this by allowing the incremental tenure effect to differ by a fixed amount for each period spent in the standard company, so that:

$$\alpha_2(\tau) = \sum_{s=1}^{\tau} \tilde{\alpha}(s) + \delta_0 STND_{\tau} + \delta_1 \sum_{s=1}^{\tau} STND_s \quad (6.1)$$

where $STND_s$ indicates that the consumer was in the standard company at tenure s , and we include a separate term for $STND_{\tau}$ to control for current placement in the standard company.

Table 6.1 presents the relevant parameters from this model. Neither δ_0 nor δ_1 are significantly different from 0, so there is limited evidence for differential treatment.⁴⁴ More importantly, the estimated heterogeneity distribution and the tenure path for the bulk of consumers – in the preferred company -- are essentially unchanged. So, the model’s basic results are robust to the inclusion of heterogeneous tenure effects, and even strengthened for those in the standard company.

B. Heterogeneity in Price Parameters

The unobserved heterogeneity in the basic model consists entirely of an additive effect on preference for the current firm. However, there may also be “random coefficients” on covariates such as price, which provide an alternate channel for selection. Fortunately, the additive error can be separately identified from such unobserved heterogeneity in price sensitivity. Under heterogeneity in match quality, a consumer who has remained with the firm in the face of high relative prices is more likely to be a good match and thus more likely to remain with the firm at any current price. Under heterogeneity in price sensitivity, however, a consumer who has remained in the face of high relative prices is more likely to be a low sensitivity type, and thus is more likely to remain if today’s price is still high, but more likely to depart if today’s price is low.

⁴² It’s easy to argue this correlation with prices in either direction. The firm may be especially good to high price consumers if they generate more profits. Or, the firm may negotiate with the regulator for low prices on those consumers that it values most. The best argument against such correlation, then, may be that these various influences net out.

⁴³ Consumers do not switch back and forth between companies. Once they’re assigned to one, they stay there.

⁴⁴ Many other possibilities for differential treatment were also considered – including age, geographic location, policy choices, etc. None were significant.

Table 6.2 presents the relevant results from a model extended to include unobserved heterogeneity in price sensitivity. This is accomplished by writing the coefficient on current price in 4.2 as $-(\alpha_1(y_t) + v)$. v is assumed to have a two-point discrete distribution, v_0 and v_1 , with the single free probability parameter, $pr(v = v_1)$ denoted by ρ_v . Because the standard deviation of $V^{p,\xi}$ depends on the price coefficient, it is modified to take the form $\sigma_y(y)\sigma_v(v)$, with the normalizations that $\sigma_y(0)$ and $\sigma_v(v_0)$ each equal 1.

With these changes, we do find some evidence for heterogeneity in price sensitivity, although the estimates are very imprecise. However, the distribution of ε and the path of tenure effects are essentially identical to those in the base model, so the paper's results are robust to this extension.

C. Search Following Class Changes

Finally, it seems likely that consumers exert more control over their search process than implied by the model's simple, exogenous offers. In particular, they may search hardest in periods when they have just entered a new pricing classification, to see if there are better options in the market. If they have found no better alternatives after a few periods in the same class, however, this incentive to search will decline.

This possibility complicates the paper's identification argument, as it suggests another role for history – a high departure probability observed for consumers shifting from a low price history to a high current price may reflect greater search effort following the class change. Fortunately, this can easily be distinguished from selection on match quality, which implies that consumers shifting from a *high* price history to a lower current price will be *less* likely to depart. Search following class changes indicates that consumers will want to pursue more options following any class change and thus will be *more* likely to depart following such a transition.

To explore this, we extend the model to include CH_t , an indicator for a change in pricing classification from $t-1$ to t , by adding $\beta_5 CH_t + \beta_6 CH_{t-1}$ to the information arrival function. The key results, presented in Table 6.3, can be summarized as follows. First, the point estimates do imply more search in the periods following a class change. But, second, these coefficients are not statistically significant. And, most importantly, results for tenure effects and heterogeneity are essentially unchanged, so the core results are robust to this extension.

VII. Conclusion

Long term consumer-firm relationships are an important feature of a growing number of industries. Consumer demand in many such markets exhibits tenure dependence – the probability

that a consumer breaks her current relationship falls over tenure. A crucial question for firm strategy is whether this tenure dependence reflects true relationship effects, so that a given consumer's preference for a firm actually increases with tenure. If so, firms face incentives to encourage consumers to remain through the early periods of the relationship, to enter new markets first in order to establish relationships, etc. However, this empirical regularity can also be explained by selection on unobserved heterogeneity. In this case, the incentives to "lock consumers in" are absent.

Empirically, the challenge is to use the data generally available in these markets – detailed records of the experiences of the consumers of one firm, with limited information about the broader market – to inform the distinction between these two explanations. This paper argues that individual specific histories with the firm are highly informative in this context. For example in auto insurance, consumers have individual specific price histories, generated by movement through rating classes over time. These price histories provide a variable which drives sample selection without directly entering the current period departure probability, and thus they allow us to trace out the heterogeneity distribution separately from the path of true tenure effects. In particular, the data contain pricing "experiments," such as two types of consumers who face the same price today, while one has faced lower prices historically. The current departure probabilities of these two consumer types are highly informative – if the low price history consumers are more likely to depart, this suggests that more bad matches have remained and thus supports sample selection; if not, it suggests that preferences have truly increased with tenure.

To implement this logic, the paper develops a structural model of consumer-firm relationships, allowing for both unobserved heterogeneity and tenure effects, and incorporating the broader market through "offers" that trigger a consumer's decision to depart. The model is estimated by maximizing the likelihood of the observed claims and departure decisions of 18,595 consumers of one auto insurance firm in Georgia. The basic findings can be summarized as follows:

1. There is evidence for economically relevant levels of both true tenure dependence and unobserved heterogeneity.
2. On several fronts, selection appears to be the more important explanation. The increase in the average match quality from period 2 to 12 is equivalent to a \$44.17 reduction in price, while the growth in tenure effects is equivalent to only a \$27.12 reduction. This increase in match quality on its own can explain nearly $\frac{3}{4}$ of the observed decline in departure probabilities, while the increase in tenure effects explains less than $\frac{1}{2}$.
3. Econometrically, the impact of ignoring selection is much more severe than ignoring true tenure effects. Ignoring selection leads us to overstate the tenure effects by more than 150%, while ignoring the tenure effects only increases the estimated standard deviation

of the heterogeneity distribution by 13.9%. And ignoring selection leads to substantial biases in the estimated coefficients on price and other covariates.

