Essays on Subcontracting, Competitive Bidding, and Dynamic Housing Demand

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Daniel Patrick Miller

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor Of Philosophy

Patrick L. Bajari

December, 2009

ACKOWLEDGEMENT

I thank my advisor Patrick Bajari for his continuous support and guidance, as well as Amil Petrin, Kyoo il Kim, Tom Holmes, Bob Town, Minjung Park, Jungwon Yeo, Steve Tadelis, Francine Lafontaine, Greg Lewis, and participants of the applied microeconomics workshop at the University of Minnesota for helpful suggestions. I am grateful to the contractors and structural engineers from Flatiron-Manson Constructors and Figg Engineering for discussions about the bridge construction industry. Financial support from the University of Minnesota Graduate Research Partners Program and Doctoral Dissertation fellowships are acknowledged.

I recognize the contributions of my co-authors on the housing market research: Patrick Bajari, Phoebe Chan, and Dirk Krueger.

Dedication

I would like to dedicate this dissertation in memory of my father Dr. Tom R. Miller Ph.D, MD. He taught me how to be a scholar and was someone with whom I could share my interests in learning. Many of our best times together were in conversations about what I had been learning in school or, later on, about the details of my research projects. He instilled in me a genuine passion and dedication to research and education.

Abstract

The first essay quantifies and compares the impact of contractual incompleteness on subcontracting and in-house contracting costs. I examine 2,200 individual construction work items (i.e. drilling, concrete, traffic striping) on 32 bridge projects procured by the California Department of Transportation through competitive bidding. I use ex-post revisions to work item contracts to construct a measure of contractual incompleteness. I model strategic bidding behavior and derive a new structural approach to estimate costs from bids. The results show that contractual incompleteness raises procurement costs, up to 12% for subcontracted work. The effect for in-house work is much smaller. The results provide one of the first pieces of quantitative evidence supporting the incomplete contracting theories of the firm.

In the U.S., macroeconomic policy makers are concerned about how consumers will respond to falling incomes, nominal home prices, falling income, rising mortgage interest rates and tightening credit standards. In the second essay, we estimate and simulate a dynamic structural model of housing demand. The model allows for realistic features of the housing market including non-convex adjustment costs from buying and selling a home and credit constraints from minimum downpayment requirements.

We use the forward simulation procedure of Bajari, Benkard and Levin (2007) to estimate the structural parameters using data from the Panel Study of Income Dynamics. Given the estimated parameters, we simulate the partial equilibrium consumption and housing and financial asset accumulation response by consumers to negative income and home price shocks and a tightening of credit constraints.

The simulation results demonstrate that many households do not adjust their stock of housing, but rather absorb the negative shocks by depleting home equity and reducing consumption. The intuition behind this result is simple— households only move two to three times before retirement. Because they are locked in, changes in housing market conditions do not influence housing decisions.

Contents

List of Tables

List of Figures

Chapter 1

Subcontracting and Competitive Bidding on Incomplete Procurement Contracts

1.1 Introduction

Subcontractors play a vital role in the construction industry. They perform 50% of the work on civil projects and 75% to 100% on a typical private construction project.[1] But, the contractual hazards of subcontracting plague the industry. Sweet (2004), an expert on the legal aspects of construction contracting, expresses this sentiment by titling his chapter on subcontracting "The Achilles Heel of Construction Management." The contractual problems he refers to are typically caused by revisions and modifications to the construction plans. The hold-up costs can be enormous; take for example Boston's notorious "Big Dig."[2] On more typical projects, changes disrupt day-to-day construction operations but can escalate to problems that require arbitration or litigation.[3]

Contractual incompleteness, the cause of these hold-up problems, is the key ingredient in leading theories of the firm. Coase (1937) laid the foundation for the incomplete contracting theories which have been expanded and formalized in the transactions cost (Williamson, 1985), property rights (Hart and Moore, 1990), and, most recently, reference point (Hart and Moore, 2008) theories. Each emphasizes different mechanisms through which hold-up costs will appear. The property rights theory emphasizes how contract renegotiations distort incentives to make ex-ante investments in cost reductions. The transactions cost theory and the more formal reference point theory emphasize the costly haggling of ex-post renegotiations.

The goal of this paper is to quantify the impact of contractual incompleteness on both subcontracting costs and the cost for the alternative arrangement of performing work in-house. Through a combination of both channels, exante distortions and ex-post haggling, costs are predicted to rise under both arrangements as the degree of incompleteness increases. But, the extent to

[1]Source: Bartholomew (1998).

[2]The "Big Dig" was a multi-billion dollar project that encountered unknown underground conditions. Modifications called for tunnel rerouting and strengthening. Much of the affected work was performed by subcontractors and several claims were filed. In the largest, the primary contractor paid $417 million in restitution for mismanagement of subcontractors. Source: Washington Post Dec 26. 2007, "On Dec. 31, It's Official: Boston's Big Dig Will Be Done."

[3]Semple et al. (1994) finds the average legal claim in their sample delays construction by 60% of the contract duration and comes with a cost equal to the value of the original contract.

which the effects differ across arrangements will, at least in part, influence contractors' decisions to hire subcontractors, or — using the term coined in the literature — their "make or buy" decisions. To provide evidence for the theory of the firm, I test whether the cost effects differ between subcontracting and in-house contracting.

The application is to 32 bridge construction projects procured by the California Department of Transportation (Caltrans) through competitive bidding. Competitively tendered public construction projects are an excellent test-bed for two reasons. First, there is a measurable dimension by which contracts are incomplete: unforeseen revisions and modifications to blueprints, plans, and specifications. I propose a proxy measure of incompleteness that captures contractors' expectations about revisions. Second, I have data on bids. Bids reflect costs. Expected costs (including hold-up costs) pass through to bids because submitted bids, not costs actually incurred, establish the terms of compensation by the department of transportation to contractors.

It is the ability to make inference on costs that distinguishes this study from prior empirical work on firm boundaries. Other studies only provide qualitative evidence. That is, they show how some feature of the theory (such as incompleteness) affects the probability that work is subcontracted. There is very little quantitative evidence about the impact of firm boundaries on economic outcomes such as cost. The lack of quantitative evidence has been a major critique in this literature (Hubbard, 2008; Lafontaine and Slade, 2007; Klein, 2005). This paper offers one of the first attempts to bring the missing data, examine the cost primitives of the theory, and, stated bluntly, to show that firm boundaries matter.

The unit of analysis is at a very detailed level. The engineer's specifications list construction work items, also called tasks. On bridge projects, tasks range from heavy engineering jobs such as installing structural concrete, steel, asphalt, and drilling to ancillary tasks such as traffic striping, fencing, and landscaping. For each task, bidding contractors make a subcontracting decision and submit a bid. Individual work item bids are aggregated according to a scoring rule to the determine the project's low bidder. Despite sampling just 32 bridges, there are over 2,200 project-tasks and over 12,000 bidder-on-a-project-task observations. The sample includes the contractors that were awarded projects as well as all losing bidders.

From the blueprints, engineers calculate an exact quantity for each item. For example, the plans could require 120 cubic meters of structural concrete. Because construction plans can be revised, the quantities actually installed can differ from the engineer's original quantities. I propose a proxy measure of contractual incompleteness based on quantity changes. Specifically, I use the difference between the work item quantity in the original blueprints and the quantity actually installed after blueprint revisions. Those tasks that experience little or no change were likely perceived by contractors to have a low degree of incompleteness while those with large changes a high degree of incompleteness. Forward looking, rational contractors will assess the likelihood that changes will occur and incorporate anticipated hold-up costs into their bids. This is an exogenous measure because contractors have little ability—both ex-ante during design and ex-post during construction, to influence installed quantities. It is typically non-contractible contingencies, such as unknown conditions at the work site, that necessitate revisions.

I adopt a structural approach. I model contractors' profit maximizing subcontracting and bidding decisions. The econometric specification collapses to an OLS regression of *cost* on the incompleteness proxy and a set of cost controls for both subcontracted and primecontracted (performed in-house by the primary contractor) tasks. But, I do not directly observe cost. Instead, I observe bids. I use the first order conditions to optimal bidding to structurally estimate costs from bids. I do not explicitly model the mechanics of the ex-post bargaining process; nor do I model any other ex-ante actions taken in anticipation of bargaining. Instead, I interpret the results in light of the relevant aspects of the theories as they pertain to construction contracting.

I need to model bidding because strategies, like the incompleteness proxy, also depend on quantity changes. These are scoring auctions. Quantities from the original blueprints are a basis for determining bids. Contractors submit a unit price bid expressed as dollars per unit of quantity. The total bid is calculated by multiplying unit price bids with original quantities, then summing those values across tasks. For each task, the winner is paid its unit price bid times the quantity actually installed. Differences in original and final quantities induce strategic bidding behavior. The basic intuition described by Athey and Levin (2001) shows that bidders skew unit price bids above cost on tasks expected to overrun on quantity and below cost on tasks expected to underrun. By skewing, a bidder earns a higher profit without affecting its total bid and hence probability of winning the contract. Bid skewing is risky. If the overrun (or, for that matter underrun) does not occur, the winning bidder suffers a lose. The principles for allocating unit price bids are analogous to concepts from modern portfolio theory (Markowitz, 1952). The specific modeling choices match industry practitioners' intuition about unit price bidding. They credit these ideas to Gates (1959). I apply results from the scoring auction literature to formally characterize equilibrium bidding behavior (Asker and Cantillon, 2008), and draw upon ideas from finance to develop a tractable estimation procedure.

Finally I add flexibility to provide a richer set of results. I make a distinction between heavy construction tasks and ancillary tasks. I also distinguish between quantity overruns and underruns. As a preview of one of the main results, I find that for heavy construction tasks, overruns—which likely generate scheduling pressures—substantially increase costs for subcontracting. Compared to a base case with no change, the largest overruns raise subcontracting costs 12 percent. In comparison, overruns have a negligible effect on the cost of in-house performance.

Note that this result does not imply those contractors choosing subcontracting are acting irrationally. There are many other factors besides contractual incompleteness affecting subcontracting decisions. To account for the endogeneity of subcontracting decisions, I use a fixed effect method which exploits the unique panel data structure. I see multiple observations of very similar transactions. For example, within a project there are multiple bidders, and across projects, the same set of construction tasks.

1.1.1 Contribution to Existing Literature

This study joins an emerging literature that quantifies the impact of firm boundaries on economic outcomes. Recent examples include Gil (2008) (*industry*: cinemas and *performance outcome*: movie run length), Ciliberto (2006) (hospitals and capital investments), Forbes and Lederman (2008) (airlines and flight delays), and Novak and Stern (2008) (automobiles and consumer quality ratings). Others—Baker and Hubbard (2004) (trucking and fuel economy) and Levin and Tadelis (2008) (municipal services and city expenditures)—consider but do not focus on performance outcomes.

