

Essays on identification and semiparametric econometrics

by

Paul Schrimpf

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of


Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

ARCHIVES



© Paul Schrimpf, MMXI. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author
Department of Economics

Certified by
Victor Chernozhukov
Professor of Economics
Thesis Supervisor

Certified by
Whitney Newey
Jane Berkowitz Carlton and Dennis William Carlton Professor of Economics
Thesis Supervisor

Accepted by ..
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Chairman, Departmental Committee on Graduate Studies

Essays on identification and semiparametric econometrics

by

Paul Schrimpf

Submitted to the Department of Economics
on May 15, 2011, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This dissertation is a collection of three independent essays in theoretical and applied econometrics. The first chapter analyzes dynamic games with continuous states and controls. There are two main contributions. First, we give conditions under which the payoff function is nonparametrically identified by the observed distribution of states and controls. The identification conditions are fairly general and can be expected to hold in many potential applications. The key identifying restrictions include that one of the partial derivatives of the payoff function is known and that there is some component of the state space that enters the policy function, but not the payoff function directly. The second contribution of the first chapter is to propose a two-step semiparametric estimator for the model. In the first step the transition densities and policy function are estimated nonparametrically. In the second step, the parameters of the payoff function are estimated from the optimality conditions of the model. We give high-level conditions on the first step nonparametric estimates for the parameter estimates to be consistent and parameters to be \sqrt{n} -asymptotically normal. Finally, we show that a kernel based estimator satisfies these conditions.

The second chapter, which is coauthored with Liran Einav and Amy Finkelstein, analyzes the welfare cost of adverse selection in the U.K. annuity market. We develop a model of annuity contract choice and estimate it using data from the U.K. annuity market. The model allows for private information about mortality risk as well as heterogeneity in preferences over different contract options. We focus on the choice of length of guarantee among individuals who are required to buy annuities. The results suggest that asymmetric information along the guarantee margin reduces welfare relative to a first best symmetric information benchmark by about 2 percent of annuitized wealth. We also find that by requiring that individuals choose the longest guarantee period allowed, mandates could achieve the first-best allocation.

The third chapter develops a test for the exogeneity assumptions of classical factor models based on the fixed interactive effects estimator of Bai (2005). The exact form of the test is given for simple linear models. Simulations are used to assess the test's performance. The application of the test to more complicated models is also considered. The test is applied to a model of education as an example.

Thesis Supervisor: Victor Chernozhukov

Title: Professor of Economics

Thesis Supervisor: Whitney Newey

Title: Jane Berkowitz Carlton and Dennis William Carlton Professor of Economics

Acknowledgments

First and foremost, I thank my advisors, Victor Chernozhukov and Whitney Newey, and coauthors, Amy Finkelstein and Liran Einav. Without their guidance and support, this thesis would not have been possible. Victor provided invaluable guidance that helped me through some difficult stages of my research. Whitney's insight, patience, and encouragement has been invaluable. Working with Amy and Liran has been extraordinarily educational. I also thank the many people who have provided feedback on various parts of this dissertation. These include Arun Chandrasekhar, Denis Chetverikov, Brad Larsen, Anna Mikusheva, Jerry Hausman, and other participants at the MIT econometrics lunch. Finally, I am very grateful to my friends and family for their support and companionship.

Contents

1	Identification and estimation of continuous dynamic games	15
1.1	Introduction	16
1.2	Model	18
1.2.1	Example	19
1.3	Identification	19
1.3.1	Proof of identification (theorem 1)	23
1.3.2	Identifying transition densities with unobserved state variables	25
1.4	Estimation	27
1.4.1	Estimating policy functions	27
1.4.2	Estimating the payoff function	28
1.5	Asymptotic Theory	30
1.5.1	Asymptotic equivalence	43
1.5.2	A kernel estimator	43
1.6	Conclusion	49
2	Optimal Mandates and The Welfare Cost of Asymmetric Information	51
2.1	Introduction	52
2.2	Data and environment	57
2.3	Model: specification, identification, and estimation	61
2.3.1	A model of guarantee choice	61
2.3.2	Modeling heterogeneity	63
2.3.3	Identification	64
2.3.4	Parameterization	69
2.3.5	Estimation	72

2.4	Estimates and fit of the baseline model	75
2.4.1	Parameter estimates	75
2.4.2	Model fit	75
2.5	Welfare estimates	78
2.5.1	Measuring welfare	78
2.5.2	Welfare in observed equilibrium	80
2.5.3	The welfare cost of asymmetric information	81
2.5.4	The welfare consequences of government mandated annuity contracts	83
2.6	Robustness	85
2.6.1	Parameter choices	86
2.6.2	Wealth portfolio outside of the compulsory annuity market	87
2.6.3	Modeling heterogeneity	89
2.6.4	Imperfect information about mortality	92
2.6.5	Departing from the neoclassical model	94
2.6.6	Estimates for a different population	95
2.7	Conclusions	96
3	Testing for Fixed Interactive vs Random Interactive Effects	113
3.1	Introduction	113
3.2	Basic Framework	114
3.2.1	Random Interactive Effects	114
3.2.2	Fixed Interactive Effects	116
3.2.3	Testing for Interactive Random vs Fixed Effects	119
3.3	Simulations of test performance	120
3.4	Testing More Elaborate Models	125
3.4.1	Example: Estimating the Return to College	126
3.5	Conclusions	132
A	Appendix to chapter 2	133
A.1	Additional details about the data	133
A.2	Proof of Proposition 2	135
A.3	Additional details about estimation	136
A.3.1	Likelihood	136

A.3.2 Guarantee indifference curves 138

List of Figures

2.1	Schematic indifference sets	98
2.2	Estimated distributions	99
2.3	Wealth contours	100
2.4	Welfare change contours	101
3.1	Size under the null hypothesis	121
3.2	Power with $\psi = 0.25$	123
3.3	Power with $\psi = 0.5$	124
3.4	Estimated Age-Earnings Profiles	130

List of Tables

2.1	Summary statistics	102
2.2	Annuity payment rates	103
2.3	Parameter estimates	103
2.4	Within-sample fit	104
2.5	Out-of-sample fit	105
2.6	Welfare estimates	106
2.6	Robustness	109
3.1	MSE and Bias with $\psi = 0$	122
3.2	MSE and Bias with $\psi = 0.25$	122
3.3	MSE and Bias with $\psi = 0.5$	122
3.4	Summary statistics	128
3.5	Earnings Estimates	129
3.6	Specification Test	131

Chapter 1

Identification and estimation of dynamic games with continuous states and controls

Abstract. This chapter analyzes dynamic games with continuous states and controls. There are two main contributions. First, we give conditions under which the payoff function is nonparametrically identified by the observed distribution of states and controls. The identification conditions are fairly general and can be expected to hold in many potential applications. The key identifying restrictions include that one of the partial derivatives of the payoff function is known and that there is some component of the state space that enters the policy function, but not the payoff function directly. The latter of these restrictions is a standard exclusion restriction and is used to identify the payoff function off the equilibrium path. By manipulating the first order condition, we can show that the payoff function satisfies an integro-differential equation. Due to the presence of the value function in the first order condition, this integro-differential equation contains a Fredholm integral operator of the second kind. Invertibility of this operator, and knowledge of one of the partial derivatives of the payoff function is used to ensure that the integro-differential equation has a unique solution.

The second contribution of this chapter is to propose a two-step semiparametric estimator for the model. In the first step the transition densities and policy function are estimated nonparametrically. In the second step, the parameters of the payoff

function are estimated from the optimality conditions of the model. Because the state and action space are continuous, there is a continuum of optimality conditions. The parameter estimates minimize the norm of these conditions. Hence, the estimator is related to recent papers on GMM in Hilbert spaces and semiparametric estimators with conditional moment restrictions. We give high-level conditions on the first step nonparametric estimates for the parameter estimates to be consistent and parameters to be \sqrt{n} -asymptotically normal. Finally, we show that a kernel based estimator satisfies these conditions.

1.1 Introduction

This chapter analyzes dynamic games with continuous states and controls. There have been many recent papers about estimating dynamic games, but few, if any, of them allow for both continuous states and continuous controls. This is a useful gap to fill since many state and action variables in real applications are approximately continuous. A generic and pervasive example is investment as a decision variable and capital stock as a state variable. This chapter gives a comprehensive econometric analysis of continuous dynamic games. There are two main results. First, we give sufficient conditions for nonparametric identification of the payoff function from the observed transition densities. Second, we propose a two-step semiparametric estimator for the payoff function and give conditions for consistency and asymptotic normality.

The first main contribution of this chapter is to give sufficient conditions for the payoff function to be identified by the transition density of the state variables and controls. It is well known (see e.g. Rust (1994)), that without some restrictions, the payoff function is unidentified. However, in the case of dynamic games with discrete controls, Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2009) give plausible restrictions sufficient for identification. This chapter develops an analogous result for dynamic games with continuous controls. The key conditions for identification of the payoff function are that one of the partial derivatives of the payoff function is known, there is some value of the control such that value of the payoff function is known for that value of the control and all possible values of the state, and a certain integral operator of the second kind is invertible. The first two of these three conditions can be expected to hold in potential applications. We illustrate their plausibility with an example. The last of these three conditions is difficult to verify analytically, but could be checked in applications.

Our proof of identification involves manipulating the first order condition for the continuous control. The first order condition is an integro-differential equation—it involves both a derivative and an integral of the payoff function. After some manipulation, we transform the first order condition into an integral equation of the second kind with the derivative of the payoff function with respect to the control as the unknown function. We state conditions sufficient for this integral equation to have a unique solution.

Identification conditions for dynamic games with discrete controls can be found in Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2009). Blevins (2009) gives identification conditions for dynamic games with both discrete and continuous controls. His result relies on the presence of discrete controls and does not apply to the purely continuous case considered in this chapter.

Although our identification result is nonparametric, we develop a semiparametric estimator of the payoff function. Given the limited sample sizes typically available for dynamic games, a fully nonparametric estimator may not be informative. Therefore, like most papers on dynamic games, we assume that the payoff function is known up to a finite dimensional parameter, but leave the transition density and policy function nonparametrically specified. Our estimation procedure has two steps. First, the transition density and policy function are nonparametrically estimated. In the second step, the fact that at the true parameters the estimated policy function should maximize the value function is used to form an objective function. Since the state and action space are continuous, there is a continuum of optimality conditions. The parameter estimates minimize the norm of these conditions.

Our asymptotic results are stated at three levels of increasing specificity. First, we analyze generic two-step semiparametric minimum distance estimators with a continuum of estimating equations. We give high-level conditions on the initial nonparametric estimates and estimating equations sufficient for the finite dimensional parameters to be \sqrt{n} asymptotically normal. Next, we specialize this result to the estimating equations that arise from the optimality conditions of dynamic games, but remain agnostic about the specific form of the initial nonparametric estimates. Finally, we describe kernel estimates of the transition density and policy function that satisfy the conditions of the previous two theorems.

The remainder of this chapter proceeds as follows. Section 1.2 describes the model. Section 1.3 gives conditions for identification. Section 1.4 describes the estimation procedure. Section 1.5 contains the asymptotic analysis. Section 1.6 concludes.

1.2 Model

We consider a dynamic decision process with state variables, x , and policy variable, i . We will refer to the decision maker as a firm. The state variables follow a controlled Markov process, and have a density which is absolutely continuous with respect to Lebesgue measure, i.e.

$$f(x_{t+1}|\mathcal{I}_t) = f_{x'|x,i}(x_{t+1}|x_t, i_t),$$

where \mathcal{I}_t is all information available at time t . Each period a firm's payoff function depends on the current state, the firm's actions, and an iid private shock, η . We denote the payoff function by $\Pi(x, i, \eta)$. The firm's value function is:

$$V(x_0, \eta_0) = \max_{\{i_t = I_t(x_t, \eta_t)\}_{t=0}^{\infty}} E \left[\sum_{t=0}^{\infty} \delta^t \Pi(x_t, i_t, \eta_t) | x_0, i_0, \eta_0 \right].$$

where $I_t(x_t, \eta_t)$ denotes the policy function at time t . Under well-known regularity conditions, see e.g. Rust (1994), a generalized version of Blackwell's Theorem ensures that the value function is the unique solution to the Bellman equation,

$$V(x, \eta) = \max_i \Pi(x, i, \eta) + \delta E [V(x', \eta') | x, i] \tag{1.2.1}$$

and that there is a time invariant measurable optimal policy function, which we will denote by $I(x, \eta)$.

Although the problem above is written for a single agent making decisions in isolation, it can be applied to dynamic games with stationary Markov perfect equilibria as well. In this case, the state variables would include information on all firms in the same market. Suppose there are N firms, each indexed by f . The transition density for x_t given the actions of all firms can be written

$$F(x_{t+1}|x_t, \{i_{ft}\}_{f=1}^N) = F(x_{t+1}|x_t, \{I(x_t, \eta_{ft})\}_{f=1}^N).$$

However, since η_{ft} is private knowledge each firm must form expectations using

$$F(x_{t+1}|x_t, i_{ft}) = E[F(x_{t+1}|x_t, \{I(x_t, \eta_{\bar{f}t})\}_{\bar{f}=1}^N) | \eta_{ft}]. \tag{1.2.2}$$

Hence, each firm's decision problem takes the form written above. In equilibrium, the policy

function must satisfy both (1.2.1) and (1.2.2). It would be necessary to take this into account if we want to perform any counterfactual analysis. However, our identification result and estimator will only rely on the individual optimality of each firm's actions. As such we will not be explicit about equilibrium considerations throughout most of this chapter.

1.2.1 Example

This section describes a concrete example that fall under the above setup. Throughout this chapter, we will return to this example to illustrate the plausibility of various assumptions.

Example 1 *Investment in natural gas pipelines.* In ongoing work, we apply the estimator developed in this chapter to natural gas pipelines. For clarity, we describe a simplified version of this application. Investment in natural gas pipelines in the United States is subject to regulatory approval. Our goal is to recover the implicit cost of investment, including costs of regulatory compliance. The state variables are revenues, r_t , operating expenses, o_t , pipeline capacity, q_t , and pipeline utilization, u_t . Each firm tracks both its own values of these four variables, as well as those of other firms in the same market. We will use an f subscript to differentiate among firms when necessary. The control variable is pipeline investment measured in dollars, i_t . We specify the profit function of firm f as,

$$\Pi_f(x_t, i_{ft}, \eta_{ft}) = r_{ft} - o_{ft} - i_{ft} - c(i_{ft}, \eta_{ft}, q_t, u_t),$$

where the final term, $c(i_{ft}, \eta_{ft}, q_t, u_t)$ represents the cost of regulatory compliance. We include capacity and utilization in this function based on the assumption that the regulator primarily looks at the rate of capacity use when deciding how costly to make approval of new pipeline projects. Also, note that although the revenues and expenses of other firms do not enter the profit function directly, they will help the firm better forecast other firms' choices of investment. Therefore, revenues and expenses of other firms will still enter the policy function of this firm.

1.3 Identification

Given observations of $\{x_t, i_t\}$, our goal is to recover the payoff function, $\Pi(x, i, \eta)$. Rust (1994) shows that this is, in general, impossible because adding any function of x to the payoff function results in the same policy function. However, it is still possible to give a limited set of additional

restrictions that can be used to fully identify the payoff function. With continuous states variables and discrete actions, Bajari, Chernozhukov, Hong, and Nekipelov (2009) show that if payoffs are additively separable in the private shock, η , and the distribution of η is known, then the payoff function can be identified up to a location normalization with an exclusion restriction. In the same setting, continuous states and discrete controls, Berry and Tamer (2006) give a converse result. They show that if the payoff function is known, then the distribution of η can be identified. In a setting very similar to this chapter's, Berry and Pakes (2001) propose a parametric estimator based on firm's first order conditions. Their estimator requires observations of realized profits. We will also consider an estimator based on firm's first-order conditions, but we will not require data on profits.

We begin by assuming that the transition density of the states and controls, $f_{x',i'|x,i}$, is identified. We then show that the optimality of observed actions imply that the payoff function must satisfy an integro-differential equation that depends on this transition density. We propose restrictions on the payoff function and transition density that guarantee that this integro-differential equation has a unique solution. We begin by stating our assumptions and main result. The assumptions will be discussed in more detail below.

Assumptions I (Payoff function identification).

I1 The transition density, $f_{x_t,i_t|x_{t-1},i_{t-1}}$, is identified.

I2 The policy function, $I(x, \eta)$, is weakly increasing in η and $f_{\eta_t|I_t} = f_{\eta_t}$ is known.

I3 The policy function satisfies the firm's first order condition.

I4 The discount factor, δ , is known.

I5 For some value of the control, i_0 , the payoff function, $\Pi(x, i_0, \eta)$, is known for all x and η .

I6 There exists a component of the state space, $x^{(k)}$ with compact support, such that:

- $\frac{\partial \Pi}{\partial x^{(k)}}(x, i, \eta)$ is known and not identically zero for all x, η and,
- The policy function is pseudo-invertible at i_0 with respect to $x^{(k)}$ in that for all $x^{(-k)}$ and η there exists a function $\chi_k(i_0, x^{(-k)}, \eta)$ such that

$$I(\chi_k(i_0, x^{(-k)}, \eta), x^{(-k)}, \eta) = i_0$$

and χ_k is measurable with respect to $x^{(-k)}$ and η .

I7 Let $\frac{\partial \Pi}{\partial i} \in \mathcal{G}$. Define the following operators $\mathcal{G} \rightarrow \mathcal{G}$,

$$\begin{aligned}\mathcal{D}(g)(x, \eta) &= \frac{\partial}{\partial i_t} E \left[\sum_{\tau=0}^{\infty} \delta^\tau g(x_{t+\tau}, \eta_{t+\tau}) \mid x_t = x, i_t = I(x, \eta) \right] \\ \mathcal{L}(g)(x, \eta) &= \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} g(\tilde{x}^{(k)}, x^{(-k)}, \eta) \frac{\partial I}{\partial x^{(k)}}(\tilde{x}^{(k)}, x^{(-k)}, \eta) d\tilde{x}^{(k)} \\ \mathcal{K}(g)(x, \eta) &= \mathcal{D}(\mathcal{L}(g))(x, \eta)\end{aligned}$$

The only solution in \mathcal{G} to

$$0 = g(x, \eta) + \mathcal{K}(g)(x, \eta) \tag{1.3.1}$$

is $g(x, \eta) = 0$.

I8 There is some component of the state x_i excluded from $\Pi(x, i, \eta)$, but entering $I(x, \eta)$.

Theorem 1 (Payoff function identification). *If conditions I1-I7 hold, then the equilibrium payoff function, $\Pi^*(x_t, \eta_t) = \Pi(x_t, I(x_t, \eta_t), \eta_t)$, and value function $V(x_t, \eta_t)$ are identified. If, additionally, condition I8 holds, then the payoff function, $\Pi(x, i, \eta)$ is identified.*

We now discuss the assumptions of this theorem. Its proof can be found in the next section. If the state vector, x_t is observed I1 is trivially satisfied. Hu and Shum (2008) give conditions that ensure I1 in the presence of unobserved components of the state vector. These conditions are discussed in detail in section 1.3.2 below.

Condition I2 is to allow $I(x, \eta)$ and the distribution of η to be recovered from the density of i and x . The requirement that the distribution of η is known is more a normalization than an assumption since η is an unobserved variable that enters unknown functions in an unrestricted manner.

Condition I3 is to ensure that the first order condition can be used to recover the payoff function. If it is not satisfied everywhere, such as with binding constraints on i or non-differentiable portions of $\Pi(\cdot)$, then the theorem can be adapted to apply for values of x and η where the first order condition is satisfied.

The first order condition can be manipulated to give an integral equation that can be solved for $\frac{\partial \Pi}{\partial i}$. Conditions I5, I6, and I7 guarantee that this integral equation has a unique solution. Conditions I5 and I6 are analogous to some of the conditions for identifying dynamic discrete games in Magnac

and Thesmar (2002) or Bajari, Chernozhukov, Hong, and Nekipelov (2009). With discrete actions, it is necessary to normalize the payoff associated with one of the alternatives to zero. Condition I5 is the continuous analog of this normalization. With discrete actions, it is also common to assume that payoff function is additively separable in private shock. Here, the assumption that one of the partial derivatives of the payoff function is known, I6, takes the place of the assumption of additively separable private shocks. To make the analogy more clear, condition I6 could be placed on η instead of $x^{(k)}$.

Condition I7 is a key identifying assumption. It ensures that the integral equation for $\frac{\partial \Pi}{\partial i}$ has a unique solution. Similar conditions appear in the literature on nonparametric identification of simultaneous equation and are referred to as completeness conditions. In that literature, the completeness condition involves an integral equation of the first kind. Inversion of such equations is an ill-posed problem, which adds to the difficulty of estimation. Here, the completeness condition involves an integral equation of the second kind, so there is no ill-posed inverse problem.¹

It is difficult to verify condition I7 from more primitive conditions. One simple sufficient condition is that the operator $\mathcal{K} : \mathcal{G} \rightarrow \mathcal{G}$ defined above has operator norm less than one. However, there is no particular reason to believe that this condition would hold in applications. Fortunately, although it is difficult to determine a priori whether I7 holds, the condition can be checked for any given dataset. Condition I7 is an assumption about the transition density and the policy function. These two things can be identified without attempting to identify the payoff function or invoking assumption I7. Therefore, condition I7 can be checked. One could even develop a formal statistical test for I7, but we will leave that task for future work.

Finally, the exclusion restriction, I8, allows us to separate the structural payoff function from the equilibrium payoff function. Without the exclusion, we would only be able to identify the payoff function along the equilibrium path, that is only at $(x, I(x, \eta), \eta)$ rather than for any combination of (x, i, η) . While identifying the payoff function only along the equilibrium path is sufficient for some purposes, such as predicting payoffs in the current environment, identifying the payoff function everywhere is essential for counterfactual analysis.

We now discuss how these assumptions relate to our working example.

Example 1 *Investment in natural gas pipelines (continued)*. All of the state variables in this example are observed, so the transition density is identified, satisfying I1. Assumption I2 is simply

¹Estimation still involves an ill-posed inverse problem in that it requires estimation of a conditional expectation operator. However, the inversion of that operator needed to solve for the value function is well-posed.

a normalization. Condition I3, which says that first order condition holds, appears reasonable. It would follow from concavity in i of $c(i, \eta, q, u)$. One might think that investment involves a fixed adjustment cost, in which the first order condition need not hold when investment is zero. However, in this case we would observed a mass point at zero investment, and there does not appear to be any such mass point in the available data. Condition I5 can be met by assuming that when investment is zero, profits are simply revenues minus expenses. In other words, $c(0, \eta, q, u) = 0$. We have revenue data, so we satisfy the first part of I6 by assuming that profits are additively separable in revenues. The second part of the assumption is also very plausible. It is reasonable to believe that with sufficiently negative revenue, the firm would choose to shut down by setting i to a large negative value. Also, if revenues are sufficiently high, the firm should choose to invest a positive amount. By continuity, there should be some intermediate amount of revenue that leads to zero investment. As mentioned above, condition I7 is difficult to verify analytically, but it can be checked using estimates of the transition density and policy function. Finally, the exclusion restriction, I8, is met by assuming that each firm's payoffs depend on the firm's own revenue, but not the revenue of other firms. The revenue of others still enters the firm's policy function because the others' revenue helps each firm to predict others' actions, which can affect future payoffs.

1.3.1 Proof of identification (theorem 1)

We first show that the policy function is identified. Assumption I1 implies that the conditional quantile function of actions given states is known. Denote it by $Q_{i|x}(\tau|x)$. Assumption I2 – that $I(x, \eta)$ is weakly increasing in η and $F\eta$ is known – imply that the policy function is

$$I(x, F_{\eta}^{-1}(\tau)) = Q_{i|x}(\tau|x).$$

Condition I3 says that the policy function satisfies the first order condition. As stated in the text, the first order condition is

$$0 = \frac{\partial \Pi}{\partial i}(x_t, I(x_t, \eta_t), \eta_t) + \frac{\partial}{\partial i} \sum_{\tau=1}^{\infty} \delta^{\tau} E [\Pi(x_{t+\tau}, I(x_{t+\tau}, \eta_{t+\tau}), \eta_{t+\tau}) | x_t, I(x_t, \eta_t)]. \quad (1.3.2)$$

We can express the payoff function in terms of its derivative,

$$\begin{aligned}
\Pi(x, I(x, \eta), \eta) &= \int_{x_0^{(k)}}^{x^{(k)}} \frac{\partial}{\partial x^{(k)}} [\Pi(x, I(x, \eta), \eta)] d\tilde{x}^{(k)} + \\
&\quad + \Pi\left(x_0^{(k)}, x^{(-k)}, I(x_0^{(k)}, x^{(-k)}, \eta), \eta\right) \\
&= \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial i}(x, I(x, \eta), \eta) \frac{\partial I}{\partial x^{(k)}}(x, \eta) + \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)} \\
&\quad + \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) \\
&= \mathcal{L}\left(\frac{\partial \Pi}{\partial i}\right)(x, \eta) + \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) + \\
&\quad + \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)} \tag{1.3.3}
\end{aligned}$$

where the first line is just the fundamental theorem of calculus, and the second line comes from setting $x_0^{(-k)} = \chi_x(i_0, x^{(-k)}, \eta)$, and expanding $\frac{\partial}{\partial x^{(k)}} [\Pi(x, I(x, \eta), \eta)]$. Let

$$\begin{aligned}
\varphi(x, \eta) &= \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) + \\
&\quad + \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)}
\end{aligned}$$

Note that by I5 and I6, $\varphi(x, \eta)$ is known. Substituting this and (1.3.3) into (1.3.2) gives,

$$\begin{aligned}
0 &= \frac{\partial \Pi}{\partial i}(x_t, I(x_t, \eta_t), \eta_t) + \frac{\partial}{\partial i} E \left[\sum_{\tau=1}^{\infty} \delta^\tau \mathcal{L}\left(\frac{\partial \Pi}{\partial i}\right)(x_{t+\tau}, \eta_{t+\tau}) + \varphi(x_{t+\tau}, \eta_{t+\tau}) \mid x_t, I(x_t, \eta_t) \right] \\
&= (\mathbf{1} + \mathcal{K}) \left(\frac{\partial \Pi}{\partial i} \right) + \mathcal{D}(\varphi)
\end{aligned}$$

We have assumed that $\frac{\partial \Pi}{\partial x^{(k)}}$ and $\Pi(x, i_0, \eta)$ are known, which implies that $\varphi(x, \eta)$ is known. Also, condition I7 ensures that the operator, $\mathbf{1} + \mathcal{K}$, in the above equation is invertible, so we can identify $\frac{\partial \Pi}{\partial i}(x, I(x, \eta), \eta)$. Integrating as in (1.3.3) gives $\Pi(x, I(x, \eta), \eta)$.

Finally, to recover the payoff function for at places where $i \neq I(x, \eta)$, we use the exclusion restriction I8. In particular this exclusion implies that the payoff function can be written as $\Pi(x^{(-i)}, i, \eta)$, so by varying $x^{(i)}$ while holding $x^{(-i)}$ and η fixed, we can identify the payoff function on the set $\{x^{(-i)}, i, \eta : i = I(x^{(i)}, x^{(-i)}, \eta)\}$. ■

1.3.2 Identifying transition densities with unobserved state variables

Theorem 1 above shows that given the transition density and some normalizations, we can identify the equilibrium payoff function, $\Pi^*(x, \eta)$. As stated above, when states are fully observed, the transition density is easily identified. When some components of the state vector are unobserved, it can still be possible to identify the transition density. Hu and Shum (2008) give conditions to identify the transition density in the presence of unobserved state variables. Their identification argument relies on the eigendecomposition of an integral operator associated with the joint density of observed states.

Suppose the state vector can be written as $x_t = (w_t, \epsilon_t)$, where w_t is observed, but ϵ_t is not. Hu and Shum (2008) show that the transition density can be identified under the following assumptions:

Assumptions U (*Transition density identification with unobserved states*).

U1 The states and controls follow a stationary first-order Markov process,

$$f_{x_t, i_t | x_{t-1}, i_{t-1}, \mathcal{I}_{t-1}} = f_{x', i' | x, i}.$$

U2 Strategies are Markovian, so that

$$f_{i_t | w_t, \epsilon_t, w_{t-1}, \epsilon_{t-1}, i_{t-1}} = f_{i_t | w_t, \epsilon_t}.$$

U3 The unobserved states are not affected by the controls in that

$$f_{\epsilon_t | w_t, i_{t-1}, w_{t-1}, \epsilon_{t-1}} = f_{\epsilon_t | w_t, \epsilon_{t-1}}.$$

U4 There is sufficient variation in i_{t-1} with ϵ_{t-1} given fixed w_t and w_{t-1} , so that there exists a known function $\omega(i)$ such that $|E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon_{t-1}]| < \infty$ and for any $\epsilon_{t-1} \neq \epsilon'_{t-1}$,

$$E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon_{t-1}] \neq E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon'_{t-1}].$$

U5 There is sufficient variation in i_t with i_{t-2} given fixed w_t and w_{t-1} so that the operators defined by

$$(L_{I_t, w_t, \omega | w_{t-1}, I_{t-2}} h)(i') = \int h(i_{t-2}) \int \omega(i_{t-1}) f_{i_t, w_t, i_{t-1} | w_{t-1}, i_{t-2}}(i', w_t, i_{t-1} | w_{t-1}, i_{t-2}) di_{t-1} di_{t-2}$$

are one-to-one for each w_t, w_{t-1} .

U6 There is variation in i_t with ϵ_{t-1} given w_t, w_{t-1} so that the operators defined by

$$(L_{I_t|w_t, w_{t-1}, \epsilon_{t-1}} h)(i') = \int h(\epsilon_{t-1}) f_{i_t|w_t, w_{t-1}, \epsilon_{t-1}}(i'|w_t, w_{t-1}, \epsilon_{t-1}) d\epsilon_{t-1}$$

are one-to-one for each w_t, w_{t-1} .

U7 There is variation in i_t with ϵ_t given w_t so that the operators defined by:

$$(L_{i_t|w_t, \epsilon_t} h)(i) = \int h(\epsilon) f_{i_t|w_t, \epsilon_t}(i|w_t, \epsilon) d\epsilon$$

are one-to-one for each w_t .

U8 For each component of ϵ_{t-1} , there exists a known functional G_k such that

$$G_k [f_{i_t|w_t, w_{t-1}, \epsilon_{t-1}}(\cdot|w_t, w_{t-1}, \epsilon_{t-1})]$$

is monotonic in $\epsilon_{t-1, k}$.

Theorem 2 (Hu and Shum). *If assumptions U1-U8, then the observed density of $(i_t, w_t, i_{t-1}, w_{t-1}, i_{t-2}, w_{t-2})$ uniquely determines the equilibrium transition density $f_{i', w', \epsilon'|i, w, \epsilon}$.*

Assumptions U1 and U2, which state that the environment is stationary and Markovian, are standard for dynamic games and have already been made above. Assumption U3 restricts the nature of the unobserved state variables. It says that the unobserved state variables are independent of the control in the previous period conditional on the previous state and current observed state variable.

Assumptions U4-U7 ensure that there is a sufficiently strong relationship between observed and unobserved variables so that the distribution of observed variables can be used to construct the distribution of unobservables. Although assumptions U5-U7 can be difficult to fully verify, simple necessary conditions for each can be expected to hold. For example, U7 requires that conditional on each observed state, the unobserved state influences firms' actions. If this were not true for at least some observed state, then the unobserved state is irrelevant and should not be part of the model. Similarly, assumptions U5 and U6 require that actions be related to unobserved states and past and current unobserved states be related. Additionally, assumptions U6 and U7 require that

the support of i_t be at least as “large” as the support of ϵ_t . For example, if ϵ_t is continuous with k dimensions, i_t must be continuous with at least k dimensions. If ϵ_t is discrete with k points of support, i_t must have at least k points of support.

Assumption U4 is easier than assumptions U5-U7 to verify. However, it also fails in some common models. In particular assumption U4 rules out models where there is a deterministic relationship between the current state, past action, and past state. For example, if the state includes the capital stock, the action is investment, and depreciation is deterministic, then $i_t = w_t - (1 - \delta)w_{t-1}$ and there is no variation with ϵ in $E[\omega(i_t)|w_t, w_{t-1}, \epsilon_t]$. Nonetheless, even in models with deterministic accumulation, the above result can be useful. Since investment can be recovered from w_{t+1} and w_t , it can be excluded from the model while identifying the transition density. Then if remaining actions satisfy the above conditions, the transition density is identified. In practice, candidates for these remaining actions include anything that satisfy assumptions U2 and U3. That is, they must be functions of only the current state and should not influence the next unobserved state. Even actions that have no dynamic implications, such as variable inputs would be suitable. Additionally, depending on the model and interpretation of ϵ_t , outcome variables such as output, revenue, or costs might be suitable.

1.4 Estimation

Although the above identification result is nonparametric, our estimation approach is semiparametric. We assume that the payoff function is known up to a finite dimensional parameter, but allow the transition distribution and policy function to be nonparametric. As in much of the dynamic game literature, we estimate the model in multiple steps. First, we estimate that transition distribution of state variables and the policy function. We then use the estimated policy function and transition distribution along with the optimality conditions of the model to estimate the payoff function.

1.4.1 Estimating policy functions

We begin by supposing we have some nonparametric estimates of the transition density, $\hat{f}_{x',i'|x,i}$, and the conditional distribution of actions given states, $\hat{F}_{i|x}$. Theorem 4 gives high level conditions on these estimates for the parameters of the payoff function to be consistent and \sqrt{n} asymptotically normal. In section 1.5.2 below, we show that certain kernel estimates satisfy these high level

conditions. We expect that sieve estimates would as well.