4. Strategically, the estimates imply that a firm would be willing to sacrifice roughly 1 single-period markup to increase a consumer's relationship value from its level at tenure 2 to its level at tenure 12. While economically relevant, this is substantially less than the 2.33 markups they'd be willing to sacrifice if the observed decline in departure rates entirely reflected relationship effects.

To generalize these results, future work should consider other relationship markets, including credit cards, internet service provision, health care, and long distance telephony, as well as insurance in less regulated states. Such extensions present two key challenges. First, we need to find covariates with sufficiently rich, individual specific, time paths. Price variation is ideal, but other experiences which impact consumers' purchase decisions in previous periods but not today – advertising promotions run in certain periods, movement between regions with different sets of competitors, entry and exit by competitors – may also work. Second, when we move to a non-regulated setting, we'll need to explicitly model the endogenous determination of these variables. However, as argued in Israel (2001), this may also allow us to use firms' behavior – in addition to consumers' reactions to this behavior – as a source of identifying information.

References

- Akerberg, D. (2002). "Advertising, Learning, and Consumer Choice in Experience Good Markets; An Empirical Examination," mimeo.
- Berger, L.A., P.R. Kleindorfer, and H. Kunreuther (1989). "A Dynamic Model of the Transmission of Price Information in Auto Insurance Markets", *The Journal of Risk and Insurance*, 56, 17-33.
- Chiappori, P. and B Salanie (2000). "Testing for Asymmetric Information in Insurance Markets", *Journal of Political Economy*, 108, 56-78.
- Chintagunta, P.K. (1998). "Inertia and Variety Seeking in a Model of Brand-Purchase Timing", *Marketing Science*, 17, 253-270.
- Dahlby, B. and D.S. West (1986). "Price Dispersion in Automobile Insurance Markets", *Journal of Political Economy*, 94, 418-438.
- Goldfarb, A. (2002). "Using Household-Specific Regressions to Estimate True State Dependence at Internet Portals," mimeo.
- Heckman, J.J (1991). "Identifying the Hand of the Past: Distinguishing State Dependence from Heterogeneity," *American Economic Review*, 81, 75-79.
- Heckman, J.J. and B. Singer (1984). "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data", *Econometrica*, 52, 271-320.
- Heckman, J.J and C. Taber (1994). "Econometric Mixture Models and More General Models for Unobservables in Duration Analysis", mimeo.
- Israel, M. (2001). "Identification of True vs. Spurious Tenure Dependence from Pricing Patterns", mimeo.
- Israel, M. (2002). "Do Consumers Anticipate Future Price Changes? Evidence from Automobile Insurance", mimeo.
- Israel, M. (2003), "Consumer Learning about Established Firms: Evidence from Automobile Insurance," mimeo.
- Jovanovic, B. (1982), "Selection and the Evolution of Industry," *Econometrica*, 50, 649-670.
- Keane, M.P. (1997). "Modeling Heterogeneity and State Dependence in Consumer Choice Behavior," *Journal of Business and Economic Statistics*, 15, 310-327.
- Klemperer, P.L. (1995). "Competition when Consumers have Switching Costs" An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade", *Review of Economic Studies*, 62, 515-539.
- Knittel, C. R. (1997). "Interstate Long Distance Rates: Search Costs, Switching Costs, and Market Power", *Review of Industrial Organization*, 12, 519-536.

- Lancaster, T (1990). *The Econometric Analysis of Transition Data* (Cambridge: Cambridge University Press).
- Moshkin, N. and R. Shachar (2000). "Switching Costs or Search Costs?", mimeo.
- Schlesinger, H. and J.M. Graf von der Schulenburg (1989). "Search Costs, Switching Costs, and Product Heterogeneity in an Insurance Market", *The Journal of Risk and Insurance*, 58, 109-119.
- Schlesinger, H. and J.M. Graf von der Schulenburg (1991). "Consumer Information and Decisions to Switch Insurers", *The Journal of Risk and Insurance*, 60, 591-615.
- Shachar, R. and J.W. Emerson (2000). "Cast Demographics, Unobserved Segments, and Heterogeneous Switching Costs in a Television Viewing Choice Model", *Journal of Marketing Research*, 37, 173-186.
- Sharpe, S.A. (1997). "The Effect of Consumer Switching Costs on Prices: A Theory and its Application to the Bank Deposit Market", *Review of Industrial Organization*, 12, 79-94.
- Shum, M. (2000). "Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast Cereals Market", mimeo.
- van den Berg, G. and J. van Ours (1996). "Unemployment Dynamics and Duration Dependence", *Journal of Labor Economics*, 14, 100-125.
- Viard, V. B. (2002). "Do Switching Costs Make Markets More or Less Competitive? The Case of 800-Number Portability", mimeo.

Appendix I: Notes on Identification

Throughout this Appendix, we simplify notation by ignoring η and referring to all variables other than tenure and match quality as x_t . Then, as noted in Section III, the two key functions are the structural departure probability conditional on a value for ε , $D(\tau_t, x_t | \varepsilon)$, and the updated probability of each match quality, $\rho_j^t(x^t)$. This appendix provides additional details on how the individual-specific histories help in identification of the model developed in Section IV.⁴⁵

If our interest was solely in the impact of observables, x_t , on the structural departure probability, the challenge would be lessened. We could include very general controls for the impact of x^t , and then rely on the independent variation in x_t to identify the effects of interest from an equation like 3.2. However, we seek to identify tenure effects. And fully general controls for x^t implicitly control for tenure with the firm through t , so this approach will not work. Rather, we need to explicitly include an error term in the departure function, model the heterogeneity distribution, and then rely on the impact of x^t on the departure probability to recover this distribution. We can then measure the “amount” of selection which has occurred for any history. True tenure effects can then be recovered, essentially by considering consumers with different tenures, but also with different histories, which together imply the same accrued amount of selection.