The structural approach offers several advantages. First, the structural model shows that I am making inference on the most relevant economic outcomes: bids and costs. The intermediate performance outcomes in the other studies do not show how firm boundaries affect the overall performance of a firm or industry. Nor do their performance outcomes significantly affect firm boundaries. Second, the auction format provides a crisp separation between ex-ante contract formation and ex-post contract execution. This allows me to capture the combined costs of both ex-ante distortions and ex-post haggling. As such, I perform an all encompassing test of the incomplete contracting theories of the firm.[4] Third, it is possible to conduct counterfactual exercises. As part of a cost-benefit analysis, one could quantify Caltrans' costs savings if it were to reduce contractual incompleteness. I could alter regulations that distort subcontracting patterns.[5]

This study of contractual incompleteness falls into the class of empirical work that considers uncertainty and complexity of a transaction. Seminal contributions include Monteverde and Teece (1982), Masten (1984), and Masten, Meehan Jr, and Snyder (1991). They find a higher degree of complexity is associated with a lower probability of subcontracting. Recent work by Gil (2007) (movies), Acemoglu, Aghion, Griffith, and Zilibotti (2009) (R&D intensity), Forbes and Lederman (2009) (airlines) Levin and Tadelis (2008) (municipal services) obtain the same result. The literature on forward integration into retailing finds mixed evidence.[6] Williamson (1985) and, in particular, his earlier work Williamson (1975), identified uncertainty as one of the key determinants of firm boundaries for empirical researchers to take to the data.[7] This factor is losing favor, in part, because of the difficulty of measuring uncertainty. Typically, studies rely on survey data of industry practitioners or measures of market volatility. These variables are notoriously measured with error.

[4]There has been a call in the literature to distinguish between the transactions cost and property rights theories (Whinston, 2003). I do not attempt to do so in the this paper.

[5]These counterfactuals are outside the scope of this paper.

[6]See Lafontaine and Slade (2007) for a survey of these results.

[7]Holmstrom and Roberts (1998) provide this textual interpretation. The other two are specificity and frequency of a transaction.

Measurement error creates severe attenuation bias that leads to statistically insignificant estimates.[8] A strength of this study is that the proxy captures a precise notion of incompleteness: changes in the construction contract measured at the detailed level of a work item transaction.

This work is related to the empirical auctions literature. There is an especially large body of work on highway procurement auctions including the contributions of Porter and Zona (1993), Hong and Shum (2002), Krasnokutskaya (2009), Jofre-Bonet and Pesendorfer (2003), Bajari, Houghton, and Tadelis (2007), Marion (2008), De Silva et al. (2008), Bajari and Lewis (2008), and Krasnokutskaya (2009). This is the first study to use work items as the unit of observation within the context of a structural auction model. [9]

Previous empirical work on bid skewing (Athey and Levin, 2001; Bajari et al., 2007) restricted attention to just one dimension of skewing. I model the bid skewing decision as a portfolio choice problem. That is, the correlation structure in quantity change risk across all tasks determines the optimal skew on any given task. The proposed empirical technique recovers the correlation structure of risk. Risk aversion in auctions has attracted attention in both the empirical auctions and experimental economics literature.[10] Adapting the method could provide field evidence on Arrow-Pratt risk aversion coefficients with millions of dollars stake.

The paper is organized as follows; section 2 describes the procurement process and follows with a discussion of the contractual incompleteness measure in section 3. Section 4 presents the model and characterizes optimal bidding and subcontracting decisions. Section 5 describes the data; section 6, the estimation procedure. Section 7 presents results. Section 8 discusses robustness; section 9 concludes.

---

[8] Klein (2005) notes the problem of measurement error in these types of studies.

[9] Marion (2008) is closely related to this paper on two dimensions. He studies a California affirmative action law regarding highway subcontracting. He uses work item observations but does not account for bid skewing.

[10] See Campo, Guerre, Perrigne, and Vuong (2003), and Harrison and Rutström (2008) for discussions.

1.2 Procurement Process

Caltrans procures projects using a method called *Design-Bid-Build*. It is the most common method used by departments of transportation. The name refers to the three stages of the process. Below I outline the timing, introduce notation, and describe the auction mechanism.

Design Caltrans' resident engineers write blueprints and specifications. They include T tasks indexed $t = 1, \ldots, T$. Examples include concrete, steel, excavation, and traffic striping. Caltrans employs a team of specialists called *quantity surveyors*. They examine the plans and calculate quantities for each task q_t^e (in vector notation $\mathbf{q}^e = [q_1^e, q_2^e, \ldots, q_T^e]$).[11] Primary contractors and subcontractors inspect blueprints and the job-site to forecast costs. They do not participate in the design process.

Sign Subcontracts At some time up until the bid letting date, bidding primary contractors, indexed $i = 1, \ldots, N$, choose subcontracting arrangements. Let c_{it}^s denote primary contractor i's unit cost to subcontract task t, and c_{it}^p denote its unit cost to perform the task in-house.

In most cases there are many potential subcontractors so the selection process is quite competitive. Matching of subcontractors with primary contractors is not exclusive. A subcontractor can sign with multiple prime contractors. By law, California's *Subletting and Subcontracting Fair Practice Act*,[12] stipulates that subcontract agreements must be signed before bidding. After a project has been awarded, the winning prime contractor cannot legally hire additional subcontractors, nor can it make substitutions. This is the time when contractors make ex-ante

[11] For convenience many variables are expressed as vectors with length T. These variables will unambiguously be denoted in boldface.

[12] California was the first state to adopt this law, 10 other states have followed. It is intended to prevent the practice of post award bid shopping; it prevents winning prime contractors from exercising monopsony power to garner price concessions from subcontractors. The law binds with the exception of cases involving the solvency of a contractor, breaches of subcontracting agreements, or if it is later deemed a contractor does not have the capabilities to perform a particular task. For transparency, the law requires bidding contractors to furnish to the public procuring agency the names, addresses, and lists of work to be performed for all subcontractors with its bid submittal. The public information disclosure is one reason why Caltrans contracts were chosen for this study.

relationship specific investments. Costs incurred and profits earned by the subcontractor pass through as costs born by the prime contractor.[13]

Bid Each bidding prime contractor submits sealed unit price bids for every task. Let b_{it} denote bidder i's unit price bid on task t (in vector notation $\mathbf{b}_i$). The total bid is calculated by multiplying a unit price bid with the corresponding quantity from the engineer's original plans then summing those values across all tasks; it is the dot product of the unit price bid vector and the vector of engineer's quantities. Call the total bid the score, $s_i = \mathbf{b}_i \cdot \mathbf{q}^e$, because this auction belongs to the class of scoring auctions. The bidder with the lowest score is awarded the entire project.

Build This is the phase when the project is constructed. During this phase contingencies can arise which will require revisions of the plans. Many of these revisions necessitate adjustments to quantities. Denote the quantity actually installed after revisions for task t as q_t^a (in vector notation $\mathbf{q}^a$). This is the ex-post phase when bargaining occurs.

Payment to the winning bidder is based on quantities actually installed: it is the dot product of the vector of unit price bids and the vector of installed quantities, $\mathbf{b}_i \cdot \mathbf{q}^a$. Losing bidders do not receive payment. It is important to note that contractors do not choose actual quantities. From their perspective, actual quantities are a random variable. Quantity surveyors re-examine the revised plans to authorize the actual quantities.

1.3 Measure of Contractual Incompleteness

Before continuing with the description of the structural model, it is useful to motivate the proxy measure of contractual incompleteness. I use the difference in original quantities, q_t^e, and quantities installed after revisions, q_t^a. There

[13]To see how costs pass through, drop subscripts and decompose subcontracting costs, c^s, into three parts: $c^s = c_{sub} + \pi_{sub} + c_{prime}$. The term, c_{sub}, is the direct cost incurred by the subcontractor; π_{sub} is the subcontractor's market rate of profit (industry professionals document very low profit margins on the order of 2 or 3 percent); c_{prime} is a cost incurred directly by the prime contractor. This last component could represent lost surplus from a hold-up situation. Prime contractors, not Caltrans, disburse payments to subcontractors. Thus, the combined cost and profits of the subcontractor, pass through to the prime contractor via the established price, or bid, b_{sub}, in the subcontract: $b_{sub} = c_{sub} + \pi_{sub}$.

Figure 1.1: Example Modification: 35W Bridge, Minneapolis

The arrow points to the unmarked drainage system. This unknown site condition led to drastic changes in the construction blueprints and increases in quantities for drilling, concrete, and steel reinforcement.

are two basic points needed to understand why quantity changes measure contractual incompleteness. First, I take as given that changes in construction plans during the build phase occur because of contractual incompleteness. This point is well established in the procurement literature.[14] Second, quantity changes are solely and unambiguously caused by changes in the construction plans.

As an example, consider the most drastic change from the construction of the replacement 35W bridge in Minneapolis. The original structure collapsed August 1st, 2007. The right pane in figure 1.1 shows a photo of the completed bridge, and the left pane shows a photo from the construction of the piers. The piers support the far side of the bridge. The engineering team failed to notice a drainage system located at the base of the piers (indicated by the arrow in construction photo). It was unmarked on sewage diagrams and hidden from sight by the wreckage of the collapsed bridge. Upon discovery, the drilling sites and piers had to be moved 20 feet. This change altered the design of substructures, and the superstructure had to be redesigned to bear the load of a longer span. Quantities changed. The revised plans required more drilling shafts and deeper piles. They also required more concrete and steel reinforcements in the superstructure.

There are a few points to emphasize. First, changes occur for external

[14] see (Bajari and Tadelis, 2001; Bajari et al., 2007, 2009)

reasons uncontrollable by contractors and Caltrans' engineers. Both parties have limited ability to influence installed quantities during the ex-post construction phase.[15] The 35W example highlights the most common cause of changes: differing site conditions. Changes also occur because of deficiencies in design that compromise structural integrity, unknown local regulatory codes, preferences changes of the buyer with regard to function or aesthetics, or because of errors, ambiguities or omissions in the design. Furthermore, contractors cannot take actions ex-ante to influence quantities because they do not participate in the design phase.

The second point regards the actions of Caltrans' engineers and quantity surveyors. The original quantities are not a guess, or expectation, of what actual quantities will turn out to be. Formally $\mathbf{q}^{\mathbf{e}} \neq E[\mathbf{q}^a]$. Engineers write plans, and, if need be, revise plans. Quantity surveyors calculate quantities from the original plans, and if need be, calculate quantities from revised plans. Quantity surveyors are trained to be exact; human error on the part of quantity surveyors does not cause quantity changes. The final point to note is that the measure is a proxy because it uses ex-post data on actual quantities. This measure captures what contractors likely perceived about incompleteness during their pre-bid inspection of the plans and job-site. It could be the case that contractors forecasted a high degree of incompleteness and by chance there happened to be no change. I address the implications of measurement error in the econometrics and results sections.

Previewing the data illustrates the richness of the contractual incompleteness proxy. The sample includes 32 projects (indexed, c) and 2201 project-tasks (indexed ct) used for analysis.[16] Table 1.1 depicts a histogram of the contractual incompleteness measure in tabular form. The categories are separated into overruns, where actual quantities are larger than original quantities

[15]Caltrans is an experienced buyer; therefore, contractors cannot exploit informational advantages to manipulate actual quantities with the intention of extracting extra profits. On the job-site, Caltrans' quantity surveyors monitor contractors and measure quantities. Contractors are only paid for the quantities that quantity surveyors authorize. In private construction, contractors are able to manipulate quantities because of information asymmetries between inexperienced buyers and contractors.

[16]I use the letter "c" to index projects because it common in the industry to call them contracts. To avoid confusion, I use the word "project" because my notion of a "contract" is at the disaggregated level of a project-task. Table 1.1 represents a truncated sample of tasks. For reasons to be discussed later, some project-tasks are excluded.