Given $\hat{F}_{i|x}$, estimation of the policy function is straightforward. We have assumed that the policy function is weakly increasing in η , that η is independent of x , and that the distribution of η is known. Therefore, an estimate of the inverse policy function can be formed from the following relationship:

$$\begin{aligned} F\eta(\eta) &= \hat{F}_{i|x}(i|x) \\ \hat{\eta}(i, x) &= F\eta^{-1} \left[\hat{F}_{i|x}(i|x) \right]. \end{aligned}$$

1.4.2 Estimating the payoff function

The payoff function is estimated from the optimality conditions of the model. To do this, we must solve for the value function given the above estimates and a candidate parameter value. This can be done by evaluating the Bellman equation at the estimated policy function using the estimated transition densities to evaluate the conditional expectation:

$$\hat{V}(x, i; \theta, \hat{f}, \hat{\eta}) = \Pi(x, i, \hat{\eta}(i, x); \theta) + \delta E_{\hat{f}} \left[\hat{V}(x', i'; \theta, \hat{f}, \hat{\eta}) | x, i \right]. \quad (1.4.1)$$

This functional equation can be used to solve for $\hat{V}(x, i; \theta)$, the value of being in state x with a private shock such that i is the estimated policy.² Let the value of deviating from the estimated policy for one period be denoted as

$$\mathcal{V}(x, \eta, \tilde{i}; \theta, \hat{f}, \hat{\eta}) = \Pi(x, \tilde{i}, \eta; \theta) + \delta E_{\hat{f}} \left[\hat{V}(x', i'; \theta, \hat{f}, \hat{\eta}) | x, \tilde{i} \right].$$

The optimality conditions of the model can be used to form moment conditions to estimate the profit function. At least three ways of using the optimality conditions of the model to estimate the payoff function have appeared in the literature. We will show that these three methods result in asymptotically equivalent estimators when the moments are weighted appropriately. Jenkins, Liu, Matzkin, and McFadden (2004) minimize the difference between observed actions and optimal

²Since, given the state, investment and the private shock contain the same information, we could have instead defined a value function over x and η . In this case, the i in the right hand side of (1.4.1) would be replaced with $\hat{I}(x, \eta)$. Although defining the value function over x and η would be more standard, we chose not to do this because estimating the inverse policy function is slightly more straightforward than estimating the policy function.

actions. This action based moment condition is:

$$0 = m_a(i, x; \theta, \hat{h}) = i - \arg \max_{\tilde{i}} \mathcal{V}(x, \hat{\eta}(i, x), \tilde{i}; \theta, \hat{h}),$$

where $\hat{h} = (\hat{\eta}, \hat{f})$. Alternatively, the first order condition for investment can be used as a moment condition:

$$0 = m_f(i, x; \theta, \hat{h}) = \frac{\partial \mathcal{V}}{\partial i}(x, \hat{\eta}(i, x), i; \theta, \hat{h}).$$

Hong and Shum (2009) use this kind of first order condition based moment. A third candidate moment condition is the difference between the value of observed actions and the maximized value function. This type of moment condition is used by Macieira (2009):

$$0 = m_v(i, x; \theta, \hat{h}) = \hat{V}(x, \eta; \theta, \hat{h}) - \max_{\tilde{i}} \mathcal{V}(x, \hat{\eta}(i, x), \tilde{i}; \theta, \hat{h}).$$

Note that each of these conditions hold for all i, x . An objective function can be formed by taking a combination of them. Since the estimates of the transition distribution and policy are estimated nonparametrically, we must use a continuum of moment conditions in the objective function to enable the payoff function parameters to be \sqrt{n} -consistent. Let $m(\cdot)$ be one of $m_a(\cdot)$, $m_f(\cdot)$, or $m_v(\cdot)$. The objective function is

$$\begin{aligned} Q_n(\theta; \hat{h}) &= \int_{\mathcal{I} \times X} \int_{\mathcal{I} \times X} m(i, x; \theta, \hat{h}) w_n(i, x; i', x') m(i', x'; \theta, \hat{h}) di dx di' dx' \\ &= \left\| B_n m(\cdot; \theta, \hat{h}) \right\|_n^2 \end{aligned}$$

where $w_n(i, x; i', x')$ is some weighting function, possibly data-dependent, and the second line is simply alternate notation for the first. In particular, $\|f\|^2 = \int f(i, x)^2 di dx$, and B_n is an integral operator with a kernel that satisfies

$$w_n(i, x; i', x') = \int b_n(i, x; \tilde{i}, \tilde{x}) b_n(i', x'; \tilde{i}, \tilde{x}) d\tilde{i} d\tilde{x}.$$

The parameters of the payoff function are estimated by minimizing the objective function,

$$\hat{\theta} = \arg \min_{\theta \in \Theta} Q_n(\theta; \hat{h}).$$

1.5 Asymptotic Theory

The above estimator of θ is similar to the estimation frameworks of Carrasco and Florens (2000); Carrasco, Chernov, Florens, and Ghysels (2007); and Ai and Chen (2003). As a result, we can derive the asymptotic distribution of $\hat{\theta}$ using a suitable adaptation of their arguments. The current situation differs from the setup of Carrasco and Florens (2000) and Carrasco, Chernov, Florens, and Ghysels (2007) in that the moment conditions here include a preliminary nonparametric estimate. Ai and Chen (2003) study conditional moment conditions with nonparametric components. They derive their results for independent data, and so their weighting operator is diagonal. Here, the data is not independent. Also, Ai and Chen (2003) consider simultaneous estimation of the nonparametric and parametric components of their model. Here, the nonparametric and parametric components are estimated in two steps.

The following theorem gives conditions on $m(\cdot; \theta, h)$ and \hat{h} sufficient for asymptotic normality. The conditions are stated for generic moment functions $m(\cdot; \theta, h)$ and nonparametric estimate \hat{h} . Theorem 4 gives more specific conditions on the payoff function and transition density that imply the generic conditions of Theorem 3 for each of the three versions of $m(\cdot; \theta, h)$ defined above. Section 1.5.2 describes kernel estimates of the transition density and inverse policy function that satisfy the generic conditions of 3. Let $\langle f, g \rangle$ denote the inner-product associated with the norm used to define Q_n , i.e.

$$\langle f, g \rangle = \int f(i, x)g(i, x)didx.$$

For any functions, f and g , let $f(x) \lesssim g(x)$ mean that there exist a constant M such that $f(x) \leq Mg(x)$. Similarly, let $f(x) \lesssim_p g(x)$ mean that $f(x) \leq Mg(x) + o_p(1)$ uniformly in x .

Assumptions A (*Asymptotic normality for semiparametric minimum distance estimators in a Hilbert space*).

A1 $\hat{\theta}$ approximately minimizes $Q_n(\theta)$, so that $Q_n(\hat{\theta}) \leq \inf_{\theta \in \Theta} Q_n(\theta) + o_p(n^{-1/2})$.

A2 Let $\|B\|_{op}$ be the usual operator norm, $\|B\|_{op} = \sup_{\|f\|=1} \|Bf\|$, assume that $\|B_n - B\|_{op} \xrightarrow{p} 0$.

A3 The nonparametric estimates of the transition density and policy function converge at a rate

faster than $n^{-1/4}$ with respect to some norm $\|\cdot\|_{\mathcal{H}}$, in that

$$\left\| B(\hat{h} - h) \right\|_{\mathcal{H}} = o_p \left(n^{-1/4} \right).$$

A4 The following derivatives exist and satisfy the following conditions:

A4.i For all θ with $\|\theta - \theta_0\| \leq \delta_n$, $m(\cdot; \theta, h)$ is pathwise differentiable with respect to h at h_0 .

We denote this derivative as $D_h^m(\theta, h_0)$. This derivative is such that

$$\sup_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}}^2$$

and

A4.ii

$$\sup_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \left(D_h^m(\theta, h_0)(\hat{h} - h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}} \|\theta - \theta_0\|.$$

A4.iii In a neighborhood of θ_0 , $m(\cdot; \theta, h_0)$ is continuously differentiable with respect to θ . This derivative is bounded away from zero in that

$$1 / \left(\inf_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \theta, h_0) \right\| \right) \lesssim_p 1,$$

and this derivative is Lipschitz continuous, so that

$$\sup_{\theta_1, \theta_2: \|\theta_i - \theta_0\| \leq \delta_n} \left\| B_n \left(\frac{\partial m}{\partial \theta}(\cdot; \theta_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \theta_2, h_0) \right) \right\| \lesssim_p \|\theta_1 - \theta_2\|.$$

A4.iv In a neighborhood of θ_0 , $\|B_n m(\cdot; \theta, h_0)\|^2$ is continuously differentiable with respect to θ with derivative

$$\frac{\partial}{\partial \theta} \|B_n m(\cdot; \theta, h_0)\|^2 = 2 \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \theta, h_0), B_n m(\cdot; \theta, h_0) \right\rangle.$$

A5 \hat{h} and B_n are bounded in the sense that there is some constant M such that with probability approaching one

$$\left\| B_n m(\cdot; \theta, \hat{h}) \right\| \leq M$$

and

$$\left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \leq M.$$

A6

$$\sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \xrightarrow{d} N(0, \Omega_B)$$

Theorem 3 (Asymptotic distribution for semiparametric minimum distance estimators in a Hilbert space). *If A1-A6 then $\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, M_B^{-1} \Omega_B M_B^{-1})$ with $M_B = \langle B \frac{\partial m}{\partial \theta}, B \frac{\partial m}{\partial \theta} \rangle^{-1}$.*

Proof. We will show that $\hat{\theta}$, which minimizes

$$Q_n(\theta) = \left\| B_n m(\cdot; \theta, \hat{h}) \right\|^2$$

is close to the minimizer of the linearized objective function,

$$L_n(\theta) = \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|^2.$$

First, we must show that $Q_n(\theta)$ is close to $L_n(\theta)$ in that

$$\sup_{\|\theta - \theta_0\| \leq \delta_n} |Q_n(\theta) - L_n(\theta)| = o_p(n^{-1/2}).$$

Note that

$$\begin{aligned} |Q_n(\theta) - L_n(\theta)| &= \left| \left\| B_n m(\cdot; \theta, \hat{h}) \right\|^2 - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|^2 \right| \\ &= \left| \left(\left\| B_n m(\cdot; \theta, \hat{h}) \right\| - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right) \times \right. \\ &\quad \left. \times \left(\left\| B_n m(\cdot; \theta, \hat{h}) \right\| + \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right) \right| \\ &\lesssim \left| \left\| B_n m(\cdot; \theta, \hat{h}) \right\| - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right| \\ &\lesssim \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \end{aligned}$$

where we used assumption A5 and the reverse triangle inequality. Note that

$$\begin{aligned} & \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \leq \\ & \leq \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta, h_0)(\hat{h} - h_0) \right) \right\| + \\ & \quad + \left\| B_n \left(D_h^m(\theta, h_0)(\hat{h} - h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \end{aligned}$$

Invoking assumptions A4.i and A4.ii gives

$$\begin{aligned} \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| & \lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}}^2 + \\ & \quad + \left\| \hat{h} - h_0 \right\|_{\mathcal{H}} \|\theta - \theta_0\| \end{aligned}$$

Using assumption A3, we have

$$\left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p o_p(n^{-1/2}) + \tag{1.5.1}$$

$$+ o_p(n^{-1/4}) \left\| \hat{\theta} - \theta_0 \right\|. \tag{1.5.2}$$

We will now show that $\left\| \hat{\theta} - \theta_0 \right\| \leq O_p(n^{-1/2})$. By assumption A4.iii, we can take a mean value expansion in θ around θ_0 , so that

$$m(\cdot; \hat{\theta}, h_0) = m(\cdot; \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0)$$

Therefore,

$$m(\cdot; \hat{\theta}, \hat{h}) - m(\cdot; \hat{\theta}, h_0) + \left(m(\cdot; \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right) = 0$$

Rearranging and using the fact that $m(\cdot; \theta_0, h_0) = 0$, we have

$$\left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right\| = \left\| B_n \left(m(\cdot; \hat{\theta}, \hat{h}) - m(\cdot; \hat{\theta}, h_0) \right) \right\|$$

By A4.i we have

$$\begin{aligned} \left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right\| & = \left\| B_n D_h^m(\hat{h} - h_0) \right\| + o_p(n^{-1/2}) \\ \left\| \hat{\theta} - \theta_0 \right\| & \lesssim \left\| B D_h^m(\hat{h} - h_0) \right\| + o_p(n^{-1/2}) \end{aligned}$$

Thus, by assumptions A4.iii and A6, we can conclude that $\|\hat{\theta} - \theta_0\| \leq O_p(n^{-1/2})$.

Combining the just obtained rate bound for $\hat{\theta}$ with equation (1.5.2) gives:

$$|Q_n(\hat{\theta}) - L_n(\hat{\theta})| = o_p(n^{-1/2}).$$

Now we show that the minimizer of $Q_n(\theta) = \|B_n m(\cdot; \theta, \hat{h})\|$ is close to the minimizer of $L_n(\theta) = \left\| B_n \left(m(\cdot, \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|$. Let $\hat{\theta}$ denote the former and $\tilde{\theta}$ denote the later. By assumption A1,

$$Q_n(\hat{\theta}) \leq Q_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

Using the expansion results above we then have,

$$L_n(\hat{\theta}) \leq L_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

On the other hand, $\tilde{\theta}$ is defined as the minimizer of $L_n(\theta)$, so $L_n(\tilde{\theta}) \leq L_n(\hat{\theta})$ and we can conclude that

$$L_n(\hat{\theta}) = L_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

By assumption A4.iv we can expand both sides around θ_0 to get

$$\begin{aligned} \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\hat{\theta} - \theta_0) &= \\ &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\tilde{\theta} - \theta_0) + o_p(n^{-1/2}) \end{aligned}$$

Rearranging gives

$$\begin{aligned} \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\tilde{\theta} - \hat{\theta}) &= \\ = \left\langle B_n \left(\frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0) \right), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\hat{\theta} - \theta_0) &+ o_p(n^{-1/2}) \end{aligned}$$

By assumption A4.iii,

$$\left\| \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0) \right\| \lesssim \|\bar{\theta}_1 - \bar{\theta}_2\|.$$

Also we have already shown that $\|\hat{\theta} - \theta_0\| \leq O_p(n^{-1/2})$, which implies that $\|\bar{\theta}_1 - \bar{\theta}_2\| \leq O_p(n^{-1/2})$. Thus, we can conclude that $\|\tilde{\theta} - \hat{\theta}\| = o_p(n^{-1/2})$.

Finally, expanding the first order condition for the linearized objective function gives,

$$\begin{aligned}
0 &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \left(m(\cdot, \tilde{\theta}, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\rangle \\
0 &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \left(\begin{array}{l} m(\cdot, \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0)(\tilde{\theta} - \theta_0) + \\ + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \end{array} \right) \right\rangle \\
(\tilde{\theta} - \theta_0) &= - \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0) \right\rangle^{-1} \times \\
&\quad \times \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \\
&= - \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\
&\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(n^{-1/2})
\end{aligned}$$

To conclude, we have

$$\begin{aligned}
\sqrt{n}(\hat{\theta} - \theta_0) &= - \sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\
&\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(1)
\end{aligned}$$

and the statement in the theorem follows from assumption A6. ■

We now give conditions on the payoff function and other model primitives such that the conditions of theorem 3 are satisfied. We begin by establishing differentiability of $m(\cdot; \theta, h)$ for each of the three candidates for $m(\cdot; \theta, h)$ introduced above. Each $m(\cdot; \theta, h)$ is simple function of $\mathcal{V}(x, \eta, i; \theta, f, \boldsymbol{\eta})$, so Lemma 1 about the derivatives of \mathcal{V} is useful. Before stating this lemma, we define some additional notation. Let $\mathcal{E}_f(g)(x, i)$ denote the following conditional expectation operator:

$$\mathcal{E}_f(g)(x, i) = E [g(x', i') | x, i]$$

where $g : X \times I \rightarrow \mathbb{R}$ and the expectation is taken with respect to the transition density, $f_{x', i' | x, i}$. Also, let \mathbf{I} denote the identity operator, $g \circ \boldsymbol{\eta} = g(x, i, \boldsymbol{\eta}(i, x))$ for $g : X \times I \times \mathcal{N} \rightarrow \mathbb{R}^k$, and let $\Pi\theta(x, i, \boldsymbol{\eta}) = \Pi(x, i, \boldsymbol{\eta}; \theta)$. Note that $(\mathbf{I} - \delta\mathcal{E}_f)^{-1}$ exists because $\|\delta\mathcal{E}_f\|_{op} = \delta < 1$. Using this

notation, we can write \mathcal{V} as

$$\mathcal{V}(x, \eta, i; \theta, \hat{f}, \hat{\eta}) = \Pi\theta(x, i, \eta) + \delta\mathcal{E}_{\hat{f}} \left((I - \delta\mathcal{E}_{\hat{f}})^{-1} \Pi\theta \circ \hat{\eta} \right) (x, i, \eta)$$

Lemma 1 (Differentiability of $\mathcal{V}(x, \eta, \tilde{i}; \theta, f, \eta)$). *If*

1. *The payoff function is twice continuously differentiable with respect to i , η , and θ for all (i, x, η) , and for all θ in a neighborhood of θ_0 . The derivative with respect to θ is bounded by a function with finite expectation with respect to f , i.e.*

$$\left| \frac{\partial \Pi}{\partial \theta} \right| \leq m(i, x)$$

with

$$\int m(i', x') f_0(i', x' | i, x) di' dx' < \infty$$

uniformly in (i, x) . The derivatives with respect to η are Lipschitz continuous in that for all θ in a neighborhood of θ_0 ,

$$\left| \frac{\partial \Pi\theta}{\partial \eta}(x, i, \eta_1) - \frac{\partial \Pi\theta}{\partial \eta}(x, i, \eta_2) \right| \leq M_1(i, x) |\eta_1 - \eta_2|$$

with $\|B_n M_1\|$ bounded. Similarly, for all θ in a neighborhood of θ_0 ,

$$\left| \frac{\partial^2 \Pi\theta}{\partial \eta \partial i}(x, i, \eta_1) - \frac{\partial^2 \Pi\theta}{\partial \eta \partial i}(x, i, \eta_2) \right| \leq M_2(i, x) |\eta_1 - \eta_2|$$

with $\|B_n M_2\|$ bounded.

2. *The conditional expectation operator, \mathcal{E}_f is twice continuously differentiable respect to i .*

then \mathcal{V} is twice continuously differentiable with respect to i and θ at $h = h_0$, for all i , and for all θ

in a neighborhood of θ_0 . Moreover, these derivatives are given by:

$$\frac{\partial \mathcal{V}}{\partial \theta}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial \Pi \theta}{\partial \theta}(x, i, \eta) + \delta \mathcal{E}_{f_0} \left((\mathbf{I} - \delta \mathcal{E}_{f_0})^{-1} \frac{\partial \Pi \theta}{\partial \theta} \circ \boldsymbol{\eta}_0 \right)(x, i) \quad (1.5.3)$$

$$\frac{\partial \mathcal{V}}{\partial i}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial \Pi \theta}{\partial i}(x, i, \eta) + \delta \frac{\partial \mathcal{E}_{f_0}}{\partial i} \left((\mathbf{I} - \delta \mathcal{E}_{f_0})^{-1} \Pi \circ \boldsymbol{\eta}_0 \right)(x, i) \quad (1.5.4)$$

$$\frac{\partial^2 \mathcal{V}}{\partial i^2}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial^2 \Pi \theta}{\partial i^2}(x, i, \eta) + \delta \frac{\partial^2 \mathcal{E}_{f_0}}{\partial i^2} \left((\mathbf{I} - \delta \mathcal{E}_{f_0})^{-1} \Pi \theta \circ \boldsymbol{\eta}_0 \right)(x, i) \quad (1.5.5)$$

$$\frac{\partial^2 \mathcal{V}}{\partial i \partial \theta}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial^2 \Pi \theta}{\partial i \partial \theta}(x, i, \eta) + \delta \frac{\partial \mathcal{E}_{f_0}}{\partial i} \left((\mathbf{I} - \delta \mathcal{E}_{f_0})^{-1} \frac{\partial \Pi \theta}{\partial \theta} \circ \boldsymbol{\eta} \right)(x, i). \quad (1.5.6)$$

Also, \mathcal{V} is pathwise differentiable with respect to f and $\boldsymbol{\eta}$ in a neighborhood of f_0 and $\boldsymbol{\eta}_0$ with derivatives

$$D_{\boldsymbol{\eta}}^{\mathcal{V}}(\theta, f, \boldsymbol{\eta})(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) = \delta \mathcal{E}_f \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\left(\frac{\partial \Pi \theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} \right) (\boldsymbol{\eta}_1(\cdot) - \boldsymbol{\eta}(\cdot)) \right) \right) \quad (1.5.7)$$

$$D_f^{\mathcal{V}}(\theta, f, \boldsymbol{\eta})(f_1 - f) = \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\delta \int \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta} \right) (x', i') \times \right. \right. \\ \left. \left. \times (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di' \right) \right] + \\ + \delta \int \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta} \right) (x', i') (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di'. \quad (1.5.8)$$

Finally, $D_{\boldsymbol{\eta}}^{\mathcal{V}}$ and $D_f^{\mathcal{V}}$ are continuously differentiable with respect to i these derivatives are equal to the pathwise derivatives of $\frac{\partial \mathcal{V}}{\partial i}$ with respect to $\boldsymbol{\eta}$ and f , which also exist. These derivatives are given by

$$\frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i}(\theta, f, \boldsymbol{\eta})(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) = \delta \frac{\partial \mathcal{E}_f}{\partial i} \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\left(\frac{\partial \Pi \theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} \right) (\boldsymbol{\eta}_1(\cdot) - \boldsymbol{\eta}(\cdot)) \right) \right) \quad (1.5.9)$$

$$\frac{\partial D_f^{\mathcal{V}}}{\partial i}(\theta, f, \boldsymbol{\eta})(f_1 - f) = \delta \frac{\partial \mathcal{E}_f}{\partial i} \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\delta \int \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta} \right) (x', i') \times \right. \right. \\ \left. \left. \times (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di' \right) \right] + \\ + \delta \int \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta} \right) (x', i') \times \\ \times \left(\frac{\partial f_1}{\partial i}(x', i'|\cdot) - \frac{\partial f}{\partial i}(x', i'|\cdot) \right) dx' di', \quad (1.5.10)$$

and these derivatives are Lipschitz continuous with respect to $\boldsymbol{\eta}$, f , and θ uniformly in a neighborhood of $\boldsymbol{\eta}_0$, f_0 , and θ_0 .

Proof. Assumption 1 allows us to apply the dominated convergence theorem to show that

$$\frac{\partial}{\partial \theta} \mathcal{E}_f \Pi(\cdot; \theta) = \mathcal{E}_f \frac{\partial \Pi}{\partial \theta}(\cdot; \theta)$$

and

$$\frac{\partial}{\partial \theta} (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi(\cdot; \theta) = (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \frac{\partial \Pi}{\partial \theta}(\cdot; \theta)$$

Given this, (1.5.3)-(1.5.6) follow directly from the assumed differentiability of the payoff function and conditional expectation.

We now show (1.5.7). Recall that

$$\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}) = \Pi \theta(x, i, \eta) + \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta} \right]$$

Let $\mathcal{V}_j = \mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}_j)$ for $j = 1, 2$. Note that,

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 &= \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta}_1 \right] - \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi \theta \circ \boldsymbol{\eta}_2 \right] \\ &= \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} (\Pi \theta \circ \boldsymbol{\eta}_1 - \Pi \theta \circ \boldsymbol{\eta}_2) \right] \end{aligned}$$

Also,

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 - D_{\boldsymbol{\eta}}^{\mathcal{V}}(\cdot; \theta, f, \boldsymbol{\eta}_1) \\ = \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\Pi \theta \circ \boldsymbol{\eta}_1 - \Pi \theta \circ \boldsymbol{\eta}_2 - \left(\frac{\partial \Pi \theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta}_1 \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2) \right) \right] \end{aligned}$$

Taking a mean value expansion in $\boldsymbol{\eta}$ of $\Pi \theta \circ \boldsymbol{\eta}_2$, along with the assumed Lipschitz condition on $\frac{\partial \Pi}{\partial \boldsymbol{\eta}}$ gives

$$\|\mathcal{V}_1 - \mathcal{V}_2 - D_{\boldsymbol{\eta}}^{\mathcal{V}}(\cdot; \theta, f, \boldsymbol{\eta}_1)\| \leq \frac{\delta}{1 - \delta} \|f\|_{\mathcal{F}} \|M_1\| \|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2\|^2$$

Thus, $D_{\boldsymbol{\eta}}^{\mathcal{V}}$ is the pathwise derivative as claimed. A similar argument would also show that (1.5.9) is the pathwise derivative of $\frac{\partial \mathcal{V}}{\partial i}$ with respect to $\boldsymbol{\eta}$.

We now show (1.5.8). Recall that $\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta})$ denotes the value of choosing i when the state is x and the private shock is η and $\hat{V}(x, i; \theta, f, \boldsymbol{\eta})$ denotes the value of being in state x and choosing i when $\eta = \boldsymbol{\eta}(i, x)$, i.e. the value of choosing i when η is such that i is optimal. \mathcal{V} can be written as

$$\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}) = \Pi \theta(x, i, \eta) + \delta \mathcal{E}_f \hat{V}(\cdot; \theta, f, \boldsymbol{\eta}),$$

where \hat{V} can be defined as the solution to the following equation:

$$\hat{V}(i, x; \theta, f, \boldsymbol{\eta}) = \Pi\theta(x, i, \boldsymbol{\eta}(i, x)) + \delta\mathcal{E}_f \hat{V}(\cdot; \theta, f, \boldsymbol{\eta}).$$

Consider the difference between \mathcal{V} for two values of f .

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 &= \delta\mathcal{E}_{f_1} \hat{V}_1 - \delta\mathcal{E}_{f_2} \hat{V}_2 \\ &= \delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' + \delta\mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2) \end{aligned}$$

where now $\mathcal{V}_j = \mathcal{V}(\cdot; \theta, f_j, \boldsymbol{\eta})$ and $\hat{V}_j = \hat{V}(\cdot; \theta, f_j, \boldsymbol{\eta})$. Similarly,

$$\hat{V}_1 - \hat{V}_2 = \delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' + \delta\mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2)$$

Although this equation is nearly identical to the previous one, here we can subtract $\delta\mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2)$ from both sides and apply $(\mathbf{I} - \delta\mathcal{E}_{f_2})^{-1}$ to obtain

$$\hat{V}_1 - \hat{V}_2 = (\mathbf{I} - \delta\mathcal{E}_{f_2})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right].$$

Now note that,

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 - D_f^{\mathcal{V}}(\cdot; \theta, f_1, \boldsymbol{\eta}) &= \\ &= \mathcal{E}_{f_2} \left((\mathbf{I} - \delta\mathcal{E}_{f_2})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) - \\ &\quad - \mathcal{E}_{f_1} \left((\mathbf{I} - \delta\mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) \\ &= \mathcal{E}_{f_2} \left(\left[(\mathbf{I} - \delta\mathcal{E}_{f_2})^{-1} - (\mathbf{I} - \delta\mathcal{E}_{f_1})^{-1} \right] \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) + \\ &\quad + (\mathcal{E}_{f_2} - \mathcal{E}_{f_1}) \left((\mathbf{I} - \delta\mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right). \end{aligned}$$

It is trivial that

$$\begin{aligned} &\|(\mathcal{E}_{f_2} - \mathcal{E}_{f_1}) \\ &\left((\mathbf{I} - \delta\mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right)\| \lesssim \\ &\lesssim \|f_2 - f_1\|_{\mathcal{F}}^2. \end{aligned}$$

Also, for any function $g(i, x)$,

$$\left\| (\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1} g - (\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1} g \right\| \leq \frac{\delta}{1 - \delta} \|g\| \|f_1 - f_2\|_{\mathcal{F}}$$

Thus, we can conclude that

$$\lim_{f_2 \rightarrow f_1} \frac{\left\| \mathcal{V}_1 - \mathcal{V}_2 - D_f^{\mathcal{V}}(\cdot; \theta, f_1, \boldsymbol{\eta})(f_1 - f_2) \right\|}{\|f_1 - f_2\|_{\mathcal{F}}} = 0$$

and $D_f^{\mathcal{V}}$ is the derivative as claimed. It is elementary to show that $D_f^{\mathcal{V}}$ has the derivative with respect to i claimed in (1.5.10). Showing that $\frac{\partial \mathcal{V}}{\partial i}$ has (1.5.10) pathwise derivative involves the same argument just used to show (1.5.8), so we omit the details. Also, showing that the pathwise derivatives are Lipschitz can be done in similar manner, so we omit the details. ■

We now give conditions on the payoff function that are sufficient for the conditions theorem 3. Let \mathcal{F} be the Sobolev like space of functions $f : (X \times \mathcal{I}) \times (X \times \mathcal{I})$ that are differentiable with respect to their last argument. Define a norm on this space by

$$\|f\|_{\mathcal{F}} = \sup_{\|Bg\|=1} \left\| B \int g(x', i') f(x', i'|\cdot) dx' di' \right\| + \sup_{\|Bg\|=1} \left\| B \int g(x', i') \frac{\partial f}{\partial i}(x', i'|\cdot) dx' di' \right\|.$$

Theorem 4 (Asymptotic distribution of the payoff function parameters). *If assumptions A1, A2, and the conditions of lemma 1 hold and,*

A'1 The estimates of $\boldsymbol{\eta}$, f , and $\frac{\partial f}{\partial i}$ converge faster than $n^{-1/4}$, i.e.

$$\begin{aligned} \|B(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0)\| &= o_p(n^{-1/4}) \\ \left\| \hat{f} - f_0 \right\|_{\mathcal{F}} &= o_p(n^{-1/4}) \end{aligned}$$

A'2 $\frac{\partial \Pi}{\partial \theta}$ and $\frac{\partial^2 \Pi}{\partial i \partial \theta}$ are Lipschitz continuous with respect to θ in a neighborhood of θ_0 in that

$$\left\| B \left(\frac{\partial \Pi}{\partial \theta}(\cdot; \theta_1) - \frac{\partial \Pi}{\partial \theta}(\cdot; \theta_2) \right) \right\| \lesssim |\theta_1 - \theta_2|$$

and

$$\left\| B \left(\frac{\partial^2 \Pi}{\partial i \partial \theta}(\cdot; \theta_1) - \frac{\partial^2 \Pi}{\partial i \partial \theta}(\cdot; \theta_2) \right) \right\| \lesssim |\theta_1 - \theta_2|$$

A'3 $\left| \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta}(x, i) \right| \leq C(x, i)$ with $(BC)(x, i)$ and $\|BC\|$ finite.

A'4 \hat{f} and $\hat{\eta}$ are bounded with probability approaching one

$$\|B_n \hat{\eta}\| \leq M$$

and

$$\|\hat{f}\|_{\mathcal{F}} \leq M.$$

A'5

$$\sqrt{n} \left[\begin{array}{l} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), BD_{\eta}^m(\theta_0, h_0)(\hat{\eta} - \eta_0) \right\rangle \\ \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), BD_f^m(\theta_0, h_0)(\hat{f} - f_0) \right\rangle \end{array} \right] \xrightarrow{d} N(0, \Omega_B)$$

Proof. We will verify the conditions of theorem 3. We have assumed A1 and A2. Also, condition A'1 is basically a restatement of A3. To verify A4, first note that the derivatives of m_a , m_v , and m_f are simply functions of the derivatives of \mathcal{V} . For m_f , it is immediate that

$$\begin{aligned} \frac{\partial m_f}{\partial \theta} &= \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta} \\ D_{\eta}^{m_f}(\eta_1 - \eta) &= \left(\frac{\partial \Pi}{\partial \eta} \circ \eta + \frac{\partial D_{\eta}^{\mathcal{V}}}{\partial i} \right) (\eta_1 - \eta) \\ D_f^{m_f}(f_1 - f) &= \frac{\partial D_f^{\mathcal{V}}}{\partial i} (f_1 - f). \end{aligned}$$

For m_a , a simple application of the implicit function theorem gives,

$$\begin{aligned} \frac{\partial m_a}{\partial \theta} &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta} \\ D_{\eta}^{m_a}(\eta_1 - \eta) &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \left(\frac{\partial \Pi}{\partial \eta} \circ \eta + \frac{\partial D_{\eta}^{\mathcal{V}}}{\partial i} \right) (\eta_1 - \eta) \\ D_f^{m_a}(f_1 - f) &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \frac{\partial D_f^{\mathcal{V}}}{\partial i} (f_1 - f). \end{aligned}$$

For m_v , using the envelope and mean value theorems gives,

$$\begin{aligned}\frac{\partial m_a}{\partial \theta} &= \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta}(i - i^*) \\ D_{\boldsymbol{\eta}}^{m_a}(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) &= \left(\frac{\partial \Pi}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} + \frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i} \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta})(i - i^*) \\ D_f^{m_a}(f_1 - f) &= \frac{\partial D_f^{\mathcal{V}}}{\partial i}(f_1 - f)(i - i^*),\end{aligned}$$

where $i^*(i, x; \theta, f, \boldsymbol{\eta}) = \arg \max_{\tilde{i}} \mathcal{V}(x, \boldsymbol{\eta}(i, x), \tilde{i}, i; \theta, f, \boldsymbol{\eta})$. Conditions A4.i and A4.ii then follow from the Lipschitz continuity of $\frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i}$ and $\frac{\partial D_f^{\mathcal{V}}}{\partial i}$ shown in lemma 1. Similarly, condition A4.iii follows directly from lemma 1 and assumption A'2.

A4.iv can be verified through repeated application of the dominated convergence theorem. In particular if $|\frac{\partial m}{\partial \theta}(x, i; \theta)| \leq C(x, i)$ with $BC(x, i) < \infty$ and $|BC| < \infty$, then

$$\begin{aligned}\frac{\partial}{\partial \theta} Bm &= \frac{\partial}{\partial \theta} \int m(x', i') b(x', i'; x, i) dx' di' \\ &= \int \frac{\partial m}{\partial \theta}(x', i') b(x', i'; x, i) dx' di' .\end{aligned}$$

Also,

$$\begin{aligned}\int \frac{\partial}{\partial \theta} (Bm)(x, i)^2 dx di &= \int 2(B \frac{\partial m}{\partial \theta})(x, i)(Bm)(x, i) dx di \\ &\leq \|BC\| \|Bm\|\end{aligned}$$

so

$$\frac{\partial}{\partial \theta} \|Bm\|^2 = 2 \left\langle B \frac{\partial m}{\partial \theta}, Bm \right\rangle .$$

Thus, we just have to show for each of the three possibilities of m that $|\frac{\partial m}{\partial \theta}(x, i; \theta)| \leq C(x, i)$ with $BC(x, i) < \infty$ and $|BC| < \infty$.

Condition A5 follows from A'4 and the following inequalities. To show the first part of A5 note that

$$\begin{aligned}\|B_n m(\cdot; \theta, \hat{h})\| &\lesssim \left\| B_n \frac{\partial \Pi}{\partial i} \circ \hat{\boldsymbol{\eta}} \right\| + \left\| B_n \delta \frac{\partial \mathcal{E}_f}{\partial i} \left(\mathbf{I} - \delta \mathcal{E}_f \right)^{-1} \Pi \circ \hat{\boldsymbol{\eta}} \right\| \\ &\lesssim \|B_n(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0)\| + \frac{\delta}{1 - \delta} \left(\|B_n\|_{op} \|\hat{f} - f\|_{\mathcal{F}} + \|B_n(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0)\| \right).\end{aligned}$$

Similarly, the second part of A5 follows from

$$\begin{aligned} \left\| B_n D_h^m(\hat{h} - h) \right\| &\lesssim \left\| B_n \frac{\partial D \mathcal{V}}{\partial i}(\hat{\eta} - \eta_0) \right\| + \left\| B_n \frac{\partial D_f \mathcal{V}}{\partial i}(\hat{f} - f_0) \right\| \\ &\lesssim \frac{\delta}{1 - \delta} \left(\|B_n(\hat{\eta} - \eta_0)\| + \|\hat{f} - f_0\|_{\mathcal{F}} \right). \end{aligned}$$

Finally, A'5 implies A6. ■

1.5.1 Asymptotic equivalence

In this subsection we show that the three types of moment conditions result in asymptotically equivalent estimators when appropriate weighting operators are used. This fact follows directly from results derived in the proofs of Theorems 3 and 4.