Note that selection effects are driven by the fact that the survival probability, $1 - D(\tau_t, x_t | \varepsilon)$ is higher for higher values of ε , increasing according to the derivative:

$$\frac{\partial(1 - D(\tau_t, x_t | \varepsilon))}{\partial \varepsilon} \tag{A1.1}$$

And the variation in selection with x^t is driven by the fact that this derivative varies with x according to:

$$\frac{\partial \frac{\partial(1 - D(\tau_t, x_t | \varepsilon))}{\partial \varepsilon}}{\partial x} \tag{A1.2}$$

So information about this derivative is the key piece in a mapping from the observable variation in selection with x^t to a heterogeneity distribution. The remainder of this section discusses sources of that information.

The simplest source – quite similar to the restrictions placed by the proportional hazard framework – would be to specify a restrictive functional form for $1 - D(\tau_t, x_t | \varepsilon)$ which determines this derivative. For example, if we simplified the model of section IV to assume that consumers

⁴⁵ We are not proving non-parametric identification results here. Rather, our goal is simply to indicate how the data are informative.

receive exactly one offer each period, then the model would basically take the form $1 - D(\tau_t, x_t | \varepsilon) = \Phi(f(\tau_t) - \alpha(x_t) + \varepsilon)$, so that:

$$\frac{\partial \frac{\partial(1 - D(\tau_t, x_t | \varepsilon))}{\partial \varepsilon}}{\partial x} = \frac{\partial \phi(f(\tau) - \alpha(x_t) + \varepsilon)}{\partial x_t} \quad (\text{A1.3})$$

In this case, the departure probability directly implies a position in the normal distribution, which in turn determines the value of this derivative as the slope of the PDF. So, in this case, while variation in the departure probability with x^t would still be informative about the heterogeneity distribution, the information would rest heavily on the normality assumption.

Under the more general model estimated, however, the survival probability is relaxed so that:

$$1 - D(\tau_t, x_t | \varepsilon) = [\Phi(f(\tau_t) - \alpha(x_t) + \varepsilon)]^{\Omega_t} \quad (\text{A1.4})$$

In essence, this means that the departure probability does not pin down a position in the normal distribution. In words, a given departure probability can be generated by a more favorable position in the distribution with more offers, or a less favorable position with fewer offers. So, the selection effects are not so heavily dependent on functional form restrictions.

Fortunately, there is information in the data to allow us to estimate this more general model. This information relies on the maintained assumption that the x_t and ε enter linearly. So, simplifying to 2 observables, we can write the departure probability as: $D(f(\tau_t) - \alpha x_t + \beta y_t + \varepsilon)$. This means that the derivative of interest is proportional to the derivative between two observables:

$$\frac{\partial \frac{\partial(1 - D(\tau_t, x_t | \varepsilon))}{\partial y}}{\partial x} \quad (\text{A1.5})$$

Hence, we can learn about the unobservable derivative of interest through the impact of one observable on the sensitivity of the departure probability to other observables, thus allowing the data to guide identification in this more flexible functional form.

Appendix II: Calculation of Expected Price Changes

This appendix details the calculation of the expected price change term, $\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta) - \Delta_c(k_t, L_{t-1}, 0, \eta)$. The key maintained assumption in this calculation is that expected price changes are made up of three additively separable components, representing changes in k , L , and τ^{nc} .⁴⁶

$$\Delta_c(k_t, L_{t-1}, \tau_t^{nc}, \eta) = \Delta_{kc}(k_t) + \Delta_{Lc}(L_{t-1}; k_t, \eta) + \Delta_{\tau c}(\tau_t^{nc}; k_t, \eta) \quad (\text{A2.1})$$

The only variable in these expressions which differs between the current and potential new firm is the accrued claims free tenure at the current firm. So, the first two terms cancel, and the calculation reduces to computing the key relationship factor – the benefit to remaining at the current firm to achieve tenure based price discounts – measured by $\Delta_{\tau c}(\tau_t^{nc}; k_t, \eta) - \Delta_{\tau c}(0; k_t, \eta)$.

Computation of this difference relies on the known, tenure-based price-discount structure, described above. The expected price-discounts after 6 and 12 periods are denoted δ_1 and δ_2 , respectively.⁴⁷ The full term, $\Delta_{\tau c}(\tau_t^{nc}; k_t, \eta)$, is defined as the average of the expected, discounted price changes for all future periods, where the average is computed by dividing by the sum of the time-discount factors, $1/(1-\beta)$. To illustrate the calculation of $\Delta_{\tau c}(\tau_t^{nc}; k_t, \eta) - \Delta_{\tau c}(0; k_t, \eta)$, suppose $\tau_t^{nc}=2$. Then, at the current firm, the consumer expects a discount of δ_1 from $t+5$ on, and a discount of δ_2 from $t+11$ on, if he has no chargeable claims. At an alternate firm, he expects these discounts from $t+7$ on, and $t+13$ on, respectively, again if he has no chargeable claims in the interim. If he has a chargeable claim at any point, τ^{nc} returns to 0 and the difference between the two firms is eliminated.

To explicitly write down this difference, we make one final simplification -- assuming that the consumer uses the current probability of a chargeable claim, $\lambda(k_t; L_0 | \eta)$, for all future periods. Then we have:

$$\begin{aligned} & \Delta_{\tau c}(2; k_t, \eta) - \Delta_{\tau c}(0; k_t, \eta) = \\ & (1-\beta)(\beta^5(1-\lambda(k_t, L_0 | \eta))^4 \delta_1 + \beta^6(1-\lambda(k_t, L_0 | \eta))^5 \delta_1 + \\ & \beta^{11}(1-\lambda(k_t, L_0 | \eta))^{10} \delta_2 + \beta^{12}(1-\lambda(k_t, L_0 | \eta))^{11} \delta_2) \end{aligned} \quad (\text{A2.2})$$

Computation for all other values of τ^{nc} follows accordingly.