Table 1.1: Quantity Overruns and Underruns

		$\frac{q_t^a - q_t^e}{q_t^e}$	Frequency
	Perfect Design	0	35.4%
	Small	(0,.1]	11.0%
Overruns	Medium	(.1,.35]	7.2%
	Large	>.35	10.0%
	Small	[-.1,0)	13.6%
Underruns	Medium	[-35,-.1)	8.4%
	Large	[-1,-.35)	14.3%
Projects			32
Project-task	Observations		2201

and underruns, where actual quantities are less than original quantities. The size of changes are divided into 4 categories: large, medium and small changes, and a category labeled, "perfect design" that corresponds to no change. First, notice that this is a continuous measure and not a crude indicator for whether or not a task experienced a change. Also notice the large degree of variation. Contractors and Caltrans consider changes of 25% to be quite large. It is remarkable that a fourth of the tasks experience changes exceeding 35%. Finally, notice the separation of overruns and underruns. This distinction allows me to distinguish between types of hold-up problems. For example, overruns can trigger hold-up problems related to scheduling and accelerating the pace of work. I elaborate in the results section.

1.4 Model of Subcontracting and Competitive Bidding

In this section, I model and characterize prime contractors' bidding and subcontracting decisions. If unit cost were directly observed, I could proceed without further discussion of the auction. But there is an important strategic reason for unit price bids to differ from unit cost. Like the incompleteness proxy, it depends on quantity changes.

Table 1.2: Bid Skewing Example

	task	bid	q^e	q^a	Score	Revenue	Cost	Profit
Bid at Cost	Mobilization	**600**	1	1	1400	1800	1800	0
	Concrete(m^3)	**2**	400	600				
Skewing Gains	Mobilization	200	1	1	1400	2000	1800	200
	Concrete(m^3)	3	400	600				
Skewing Loss	Mobilization	200	1	1	1400	500	800	-300
	Concrete(m^3)	3	400	100				

1.4.1 Bid Skewing Strategies: Example

I present the intuition behind bidding strategies with an example that construction professionals credit to Gates (1959). His insight motivates the remaining modeling choices.

Consider a project with 2 tasks: 1 unit of Mobilization (setting up shop at the work-site) and the installation of 400 cubic meters of concrete. The above table depicts two possible bidding strategies and two realizations of final quantities. In one situation, concrete quantities overrun by 200 cubic meters, in the other they underrun by 300; the quantity on Mobilization is fixed at 1 in both cases. In the baseline example labeled “Bid at Cost” the bidder submits unit price bids equal to unit costs. When concrete quantities overrun, revenue increases above the score, but profits are zero. In case of an underrun (not depicted in table), revenue decreases, but profits remain zero. Suppose the bidder forecasts concrete quantities to overrun. In the row labeled “Skewing Gains” the bidder skews the unit price bid on concrete above unit cost, and below cost on Mobilization. Revenue and profit increase as compared to a strategy of bidding at cost, yet, the bidder maintains the same score, thereby not diminishing its chances of winning the auction. Such a strategy comes with risk as the row labeled “Skewing Loss” illustrates. If the concrete quantity underruns, the bidder suffers a loss of profits. Gates (1959) intuition is about the risk-reward tradeoff of skewing bids away from cost.

1.4.2 Continued Auction Description

I model the auction in the private values paradigm. Actual quantities are drawn from the joint density $g(\mathbf{q}^a)$. Realization of the draw occurs after bidding. Ex-ante bidders have symmetric information about this density.[17]

Bidder i's costs are drawn from the joint density, $f_i(\mathbf{c}_i^s, \mathbf{c}_i^p)$. Densities may differ across bidders. A bidder knows it own realization from the cost draw, but only knows the distributions its rivals are drawing from. Note that this realistically captures an implicit assumption that bidders do not know the subcontracting choices of their opponents.

Following the intuition of Gates' example, bidders are modeled to exhibit risk aversion over profits.[18] Let bidders have identical utility over profits, represented by a twice continuously differentiable, increasing, weakly concave utility function $u(\cdot)$.

The expected utility of a bidder if it wins the project is $E_{\mathbf{q}^a}\left[u\left(\pi(\mathbf{b}_i, \mathbf{c}_i; \mathbf{q}^a)\right)\right]$ where the expectation is integrated across possible realizations of actual quantities. This value is normalized to zero for losing bidders. Bidder i's payoff is its expected utility if it wins the auction times the probability it wins the auction.

$$E_{\mathbf{q}^a}\left[u\left(\pi(\mathbf{b}_i, \mathbf{c}_i, \mathbf{q}^a)\right)\right] \times Pr(s_i < s_j \quad \forall j \neq i)$$

1.4.3 Characterization of Bidding

The equilibrium concept for this auction is a Bayesian Nash equilibrium. Bidders select unit price bids and subcontracting arrangements that are best

[17]Athey and Levin (2001) analyze timber auctions where a similar phenomenon occurs. The amount of timber species bid on and harvested differs. A model with affiliated signals about actual quantities is appropriate in their setting because bidders sample only small portions of a forest. In this application all bidders inspect exactly the same plans and job site so it is not clear why they would receive different signals about installed quantities. Furthermore modeling a general affiliated values auction makes formal analysis intractable with more than two tasks.

[18]That contractors are risk averse is well justified. First, industry practitioners, such as Gates, acknowledge their risk aversion. Second, contractors are highly leveraged, have small profit margins, and operate at a small scale. Risky bidding behavior could cause bankruptcy. To be somewhat tongue and cheek, contractors, unlike large, risk neutral insurers (e.g. AIG) are not too big to fail. Soon, I will present evidence to dismiss bid rejections as an alternative that could mimic the role of risk aversion in bid skewing.

responses to the equilibrium distribution of opponents' bids. A best response solves the following maximization problem

$$\begin{aligned}
\max_{\mathbf{b}_i,\mathbf{c}_i} \quad & E_{\mathbf{q}^a}\left[u\left(\pi(\mathbf{b}_i,\mathbf{c}_i,\mathbf{q}^a)\right)\right] \times Pr(s_i < \hat{s}_j \quad \forall j \neq i) \\
s.t. & \mathbf{b}_i \in \mathbf{R}_+^T \\
& c_{it} \in \{c_{it}^s, c_{it}^p\}
\end{aligned}$$

where $\hat{s}_j$ is the equilibrium score for some type of opponent j. It is a strictly dominating strategy for contractors to choose the lowest cost subcontracting arrangement. Following the general result on scoring auctions in Asker and Cantillon (2008), bidding behavior is characterized by separating the problem into two parts. In the first part, bidders choose an optimal score, $\hat{s}_i$, that is a best response to their rivals' scores. In the second part, bidders allocate unit price bids subject to the constraint that they sum to the chosen score, $\mathbf{b}_i \cdot \mathbf{q}^e = \hat{s}_i$.

Consider the second stage problem of allocating unit price bids. There is a special type of task: lump sum tasks. Actual quantity do not differ from the engineer's quantity for lump sum tasks. Stated precisely, this implies that $Pr(q_t^a = q_t^e) = 1$. Mobilization is an example. Designate the first $t = 1, \ldots, L$ tasks as lump sum, and normalize their unit of measurement to one, so that $q_t^a = q_t^e = 1$. The remaining $t = L+1, \ldots, T$ tasks are called variable quantity tasks. The unit price bid allocation problem is,

$$\max_{\mathbf{b}_i} \quad E_{\mathbf{q}^a}\left[u\left(\sum_{t=1}^{L}(b_{it} - c_{it}) + \sum_{t=L+1}^{T}(b_{it} - c_{it})\, q_t^a\right)\right] \tag{1.1}$$

$$s.t. \qquad \sum_{t=1}^{L} b_{it} + \sum_{t=L+1}^{T} b_{it} q_t^e = \hat{s}_i \tag{1.2}$$

For emphasis, lump sum and variable quantity tasks are separated.

This problem mimics the standard portfolio choice problem in finance where a risk averse investor chooses how to allocate wealth amongst risk-free and risky assets. The engineer's quantity, q_t^e, corresponds to the price of an asset today; the actual quantity, q_t^a, tomorrow's price; the unit price bid, b_{it}, the number of shares that an investor purchases; the score, $\hat{s}_i$, wealth. Because actual quantities do not vary from the engineer's quantity for lump sum tasks and do vary for variable quantity task, the two classes of tasks correspond to risk-free and risky assets.

Two immediate observations are available by inspecting the maximization problem. First, a bidder earns risk free profits by submitting a bid at cost on the variable quantity tasks, and allocating the remainder of the bid to the lump sum tasks. The markup appears on the bids for lump sum tasks. Second, the allocation of bids across lump sum tasks is indeterminant. This is analogous to the notion in finance that owning a variety of issues of risk free assets is redundant. The importance of this remark is that observation of a submitted bid on any individual lump sum task is uninformative of cost unless other bidding motivations are taken into consideration.[19]

To further characterize bidding behavior take the first order conditions for an interior solution:

$$E\left[u'\left(\pi(\mathbf{b}_i)\right)\right] = \lambda \qquad \forall t = 1, \ldots, L \tag{1.3}$$

$$E\left[u'\left(\pi(\mathbf{b}_i)\right)\right] q_t^e = E\left[u'\left(\pi(\mathbf{b}_i)\right) q_t^a\right] \qquad \forall t = L+1, \ldots, T \tag{1.4}$$

where $\pi(\mathbf{b}_i) = (\mathbf{b}_i - \mathbf{c}_i) \cdot \mathbf{q}^a$ is profit as a function of the bid vector and λ is the Lagrange multiplier on the constraint. To obtain an analytical closed form solution, I parameterize the utility function with a constant absolute risk aversion (CARA) representation: $u(x) = \frac{-1}{\gamma} e^{-\gamma x} + k$ where k is a constant that normalizes a losing bidder's payoff to zero. I perform a first order Taylor series expansion around bids at cost, ($\mathbf{b}_i = \mathbf{c}_i$ for the variable quantity task) to obtain a linearized unit price bidding equation that maps unit costs into unit price bids. See appendix for derivation.

[19] Caltrans may reject a bid if it is deemed irregular. A bid of zero is grounds for rejection. The data indicate Caltrans only enforces this clause in the absolute most extreme cases; I have only observed rejection because of an irregular bid when a large lump sum task has a bid of zero. Bids on variable quantity tasks are highly skewed (see data section) without rejection. Some lump sum tasks have maximum bid limits. For example, the bid for Mobilization is often capped at 10% of the total bid to prevent "front end loading": over bidding on work paid near the beginning of the project.

If there is only one variable quantity task, the unit price bidding equation is,

$$b_{it} = c_{it} + \frac{1}{\gamma}\frac{E[q_t^a] - q_t^e}{E[(q_t^a - q_t^e)^2]} \tag{1.5}$$

A unit price bid is a linear function of unit cost and a skewing term that depends on the difference in a bidder's expectation of the actual quantity and the engineer's original quantity. The comparative statics match the intuition presented in Gate's (1959) example. If a bidder expects quantity to overrun, it skews its bid above cost, and skews below cost if an underrun is expected. Skewing aggressiveness increases for larger quantity differences. Risk considerations temper skewing aggressiveness. The term in the denominator, γ, is the Arrow-Pratt measure of absolute risk aversion; a more risk averse bidder skews less aggressively. The variance-like expression $E[(q_t^a - q_t^e)^2]$ is the riskiness of a quantity change; bidders skew less aggressively if quantity changes are expected to be more volatile. If a bidder forecasts no change and expects volatility, it would submit a bid equal to cost.