Corollary 1 (Asymptotic equivalence). *Suppose the conditions of Theorem 4 hold. Let $\hat{\theta}_f(B)$ be the estimate from using moment conditions m_f and weighting operator B . Similarly, define $\hat{\theta}_a(B)$ and $\hat{\theta}_v(B)$. Then*

$$\left\| \hat{\theta}_f(B) - \hat{\theta}_a(B \frac{\partial^2 \mathcal{V}}{\partial i}) \right\| = o_p(n^{-1/2}) \left\| \hat{\theta}_f(B) - \hat{\theta}_v(B(i^* - i)^{-1}) \right\| = o_p(n^{-1/2}).$$

Proof. The proof of Theorem 3 shows that

$$\begin{aligned} \sqrt{n}(\hat{\theta} - \theta_0) &= -\sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\ &\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(1). \end{aligned}$$

Substituting in the formulas for D_h^m and $\frac{\partial m}{\partial \theta}$ given in the proof of theorem 4 verifies the proposition. ■

1.5.2 A kernel estimator

This section describes kernel estimates of the transition density and policy function that satisfy the conditions for asymptotic normality given in theorem 4. We assume that our data comes from a balanced panel of length T . There are M independent markets, each with F firms. We will state results for M and F fixed with $T \rightarrow \infty$, but our proofs could easily be adapted to situations with $M \rightarrow \infty$ and T either growing or fixed, provided that at least $MT \rightarrow \infty$. We will let $n = MTF$.

When the distinction among M , T , and F is unimportant, we will let

$$\sum_{j=1}^n x_j = \sum_{t=1}^T \sum_{m=1}^M \sum_{f=1}^F x_{t,m,f}.$$

Our estimate of the policy function is

$$\hat{\eta}(x, i) = F_{\eta}^{-1} \left(\frac{\sum_{j=1}^n \mathbf{1}(i_j \leq i) k((x_j - x)/h_n)}{\sum_{j=1}^n k((x_j - x)/h_n)} \right),$$

where k is a kernel satisfying the conditions states below and h_n is the bandwidth. We will let

$$F(i|\widehat{x})f(x) = \frac{1}{h_n n} \sum_{j=1}^n \mathbf{1}(i_j \leq i) k((x_j - x)/h_n)$$

and

$$\hat{f}(x) = \frac{1}{h_n n} \sum_{j=1}^n k((x_j - x)/h_n).$$

Let $z = (x, i)$. The transition density is estimated by

$$\hat{f}(z'|z) = \frac{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t+1,m,f} - z', z_{t,m,f} - i)/h_n)}{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t,m,f} - z)/h_n)},$$

and the estimate of its derivative is simply the derivative of the estimate of the transition density,

$$\frac{\partial \hat{f}}{\partial i}(z'|z) = \frac{\partial \sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t+1,m,f} - z', z_{t,m,f} - z)/h_n)}{\partial i \sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t,m,f} - z)/h_n)}.$$

We will denote the joint density of z_{t+1} and z_t by f_{zz} , the marginal density of z by f_z , and the marginal density of x by f_x . We denote kernel estimates of these densities by \hat{f}_{zz} , \hat{f}_z , and \hat{f}_x .

To apply Theorem 4 to conclude that $\hat{\theta}$ is \sqrt{n} asymptotically normal; we must verify that the kernel estimates above satisfy the rate condition A'1 and the asymptotic normality condition A'5. The rate condition can be verified by using available results on uniform convergence rates for kernel estimators. We employ the results of Hansen (2008) for this purpose. Showing condition A'5 is more involved. In independent work, Srisuma and Linton (2010) propose a similar kernel estimator for the case with continuous states and discrete actions, and Srisuma (2010) develops a similar estimator for models with discrete states and continuous controls. These papers employ U-statistics to show \sqrt{n} asymptotic normality of the parameters of the payoff function. We take a different

approach. We extend the uniform central limit theorem for smoothed empirical processes of van der Vaart (1994) and Giné and Nickl (2007) to allow for dependent data. We then utilize this theorem to verify condition A'5.

Condition A'5 is that

$$\sqrt{n} \left\langle B^* B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \xrightarrow{d} N(0, \Omega_B).$$

To verify this condition, we first observe that this inner-product has the following integral form:

$$\begin{aligned} \left\langle B^* B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle &= \\ &= \int c_1(z', z) \left(\hat{f}_{z',z}(z', z) - f_{z',z}(z', z) \right) dz' dz + \\ &+ \int c_2(z', z) \left(\hat{f}_z(z) - f_z(z) \right) dz' dz + \\ &+ \int c_3(z', z) \left(\frac{\partial \hat{f}_{z',z}(z', z)}{\partial i} - \frac{\partial f_{z',z}(z', z)}{\partial i} \right) dz' dz + \\ &+ \int c_4(z', z) \left(\frac{\partial \hat{f}_z(z)}{\partial i} - \frac{\partial f_z(z)}{\partial i} \right) dz' dz + \\ &+ \int c_5(z', z) \left(\hat{f}_x(x) - f_x(x) \right) dz' dz + \\ &+ \int c_6(z', z) \left(F(\widehat{i|x})f(x) - F(i|x)f(x) \right) dz' dz + o_p(n^{-1/2}), \end{aligned} \quad (1.5.11)$$

where $c_j(\cdot)$ are functions that depend on B , $\frac{\partial m}{\partial \theta}$ and D_h^m . The exact form of $c_j(\cdot)$ and formal verification of (1.5.11) are given in Lemma 2 below. Given (1.5.11), we can then apply the following theorem, which is an extension of van der Vaart (1994) and Giné and Nickl (2007) to allow for dependent data and estimates of derivatives of densities as well as densities themselves. This result may be of independent interest.

Theorem 5 (Extension of van der Vaart (1994) and Giné and Nickl (2007)). *Let $X = (X_i)_{i \in \mathbb{Z}}$ be a stationary sequence with marginal distribution \mathbb{P} and β -mixing coefficients $\beta(k)$. Let $\mathcal{F} \subset \mathcal{L}^2(\mathbb{P})$ be a translation invariant class of functions. Suppose there exists $2 < p < \infty$ such that $\sum_{k=1}^{\infty} k^{2/(2-p)} \beta(k) < \infty$ and \mathcal{F} is not too complex, namely the bracketing entropy integral is finite, $J_{[]}(\infty, \mathcal{F}, \mathcal{L}_p(\mathbb{P})) < \infty$. Let $\|f\|_{LR} = \sum_{i \in \mathbb{Z}} E[f(z_0)f(z_i)]$. Let $\{\mu_n\}_{n=1}^{\infty}$ converge weakly to the Dirac measure at zero, δ_0 , or one of its derivatives, $\partial^\alpha \delta_0$. Also assume that for all n , $\mu_n(\mathbb{R}^d) \leq C$,*

$\mathcal{F} \subseteq \mathcal{L}^1(|\mu_n|)$, and $\int \|f(\cdot - y)\|_{LR} d|\mu_n|(y) < \infty$ for all $f \in \mathcal{F}$. If

$$\sup_{f \in \mathcal{F}} \sum_{i \in \mathbb{Z}} E \left[\left(\int f(X_0 + y) - f(X_0) d\mu_n(y) \right) \left(\int f(X_i + y) - f(X_i) d\mu_n(y) \right) \right] \rightarrow 0 \quad (1.5.12)$$

and

$$\sup_{f \in \mathcal{F}} \left| E \int f(X + y) - f(X) d\mu_n(y) \right| \rightarrow 0 \quad (1.5.13)$$

then

$$\sqrt{n} (\mathbb{P}_n * \mu_n - \mathbb{P} * \mu_\infty) \rightsquigarrow \mathcal{G} \text{ in } \ell^\infty(\mathcal{F}),$$

where \mathcal{G} is the \mathbb{P} -Brownian bridge indexed by \mathcal{F} and μ_∞ is either δ_0 or $\partial^\alpha \delta_0$.

Proof. This is Theorem 2 of Giné and Nickl (2007), except it allows for dependent data and $\mu_n \rightarrow \partial^\alpha \delta_0$ in addition to $\mu_n \rightarrow \delta_0$. Allowing for $\partial^\alpha \delta_0$ requires no modification of the proof. Allowing for dependent data simply requires substituting the empirical process results for independent data used by van der Vaart (1994) and Giné and Nickl (2007) with a suitable empirical process result for dependent data. A variety of such results are available; we use Theorem 11.22 of Kosorok (2008). This theorem says that given the condition on β -mixing coefficients and bracketing entropy above

$$\sqrt{n} (\mathbb{P}_n - \mathbb{P}) \rightsquigarrow \xi \text{ in } \ell^\infty(\mathcal{F}),$$

where ξ is a tight, mean zero Gaussian process with covariance

$$V(f, g) = \sum_{i \in \mathbb{Z}} \text{cov}(f(X_0)g(X_i)).$$

Then, by following the arguments used to prove Theorem 1 and Lemmas 1 and 2 of Giné and Nickl (2007), but with $\|\cdot\|_{LR}$ in place of $\|\cdot\|_{2, \mathbb{P}}$, we can show that

$$\mathbb{G}_n(f, \mu) = \sqrt{n} (\mathbb{P}_n - \mathbb{P}) \int f(X + y) d\mu(y)$$

converges in distribution in $\ell^\infty(\mathcal{F} \times \mathcal{M})$. Finally, the proof of van der Vaart (1994), with his

semi-metric d redefined as

$$d^2((f, \mu), (g, \nu)) = \sum_{i \in \mathbb{Z}} E \left[\left(\int f(X_0 + y) d\mu(y) - \int g(X_0 + y) d\nu(y) \right) \times \right. \\ \left. \times \left(\int f(X_i + y) d\mu(y) - \int g(X_i + y) d\nu(y) \right) \right],$$

leads to the conclusion. ■

We now state regularity conditions that are sufficient for (1.5.11) to hold and Theorem 5 to apply.

Assumptions K (Kernel estimator).

K1 The densities $f(z', z)$ and $f(z)$ are a ≥ 2 -times continuously differentiable.

K2 The data is comes from M markets with F firms and $T \rightarrow \infty$ periods. $\{z_{t,m,f}\}_{t \in \mathbb{Z}}$ is stationary and strongly mixing with strong mixing coefficients α_m that decay at rate β , i.e. $\alpha_m \lesssim m^{-\beta}$ for some $m > 1$. There exists $q > 0$ such that $\beta > 1 + \frac{d}{q} + d$ and

$$\frac{\beta - 1 - \frac{d}{q} - d}{d(\beta + 3 - d)} > \frac{1}{4r}.$$

K3 With probability approaching one,

$$\int \left| \left[B_n^* B_n D_h^m (\hat{h}(z', z) - h(z', z)) \right] (z_0) \right| dz' dz < C$$

uniformly in z_0 .

K4 The dimension of $z = (x, i)$ is d .

K5 $c_j(z', z)$ is $s > d/2$ times differentiable with uniformly bounded derivatives.

K6 The kernel is of order $r = a + s - k > a + d/2$.

K7 The bandwidth is such that $h_n^{a+s-k} n^{d/2} \rightarrow \infty$

K8 The weighting operator is such that $\int b_n(z_0, z_1) b_n(z, z_1) dz_1$ has compact convex support and vanishes at the boundary of its support.

Lemma 2 shows that these conditions are sufficient for (1.5.11).

Lemma 2 (Verification of (1.5.11)). *If K1-K7 hold, then the following functions exist*

$$\begin{aligned} d_1(z', z; z_0) &= \delta \left[\frac{\partial \mathcal{E}_f}{\partial i} \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \right) (z|z_0) [\delta (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_{\theta_0} \circ \boldsymbol{\eta}_0] (z') \right. \\ d_2(z', z; z_0) &= \left. [\delta (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_{\theta_0} \circ \boldsymbol{\eta}_0] (z') \right. \\ d\eta(z', z; z_0) &= \left. \frac{\partial^2 \Pi_{\theta_0}}{\partial \eta \partial i} \circ \boldsymbol{\eta}_0(z_0) + \delta \left[\mathcal{E}_f \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\frac{\partial \Pi_{\theta_0}}{\partial \eta} \circ \boldsymbol{\eta}_0 \right) \right) \right] (z), \right. \end{aligned}$$

and equation (1.5.11) holds with

$$\begin{aligned} c_j(z', z) &= \left\langle \frac{\partial m}{\partial \theta}, (B^* B \tilde{c}_j(z', \cdot))(z, \cdot) \right\rangle \\ &= \int \frac{\partial m}{\partial \theta}(z_0) \left(\int b(z_0, z_1) b(z, z_1) dz_1 \right) \tilde{c}(z', z; z_0) dz_0 \end{aligned}$$

where

$$\begin{aligned} \tilde{c}_1(z', z; z_0) &= d_1(z', z; z_0) \frac{1}{f_z(z)} - d_2(z', z; z_0) \frac{\partial f_{zz} / \partial i(z', z)}{f_z(z)^2} \\ \tilde{c}_2(z', z; z_0) &= -d_1(z', z; z_0) \frac{f_{zz}(z', z)}{f_z(z)^2} - \\ &\quad - d_2(z', z; z_0) \left(\frac{\partial f_{zz} / \partial i(z', z)}{f_{zz}(z', z)} + 2 \frac{f_{zz}(z', z) \partial f_z / \partial i(z)}{f_z(z)^2} \right) \\ \tilde{c}_3(z', z; z_0) &= d_2(z', z; z_0) \frac{1}{f_z(z)} \\ \tilde{c}_4(z', z; z_0) &= d_2(z', z; z_0) \frac{f_{zz}(z', z)}{f_z(z)} \\ \tilde{c}_5(z', z; z_0) &= d\eta(z', z; z_0) \frac{1}{f_x(x)} \\ \tilde{c}_6(z', z; z_0) &= -d\eta(z', z; z_0) \end{aligned}$$

Proof. Equation (1.5.11) is derived by repeatedly applying the mean value theorem to write $\hat{h} - h$ as a sum of $\hat{f}_{zz} - f_{zz}$, $\hat{f}_z - f_z$, $\hat{f}_x - f_x$, and their derivatives. Condition K3 ensures that we can freely interchange the order of integration while doing so. The definitions of d_1 , d_2 , and $d\eta$ are then verified by straightforward calculation. Condition K7 and K2, along with the uniform convergence rate results of Hansen (2008), are sufficient to show that the remainder in (1.5.11) is $o_p(n^{-1/2})$. ■

Combining this lemma with Theorems 5 and 4 shows that the kernel estimator of $\hat{\theta}$ is asymptotically normal. Theorem 6 states this result.

Theorem 6 (Asymptotic distribution for kernel estimator). *If conditions A'2-A'4 and K1-K7*

hold, then

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, M_B^{-1} \Omega_B M_B^{-1}),$$

where

$$\Omega_b = \sum_{t=-\infty}^{\infty} \sum_{f_2, f_1=1}^F \left[\begin{array}{l} \text{cov}(c_1(z_{0+1, m_1, f}, z_{0, m, f_1}), c_1(z_{t+1, m_2, f}, z_{t, m, f_2})) + \\ \text{cov}\left(\int c_2(z', z_{0, m, f_1}) dz', \int c_2(z', z_{t, m, f_2}) dz'\right) + \\ \text{cov}\left(\frac{\partial c_3}{\partial i}(z_{0+1, m_1, f}, z_{0, m, f_1}), \frac{\partial c_3}{\partial i}(z_{t+1, m_2, f}, z_{t, m, f_2})\right) + \\ \text{cov}\left(\int \frac{\partial c_4}{\partial i}(z', z_{0, m, f_1}) dz', \int \frac{\partial c_4}{\partial i}(z', m_2, f, z_{t, m, f_2}) dz'\right) + \\ \text{cov}\left(\int c_5(z', (x_{0, m, f_1}, i)) dz' di, \int c_5(z', (x_{t, m, f_2}, i)) dz' di\right) + \\ \text{cov}\left(\int \frac{\partial c_6}{\partial i}(z', z_{0, m, f_1}) dz', \int \frac{\partial c_6}{\partial i}(z', m_2, f, z_{t, m, f_2}) dz'\right) \end{array} \right].$$

Proof. As in the proof of Lemm 2, conditions K1 and K7, along with the uniform convergence rate results of Hansen (2008), are sufficient to show that the rate assumption, A'1, holds. We have assumed A'2-A'4. All that remains is to show A'5 holds with the stated Ω_B . By Lemma 2, equation (1.5.11) holds. Applying Theorem 5 to (1.5.11) yields the desired conclusion. All that remains is to show that the conditions of Theorem 5 are satisfied.

Let $\mathcal{F} = \{c_j(z', z)\}_{j=1}^6$. This is finite class and trivially satisfies $J_{\square}(\infty, \mathcal{F}, \mathcal{L}_p(\mathbb{P})) < \infty$. Conditions K1 and K5 are sufficient to apply Proposition 1 of Giné and Nickl (2007) to show that (1.5.12) holds for c_1 - c_6 . Conditions K5-K7 are sufficient to apply Theorem 6 of Giné and Nickl (2007) to show that (1.5.13) holds for c_1 , c_2 , c_5 , and c_6 . These same conditions along with K8 and a slight modification of Theorem 6 of Giné and Nickl (2007) shows (1.5.13) for c_3 and c_4 . ■

1.6 Conclusion

We have shown that the payoff function in dynamic games with continous states and controls is nonparametrically identified by the observed distribution of states and controls. The key identifying restrictions include that one of the partial derivatives of the payoff function is known, that there is some component of the state space that enters the policy function, but not the payoff function directly, and that a certain integral operator of the second kind is invertible. We have also developed a semiparametric estimator for the model. In the first step the transition densities and policy function are estimated nonparametrically. In the second step, the parameters of the payoff function

are estimated from the optimality conditions of the model. we gave high-level conditions on the first step nonparametric estimates for the parameter estimates to be consistent and parameters to be \sqrt{n} asymptotically normal, and we have shown that a kernel based estimator satisfies these conditions.

Chapter 2

Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from The U.K. Annuity Market, coauthored with Liran Einav and Amy Finkelstein

Abstract. Much of the extensive empirical literature on insurance markets has focused on whether adverse selection can be detected. Once detected, however, there has been little attempt to quantify its welfare cost, or to assess whether and what potential government interventions may reduce these costs. To do so, we develop a model of annuity contract choice and estimate it using data from the U.K. annuity market. The model allows for private information about mortality risk as well as heterogeneity in preferences over different contract options. We focus on the choice of length of guarantee among individuals who are required to buy annuities. The results suggest that asymmetric information along the guarantee margin reduces welfare relative to a first best symmetric information benchmark by about £127 million per year, or about 2 percent of annuitized wealth. We also find that by requiring that individuals choose the longest guarantee period allowed, mandates could achieve the first-best allocation. However, we estimate that other mandated guarantee lengths would have detrimental

effects on welfare. Since determining the optimal mandate is empirically difficult, our findings suggest that achieving welfare gains through mandatory social insurance may be harder in practice than simple theory may suggest.

2.1 Introduction

Ever since the seminal works of Akerlof (1970) and Rothschild and Stiglitz (1976), a rich theoretical literature has emphasized the negative welfare consequences of adverse selection in insurance markets and the potential for welfare-improving government intervention. More recently, a growing empirical literature has developed ways to detect whether asymmetric information exists in particular insurance markets (Chiappori and Salanie (2000), Finkelstein and McGarry (2006)). Once adverse selection is detected, however, there has been little attempt to estimate the magnitude of its efficiency costs, or to compare welfare in the asymmetric information equilibrium to what would be achieved by potential government interventions. In an attempt to start filling this gap, this paper develops an empirical approach that can quantify the efficiency cost of asymmetric information and the welfare consequences of government intervention.¹

We apply our approach to the semi-compulsory market for annuities in the United Kingdom. Individuals who have accumulated funds in tax-preferred retirement saving accounts (the equivalents of an IRA or 401(k) in the United States) are required to annuitize their accumulated lump sum balances at retirement. These annuity contracts provide a survival-contingent stream of payments. As a result of these requirements, there is a sizable volume in the market. In 1998, new funds annuitized in this market totalled £6 billion (Association of British Insurers (1999)).

⁰We are grateful to three anonymous referees and Steve Berry (the Editor) for many useful comments and suggestions. We also thank James Banks, Richard Blundell, Jeff Brown, Peter Diamond, Carl Emmerson, Jerry Hausman, Igal Hendel, Wojciech Kopczuk, Jonathan Levin, Alessandro Lizzeri, Ben Olken, Casey Rothschild, and many seminar participants for helpful comments, and to several patient and helpful employees at the firm whose data we analyze. Financial support from the National Institute of Aging grant #R01 AG032449 (Einav and Finkelstein), the National Science Foundation grant #SES-0643037 (Einav), the Social Security Administration grant #10-P-98363-3 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium (Einav and Finkelstein), and the the Alfred P. Sloan Foundation (Finkelstein) is gratefully acknowledged. Einav also acknowledges the hospitality of the Hoover Institution. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

¹More recently, several new working papers have presented additional attempts to quantify the efficiency cost of adverse selection in annuities (Hosseini (2008)) and in health insurance (Carlin and Town (2007), Bundorf, Levin, and Mahoney (2008), Einav, Finkelstein, and Cullen (2008), and Lustig (2008)).

Although they are required to annuitize their balances, individuals are allowed choice in their annuity contract. In particular, they can choose from among guarantee periods of 0, 5, or 10 years. During a guarantee period, annuity payments are made (to the annuitant or to his estate) regardless of the annuitant's survival. The choice of a longer guarantee period comes at the cost of lower annuity payments while alive. When annuitants and insurance companies have symmetric information about an annuitant's mortality rate, a longer guarantee is more attractive to an annuitant who cares more about their wealth when they die relative to consumption while alive; as a result, the first-best guarantee length may differ across annuitants. When annuitants have private information about their mortality rate, a longer guarantee period is also more attractive, all else equal, to individuals who are likely to die sooner. This is the source of adverse selection, which can affect the equilibrium price of guarantees and thereby distort guarantee choices away from the first-best symmetric information allocation.

The pension annuity market provides a particularly interesting setting in which to explore the welfare costs of asymmetric information and the welfare consequences of potential government intervention. Annuity markets have attracted increasing attention and interest as Social Security reform proposals have been advanced in various countries. Some proposals call for partly or fully replacing government-provided defined benefit, pay-as-you-go retirement systems with defined contribution systems in which individuals would accumulate assets in individual accounts. In such systems, an important question concerns whether the government would require individuals to annuitize some or all of their balance, and whether it would allow choice over the type of annuity product purchased. The relative attractiveness of these various options depends critically on consumer welfare in each alternative allocation.

In addition to their substantive interest, several features of annuities make them a particularly attractive setting for our purpose. First, adverse selection has already been detected and documented in this market along the choice of guarantee period, with private information about longevity affecting both the choice of contract and its price in equilibrium (Finkelstein and Poterba (2004, 2006)). Second, annuities are relatively simple and clearly defined contracts, so that modeling the contract choice requires less abstraction than in other insurance settings. Third, the case for moral hazard in annuities is arguably less compelling than for other forms of insurance; our ability to assume away moral hazard substantially simplifies the empirical analysis.

We develop a model of annuity contract choice and use it, together with individual-level data on annuity choices and subsequent mortality outcomes from a large annuity provider, to recover the

joint distribution of individuals' (unobserved) risk and preferences. Using this joint distribution and the annuity choice model, we compute welfare at the observed allocation, as well as allocations and welfare for counterfactual scenarios. We compare welfare under the observed asymmetric information allocation to what would be achieved under the first-best, symmetric information benchmark; this comparison provides our measure of the welfare cost of asymmetric information. We also compare equilibrium welfare to what would be obtained under mandatory social insurance programs; this comparison sheds light on the potential for welfare improving government intervention.

Our empirical object of interest is the joint distribution of risk and preferences. To estimate it, we rely on two key modeling assumptions. First, to recover mortality risk we assume that mortality follows a mixed proportional hazard model. Individuals' mortality tracks their own individual-specific mortality rates, allowing us to recover the extent of heterogeneity in (ex-ante) mortality rates from (ex-post) information about mortality realization. Second, to recover preferences, we use a standard dynamic model of consumption by retirees. In our baseline model we assume that retirees perfectly know their (ex-ante) mortality rate, which governs their stochastic time of death. This model allows us to evaluate the (ex-ante) value-maximizing choice of a guarantee period as a function of ex ante mortality rate and preferences for wealth at death relative to consumption while alive.

Given the above assumptions, the parameters of the model are identified from the variation in mortality and guarantee choices in the data, and in particular from the correlation between them. However, no modeling assumptions are needed to establish the existence of private information about the individual's mortality rate. This is apparent from the existence of (conditional) correlation between guarantee choices and ex post mortality in the data. Given the annuity choice model, rationalizing the observed choices with only variation in mortality risk is hard. Indeed, our findings suggest that unobserved mortality risk and preferences are both important determinants of the equilibrium insurance allocations.

We measure welfare in a given annuity allocation as the average amount of money an individual would need to make him as well off without the annuity as with his annuity allocation and his pre-existing wealth. We also examine the optimal government mandate among the currently existing guarantee options of 0, 5, or 10 years. In a standard insurance setting – that is, when all individuals are risk averse, the utility function is state-invariant, and there are no additional costs of providing insurance – it is well-known that mandatory (uniform) full insurance can achieve the first best allocation, even when individuals vary in their preferences. In contrast, we naturally view annuity

choices as governed by two different utility functions, one from consumption when alive and one from wealth when dead. In such a case, whether and which mandatory guarantee can improve welfare gains relative to the adverse selection equilibrium is not clear without more information on the cross-sectional distribution of preferences and mortality risk. The investigation of the optimal mandate – and whether it can produce welfare gains relative to the adverse selection equilibrium – therefore becomes an empirical question.

While caution should always be exercised when extrapolating estimates from a relatively homogeneous subsample of annuitants of a single firm to the market as a whole, our baseline results suggest that a mandatory social insurance program that required individuals to purchase a 10 year guarantee would increase welfare by about £127 million per year or £423 per new annuitant, while one that requires annuities to provide no guarantee would reduce welfare by about £107 million per year or £357 per new annuitant. Since determining which mandates would be welfare improving is empirically difficult, our results suggest that achieving welfare gains through mandatory social insurance may be harder in practice than simple theory would suggest. We also estimate welfare in a symmetric information, first-best benchmark. We find that the welfare cost of asymmetric information within the annuity market along the guarantee margin is about £127 million per year, £423 per new annuitant, or about two percent of the annuitized wealth in this market. Thus, we estimate that not only is a 10 year guarantee the optimal mandate, but also that it achieves the first best allocation.

To put these welfare estimates in context given the margin of choice, we benchmark them against the maximum money at stake in the choice of guarantee. This benchmark is defined as the additional (ex-ante) amount of wealth required to ensure that if individuals were forced to buy the policy with the least amount of insurance, they would be at least as well off as they had been. We estimate that the maximum money at stake in the choice of guarantee is only about 8 percent of the annuitized amount. Our estimates therefore imply that the welfare cost of asymmetric information is about 25 percent of this maximum money at stake.

Our welfare analysis is based on a model of annuity demand. This requires assumptions about the nature of the utility functions that govern annuity choice, as well as assumptions about the expectation individuals form regarding their subsequent mortality outcomes. Data limitations, particularly lack of detail on annuitant's wealth, necessitate additional modeling assumptions. Finally, our approach requires several other parametric assumptions for operational and computational reasons. The assumptions required for our welfare analysis are considerably stronger than those that

have been used in prior work to test whether or not asymmetric information exists. This literature has tested for the existence of private information by examining the correlation between insurance choice and ex-post risk realization (Chiappori and Salanie (2000)). Indeed, the existing evidence of adverse selection along the guarantee choice margin in our setting comes from examining the correlation between guarantee choice and ex-post mortality (Finkelstein and Poterba (2004)). By contrast, our effort to move from testing for asymmetric information to quantifying its welfare implications requires considerably stronger modeling assumptions. Our comfort with this approach is motivated by a general “impossibility” result which we illustrate in the working paper version (Einav, Finkelstein, and Schripf (2007)): even when asymmetric information is known to exist, the reduced form equilibrium relationship between insurance coverage and risk occurrence does not permit inference about the efficiency cost of this asymmetric information without strong additional assumptions.

Of course, a critical question is how important our particular assumptions are for our central results regarding welfare. We therefore explore a range of possible alternatives, both for the appropriate utility model and for our various parametric assumptions. We are reassured that our central results are quite stable. In particular, the finding that the 10 year guarantee is the optimal mandate, and achieves virtually the same welfare as the first best outcome, persists under all the alternative specifications that we have tried. However, the quantitative estimates of the welfare cost of adverse selection can vary with the modeling assumptions by a non trivial amount; more caution should therefore be exercised in interpreting these quantitative estimates.

The rest of the paper proceeds as follows. Section 2.2 describes the environment and the data. Section 2.3 describes the model of guarantee choice, presents its identification properties, and discusses estimation. Section 2.4 presents our parameter estimates and discusses their in-sample and out-of-sample fit. Section 2.5 presents the implications of our estimates for the welfare costs of asymmetric information in this market, as well as the welfare consequences of potential government policies. The robustness of the results is explored in Section 2.6. Section 2.7 concludes by briefly summarizing our findings and discussing how the approach we develop can be applied in other insurance markets, including those where moral hazard is likely to be important.

2.2 Data and environment

Environment. All of the annuitants we study are participants in the semi-compulsory market for annuities in the U.K.. In other words, they have saved for retirement through tax-preferred defined contribution private pensions (the equivalents of an IRA or 401(k) in the United States) and are therefore required to annuitize virtually all of their accumulated balances.² They are however offered choice over the nature of their annuity product. We focus on the choice of the length of the guarantee period, during which annuity payments are made (to the annuitant or to his estate) regardless of annuitant survival. Longer guarantees therefore trade off lower annuity payments in every period the annuitant is alive in return for payments in the event that the annuitant dies during the guarantee period.

The compulsory annuitization requirement is known to individuals at the time (during working age) that they make their pension savings contributions, although of course the exact nature of the annuity products (and their pricing) that will be available when they have to annuitize is uncertain. Choices over annuity products are only made at the time of conversion of the lump-sum defined contribution balances to an annuity and are based on the products and annuity rates available at that time.

All of our analysis takes the pension contribution decisions of the individual during the accumulation phase (as well as their labor supply decisions) as given. In other words, in our analysis of welfare under counterfactual pricing of the guarantee options, we do not allow for the possibility that the pre-annuitization savings and labor supply decisions may respond endogenously to the change in guarantee pricing. This is standard practice in the annuity literature (Brown (2001), Davidoff, Brown, and Diamond (2005), and Finkelstein, Poterba, and Rothschild (2009)). In our context, we do not think it is a particularly heroic assumption. For one thing, as we will discuss in more detail in Section 2.5.1, the maximum money at stake in the choice over guarantee is only about 8 percent of annuitized wealth under the observed annuity rates (and only about half that amount under the counterfactual rates we compute); this should limit any responsiveness of pre-annuitization decisions to guarantee pricing. Moreover, many of these decisions are made decades before annuitization and therefore presumably factor in considerable uncertainty (and discounting) of future guarantee prices.

Data and descriptive statistics. We use annuitant-level data from one of the largest annuity

²For more details on these rules, see Appendix A and Finkelstein and Poterba (2002).

providers in the U.K. The data contain each annuitant's guarantee choice, several demographic characteristics (including everything on which annuity rates are based), and subsequent mortality. The data consist of all annuities sold between 1988 and 1994 for which the annuitant was still alive on January 1, 1998. We observe age (in days) at the time of annuitization, the gender of the annuitant, and the subsequent date of death if the annuitant died before the end of 2005.

For analytical tractability, we make a number of sample restrictions. In particular, we restrict our sample to annuitants who purchase at age 60 or 65 (the modal purchase ages), and who purchased a single life annuity (that insures only his or her own life) with a constant (nominal) payment profile.³ Finally, the main analysis focuses on the approximately two-thirds of annuitants in our sample who purchased an annuity with a pension fund that they had accumulated within our company; in Section 2.6 we re-estimate the model for the remaining individuals who had brought in external funds. Appendix A.1 discusses these various restrictions in more detail; they are made so that we can focus on the purchase decisions of a relatively homogenous subsample.

Table 2.1 presents summary statistics for the whole sample and for each of the four age-gender combinations. Our baseline sample consists of over 9,000 annuitants. Sample sizes by age and gender range from a high of almost 5,500 for 65 year old males to a low of 651 for 65 year old females. About 87 percent of annuitants choose a 5 year guarantee period, 10 percent choose no guarantee, and only 3 percent choose the 10 year guarantee. These are the only three options available to annuitants in our sample and the focus of our subsequent analysis.

Given our sample construction described above, our mortality data are both left-truncated and right-censored, and cover mortality outcomes over an age range of 63 to 83. About one-fifth of our sample dies between 1998 and 2005. As expected, death is more common among men than women, and among those who purchase at older ages.

There is a general pattern of higher mortality among those who purchase 5 year guarantees than those who purchase no guarantees, but no clear pattern (possibly due to the smaller sample size) of mortality differences for those who purchase 10 year guarantees relative to either of the other two options. This mortality pattern as a function of guarantee persists in more formal hazard modeling that takes account of the left truncation and right censoring of the data (not shown).⁴

³Over 90 percent of the annuitants in our firm purchase policies that pay a constant nominal payout (rather than policies that escalate in nominal terms). This is typical of the market as a whole. Although escalating policies (including inflation-indexed policies) are offered by some firms, they are rarely purchased (Murthi, Orszag, and Orszag (1999), and Finkelstein and Poterba (2004)).

⁴Specifically, we estimated Gompertz and Cox proportional hazard models in which we included indicator variables for age at purchase and gender, as well as indicator variables for a 5 year guarantee and a 10 year guarantee. In both models, we found that the coefficient on the 5 year guarantee dummy was significantly different from that on the 0

As discussed in the introduction, the existence of a (conditional) correlation between guarantee choice and mortality – such as the higher mortality experienced by purchasers of the 5 year guarantee relative to purchasers of no guarantee – indicates the presence of private information about individual mortality risk in our data, and motivates our exercise. That is, this correlation between mortality outcomes and guarantee choices rules out a model in which individuals have no private information about their idiosyncratic mortality rates, and guides our modeling assumption in the next section that allow individuals to make their guarantee choices based on information about their idiosyncratic mortality rate.

Annuity rates. The company supplied us with the menu of annuity rates, that is the annual annuity payments per £1 of the annuitized amount. These rates are determined by the annuitant’s gender, age at the time of purchase, and the date of purchase; there are essentially no quantity discounts.⁵ All of these components of the pricing structure are in our data.

Table 2.2 shows the annuity rates by age and gender for different guarantee choices from January 1992; these correspond to roughly the middle of the sales period we study (1988-1994) and are roughly in the middle of the range of rates over the period. Annuity rates decline, of course, with the length of guarantee. Thus, for example, a 65 year old male in 1992 faced a choice among a 0 guarantee with an annuity rate of 0.133, a 5 year guarantee with a rate of 0.1287, and a 10 year guarantee with a rate of 0.1198. The magnitude of the rate differences across guarantee options closely tracks expected mortality. For example, our mortality estimates (discussed later) imply that for 60 year old females the probability of dying within a guarantee period of 5 and 10 years is about 4.3 and 11.4 percent, respectively, while for 65 year old males these probabilities are about 7.4 and 18.9 percent. Consequently, as shown in Table 2.2, the annuity rate differences across guarantee periods are much larger for 65 year old males than they are for 60 year old females.

The firm did not change the formula by which it sets annuity rates over our sample of annuity sales. Changes in nominal payment rates over time reflect changes in interest rates. To use such variation in annuity rates for estimation would require assumptions about how the interest rate that enters the individual’s value function covaries with the interest rate faced by the firm, and whether the individual’s discount rate covaries with these interest rates. Absent any clear guidance on these issues, we analyze the guarantee choice with respect to one particular menu of annuity

year guarantee dummy; however, the standard error on the coefficient on the 10 year guarantee dummy was high, so it wasn’t estimated to be significantly different from the 5 year guarantee dummy (or from the 0 year guarantee dummy as well).