⁴⁶ Note that expected changes in L and τ depend on the probability of claims, which in turn depends on k and η . So k and η are included in the second and third terms of this expression.

⁴⁷ For simplicity, we are estimating a constant expected discount for all consumers. In fact, the actual discounts vary by consumer type, but consumers are assumed to have knowledge only of the average discount, not the full menu of discounts by consumer type and thus not the discount which will be in effect when they reach the relevant tenure. This leaves the problem tractable while capturing the key feature – consumers may remain with their current firm to achieve discounts, rather than for any relationship value or match quality.

$\lambda(k_i; L_0 | \eta)$ is specified as a Normal CDF, with an index that is linear in k , L , and η , and estimated jointly with the departure probability using the likelihood function given in equations 4.13 – 4.15. The estimated parameters are:

Parameter	Variable Description	Estimate	Standard Error
γ_1	Intercept	-2.218	0.063
γ_2	Purchases Collision Coverage	0.265	0.031
γ_3	Placement In Alternate Company	0.118	0.018
γ_4	Primary Driver age 65 or over	0.063	0.104
γ_5	Occasionally Driven by Unmarried Female Under 25	0.386	0.133
γ_6	Primary Driver Unmarried Female Under 25	0.276	0.090
γ_7	Primary Driver Married Male Under 25	0.155	0.113
γ_8	Occasionally Driven by Unmarried Male Under 25	0.349	0.122
γ_9	Primary Driver Married Male Under 26	0.389	0.092
γ_{10}	Driver Under 21	0.271	0.054
γ_{11}	Vehicle Used for Commute to Work	0.046	0.015
γ_{12}	Vehicle Driven 7,500+ Miles per Year	0.150	0.065
γ_{13}	Vehicle less than 1 Year Old	0.122	0.039
γ_{14}	Vehicle 8+ Years Old	-0.099	0.030
γ_{15}	Compact Car	0.113	0.044
γ_{16}	Light Truck/Recreational Vehicle	0.063	0.034
γ_{17}	Heavy Truck	0.001	0.031
γ_{18}	(Vehicles per Square Mile in Zip Code/1000)	0.042	0.015
γ_{19}	(Vehicles per Square Mile in Zip Code/1000) ²	-0.003	0.002
γ_{20}	Vehicle Garaged in Atlanta	0.032	0.022
γ_{21}	Claim in 3 Years Prior to Joining Firm	0.047	0.035
ρ_η	Probability of High Risk Error	0.261	0.090
$\bar{\eta}$	High Risk Error	0.574	0.082
$\underline{\eta}$	Low Risk Error	-0.203	0.051

Appendix III: Calculation of Market Average Prices

The market average price for each class is defined as the average price among the member firms of the Insurance Services Organization (ISO) for fiscal 1994-1998 by territory, rating class, and claims history. Average prices are computed separately by coverage. Each average is computed as total earned premiums divided by the total number of policies with the specified coverage in force

among the member firms. The average price facing a given consumer is the sum of the prices for the coverages he currently purchases.

Average prices are assumed to be stable for the entire study period, save for growth affecting all firms measured by the motor vehicle insurance CPI. So, premiums are converted to their 1994 equivalents by deflating by this CPI, and all the premium and exposure data are combined for all years. The value for each year is then re-inflated by the appropriate motor vehicle insurance CPI. Finally, the insurance price is converted to real dollars by deflating by the overall CPI.

While this measure is better than any used in existing insurance literature, it still faces several limitations. First, ISO firms sell only roughly 25% of the policies in Georgia not sold by the study firm. Second, the computed price refers to all policies in force, rather than just the relevant policies for new consumers. So, for all consumer types, the computed market average price is assumed to be proportional to the average for new consumers in the overall market. The factor of proportionality, μ_0 , is estimated as part of the likelihood function, and reported in the table below.

In addition, the available data fail to include several sanctioned rating factors. These are:

1. Whether the vehicle is driven over 7,500 miles per year.
2. Whether the consumer has completed a defensive driving course.
3. Whether the driver is a student who commutes over 150 miles to college.
4. Whether the vehicle is used on a farm.
5. Whether the vehicle is used primarily in the operation of a business.
6. Whether a consumer between 21-24 is 21-22 or 23-24 years old.

Together, these account for less than 1% of the variance in prices at the study firm, so the impact of excluding them should be small. But to be complete, parameters are included to adjust for these omitted classifications. To accomplish this, each excluded factor is assumed to have a multiplicative impact on price that applies to all observed consumer types at all firms. For example, consider the impact of mileage. The observed average price is a combination of high and low mileage consumers:

$$\tilde{\mu}_p = q_1 \mu_1 \underline{p} + (1 - q_1) \underline{p} \quad (\text{A3.1})$$

where $\tilde{\mu}_p$ is the observed average price, q_1 is the percentage of high mileage consumers, μ_1 is the multiplicative price increase for high mileage consumers, and \underline{p} is the average price for a low mileage consumer.⁴⁸ Solving for \underline{p} yields the average price for low-mileage consumers:

$$\underline{p} = \frac{\tilde{\mu}_p}{q_1 \mu_1 + (1 - q_1)} \quad (\text{A3.2})$$

⁴⁸ In this discussion, we are assuming that the observed average price is already adjusted for the base proportionality factor, μ_0 . Also note the implicit assumption that mileage status is uncorrelated with all other rating categories, both observed and unobserved. In this way, the only difference between the average price for low and high mileage consumers results from the multiplicative pricing factor, μ_1 .

The price for high-mileage consumers, then, is simply $\mu_1 p$. We assume that the proportion of high-mileage consumers observed at the study firm applies to the market as a whole, so q_1 is observable. This leaves only μ_1 to be estimated as part of the likelihood function. The same logic is followed for $\mu_2 - \mu_6$ as well.