With more than one variable quantity task, the unit price bidding equation, stacked by task, is

$$\mathbf{b}_i = \mathbf{c}_i + \frac{1}{\gamma}E\left[(\mathbf{q}^a - \mathbf{q}^e)(\mathbf{q}^a - \mathbf{q}^e)'\right]^{-1}(E[\mathbf{q}^a] - \mathbf{q}^e) \tag{1.6}$$

Boldface vectors have length equal to the number of variable quantity tasks. Unit price bids are a linear function of unit cost and a skewing term that involves an expression resembling the inverse of a covariance matrix of quantity overruns. Not only does the expected overrun on a quantity affect the direction and aggressiveness of a skew, skewing also depends on the covariances and expected overruns of other tasks. This is analogous to the idea in finance that an optimal portfolio depends on the correlation of returns across all assets. Like modern investors, contractors use sophisticated computer algorithms that take into account the correlation structure of risk.[20] The econometric procedure places restrictions on the covariance-like matrix $E\left[(\mathbf{q}^a - \mathbf{q}^e)(\mathbf{q}^a - \mathbf{q}^e)'\right]$ to recover unit costs from unit price bids.

An especially convenient feature of this derivation is that unit price bids

[20] Cattell et al. (2007) reviews these methods.

do not depend on the bidding behavior of rivals. Intuition of first price auctions suggests a unit price bid should include a profit markup that depends on the competitiveness of the auction. Strategic interaction only matters for the choice of the score and consequently, the lump sum component of the bid. Interest centers on variable quantity tasks because I cannot measure incompleteness for lump sum task. Therefore it is sufficient to only consider this bidding equation.[21] Characterization (unreported) of the first stage problem, choosing a score, follows the standard derivation for a first-price auction with bidder types defined by a pseudo-type, $\mathbf{c}_i \cdot \mathbf{q}^e$.

As a caveat, notice that unit costs are not modeled to vary with quantity: $c(q) = c$. This assumption admits a tractable characterization. Within a project, differences in quantities are small enough that unit costs are locally well approximated by a constant unit cost. But, for large differences in quantities, there are economies of scale. Across projects, quantities vary by degrees of magnitude. In the econometric specification, I will account for scale economies using variation across projects.

1.5 Data and Descriptive Evidence

The data were collected from public records of construction projects procured by the California Department of Transportation. The sample includes all 32 bridge projects bid on and built between 2002 and 2005 using the *design-bid-build* procurement method. Most of the variables from the modeling section are found directly in the bidding documents. They list construction tasks, engineer's quantities, unit price bids, and a complete record of subcontracted tasks for all bidders. [22] In addition, they contain identifying information for bidders (names, addresses, phone numbers), and the names of, and tasks to be

[21]This lack of dependence on strategic interaction is an artifact of the CARA utility. A bidder with decreasing absolute risk aversion skews more aggressively as wealth increases. In the model, wealth is higher for larger values of $\hat{s}_i - \mathbf{c}_i \cdot \mathbf{q}^e$. This can be demonstrated by following the same derivation with constant relative risk aversion—a representation with decreasing absolute risk aversion. In unreported work I perform the tests proposed in Athey and Levin (2001) and reject the hypothesis that bidders behave in accordance with decreasing absolute risk aversion. I conclude a CARA representation is appropriate.

[22]To comply with the *Subletting and Subcontracting Fair Practice Act*, contractors name all subcontractors that perform more than 1/2% of the work.

performed by, each subcontractor hired by a prime contractor. Subcontracting information is available for all bidders, even those that lose the auction. In total there are 17,018 observations of prime contractors' "make or buy" decisions for individual work items.

The sample includes 331 contractors. Of these, 74 participate as a prime contractor. In total they submit 178 bids on the 32 projects. On average, a project receives 5.6 bids with a minimum of 2 and maximum of 13. 274 contractors participate as subcontractors. 17 contractors participate at least once as a prime contractor and at least once as a subcontractor. The average project-bidder subcontracts 37% of project work by value.

These are not monumental bridges. Engineers estimate a "fair and reasonable" cost. By this measure, project range in size from $700,000 to $22,700,000 with an average of $7,000,000. For comparison, the replacement 35W bridge is valued at $234 million and the new east span of the San Francisco Oakland Bay Bridge into the billions of dollars. From visual inspection of Google Earth satellite images[23] I see the sample includes interstate overpasses, ramps and exchanges, and state highway bridges spanning rivers, creeks, washes, and hillside ravines.[24] For the average project there are 94 tasks; by comparison, Caltrans highway paving projects average 33 tasks.[25] Bridges require more tasks because a free standing structure is built in addition to some paving work that all bridge projects require.

1.5.1 Construction Tasks

Across the 32 projects, there are 2,511 variable quantity and 482 lump sum project-tasks. Caltrans classifies tasks using a coding system similar to NAIC industrial classifications. The sample includes 982 distinct tasks (indexed t).

[23] A bridge's precise location was found by cross referencing location data in the bidding documents with information from the Federal Highway Administration's National Bridge Inventory.

[24] Sample inclusion requires a project to use the task labeled "Structural Concrete, Bridge." All bridges, even steel bridges, require concrete. Many projects let during this time period were sound walls where a bridge member is built for a section of masonry wall. I only included bridges designed for vehicular traffic. Also excluded are very large interstate construction projects with project values on the magnitude of hundreds of millions of dollars into the billions.

[25] Statistic from Bajari et al. (2007).

On average a task, t, is used on 3 of the 32 projects. Table 1.3 lists 26, Caltrans defined, categories of construction tasks at an aggregated level of grouping (hereafter referred to as industries and indexed τ). The table also lists the dollar-valued amount of work performed in each industry and the fraction of work subcontracted. Industry literature and primary sources list the types of tasks that are considered ancillary and heavy construction tasks.[26] Taking this into account, I separate the industries into the two groups. I separate the industries because prime contractors are heavy civil engineering construction firms, not ancillary task firms. This assertion is evident in the data. A contractor only sells its services on the subcontract market if it has capabilities in that task.[27] Prime contractors have a 6% share of subcontracts for heavy construction tasks, and there are only a few exceptions where a prime contractor serves as a subcontractor on ancillary tasks.[28] That an industry, such as “Reinforcement” (installing rebar), is always subcontracted does not indicate prime contractors lack capabilities. Historic iron worker union rules prohibit employment by a firm that hires any non-union workers on a particular project. Thus a prime contractor must subcontract reinforcement. There are cases where a prime contractor subcontracts reinforcement, yet serves as a subcontractor for a rival bidder.

Caltrans provides blue book prices for all standardized tasks in its *California Cost Data Book*. This booklet is published annually. Blue book prices are based on unit price bids for all projects awarded by Caltrans.

[26]Sources include scholarly articles and books from the construction management literature (Arditi and Chotibhongs, 2005; Hinze, 1993; Sweet, 2004). Primary sources include discussions with contractors on the 35W bridge project and the annual investor report of Granite Construction, the only publicly traded company in the sample. Caltrans' documents do not separate heavy and ancillary tasks.

[27]Industry sources and practitioners cite that prime contractors do achieve a minimum efficient scale to maintain divisions in ancillary tasks (Hinze, 1993; Sweet, 2004; Arditi and Chotibhongs, 2005). I define ”having capabilities” to mean a firm achieves the minimum efficient scale.

[28]Modern Alloys, a large firm that installs metal beam guard rails and concrete barriers once participated as a prime contractor. Granite Construction, the largest firm in the sample did a $300,000 subcontract for concrete sidewalks. Twice a prime contractor performed clearing and grubbing, once removed a tree, and once installed a small concrete barrier.

1.5.2 Incompleteness

Actual installed quantity data, q^a_{ct}, were collected from final payment forms, administered by Caltrans' finance department. Incompleteness, inc_{ct} is measured as deviations in installed quantities from those specified in the original plans. It has a nested structure: overruns $inc_{ct} = \left|\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right| \mathbf{1}(q^a_{ct} > q^e_{ct})$, underruns $inc_{ct} = \left|\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right| \mathbf{1}(q^a_{ct} \leq q^e_{ct})$, and the nested variant with no directional distinction $inc_{ct} = \left|\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right|$. I only use variable quantity tasks. For lump sum tasks it is not possible to quantify a task specific measure of incompleteness because, by definition, quantities cannot vary. The main point from the prior discussion about incompleteness is that there is large degree of variation (see table 1.1). Also note there is no systematic tendency for tasks to be perfectly designed.[29] There is large variation in the overall degree to which a project is completely specified. Weighted by dollar value, the most completely designed project averages quantity deviations in absolute value of just 1%; the most incomplete, 45%.

1.5.3 Prime Contractors and Subcontractors

The industry is localized for both subcontractors and prime contractors. The average prime contractor enters bids on 2.4 of the 32 projects; the average subcontractor enters subcontracting agreements on 2.5 projects. On projects for which a contractor participates, the average distance between a firm's nearest construction office and the job-site is 98 miles for prime contractors and 94 miles for subcontractors.[30] Some of the contractors operate in the California wide market. This generates skew in the size distribution of firms. I classify any bidder that enters bids on fewer than 4 projects as a "fringe" firm. Ten of the 74 prime contractors are classified as non-fringe. Studies of

[29] The concern is that tasks measured in discrete units, such as a stoplight where the original quantity is a number like 4, would be perfectly designed whereas tasks measured in continuous units, such as cubic meters of concrete, would have at least a minor change. This is not the case.

[30] Construction office locations found by cross referencing address data from the bidding documents with address information provided by companies own websites, and Google maps "find business" directory. I visually verified that the address is a construction office and not equipment yard or private residence. For contractors with multiple offices, the nearest office is paired to a project.

highway procurement commonly make this distinction.[31]

Subcontractors are more specialized than prime contractors. The average prime contractor performs work in 2 of the 26 industries; the average prime contractor, in 16. Finally, as anecdotal evidence, 149 of the 274 (54%) subcontractors' business names references a construction specialty whereas only 5 of 74 for prime contractors. [32]

1.5.4 Unit Price Bids

The median bottom line bid, s_i is 5% above the engineer's cost estimate and the standard deviation of total bids divided by engineer estimates is 0.22. The moderate bid dispersion represents heterogeneity in bidders' total costs (in the model, heterogeneity in pseudo-types $\mathbf{c}_i \cdot \mathbf{q}^e$). The median unit price bid, b_{cit}, is 22% above the blue book value and the standard deviation is very large, 3.25. The large bid dispersion cannot be fully attributed to cost heterogeneity. Bids are skewed.

1.5.5 Excluded Data

For estimation, lump sum tasks are excluded because incompleteness cannot be measured and the bid need not reflect cost by the indeterminacy result. Non-standard tasks are excluded because they do not fall into industry classifications and blue book costs are unavailable. These exclusions do not pose problems. Many lump sum tasks are administrative duties or do not involve a construction service (pollution permits, warranties, mobilization). Nonstandard tasks are customized items like those labeled "drinking fountain", "bat habitat", "San Francisco manhole". Exclusions shrink the sample size from 17,018 observations to 12,354—7,114 ancillary tasks and 5,240 heavy construction tasks.