⁵A rare exception on quantity discounts is made for individuals who annuitize an extremely large amount.

rates. For our baseline model we use the January 1992 menu shown in Table 2.2. In the robustness analysis, we show that the welfare estimates are virtually identical if we choose pricing menus from other points in time; this is not surprising since the *relative* payouts across guarantee choices is quite stable over time. For this reason, the results hardly change if we instead estimate a model with time-varying annuity rates, but constant discount factor and interest rate faced by annuitants (not reported).

Representativeness. Although the firm whose data we analyze is one of the largest U.K. annuity sellers, a fundamental issue when using data from a single firm is how representative it is of the market as a whole. We obtained details on market-wide practices from Moneyfacts (1995), Murthi, Orszag, and Orszag (1999), and Finkelstein and Poterba (2002).

On all dimensions we are able to observe, our sample firm appears typical of the industry as a whole. The types of contracts it offers are standard for this market. In particular, like all major companies in this market during our time period, it offers a choice of 0, 5, and 10 year guaranteed, nominal annuities.

The pricing practices of the firm are also typical. The annuitant characteristics that the firm uses in setting annuity rates (described above) are standard in the market. In addition, the level of annuity rates in our sample firm's products closely match industry-wide averages.

While market-wide data on characteristics of annuitants and the contracts they choose are more limited, the available data suggest that the annuitants in this firm and the contracts they choose are typical of the market. In our sample firm, the average age of purchase is 62, and 59 percent of purchasers are male. The vast majority of annuities purchased pay a constant nominal payment stream (as opposed to one that escalates over time), and provide a guarantee, of which the 5 year guarantee is by far the most common.⁶ These patterns are quite similar to those in another large firm in this market analyzed by Finkelstein and Poterba (2004), as well as to the reported characteristics of the broader market as described by Murthi, Orszag, and Orszag (1999).

Finally, the finding in our data of a higher mortality rate among those who choose a 5 year guarantee than those who choose no guarantee is also found elsewhere in the market. Finkelstein and Poterba (2004) present similar patterns for another firm in this market, and Finkelstein and Poterba (2002) present evidence on annuity rates that is consistent with such patterns for the

⁶These statistics are reported in Finkelstein and Poterba (2006) who also analyze data from this firm. These statistics refer to single life annuities, which are the ones we analyze here, but are (obviously) computed prior to the additional sample restrictions we make here (e.g., restriction to nominal annuities purchased at ages 60 or 65).

industry as a whole.

Thus, while caution must always be exercised in extrapolating from a single firm, the available evidence suggests that the firm appears to be representative – both in the nature of the contracts it offers and its consumer pool – of the entire market.

2.3 Model: specification, identification, and estimation

we start by discussing a model of guarantee choice for a particular individual. We then complete the empirical model by describing how (and over which dimensions) we allow for heterogeneity. We finish this section by discussing the identification of the model, our parameterization, and the details of the estimation.

2.3.1 A model of guarantee choice

We consider the utility-maximizing guarantee choice of a fully rational, forward looking, risk averse, retired individual, with an accumulated stock of wealth, stochastic mortality, and time-separable utility. This framework has been widely used to model annuity choices (Kotlikoff and Spivak (1981), Mitchell, Poterba, Warshawsky, and Brown (1999), Davidoff, Brown, and Diamond (2005)). At the time of the decision, the age of the individual is t_0 , and he expects a random length of life⁷ characterized by a mortality hazard κ_t during period $t > t_0$.⁸ We also assume that there exists time T after which individual i expects to die with probability one.

Individuals obtain utility from two sources. When alive, they obtain flow utility from consumption. When dead, they obtain a one-time utility that is a function of the value of their assets at the time of death. In particular, if the individual is alive as of the beginning of period $t \leq T$, his period t utility, as a function of his current wealth w_t and his consumption plan c_t , is given by

$$v(w_t, c_t) = (1 - \kappa_t) u(c_t) + \kappa_t b(w_t), \tag{2.3.1}$$

where $u(\cdot)$ is his utility from consumption and $b(\cdot)$ is his utility from wealth remaining after death. A positive valuation for wealth at death may stem from a number of possible underlying structural

⁷As might be expected, we can rule out a model with deterministic length of life and perfect foresight. Most individuals in the data choose a positive guarantee length and are alive at the end of it, thus violating such a model.

⁸Of course, one would expect some relationship between the individual's expectation and the actual underlying risk which governs the (stochastic) mortality outcome. We make specific assumptions about this relationship later, but for the purpose of modeling guarantee choice this is not important.

preferences, such as a bequest motive (Sheshinski (2006)) or a “regret” motive (Braun and Muermann (2004)). Since the exact structural interpretation is not essential for our goal, we remain agnostic about it throughout the paper.

In the absence of an annuity, the optimal consumption plan can be computed by solving the following program:

$$\begin{aligned} V_t^{NA}(w_t) &= \max_{c_t \geq 0} \left[(1 - \kappa_t)(u(c_t) + \delta V_{t+1}^{NA}(w_{t+1})) + \kappa_t b(w_t) \right] \\ \text{s.t. } w_{t+1} &= (1 + r)(w_t - c_t) \geq 0 \end{aligned} \quad (2.3.2)$$

where δ is the per-period discount rate and r is the per-period real interest rate. That is, we make the standard assumption that, due to mortality risk, the individual cannot borrow against the future. Since death is expected with probability one after period T , the terminal condition for the program is given by $V_{T+1}^{NA}(w_{T+1}) = b(w_{T+1})$.

Suppose now that the individual annuitizes a fraction η of his initial wealth, w_0 . Broadly following the institutional framework discussed earlier, individuals take the (mandatory) annuitized wealth as given. In exchange for paying ηw_0 to the annuity company at $t = t_0$, the individual receives a per-period real payout of z_t when alive. Thus, the individual solves the same problem as above, with two small modifications. First, initial wealth is given by $(1 - \eta)w_0$. Second, the budget constraint is modified to reflect the additional annuity payments z_t received every period.

For a given annuitized amount ηw_0 , consider a choice from a set $G \subseteq [0, T]$ of possible guarantee lengths; during the guaranteed period, the annuity payments are not survival-contingent. Each guarantee length $g \in G$ corresponds to a per-period payout stream of $z_t(g)$, which is decreasing in g ($\frac{\partial z_t(g)}{\partial g} < 0$ for any $t \geq t_0$). For each g , the optimal consumption plan can be computed by solving

$$\begin{aligned} V_t^{A(g)}(w_t) &= \max_{c_t \geq 0} \left[(1 - \kappa_t)(u(c_t) + \delta V_{t+1}^{A(g)}(w_{t+1})) + \kappa_t b(w_t + Z_t(g)) \right] \\ \text{s.t. } w_{t+1} &= (1 + r)(w_t + z_t(g) - c_t) \geq 0 \end{aligned} \quad (2.3.3)$$

where $Z_t(g) = \sum_{\tau=t}^{t_0+g} \left(\left(\frac{1}{1+r} \right)^{\tau-t} z_\tau(g) \right)$ is the present value of the remaining guaranteed payments. As before, since after period T death is certain and guaranteed payments stop for sure (recall, $G \subseteq [0, T]$), the terminal condition for the program is given by $V_{T+1}^{A(g)}(w_{T+1}) = b(w_{T+1})$.

The optimal guarantee choice is then given by

$$g^* = \arg \max_{g \in G} \left\{ V_{t_0}^{A(g)}((1 - \eta)w_0) \right\}. \quad (2.3.4)$$

Information about the annuitant's guarantee choice combined with the assumption that this choice was made optimally thus provides information about the annuitant's underlying preference and expected mortality parameters. Intuitively, everything else equal, a longer guarantee will be more attractive for individuals with higher mortality rate and for individuals who obtain greater utility from wealth after death. We later check that this intuition in fact holds in the context of the specific parametrized model we estimate.

2.3.2 Modeling heterogeneity

To obtain our identification result in the next section, we make further assumptions that allow only one-dimensional heterogeneity in mortality risk and one-dimensional heterogeneity in preferences across different individuals in the above model.

We allow for one-dimensional heterogeneity in mortality risk by using a mixed proportional hazard (MPH) model. That is, we assume that the mortality hazard rate of individual i at time t is given by

$$\theta_{it} \equiv \lim_{dt \rightarrow 0} \frac{P(m_i \in [t, t + dt] | x_i, m_i \geq t)}{dt} = \alpha_i \theta_0(x_i) \psi(t) \quad (2.3.5)$$

where m_i denotes the realized mortality date, $\psi(t)$ the baseline hazard rate, x_i is an observable that shifts the mortality rate, and $\alpha_i \in R_+$ represents unobserved heterogeneity. We also assume that individuals have perfect information about this stochastic mortality process; that is, we assume that individuals know their θ_{it} 's. This allows us to integrate over this continuous hazard rate to obtain the vector $\kappa^i \equiv (\kappa_t^i)_{t=t_0}^T$ that enters the guarantee choice model.

We allow for one-dimensional heterogeneity in preferences by assuming that $u^i(c)$ is homogeneous across all individuals and that $b^i(w)$ is the same across individuals up to a multiplicative factor. Moreover, we assume that

$$u^i(c) = \frac{c^{1-\gamma}}{1-\gamma} \quad (2.3.6)$$

and

$$b^i(w) = \beta_i \frac{w^{1-\gamma}}{1-\gamma}. \quad (2.3.7)$$

That is, we follow the literature and assume that all individuals have a (homogeneous) CRRA

utility function, but, somewhat less standard, we specify the utility from wealth at death using the same CRRA form with the same parameter γ , and allow (proportional) heterogeneity across individuals in this dimension, captured by the parameter β_i . One can interpret β_i as the relative weight that individual i puts on wealth when dead relative to consumption while alive. All else equal, a longer guarantee is therefore more attractive when β_i is higher. We note, however, that since $u(\cdot)$ is defined over a flow of consumption while $b(\cdot)$ is defined over a stock of wealth, it is hard to interpret the level of β_i directly. We view this form of heterogeneity as attractive both for intuition and for computation; in Section 2.6 we investigate alternative assumptions regarding the nature of preference heterogeneity.

Since we lack data on individuals' initial wealth w_0^i , we chose the utility function above to enable us to ignore w_0^i . Specifically, our specification implies that preferences are homothetic, and -- combined with the fact that guarantee payments are proportional to the annuitized amount (see Section 2.2) -- that an individual's optimal guarantee choice g_i^* is invariant to initial wealth w_0^i . This simplifies our analysis, as it means that in our baseline specification unobserved heterogeneity in initial wealth w_0^i is not a concern. It is, however, potentially an unattractive modeling decision, since it is not implausible that wealthier individuals care more about wealth after death. In Section 2.6 we explore specifications with non-homothetic preferences, but this requires us to make an additional assumption regarding the distribution of initial wealth. With richer data that included w_0^i we could estimate a richer model with non-homothetic preferences.

Finally, we treat a set of other parameters that enter the guarantee choice model as observable (known) and identical across all annuitants. Specifically, as we describe later, we use external data to calibrate the values for risk aversion γ , the discount rate δ , the fraction of wealth which is annuitized η , and the real interest rate r . While in principle we could estimate some of these parameters, they would be identified solely by functional form assumptions. We therefore consider it preferable to choose reasonable calibrated values, rather than impose a functional form that would generate these reasonable values. Some of these calibrations are necessitated by the limitations of our existing data. For example, we observe the annuitized amount so with richer data on wealth we could readily incorporate heterogeneity in η_i into the model.

2.3.3 Identification

In order to compute the welfare effect of various counterfactual policies, we need to identify the distribution (across individuals) of preferences and mortality rates. Here we explain how the as-

sumptions we made allow us to recover this distribution from the data we observe about the joint distribution of mortality outcomes and guarantee choices. We make the main identification argument in the context of a continuous guarantee choice set, a continuous mortality outcome, and no truncation or censoring. In the end of the section we discuss how things change with a discrete guarantee choice and mortality outcomes that are left truncated and right censored, as we have in our setting. This requires us to make additional assumptions, which we discuss later.

Identification with a continuous guarantee choice (and uncensored mortality outcomes). To summarize briefly, our identification is achieved in two steps. In the first step we identify the distribution of mortality rates from the observed marginal (univariate) distribution of mortality outcomes. This is possible due to the mixed proportional hazard model we assumed. In the second step we use the model of guarantee choice and the rest of the data – namely, the distribution of guarantee choices conditional on mortality outcomes – to recover the distribution of preferences and how it correlates with the mortality rate. The key conceptual step here is an exclusion restriction, namely that the mortality process is not affected by the guarantee choice. We view this “no moral hazard” assumption as natural in our context.

We start by introducing notation. The data about individual i is (m_i, g_i, x_i) , where m_i is his observed mortality outcome, $g_i \in G$ his observed guarantee choice, and x_i is a vector of observed (individual) characteristics. The underlying object of interest is the joint distribution of unobserved preferences and mortality rates $F(\alpha, \beta|x)$, as well as the baseline mortality hazard rate $(\theta_0(x_i)$ and $\psi(t))$. Identification requires that, with enough data, these objects of interest can be uniquely recovered.

At the risk of repetition, let us state four important assumptions that are key to the identification argument.

Assumption 1. *Guarantee choices are given by $g_i = g((\kappa_i^i)_{t=t_0}^T, \beta_i|x_i)$, which comes from the solution to the guarantee choice model of Section 2.3.1.*

Assumption 2 (MPH). *Mortality outcomes are drawn from a mixed proportional hazard (MPH) model. That is, $\theta_{it} = \alpha_i \theta_0(x_i) \psi(t)$ with $\alpha_i \in R_+$.*

Assumption 3 (No moral hazard). *m_i is independent of β_i , conditional on α_i .*

Assumption 4 (Complete information).

$$\kappa_t^i = \left(\exp \left(- \int_0^{t-1} \theta_{i\tau} d\tau \right) - \exp \left(- \int_0^t \theta_{i\tau} d\tau \right) \right) / \exp \left(- \int_0^{t-1} \theta_{i\tau} d\tau \right).$$

The first assumption simply says that all individuals in the data make their guarantee choices using the model. It is somewhat redundant, as it is only the model that allows us to define κ_i and β_i , but we state it for completeness. The second assumption (MPH) is key for the first step of the identification argument. This assumption will drive our ability to recover the distribution of mortality rates from mortality data alone. Although this is a non-trivial assumption, it is a formulation which is broadly used in much of the duration data literature (Van den Berg (2001)). We note that assuming that α_i is one-dimensional is not particularly restrictive, as any multi-dimensional α_i could be summarized by a one-dimensional statistic in the context of the MPH model.

The third assumption formalizes our key exclusion restriction. It states that θ_{it} is a sufficient statistic for mortality, and although β_i may affect guarantee choices g_i , this in turn doesn't affect mortality. In other words, if individuals counterfactually change their guarantee choice, their mortality experience will remain unchanged. This seems a natural assumption in our context. We note that, unconditionally, β_i could be correlated with mortality outcomes indirectly, through a possible cross-sectional correlation between α_i and β_i .

The fourth and final assumption states that individuals have perfect information about their mortality process; that is, we assume that individuals know their θ_{it} 's. This allows us to integrate over this continuous hazard rate to obtain the vector $\kappa^i \equiv (\kappa_t^i)_{t=t_0}^T$ that enters the guarantee choice model, so we can write $g(\alpha_i, \beta_i)$ instead of $g((\kappa_t^i)_{t=t_0}^T, \beta_i | x_i)$. This is however a very restrictive assumption, and its validity is questionable. Fortunately, we note that any other information structure – that is, any known (deterministic or stochastic) mapping from individuals' actual mortality process θ_{it} to their perception of it κ^i – would also work for identification. Indeed, we investigate two such alternative assumptions in Section 2.6.4. Some assumption about the information structure is required since we lack data on individuals' ex ante expectations about their mortality.

Before deriving our identification results, we should point out that much of the specification decisions, described in the previous section, were made to facilitate identification. That is, many of the assumptions were made so that preferences and other individual characteristics are known up to a one-dimensional unobservable β_i . This is a strong assumption, which rules out interesting cases of, for example, heterogeneity in both risk aversion and utility from wealth after death.

We now show identification of the model in two steps, in Proposition 1 and Proposition 2.

Proposition 1. *If (i) Assumption 2 holds; (ii) $E[\alpha] < \infty$; and (iii) $\theta_0(x_i)$ is not a constant, then the marginal distribution of α_i , $F_\alpha(\alpha_i)$, as well as $\theta_0(x_i)$ and $\psi(t)$, are identified – up to the*

normalizations $E[\alpha] = 1$ and $\theta_0(x_i) = 1$ for some i – from the conditional distribution of $F_m(m_i|x_i)$.

This proposition is the well known result that MPH models are non-parameterically identified. It was first proven by Elbers and Ridder (1982). Heckman and Singer (1984) show a similar result, but instead of assuming that α has a finite mean, they make an assumption about the tail behavior of α . Ridder (1990) discusses the relationship between these two assumptions, and Van den Berg (2001) reviews these and other results. The key requirement is that x_i (such as a gender dummy variable in our context) shifts the mortality distribution.

We can illustrate the intuition for this result using two values of $\theta_0(x_i)$, say θ_1 and θ_2 . The data then provides us with two distributions of mortality outcomes, $H^j(m) = F(m|\theta_0(x_i) = \theta_j)$ for $j = 1, 2$. With no heterogeneity in α_i , the MPH assumption implies that the hazard rates implied by $H^1(m)$ and $H^2(m)$ should be a proportional shift of each other. Once α_i is heterogeneous, however, the difference between θ_1 and θ_2 leads to differential composition of survivors at a given point in time. For example, if θ_1 is less than θ_2 , then high α_i people will be more likely to survive among those with θ_1 . Loosely, as time passes, this selection will make the hazard rate implied by Z_m^1 closer to that implied by Z_m^2 . With continuous (and uncensored) information about mortality outcomes, these differential hazard rates between the two distributions can be used to back out the entire distribution of α_i , $F_\alpha(\alpha_i)$, which will then allow us to know $\theta_0(x_i)$ and $\psi(t)$.

This result is useful because it shows that we can obtain the (marginal) distribution of α_i (and the associated $\theta_0(x_i)$ and $\psi(t)$ functions) from mortality data alone, i.e. from the marginal distribution of m_i . We now proceed to the second step, which shows that given $\theta_0(x_i)$, $\psi(t)$, and $F_\alpha(\cdot)$, the joint distribution $F(\alpha, \beta|x)$ is identified from the observed joint distribution of mortality and guarantee choices. Although covariates were necessary to identify $\theta_0(x_i)$, $\psi(t)$, and $F_\alpha(\cdot)$, they will play no role in what follows, so we will omit them for convenience for the remainder of this section.

Proposition 2. *If Assumptions 1-4 hold, then the joint distribution of mortality outcomes and guarantee choices identifies $P(g(\alpha, \beta) \leq y|\alpha)$. Moreover, if, for every value of α , $g(\alpha, \beta)$ is invertible with respect to β then $F_{\beta|\alpha}$ is identified.*

The proof is provided in Appendix A.2. Here we provide intuition, starting with the first part of the proposition. If we observed α_i , identifying $P(g(\alpha, \beta) \leq y|\alpha)$ would have been trivial. We could simply estimate the cumulative distribution function of g_i for every value of α_i off the data. While in practice we can't do exactly this because α_i is unobserved, we can almost do this using

the mortality information m_i and our knowledge of the mortality process (using Proposition 1). Loosely, we can estimate $P(g(\alpha, \beta) \leq y|m)$ off the data, and then “invert” it to $P(g(\alpha, \beta) \leq y|\alpha)$ using knowledge of the mortality process. That is, we can write

$$P(g(\alpha, \beta) \leq y|m) = \left(\int_{\alpha} f_m(m|\alpha) dF_{\alpha}(\alpha) \right)^{-1} \int_{\alpha} P(g(\alpha, \beta) \leq y|\alpha) f_m(m|\alpha) dF_{\alpha}(\alpha) \quad (2.3.8)$$

where the left hand side is known from the data, and $f_m(m|\alpha)$ (the conditional density of mortality date) and $F_{\alpha}(\alpha)$ are known from the mortality data alone (Proposition 1). The proof (in Appendix A.2) simply verifies that this integral can be inverted.

The second part of Proposition 2 is fairly trivial. If $P(g(\alpha, \beta) \leq y|\alpha)$ is identified for every α , and $g(\alpha, \beta)$ is invertible (with respect to β) for every α , then it is straightforward to obtain $P(\beta \leq y|\alpha)$ for every α . This together with the marginal distribution of α , which is identified through Proposition 1, provides the entire joint distribution.

One can see that the invertibility of $g(\alpha, \beta)$ (with respect to β) is important. The identification statement is stated in such a way because, although intuitive, proving that the guarantee choice is monotone (and therefore invertible) in β is difficult. The difficulty arises due to the dynamics and non-stationarity of the guarantee choice model, which require its solution to be numerical and make general characterization of its properties difficult. One can obtain analytic proofs of this monotonicity property in simpler (but empirically less interesting) environments (e.g., in a two period model, or in an infinite horizon model with log utility). We note, however, that we are reassured about our simple intuition based on numerical simulations; the monotonicity result holds for any specification of the model and/or values of the parameters that we have tried, although absent an analytical proof some uncertainty must remain regarding identification.

Implications of a discrete guarantee choice and censored mortality outcomes. In many applications the (guarantee) choice is discrete, so – due to its discrete nature – $g(\beta|\alpha)$ is only weakly monotone in β , and therefore not invertible. In that case, the first part of Proposition 2 still holds, but $P(\beta \leq y|\alpha)$ is identified only in a discrete set of points, so some parametric assumptions will be needed to recover the entire distribution of β , conditional on α . In our specific application, there are only three guarantee choices, so we can only identify the marginal distribution of α , $F(\alpha)$, and, for every value of α , two points of the conditional distribution $F_{\beta|\alpha}$. We therefore recover the entire joint distribution by making a parametric assumption (see below) that essentially allows us to interpolate $F_{\beta|\alpha}$ from the two points at which it is identified to its entire support. We note

that, as in many discrete choice models, if we had data with sufficiently rich variation in covariates or variation in annuity rates that was exogenous to demand, we would be non-parametrically identified even with a discrete choice set.

Since our data limitations mean that we require a parametric assumption for $F_{\beta|\alpha}$ we try to address concerns about such (ad hoc) parametric assumptions by investigating the sensitivity of the results to several alternatives in Section 2.6. An alternative to a parametric interpolation is to make no attempt at interpolation, and to simply use the identified points as bounds on the cumulative distribution function. In Section 2.6 we also report such an exercise.

A second property of our data that makes it not fully consistent with the identification argument above is the censoring of mortality outcomes. Specifically, we do not observe mortality dates for those who are alive by the end of 2005, implying that we have no information in the data about mortality hazard rates for individuals older than 83. While we could identify and estimate a non-parametric baseline hazard for the periods for which mortality data are available (as well as a non-parametric distribution of α_i), there is obviously no information in the data about the baseline hazard rate for older ages. Because evaluating the guarantee choice requires knowledge of the entire mortality process (through age T , which we assume to be 100), some assumption about this baseline hazard is necessary. We therefore make (and test for) a parametric assumption about the functional form of the baseline hazard.

2.3.4 Parameterization

Mortality process. As we have just mentioned, due to the censored mortality data, we make a parametric assumption about the mortality hazard rate. Specifically, we assume that the baseline hazard rate follows a Gompertz distribution with shape parameter λ . That is, the baseline hazard rate is given by $\psi(t) = e^{\lambda t}$ and individual i 's mortality hazard at time $t = age_i - 60$ is therefore given by $\psi^i(t) = \alpha_i e^{\lambda t}$. We can test the Gompertz assumption in our sample against more flexible alternatives by focusing on individuals' mortality experience prior to the age of 83. We are reassured that the Gompertz assumption cannot be rejected by our (censored) mortality data.⁹ We also note that the Gompertz distribution is widely used in the actuarial literature that models mortality (Horiuchi and Coale (1982)).

⁹Specifically, we use likelihood-ratio tests of the baseline Gompertz model against more general alternatives where λ is allowed to vary with time. We divide the period of observation over which we observe mortality outcomes (21 years) into two and three evenly spaced intervals and let λ vary across intervals. The p -value of these tests are 0.938 and 0.373, respectively.

We model mortality as a continuous process and observe mortality at the daily level. However, since the parameterized version of the guarantee choice model is solved numerically, we work with a coarser, annual frequency, reducing the computational burden. In particular, given the above assumption, let

$$S(\alpha, \lambda, t) = \exp\left(\frac{\alpha}{\lambda} (1 - e^{\lambda t})\right) \quad (2.3.9)$$

be the Gompertz survival function, and the discrete (annual) hazard rate at year t is given by $\kappa_t^i = \frac{S(\alpha_i, \lambda, t) - S(\alpha_i, \lambda, t+1)}{S(\alpha_i, \lambda, t)}$.

Unobserved heterogeneity. An individual in our data can be characterized by an individual-specific mortality parameter α_i and an individual-specific preference parameter β_i . Everything else is assumed common across individuals. Although, as we showed, the joint distribution $F(\alpha, \beta)$ is non-parameterically identified with continuous guarantee choice, in practice only three guarantee lengths are offered, so we work with a parametrized distribution.

In the baseline specification we assume that α_i and β_i are drawn from a bivariate lognormal distribution

$$\begin{pmatrix} \log \alpha_i \\ \log \beta_i \end{pmatrix} \sim N \left(\begin{bmatrix} \mu_\alpha \\ \mu_\beta \end{bmatrix}, \begin{bmatrix} \sigma_\alpha^2 & \rho\sigma_\alpha\sigma_\beta \\ \rho\sigma_\alpha\sigma_\beta & \sigma_\beta^2 \end{bmatrix} \right). \quad (2.3.10)$$

In Section 2.6 we explore other distributional assumptions.

Calibrated values for other parameters. As mentioned, we treat a set of other parameters – γ , δ , η , and r – as observables, and calibrate their values. Here, we list the calibrated values and their source; in Section 2.6 we assess the sensitivity of the results to these values.

Since the insurance company does not have information on the annuitant’s wealth outside of the annuity, we calibrate the fraction of wealth annuitized (η) based on Banks and Emmerson (1999), who use market-wide evidence from the Family Resources Survey. They report that for individuals with compulsory annuity payments, about one-fifth of income (and therefore presumably of wealth) comes from the compulsory annuity. We therefore set $\eta = 0.2$. In Section 2.6 we discuss what the rest of the annuitants’ wealth portfolio may look like and how this may affect our counterfactual calculations.

We use $\gamma = 3$ as the coefficient of relative risk aversion. A long line of simulation literature uses this value (Hubbard, Skinner, and Zeldes (1995), Engen, Gale, and Uccello (1999), Mitchell, Poterba, Warshawsky, and Brown (1999), Scholz, Seshadri, and Khitatrakun (2003), Davis, Kubler, and Willen (2006)). Although a substantial consumption literature, summarized in

Laibson, Repetto, and Tobacman (1998), has found risk aversion levels closer to 1, as did Hurd’s (1989) study among the elderly, other papers report higher levels of relative risk aversion (Barsky, Kimball, Juster, and Shapiro (1997), Palumbo (1999)).

For r we use the real interest rate corresponding to the inflation-indexed zero-coupon ten-year Bank of England bond, as of the date of the pricing menu we use (January 1, 1992, in the baseline specification). This implies a real interest rate r of 0.0426. We also assume that the discount rate δ is equal to the real interest rate r .

Finally, since the annuities make constant nominal payments, we need an estimate of the expected inflation rate π to translate the initial nominal payment rate shown in Table 2.2 into the real annuity payout stream z_t in the guarantee choice model. We use the difference between the real and nominal interest rates on the zero-coupon ten year Treasury bonds on the same date to measure the (expected) inflation rate. This implies an (expected) inflation rate π of 0.0498.¹⁰

Summary and intuition. Thus, to summarize, in the baseline specification we estimate six remaining structural parameters: the five parameters of the joint distribution of α_i and β_i , and the shape parameter λ of the Gompertz distribution. We also allow for observable shifters to the means of the distribution. Specifically, we allow μ_α and μ_β to vary based on the individual’s gender and age at the time of annuitization. We do this because annuity rates vary with these characteristics, presumably reflecting differential mortality by gender and age of annuitization; so that our treatment of preferences and mortality is symmetric, we also allow mean preferences to vary on these same dimensions.

To gain intuition, note that one way to summarize the mortality data is by a graph of the log hazard mortality rate with respect to age. The Gompertz assumption implies that, without heterogeneity, this graph is linear with a slope of λ . Heterogeneity implies a concave graph, as over time lower mortality individuals are more likely to survive. Thus, loosely, the level of this graph affects the estimate of μ_α , the average slope affects the estimate of λ , and the concavity affects the estimate of σ_α . Since σ_α is a key parameter (which determines the extent of adverse selection), in Section 2.6 we explore the sensitivity of the results to more and less concave baseline hazard models.

¹⁰We ignore inflation uncertainty, which may lead us to over-state the welfare value of the nominal annuities we analyze. We make this abstraction for computational simplicity, and because prior work has found that incorporating uncertain inflation based on historical inflation patterns in the U.S. has a small quantitative effect (of about 1-2 percent) on the welfare gain from annuitization (Mitchell, Poterba, Warshawsky, and Brown (1999)). Since the U.K. inflation experience has been broadly similar, it seems natural to expect a qualitatively similar (small) effect in our context too.

Consider now the data on guarantee choices, and its relationship to mortality outcomes. Suppose first that there was no heterogeneity in mortality rates ($\sigma_\alpha = 0$). In such a case, the guarantee choice model would reduce to a standard ordered probit with three choices (see equation (2.3.14) below), and the thresholds would be known from the guarantee choice model and estimates of μ_α and λ . In this simple case the mean and variance of β would be directly estimated off the observed shares of the three different guarantee choices.

It is the presence of unobserved heterogeneity in mortality risk ($\sigma_\alpha > 0$) that makes intuition more subtle. The guarantee choice is still similar to an ordered probit, but the thresholds (which depend on α_i) are now unobserved. Therefore, the model is similar to an ordered probit with random effects. This is where the relationship between mortality and guarantee choices is crucial. By observing m_i , we obtain information about the unobserved α_i . Although this information is noisy (due to the inherent randomness of any hazard model), it is still useful in adjusting the weights $P(m_i|\alpha, \lambda)$ in the integral in equations (2.3.13) and (2.3.14) below. Loosely, individuals who (ex post) die earlier are more likely (from the econometrician's perspective) to be of higher (ex ante) mortality rate α_i . Therefore, the mortality data is used as a stochastic shifter of the individual random effects. This allows separate identification of σ_β and the correlation parameter ρ .

2.3.5 Estimation

For computational convenience, we begin by estimating the shape parameter of the Gompertz hazard λ using only mortality data. We then use the guarantee choice and mortality data together to estimate the parameters of the joint distribution $F(\alpha, \beta)$. We estimate the model using maximum likelihood. Here we provide a general overview; more details are provided in Appendix A.3.

Estimation of the parameters of the baseline hazard rate (λ). We observe mortality in daily increments, and treat it as continuous for estimation. We normalize $t_i = age_i - 60$ (as 60 is the age of the youngest individual who makes a guarantee choice in our sample). For each individual i , the mortality data can be summarized by $m_i = (c_i, t_i, d_i)$ where c_i is the (normalized) age at which individual i entered the sample (due to left truncation) and t_i is the age at which he exited the sample (due to death or censoring). d_i is an indicator for whether the person died ($d_i = 1$) or was censored ($d_i = 0$).

Conditional on α , the likelihood of observing m_i is

$$P(m_i = (c_i, t_i, d_i) | \alpha, \lambda) = \frac{1}{S(\alpha, \lambda, c_i)} (s(\alpha, \lambda, t_i))^{d_i} (S(\alpha, \lambda, t_i))^{1-d_i}, \quad (2.3.11)$$

where $S(\cdot)$ is the Gompertz survival function (see equation (2.3.9)) and $s(\cdot) = \frac{\partial S(\alpha, \lambda, t)}{\partial t}$ is the Gompertz density. Our incorporation of c_i into the likelihood function accounts for the left truncation in our data.

We estimate λ using only mortality data. We do so by using equation (2.3.11) and integrating over α_i . That is, we maximize the following likelihood

$$L^M(\lambda, \mu_\alpha, \sigma_\alpha | (m_i)_{i=1}^N) = \sum_{i=1}^N \log \left(\int P(m_i | \alpha, \lambda) \frac{1}{\sigma_\alpha} \phi \left(\frac{\log \alpha - \mu_\alpha}{\sigma_\alpha} \right) d\alpha \right) \quad (2.3.12)$$

to obtain a consistent estimate of λ .¹¹

Estimation of the parameters of $F(\alpha, \beta)$. Having estimated λ , we can then use the guarantee choice model to numerically compute the optimal guarantee choice for each combination of α_i and β_i . This choice is also a function of the other (calibrated) parameters of the model and of the observed annuity rates. Consistent with intuition, the numerical solution to the model has the property that the relative value that individual i obtains from a (longer) guarantee is increasing in both α_i and β_i . Recall that this monotonicity property is important for identification; specifically, it is key to proving the second part of Proposition 2. This implies that for any value of α_i , the guarantee choice can be characterized by two cutoff points: $\beta_{0/5}^*(\alpha_i)$ and $\beta_{5/10}^*(\alpha_i)$. The former is the value of β_i that makes an individual (with parameter α_i) indifferent between choosing no guarantee and a 5 year guarantee, while the latter is the value of β_i that makes an individual (with parameter α_i) indifferent between choosing a 5 year and a 10 year guarantee. For almost all relevant values of α_i the baseline model – as well as other variants we estimated – and its specification results in $\beta_{0/5}^*(\alpha_i) < \beta_{5/10}^*(\alpha_i)$, implying that there exists a range of β_i 's that implies a choice of a 5 year guarantee (the modal choice in the data). For some extreme values of α_i this does not hold, but because α_i is unobserved this does not create any potential problem. Figure 2.1 illustrates the optimal guarantee choice in the space of α_i and β_i , in the context of the baseline specification and the mortality data (which were used to estimate λ).

¹¹Note that all three parameters – λ , μ_α , σ_α – are in fact identified and estimated. However, we later re-estimate μ_α and σ_α using the entire data (that contain the guarantee choices), which is more efficient. As will be clear below, estimating λ using the entire data is computationally more demanding.

Keeping λ fixed at its estimate, we then estimate the parameters of $F(\alpha, \beta)$ by maximizing the likelihood of guarantee choices and mortality. The likelihood depends on the observed mortality data m_i and on individual i 's guarantee choice $g_i \in \{0, 5, 10\}$. We can write the contribution of individual i to the likelihood as

$$l_i(m_i, g_i; \mu, \Sigma, \lambda) = \int \mathbf{P}(m_i | \alpha, \lambda) \left(\int \mathbf{1} \left(g_i = \arg \max_g V_0^{A(g)}(\beta, \alpha, \lambda) \right) dF(\beta | \alpha; \mu, \Sigma) \right) dF(\alpha; \mu, \Sigma) \quad (2.3.13)$$

where $F(\alpha; \mu, \Sigma)$ is the marginal distribution of α_i , $F(\beta | \alpha; \mu, \Sigma)$ is the conditional distribution of β_i , λ is the Gompertz shape parameter, $\mathbf{P}(m_i | \alpha, \lambda)$ is given in equation (2.3.11), $\mathbf{1}(\cdot)$ is the indicator function, and the value of the indicator function is given by the guarantee choice model discussed in Section 2.3.1.