The estimated market price adjustment factors are:

Parameter	Classification	Estimate	Standard Error
μ_0	Proportionality Factor	0.891	0.112
μ_1	High Mileage	1.281	0.231
μ_2	Defensive Driving	0.967	0.441
μ_3	Commuting Student	1.155	0.412
μ_4	Farm	0.821	0.412
μ_5	Business	1.254	0.406
μ_6	23-24	0.929	0.315

Figure 1.1: Departure Probability by Tenure




Figure 3.1: Departure Probability -- Men vs. Women




Figure 3.2: Departure Probability -- Over 55 vs. Under 55




Figure 3.3: Departure Probability -- Atlanta vs. Rest of State




Figure 5.1: Estimated Tenure Effects




Figure 5.2: Predicted Departure Probabilities

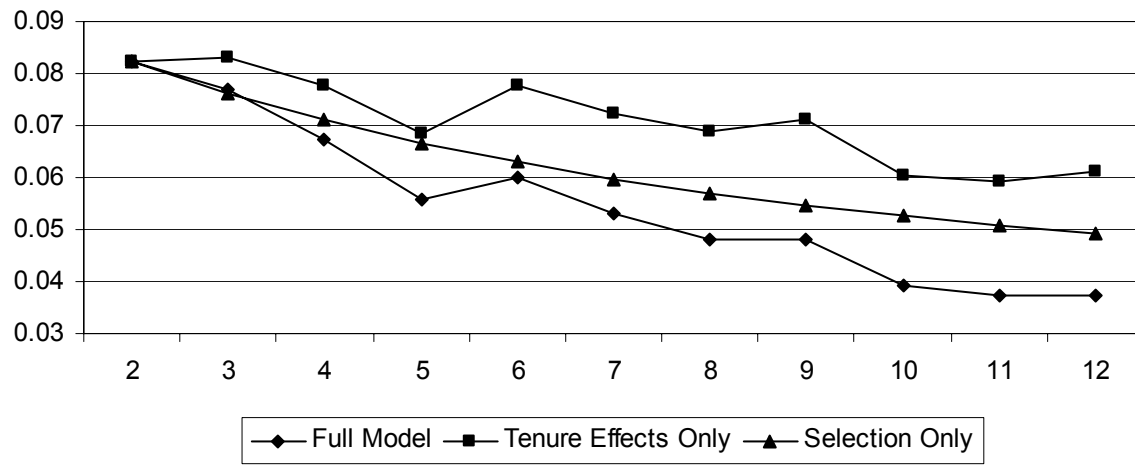


Table 3.1: Entry and Exit from Sample

	Count	Percent
Total Policies	18,595	100.0
<u>Year of Entry</u>		
1991	509	2.7
1992	2,336	12.6
1993	2,475	13.3
1994	2,648	14.2
1995	2,709	14.6
1996	2,820	15.2
1997	2,642	14.2
1998	2,456	13.2
<u>Cause of Exit</u>		
Voluntary Departure	5,854	31.5
Right Censored at Dec 1998	7,976	42.9
Randomly Censored before Dec 1998	4,765	25.6

Table 3.2: Descriptive Statistics for Price

Statistic	Study Firm Price	Market Average Price
Mean	182.38	202.77
Standard Deviation	92.37	113.13
25 th Percentile	131.94	135.99
50 th Percentile	158.56	171.33
75 th Percentile	198.55	225.49

Table 3.3: Frequency Tabulation of Discrete Variables (Page 1)

Variable	Value	Frequency	Percent
Total Consumer-Periods	N/A	86,561	100.0
Tenure (τ)	1	19,225	22.2
	2	14,484	16.7
	3	11,640	13.5
	4	9,292	10.7
	5	7,484	8.7
	6	6,099	7.1
	7	4,864	5.6
	8	3,800	4.4
	9	2,956	3.4
	10	2,281	2.6
	11	1,718	2.0
	≥ 12	2,718	3.1
Claims Free Tenure (τ^{nc})	0	3,305	3.8
	1	21,033	24.3
	2	15,127	17.5
	3	11,666	13.5
	4	8,895	10.3
	5	6,875	7.9
	6	5,375	6.2
	7	4,101	4.7
	8	3,081	3.6
	9	2,295	2.7
	10	1,705	2.0
	11	1,244	1.4
	12	895	1.0
	≥ 13	964	1.1

Table 3.3: Frequency Tabulation of Discrete Variables (Page 2)

Variable	Value	Frequency	Percent
Limit Class (<i>y</i>)	0	23,253	26.9
	1	27,054	31.3
	2	25,219	29.1
	3	11,035	12.8
Multiple Policies (<i>MCAR</i>)	0	18,955	21.9
	1	67,606	78.1
First Party Coverage (<i>FP</i>)	0	30,117	34.8
	1	56,444	65.2
Claim in Previous Period (<i>ICLM</i>)	0	80,694	93.2
	1	5,867	6.8
Agent Class (<i>AGT</i>)	0	20,469	23.7
	1	23,255	26.9
	2	17,510	20.2
	3	25,327	29.3
Youthful (<i>YNG</i>)	0	77,130	89.1
	1	9,431	10.9
Commute to Work (<i>WK</i>)*	0	36,106	46.8
	1	41,024	53.2
Good Student (<i>GS</i>)**	0	7,396	78.4
	1	2,035	21.6

* Only defined for 77,130 non-youthful drivers.

** Only defined for 9,431 youthful drivers.

Table 3.4: Impact of Past Prices

Variable	Coefficient	Standard Error
Current Relative Price (\$100)	0.612	0.111
Relative Price at t-1 (\$100)	-0.103	0.052
Relative Price at t-2 (\$100)	-0.141	0.038
Relative Price at t-3 (\$100)	-0.112	0.040
Tenure = 4	-1.291	0.031
Tenure = 5	-1.364	0.034
Tenure = 6	-1.299	0.035
Tenure = 7	-1.351	0.044
Tenure = 8	-1.375	0.041
Tenure = 9	-1.360	0.046
Tenure = 10	-1.457	0.050
Tenure = 11	-1.472	0.048
Tenure \geq 12	-1.463	0.053