Table 3 lists additional summary statistics. Tables 4 and 5 list summary statistics for all variables used in estimation. The bid skewing variables will

[31] See Bajari et al. (2007), Krasnokutskaya (2009).

[32] Examples of subcontractor specialty titles: Mike Brown Electric, West Coast Demolition, Pisor Fence. Prime contractor specialty titles: Lees Paving, Security Paving, Parnum Paving, Modern Alloys, Benco Bridges. Example non-specialty names: Shasta Constructors, Sterndahl Enterprises, Kiewit Pacific.

be constructed in the following section.

1.6 Estimation

The goal of estimation is to determine the effect of contractual incompleteness on unit cost for both primecontracting and subcontracting. In this section I describe the parametrization of cost functions. Then, I discuss the fixed effect identification strategy to account for self-selection bias. I also introduce a source of endogeneity for incompleteness and discuss why the fixed effects control for this source of endogeneity. Finally, I present the technique to correct bid skewing.

1.6.1 Parametrization of Cost Functions

Rather than working directly with unit costs, that, across types of construction tasks, do not share a common unit of measurement, I instead use a normalized measure of unit costs. To normalize divide unit costs by a task's blue book value. Normalization admits a natural interpretation of bids and costs as deviations from blue book value. It also controls for some of the unobserved cost heterogeneity across types of tasks.

Observations are indexed by project c, bidder i, and task t. Let c^*_{cits} and c^*_{citp} be the (normalized) unit cost of subcontracting and primecontracting and c^*_{cit} be the cost of the chosen arrangement. According to the subcontracting decision rule, prime contractors choose the lowest cost alternative. Thus,

$$sub_{cit} = \begin{cases} 1 & \text{if} \quad c^*_{cits} < c^*_{citp} \\ 0 & \text{if} \quad c^*_{cits} \geq c^*_{citp} \end{cases}$$

where the variable sub_{cit} indicates whether or not subcontracting is observed. The asterisks on primecontracting and subcontracting cost variables emphasize that neither cost is directly observed by the econometrician (they are both known to the bidder). Bids are observed.

I specify cost functions linearly:

$$c^*_{citp} = \mathbf{x}'_{citp}\beta_{0p} + \beta^{inc}_{0p} inc_{ct} + e_{citp} \quad (1.7)$$

$$c^*_{cits} = \mathbf{x}'_{cits}\beta_{0s} + \beta^{inc}_{0s} inc_{ct} + e_{cits} \quad (1.8)$$

The main variable of interest is contractual incompleteness, inc_{ct}. The other covariates are $\mathbf{x}_{cit}$; e_{cits} and e_{citp} are error terms. The parameters, β_{0s} and β_{0p} capture marginal effects on subcontracting and primecontracting costs, respectively. Both β_{0p}^{inc} and β_{0s}^{inc} are expected to be positive. The differential effect, $\beta_{0s}^{inc} - \beta_{0p}^{inc}$, is also relevant to test the theory of the firm.

Other covariates include the log of the distance between the job-site and the prime contractor's nearest construction office and an indicator variable for whether the prime contractor is a fringe firm. Distant prime contractors—infrequent participants in the local market, unfamiliar with local ordinances, and facing higher costs to transport their own equipment — are predicted to have higher primecontracting costs. For reputation reasons, distant contractors are predicted to have higher subcontracting costs. Overall, its ambiguous how distance affects the relative costs of subcontracting and primecontracting. For scale economies reasons, non-fringe contractors may have large enough logs of work to warrant maintaining divisions in a broad scope of construction activities. Thus fringe firms are predicted to have higher costs for primecontracting. This margin may be particularly relevant for ancillary tasks. Fringe status also captures reputation factors for both subcontracting and primecontracting.

Industry sources indicate project-task scale economies are an important determinant of unit cost. They arise by spreading out project-task fixed costs across more work and because of learning-by-doing. I include a normalized measure of the quantity: $\frac{q_{ct}^e}{\bar{q}_t^e}$. The denominator is the sample average engineer's original quantity for task, t.

1.6.2 Self-Selection Bias and Endogeneity of Contractual Incompleteness

For the moment, suppose unit price bids equal unit cost. There is a potential self-selection bias in OLS estimation of the separate cost equations. Inconsistency arises if an omitted variable has a different effect on subcontracting costs and primecontracting costs. There is also the potential for inconsistency due to the endogeneity of contractual incompleteness. The identification strategy exploits the panel data structure to control for both self-selection bias and incompleteness endogeneity.

The specifications include individual effects to capture industry, contract, and bidder specific characteristics. Expand the composite error term:[33]

$$
\begin{aligned}
e_{citp} &= \alpha_{\tau p} + \alpha_{cp} + \alpha_{ip} + \epsilon_{citp} \\
e_{cits} &= \alpha_{\tau s} + \alpha_{cs} + \alpha_{is} + \epsilon_{cits}
\end{aligned}
$$

The individual effects $\alpha_{\tau s}$ and $\alpha_{\tau p}$ capture factors specific to industries [34] for both primecontracting and subcontracting; the terms α_{cs} and α_{cp}, factors specific to the project; and α_{is} and α_{ip} bidder specific factors. Individual effects are allowed to be fixed: correlated with regressors. They can be represented by dummy variables. Define $\mathbf{z}_{cits} = [\mathbf{x}_{cit}, \alpha_{\tau s}, \alpha_{cs}, \alpha_{is}]$ and $\mathbf{z}_{citp} = [\mathbf{x}_{cit}, \alpha_{\tau p}, \alpha_{cp}, \alpha_{ip}]$. With dozens to hundreds of observations in each cluster, dummy variable estimation will not be inconsistent due to an incidental parameters problem.

The residual error terms , ϵ_{cits} and ϵ_{citp}, are assumed to capture exogenous cost shocks that are common across primecontracting and subcontracting. They represent idiosyncratic input cost shocks incurred by any subcontractor or prime contractor on a project-task (i.e. labor wages, equipment rental rates, material costs).[35] They also represents a bidder's private information about managerial and oversight costs. Formally, conditional on the fixed effect dummies, observed covariates, and subcontracting choices,

$$
\begin{aligned}
E[\epsilon_{cits} - \epsilon_{citp} | \mathbf{z}_{citp}, \mathbf{z}_{cits}, sub_{cit} = 1] &= 0 \\
E[\epsilon_{citp} - \epsilon_{cits} | \mathbf{z}_{citp}, \mathbf{z}_{cits}, sub_{cit} = 0] &= 0.
\end{aligned}
$$

The use of fixed effects for identification requires justification. In general, fixed effect specifications are attractive because they control for a lot of the unobserved heterogeneity. I will not attempt to exhaustively list all omitted

[33] Negative costs are in the support of the distribution because the specifications is in levels. The alternative, taking the logarithm of unit costs, is an unattractive specification. Doing so, without the blue book normalization, requires an individual effect for each type of task. This creates an incidental parameters problem. Moreover, when I introduce bids, the bid skewing equation 1.6 is characterized in levels, not logs.

[34] τ subscripts refer to industry clusters, distinct from the t subscripts referencing tasks. See table 1.3 for the list of industries.

[35] This generates correlation in cost shocks across observations within a project-task cluster. I cluster standard errors by project-task.

variables that could generate self-selection bias. Instead I will discuss a few that are important to the construction industry and other margins of the incomplete contracting theories of the firm besides contractual incompleteness.

Both the property rights and transaction costs theory predict asset characteristics are an important determinant of subcontracting decisions: examples include whether performance of a task is complementary to the management activities of a prime contractor and whether equipment is specific to the task. Industry fixed effects control for asset characteristics. I also consider agency (Holmstrom and Milgrom, 1991), and relational contracting (Baker et al., 2002) theories of the firm. Agency theory regards the importance of monitoring employees. For some tasks, such as drilling, the care and maintenance of equipment requires attentive monitoring. Project fixed effects also capture the importance monitoring. Monitoring is important for projects with limited construction zone accessibility.[36] Bidder fixed effects and fringe status capture the reputation status of prime contractors. Industry fixed effects also capture institutional features such as the labor union provisions for iron workers (discussed in the data section).

Endogeneity of contractual incompleteness (correlation between cost shocks and the incompleteness measure) is a concern for two reasons. The first regards the job-site environment; the second, quality and workmanship standards. Some job-sites, such as a mountainous ravine, present significant engineering challenges. Such an environment would increase the cost of construction and make it more costly to write a complete design. Projects with high unobserved quality standards are costly and modifications exacerbate hold-up problems. As such, engineers find it worthwhile to write more complete plans and specifications on high quality projects. The industry's typical example of a high quality, highly complete design, is an airport paving project. Contract fixed effects control for unobserved quality standards and the job-site environment.

In the robustness section, I consider the Heckman (1979) control function approach as an alternative identification strategy.

[36]On the 35W bridge separate contractors performed work on the north and south end of the bridge. Contractors could not monitor both ends of the bridge because the alternative river crossings were clogged with traffic. A bridge in my sample was along a mountain pass where materials and contractors came from distant towns on opposite ends of the pass.

1.6.3 Bid Skewing Correction

In this section I describe the technique to correct bid skewing.

The general bidding equation with all contract-task engineer quantities normalized to one,

$$\mathbf{b}_{ci} = \mathbf{c}_{ci} + \frac{1}{\gamma} E\left[(\mathbf{q}_c^a - \mathbf{1})(\mathbf{q}_c^a - \mathbf{1})'\right]^{-1} \left(E[\mathbf{q}_c^a] - \mathbf{1}\right)$$

is an additive function of unit cost and a skew term. For estimation, I place three restrictions on the covariance-like expression to exploit its symmetry. Rename the matrix $\mathbf{V}_c := \gamma E\left[(\mathbf{q}_c^a - \mathbf{1})(\mathbf{q}_c^a - \mathbf{1})'\right]$ with the i, j cell denoted as v_{ij}^c. The superscript c references the project.

Diagonal elements capture the variance of task overruns. The first restriction assumes tasks within an industry, τ (i.e. roadwork, bridgework, ancillary tasks) share a common variance.

Restriction 1 $v_{ii}^c := v_\tau^c$ *for all* $i \in \tau$

Off diagonal terms capture covariances in task overruns. The second restriction assumes pairwise task covariances within industries and across industries are the same no matter which task pair is considered.

Restriction 2 $v_{ij}^c := v_{\tau\tau'}^c$ *for any* $i \in \tau$ *and* $j \in \tau'$ *and* $(i \neq j)$

Risk of quantity changes varies across projects. Recall the average deviation of quantities across tasks is 45% for the riskiest project and just 1% for the least risky. The third restriction preserves the relative variances and covariances of industries and scales these values by a project's overall quantity change risk. Let the scalar η_c represent the quantity change risk for project c.

Restriction 3 $\eta_c \mathbf{V} = \mathbf{V}_c$

The following example $\mathbf{V}$ matrix illustrates the symmetric structure. There are 6 tasks and 2 industries: 3 (r)oad tasks and 3 (b)ridge tasks ($\tau \in \{r, b\}$).

$$\begin{pmatrix} v_r & v_{rr} & v_{rr} & v_{rb} & v_{rb} & v_{rb} \\ v_{rr} & v_r & v_{rr} & v_{rb} & v_{rb} & v_{rb} \\ v_{rr} & v_{rr} & v_r & v_{rb} & v_{rb} & v_{rb} \\ v_{rb} & v_{rb} & v_{rb} & v_b & v_{bb} & v_{bb} \\ v_{rb} & v_{rb} & v_{rb} & v_{bb} & v_b & v_{bb} \\ v_{rb} & v_{rb} & v_{rb} & v_{bb} & v_{bb} & v_b \end{pmatrix}$$

Placing restrictions by industries follows the finance analogy. It is like assuming stock returns within a sector share the same variance, and covariances are common amongst stocks within and across sectors. That overall quantity change risk varies across projects is analogous to the notion that market riskiness varies across countries.