Given the monotonicity of the optimal guarantee choice in β_i (and ignoring – for presentation only – the rare cases of $\beta_{0/5}^*(\alpha_i) > \beta_{5/10}^*(\alpha_i)$), we can rewrite equation (2.3.13) as

$$l_i(m_i, g_i; \mu, \Sigma, \lambda) = \begin{cases} \int \mathbf{P}(m_i | \alpha, \lambda) \left(F(\beta_{0/5}^*(\alpha) | \alpha; \mu, \Sigma) \right) dF(\alpha; \mu, \Sigma) & \text{if } g_i = 0 \\ \int \mathbf{P}(m_i | \alpha, \lambda) \left(\begin{array}{l} F(\beta_{5/10}^*(\alpha) | \alpha; \mu, \Sigma) - \\ - F(\beta_{0/5}^*(\alpha) | \alpha; \mu, \Sigma) \end{array} \right) dF(\alpha; \mu, \Sigma) & \text{if } g_i = 5 \\ \int \mathbf{P}(m_i | \alpha, \lambda) \left(1 - F(\beta_{5/10}^*(\alpha) | \alpha; \mu, \Sigma) \right) dF(\alpha; \mu, \Sigma) & \text{if } g_i = 10 \end{cases} \quad (2.3.14)$$

That is, the inner integral in equation (2.3.13) becomes an ordered probit, where the cutoff points are given by the location in which a vertical line in Figure 2.1 crosses the two curves.

The primary computational challenge in maximizing the likelihood is that, in principle, each evaluation of the likelihood requires us to resolve the guarantee choice model and compute these cutoff points for a continuum of values of α . Since the guarantee choice model is solved numerically, this is not trivial. Therefore, instead of recalculating these cutoffs at every evaluation of the likelihood, we calculate the cutoffs on a large grid of values of α only once and then interpolate to evaluate the likelihood. Unfortunately, since the cutoffs also depend on λ , this method does not allow us to estimate λ jointly with all the other parameters. We could calculate the cutoffs on a grid of values of both α and λ , but this would increase computation time substantially. This is why, at some loss of efficiency but not of consistency, we first estimate λ using only the mortality portion of the likelihood, fix λ at this estimate, calculate the cutoffs, and estimate the remaining parameters from the full likelihood above. To compute standard errors, we use a nonparametric bootstrap.

2.4 Estimates and fit of the baseline model

2.4.1 Parameter estimates

Table 2.3 reports the parameter estimates. We estimate significant heterogeneity across individuals, both in their mortality and in their preference for wealth after death. We estimate a positive correlation (ρ) between mortality and preference for wealth after death. That is, individuals who are more likely to live longer (lower α) are likely to care less about wealth after death. This positive correlation may help to reduce the magnitude of the inefficiency caused by private information about risk; individuals who select larger guarantees due to private information about their mortality (i.e. high α individuals) are also individuals who tend to place a relatively higher value on wealth after death, and for whom the cost of the guarantee is not as great as it would be if they had relatively low preferences for wealth after death.

For illustrative purposes, Figure 2.2 shows random draws from the estimated distribution of $\log \alpha$ and $\log \beta$ for each age-gender cell, juxtaposed over the estimated indifference sets for that cell. The results indicate that both mortality and preference heterogeneity are important determinants of guarantee choice. This is similar to recent findings in other insurance markets that preference heterogeneity can be as or more important than private information about risk in explaining insurance purchases (Finkelstein and McGarry (2006), Cohen and Einav (2007), Fang, Keane, and Silverman (2008)). As discussed, we refrain from placing a structural interpretation on the β parameter, merely noting that a higher β reflects a larger preference for wealth after death relative to consumption while alive. Nonetheless, our finding of heterogeneity in β is consistent with other estimates of heterogeneity in the population in preferences for leaving a bequest (Laitner and Juster (1996), Kopczuk and Lupton (2007)).

2.4.2 Model fit

Table 2.4 and Table 2.5 present some results on the in-sample and out-of-sample fit of the model, respectively. We report results both overall and separately for each age-gender cell. Table 2.4 shows that the model fits very closely the probability of choosing each guarantee choice, as well as the observed probability of dying within our sample period. The model does, however, produce a monotone relationship between guarantee choice and mortality rate, while the data show a non-monotone pattern, with individuals who choose a 5 year guarantee period associated with highest mortality. As previously discussed (see footnote 4), the non-monotone pattern in the data may

merely reflect sampling error; we are unable to reject the null that the 5 and 10 year guarantees have the same mortality rate.

Table 2.5 compares our mortality estimates to two different external benchmarks. These speak to the out-of-sample fit of our model in two regards: the benchmarks are not taken from the data, and the calculations use the entire mortality distribution based on the estimated Gompertz mortality hazard, while our mortality data are right censored. The top panel of Table 2.5 reports the implications of our estimates for life expectancy. As expected, men have lower life expectancies than women. Men who purchase annuities at age 65 have higher life expectancies than those who purchase at age 60, which is what we would expect if age of annuity purchase were unrelated to mortality. Women who purchase at 65, however, have lower life expectancy than women who purchase at 60, which may reflect selection in the timing of annuitization, or the substantially smaller sample size available for 65 year old women. As one way to gauge the magnitude of the mortality heterogeneity we estimate, Table 2.5 indicates that in each age-gender cell, there is about a 1.4 year difference in life expectancy, at the time of annuitization, between the 5th and 95th percentile.

The fourth row of Table 2.5 contains life expectancy estimates for a group of U.K. pensioners whose mortality experience may serve as a rough proxy for that of U.K. compulsory annuitants.¹² We would not expect our life expectancy estimates – which are based on the experience of actual compulsory annuitants in a particular firm – to match this rough proxy exactly, but it is reassuring that they are in a similar ballpark. Our estimated life expectancy is about 2 years higher. This difference is not driven by the parametric assumptions, but reflects higher survival probabilities for our annuitants than our proxy group of U.K. pensioners; this difference between the groups exists even within the range of ages for which we observe survival in our data and can compare the groups directly (not shown).

The bottom of Table 2.5 presents the average expected present discounted value (EPDV) of annuity payments implied by our mortality estimates and our assumptions regarding the real interest rate and the inflation rate. Since each individual's initial wealth is normalized to 100, of which 20 percent is annuitized, an EPDV of 20 would imply that the company, if it had no transaction costs, would break even. Note that nothing in our estimation procedure guarantees that we arrive at reasonable EPDV payments. It is therefore encouraging that for all the four cells, and for all

¹²Exactly how representative the mortality experience of the pensioners is for that of compulsory annuitants is not clear. See Finkelstein and Poterba (2002) for further discussion of this issue.

guarantee choices within these cells, the expected payout is fairly close to 20; it ranges across the age-gender cells from 19.74 to 20.66. One might be concerned by an average expected payment that is slightly above 20, which would imply that the company makes negative profits. Note, however, that if the effective interest rate the company uses to discount its future payments is slightly higher than the risk-free rate of 0.043 that we use in the individual's guarantee choice model, the estimated EPDV annuity payments would all fall below 20. It is, in practice, likely that the insurance company receives a higher return on its capital than the risk free rate, and the bottom row of Table 2.5 shows that a slightly higher interest rate of 0.045 would, indeed, break even. In Section 2.6 we show that our welfare estimates are not sensitive to using an interest rate that is somewhat higher than the risk free rate used in the baseline model.

As another measure of the out-of-sample fit, we examined the optimal consumption trajectories implied by our parameter estimates and the guarantee choice model. These suggest that most of the individuals are saving in their retirement (not shown). This seems contrary to most of the empirical evidence (e.g., Hurd (1989)), although there is evidence consistent with positive wealth accumulation among the very wealthy elderly (Kopczuk (2007)), and evidence, more generally, that saving behavior of high wealth individuals may not be representative of the population at large (Dynan, Skinner, and Zeldes (2004)); individuals in this market are higher wealth than the general U.K. population (Banks and Emmerson (1999)). In light of these potentially puzzling wealth accumulation results, we experimented with a variant of the baseline model that allows individuals to discount wealth after death more steeply than consumption while alive. Specifically, we modified the consumer per-period utility function (as shown in equation (2.3.1)) to be

$$v^i(w_t, c_t) = (1 - \kappa_t^i) u^i(c_t) + \zeta^t \kappa_t^i b^i(w_t), \quad (2.4.1)$$

where ζ is an additional parameter to be estimated. Our benchmark model corresponds to $\zeta = 1$. Values of $\zeta < 1$ imply that individuals discount wealth after death more steeply than consumption while alive. Such preferences might arise if individuals care more about leaving money to children (or grandchildren) when the children are younger than when they are older. We find that the maximum likelihood value of ζ is 1. Moreover, when we re-estimate the model imposing values of ζ relatively close to 1 (such as $\zeta = 0.95$), we are able to produce more sensible wealth patterns in retirement, but do not have a noticeable effect on our core welfare estimates.

2.5 Welfare estimates

We now take our parameter estimates as inputs in calculating the welfare consequences of asymmetric information and government mandates. We start by defining the welfare measure we use, and calculating welfare in the observed, asymmetric information equilibrium. We then perform two counterfactual exercises in which we compare equilibrium welfare to what would arise under a mandatory social insurance program that does not permit choice over guarantee, and under symmetric information. Although we focus primarily on the average welfare, we also briefly discuss distributional implications.

2.5.1 Measuring welfare

A useful monetary metric for comparing utilities associated with different annuity allocations is the notion of *wealth-equivalent*. The wealth-equivalent denotes the amount of initial wealth that an individual would require in the absence of an annuity, in order to be as well off as with his initial wealth and his annuity allocation. The wealth-equivalent of an annuity with guarantee period g and initial wealth of w_0 is the implicit solution to

$$V_0^{A(g)}(w_0) \equiv V_0^{NA}(\text{wealth} - \text{equivalent}), \quad (2.5.1)$$

where both $V_0^{A(g)}(\cdot)$ and $V_0^{NA}(\cdot)$ are defined in Section 2.3. This measure is commonly used in the annuity literature (Mitchell, Poterba, Warshawsky, and Brown (1999), Davidoff, Brown, and Diamond (2005)).

A higher value of wealth-equivalent corresponds to a higher value of the annuity contract. If the wealth equivalent is less than initial wealth, the individual would prefer not to purchase an annuity. More generally, the difference between the wealth-equivalent and the initial wealth is the amount an individual is willing to pay in exchange for having access to the annuity contract. This difference is always positive for a risk averse individual who does not care about wealth after death and faces an actuarially fair annuity rate. It can take negative values if the annuity contract is over-priced (compared to the individual-specific actuarially fair rate) or if the individual sufficiently values wealth after death.

Our estimate of the average wealth-equivalent in the observed equilibrium provides a monetary measure of the welfare gains (or losses) from annuitization given equilibrium annuity rates and individuals' contract choices. The difference between the average wealth equivalent in the observed

equilibrium and in a counterfactual allocation provides a measure of the welfare difference between these allocations.

We provide two ways to quantify these welfare differences. The first provides an absolute monetary estimate of the welfare gain or loss associated with a particular counterfactual scenario. To do this, we scale the difference in wealth equivalents by the £6 billion which are annuitized annually (in 1998) in the U.K. annuity market (Association of British Insurers (1999)). Since the wealth equivalents are reported per 100 units of initial wealth and we assume that 20 percent of this wealth is annuitized, this implies that each unit of wealth-equivalent is equivalent, at the aggregate, to £300 million annually. We also occasionally refer to a per-annuitant welfare gain, which we compute by dividing the overall welfare effect by 300,000, which is our estimate of new annuitants in the U.K. market in 1998.¹³ Of course, one has to be cautious about these specific numbers, as they rely on extrapolating our estimates from our specific sample to the entire market.

While an absolute welfare measure may be a relevant benchmark for policies associated with the particular market we study, a relative measure may be more informative when considering using our estimates as a possible benchmark in other contexts, or examining the quantitative sensitivity of our estimates. For example, if we considered the decision to buy a one month guarantee, we would not expect efficiency costs associated with this decision to be large relative to life-time wealth. A relative welfare estimate essentially requires a normalization factor.

Therefore, to put these welfare estimates in perspective, we measure the welfare changes relative to how large this welfare change could have been, given the observed annuity rates. We refer to this maximum potential welfare change as the “Maximum Money at Stake” (MMS). We define the MMS as the minimum lump sum that individuals would have to receive to insure them against the possibility that they receive their least-preferred allocation in the observed equilibrium, given the observed equilibrium pricing. The MMS is therefore the additional amount of pre-existing wealth an individual requires so that they receive the same annual annuity payment if they purchase the maximum guarantee length (10 years) as they would receive if they purchase the minimum guarantee length (0 years).

The nature of the thought experiment behind the MMS is that the welfare loss from buying a 10 year guarantee is bounded by the lower annuity payment that the individual receives as a result. This maximum welfare loss would occur in the worst case scenario, in which the individual had

¹³We obtain it by dividing the £6 billion figure we have just referred to by the average annuitized amount (in 1998) in our full company data (rather than the sample we use for estimation; see Appendix A.1), which is £20,000.

no chance of dying during the first 10 years (or alternatively, no value of wealth after death). We report the MMS per 100 units of initial wealth (i.e., per 20 units of the annuitized amount)

$$MMS \equiv 20 \left(\frac{z_0}{z_{10}} - 1 \right), \quad (2.5.2)$$

where z_0 and z_{10} denote the annual annuity rates for 0 and 10 year guarantees, respectively (see Table 2.2). A key property of the MMS is that it depends only on annuity rates, but not on our estimates of preferences or mortality risk. Converting this to absolute amounts, the MMS is just over £500 million annually, just below £1,700 per new annuitant, or about 8 percent of the market as a whole.

2.5.2 Welfare in observed equilibrium

The first row of Table 2.6 shows the estimated average wealth equivalents per 100 units of initial wealth in the observed allocations implied by our parameter estimates. The average wealth equivalent for our sample is 100.16, and ranges from 99.9 (for 65 year old males) to 100.4 (for 65 year old females). An average wealth equivalent of less than 100 indicates an average welfare loss associated with the equilibrium annuity allocations relative to a case in which wealth is not annuitized; conversely, an average wealth equivalent of more than 100 indicates an average welfare gain from annuitization at the observed rates. Note that because annuitization of some form is compulsory, it is possible that individuals in this market would prefer not to annuitize.¹⁴

Figure 2.3 shows the distribution across different types of the welfare gains and losses from annuitization at the observed annuity rates, relative to no annuities. This figure super-imposes iso-welfare contour lines over the same scatter plots presented in Figure 2.2. It indicates that, as expected, the individuals who benefit the most from the annuity market are those with low mortality (low α) and weak preference for wealth after death (low β). The former are high (survival) risk, who face better than actuarially fair annuity rates when they are pooled with the rest of the annuitants.

¹⁴Our average wealth equivalent is noticeably lower than what has been calculated in the previous literature (Mitchell, Poterba, Warshawsky, and Brown (1999), Davidoff, Brown, and Diamond (2005)). The high wealth equivalents in these papers in turn implies a very high rate of voluntary annuitization, giving rise to what is known as the “annuity puzzle” since, empirically, very few individuals voluntarily purchase annuities (Brown, Mitchell, Poterba, and Warshawsky (2001)). Our substantially lower wealth equivalents – which persist in the robustness analysis (see Table 2.6) – arise because of the relatively high β that we estimate. Previous papers have calibrated rather than estimated β and assumed it to be 0. If we set $\log \alpha = \mu_\alpha$ and $\beta = 0$, and also assume – like these other papers – that annuitization is full (i.e., 100 percent vs. 20 percent in our baseline), then we find that the wealth equivalent of a 0 year guarantee for a 65 year old male rises to 135.9, which is much closer to the wealth equivalent of 156 reported by Davidoff, Brown, and Diamond (2005).

The latter are individuals who get less disutility from dying without much wealth, which is more likely to occur with than without annuities.

2.5.3 The welfare cost of asymmetric information

In the counterfactual symmetric information equilibrium, each person faces an actuarially fair adjustment to annuity rates depending on her mortality. Specifically, we offer each person payment rates such that the EPDV of payments for that person for each guarantee length is equal to the equilibrium average EPDV of payments. This ensures that each person faces an individual-specific actuarially fair reductions in payments in exchange for longer guarantees. Note that this calculation is (expected) revenue neutral, preserving any average load (or subsidy) in the market.

Figure 2.2 may provide a visual way to think about this counterfactual. In the counterfactual exercise, the points in Figure 2.2, which represent individuals, are held constant, while the indifference sets, which represent the optimal choices at a given set of annuity rates, shift. Wealth equivalents are different at the new optimal choices both because of the direct effect of the different annuity rates and because these rates in turn affect optimal contract choices.

We note that our welfare analysis of the impact of adverse selection considers only the impact of selection on the pricing of the observed contracts. Adverse selection may also affect the set of contracts offered, and this may have non trivial welfare costs. Our analysis however treats the contract set (of 0, 5, and 10 year guarantees) as given; that is, we assume that the contract space does not change in the counterfactual of symmetric information. The most important reason for this assumption is that incorporating the impact of adverse selection on the contract space would require a model of guarantee lengths in which the current offered guarantee lengths are optimal. This seems difficult given that the three offered guarantee lengths are fixed over time, across the annuity providers in the market, and perhaps most surprisingly over different age and gender combinations, which are associated with different mortality profiles.

The second panel of Table 2.6 presents our estimates of the welfare cost of asymmetric information. The first row shows our estimated wealth-equivalents in the symmetric information counterfactual. As expected, welfare is systematically higher in the counterfactual world of symmetric information. For 65 year old males, for example, the estimates indicate that the average wealth equivalent is 100.74 under symmetric information, compared to 100.17 under asymmetric information. This implies that the average welfare loss associated with asymmetric information is equivalent to 0.57 units of initial wealth. For the other three age-gender cells, this number ranges

from 0.14 to 0.27. Weighting all cells by their relative sizes, we obtain the overall estimate reported in the introduction of annual welfare costs of £127 million, £423 per new annuitant, or about 2 percent of annuitized wealth. This also amounts to 0.25 of the concept of maximum money at stake (MMS) introduced earlier.

What is the cause of this welfare loss? It arises from the distortion in the individual's choice of guarantee length relative to what he would have chosen under symmetric information pricing. Despite preference heterogeneity, we estimate that under symmetric information all individuals would choose 10 year guarantees (not shown). However, in the observed equilibrium only about 3 percent of individuals purchase these annuities. This illustrates the distortions in optimal choices in the observed equilibrium.

To illustrate the impact on different individuals, Figure 2.4 presents contour graphs of the changes in wealth equivalents associated with the change to symmetric information. That is, as before, for each age-gender cell we plot the individuals as points in the space of $\log \alpha$ and $\log \beta$, and then draw contour lines over them. All the individuals along a contour line are predicted to have the same absolute welfare change as a result of the counterfactual. Figure 2.4 indicates that, while almost all individuals benefit from a move to the first best, there is significant heterogeneity in the welfare gains arising from individual-specific pricing. The biggest welfare gains accrue to individuals with high mortality (high α) and high preferences for wealth after death (high β).

Two different factors work in the same direction to produce the highest welfare gains for high α , high β individuals. First, a standard one-dimensional heterogeneity setting would predict that symmetric information would improve welfare for low risk (high α) individuals relative to high risk (low α) individuals. Second, the asymmetric information equilibrium involves cross-subsidies from higher guarantees to lower guarantees (the EPDV of payout decreases with the length of the guarantee period, as shown in Table 2.5);¹⁵ by eliminating these cross-subsidies, symmetric information also improves the welfare of high β individuals, who place more value on higher guarantees. Since we estimate that α and β are positively correlated, these two forces reinforce each other.

A related question concerns the extent to which our estimate of the welfare cost of asymmetric

¹⁵The observed cross-subsidies across guarantee choices may be due to asymmetric information. For example, competitive models of pure adverse selection (with no preference heterogeneity), such as Miyazaki (1977) and Spence (1978), can produce equilibria with cross-subsidies from the policies with less insurance (in our context, longer guarantees) to those with more insurance (in our context, shorter guarantees). We should note that the observed cross-subsidies may also arise from varying degrees of market power in different guarantee options. In such cases, symmetric information may not eliminate cross-subsidies, and our symmetric information counterfactual would therefore conflate the joint effects of elimination of informational asymmetries and of market power. Our analysis of the welfare consequences of government mandates in the next subsection does not suffer from this same limitation.

information is influenced by re-distributional effects. As just discussed, symmetric information produces different welfare gains for individuals with different α and β . To investigate the extent to which our welfare comparisons are affected by the changes in cross-subsidy patterns, we recalculated wealth-equivalents in the symmetric information counterfactual under the assumption that each individual faces the same expected payments for each option in the choice set of the counterfactual as she receives at her choice in the observed equilibrium. The results (not shown) suggest that, in all the age-gender cells, our welfare estimates are not, in practice, affected by redistribution.

2.5.4 The welfare consequences of government mandated annuity contracts

Although symmetric information is a useful conceptual benchmark, it may not be relevant from a policy perspective since it ignores the information constraints faced by the social planner. We therefore consider the welfare consequences of government intervention in this market. Specifically, we consider the consequences of government mandates that each individual purchases the same guarantee length, eliminating any contract choice; as noted previously, such mandates are the canonical solution to adverse selection in insurance markets (Akerlof (1970)). To evaluate welfare under alternative mandates, we calculate average wealth equivalents when all people are forced to have the same guarantee period and annuity rate, and compare them to the average wealth equivalents in the observed equilibrium. We set the payment rate such that average EPDV of payments is the same as in the observed equilibrium; this preserves the average load (or subsidy) in the market.

Before presenting the results, it is useful to note a contrast between our setting and the standard or canonical insurance model. As mentioned in the introduction, unlike in a standard insurance setting, the optimal mandatory annuity contract cannot be determined by theory alone. In the canonical insurance model – that is, when all individuals are risk averse, the utility function is state-invariant, and there are no additional cost of providing insurance – it is well-known that mandatory (uniform) full insurance can achieve the first best allocation, even when individuals vary in their preferences. Since adverse selection reduces insurance coverage away from this first-best, no estimation is required in this standard context to realize that the optimal mandate is full insurance. In contrast, our model of annuity choices is governed by two different utility functions, one from consumption when alive, $u(\cdot)$, and one from wealth when dead, $b(\cdot)$ (see equation (2.3.1)). Therefore optimal (actuarially fair) guarantee coverage will vary across individuals depending on their relative preference for wealth at death vis-a-vis consumption while alive. In such a case,

whether and which mandatory guarantee can improve welfare gains relative to the adverse selection equilibrium is not a-priori clear.¹⁶ The investigation of the optimal mandate – and whether it can produce welfare gains relative to the adverse selection equilibrium – therefore becomes an empirical question.

The results are presented in the bottom panels of Table 2.6. In all four age-gender cells, welfare is lowest under a mandate with no guarantee period, and highest under a mandate of a 10 year guarantee. Welfare under a mandate of a 5 year guarantee is similar to welfare in the observed equilibrium.

The increase in welfare from a mandate of 10 year guarantee is virtually identical to the increase in welfare associated with the first best, symmetric information outcome reported earlier. This mandate involves no allocative inefficiency, since we estimated that a 10 year guarantee is the first best allocation for all individuals. Although it does involve transfers (through the common pooled price) across individuals of different mortality risk, these do not appear to have much effect on our welfare estimate.¹⁷ Consistent with this, when we recalculated wealth-equivalents in each counterfactual under the assumption that each individual faces the same expected payments in the counterfactual as she receives from her choice in the observed equilibrium, our welfare estimates were not noticeably affected (not shown). As with the counterfactual of symmetric information, there is heterogeneity in the welfare effects of the different mandates for individuals with different α and β . Not surprisingly, high β individuals benefit relatively more from the 10 year mandate and lose relatively more from the 0 year mandate (not shown).

Our findings highlight both the potential benefits and the potential dangers from government mandates. Without estimating the joint distribution of risk and preferences, it would not have been apparent that a 10 year guarantee is the welfare-maximizing mandate, let alone that such a mandate comes close to achieving the first best outcome. Were the government to mandate no guarantee, it would reduce welfare by about £107 million per year (£357 per new annuitant), achieving a welfare loss of about equal and opposite magnitude to the £127 million per year (£423 per new annuitant) welfare gain from the optimal 10 year guarantee mandate. Were the government to

¹⁶This is somewhat analogous to an insurance market with a state-dependent utility function. In such a case, the optimal mandate could be either full, partial, or no insurance (and analogously longer or shorter guarantee). For more details, see Sections 2 and 3.1 of the working paper version (Einav, Finkelstein, and Schrimpf (2007)).

¹⁷We estimate that welfare is slightly higher under the 10 year mandate than under the symmetric information equilibrium (in which everyone chooses the 10 year guarantee). This presumably reflects the fact that under the mandated (pooling) annuity payout rates, consumption is higher for low mortality individuals and lower for high mortality individuals than it would be under the symmetric information annuity payout rates. Since low mortality individuals have lower consumption in each period and hence higher marginal utility of consumption, this transfer improves social welfare (given the particular social welfare measure we use).

pursue the naive approach of mandating the currently most popular choice (5 year guarantees) our estimates suggest that this would raise welfare by only about £2 million per year or less than £7 per new annuitant, foregoing most of the welfare gains achievable from the welfare maximizing 10 year mandate. These results highlight the practical difficulties involved in trying to design mandates to achieve social welfare gains.

2.6 Robustness

In this section, we explore the robustness of our welfare findings. Our qualitative welfare conclusions are quite stable across a range of alternative assumptions. In particular, the finding that the welfare maximizing mandate is a 10 year guarantee, and that this mandate achieves virtually the same welfare as the first best outcome, persists across all alternative specifications. The finding of welfare gains from a 10 year guarantee mandate but welfare losses from mandating no guarantee is also robust.

However, the quantitative estimates of the welfare cost of asymmetric information can vary non-trivially across specifications, and as a result needs to be interpreted with more caution. It is £127 million per year (i.e. 25 percent of the MMS) in our baseline specification. It ranges from £111 million per year to £244 million per year (or from 22 percent to about 50 percent of the MMS) across the alternative specifications. Our bounds exercise, which we discuss below, produces similar conclusions concerning the robustness of our findings concerning the optimal guarantee mandate and its ability to achieve close to the first best outcome, as well as the greater uncertainty about our quantitative welfare estimates of the gains from symmetric information.

Finally, we note that our robustness discussion focuses on the (qualitative and quantitative) sensitivity of our welfare estimates, rather than the estimates of the underlying parameters (e.g., the magnitude of the average β). The underlying parameters change quite a bit under many of the alternative models. This is important for understanding why, as we vary certain assumptions, it is not a-priori obvious how our welfare estimates will change (in either sign or magnitude). For example, although it may seem surprising that welfare estimates are not very sensitive to our assumption about the risk aversion parameter, recall that the estimated parameters also change with the change in the assumption about risk aversion.

The change in the estimated parameters across specifications is also important for the overall interpretation of our findings. One reason we hesitate to place much weight on the structural

interpretation of the estimated parameters (or the extent of heterogeneity in these parameters) is that their estimates will be affected by our assumptions about other parameters (such as risk aversion or discount rate). This is closely related to the identification result in Section 2.3.

The remainder of this section describes the alternative specifications we explored. Table 2.6 provides a summary of the main results.

2.6.1 Parameter choices

Following our discussion of the baseline model in Section 2.3, although we estimate the average level and heterogeneity in mortality (α_i) and in preferences for wealth after death (β_i), we choose values for a number of other parameters based on external information. While we could, in principle, estimate some of these parameters, they would be identified solely by functional form assumptions. Therefore, we instead chose to explore how our welfare estimates are affected by alternative choices for these parameters.

Choice of risk aversion coefficient (γ). Our baseline specification (reproduced in row 1 of Table 2.6) assumes a (common) CRRA parameter of $\gamma = 3$ for both the utility from consumption $u(c)$ and from wealth after death $b(w)$. Rows 2 and 3 of Table 2.6 show the results if instead we assume $\gamma = 5$ or $\gamma = 1.5$.

Rows 4 and 5 report specifications in which we hold constant the CRRA parameter in the utility from consumption (at $\gamma = 3$) but vary the CRRA parameter in the utility from wealth after death. Specifically, we estimate the model with $\gamma = 1.5$ or $\gamma = 5$ for the utility from wealth after death $b(w)$.

A downside of the specifications reported in rows 4 and 5 is that they give rise to non-homothetic preferences and are therefore no longer scalable in wealth. This implies that heterogeneity in initial wealth may confound the analysis. Therefore, in row 6, we also allow for heterogeneity in initial wealth. As in row 5, we assume that $\gamma = 3$ for utility from consumption, but that $\gamma = 1.5$ for the utility from wealth after death. This implies that wealth after death acts as a luxury good, with wealthier individuals caring more, at the margin, about wealth after death. Such a model is consistent with the hypothesis that bequests are a luxury good, which may help explain the higher rate of wealth accumulation at the top of the wealth distribution (Dynan, Skinner, and Zeldes (2004), Kopczuk and Lupton (2007)). Unfortunately, we do not have data on individual's initial wealth w_0^i , which would allow us to incorporate it directly into the model. Instead, to

allow for heterogeneity in initial wealth, we calibrate the distribution of wealth based on Banks and Emmerson (1999) and integrate over this (unobserved) distribution.¹⁸ We also let the means (but not variances) of $\log \alpha$ and $\log \beta$ to vary with unobserved wealth. The welfare estimates are normalized to be comparable with the other exercises.

Choice of other parameters. We also reestimated the model assuming a higher interest rate than in the baseline case. As already mentioned, our estimates suggest that a slightly higher interest rate than the risk free rate we use in the individual's value function is required to have the annuity company not lose money. Thus, rather than the baseline which uses the risk free rate as of 1992 ($r = \delta = 0.043$), in row 7 we allow for the likely possibility that the insurance company receives a higher rate of return, and reestimate the model with $r = \delta = 0.05$. This in turn implies an average load on policies of 3.71 percent.

In row 8 we use a different set of annuity rates. Since the choice of 1992 pricing for our baseline model was arbitrary, we report results for a different set of annuity rates, from 1990, with the corresponding inflation and interest rates.

2.6.2 Wealth portfolio outside of the compulsory annuity market

As noted, our data do not contain information on the annuitant's wealth portfolio outside of the compulsory market. This is an important limitation to the data. In our baseline specification we used survey data reported by Banks and Emmerson (1999) to assume that 20 percent of the annuitants' financial wealth is in the compulsory annuity market ($\eta = 0.2$), and the rest is in liquid financial wealth. Rows 9 and 10 report results under different assumptions of the fractions of wealth annuitized in the compulsory market (we tried values of 0.1 and 0.3 of η).

In row 11 we report results in which we allow for heterogeneity in η . We calibrate the distribution of η and integrate over this unobserved distribution.¹⁹ We allow the means (but not variances) of $\log \alpha$ and $\log \beta$ to vary with this unobserved η .

In row 12, we assume that 50 percent of wealth is annuitized (at actuarially fair annuity rates) through the public Social Security program.²⁰ We then consider the welfare cost of asymmetric

¹⁸Banks and Emmerson (1999) report that the quartiles of the wealth distribution among 60-69 pensioners are 1,750, 8,950, and 24,900 pounds. We assume that the population of retirees is drawn from these three levels, with probability 37.5%, 25%, and 37.5%, respectively.

¹⁹Banks and Emmerson (1999) report an average η of 20 percent and a median of 10 percent. We therefore calibrate heterogeneity in η by assuming it can obtain one of three values – 0.1, 0.2, and 0.4 – with probabilities of 0.5, 0.25, and 0.25, respectively.

²⁰On average in the U.K. population, about 50 percent of retirees' wealth is annuitized through the public Social

information for the 20 percent of wealth annuitized in the compulsory market. As can be seen in Table 2.6, this alternative assumption has by far the biggest effect on our estimate of the welfare cost of asymmetric information, raising it from £127 million per year (or about 25 percent of the MMS) in the baseline specification to £244 million per year (or about 50 percent of the MMS).

As we noted at the outset of this section, it is difficult to develop good intuition for the comparative statics across alternative models since the alternative models also yield different estimated parameters. However, one potential explanation for our estimate of a larger welfare cost when 50 percent of wealth is in the public annuity may be that the individual now only has 30 percent of his wealth available to “offset” any undesirable consumption path generated by the 70 percent of annuitized wealth.

A related issue is the possibility that annuitants may adjust their non-annuitized financial wealth portfolio in response to the changes in guarantee prices created by our counterfactuals. Our analysis assumes that individuals do not adjust the rest of their portfolio in response to changes in their guarantee length or price. If individuals could purchase actuarially fair life insurance policies with no load, and without incurring any transaction costs in purchasing these policies, they could in principle undo much of the efficiency cost of annuitization in the current asymmetric information equilibrium. More generally, this issue fits into the broader literature that investigates the possibility and extent of informal insurance to lower the welfare benefits from government interventions or private insurance (Golosov and Tsyvinski (2007))

Of course, in practice the ability to offset the equilibrium using other parts of the financial portfolio will be limited by factors such as loads and transaction costs. Given that the maximum money at stake in the choice of guarantee is only about 8 percent of annuitized wealth under the observed annuity rates (and only about 4 percent (on average) under the counterfactual symmetric information rates), even relatively small transaction costs could well deter individuals from re-optimizing their portfolios in response to changes in guarantee prices. Re-optimization will also be limited by the fact that much of individuals’ wealth outside of the compulsory annuity market is tied up in relatively illiquid forms such as the public pension. Indeed, the data suggest that for individuals likely to be in the compulsory annuity market, only about 10 to 15 percent of their total wealth is in the form of liquid financial assets (Banks, Emmerson, Oldfield, and Tetlow (2005)).

Security program, although this fraction declines with retiree wealth (Office of National Statistics (2006)). Compulsory annuitants tend to be of higher than average socio-economic status (Banks and Emmerson (1999)) and may therefore have on average a lower proportion of their wealth annuitized through the public Social Security program. However, since our purpose is to examine the sensitivity of our welfare estimates to accounting for publicly provided annuities, we went with the higher estimate to be conservative.

A rigorous analysis of this is beyond the scope of the current work, and would probably require better information than we have on the asset allocation of individual annuitants. With richer data that included information on the life insurance holdings in each individual’s portfolio, we could potentially expand our model to include a model of life insurance demand and thereby use our estimates to examine how this aspect of the portfolio would respond to our counterfactual annuity rates, and how this in turn it would affect the welfare estimates of these counterfactuals. We hope that further research with hopefully richer data will build on the model and identification results here to extend the analysis in this important dimension.

2.6.3 Modeling heterogeneity

Different distributional assumptions of heterogeneity. We explored the sensitivity of our welfare estimates to the parameterization of unobserved heterogeneity. One potential issue concerns our parametric assumption regarding the baseline mortality distribution at the individual level. As discussed in the end of Section 2.3, our assumption about the shape of the individual mortality hazard affects our estimate of unobserved mortality heterogeneity (i.e., σ_α). To explore the importance of our assumption, row 13 presents results under a different assumption about the mortality distribution at the individual level. In particular, we assume a mortality distribution at the individual level with a hazard rate of $\alpha_i \exp(\lambda(t - t_0)^h)$ with $h = 1.5$, which increases faster over time than the baseline Gompertz specification (which has the same form, but $h = 1$). This, by construction, leads to a higher estimated level of heterogeneity in mortality, since the baseline hazard is more convex at the individual level.