Table 5.1: Non-Relationship Parameter Estimates

Category	Parameter	Description	Estimate	Standard Error
Sensitivity to Current Prices (\$100)	$\alpha_1(0)$	Limit Class 0	0.671	0.068
	$\alpha_1(1)$	Limit Class 1	0.585	0.059
	$\alpha_1(2)$	Limit Class 2	0.551	0.064
	$\alpha_1(3)$	Limit Class 3	0.574	0.095
Standard Error of $V^{p,\xi}$	$\sigma(1)$	Limit Class 1	0.961	0.141
	$\sigma(2)$	Limit Class 2	0.942	0.129
	$\sigma(3)$	Limit Class 3	0.959	0.222
Sensitivity to Future Prices (\$100)	$\alpha_3(0)$	Limit Class 0	0.599	0.146
	$\alpha_3(1)$	Limit Class 1	0.501	0.121
	$\alpha_3(2)$	Limit Class 2	0.490	0.129
	$\alpha_3(3)$	Limit Class 3	0.512	0.203
Switching Costs	α_5	Multiple Cars	0.271	0.045
	α_6	First-Party Coverage	0.055	0.026
	α_7	Claim Last Period	0.291	0.113
Offer Arrival	β_0	Constant	0.711	0.088
	$\beta_1(1)$	Agent Class 1	0.089	0.031
	$\beta_1(2)$	Agent Class 2	0.102	0.040
	$\beta_1(3)$	Agent Class 3	0.118	0.028
	β_2	Youthful	0.114	0.440
	β_3	Commute to Work	-0.204	0.033
	β_4	Good Student	-0.101	0.051

Recall that we have normalized $(\xi_c - \mu_\xi)$ and α_4 to 0.

Table 5.2 Relationship Parameter Estimates

Category	Parameter	Description	Estimate	Standard Error
True Tenure Effects	$\alpha_2(1)$	Constant Term	1.161	0.007
	$\tilde{\alpha}_2(2)$	Marginal Effect of Tenure 2	-0.139	0.010
	$\tilde{\alpha}_2(3)$	Marginal Effect of Tenure 3	-0.005	0.015
	$\tilde{\alpha}_2(4)$	Marginal Effect of Tenure 4	0.037	0.018
	$\tilde{\alpha}_2(5)$	Marginal Effect of Tenure 5	0.071	0.019
	$\tilde{\alpha}_2(6)$	Marginal Effect of Tenure 6	-0.070	0.024
	$\tilde{\alpha}_2(7)$	Marginal Effect of Tenure 7	0.040	0.027
	$\tilde{\alpha}_2(8)$	Marginal Effect of Tenure 8	0.027	0.030
	$\tilde{\alpha}_2(9)$	Marginal Effect of Tenure 9	-0.018	0.034
	$\tilde{\alpha}_2(10)$	Marginal Effect of Tenure 10	0.085	0.043
	$\tilde{\alpha}_2(11)$	Marginal Effect of Tenure 11	0.010	0.044
	$\tilde{\alpha}_2(12)$	Marginal Effect of Tenure 12	-0.015	0.039
Heterogeneity Distribution	ε_1	Points of Support	-0.955	0.104
	ε_2		-0.131	0.111
	ε_3		-0.039	0.030
	ε_4		0.150	0.108
	ε_5		0.806	0.131
	ρ_1	Probabilities	0.143	0.044
	ρ_2		0.362	0.071
	ρ_3		0.141	0.040
	ρ_4		0.146	0.039
	ρ_5		0.208	0.063
Price Discount Expectations (\$100)	δ_1	Price Change after 6 periods	-0.171	0.041
	δ_2	Price Change after 12 periods	-0.111	0.056
	β	Discount Factor	0.564	0.099

Standard Error of ε_3 and ρ_3 computed via the Delta Method.

Table 5.3: Impact of Expected Price Discounts

τ^{nc}	Expected Price Change (\$) (Relative to $\tau^{nc} = 0$)	Departure Probability
0	0.00	0.071
1	-0.25	0.071
2	-0.68	0.070
3	-1.45	0.070
4	-2.82	0.069
5	-5.24	0.068
6	-9.53	0.065
7	-0.04	0.071
8	-0.32	0.071
9	-0.81	0.070
10	-1.67	0.070
11	-3.21	0.069
12	-5.94	0.067
13	0.32	0.071

The expected price change is the discounted average of all expected, future, tenure-based price discounts -- relative to this average at $\tau^{nc} = 0$ -- computed as described in Appendix II. Departure probabilities are computed with tenure equal to 1, price equal to the market average, and all other variables set to sample averages.

Table 5.4: Cumulative Utility Changes

Period	Change in Tenure Effect	Dollar Equivalent	Change in Expected Value of ε	Dollar Equivalent
2 (Base Value)	1.022	171.07	0.035	5.93
3	-0.005	-0.84	0.039	6.41
4	0.032	5.36	0.073	12.10
5	0.103	17.24	0.101	16.89
6	0.033	5.52	0.125	20.76
7	0.073	12.22	0.148	24.65
8	0.100	16.74	0.168	27.99
9	0.082	13.73	0.185	30.94
10	0.167	27.95	0.202	33.77
11	0.177	29.63	0.216	36.08
12	0.162	27.12	0.229	44.17

All computations are for a “base case” consumer – with price equal to the market average, no accrued claims free tenure, and all other variables at their sample means. Dollar equivalents are found by dividing by the sample average of $\alpha_1(y_t)$.

Table 5.5: Departure Probabilities by Type

τ	Departure Probability by Type					Updated Probability of Type					Overall Departure Prob
	ε_1	ε_2	ε_3	ε_4	ε_5	ρ_1^ε	ρ_2^ε	ρ_3^ε	ρ_4^ε	ρ_5^ε	
1	0.232	0.064	0.054	0.037	0.007	0.143	0.362	0.141	0.146	0.208	0.071
2	0.273	0.083	0.070	0.049	0.011	0.118	0.365	0.144	0.151	0.222	0.082
3	0.275	0.084	0.071	0.049	0.011	0.094	0.364	0.145	0.157	0.240	0.077
4	0.263	0.078	0.066	0.046	0.010	0.074	0.362	0.146	0.162	0.257	0.067
5	0.242	0.069	0.058	0.039	0.008	0.058	0.357	0.147	0.165	0.272	0.056
6	0.263	0.078	0.066	0.046	0.010	0.047	0.353	0.146	0.168	0.286	0.060
7	0.251	0.073	0.061	0.042	0.009	0.037	0.346	0.145	0.171	0.301	0.053
8	0.243	0.069	0.058	0.040	0.008	0.029	0.339	0.144	0.173	0.316	0.048
9	0.248	0.071	0.060	0.041	0.009	0.023	0.331	0.143	0.174	0.329	0.048
10	0.224	0.061	0.051	0.034	0.007	0.018	0.323	0.141	0.176	0.342	0.039
11	0.221	0.060	0.050	0.034	0.007	0.015	0.316	0.139	0.177	0.354	0.037
12	0.225	0.061	0.051	0.035	0.007	0.012	0.308	0.137	0.177	0.365	0.037

All computations are for a “base case” consumer.