In the context of construction, these restrictions are reasonable. Events can have an effect that ripple across tasks. For example, an event could increase the length of road by 10% which would require a 10% increase in roadwork quantities. For another example, consider the modification to the 35W bridge. When the piers were moved 20 feet, extra concrete, drilling and steel reinforcements were needed to bear the additional load of a longer span.

Inversion preserves symmetry. With $\mathbf{A} := \mathbf{V}^{-1}$, the unit price bidding equation becomes, $\mathbf{b}_{ci} = \mathbf{c}^*_{ci} + \frac{1}{\eta_c} AE[(\mathbf{q}^a_c - \mathbf{1})]$. Expanding around some contract-task ct in industry τ provides a map between a unit price bid and unit cost with a linear skewing correction:

$$b_{cit} = c^*_{cit} + \frac{1}{\eta_c}\left[a_\tau(E[q^a_{ct} - 1]) + \sum_{\tau'}\left(a_{\tau\tau'}\left(\sum_{ct' \in \tau' \cup c \; ct' \neq ct}(E[q^a_{ct'} - 1])\right)\right)\right] + \psi_{cit} \tag{1.9}$$

The diagonal coefficients, a_τ, and off diagonal coefficients, $a_{\tau\tau'}$ (as a collection $\mathbf{a}$) are called bid-skew parameters. They can be estimated from bid data. Actual quantities, q^a_{ct} serve as a measure for expected actual quantities, $E[q^a_{ct}]$. Rather than estimating contract riskiness terms from bid data, I use the average squared quantity changes on a project, $\frac{1}{T_c}(\mathbf{q}^a_c - \mathbf{1})'(\mathbf{q}^a_c - \mathbf{1})$, as an estimate of project riskiness, η_c. The error term, ψ_{cit}, represents expectational error and idiosyncratic variations in bid skewing. Tasks are grouped by 3 industries: roadwork, bridgework and ancillary tasks. There are a total of 9 bid-skew parameters. [37]

The specification is a system of four cost equations: on the subsamples of primecontracted and subcontracted tasks for both heavy industries and ancillary tasks. Cross equation restrictions are placed on the 9 covariates related to bid skewing.

Measurement error presents a limitation. The incompleteness proxy is measured with error and there is expectational error in the bid-skewing variables.[38] I consider the implications of measurement error while interpreting results and further discuss in the robustness section.

1.7 Results

I first present the results on contractual incompleteness. I separately discuss heavy industries and ancillary tasks. Next, I consider the other covariates (fringe status, distance, and scale economies). Then, I present the results on bid skewing and discuss the bias from misspecification.

1.7.1 Heavy Industry Contractual Incompleteness

Table 1.7 reports the estimates for the most preferred specification. Notice that I restrict the coefficients on contractual incompleteness to be the same for overruns and underruns on primecontracted work. In the unrestricted specification, the coefficients are not statistically different from one another. Later, I discuss the econometric reasons for preferring the restricted specification in light of measurement error and bid skewing.

[37] Alternatively the covariance structure could have been estimated from quantity data. This creates additional non-linear measurement error because it is multiplied with the measurement already present in the overrun term. As a robustness check I consider nested variants of the restrictions.

[38] Variables in a bidder's current information set could serve as valid instruments for expectational error. Hansen and Singleton (1982) used lagged assets returns as instruments. Using the analog of their instruments, lagged quantity overruns, does not work because overruns are not correlated across projects located in geographically distinct locations.

First consider heavy industries. (refer to the bold-face estimates in table 1.7 under the columns for heavy industries). Contractual incompleteness on overruns has a significant effect on subcontracting costs. Interpreted as a dollar-valued elasticity, a 10 percentage point increase in quantity deviation leads to a 4.5 percentage point increase in unit cost relative to the blue book value. In the range from perfect design (no measured incompleteness) to overruns of 35%, subcontracting costs increase about 12 percentage points. The range roughly corresponds to a movement from the 35 percentile to the 80 percentile of the incompleteness proxy. In contrast, contractual incompleteness from underruns has zero effect on subcontracting costs. Thus, hold-up costs exist and are economically large for subcontracting, but only in cases where contractors likely expect overruns. In contrast to subcontracting, incompleteness has no effect on primecontracting costs. This is evidence that hold-up costs are mitigated if heavy construction work is performed in-house by the primary contractor.

The distinction between overruns and underruns can indicate the type of hold-up that is occurring. A leading source of dispute regards scheduling and acceleration (working overtime and mobilizing additional workers and equipment). Quantity overruns are likely to trigger this sort of ex-post haggling. In addition, ex-ante relationship specific commitments to pre-coordinate schedules on the part of subcontractors and prime contractors could reap large benefits. The threat of overruns and renegotiating schedules could reduce the amount of coordination. I focus this interpretation on heavy construction tasks because they are the ones typically included in schedule planning; ancillary tasks, which typically occur towards the end of a project, are excluded.[39] Underruns, which reduce the amount of work, alleviate scheduling pressures. In part, this explains why underrun contractual incompleteness has no impact on subcontracting costs.

1.7.2 Ancillary Task Contractual Incompleteness

The results are quite different for ancillary tasks. First consider subcontracting (refer the bold-faced entries under ancillary tasks in table 1.7). For underruns,

[39] Contractors use a computerized system called the *critical path method* to schedule construction. As noted by Levin (1998) ancillary tasks are omitted.

contractual incompleteness has no effect on subcontracting costs. Yet for overruns there is a very large impact on subcontracting costs. In the range from perfect design to underruns of 35%, subcontracting costs increase about 23 percentage points. Scheduling hold-up does not seem to be important for ancillary tasks. Again, scheduling is not expected to be important because, as tasks performed at the end, up to two years after bidding, they are excluded from the pre-planning of schedules.

There is an explanation for the large hold-up stemming from underruns. Industry sources state that quantity underruns cause disputes about lost revenue (Sweet, 2004). Subcontractors lose business and cannot recoup project specific fixed costs. This hold-up is best interpreted in light of the reference point theory (Hart and Moore, 2008). In the norm, subcontractors expect to cover fixed costs. This is in spite of the terms of their original pricing contracts that adjust payment based on quantities. The norm serves as their threat point for holding-up prime contractors.

Next consider primecontracting of ancillary tasks. Contractual incompleteness has a large effect on primecontracting costs. In the previously discussed range, costs increase 15 percentage points. At first glance, this result seems at odds with the theory of the firm; performing work in-house typically decreases the hazards of contractual incompleteness. But, in fact, the theory accommodates this result. As touched on earlier, prime contractors are heavy civil engineering firms, not ancillary task firms.[40] They do not maintain divisions in ancillary tasks because they do not achieve an efficient scale. As such, when they elect to perform these tasks themselves, they must transact in spot markets for inputs: laborers, equipment rental yards, and material suppliers. The evidence points to severe contractual hold-up in these one-off input market transactions. Specialized ancillary task subcontractors, operating in local markets, have frequent, repeated interactions with local input markets. The repeated transacting mitigates hold-up problems according to relational contracting theories of the firm (Baker et al., 2002). The distinction is not present for heavy construction tasks because both prime contractors and subcontractors are heavy construction firms.

There is a tentative explanation for why underruns cause payment disputes

[40]Industry practitioners and sources make this point, and I provide evidence in the data section to support this claim.

for ancillary tasks but not heavy construction tasks. It regards a government program about subcontracting to Disadvantaged Business Enterprises (DBEs). DBEs are firms owned by racial minorities and women. A two-part subcontract, with a fixed and marginal pricing component would alleviate payment frictions. A subcontractor could recoup its fixed costs. For heavy construction tasks there is leniency on the part of Caltrans to allow flexible subcontracts. Ancillary tasks are more closely overseen by Caltrans because they involve DBEs. The DBE compliance form explicitly restricts subcontracts to be one-part. In the data, 30% of all subcontracts for ancillary tasks are with DBEs, just 4% for heavy construction tasks.

The typical study in the make-or-buy literature starts with cost functions similar to those in equation 1.7. With distributional assumptions on the error term and data on subcontracting decisions, a researcher can use a binary discrete choice estimator, such as probit, to infer differential marginal effects, $\beta_{sub} - \beta_{prime}$, but only up to a scale normalization. I can identify much more: the dollar valued magnitude of the difference and the separate coefficients on β_{sub} and β_{prime}.[41] Despite the drawbacks, it is still useful to consider probit estimates. See the estimates for contractual incompleteness in table 1.8. The probit results corroborate the cost based estimates, which is reassuring. The probit model does not have the complications of bid skewing or self selection and suggests the cost based estimates are not suffering from misspecification.

1.7.3 Fringe Status, Distance, Scale Economies

Contractual incompleteness is the primary focus, but the other cost controls—fringe status, distance, and scale economies—are also of interest. The coefficients on fringe status and distance generally have the expected signs (see the preferred specification in table 1.7). As could be reasonably expected, fringe and distant firms have higher primecontracting costs for all types of tasks. The exception is a tightly estimated zero effect of fringe status on ancillary task primecontracting costs. The zero result is meaningful because it further supports the claim that no prime contractors, not even the large firms,

[41]See Masten et al. (1991) for further discussion of the shortcomings of a discrete choice methodology in make-or-buy studies. Also see Lafontaine and Slade (2007) for a comprehensive survey of make-or-buy studies using discrete choice methods.

realize an efficient scale on ancillary tasks. For subcontracting, fringe firms have higher costs on both types of tasks. This suggests reputation matters. Distance has a minimal effect.

Although the cost based estimates are statistically insignificant, probit estimates for fringe status and distance reveal an interesting pattern. Distant and larger firms are more likely to subcontract heavy construction tasks. The results are exactly the opposite for ancillary tasks. Distant and larger firms are less like to subcontract. It would be interesting to focus a study on this pattern, perhaps in light of reputation factors and the decisions of distant prime contractors to participate in an auction.

Task economies of scale are significant, and explain more variation in unit price bids than any other regressor (see table 1.7). It is notable that scale economies are larger in magnitude for primecontracting. Learning-by-doing is important on the job site and could possibly explain the difference. Subcontractors experience less learning-by-doing on any given project because they typically work on many more projects in a construction season than prime contractors. The difference in scale economies suggests prime contractors close the learning gap with more work.

1.7.4 Bid Skewing

A lot of care has been put into bid skewing. First consider the raw estimates on the bid-skew parameters. The top panel of table 1.9 presents estimates of the 9 bid-skew parameters for the preferred specification. The diagonal terms are positive for all three industries. As expected, this indicates bidders skew up on overruns and down on underruns. The bottom panel reports point estimates of the variance matrix $\mathbf{V}$—the calculated inverse of the estimates for $\mathbf{A}$.