We also investigated the sensitivity of the results to alternative joint distributional assumptions than our baseline assumption that α and β are joint lognormally distributed. Due to our estimation procedure, it is convenient to parameterize the joint distribution of α and β in terms of the marginal distribution of α and the conditional distribution of β . It is common in hazard models with heterogeneity to assume a gamma distribution (Han and Hausman (1990)). Accordingly, we estimate our model assuming that α follows a gamma distribution. We assume that conditional on α , β is distributed either lognormally (row 14) or gamma (row 15). Specifically, let a_α be the shape parameter and b_α be the scale parameter of the marginal distribution of α . When β is conditionally log-normally distributed, its distribution is parameterized by

$$\log(\beta)|\alpha \sim N(\mu_\beta + \rho(\log(\alpha) - \log(b_\alpha)), \sigma_\beta^2). \tag{2.6.1}$$

When β is conditionally gamma distributed, its shape parameter is simply a_β , and its conditional scale parameter is $b_\beta = \exp(\mu_\beta + \rho(\log(\alpha) - \log(b_\alpha)))$. These specifications allow thinner tails, compared to the bivariate lognormal baseline.

In unreported specifications, we have also experimented with discrete mixtures of lognormal distributions, in an attempt to investigate the sensitivity of our estimates to the one-parameter correlation structure of the baseline specification. These mixtures of lognormal distributions almost always collapsed back to the single lognormal distribution of the baseline estimates, trivially leading to almost identical welfare estimates.

Bounds. As mentioned earlier, an alternative to a parametric interpolation is to make no attempt at interpolation, and to simply use the identified points as bounds on the cumulative distribution function. To do so, we fix μ_α and σ_α (and λ) at our baseline estimates, and then use semiparametric Maximum Likelihood to obtain estimates for $P(g(\alpha, \beta) = y|\alpha)$, where $y = 0, 5, 10$. As shown in Proposition 2, this conditional guarantee choice is identified even when the choice set is discrete. Using the guarantee choice model and the fact that the guarantee choice is (weakly) monotone in β in our model, these conditional guarantee choices can be mapped to bounds on the conditional distribution $F_{\beta|\alpha}$ (see our discussion of $\beta_{0/5}^*(\alpha_i)$ and $\beta_{5/10}^*(\alpha_i)$ in the end of Section 2.3). We can then use these bounds to compute bounds on any object of interest.

To be more precise, let $h(\alpha, \beta)$ be an object of interest (e.g., welfare), and consider the case in which we wish to bound its population average. We then compute an upper bound by:

$$\overline{Eh} = \int \left(\begin{aligned} & \left(\sup_{\beta < \beta_{0/5}^*(\alpha)} h(\alpha, \beta) \right) P\left(\beta < \beta_{0/5}^*(\alpha)\right) + \\ & + \left(\sup_{\beta \in [\beta_{0/5}^*(\alpha), \beta_{5/10}^*(\alpha)]} h(\alpha, \beta) \right) P\left(\beta \in [\beta_{0/5}^*(\alpha), \beta_{5/10}^*(\alpha)]\right) + \\ & + \left(\sup_{\beta > \beta_{5/10}^*(\alpha)} h(\alpha, \beta) \right) P\left(\beta > \beta_{5/10}^*(\alpha)\right) \end{aligned} \right) dF(\alpha), \quad (2.6.2)$$

and similarly for the lower bound (with sup replaced by inf). We focus on bounding the welfare change from the different counterfactuals. To do this, we first compute the expected annuity payments in the observed equilibrium (these are point identified, as they are a function of the conditional guarantee choice, $P(g(\alpha, \beta) = y|\alpha)$), and use this to compute annuity rates in each of the counterfactuals. We then follow the procedure above to obtain bounds on the welfare change for each of the counterfactuals (a symmetric information case, and each of the three mandates we explored), for each of the age and gender combination separately.

The results from this exercise (not shown) imply that across all age and gender combinations,

the welfare ranking of the different mandates is the same as in our baseline case. In all age-gender cases, the welfare effect of the different mandates can be unambiguously ranked in the sense that their bounds do not overlap. In particular, a 10 year guarantee mandate results in a positive welfare gain which even at its lower bound is always higher than the upper bound of the welfare gain from any other mandate. The no guarantee mandate always produces a negative effect on welfare (even at the upper bound), and a 5 year guarantee mandate results in a small and mostly negative welfare effect (in two of the four age-gender combinations the upper bound of the welfare is positive, but very small). As in the baseline model, the welfare gain of the symmetric information equilibrium is similar to that of a 10 year guarantee mandate in the sense that the ranges of these welfare gains largely overlap (although in most cases the symmetric equilibrium outcome results in slightly tighter bounds). Consistent with the baseline results, in all cases we also obtain the result that the vast majority of individuals choose the 10-year guarantee contract in the symmetric information counterfactual. To check robustness, we also use the same procedure to bound the difference in welfare between one counterfactual to each of the others. Given that the bounds on the welfare change do not overlap, it may not be surprising that the bounds on the welfare differences also give rise to the same ranking of guarantee mandates. That is, zero is never within these bounds, so each mandate can be unambiguously ranked with respect to each of the alternatives.

In contrast to the robust ranking, the bounds on the estimated magnitude of the welfare gains (from either symmetric information or from the 10-year guarantee mandate) are not tight. For example, in the largest age-gender cell (65 year old males), we estimate the lower bound on the welfare gain from symmetric information to be as low as 30 percent of our baseline estimate, and in another cell (60 year old males) the upper bound on the welfare change from symmetric information is 56% higher than our baseline estimate. We view these results as largely consistent with the rest of the sensitivity analysis in this section; the results regarding the optimal mandate, as well as the similarity of the welfare gains from the optimal mandate and symmetric information are quite robust, but the quantitative estimates of the welfare gains are more sensitive to various assumptions.

Allowing heterogeneity in other parameters. While we allow for heterogeneity in mortality (α) and in preference for wealth after death (β), our baseline specification does not allow for heterogeneity in other determinants of annuity choice, such as risk aversion and discount rate. Since the various parameters are only identified up to a single dimension (see Section 2.3), except by functional form, more flexible estimation of α and β is analogous to a specification which frees

up these other parameters.

One way to effectively allow for more flexible heterogeneity is to allow the mean of β and α to depend on various observable covariates. In particular, one might expect both mortality and preferences for wealth after death to vary with an individual's socioeconomic status. We observe two proxies for the annuitant's socioeconomic status: the amount of wealth annuitized and the geographic location of the annuitant residence (his or her ward) if the annuitant is in England or Wales (about 10 percent of our sample is from Scotland). We link the annuitant's ward to ward-level data on socioeconomic characteristics of the population from the 1991 U.K. Census; there is substantial variation across wards in average socioeconomic status of the population (Finkelstein and Poterba (2006)). Row 16 shows the results of allowing the mean of both parameters to vary with the annuitized amount and the percent of the annuitant's ward that has received the equivalent of a high school degree or higher; both of these covariates may proxy for the socioeconomic status of the annuitant.

We also report results from an alternative model in which – in contrast to our baseline model – we assume that individuals are homogenous in their β but heterogeneous in their consumption γ . Rows 17 and 18 report such a specification. In row 17 we fix β at its estimated conditional median from the baseline specification (Table 2.3) and assume that α and the coefficient of risk aversion for utility from consumption are heterogeneous and (bivariate) lognormally distributed. The γ coefficient in the utility from wealth after death $b(w)$ is fixed at 3. As in row 6, this specification gives rise to non-homothetic preferences, so we use the median wealth level from Banks and Emmerson (1999) and later renormalize, so the reported results are comparable.

Row 18 allows for preference heterogeneity in both β and γ . For computational reason, we assume that γ is drawn from a discrete support (of 1.5, 3, and 4.5). We assume that α and β are (as in the baseline model) joint lognormally distributed, but we allow γ (which is unobserved) to shift their means. We note that this specification of heterogeneity in both β and γ is only identified by functional form, cautioning against structural interpretation of the estimated distribution of heterogeneity.

2.6.4 Imperfect information about mortality

Throughout we made a strong assumption that individuals have perfect information about their actual mortality rate α_i . This is consistent with empirical evidence that individuals' perceptions about their mortality probabilities covary in sensible ways with known risk factors, such as age,

gender, smoking, and health status (Hamermesh (1985), Smith, Taylor, and Sloan (2001), Hurd and McGarry (2002)). Of course, such work does not preclude the possibility that individuals also make some form of an error in forecasting their mortality.

We therefore investigate other assumptions about the information structure. Recall that while we make a perfect information assumption in order to establish identification, we can identify the model using alternative assumptions about the information structure. We report two such exercises here.

Before reporting the exercises, we note at the outset two potential complications with models of imperfect information, which are why we prefer to work with perfect information in our baseline specification. First, the dynamic nature of our model gives rise to potential learning. As individuals survive longer they may update their prior about their true underlying mortality process. While such learning can no longer affect their (past) guarantee choice, it could affect their consumption decisions. If forward looking individuals anticipate this possibility for learning, they may take this into account and it could alter their guarantee choice. We do not account for such learning in the exercises we report below. Second, once information is imperfect, the notion of welfare may be less obvious. One could measure “perceived” welfare which is measured with respect to the individual’s information, or “true” welfare which is measured with respect to the true mortality process. We choose to report perceived welfare, which is more consistent with our notion of wealth equivalence.

Throughout, we assume that individuals have perfect information about the mortality process, except for their idiosyncratic risk characterized by α_i . With some abuse of notation, we denote by $\kappa(\alpha_i)$ the perceived mortality risk by individual i . Our first set of exercises assumes that individuals have biased beliefs about their mortality risk. In particular, individuals know that

$$\log \kappa(\alpha_i) = \mu_\alpha(x_i) + \theta (\log \alpha_i - \mu_\alpha(x_i)), \quad (2.6.3)$$

where α_i is the true mortality rate of individual i , μ_α is the population mean of $\log \alpha_i$ (estimated in Table 2.3), and $\kappa(\alpha_i)$ is the mortality rate perceived by individuals when they make their guarantee choice and subsequent consumption decisions. θ is a free parameter. When $\theta = 1$ individuals have correct beliefs and the above assumption reduces to our baseline model. When $\theta < 1$ individuals perceive their mortality process as closer to the mean, while $\theta > 1$ is the case where individuals over-weight idiosyncratic information. Results for the cases of $\theta = 0.5$ and $\theta = 2$ are summarized in rows 19 and 20 of Table 2.6.

The second set of exercises assumes that individuals have correct, but uncertain beliefs about their mortality risk. In particular, let

$$\log \kappa(\alpha_i) \sim N(\log \alpha_i, \sigma_\varepsilon^2). \quad (2.6.4)$$

Our baseline model is the special case of $\sigma_\varepsilon = 0$. The case of $\sigma_\varepsilon > 0$ represents specifications where individuals are more uncertain about their mortality realization. We model the guarantee choices by having individuals form expected value functions by integrating over this additional uncertainty. In rows 21 and 22 we summarize results for the cases of $\sigma_\varepsilon = 0.027$ and $\sigma_\varepsilon = 0.108$, which are half and twice our estimate of σ_α (see Table 2.3).

2.6.5 Departing from the neoclassical model

Our baseline model is a standard neoclassical model with fully rational individuals. It is worth briefly discussing various “behavioral” phenomena that our baseline model (or extensions to it) can accommodate.

A wide variety of non-standard preferences may be folded into the interpretation for the preference for wealth after death parameter β . As previously noted, this preference may reflect a standard bequest motive, or some version of “regret” or “peace of mind” that have been discussed in the behavioral literature (Braun and Muermann (2004)).

Another possibility we considered is non-traditional explanations for the high fraction of individuals in our data who choose the 5 year guarantee option. One natural possibility that can be ruled out is that this reflects an influence of the 5 year guarantee as the default option. In practice there is no default for individuals in our sample, all of whom annuitized at age 60 or 65. Individuals in this market are required to annuitize by age 70 (for women) or 75 (for men). To annuitize before that age, they must actively fill a form when they decide to annuitize, and must check a chosen guarantee length. Failure to complete such an active decision would simply delay annuitization until the maximum allowed age.

Another natural possibility is that the popularity of the 5 year guarantee may partly reflect the well-known phenomenon in the marketing literature that individuals are more likely to “choose the middle” (Simonson and Tversky (1992)). We therefore estimated a specification of the model in which we allow for the possibility that some portion of individuals “blindly” choose the middle, that is the 5 year guarantee option. We allow such individuals to also differ in the mean mortality

rate. Row 23 summarizes the results from such a specification.²¹

2.6.6 Estimates for a different population

As a final robustness exercise, we re-estimated the baseline model on a distinct sample of annuitants. As mentioned briefly in Section 2.2 and discussed in more detail in Appendix A.1, in our baseline estimates we limit the annuitant sample to the two-thirds of individuals who have accumulated their pension fund with our company. Annuitants may choose to purchase their annuity from an insurance company other than the one in which their funds have been accumulating, and about one-third of the annuitants in the market choose to do so. As our sample is from a single company, it includes those annuitants who accumulated their funds with the company and stayed with the company, as well as those annuitants who brought in external funds. Annuitants who approach the company with external funds face a different pricing menu than those who buy internally. Specifically, the annuity payment rates are lower by 2.5 pence per pound of the annuitized amount than the payment rates faced by “internal” annuitants.²² Annuitants who approach the company with external funds may also be drawn from a different distribution of risk and preferences, which is why we do not include them in our main estimates. The estimated parameters for this population are, indeed, quite different from the estimates we obtain for the internal individuals (not shown).

Row 24 shows the results of estimating the model separately for this distinct group of individuals, using their distinct pricing menu. We continue to find that the welfare minimizing mandate is of no guarantee and that the welfare maximizing mandate is a 10 year guarantee, and it can get very close to the welfare level of the first best outcome. The welfare cost of asymmetric information is also quite similar: £137 in this “external” annuitant sample, compared to our baseline estimate of £127 in our sample of annuitants who are “internal” to our firm. This gives us some confidence that our results may be more broadly applicable to the U.K. annuitant population as a whole and are not idiosyncratic to our particular firm and its pricing menu.

²¹Welfare of individuals who always choose the middle is not well defined, and the reported results only compute the welfare for those individuals who are estimated to be “rational” and to choose according to the baseline model. For comparability with the other specifications, we still scale the welfare estimates by the overall annuitized amount in the market.

²²We found it somewhat puzzling that payout rates are lower for individuals who approach the company with external funds, and who therefore are more likely to be actively searching across companies. According to the company executives, some of the explanation lies in the higher administrative costs associated with transferring external funds, also creating higher incentives to retain internal individuals by offering them better rates.

2.7 Conclusions

This paper represents, to our knowledge, one of the first attempts to empirically estimate the welfare costs of asymmetric information in an insurance market and the welfare consequences of mandatory social insurance. We have done so in the specific context of the semi-compulsory U.K. annuity market. In this market, individuals who save for retirement through certain tax-deferred pension plans are required to annuitize their accumulated wealth. They are allowed, however, to choose among different types of annuity contracts. This choice simultaneously opens up scope for adverse selection as well as selection based on preferences over different contracts. We estimate that both private information about risk and preferences are important in determining the equilibrium allocation of contracts across individuals. We use our estimates of the joint distribution of risk and preferences to calculate welfare under the current allocation and to compare it to welfare under various counterfactual allocations.

We find that government mandates that eliminate any choice among annuity contracts do not necessarily improve on the asymmetric information equilibrium. We estimate that a mandated annuity contract could increase welfare relative to the current equilibrium by as much as £127 million per year, or could reduce it by as much as £107 million per year, depending on what contract is mandated. Moreover, the welfare maximizing choice for a mandated contract would not be apparent to the government without knowledge of the joint distribution of risk and preferences. Our results therefore suggest that achieving welfare gains through mandatory social insurance may be harder in practice than simple theory would suggest.

Our results also suggest that, relative to a first-best symmetric information benchmark, the welfare cost of asymmetric information along the dimension of guarantee choice is about 25 percent of the maximum money at stake in this choice. These estimates account for about £127 million annually, or about 2 percent of annual premia in the market. However, these quantitative results are less robust to some of the modeling assumptions than the results concerning the optimal mandate.

Although our analysis is specific to the U.K. annuity market, the approach we take can be applied in other insurance markets. As seen, the data requirements for recovering the joint distribution of risk and preferences are data on the menu of choices each individual faces, the contract each chooses, and a measure of each individual's ex-post risk realization. Such data are often available from individual surveys or from insurance companies. These data are now commonly used to test for the presence of asymmetric information in insurance markets, including automobile


insurance (Chiappori and Salanie (2000), Cohen and Einav (2007)), health insurance (Cardon and Hendel (2001)), and long term care insurance (Finkelstein and McGarry (2006)), as well as annuity markets. This paper suggests that such data can now also be used to estimate the welfare consequences of any asymmetric information that is detected, or of imposing mandatory social insurance in the market.

Our analysis was made substantially easier by the assumption that moral hazard does not exist in annuity markets. As discussed, this may be a reasonable assumption for the annuity market. It may also be a reasonable assumption for several other insurance markets. For example, Cohen and Einav (2007) argue that moral hazard is unlikely to be present over small deductibles in automobile insurance. Grabowski and Gruber (2005) present evidence that suggests that there is no detectable moral hazard effect of long term care insurance on nursing home use. In such markets, the approach in this paper can be straightforwardly adopted.

In other markets, such as health insurance, moral hazard is likely to play an important role. Estimation of the efficiency costs of asymmetric information therefore requires some additional source of variation in the data to separately identify the incentive effects of the insurance policies. One natural source would be exogenous changes in the contract menu. Such variation may occur when regulation requires changes in pricing, or when employers change the menu of health insurance plans from which their employees can choose.²³ Non-linear experience rating schemes may also introduce useful variation in the incentive effects of insurance policies (Abbring, Chiappori, and Pinquet (2003), Abbring, Heckman, Chiappori, and Pinquet (2003), Israel (2004)). We consider the application and extension of our approach to other markets, including those with moral hazard, an interesting and important direction for further work.


²³See also Adams, Einav, and Levin (2009) for a similar variation in the context of credit markets.

Figure 2.1: Schematic indifference sets




The figure provides an illustration of the pairs of points (α, β) which would make individuals indifferent between choosing 0 year guarantee and 5 year guarantee (lower left curve) and between 5 year guarantee and 10 year guarantee (upper right curve). These particular curves are computed based on our baseline estimate of λ and the annuity rates faced by 65 year old males; the sets are not a function of the other estimated parameters. Individuals are represented as points in this space, with individuals between the curves predicted to choose 5 year guarantee, and individuals below (above) the lower (upper) curve predicted to choose 0 (10) year guarantee.

Figure 2.2: Estimated distributions



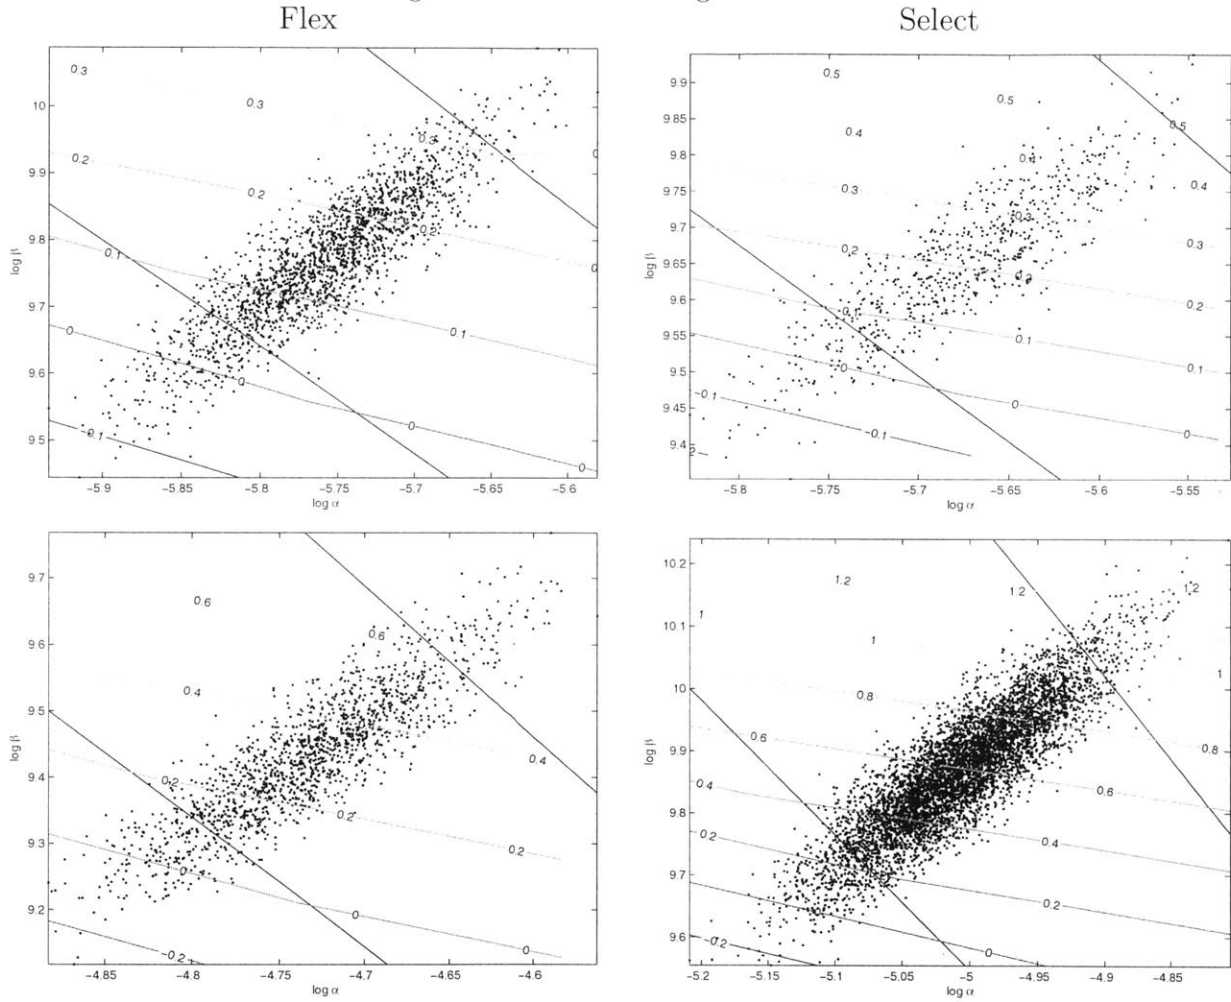
The figure presents the estimated indifference sets for each age-gender cell, with a scatter plots from the estimated joint distribution of $(\log\alpha, \log\beta)$ super-imposed; each point is a random draw from the estimated distribution in the baseline specification. The estimated indifference sets for the 65 year old males are given by the pair of dark dashed lines, for the 60 year old males by the pair of lighter dashed lines, for the 65 year old females by the pair of dotted lines, and for the 60 year old females by the pair of solid lines. The estimated indifference sets for the 65 year old males are the same as those shown in Figure 1 (but a “close up” and in log scale).

Figure 2.3: Wealth contours



The figure super-imposes iso-welfare (wealth equivalent) contour lines on the previous Figure 2. Individuals with wealth equivalent greater than 100 would voluntarily annuitize, while individuals with wealth equivalent less than 100 would not. Each panel represents a different age-gender cell: 60 year old females (upper left), 65 year old females (upper right), 60 year old males (lower left), and 65 year old males (lower right).

Figure 2.4: Welfare change contours



The figure presents Figure 2, with contour lines that present the change in welfare (wealth equivalent) from the counterfactual exercise of symmetric information. Individuals with positive (negative) welfare change are estimated to gain (lose) from symmetric information, compared to their welfare in the observed asymmetric information equilibrium. Each panel represents a different age-gender cell: 60 year old females (upper left), 65 year old females (upper right), 60 year old males (lower left), and 65 year old males (lower right).

Table 2.1: Summary statistics

	60 Females	65 Females	60 Males	65 Males	All
No. of obs.	1,800	651	1,444	5,469	9,364
Fraction choosing 0 year guarantee	14.0	16.0	15.3	7.0	10.2
Fraction choosing 5 year guarantee	83.9	82.0	78.7	90.0	86.5
Fraction choosing 10 year guarantee	2.1	2.0	6.0	3.0	3.2
Fraction who die within observed mortality period:					
Entire sample	8.4	12.3	17.0	25.6	20.0
Among those choosing 0 year guarantee	6.7	7.7	17.7	22.8	15.7
Among those choosing 5 year guarantee	8.7	13.3	17.0	25.9	20.6
Among those choosing 10 year guarantee	8.1	7.7	16.1	22.9	18.5

Recall that we only observe individuals who are alive as of January 1, 1998, and we observe mortality only for individuals who die before December 31, 2005.

Table 2.2: Annuity payment rates

Guarantee Length	60 Females	65 Females	60 Males	65 Males
0	0.1078	0.1172	0.1201	0.1330
5	0.1070	0.1155	0.1178	0.1287
10	0.1049	0.1115	0.1127	0.1198

These are the rates from January 1992, which we use in our baseline specification. A rate is per pound annuitized. For example, a 60 year old female who annuitized X pounds and chose a 0 year guarantee will receive a nominal payment of 0.1078X every year until she dies.

Table 2.3: Parameter estimates

		Estimate	Std. Error
μ_α	60 Females	-5.76	(0.165)
	65 Females	-5.68	(0.264)
	60 Males	-4.74	(0.223)
	65 Males	-5.01	(0.189)
σ_α		0.054	(0.019)
λ		0.110	(0.015)
μ_β	60 Females	9.77	(0.221)
	65 Females	9.65	(0.269)
	60 Males	9.42	(0.300)
	65 Males	9.87	(0.304)
σ_β		0.099	(0.043)
ρ		0.881	(0.415)
No. of Obs.		9,364	

These estimates are for the baseline specification described in the text. Standard errors are in parentheses. Since the value of λ is estimated separately, in a first stage, we bootstrap the data to compute standard errors using 100 bootstrap samples.

Table 2.4: Within-sample fit

	60 Females		65 Females		60 Males		65 Males		Overall	
	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted
Fraction 0 year guarantee	14.00	14.42	15.98	15.32	15.30	14.49	6.99	7.10	10.24	10.22
Fraction 5 year guarantee	83.94	83.16	82.03	83.21	78.67	80.27	89.98	89.75	86.52	86.57
Fraction 10 year guarantee	2.06	2.42	2.00	1.47	6.03	5.25	3.04	3.15	3.24	3.22
Fraction die while observed:										
Entire sample	8.44	7.56	12.29	14.23	17.04	19.73	25.56	25.80	20.03	20.20
0 year guarantee	6.75	6.98	7.69	13.21	17.65	18.32	22.77	23.14	15.75	18.60
5 year guarantee	8.74	7.63	13.30	14.39	16.99	19.86	25.87	25.31	20.60	20.31
10 year guarantee	8.11	8.48	7.69	16.05	16.09	21.67	22.89	27.88	18.48	22.37

This table summarizes the fit of our estimates within sample. For each age-gender cell, we report the observed quantity (identical to Table I) and the corresponding quantity predicted by the model. To construct the predicted death probability, we account for the fact that our mortality data is both censored and truncated, by computing predicted death probability for each individual in the data conditional on the date of annuity choice, and then integrating over all individuals.

Table 2.5: Out-of-sample fit

	60 Females	65 Females	60 Males	65 Males	Overall
Life Expectancy:					
5th percentile	87.4	86.7	79.4	81.4	79.8
Median individual	88.1	87.4	80.0	82.1	82.2
95th percentile	88.8	88.2	80.7	82.8	88.4
U.K. mortality table	82.5	83.3	78.9	80.0	80.5
Expected value of payments:					
0 year guarantee	19.97	20.34	20.18	21.41	20.63
5 year guarantee	19.77	20.01	19.72	20.64	20.32
10 year guarantee	19.44	19.49	19.12	19.61	19.45
Entire sample	19.79	20.05	19.74	20.66	20.32
Break-even interest rate	0.0414	0.0430	0.0409	0.0473	0.0448

This table summarizes the fit of our estimates out of sample. The top panel report life expectancies for different percentiles of the mortality distribution, using the parametric distribution on mortality to predict mortality beyond our mortality observation period. The bottom row of this panel presents the corresponding figures for the average pensioner, based on the PFL/PML 1992 period tables for “life office pensioners” (Institute of Actuaries (1992)). While the predicted life expectancy is several years greater, this is not a problem of fit; a similar difference is also observed for survival probabilities within sample. This simply implies that the average “life office pensioner” is not representative of our sample of annuitants. The bottom panel provides the implications of our mortality estimates for the profitability of the annuity company. These expected payments should be compared with 20, which is the amount annuitized for each individual in the model. Of course, since the payments are spread over a long horizon of several decades, the profitability is sensitive to the interest rate we use. The reported results use our baseline assumption of a real, risk-free interest rate of 0.043. The bottom row provides the interest rate that would make the annuity company break even (net of various fixed costs).

Table 2.6: Welfare estimates

	60 Females	65 Females	60 Males	65 Males	Average
Observed equilibrium:²⁴					
Average wealth-equivalent	100.24	100.40	99.92	100.17	100.16
Maximum Monet at Stake (MMS)	0.56	1.02	1.32	2.20	1.67
Mandate 0 year guarantee counterfactual:²⁵					
Average wealth-equivalent	100.14	100.22	99.67	99.69	99.81
Absolute welfare difference (M pounds)	-30.1	-53.2	-73.7	-146.1	-107.3
Relative welfare difference (as a fraction of MMS)	-0.18	-0.17	-0.19	-0.22	-0.21
Mandate 5 year guarantee counterfactual:					
Average wealth-equivalent	100.25	100.42	99.92	100.18	100.17
Absolute welfare difference (M pounds)	2.8	6.0	1.7	1.6	2.1
Relative welfare difference (as a fraction of MMS)	0.02	0.02	0.004	0.002	0.006
Mandate 10 year guarantee counterfactual:					

Average wealth-equivalent	100.38	100.64	100.19	100.74	100.58
Absolute welfare difference (M pounds)	43.7	72.1	82.3	170.0	126.7
Relative welfare difference (as a fraction of MMS)	0.26	0.23	0.21	0.26	0.25

Symmetric information counterfactual:²⁶

Average wealth-equivalent	100.38	100.64	100.19	100.74	100.58
Absolute welfare difference (M pounds)	43.7	72.0	82.1	169.8	126.5
Relative welfare difference (as a fraction of MMS)	0.26	0.23	0.21	0.26	0.25

Table 2.6: Robustness

Specification	Avg. wealth equivalent	Average absolute welfare difference (million pounds)			
		Mandate 0	Mandate 5	Mandate 10	Symm. info.
1 Baseline specification	100.16	-107.3	2.1	126.7	126.5
Different choices of γ's:					
2 Consumption $\gamma=5$, Wealth after death $\gamma=5$	100.51	-117.0	0.0	111.0	111.0
3 Consumption $\gamma=1.5$, Wealth after death $\gamma=1.5$	99.92	-102.0	0.6	133.2	133.2
4 Consumption $\gamma=3$, Wealth after death $\gamma=5$	100.47	-123.0	3.0	120.0	120.0
5 Consumption $\gamma=3$, Wealth after death $\gamma=1.5$	99.94	-96.9	2.1	135.3	135.3
6 Row 5 + allow heterogeneity in initial wealth ²⁸	101.18	?148.3	-32.9	128.8	127.4
Other parameter choices:					
7 $r=0.05$ and $\delta=0.05$	99.29	-97.5	5.7	119.4	119.4
8 January 1990 annuity rates	100.16	-112.5	0.0	123.0	123.0
Wealth portfolio outside of compulsory annuity:					
9 Fraction annuitized (η) = 0.3	100.65	-118.0	0.0	114.0	114.0
10 Fraction annuitized (?) = 0.1	99.93	?108.0	?4.2	135.0	135.0

²⁷ The table reports summary results – average wealth equivalent and average welfare effects – from a variety of specifications of the model. Each specification is discussed in the text in more detail. Each specification is shown on a separate row of Table VII and differs from the baseline specification of Table VI (which is reproduced in the first row of Table VII) in only one dimension, keeping all other assumptions as in the baseline case.

²⁸ ^a See text for the parameterization of the unobserved wealth distribution. For comparability, the average wealth-equivalent is normalized to be out of 100 so that it is on the same scale as in the other specifications.

11	Allow heterogeneity in η ²⁹	100.22	-113.7	2.5	132.4	141.3
12	Half of initial wealth in public annuity ³⁰	99.95	-426.3	-34.2	243.6	255.6

Parameterization of heterogeneity:

13	Non-Gompertz mortality distribution ³¹	100.06	-100.8	6.0	144.0	144.0
14	α dist. Gamma, β dist. Lognormal	100.20	-111.6	3.0	132.0	132.0
15	α dist. Gamma, β dist. Gamma	100.14	-105.6	3.0	123.0	123.0
16	Allow covariates ³²	100.17	-110.1	3.0	132.0	132.0
17	β fixed, Consumption γ heterogeneous	100.55	-110.0	2.1	129.4	129.3
18	Heterogeneity in both β and γ	100.05	-117.0	-5.9	129.0	131.9

Different information structure

19	Biased belief: $\theta = 0.5$	100.16	-104.0	3.0	122.9	122.9
20	Biased beliefs: $\theta = 2$	100.19	-101.6	5.9	126.0	126.0
21	Uncertain α : $\sigma\epsilon = 0.027$	100.15	-104.7	5.9	128.9	128.9
22	Uncertain α : $\sigma\epsilon = 0.108$	100.17	-105.9	3.0	126.0	126.0

Departure from neo-classical model:

²⁹See text for the parameterization of the unobserved fraction of non-annuitized wealth (η) distribution.

³⁰We assume the public annuity is constant, nominal, and actuarially fair for each person.

³¹This specification uses hazard rate of $\alpha_i \exp(\lambda(t - t_0)^h)$ with $h = 1.5$ (Gompertz, as in the baseline, has $h = 1$).

³²Covariates (for the mean of both α and β) consist of the annuitized amount and the education level at the individual's ward.

³² β is fixed at the estimated μ_β (see Table III). Since the resulting utility function is non-homothetic, we use the average wealth in the population and renormalize, as in row 6. See text for more details.

23	Some individuals always "pick the middle" ³³	100.22	-99.9	9.0	132.0	132.0
Different sample:						
24	"External" individuals ³⁴	95.40	-134.4	-16.8	137.7	137.4

³³The welfare estimates from this specification only compute welfare for the "rational" individuals, ignoring the individuals who are assumed to always pick the middle.

³⁴"External" individuals are individuals who did not accumulated their annuitized funds with the company whose data we analyze. These individuals are not used in the baseline analysis (see Appendix B).

Chapter 3

Testing for Fixed Interactive vs Random Interactive Effects

Abstract. This chapter develops a test for the exogeneity assumptions of classical factor models based on the fixed interactive effects estimator of Bai (2005). The exact form of the test is given for simple linear models. Simulations are used to assess the test's performance. The application of the test to more complicated models is also considered. The test is applied to a model of education as an example.