Table 5.6: Cumulative Decline in Departure Probability by Source

Period	Total Change	Tenure Effects Alone	Percent of Change	Selection Alone	Percent of Change
2 (Base)	0.0824	0.0824	N/A	0.0824	N/A
2→3	-0.0056	0.0007	-12.5%	-0.0062	110.7%
2→4	-0.0152	-0.0046	30.3%	-0.0114	75.0%
2→5	-0.0266	-0.0141	53.0%	-0.0158	59.4%
2→6	-0.0225	-0.0048	21.3%	-0.0195	86.7%
2→7	-0.0294	-0.0102	34.7%	-0.0226	76.9%
2→8	-0.0342	-0.0137	40.1%	-0.0253	74.0%
2→9	-0.0344	-0.0114	33.1%	-0.0277	80.5%
2→10	-0.0431	-0.0219	50.8%	-0.0297	68.9%
2→11	-0.0450	-0.0230	51.1%	-0.0316	70.2%
2→12	-0.0450	-0.0212	47.1%	-0.0332	73.8%

All computations are for a base case consumer. Tenure effects alone are computed by fixing $\varepsilon = -0.1275$ to match the period 2 departure probability from the full model. Selection alone is computed by setting $\alpha_1(\tau) = \alpha_1(2)$ for all $\tau \geq 2$.

Table 5.7: Change in Departure Probability by Price History

Period	Previous Relative Prices = -\$100	Previous Relative Prices = \$0	Previous Relative Prices = \$100
1 (Base)	0.0710	0.0710	0.0710
1 → 2	0.0149	0.0114	0.0059
2→3	0.0126	0.0058	-0.0036
2→4	0.0052	-0.0038	-0.0150
2→5	-0.0053	-0.0152	-0.0266
2→6	0.0008	-0.0111	-0.0244
2→7	-0.0057	-0.0180	-0.0312
2→8	-0.0103	-0.0228	-0.0359
2→9	-0.0098	-0.0230	-0.0369
2→10	-0.0196	-0.0327	-0.0442
2→11	-0.0215	-0.0336	-0.0463
2→12	-0.0211	-0.0336	-0.0470

Computations are for a base case consumer, with the single exception that they have faced the specified relative price in all previous periods (in the current period, relative price is 0).

Table 5.8: Estimates without Heterogeneity or Without Tenure Effects

	Base Model	No Heterogeneity	No Tenure Effects
<u>Tenure Effects</u> Total Increase Period 2 – 12	0.162	0.406	0
$\alpha_2(2)$	1.022	0.895	1.091
$\tilde{\alpha}_2(3)$	-0.005	0.039	0
$\tilde{\alpha}_2(4)$	0.037	0.072	0
$\tilde{\alpha}_2(5)$	0.071	0.098	0
$\tilde{\alpha}_2(6)$	-0.070	-0.037	0
$\tilde{\alpha}_2(7)$	0.040	0.062	0
$\tilde{\alpha}_2(8)$	0.027	0.047	0
$\tilde{\alpha}_2(9)$	-0.018	0.004	0
$\tilde{\alpha}_2(10)$	0.085	0.097	0
$\tilde{\alpha}_2(11)$	0.010	0.023	0
$\tilde{\alpha}_2(12)$	-0.015	0.001	0
<u>Heterogeneity</u> <u>Distribution</u> $\sigma(\varepsilon)$	0.525	0	0.598
ε_1	-0.955	N/A	-0.949
ε_2	-0.131	N/A	-0.265
ε_3	-0.039	N/A	0.123
ε_4	0.150	N/A	0.211
ε_5	0.806	N/A	0.841
ρ_1	0.143	N/A	0.171
ρ_2	0.362	N/A	0.329
ρ_3	0.141	N/A	0.111
ρ_4	0.146	N/A	0.145
ρ_5	0.208	N/A	0.244

Table 5.9: Price Parameters with no Heterogeneity

Category	Parameter	New Estimate	New Standard Error	Base Estimate	Base Standard Error
Sensitivity to Current Prices (\$100)	$\alpha_1(0)$	0.489	0.066	0.671	0.068
	$\alpha_1(1)$	0.432	0.061	0.585	0.059
	$\alpha_1(2)$	0.404	0.063	0.551	0.064
	$\alpha_1(3)$	0.419	0.080	0.574	0.095

Table 5.10: Firm's Willingness to Pay for Higher Tenure In "Markup" Units

Tenure	Base Tenure Effects	"Full" Tenure Effects	Base as Percentage of Full
3	0.16	0.48	33.3%
4	0.35	0.93	37.6%
5	0.49	1.28	38.3%
6	0.53	1.48	35.8%
7	0.69	1.78	38.8%
8	0.79	1.99	39.7%
9	0.87	2.14	40.7%
10	0.99	2.30	43.0%
11	0.98	2.33	42.1%
12	0.96	2.33	41.2%

Computations are for a base case consumer with ϵ set to 0.