For ancillary tasks, the diagonal term in $\mathbf{A}$ is small relative to bridge and road tasks. This implies the expected variance in quantity changes are high (see corresponding diagonal term in $\mathbf{V}$), and hence, risk averse bidders do not skew aggressively on ancillary tasks. The covariances amongst ancillary tasks, and with road and bridge tasks are small in comparison to the variance. The diagonal terms in $\mathbf{A}$ are large for road and bridge tasks. This implies the expected variance is small, and hence, bidders skew more aggressively on these tasks. For a sense the magnitude, the predicted skew is 12.6% above cost for a one standard deviation overrun on a bridge task if the covariance terms are ignored.[42] The covariances amongst and across road and bridge tasks are significant, which indicates that bidders do in fact make sophisticated portfolio allocation decisions.

I need to account for bid skewing to avoid bias in the contractual incompleteness estimates. Table 1.11 illustrates the bias from misspecification. The table shows three specifications ordered from the most biased specification, with no bid-skew correction terms and no restrictions on the contractual incompleteness variable, to a moderately biased specification with the bid-skew correction and unrestricted coefficients, to the most preferred specification. The arrows indicate the direction of bias on the contractual incompleteness estimates relative to the preferred specification. First consider the differences between the *most biased* estimates and the *moderately biased* estimates. This source of bias is easy to interpret. Recall the basic intuition about bid-skewing; bidders skew up on overruns and down on underruns. Without the bid skew correction, the overrun contractual incompleteness estimate is larger because the estimated coefficient is picking up the effect of bidders skewing up, and the underrun incompleteness is smaller because it is picking up downward skewing.

The bias in the *moderately biased*, unrestricted specification relative to the *preferred*, restricted specification, again, boils down to a failure to fully account for bid skewing. The source of this failure is subtle. It is caused by contamination bias from measurement error. Classical measurement error, in a simple econometric model with a single regressor that is orthogonal to the other regressors only causes attenuation bias. This econometric model is more complicated. There is measurement error in two types of variables: the contractual incompleteness variables and the bid-skew variables.[43] Moreover, these variables are highly correlated; the bid-skew terms are positively correlated with overrun incompleteness and negatively correlated with underrun incompleteness. The correlation creates contamination bias. By placing the

[42]A one standard deviation change in the diagonal bid-skew variable for bridge tasks is 1.423 (see table 1.6).

[43]In the unrestricted specification, there are 8 contractual incompleteness variables and 9 bid skew terms.

restrictions on primecontracting incompleteness, I am using an incompleteness variable $\left|\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right|$ that is less correlated with the bid-skew terms. To see how the restrictions effect the estimated bid-skew terms, compare the diagonal elements in **A** from the preferred, restricted specification (table 1.9) to those in the unrestricted specification (table 1.10). The diagonal terms are smaller in the unrestricted specification. Essentially, the unrestricted model underestimates the degree of bid skewing which, in turn, biases the estimates on contractual incompleteness.

The magnitude of the bias is large. Across the 8 incompleteness categories, a quick calculation shows the coefficients in the worst specification are biased by at least 17% and up to a factor of 8. I calculate the bias as $\left|\frac{\beta_{pref}-\beta_{biased}}{\beta_{pref}}\right|$. For underruns on heavy industries, the bias is so great that the coefficient signs change to negative values.

1.8 Robustness

The identification strategy relies on fixed effects to control for selection. An alternative specification considers a switching regression technique that uses a control function approach (Heckman, 1979). In the Roy (1951) model framework, where selection is based solely on cost outcomes and not on non-monetary considerations, finding valid exclusion restrictions is challenging. This is especially challenging because the excluded variables should be at the detailed level of a project-task.[44] I tested a specification without exclusion restrictions. The estimates differ negligibly, and the coefficients on the control functions are statistically insignificant which is not surprising given the fixed effects already control for the key unobserved variables.

I also consider a specification that relaxes the Roy (1951) model framework where I do have a justifiable exclusion restriction. The exclusion restriction is based on two regulations: the requirement to make a "good faith effort" to hire disadvantaged business enterprise (DBE) subcontractors and subcontracting limits. The DBE requirement (stated as a percentage of total contract value) could be a plausible exogenous shifter of subcontracting probabilities that has

[44] To identify subcontracting costs the excluded variables should affect primecontracting costs and not subcontracting costs. To identify primecontracting costs the excluded variables should affect subcontracting costs and not primecontracting costs.

no effect on either subcontracting or primecontracting costs. In the data, a higher DBE goal increases the amount of ancillary tasks subcontracted, and when the cap bids, there is a reduction in the amount of heavy construction tasks subcontracted. The data show DBE's almost exclusively perform ancillary tasks (30% of all ancillary task subcontracts and 4% of heavy construction subcontracts). The DBE subcontractors' specialty in ancillary task explains the positive correlation between subcontracting probabilities and the DBE goal. The binding cap explains the negative correlation for heavy construction tasks. Estimates preserve results, but magnitudes differ. There is an important tradeoff in using this method. I must drop project fixed effects which is why I find this an undesirable specification.

Unit price bids display an extreme amount of dispersion. Beeston (1975) first documented this fact, and I find it in this data. For example, within the same project-task the highest bid is on average 7.49 times greater than the lowest. In situations when the same subcontractor is hired for a project-task—presumably costs are nearly identical—this ratio is 3.49. Other measures of bid dispersion show the same pattern. This dispersion cannot be reconciled by heterogeneity in cost or risk aversion. There are 4 possible explanations. First, pairwise task correlation could be quite high in which case those tasks are skewed as a pair rather than individually. Allocation within the pair is arbitrary. Second, if bidders do not expect volatility in quantity, the task resembles a lump sum task. Third, with an average of 94 tasks, where quantities average out to zero overrun, there is little aggregate risk to arbitrarily allocating unit price bids. Fourth, industry sources state that contractors have an incentive to conceal their true unit costs from Caltrans and rivals by randomizing bids. The econometric technique corrects for systematic bid skewing, but it cannot fully control for the full amount of bid dispersion. As a result, bid dispersion adversely impacts estimation efficiency. This fact is actually reassuring because the relatively large standard errors are primarily an artifact of bid dispersion, not an indication that the effects of interest are insignificant.

1.9 Conclusion

This paper provides empirical evidence that contractual incompleteness has a substantial impact on costs in the construction industry. As the incomplete contracting theories of the firm predict, the effect differs for primecontracting (in-house) and subcontracting (market) transactions. Unlike traditional studies of the "make or buy" decision that provide qualitative evidence about firm boundaries, this study quantifies the dollar-valued cost of those decisions. Such a research design is made possible because contractors' bids, after correcting for bid skewing strategies, reflect cost.

I developed a model of the subcontracting and competitive bidding process for Caltrans bridge projects. At the detailed level of a work item, prime contractors submit a unit price bid and decide whether to perform work themselves or hire a subcontractor. I proposed a measure of incompleteness—the difference in original and revised work item quantities. I characterized bidding behavior which shows how bidders forecast quantity changes to allocate unit price bids based on principles analogous to those in modern portfolio theory. Finally, I embed this characterization into the econometric procedure to get an unbiased estimate of how contractual incompleteness affects both subcontracting and primecontracting costs.

The results show that contractual incompleteness not only matters, but also has a big impact on firm boundaries. For heavy construction tasks, I find no effect on primecontracted work, which indicates hold-up problems are mitigated within the firm. Subcontracting involves significant hold-up costs. In the range from perfect design, that occur about a third of the time, to quantity overruns of 35% or more, again not uncommon, subcontracting costs increase 12%. The asymmetric result, that costs increase in anticipation of extra work, indicates scheduling pressures and conflict over acceleration are the source of hold-up.

For ancillary tasks, incompleteness impacts cost under both forms of organization. The effect for subcontracting is modest, except one case. In anticipation of quantity underruns, there are severe frictions likely stemming from disputes about lost revenue. In the mentioned range, cost increases 23%. The effect for primecontracting is quite large. Cost increases 15%. The theory of the firm can accommodate this seemingly contradictory result. Heavy civil engineering prime contractors transact in input markets on a one-shot basis. In contrast, specialty subcontractors mitigate input market hold-up because of frequent, repeated transacting.

This paper also provides a glimpse of evidence about relational determinants of firm boundaries by considering the fringe status of contractors and their proximity to the job-site. Reputation factors appear to differ for heavy and ancillary tasks. There is interest in further studies of relational contracting in this industry.[45].

In conclusion, I offer practical thoughts on procurement practices and motivate policy relevant research. Public agencies are typically concerned about the competitiveness of the procurement process. As such, they try to promote competition with the goal of reducing bidder markups. Markups are already small, estimated to be around 4% in highway contracts.[46] This suggests there are negligible gains available from promoting further competition. This paper suggests an alternative, and perhaps, more effective means to achieve cost savings. As part of a cost-benefit analysis, the most stark prescription would be for engineers to prepare more complete plans and for buyers to avoid making ex-post changes. As a qualification, I can only comment on the benefit of writing more complete plans, not the cost.

Using the nuanced results on overruns and underruns, and heavy construction and ancillary tasks I can offer more directed thoughts. For example, on heavy construction tasks, avoid increasing the amount of work or making changes that generate scheduling pressures. On ancillary tasks, interacting this paper's results with a more detailed analysis of the *Subcontracting and Subletting Fair Practice Act*, could uncover opportunities for additional cost savings. The act requires prime contractors to name all subcontractors in the pre-bidding phase, as opposed to the time of construction when plans could be more complete. There is also scope to explore how subcontracting limits and minority subcontracting requirements distort subcontracting decisions, and how those distortions interact with the results on contractual incompleteness.

Construction contracting is particularly important in times of economic slowdown. It is a leading tool of fiscal stimulus. Recently, there has been a

[45] An example is Gil and Marion (2009)

[46] Estimate taken from Bajari et al. (2007). Their survey of the literature finds similar markups.

lot of publicity about shovel-ready projects. Contractual incompleteness and the associated wasteful costs would be modest. But, since fiscal policy needs to act quickly, the next round of fast-tracked projects are prone to be quite incomplete. With this in mind, policy makers should think about how many bridges they can buy with their stimulus dollars.

1.10 Tables

Table 1.3: Construction Industry List

Percent of Work Subcontracted	Industry	Value of Work Performed($)
Heavy Industries		
12.0%	Concrete Structures	38,000,000
12.0%	Bituminous Seals	5,718,948
15.8%	Miscellaneous Metal	878,599
16.0%	Pipes Sewers Drainage	2,205,067
17.0%	Steel Structures	4,198,709
22.8%	Slope Protection	1,741,642
23.0%	Earthwork	16,200,000
42.7%	Shotcrete	130,040
46.3%	Asphalt Concrete	13,200,000
47.8%	Portland Cement Pavement	2,899,470
85.5%	Piling	10,700,000
87.0%	Prestressing Concrete	1,933,822
99.3%	Reinforcement	14,900,000
	Total Heavy Industry Work	112,706,297
Ancillary Industries		
13.5%	Markers And Delineators	31,296
18.1%	Traffic Control Devices	8,763,694
43.5%	Concrete Curbs And Sidewalks	1,248,134
47.7%	Clearing And Grubbing	620,460
50.3%	Waterproofing	205,623
59.3%	Existing Highway Facilities	7,460,883
67.5%	Railings And Barriers	4,957,768
67.6%	Signs	1,384,097
71.6%	Erosion Control And Planting	3,118,165
83.4%	Traffic Stripes And Markings	642,315
86.7%	Fences	303,895
98.2%	Electrical	3,445,871
98.9%	Painting	3,975,775
	Total Ancillary Work	36,157,976
	Total All Work	148,864,273

Industry size based on winning bids. Subcontracting percentage based on submitted bids. Nonstandard tasks excluded. Non-administrative lump sum tasks included. Noteworthy lump sum task: "Remove Existing Bridge" in industry "Existing Highway Facilities". Remaining tasks are the removal of signs, barriers and traffic striping.