3.1 Introduction

The use of models with a multifactor error structure has received considerable attention recently. Examples from labor economics include Carneiro, Hanson, and Heckman (2003) and Heckman and Navarro (2005) both of which employ a factor error structure to study individuals' education decisions. Additionally, Andr n and Andr n (2004) use a single factor model to assess a training program, and Coelli, Green, and Warburton (2004) investigate the effect of education on welfare receipt with a factor model. Factor models have also been employed in finance. Recent examples include Favero, Marcellino, and Neglia (2002) and Stock and Watson (2002). Traditional factor analysis, as exemplified by Goldberger (1972), assumes that the individual effects are independent of included regressors. Bai (2005) develops an interactive fixed effects estimator that relaxes this assumption. This chapter considers using the approach of Hausman (1978) to test the exogeneity assumption of traditional, random interactive effects, factor models. Section 3.2 describes the random and fixed interactive effects estimators. Section 3.3 presents a Monte-Carlo study of the

performance of the estimators and the test of exogeneity. Section 3.4 discusses using a similar approach for more complicated factor models. Section 3.4.1 applies the proposed test to a model of education similar to the one found in Cunha, Heckman, and Navarro (2005). Section 3.5 concludes.

3.2 Basic Framework

This chapter is concerned with linear factor models of the form:

$$y_{it} = x_{it}\beta + \theta_i\alpha_t + \epsilon_{it} \quad (3.2.1)$$

y_{it} is an observable outcome. x_{it} are observed covariates, which need not vary over time. θ_i is a $1 \times M$ vector of unobserved individual effects, which have time varying coefficients, α_t . This chapter assumes that ϵ_i is independently and identically distributed with diagonal covariance matrix Σ_{ϵ_i} and uncorrelated with x_{it} .

There is some confusion of language in the literature regarding the terms used to refer to θ_i and α_t . Some authors, such as Pesaran (2004) refer to α_t as the factors and θ_i as loadings. Others, like Cunha, Heckman, and Navarro (2005) call θ_i the factors and α_t the loadings. In part, this difference in terminology reflects slightly different interpretations of equation 3.2.1. When α_t is called the vector of factors, the focus is often on common time shocks, which affect different units of observation differently according to their individual loadings θ_i . On the other hand, when the factors are called θ_i , they are interpreted as individual characteristics that have time varying loadings. These two interpretations differ only in flavor and have no substantial difference. Nonetheless, to avoid confusion, this chapter will avoid using the terms “factors” and “loadings.” Instead, θ_i will be referred to as individual effects, and α_t will be called time effects.

3.2.1 Random Interactive Effects

Classical factor analysis assumes that $\theta_i\alpha_t + \epsilon_{it}$ is uncorrelated with x_{it} . In this case equation 3.2.1 can be estimated by generalized least squares. Let:

$$Y = X\beta + U \quad (3.2.2)$$

be equation 3.2.1 in stacked form. In other words, $Y = \begin{bmatrix} y_{11} \\ y_{12} \\ \vdots \\ y_{NT} \end{bmatrix}$ is a $NT \times 1$ matrix, where

there are N individuals and T time periods. Similarly, X is a $NT \times K$ matrix, where there are K

covariates. $U = \begin{bmatrix} \theta_1\alpha_1 + \epsilon_{11} \\ \theta_1\alpha_2 + \epsilon_{12} \\ \vdots \\ \theta_N\alpha_T + \epsilon_{NT} \end{bmatrix}$ is an $NT \times 1$ matrix. It is assumed that θ_i is independent of θ_j .

The GLS estimator of β is then:

$$\hat{\beta}^{GLS} = (X'(I_N \otimes \Sigma_U^{-1})X)^{-1}X'(I_N \otimes \Sigma_U^{-1})Y \quad (3.2.3)$$

In this equation, Σ_U is the covariance matrix of U_i and is equal to:

$$\begin{aligned} E(U_i U_i') &= E(\alpha' \theta' \theta \alpha + \epsilon) \\ &= \alpha' \Sigma \theta \alpha + \Sigma_{\epsilon_i} \end{aligned} \quad (3.2.4)$$

Since ϵ and θ are uncorrelated with X , β can be consistently estimated with least squares. The residuals can be used to estimate Σ_U , which can then be used to form the GLS estimator. Although it results in no gain in asymptotic efficiency, if desired, this procedure can be iterated to convergence to produce the maximum likelihood estimator.

Σ_U does not uniquely determine θ and α without further restrictions. This chapter will follow Bai (2005) and impose the following normalizations:

$$\alpha' \alpha / T = I_M \quad (3.2.5)$$

$$\Sigma \theta = \text{diagonal} \quad (3.2.6)$$

The choice of normalization does not affect the estimate of β , and these particular normalizations are only made for convenience. If desired, the estimates of θ and α can be rotated to allow a structural interpretation.

If the number of individual effects, M is less than $\frac{T-1}{2}$, an unrestricted estimate of Σ_U might not satisfy equation 3.2.4. In this case, a restricted estimate of Σ_U can be recovered in a number of ways. For example, principal components analysis can be used to extract estimates of α , $\Sigma \theta$, and Σ_{ϵ_i} . These can be used to construct an estimate of Σ_U that satisfies equation 3.2.4.

3.2.2 Fixed Interactive Effects

A number of authors have developed estimators that relax the assumption that θ_i is uncorrelated with x_{it} . For example, Hotz-Eakin, Newey, and Rosen (1988) estimate equation 3.2.1 by quasi-differencing and using lagged variables as instruments. However, this approach rules out time constant regressors. Coakley, Fuertes, and Smith (2002) propose estimating equation 3.2.1 by augmenting the regression of y on x with the principal components of the OLS residuals. Pesaran (2004) shows that this method is inconsistent unless the correlation between x_{it} and θ_i tends to zero or one as $N \rightarrow \infty$. As an alternative, Pesaran (2004) develops a “correlated common effects” estimator, where equation 3.2.1 is augmented by the cross-sectional averages of x_{it} . This estimator is consistent, but it does not allow for time-invariant individual regressors. Ahn, Lee, and Schmidt (2001) develop a GMM estimator for equation 3.2.1. Their estimator is more efficient than least squares under fixed T . However, identification of their estimator requires that x_{it} is correlated with θ_i , making testing the interactive random effects assumption impossible.

This chapter focuses on the interactive fixed effects estimator proposed by Bai (2005). Bai shows that this estimator is \sqrt{NT} consistent and provides monte carlo evidence that it performs well in panels as small as $N = 100$ and $T = 10$. Bai’s estimator is the minimizer of the least squares objective function,

$$SSR(\beta, \alpha, \theta) = \sum_{i=1}^N (y_i - x_i\beta - \theta_i\alpha)'(y_i - x_i\beta - \theta_i\alpha)$$

subject to the normalizations above, with the second normalization (3.2.6) replaced by $\theta'\theta = \text{diagonal}$, since θ is now treated as an incidental parameter instead of a random variable. Bai shows that the least squares estimators of β and α are the solution of the following nonlinear system of equations:

$$\hat{\beta}^{BAI} = (X'(I_N \otimes M_{\hat{\alpha}^{BAI}})X)^{-1}X'(I_N \otimes M_{\hat{\alpha}^{BAI}})Y \quad (3.2.7)$$

$$\hat{\alpha}^{BAI}V_{NT} = \left(\frac{1}{NT} \sum_{i=1}^N (y_i - x_i\hat{\beta}^{BAI})(y_i - x_i\hat{\beta}^{BAI})' \right) \hat{\alpha}^{BAI} \quad (3.2.8)$$

where $M\alpha = I_T - \alpha(\alpha'\alpha)^{-1}\alpha'$ is a projection matrix, and V_{NT} is a diagonal matrix consisting of the M largest eigenvalues of $\left(\frac{1}{NT} \sum_{i=1}^N (y_i - x_i\hat{\beta}^{BAI})(y_i - x_i\hat{\beta}^{BAI})' \right)$ arranged in decreasing order. The solution to this system of equations can be obtained by iteration. As noted by Bai (2005), the iterated solution is somewhat sensitive to starting values. Bai proposes using either the least squares estimator of β or the principal components estimate of α to begin. Like Bai, we find that

using the principal components estimator as initial values is more robust. Generally, the presence of poor initial values is quite obvious and results in an exceptionally large number of iterations. Among the simulation results reported in section 3.3, in no case did both sets of initial values fail to converge to an acceptable solution.

Conditions for identification

Bai (2005) gives detailed conditions for identification. The existence of M interactive effects requires that $\text{plim}_{T \rightarrow \infty} \alpha\alpha'/T$ and $\text{plim}_{N \rightarrow \infty} \theta'\theta/N$ are positive definite. It is also assumed that ϵ_{it} is independent of x_{js} , θ_j , and α_s for all i, t, j , and s . Bai allows for the possibility of weak serial and cross-sectional correlation and time and cross-sectional heteroskedasticity in ϵ_{it} . However, to match the random interactive effects model above, this chapter will assume that there is no serial or cross-sectional correlation in ϵ_{it} . Also, only heteroskedasticity across time will be considered. Thus, if $\epsilon_i = [\epsilon_{i1}, \dots, \epsilon_{iT}]'$, the assumption is that ϵ_i is i.i.d., and $E(\epsilon_i\epsilon_i')$ is diagonal.

The key condition for the identification of β is that the following matrix is positive definite at the true value of α :

$$D(\alpha) = \frac{1}{NT} \sum_{i=1}^N x_i' M \alpha x_i - \frac{1}{T} \left(\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N x_i' M \alpha x_j \theta_i (\theta' \theta / N)^{-1} \theta_j' \right) \quad (3.2.9)$$

With time varying regressors, this condition will be met provided there is sufficient variation in x_{it} . However, one advantage of the interactive fixed effects estimator over the traditional fixed effects estimator is that the former allows for regressors that are constant over time. In the presence of time-invariant regressors, for $D(\alpha)$ to be positive definite, the individual effects must genuinely be interactive. In other words, there cannot be a fixed effect, m , with $\alpha_{tm} = \alpha_{sm}$ for all t and s . Similarly, there cannot be an m such that $\theta_{im} = \theta_{jm}$ for all i and j . Alternatively, one can think of this as ruling out the inclusion of time or individual dummies among x_{it} .

Bai (2005) shows that $\hat{\beta}^{BAI}$ is consistent as N and T approach infinity with T/N converging to zero. If $N/T \rightarrow \rho > 0$, then the estimator is not consistent. Bai develops a biased corrected estimator for this case. Let

$$\hat{B} = \frac{1}{NT} \left(\sum_{i=1}^N \left(x_i - \left(\sum_{j=1}^N x_j \hat{\theta}_i (\hat{\theta}' \hat{\theta} / N)^{-1} \hat{\theta}_j' \right) \right) \hat{\alpha} (\hat{\theta}' \hat{\theta})^{-1} \hat{\theta}_i \hat{\sigma} \hat{\epsilon}^2 \right)$$

$$\hat{C} = \frac{1}{T^2} \sum_{i=1}^N x_i' M \hat{\alpha} \hat{\Sigma} \hat{\epsilon} \hat{\alpha} (\hat{\theta}' \hat{\theta})^{-1} \hat{\theta}_i$$

Then the bias corrected estimator is

$$\hat{\beta}_{BC}^{BAI} = \hat{\beta}^{BAI} + D(\hat{\alpha})^{-1} (\hat{B} + \hat{C}) \quad (3.2.10)$$

This estimator is consistent if ϵ_{it} is not serially or cross-sectionally correlated and both $T/N^2 \rightarrow 0$ and $N/T^2 \rightarrow 0$ as N and $T \rightarrow \infty$. All simulations and results reported in this chapter use the bias corrected version of the estimator.

Asymptotic Distribution

Let

$$D_E = \frac{1}{NT} \sum_{i=1}^N x_i' \Sigma \epsilon M \alpha x_i - \frac{1}{T} \left(\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N x_i' \Sigma \epsilon M \alpha x_j \theta_i (\theta' \theta / N)^{-1} \theta_j' \right)$$

where the only difference from equation 3.2.9 is the inclusion of $\Sigma \epsilon$. Bai (2005) shows that as N and $T \rightarrow \infty$ with $T/N \rightarrow 0$,

$$\sqrt{NT}(\hat{\beta}^{BAI} - \beta) \rightarrow^d N(0, D(\alpha)^{-1} D_E D(\alpha)^{-1}) \quad (3.2.11)$$

Bai shows that D_z and $D(\alpha)$ can be consistently estimated by replacing the parameters in the expressions above with their estimates. Similarly, if $T/N^2 \rightarrow 0$ and $N/T^2 \rightarrow 0$ as N and $T \rightarrow \infty$, then

$$\sqrt{NT}(\hat{\beta}_{BC}^{BAI} - \beta) \rightarrow^d N(0, D(\alpha)^{-1} D_E D(\alpha)^{-1}) \quad (3.2.12)$$

Applying to Unbalanced Panels

Bai (2005) only gives formulas for his estimator for balanced panels. Many panels are not balanced, including the one used in our empirical application below. Therefore, this section derives the formula for unbalanced panels, assuming random attrition. Equation 3.2.7 comes from minimizing the least squares objective with respect to β for a given α , while recognizing that $\hat{\theta}_i = (\hat{\alpha}' \hat{\alpha})^{-1} \hat{\alpha}' (y_i - x_i \hat{\beta})$ is the estimate for θ_i given α and β . In an unbalanced panel, this expression must be modified slightly. Suppose a person is observed at times $[t_i^1, t_i^2, \dots, t_i^{T_i}]$. Redefine y_i to be a $T_i \times 1$ vector consisting of the outcomes that were actually observed. Similarly redefine x_i to be $T_i \times K$. Let $\hat{\alpha}_i$ denote the $T_i \times M$ matrix consisting of the elements of $\hat{\alpha}$ associated with the time periods that person i was in the panel. The formula for $\hat{\theta}_i$ given β and α is then:

$$\hat{\theta}_i(\beta, \alpha) = (\alpha_i' \alpha_i)^{-1} \alpha_i' (y_i - x_i \beta) \quad (3.2.13)$$

The estimate for β given α simply becomes:

$$\hat{\beta}(\alpha) = \left(\sum_i^N x_i' M_{\alpha_i} x_i \right)^{-1} \left(\sum_i^N x_i' M_{\alpha_i} y_i \right) \quad (3.2.14)$$

which is just a trivial variation of the formula for a balanced panel.

However, estimating α given β is not as simple in an unbalanced panel. As discussed by Stock and Watson (1998), the solution to the least squares problem for α cannot be expressed in terms of eigenvectors and eigenvalues when the panel is not balanced. One could numerically solve:

$$\begin{aligned} \hat{\alpha}(\beta) &= \arg \min_{\alpha} \sum_{i=1}^N (y_i - x_i \beta - \theta_i \alpha_i)' (y_i - x_i \beta - \theta_i \alpha_i) \\ \text{s.t. } \theta_i &= (\alpha_i' \alpha_i)^{-1} \alpha_i' (y_i - x_i \beta) \end{aligned}$$

However, computing the solution to this problem is rather time consuming. Instead, Stock and Watson (1998) propose using the EM algorithm to calculate $\hat{\alpha}$. We adopt this approach. This strategy proceeds iteratively. Given initial estimates of α and θ , the data is completed by setting $y_{it} - x_{it} \beta = \theta_i \alpha_t$ when it is missing. From the completed data, $\hat{\alpha}$ is estimated as the square root of T times the M eigenvectors associated with the M largest eigenvalues of $\left(\frac{1}{NT} \sum_{i=1}^N (y_i - x_i \hat{\beta}^{BAI}) (y_i - x_i \hat{\beta}^{BAI})' \right)$, as in the balanced panel case. New estimates of θ are then calculated from 3.2.13 using only data that is actually observed. This process is iterated to convergence. Thus, with an unbalanced panel, calculation of $\hat{\beta}^{BAI}$ and $\hat{\alpha}^{BAI}$ involves two nested iterative process. The inner iteration calculates $\hat{\alpha}(\beta)$ using the EM algorithm just described. The outer iteration alternately computes $\hat{\beta}(\alpha)$ and $\hat{\alpha}(\beta)$ until convergence. Despite all this, computation time is very reasonable. In the empirical example below, with $N = 1343$ and $T = 24$, the entire estimation only takes twenty seconds. The EM algorithm typically converges very quickly, typically in five or fewer iterations. The outer iteration seems to take between twenty and thirty repetitions before convergence.

3.2.3 Testing for Interactive Random vs Fixed Effects

If the identification conditions in section 3.2.2 hold, and θ_i is independent of x_{it} , then the random interactive effects (GLS) estimator is asymptotically efficient, as shown by Chamberlain (1987). Furthermore, under these conditions, Bai's interactive fixed effects estimator will also be consistent, but less efficient. Bai's estimator remains consistent even if θ_i is not independent of x_{it} . However, in this case the random interactive effects estimator will not be consistent. Therefore, the principal of the Hausman (1978) test applies. Under the null hypothesis that θ_i is independent of x_{it} ,

$\hat{\beta}^{GLS}$ is orthogonal to $\hat{\beta}^{BAI}$, since the former is efficient. Therefore, the asymptotic variance of $\hat{\beta}^{GLS} - \hat{\beta}^{BAI}$ is simply $V(\hat{\beta}^{BAI}) - V(\hat{\beta}^{GLS})$. Thus,

$$NT(\hat{\beta}^{BAI} - \hat{\beta}^{GLS})'V(\hat{\beta}^{BAI}) - V(\hat{\beta}^{GLS})^{-1}(\hat{\beta}^{BAI} - \hat{\beta}^{GLS}) \chi_K^2 \quad (3.2.15)$$

where $V(\hat{\beta}^{BAI}) = NTD(\alpha)^{-1}D_E D(\alpha)^{-1}$ and $V(\hat{\beta}^{GLS}) = NT(X'\hat{\Sigma}_U X)^{-1}$. Although the above exposition has presented this as a test of whether θ is independent of x , it is important to keep in mind that rejection of the null could result from any sort of specification error. Therefore, care must be taken in interpreting rejection in real data.

3.3 Simulations of test performance

We conduct a number of Monte Carlo experiments to assess the finite the finite sample properties of the Hausman test for random versus fixed interactive effects. We present results using the graphical approach advocated by Davidson and MacKinnon (1998). P-value plots are used to gauge the agreement of the finite sample distribution of the test statistic and its theoretic asymptotic distribution. The relationship among sample size, true alternative hypothesis, and test performance is assessed with size-power curves.

Throughout this section, the data-generating process is made as simple as possible, while retaining the key features of the model. We set $M = K = 1$. θ_i and ϵ_{it} are generated as standard normal random variables. The regressor is constant through time and is created according to

$$x_{it} = \psi\theta_i + u_i$$

where u_i is a standard normal random variable. ψ is the covariance between θ and x . When the null hypothesis is true, $\psi = 0$. The coefficient on θ at time t is equal to t/T . β is set to one.

Figure 3.1 plots the theoretical size of the test against the simulated size when the null hypothesis is true ($\psi = 0$). Data is generated and the test performed 1,000 times. The figure plots the empirical cumulative distribution function of the test's p-value. If the small sample distribution of the test statistic matches its theoretical distribution, then the size of the test should be equal to the chosen critical value. In this case, the plotted empirical distribution function should coincide with the 45 degree line. If the test rejects too often, the empirical distribution function will lie above the 45 degree line. Here, we see that the empirical distribution function is below the 45 degree line. The test fails to reject too frequently. For example, with a critical value of 10%, the test would only

reject the null hypothesis 5% of the time when the null hypothesis is true. However, the discrepancy between the empirical and theoretical distribution of the p-value is not overly large. Furthermore, in another set of simulations not shown here, with $K = 2$, and $\alpha_t = t^2/T^2$, the samples with $T = 5$ rejected too often, but the samples with $T = 10$ fit the theoretic distribution nearly perfectly.




Figure 3.1: Size under the null hypothesis

This figure shows the empirical distribution function of the test's p-value. The empirical distribution function is defined as:

$$\hat{F}(x) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}(\hat{p}_s < x)$$

where there were S simulations, and \hat{p}_s is the p-value from the s th simulation.

To assess the performance of the estimators and the tests under the alternative hypothesis, we repeat the simulations described in the previous paragraph for $\psi = 0.25$ and $\psi = 0.5$. Tables 3.1-3.3 shows the average and mean squared errors of the two estimators in each set of simulations. Table 3.1 shows that both estimators are unbiased under the null hypothesis. The GLS estimator has a

lower MSE since it is more efficient. Tables 3.2 and 3.3 show that GLS is biased when ψ is not zero. Also, the fixed interactive effects estimator appears to be slightly biased in these small samples. The bias of both estimators increase as sample size decreases. The GLS estimator generally has a larger MSE due to its larger bias.

Table 3.1: MSE and Bias with $\psi = 0$

N	T	$\hat{\beta}_{BC}^{BAI}$		$\hat{\beta}^{GLS}$	
		Bias	MSE	Bias	MSE
100	10	0.001008	0.006534	0.00143	0.003939
100	5	0.006121	0.01517	0.004225	0.006078
50	10	0.007937	0.0152	0.00783	0.009185
50	5	-0.00865	0.0453	-0.007938	0.01401

Table 3.2: MSE and Bias with $\psi = 0.25$

N	T	$\hat{\beta}_{BC}^{BAI}$		$\hat{\beta}^{GLS}$	
		Bias	MSE	Bias	MSE
100	10	0.01535	0.00671	0.08225	0.01315
100	5	0.0252	0.01735	0.1191	0.02575
50	10	0.02736	0.01507	0.08614	0.01868
50	5	0.03481	0.03866	0.1175	0.03024

Table 3.3: MSE and Bias with $\psi = 0.5$

N	T	$\hat{\beta}_{BC}^{BAI}$		$\hat{\beta}^{GLS}$	
		Bias	MSE	Bias	MSE
100	10	0.02427	0.008129	0.1377	0.02743
100	5	0.03485	0.02662	0.1928	0.05198
50	10	0.03529	0.03078	0.1439	0.03323
50	5	0.06862	0.06358	0.1978	0.0599

Figures 3.2 and 3.3 illustrate the power of the test with $\psi = 0.25$ and $\psi = 0.5$, respectively. They show the relationship between the critical value and power of the test under a particular alternative. The closer the curve is to the upper left corner of the figure, the more powerful the test. From the first of these figures, we see that with $\psi = 0.25$, even with $N = 100$ and $T = 10$, the test has fairly low power. At the usual 95% critical value, we would only reject the null about 30% of the time. Given the relatively small sample size and the fact that the random effects estimator is only biased by 10%, the low power of the test is not very surprising or concerning. Most real panel datasets are larger, and a difference in estimates of 10% is often not economically meaningful. Figure 3.3 shows that the performance of the test improves when $\psi = 0.5$. With greater correlation




Figure 3.2: Power with $\psi = 0.25$

This figure shows the power of the test under the alternative hypothesis with $\psi = 0.25$. We plot the empirical distribution function of the p-value, which is defined as:

$$\hat{F}(x) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}(\hat{p}_s < x)$$

where there were S simulations, and \hat{p}_s is the p-value from the s th simulation.




Figure 3.3: Power with $\psi = 0.5$

This figure shows the power of the test under the alternative hypothesis with $\psi = 0.25$. We plot the empirical distribution function of the p-value, which is defined as:

$$\hat{F}(x) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}(\hat{p}_s < x)$$

where there were S simulations, and \hat{p}_s is the p-value from the s th simulation.

between θ and x , the test rejects the null at the 95% critical value three-quarters of the time.

3.4 Testing More Elaborate Models

Typically, empirical work that invokes the interactive random effects assumption is not based on a model as simple as equation 3.2.1. Instead, these papers recognize that one or more components of x_{it} may be correlated with θ_i and include that relationship in the model. To formalize this idea, redefine x_{it} as only those regressors that are thought to be uncorrelated with θ , and let z_{it} be a set of regressors that are related to θ . Equation 3.2.1 becomes:

$$y_{it} = x_{it}\beta + z_{it}\delta + \theta_i\alpha_t + \epsilon_{it} \quad (3.4.1)$$

This equation is augmented by a model of z_{it} :

$$z_{itj} = f_j(w_i, \theta_i, v_{ij})$$

where z_{itj} is the j th component of z_{it} , w_i is vector of regressors uncorrelated with θ_i , v_{ij} is a random variable independent of all other components of the model, and $f_j(\cdot)$ is some function. w_i contain some or all of the components of x_{it} or z_{it} .

Examples of models of this sort include Carneiro, Hanson, and Heckman (2003), Andrén and Andrén (2004), and Coelli, Green, and Warburton (2004). Carneiro, Hanson, and Heckman (2003) use a factor model to estimate the distribution of the return to education. Andrén and Andrén (2004) employ a single factor model to estimate the treatment effect of a vocational training program in Sweden. Coelli, Green, and Warburton (2004) estimate the effect of education on welfare receipt using a factor model.

Even in more complicated models such as these, the fixed interactive effects estimator can form the basis of a specification test. The model must contain a portion that resembles equation 3.4.1. In other words, the model should contain a panel of outcomes that are a linear function of regressors and individual effects. The coefficients on the regressors must be constant across time, and the coefficients on the individual effects must change through time. The application of the test will vary from case to case. Of particular concern is the variance of the difference between the estimate of β from the full model and $\hat{\beta}^{BAI}$. Unless the full model estimate of β is efficient, the variance of the difference is not a simple difference of variances. In this case, the bootstrap can be used to estimate the difference, provided that computation of the estimates is not overly burdensome.

3.4.1 Example: Estimating the Return to College

This section applies the fixed interactive effects specification test to a model of education choice. The model is closely based on Cunha, Heckman, and Navarro (2005). Cunha, Heckman, and Navarro (2005) use a factor structure to endogeneously model individuals' decisions between high school and college and their subsequent earning. Some deviations from Cunha, Heckman, and Navarro's setup are needed to accommodate the fixed interactive effects estimator. These differences will be highlighted below.

Model

The model includes two possible education levels. High school graduates are denoted by $s = 0$, and college graduates are labeled $s = 1$. Earnings at time t for individual i in education level s are given by:

$$y_{its} = x_i\beta_s + \theta_i\alpha_t + \epsilon_{it} \quad (3.4.2)$$

It is assumed that θ_i is independent of x_i , and ϵ_{it} is independent of x_i and θ_i . The number of factors, M , is set to two. In Cunha, Heckman, and Navarro (2005), time periods are defined as 10 years, and y_{its} represents the present value of earnings during that time. Here, a time period will be a single year. Also, Cunha, Heckman, and Navarro (2005) allow the distribution of ϵ_{it} to differ by education level, whereas in this chapter it only differs across time.

More importantly, Cunha, Heckman, and Navarro (2005) specify that α_t varies with s . This is not possible in the fixed interactive effects model above. Operationally, the fixed interactive effects estimator views equation 3.4.2 as:

$$y_{it} = \begin{bmatrix} x_i * (1 - s_i) & x_i * s_i \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} + \theta_i\alpha_t + \epsilon_{it} \quad (3.4.3)$$

Thus, the fixed interactive effects estimator assumes that an individual's θ_i may be correlated with their schooling choice, but has the same effect on the person's earnings regardless of their education. In other words, the estimator makes what Cunha, Heckman, and Navarro (2005) call the 'common coefficient' returns to schooling assumption and rules out unobserved heterogeneity in returns. Cunha, Heckman, and Navarro (2005) strongly criticize this assumption and their results, as well as those of Carneiro, Hanson, and Heckman (2003) violate it. Clearly, this makes the fixed interactive effects estimator less than ideal for this application. A useful generalization of the estimator would be to allow α_t to be a function of x_{it} . Nonetheless, we proceed, keeping in a mind

that a rejection of the null hypothesis in the specification test may simply indicate a failure of the common coefficient assumption.

In the model of Cunha, Heckman, and Navarro (2005), education is modeled with a latent index. Let:

$$s_i^* = z_i\gamma + \theta_i\psi + \eta_i \tag{3.4.4}$$

where z_i is a vector of observed variables, η_i is an unobserved random variable, and γ and ψ are coefficients to be estimated. A person does not to college if ($s_i = 0$) if $s_i^* < 0$, and graduates from college otherwise. As in the earnings equation, θ_i is independent of z_i , and η_i is independent of θ_i , z_i , x_i , and ϵ_{it} .

To aid identification and interpretation of θ_i , the model is augmented with information on test scores. The test scores come from the Armed Service Vocational Aptitude Battery (ASVAB). The ASVAB consists of ten subtests. The same five of these ten tests are used as in Cunha, Heckman, and Navarro (2005). They are a test on coding speed and the four tests that form the Armed Forces Qualifying Test (AFQT), Arithmetic Reasoning, Paragraph Composition, Word Knowledge, and Math Knowledge. These five test scores are meant to be measures of cognitive ability. Let M_{ij} for $j = 1, \dots, 5$ denote the test scores. We estimate:

$$M_{ij} = w_i\delta_j + \theta_i\zeta_j + v_{ij} \tag{3.4.5}$$

where w_i is a vector of observable characteristics, δ_j and ζ_j are parameters, and v_{ij} is a mean zero unobserved random variable. As above, v_{ij} is independent of all other regressors.

Data

A subset of the NLSY79 was used to estimate the above model. Specifically, white males whose highest education level complete was either high school or college and were not missing any necessary variables were included in the analysis. Cunha, Heckman, and Navarro (2005) are interested in estimating the portion of lifetime earnings that was known to individuals at age 18, so they combine the NLSY and PSID to obtain data on earnings at all ages. Estimating an earnings function for the entire life-cycle is not necessary in this chapter, so the PSID is not used. Instead, we just use the NLSY to estimate earnings from age 18 to 41. Table 3.4 shows summary statistics for each variable and indicates which portion(s) of the model use that variable.

Table 3.4: Summary statistics

Variable	Mean	Std. Dev.	Tests	Included in:	
				Schooling	Earnings
Arithmetic Reasoning ¹	0.717	0.955	Yes	-	-
Word Knowledge ¹	0.526	0.728	Yes	-	-
Paragraph Comprehension ¹	0.428	0.8	Yes	-	-
Coding Speed ¹	0.214	0.841	Yes	-	-
Math Knowledge ¹	0.634	1.035	Yes	-	-
Mother's Education ²	12.393	2.279	Yes	Yes	-
Father's Education ²	12.802	3.189	Yes	Yes	-
Family Income ³	23.889	13.799	Yes	Yes	-
Divorced Parents	0.157	0.364	Yes	Yes	-
Number of Siblings	2.824	1.812	Yes	Yes	-
Grew up in the South	0.211	0.408	Yes	Yes	-
Grew up in Urban Area	0.748	0.434	Yes	Yes	-
Age at Test Date	19.255	2.19	Yes	-	-
In School at Test Date	0.572	0.495	Yes	-	-
Education at Test Date ²	12.076	1.667	Yes	-	-
Dist. from Nearest College	7.946	19.172	-	Yes	-
Local College Tuition ³	2.151	0.826	-	Yes	-
High School	0.543	0.498	-	Yes	Yes
College Graduate	0.457	0.498	-	Yes	Yes
N		1343			

¹Standardized in the full NLSY

²Measured in years of schooling

³In thousands of year 2000 dollars

Estimates

Table 3.5 shows the estimates of β from both the full random interactive effects model described above and the fixed interactive effects model of just earnings. The estimates of the full model are computed using Markov Chain Monte Carlo methods described in Carneiro, Hanson, and Heckman (2003). This is the same method used as in Cunha, Heckman, and Navarro (2005). In fact, the same program is used to compute the estimates. The Markov Chain was run for 75,000 iterations. The first 50,000 iterations were discarded. Every fifth of the last 25,000 iterations were saved. The reported estimates are the average of these 5,000 saved iterations. The standard errors are the standard errors of the 5,000 saved iterations. The interactive fixed effects estimates were computed using the unbalanced panel estimator described in section 3.2.2. The standard errors were computed by bootstrap with fifty repetitions. The estimates appear roughly similar, but their standard errors are small enough to suggest that they may be statistically different. This will be formally tested below. To get a better idea of the economic significance of the difference in the estimates, figure 3.4 shows the estimated age-earnings profiles for high school and college with $\theta_i = 0$ and $\epsilon_{it} = 0$. With both estimators, college earnings begin below high school earnings, but college earnings increase more rapidly. However, with the fixed interactive effects estimates, college earnings increase more rapidly and reach a higher level.

Table 3.5: Earnings Estimates

Variable	Interactive FE		Full Model	
	Estimate	(Std. Err.)	Estimate	(Std.Err)
HS	3.82	0.22	4.89	0.26
HS*Age	15.8	0.64	12.7	0.798
HS*Age ²	-9.49	0.458	-7.35	0.522
College	-0.826	0.305	-0.312	0.342
Col*Age	25.6	0.889	24	1.07
Col*Age ²	-13.9	0.631	-13.1	0.691

Test

A Hausman specification test can be applied to the null hypothesis that the variant of the Cunha, Heckman, and Navarro (2005) model is correctly specified. Of particular concern is the assumption that θ is independent of the regressors included in the test equations (3.4.5) and the schooling equation (3.4.4). For example, θ_i is commonly interpreted as representing unobserved ability, and it seems unlikely that ability would be independent of parent's education and income, which are




Figure 3.4: Estimated Age-Earnings Profiles

This figure shows the estimated age-earnings profiles for both the full (CHN) model and the fixed interactive effects (BAI) estimator. We set $\theta_i = 0$ and $\epsilon_{it} = 0$.

included in the test and schooling equations. If θ_i is not independent of these regressors, then the full model is misspecified and estimates of the earnings equation will be inconsistent. On the other hand, the fixed interactive effects estimator will remain consistent, so a Hausman test should be able to detect this sort of misspecification (in the sense that as sample size approaches infinity, the power of the test approached one). As mentioned above, another concern with the model is that α_t might vary with schooling. In this case, both models are misspecified. If this misspecification causes the estimates of the two models to have different probability limits, then the Hausman test might detect it as well.

The test statistic is:

$$t = (\hat{\beta}^{BAI} - \hat{\beta}^{CHN})' V(\hat{\beta}^{BAI} - \hat{\beta}^{CHN})^{-1} (\hat{\beta}^{BAI} - \hat{\beta}^{CHN}) \chi_K^2$$

The Markov Chain Monte Carlo estimate of β is not necessarily the most efficient estimate. Therefore, the variance of the difference of the estimates is not just the difference of variances. Instead, we compute the variance of the difference by bootstrap with fifty replications. For each bootstrap sample, the estimate of the $\hat{\beta}^{CHN}$ using the full model and $\hat{\beta}^{BAI}$ using the fixed interactive effects model are computed as above.

Table 3.6 shows the test statistic and p-value. As shown, we strongly reject the null hypothesis. This may be a sign that the exogeneity assumptions of the full model fail to hold. Specifically, it is likely that parents' education and income are correlated with θ_i . Alternatively, the rejection of the null hypothesis may indicate that both models are misspecified. For example, α_t might depend education level. This seems especially likely given the findings of Cunha, Heckman, and Navarro (2005). They strongly reject the common coefficient returns to schooling assumption imposed by making α_t the same for all schooling levels. Thus, although this test provides evidence that the full model estimated here is misspecified, it does not tell us whether the misspecification is because of an incorrect exogeneity assumption or something else.

Table 3.6: Specification Test

Test Statistic	28.82
P-Value	6.569e-05

3.5 Conclusions

This chapter has presented a method for testing the exogeneity assumptions of classical factor models. The fixed interactive effects estimator of Bai (2005) is consistent whether or not individual effects are correlated with regressors. For simple linear models, GLS is the efficient estimator when individual effects are uncorrelated with the regressors. This allows for easy application of the Hausman (1978) test. The simulations above show that the test performs reasonably well in small samples. Admittedly, such simple linear models are of limited interest. However, the fixed interactive effects estimator can also be used for specification tests of more complicated models. As discussed in section 3.4, as long as the model includes a panel of linear outcomes, the test can be applied. Application of the test to a model of college attendance leads to a rejection of the null hypothesis. However, as discussed above, it is unclear whether the misspecification lies in a violation of the exogeneity assumptions or an incorrect restriction on the factor loadings. A useful generalization of the fixed interactive effects estimator would be to allow factor loadings to depend upon observed regressors.