Table 6.1: Extension 1
Heterogeneity in Tenure Effects Across “Companies”

Category	Parameter	New Estimate	New Standard Error	Original Estimate	Original Standard Error
Alternate Company Effect	δ_0	0.039	0.112	N/A	N/A
	δ_1	-0.014	0.010	N/A	N/A
True Tenure Effects	$\alpha_2(1)$	1.168	0.009	1.161	0.007
	$\tilde{\alpha}_2(2)$	-0.142	0.012	-0.139	0.010
	$\tilde{\alpha}_2(3)$	-0.006	0.021	-0.005	0.015
	$\tilde{\alpha}_2(4)$	0.035	0.019	0.037	0.018
	$\tilde{\alpha}_2(5)$	0.067	0.022	0.071	0.019
	$\tilde{\alpha}_2(6)$	-0.073	0.025	-0.070	0.024
	$\tilde{\alpha}_2(7)$	0.039	0.030	0.040	0.027
	$\tilde{\alpha}_2(8)$	0.026	0.032	0.027	0.030
	$\tilde{\alpha}_2(9)$	-0.020	0.037	-0.018	0.034
	$\tilde{\alpha}_2(10)$	0.082	0.044	0.085	0.043
	$\tilde{\alpha}_2(11)$	0.010	0.044	0.010	0.044
	$\tilde{\alpha}_2(12)$	-0.019	0.043	-0.015	0.039
Heterogeneity Distribution	ε_1	-0.950	0.171	-0.955	0.104
	ε_2	-0.130	0.189	-0.131	0.111
	ε_3	0.049	N/A	-0.039	N/A
	ε_4	0.141	0.212	0.150	0.108
	ε_5	0.802	0.230	0.806	0.131
	ρ_1	0.150	0.091	0.143	0.044
	ρ_2	0.361	0.130	0.362	0.071
	ρ_3	0.141	N/A	0.141	N/A
	ρ_4	0.146	0.084	0.146	0.039
	ρ_5	0.202	0.091	0.208	0.063

**Table 6.2: Extension 2
Heterogeneity in Price Parameters**

Category	Parameter	New Estimate	New Standard Error	Original Estimate	Original Standard Error
Price Sensitivity Error	ρ_v	0.391	0.244	N/A	N/A
	v_1	0.461	0.388	N/A	N/A
	v_0	-0.296	N/A	N/A	N/A
	$\sigma_v(v_1)$	1.213	0.434	N/A	N/A
True Tenure Effects	$\alpha_2(1)$	1.170	0.009	1.161	0.007
	$\tilde{\alpha}_2(2)$	-0.144	0.012	-0.139	0.010
	$\tilde{\alpha}_2(3)$	-0.006	0.021	-0.005	0.015
	$\tilde{\alpha}_2(4)$	0.033	0.019	0.037	0.018
	$\tilde{\alpha}_2(5)$	0.068	0.022	0.071	0.019
	$\tilde{\alpha}_2(6)$	-0.074	0.025	-0.070	0.024
	$\tilde{\alpha}_2(7)$	0.043	0.030	0.040	0.027
	$\tilde{\alpha}_2(8)$	0.031	0.032	0.027	0.030
	$\tilde{\alpha}_2(9)$	-0.019	0.037	-0.018	0.034
	$\tilde{\alpha}_2(10)$	0.082	0.044	0.085	0.043
	$\tilde{\alpha}_2(11)$	0.012	0.044	0.010	0.044
	$\tilde{\alpha}_2(12)$	-0.015	0.043	-0.015	0.039
Heterogeneity Distribution	ε_1	-0.952	0.171	-0.955	0.104
	ε_2	-0.133	0.189	-0.131	0.111
	ε_3	0.006	N/A	-0.039	N/A
	ε_4	0.144	0.212	0.150	0.108
	ε_5	0.811	0.230	0.806	0.131
	ρ_1	0.151	0.091	0.143	0.044
	ρ_2	0.358	0.130	0.362	0.071
	ρ_3	0.138	N/A	0.141	N/A
	ρ_4	0.141	0.084	0.146	0.039
	ρ_5	0.212	0.091	0.208	0.063

Table 6.3: Extension 3
Search Following Class Changes

Category	Parameter	New Estimate	New Standard Error	Original Estimate	Original Standard Error
Information Arrival	β_0	0.503	0.095	0.711	0.088
	$\beta_1(1)$	0.085	0.033	0.089	0.031
	$\beta_1(2)$	0.098	0.038	0.102	0.040
	$\beta_1(3)$	0.111	0.029	0.118	0.028
	β_2	0.101	0.298	0.114	0.440
	β_3	-0.194	0.031	-0.204	0.033
	β_4	-0.093	0.049	-0.101	0.051
	β_5	0.441	0.331	N/A	N/A
	β_6	0.141	0.156	N/A	N/A
	True Tenure Effects	$\alpha_2(1)$	1.170	0.011	1.161
$\tilde{\alpha}_2(2)$		-0.146	0.012	-0.139	0.010
$\tilde{\alpha}_2(3)$		-0.006	0.023	-0.005	0.015
$\tilde{\alpha}_2(4)$		0.031	0.018	0.037	0.018
$\tilde{\alpha}_2(5)$		0.065	0.021	0.071	0.019
$\tilde{\alpha}_2(6)$		-0.078	0.028	-0.070	0.024
$\tilde{\alpha}_2(7)$		0.041	0.028	0.040	0.027
$\tilde{\alpha}_2(8)$		0.030	0.034	0.027	0.030
$\tilde{\alpha}_2(9)$		-0.022	0.035	-0.018	0.034
$\tilde{\alpha}_2(10)$		0.080	0.047	0.085	0.043
$\tilde{\alpha}_2(11)$		0.011	0.041	0.010	0.044
$\tilde{\alpha}_2(12)$		-0.016	0.046	-0.015	0.039
Heterogeneity Distribution	ε_1	-0.947	0.155	-0.955	0.104
	ε_2	-0.130	0.158	-0.131	0.111
	ε_3	0.006	N/A	-0.039	N/A
	ε_4	0.139	0.205	0.150	0.108
	ε_5	0.822	0.212	0.806	0.131
	ρ_1	0.145	0.082	0.143	0.044
	ρ_2	0.350	0.111	0.362	0.071
	ρ_3	0.135	N/A	0.141	N/A
	ρ_4	0.145	0.077	0.146	0.039
	ρ_5	0.225	0.080	0.208	0.063