Appendix A

Appendix to chapter 2

A.1 Additional details about the data

As mentioned in the text, we restrict our sample in several ways:

- As is common in the analysis of annuitant choices, we limit the sample to the approximately sixty percent of annuities that insure a single life. The mortality experience of the single life annuitant provides a convenient ex-post measure of risk; measuring mortality risk of a joint life policy which insures multiple lives is less straightforward (Mitchell, Poterba, Warshawsky, and Brown (1999), Finkelstein and Poterba (2004, 2006)).
- We also restrict the sample to the approximately eighty percent of annuitants who hold only one annuity policy, since characterizing the features of the total annuity stream for individuals who hold multiple policies is more complicated. Finkelstein and Poterba (2006) make a similar restriction.
- We focus on the choice of guarantee period and abstract from a number of other dimensions of individuals' choices.
 - Individuals can choose the timing of their annuitization, although they cannot annuitize before age 50 (45 for women) or delay annuitizing past age 75 (70 for women). We allow average mortality and preferences for wealth after death to vary with age at purchase (as well as gender), but do not explicitly model the timing choice.
 - Annuitants may also take a tax-free lump sum of up to 25 percent of the value of the accumulated assets. We do not observe this decision – we observe only the amount annuitized – and therefore do not model it. However, because of the tax advantage of

the lump sum – income from the annuity is treated as taxable income – it is likely that most individuals fully exercise this option, and ignoring it is therefore unlikely to be a concern.

- To simplify the analysis, we analyze policies with the same payment profile, restricting our attention to the 90 percent of policies that pay a constant nominal payout (rather than payouts that escalate in nominal terms). As an ancillary benefit, this may make our assumption that individuals all have the same discount rate more plausible.
- We limit our sample of annuitants to those who purchased a policy between January 1, 1988 and December 31, 1994. Although we also have data on annuitants who purchased a policy between January 1, 1995 and December 31, 1998, the firm altered its pricing policy in 1995. An exogenous change in the pricing menu might provide a useful source of variation in estimating the model. However, if the pricing change arose due to changes in selection of individuals into the firm – or if it affects subsequent selection into the firm – using this variation without allowing for changes in the underlying distribution of the annuitant parameters (i.e., in the joint distribution of α and β) could produce misleading estimates. We therefore limit the sample to the approximately one-half of annuities purchased in the pre-1995 pricing regime. In principle, we could also separately estimate the model for the annuities purchased in the post-1995 pricing regime. In practice, the small number of deaths among these more recent purchasers created problems for estimation in this sample.
- Annuitants may choose to purchase their annuity from an insurance company other than the one in which their fund has been accumulating, and about one-third of annuitants market-wide choose to do so. As our sample is from a single company, it includes both annuitants who accumulated their fund with the company and stayed with the company, as well as those annuitants who brought in external funds. We limit our main analysis to the approximately two-thirds of individuals in our sample who purchased an annuity with a pension fund that they had accumulated within our company. In the robustness section, we re-estimate the model for the one-third of individuals who brought in external funds, and find similar welfare estimates.
- The pricing of different guarantees varies with the annuitant's gender and age at purchase. We limit our sample of annuitants to those who purchased at the two most common ages of 60 or 65. About three-fifths of our sample purchased their annuity at 60 or 65.

A.2 Proof of Proposition 2

We can write the observed distribution of mortality outcomes and guarantee choices in terms of the unobservables as

$$P(g(\alpha, \beta) \leq y | m_i \leq m) P(m_i \leq m) = \int_0^\infty P(g(\alpha, \beta) \leq y | \alpha) P(m_i \leq m | \alpha) dF_\alpha(\alpha) \quad (\text{A.2.1})$$

The left side of this equation is known from $Z(g, m)$. From Proposition 1 we know that $P(m_i \leq m | \alpha)$ and $F_\alpha(\alpha)$ can be identified from mortality data. Thus, all we need to show is that this equation can be uniquely solved for $P(g(\alpha, \beta) \leq y | \alpha)$. We will use the fact that mortality follows an MPH model to derive an explicit expression for $P(g(\alpha, \beta) \leq y | \alpha)$ in terms of the inverse Laplace transform.¹

Since $P(m_i \leq m | \alpha)$ comes from an MPH model, we can write it as

$$P(m_i \leq m | \alpha) = 1 - e^{-\alpha \Lambda(m)}, \quad (\text{A.2.2})$$

where $\Lambda(m) = \int_0^m \psi(t) dt$ is the integrated hazard function, which increases from 0 to ∞ . Substituting equation (A.2.2) into equation (A.2.1) and rearranging yields

$$\begin{aligned} P(g(\alpha, \beta) \leq y, m_i \leq m) &= \int_0^\infty P(g(\alpha, \beta) \leq y | \alpha) (1 - e^{-\alpha \Lambda(m)}) dF_\alpha(\alpha) = \\ &= \int_0^\infty P(g(\alpha, \beta) \leq y | \alpha) dF_\alpha(\alpha) - \int_0^\infty P(g(\alpha, \beta) \leq y | \alpha) e^{-\alpha \Lambda(m)} dF_\alpha(\alpha) = \\ &= P(g(\alpha, \beta) \leq y) - \int_0^\infty P(g(\alpha, \beta) \leq y | \alpha) e^{-\alpha \Lambda(m)} dF_\alpha(\alpha). \end{aligned} \quad (\text{A.2.3})$$

The first part of the right side of this equation is simply the unconditional cumulative distribution function of g and is known. The remaining integral on the right side is the Laplace transform of $P(g(\alpha, \beta) \leq y | \alpha) f_\alpha(\alpha)$ evaluated at $\Lambda(m)$. It is well known that the Laplace transform is unique and can be inverted. If we let $\mathcal{L}^{-1}\{h(\cdot)\}(\alpha)$ denote the inverse Laplace transform of $h(\cdot)$ evaluated at α , then

$$P(g(\alpha, \beta) \leq y | \alpha) = \frac{1}{f_\alpha(\alpha)} \mathcal{L}^{-1}\{P(g(\alpha, \beta) \leq y) - P(g(\alpha, \beta) \leq y, m_i \leq \Lambda(\cdot))\}(\alpha). \quad (\text{A.2.4})$$

This equation provides an explicit expression for $P(g(\alpha, \beta) \leq y | \alpha)$, so it is identified.

Given $P(g(\alpha, \beta) \leq y | \alpha)$ we can recover $F_{\beta|\alpha}$ if $g(\alpha, \beta)$ is invertible with respect to β , for every

¹Alternatively, we could proceed by noting that for each x , equation (A.2.1) is a Fredholm integral equation of the first kind with kernel $P(m_i \leq m | \alpha)$. We could appeal to the theory of integral equations and linear operators to show that the equation has a unique solution when $P(m_i \leq m | \alpha)$ satisfies an appropriate condition. Proving the proposition in this way would be slightly more general, but it would lead to a highly implicit function that defines $P(g(\alpha, \beta) \leq x | \alpha)$.

α . With invertibility, we can write:

$$\mathbb{P}(g(\alpha, \beta) \leq y | \alpha) = \mathbb{P}\left(\beta \leq g_\beta^{-1}(\alpha, y) | \alpha\right) = F_{\beta|\alpha}(g_\beta^{-1}(\alpha, y) | \alpha). \quad (\text{A.2.5})$$

Thus, we identify $F_{\beta|\alpha}$.

A.3 Additional details about estimation

A.3.1 Likelihood

For each individual we observe mortality data, $m_i = (c_i, t_i, d_i)$, where c_i is the time at which person i entered the sample, t_i is the time at which the person left the sample, and d_i indicates whether the person died ($d_i = 1$) or was censored ($d_i = 0$). The contribution of an individual's mortality to the likelihood, conditional on α_i , is:

$$\begin{aligned} \mathbb{P}(m_i = (c_i, t_i, d_i) | \alpha, \lambda) &= \mathbb{P}(t = t_i | t > c_i, \alpha, \lambda)^{d_i} \mathbb{P}(t \geq t_i | t > c_i, \alpha, \lambda)^{1-d_i} = \\ &= \frac{1}{S(\alpha, \lambda, c_i)} (s(\alpha, \lambda, t_i))^{d_i} (S(\alpha, \lambda, t_i))^{1-d_i}, \end{aligned} \quad (\text{A.3.1})$$

where $S(\alpha, \lambda, t) = \exp\left(\frac{1}{\lambda}(1 - e^{\lambda t})\right)$ is the Gompertz survival function, and $s(\alpha, \lambda, t) = \alpha e^{\lambda t} \exp\left(\frac{1}{\lambda}(1 - e^{\lambda t})\right)$ is the Gompertz density. The log likelihood of the mortality data is computed by integrating equation (A.3.1) over α , and adding up all individuals:

$$L^M(\lambda, \mu_\alpha, \sigma_\alpha | (m_i)_{i=1}^N) = \sum_{i=1}^N \log \left(\int \mathbb{P}(m_i | \alpha, \lambda) \frac{1}{\sigma_\alpha} \phi\left(\frac{\log \alpha - \mu_\alpha}{\sigma_\alpha}\right) d\alpha \right). \quad (\text{A.3.2})$$

We maximize equation (A.3.2) over λ , μ_α , and σ_α to obtain an estimate of λ . The initial estimates of μ_α and σ_α are not used, as we obtain a more efficient estimate of these parameters in the next step (described below).

The contribution of an individual's guarantee choice to the likelihood is based on the guarantee choice model above. Recall that the value of a given guarantee depends on preference for wealth after death β , and annual mortality hazard, which depends on λ and α . Some additional notation will be necessary to make this relationship explicit. Let $V_0^{A(g)}(w_0; \beta, \alpha, \lambda)$ be the value of an annuity with guarantee length g to someone with initial wealth w_0 , Gompertz parameter λ , mortality rate α , and preference for wealth after death β . Conditional on α , the likelihood of choosing a guarantee of length g_i is:

$$\mathbb{P}(g_i | \alpha, \lambda) = \int \mathbf{1}\left(g_i = \arg \max_g V_0^{A(g)}(w_0; \beta, \alpha, \lambda)\right) dF_{\beta|\alpha}(\beta | \alpha) \quad (\text{A.3.3})$$

where $\mathbf{1}(\cdot)$ is an indicator function. As mentioned in the text, we numerically verified that the relative value of a longer guarantee increases with β . Therefore, we know that for each α there is some interval, $[0, \beta_{0,5}^*(\alpha, \lambda))$, such that the zero year guarantee is optimal for all β in that interval. $\beta_{0,5}^*(\alpha, \lambda)$ is the value of β that makes someone indifferent between choosing a 0 and 5 year guarantee. Similarly, there are intervals, $[\beta_{0,5}^*(\alpha, \lambda), \beta_{5,10}^*(\alpha, \lambda))$, where the five year guarantee is optimal, and $[\beta_{5,10}^*(\alpha, \lambda), \infty)$, where the ten year guarantee is optimal.²

We can express the likelihood of an individual's guarantee choice in terms of these indifference cutoffs as:

$$P(g_i|\alpha, \lambda) = \begin{cases} F_{\beta|\alpha}(\beta_{0,5}^*(\alpha, \lambda)) & \text{if } g = 0 \\ F_{\beta|\alpha}(\beta_{5,10}^*(\alpha, \lambda)) - F_{\beta|\alpha}(\beta_{0,5}^*(\alpha, \lambda)) & \text{if } g = 5 \\ 1 - F_{\beta|\alpha}(\beta_{5,10}^*(\alpha, \lambda)) & \text{if } g = 10 \end{cases} \quad (\text{A.3.4})$$

Given our lognormality assumption, the conditional cumulative distribution function $F_{\beta|\alpha}(\cdot)$ can be written as:

$$F_{\beta|\alpha}(\beta^*(\alpha, \lambda)) = \Phi\left(\frac{\log(\beta^*(\alpha, \lambda)) - \mu_{\beta|\alpha}}{\sigma_{\beta|\alpha}}\right) \quad (\text{A.3.5})$$

where $\Phi(\cdot)$ is the normal cumulative distribution function, $\mu_{\beta|\alpha} = \mu_\beta + \frac{\sigma_{\alpha,\beta}}{\sigma_\alpha^2}(\log \alpha - \mu_\alpha)$ is the conditional mean of β , and $\sigma_{\beta|\alpha} = \sqrt{\sigma_\beta^2 - \frac{\sigma_{\alpha,\beta}^2}{\sigma_\alpha^2}}$ is the conditional standard deviation of β . The full log likelihood is obtained by combining $P(g_i|\alpha, \lambda)$ and $P(m_i|\alpha, \lambda)$, integrating over α , taking logs, and adding up over all individuals:

$$L(\mu, \Sigma, \lambda) = \sum_{i=1}^N \log \int P(m_i|\alpha, \lambda)P(g_i|\alpha, \lambda) \frac{1}{\sigma_\alpha} \phi\left(\frac{\log \alpha - \mu_\alpha}{\sigma_\alpha}\right) d\alpha. \quad (\text{A.3.6})$$

We calculate the integral in equation (A.3.6) by quadrature. Let $\{x_j\}_{j=1}^M$ and $\{w_j\}_{j=1}^M$ be M quadrature points and weights for integrating from $-\infty$ to ∞ . Person i 's contribution to the likelihood is:

$$L_i(\mu, \Sigma, \lambda) = \sum_{j=1}^M P(m_i|\alpha = e^{x_j\sigma_\alpha + \mu_\alpha}, \lambda)P(g_i|\alpha = e^{x_j\sigma_\alpha + \mu_\alpha}, \lambda)\phi(x_j)w_j. \quad (\text{A.3.7})$$

We maximize the likelihood using a gradient based search. Specifically, we use the modeling language AMPL along with the SNOPT sequential quadratic programming algorithm (Gill, Murray, and Saunders (2002)) for maximization.

²Note that it is possible that $\beta_{0,5}^*(\alpha, \lambda) > \beta_{5,10}^*(\alpha, \lambda)$. In this case there is no interval where the five year guarantee is optimal. Instead, there is some $\beta_{0,10}^*(\alpha, \lambda)$ such that a 0 year guarantee is optimal if $\beta < \beta_{0,10}^*(\alpha, \lambda)$ and a 10 year guarantee is optimal otherwise. This situation (which does not create potential estimation problems, but simply implies that a 5 year guarantee is never optimal) only arises for high values of α 's that are well outside the range of our mortality data.

A.3.2 Guarantee indifference curves

As mentioned in the text, the most difficult part of calculating the likelihood is finding the points where people are indifferent between one guarantee option and another, that is finding $\beta_{0,5}^*(\alpha, \lambda)$ and $\beta_{5,10}^*(\alpha, \lambda)$. To find these points we need to compute the expected utility associated with each guarantee length.

The value of a guarantee of length g with associated annual payments $z_t(g)$ is

$$V^{A(g)}(w_0; \alpha, \beta) = \max_{c_t, w_t} \sum_{t=0}^T a_t(\alpha) \delta^t \frac{c_t^{1-\gamma}}{1-\gamma} + \beta f_t(\alpha) \delta^t \frac{(w_t + Z_t(g))^{1-\gamma}}{1-\gamma} \quad (\text{A.3.8})$$

$$\text{s.t. } w_{t+1} = (1+r)(w_t + z_t(g) - c_t) \geq 0$$

where δ is the discount factor, r is the interest rate, and $Z_t(g) = \sum_{\tau=t}^{t_0+g} \left(\frac{1}{1+r}\right)^{\tau-t} z_\tau(g)$ is the present discounted value of guaranteed future payments at time t . Also, $a_t(\alpha) = \prod_{\tau=1}^t (1 - \kappa_\tau(\alpha))$ is the probability of being alive at time t and $f_t(\alpha) = \kappa_t(\alpha) \prod_{\tau=1}^{t-1} (1 - \kappa_\tau(\alpha))$ is the probability of dying at time t . Note that a person who dies at time t , dies before consuming c_t or receiving $z_t(g)$. Technically, there are also no borrowing constraints and non-negativity constraints on wealth and consumption. However, it is easy to verify that these constraints never bind, the former due to the fact that the individuals are retirees who do not accumulate new income, and the latter due to the form of the utility functions.

We used the first order conditions from equation (A.3.8) to collapse the problem to a numerical optimization over a single variable, consumption at time zero. The first order conditions for equation (A.3.8) are

$$\delta^t a_t(\alpha) c_t^{-\gamma} = \psi_t \quad \forall t \in \{0, 1, \dots, T\} \quad (\text{A.3.9})$$

$$\delta^t f_t(\alpha) \beta (w_t + G_t^g)^{-\gamma} = -\psi_t + \frac{1}{1+r} \psi_{t-1} \quad \forall t \in \{1, 2, \dots, T\} \quad (\text{A.3.10})$$

$$(w_t + z_t - c_t)(1+r) = w_{t+1} \quad \forall t \in \{0, 1, \dots, T-1\} \quad (\text{A.3.11})$$

where ψ_t is the Lagrange multiplier on the budget constraint at time t . Initial wealth w_0 is taken as given. It is not possible to completely solve the first order conditions analytically. However, suppose we knew c_0 . Then from the budget constraint (equation (A.3.11)), we can calculate w_1 . From the first order condition for c_0 (equation (A.3.9)), we can find ψ_0 :

$$\psi_0 = s_0(\alpha) \delta^0 c_0^{-\gamma}. \quad (\text{A.3.12})$$

We can then use the first order condition for w_1 to solve for ψ_1

$$\psi_1 = -f_1(\alpha)\delta^1\beta(w_1 + G_1^g)^{-\gamma} + \frac{1}{1+r}\psi_0. \quad (\text{A.3.13})$$

Then, ψ_1 and the first order condition for c_t gives c_1 :

$$c_1 = \left(\frac{\psi_1}{\delta^1 a_1(\alpha)} \right)^{-1/\gamma}. \quad (\text{A.3.14})$$

Continuing in this way, we can find the whole path of optimal c_t and w_t associated with the chosen c_0 . If this path satisfies the non-negativity constraints on consumption and wealth, then we have defined a value function of c_0 , $\tilde{V}(c_0, g, \alpha, \beta)$. Thus, we can reformulate the optimal consumption problem as an optimization problem over one variable.

$$\max_{c_0} \tilde{V}(c_0, g, \alpha, \beta). \quad (\text{A.3.15})$$

Numerically maximizing a function of a single variable is a relatively easy problem and can be done quickly and robustly. We solve the maximization problem in equation (A.3.15) using a simple bracket and bisection method. To check our program, we compared the value function as computed in this way and by an earlier version of the program that used a discretization and backward induction approach. They agreed up to the expected precision.

Finally, the guarantee cutoffs, $\beta_{0,5}^*(\alpha, \lambda)$ and $\beta_{5,10}^*(\alpha, \lambda)$, are defined as the solution to

$$V^{A(0)}(w_0; \alpha, \beta_{0,5}^*(\alpha, \lambda)) = V^{A(5)}(w_0; \alpha, \beta_{0,5}^*(\alpha, \lambda)) \quad (\text{A.3.16})$$

$$V^{A(5)}(w_0; \alpha, \beta_{5,10}^*(\alpha, \lambda)) = V^{A(10)}(w_0; \alpha, \beta_{5,10}^*(\alpha, \lambda)) \quad (\text{A.3.17})$$

For each α , we solve for these cutoff points using a simple bisection search. Each evaluation of the likelihood requires knowledge of $\beta_{0,5}^*(\alpha(x_j), \lambda)$ and $\beta_{5,10}^*(\alpha(x_j), \lambda)$ at each integration point x_j . Maximizing the likelihood requires searching over μ_α and σ_α , which will shift $\alpha(x_j)$. As mentioned in the text, rather than recomputing these cutoff points each time $\alpha(x_j)$ changes, we initially compute them on a dense grid of values of α , and log-linearly interpolate as needed.

Bibliography

- ABBRING, J. H., P.-A. CHIAPPORI, J. J. HECKMAN, AND J. PINQUET (2003): “Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?,” *Journal of the European Economic Association*, 1(2/3), 512–521.
- ABBRING, J. H., P.-A. CHIAPPORI, AND J. PINQUET (2003): “Moral Hazard and Dynamic Insurance Data,” *Journal of the European Economic Association*, 1(4), 767–820.
- ADAMS, W., L. EINAV, AND J. LEVIN (2009): “Liquidity Constraints and Imperfect Information in Subprime Lending,” *American Economic Review*, 99, 49–84.
- AHN, S., Y.-H. LEE, AND P. SCHMIDT (2001): “GMM Estimation of Linear Panel Data Models with Time-varying Individual Effects,” *Journal of Econometrics*.
- AI, C., AND X. CHEN (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71(6), 1795–1843.
- AKERLOF, G. A. (1970): “The Market for ”Lemons”: Quality Uncertainty and the Market Mechanism,” *The Quarterly Journal of Economics*, 84(3), 488–500.
- ANDRÉN, T., AND D. ANDRÉN (2004): “Assessing The Employment Effects Of Vocational Training Using A One-Factor Model,” (133), available at <http://ideas.repec.org/p/hhs/gunwpe/0133.html>.
- BAI, J. (2005): “Panel Data Models with Interactive Fixed Effects,” .
- BAJARI, P., V. CHERNOZHUKOV, H. HONG, AND D. NEKIPELOV (2009): “Nonparametric and Semiparametric Analysis of a Dynamic Discrete Game,” Discussion paper.
- BANKS, J., AND C. EMMERSON (1999): “U.K. Annuitants - IFS Briefing Note,” Discussion paper, Institute for Fiscal Studies, London.
- BANKS, J., C. EMMERSON, Z. OLDFIELD, AND G. TETLOW (2005): “Prepared for Retirement? The Adequacy and Distribution of Retirement Resources in England,” Discussion paper, Institute for Fiscal Studies, London.
- BARSKY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study*,” *Quarterly Journal of Economics*, 112(2), 537–579.
- BERRY, S., AND A. PAKES (2001): “Estimation from the Optimality Conditions for Dynamic Controls,” .

- BERRY, S., AND E. TAMER (2006): "Identification in Models of Oligopoly Entry," in *Advances in economics and econometrics: theory and applications, ninth world congress*, ed. by R. Blundell, W. K. Newey, and T. Persson. Cambridge University Press.
- BLEVINS, J. (2009): "Nonparametric Identification of Dynamic Games with Discrete and Continuous Choices," Discussion paper, Duke University.
- BRAUN, M., AND A. MUERMANN (2004): "The Impact of Regret on the Demand for Insurance," *The Journal of Risk and Insurance*, 71(4), 737–767.
- BROWN, J., O. MITCHELL, J. POTERBA, AND M. WARSHAWSKY (2001): *The Role of Annuity Markets in Financing Retirement*. MIT Press, Cambridge, MA.
- BROWN, J. R. (2001): "Private pensions, mortality risk, and the decision to annuitize," *Journal of Public Economics*, 82(1), 29–62.
- BUNDORF, M. K., J. D. LEVIN, AND N. MAHONEY (2008): "Pricing and Welfare in Health Plan Choice," *National Bureau of Economic Research Working Paper Series*, No. 14153.
- CARDON, J. H., AND I. HENDEL (2001): "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *The RAND Journal of Economics*, 32(3), 408–427.
- CARLIN, C., AND R. J. TOWN (2007): "Adverse Selection, Welfare and Optimal Pricing of Employer-Sponsored Health Plans," Discussion paper, University of Minnesota.
- CARNEIRO, P., K. T. HANSON, AND J. J. HECKMAN (2003): "Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice," *International Economic Review*.
- CARRASCO, M., M. CHERNOV, J.-P. FLORENS, AND E. GHYSELS (2007): "Efficient estimation of general dynamic models with a continuum of moment conditions," *Journal of Econometrics*, 140(2), 529 – 573.
- CARRASCO, M., AND J. FLORENS (2000): "Generalization of Gmm to a Continuum of Moment Conditions," *Econometric Theory*, 16(06), 797–834.
- CHAMBERLAIN, G. (1987): "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*.
- CHIAPPORI, P., AND B. SALANIE (2000): "Testing for Asymmetric Information in Insurance Markets," *The Journal of Political Economy*, 108(1), 56–78.
- COAKLEY, J., A.-M. FUERTES, AND R. SMITH (2002): "A Principal Components Approach to Cross-Section Dependence in Panels," (B5-3), available at <http://ideas.repec.org/p/cpd/pd2002/b5-3.html>.
- COELLI, M., D. GREEN, AND W. WARBURTON (2004): "Breaking the cycle? The effect of education on welfare receipt among children of welfare recipients," (W04/14).
- COHEN, A., AND L. EINAV (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97, 745–788.

- CUNHA, F., J. J. HECKMAN, AND S. NAVARRO (2005): "Separating Uncertainty from Heterogeneity in Life Cycle Earnings," (11024), available at <http://ideas.repec.org/p/nbr/nberwo/11024.html>.
- DAVIDOFF, T., J. R. BROWN, AND P. A. DIAMOND (2005): "Annuities and Individual Welfare," *American Economic Review*, 95, 1573–1590.
- DAVIDSON, R., AND J. G. MACKINNON (1998): "Graphical Methods for Investigating the Size and Power of Hypothesis Tests," *The Manchester School of Economic & Social Studies*, 66(1), 1–26, available at <http://ideas.repec.org/a/bla/manch2/v66y1998i1p1-26.html>.
- DAVIS, S. J., F. KUBLER, AND P. WILLEN (2006): "OVER THE LIFE CYCLE," *Review of Economics and Statistics*, 88(May), 348–362.
- DYNAN, K., J. SKINNER, AND S. ZELDES (2004): "Do the Rich Save More?," *The Journal of Political Economy*, 112(2), 397–444.
- EINAV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125(3), 877–921.
- EINAV, L., A. FINKELSTEIN, AND P. SCHRIMPF (2007): "The Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market," Discussion paper, NBER working paper No. 13228.
- ELBERS, C., AND G. RIDDER (1982): "True and Spurious Duration Dependence: The Identifiability of the Proportional Hazard Model," *The Review of Economic Studies*, 49(3), 403–409.
- ENGEN, E. M., W. G. GALE, C. E. UCCELLO, C. D. CARROLL, AND D. I. LAIBSON (1999): "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999(2), 65–187
CR – Copyright © 1999 The Brookings Institution.
- FANG, H., M. P. KEANE, AND D. SILVERMAN (????): "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market," .
- FAVERO, C. A., M. MARCELLINO, AND F. NEGLIA (2002): "Principal components at work: The empirical analysis of monetary policy with large datasets," (223), available at <http://ideas.repec.org/p/igi/igierp/223.html>.
- FELDSTEIN, M. (2005): "Rethinking Social Insurance," Discussion paper, NBER Working Paper No. 11250.
- FINKELSTEIN, A., AND K. MCGARRY (2006): "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *The American Economic Review*, 96(4), 938–958.
- FINKELSTEIN, A., AND J. POTERBA (????a): "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market," .
- (????b): "Selection Effects in the United Kingdom Individual Annuities Market," .
- FINKELSTEIN, A., AND J. POTERBA (2006): "Testing for Adverse Selection with 'Unused Observables'," Discussion paper, NBER Working Paper No. 12112.

- FINKELSTEIN, A., J. POTERBA, AND C. ROTHSCHILD (2009): “Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities,” *Journal of Financial Economics*, 91(1), 38–58.
- GILL, P., W. MURRAY, AND M. A. SAUNDERS (2002): “SNOPT: An SQP algorithm for large-scale constrained optimization,” *SIAM Journal of Optimization*, 12, 979–1006.
- GINÉ, E., AND R. NICKL (2007): “Uniform central limit theorems for kernel density estimators,” *Probability Theory and Related Fields*, 141(3-4), 333–387.
- GOLDBERGER, A. S. (1972): “Maximum-Likelihood Estimation of Regressions Containing Unobservable Independent Variables,” *International Economic Review*.
- GOLOSOV, M., , AND A. TSYVINSKI (2007): “Optimal Taxation with Endogenous Insurance Markets,” *Quarterly Journal of Economics*, 122, 487–534.
- GRABOWSKI, D. C., , AND J. GRUBER (2005): “Moral Hazard in Nursing Home Use,” *NBER Working Paper No. 11723*.
- HAMERMESH, D. (1985): “Expectations, Life Expectancy, and Economic Behavior,” *Quarterly Journal of Economics*, pp. 389–408.
- HAN, A., , AND J. HAUSMAN (1990): “Flexible Parametric Estimation of Duration and Competing Risk Models,” *Journal of Applied Econometrics*, 5, 1–28.
- HANSEN, B. E. (2008): “Uniform Convergence Rates for Kernel Estimation with Dependent Data,” *Econometric Theory*, 24(03), 726–748.
- HAUSMAN, J. (1978): “Specification Tests in Econometrics,” *Econometrica*.
- HECKMAN, J. J., , AND B. SINGER (1984): “The Identifiability of the Proportional Hazard Model,” *Review of Economic Studies*, 51, 231–241.
- HECKMAN, J. J., AND S. NAVARRO (2005): “Dynamic Discrete Choice and Dynamic Treatment Effects,” (0316), available at <http://ideas.repec.org/p/nbr/nberte/0316.html>.
- HONG, H., AND M. SHUM (2009): “Pairwise-Difference Estimation of a Dynamic Optimization Model,” *Review of Economic Studies*, 77(1), 273–304.
- HORIUCHI, S., , AND A. COALE (1982): “A simple Equation for Estimating the Expectation of Life at Old Ages,” *Population Studies*, 36, 317–326.
- HOSSEINI, R. (2008): “Adverse Selection in the Annuity Market and the Role for Social Security,” Discussion paper, Arizona State University.
- HOTZ-EAKIN, D., W. NEWEY, AND H. S. ROSEN (1988): “Estimating Vector Autoregressions with Panel Data,” *Econometrica*.
- HU, Y., AND M. SHUM (2008): “Identifying Dynamic Games with Serially-Correlated Unobservables,” Discussion paper, The Johns Hopkins University, Department of Economics.
- HUBBARD, G., J. SKINNER, AND S. ZELDES (1995): “Precautionary Savings and Social Insurance,” *Journal of Political Economy*, 103, 360–99.

- HURD, M. (1989): “Mortality Risk and Bequests,” *Econometrica*, 57, 779–813.
- HURD, M. D., , AND K. MCGARRY (2002): “The Predictive Validity of Subjective Probabilities of Survival,” *Economic Journal*, 112, 966–985.
- ISRAEL, M. (2004): “Do We Drive More Safely When Accidents,” Discussion paper.
- JENKINS, M., P. LIU, R. L. MATZKIN, AND D. L. MCFADDEN (2004): “The Browser War — Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects,” Discussion paper.
- KOPCZUK, W., , AND J. LUPTON (2007): “To Leave or not To Leave: The Distribution of Bequest Motives,” *Review of Economic Studies*, 74, 207–235.
- KOPCZUK, W. (2007): “Bequest and Tax Planning: Evidence from Estate Tax Returns,” *Quarterly Journal of Economics*, 122, 1801–1854.
- KOSOROK, M. R. (2008): *Introduction to empirical processes and semiparametric inference*. Springer Science.
- KOTLIKOFF, L. J., , AND A. SPIVAK (1981): “The Family as an Incomplete Annuities Market,” *Journal of Political Economy*, 89, 372–391.
- LAIBSON, D. I., A. REPETTO, AND J. TOBACMAN (1998): “Self-Control and Saving for Retirement,” *Brookings Papers on Economic Activity*, 1, 91–196.
- LAITNER, J., , AND J. F. THOMAS (1996): “New Evidence on Altruism: A Study of TIAA-CREF Retirees,” *American Economic Review*, 86, 893–908.
- LUSTIG, J. (2008): “The Welfare Effects of Adverse Selection in Privatized Medicare,” Discussion paper, Boston University.
- MACIEIRA, J. (2009): “A Dynamic Model of Innovation and Technological Competetion in the Supercomputer Industry,” Discussion paper.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70(2), 801–816.
- MITCHELL, O. S., J. POTERBA, M. WARSHAWSKY, AND J. R. BROWN (1999): “New Evidence on the Money’s Worth of Individual Annuities,” *American Economic Review*, 89, 1299–1318.
- MIYAZAKI, H. (1977): “The Rat Race and Internal Labor Markets,” *Bell Journal of Economics*, 8, 394–418.
- MONEYFACTS, ., AND A. DAILY (1995): “Annuities Daily Update January 4,” *Moneyfacts Publications*.
- MURTHI, M., J. M. ORSZAG, AND P. R. ORSZAG (1999): “The Value for Money of Annuities in the UK: Theory, Experience and Policy,” Discussion paper.
- OF ACTUARIES, I., AND C. M. I. C. FACULTY OF ACTUARIES (1999): “Continuous Mortality Investigation Reports, Number 17,” Discussion paper.
- OF BRITISH INSURERS, A. (1999): “Insurance Statistics Year Book: 1986-1998,” Discussion paper.

- OF NATIONAL STATISTICS, O. (2006): "The Pensioners' Incomes Series 2004/5," *Pensions Analysis Directorate*.
- PALUMBO, M. G. (1999): "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies*, 66, 395–421.
- PESARAN, M. H. (2004): "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure," (CESifo Working Paper No.), available at http://ideas.repec.org/p/ces/ceswps/_1331.html.
- RIDDER, G. (1990): "The Non-Parametric Identification of Generalized Hazard Models," *Review of Economic Studies*, 57, 167–182.
- ROTHSCHILD, M., , AND J. E. STIGLITZ (1976): "Equilibrium in Information," *Quarterly Journal of Economics*, 90, 630–649.
- RUST, J. (1994): "Chapter 51 Structural estimation of markov decision processes," vol. 4 of *Handbook of Econometrics*, pp. 3081–3143. Elsevier.
- SCHOLZ, J. K., A. SESHADRI, AND S. KHITATRAKUN (2006): "Are Americans Saving 'Optimally' for Retirement?," *Journal of Political Economy*, 114, 607–643.
- SHESHINSKI, E. (2006): "Differentiated Annuities in a Pooling Equilibrium," Discussion paper, Hebrew University of Jerusalem.
- SIMONSON, I., , AND A. TVERSKY (1992): "Choice in Context: Trade-off Contrast and Extremeness Aversion," *Journal of Marketing Research*, 29, 281–295.
- SMITH, V. K., D. TAYLOR, AND F. SLOAN (2001): "Longevity Demise?," *American Economic Review*, 91, 1126–1134.
- SPENCE, M. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics*, 10, 427–447.
- SRISUMA, S. (2010): "Minimum distance estimation for a class of Markov decision processes," Discussion paper, London School of Economics and Political Science.
- SRISUMA, S., AND O. B. LINTON (2010): "Semiparametric Estimation of Markov Decision Processes with Continuous State Space," *SSRN eLibrary*.
- STOCK, J. H., AND M. W. WATSON (1998): "Diffusion Indexes," (6702), available at <http://ideas.repec.org/p/nbr/nberwo/6702.html>.
- (2002): "Forecasting Using Principal Components From a Large Number of Predictors," *Journal of the American Statistical Association*.
- VAN DEN BERG, G. J. (2001): "Duration Models: Specification, Identification and Multiple Durations," *Handbook of Econometrics*, 5.
- VAN DER VAART, A. (1994): "Weak Convergence of Smoothed Empirical Processes," *Scandinavian Journal of Statistics*, 21(4), 501–504.