Table 1.4: Summary Statistics

	Obs	Mean	Std. Dev.	Min	Max
Across Projects					
Engineer's Cost Estimate($)	32	6,985,134	6,414,246	701,000	22,200,000
Number of Bidders	32	5.56	2.38	2	13
Number of Line Item Tasks	32	93.53	42.45	37	221
Weighted Average Incompleteness	32	0.109	0.097	0.012	0.448
Across Industries					
Percent Contract Value Subcontracted	26	49.46%	31.61	1.77%	99.28%
Value of Industry Work($)	26	5,742,537	7,741,581	26,457	35,800,000
Number of Distinct Tasks	26	25.3	29.7	1	131
Occurrences on Projects(out of 32)	26	21.9	9.6	2	32
Across Project Bidders					
Total Bid/Engineer's Cost Estimate	178	1.06	0.22	0.62	2.10
Distance to Project(mi)	178	107.3	126.0	0.3	663.0
Percent Contract Value Subcontracted	178	37.22%	13.55	8.92%	77.38%
Across Project Tasks (Not Lump Sum)					
Absolute Quantity Change $\frac{\lvert q_{ct}^a - q_{ct}^e \rvert}{q_{ct}^e}$	2511	0.302	0.749	0	14.487
Quantity Change $\frac{q_{ct}^a - q_{ct}^e}{q_{ct}^e}$	2511	0.035	0.807	-1	14.487
Across Prime Contractors					
Total Value Primecontracts Tendered($)	74	18,400,000	20,600,000	1,083,000	93,700,000
Contract Participation(out of 32)	74	2.405	2.013	1	11
Contracts Awarded(out of 32)	74	0.432	0.684	0	3
Industry Performance(out of 26)	74	16.351	4.001	7	26
Across Subcontractors					
Total Value Subcontracts Tendered($)	274	1,856,867	4,362,414	3,072	46,400,000
Contract Participation(out of 32)	274	2.460	2.767	1	19
Prime Contractor Partners(out of 74)	274	5.062	5.587	1	37
Industry Performance(out of 26)	274	1.985	1.526	0.0	11

Weighted average incompleteness: absolute quantity change weighted by bluebook values. Value of industry work based on winning bids.

Table 1.5: Summary Statistics: Variables Used in Estimation

	Obs	Mean	Std. Dev.	Min	Max
Heavy Subcontracted					
Incompleteness					
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert$	1577	0.169	0.460	0	7.681
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	1577	0.075	0.413	0	7.681
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	1577	0.094	0.236	0	1
Fringe (indicator)	1577	0.564	0.496	0	1
Log Distance to Project	1577	4.099	1.276	-1.346	6.497
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	1577	1.095	1.035	0.0002	10.113
Normalized Unit Price Bid $\frac{b_{cit}}{c^{bb}_t}$	1577	1.901	2.362	0.001	32.648
Heavy Primecontracted					
Incompleteness					
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert$	3663	0.251	0.838	0	12.444
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	3663	0.146	0.819	0	12.444
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	3663	0.105	0.251	0	1
Fringe (indicator)	3663	0.542	0.498	0	1
Log Distance to Project	3663	3.937	1.223	-1.346	6.497
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	3663	1.083	1.225	0.0009	10.113
Normalized Unit Price Bid $\frac{b_{cit}}{c^{bb}_t}$	3663	2.016	2.682	0.001	50.036
Ancillary Subcontracted					
Incompleteness					
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert$	3597	0.379	0.754	0	10.735
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	3597	0.219	0.739	0	10.735
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	3597	0.159	0.303	0	1
Fringe (indicator)	3597	0.535	0.499	0	1
Log Distance to Project	3597	3.854	1.192	-1.346	6.497
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	3597	1.037	1.069	0.0002	10.986
Normalized Unit Price Bid $\frac{b_{cit}}{c^{bb}_t}$	3597	1.689	2.286	0.007	53.227
Ancillary Primecontracted					
Incompleteness					
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert$	3517	0.317	0.539	0	6.667
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	3517	0.159	0.497	0	6.667
$\left\lvert \frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	3517	0.158	0.305	0	1
Fringe (indicator)	3517	0.560	0.497	0	1
Log Distance to Project	3517	3.997	1.213	-1.346	6.497
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	3517	0.973	0.929	0.0002	7.423
Normalized Unit Price Bid $\frac{b_{cit}}{c^{bb}_t}$	3517	1.980	4.678	0.001	189.251

Table 1.6: Summary Statistics: Bid Skewing Correction Variables

	Obs	Mean	Std. Dev.	Min	Max
Diagonal Skew Variables $\frac{1}{\hat{\eta}_c} q^a_{ct} - 1$ in industry:					
bridge	3422	-0.024	1.423	-11.936	11.277
road	1818	-0.320	1.808	-10.980	10.353
ancillary	7114	-0.144	2.528	-15.084	40.630
Off Diagonal Skew Variables $\frac{1}{\hat{\eta}_c}\left(\sum_{ct' \in \tau' \cup c\ ct' \neq ct}(q^a_{ct'} - 1)\right)$ across industries:					
bridge-bridge	3422	6.890	60.654	-120.934	103.495
road-road	1818	-32.501	84.012	-295.461	91.951
ancillary-ancillary	7114	-16.459	120.918	-386.906	318.106
bridge-road	5240	-12.707	65.584	-295.461	103.495
bridge-ancillary	10536	-3.684	90.894	-386.906	318.106
road-ancillary	8932	-29.361	92.046	-386.906	318.106
Project Riskiness η_c					
$\frac{1}{T_c}\left(\mathbf{q}^a_c - \mathbf{1}\right)'\left(\mathbf{q}^a_c - \mathbf{1}\right)$	32	0.524	0.669	0.019	2.708

See estimation section for definition and construction of variables.

Table 1.7: Cost Estimates: Preferred Specification

	Heavy Industries		Ancillary Tasks	
	Sub	Prime	Sub	Prime
Overrun				
Incompleteness $\left\lvert\frac{q^a_{ct} - q^e_{ct}}{q^e_{ct}}\right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	**0.437**	**-0.037**	**0.074**	**0.535**
	(0.188)	(0.040)	(0.087)	(0.183)
Underrun				
Incompleteness $\left\lvert\frac{q^a_{ct} - q^e_{ct}}{q^e_{ct}}\right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	**0.044**	**-0.037**	**0.792**	**0.535**
	(0.281)	(0.040)	(0.268)	(0.183)
Other Cost Controls				
Fringe (indicator)	0.159	0.101	0.074	-0.005
	(0.117)	(0.097)	(0.086)	(0.140)
Log Distance to Project	0.034	0.087	-0.027	0.038
	(0.039)	(0.045)	(0.041)	(0.049)
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	-0.307	-0.391	-0.244	-0.369
	(0.041)	(0.032)	(0.038)	(0.058)
Industry Effect(N=13)	Fixed	Fixed	Fixed	Fixed
Project Effect(N=32)	Fixed	Fixed	Fixed	Fixed
Bidder Effect(N=74)	None	None	None	None
Contract Task Observations	5240	5240	7114	7114
Subcontracted Tasks	1577	1577	3597	3597

Regressand: Unit price bid/Blue book unit cost
Coefficients on primecontracting for both heavy industries and ancillary tasks are restricted to be the the same for underruns and overruns.
Bid-skew terms not shown

Table 1.8: Probit Estimates: Likelihood of Subcontracting

	Heavy Industries	Ancillary Tasks
Overrun		
Incompleteness $\left\lvert \frac{q^a_{ct} - q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$	**-0.087**	**0.120**
	(0.037)	(0.030)
Underrun		
Incompleteness $\left\lvert \frac{q^a_{ct} - q^e_{ct}}{q^e_{ct}} \right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$	**0.005**	**-0.064**
	(0.092)	(0.057)
Other Cost Controls		
Fringe(indicator)	0.127	-0.041
	(0.047)	(0.057)
Log Distance to Project	0.092	-0.054
	(0.020)	(0.016)
Task Scale Economies $\frac{q^e_{ct}}{\bar{q}^e_t}$	-0.036	0.021
	(0.021)	(0.018)
Industry Effect(N=13)	Fixed	Fixed
Project Effect(N=32)	Fixed	Fixed
Bidder Effect(N=74)	None	None
Contract Task Observations	5240	7114
Subcontracted Tasks	1577	3597

Regressand: subcontracting indicator

Negative coefficients imply relatively more costly to subcontract

Table 1.9: Skew Matrix Estimates: Restricted & Preferred Specification

$(\mathbf{A} = \mathbf{V}^{-1})$ **Inverse Covariance Matrix**			
	Bridge	Road	Ancillary
Bridge (diagonal)	0.0804		
	(0.0263)		
Bridge (off diagonal)	0.0196		
	(0.0043)		
Road (diagonal)		0.0754	
		(0.0445)	
Road (off diagonal)	0.0100	0.0067	
	(0.0037)	(0.0037)	
Ancillary (diagonal)			0.0398
			(0.0289)
Ancillary (off diagonal)	-0.0026	0.0004	0.0005
	(0.0041)	(0.0041)	(0.0015)
$(\mathbf{V})$ **Covariance Matrix**			
Bridge (diagonal)	13.52		
Bridge (off diagonal)	-2.93		
Road (diagonal)		13.69	
Road (off diagonal)	-1.30	-0.87	
Ancillary (diagonal)			25.23
Ancillary (off diagonal)	0.71	-0.29	-0.22

standard errors in parentheses.
Matrix $(\mathbf{A})$ estimated from bid data. Matrix $(\mathbf{V})$ inverse of matrix $(\mathbf{A})$ point estimates.
See estimation section for description of matrices and estimation procedure.
Preferred specification with coefficient restrictions on contractual incompleteness for primecontracting.

Table 1.10: Skew Matrix Estimates: Unrestricted & Biased Specification

$(\mathbf{A} = \mathbf{V}^{-1})$ **Inverse Covariance Matrix**			
	Bridge	Road	Ancillary
Bridge (diagonal)	0.0655		
	(0.0304)		
Bridge (off diagonal)	0.0196		
	(0.0043)		
Road (diagonal)		0.0600	
		(0.0427)	
Road (off diagonal)	0.0100	0.0068	
	(0.0037)	(0.0037)	
Ancillary (diagonal)			0.0248
			(0.0323)
Ancillary (off diagonal)	-0.0026	0.0004	0.0005
	(0.0041)	(0.0041)	(0.0015)
$(\mathbf{V})$ **Covariance Matrix**			
Bridge (diagonal)	17.31		
Bridge (off diagonal)	-4.48		
Road (diagonal)		17.47	
Road (off diagonal)	-1.94	-1.33	
Ancillary (diagonal)			40.64
Ancillary (off diagonal)	1.38	-0.65	-0.51

standard errors in parentheses.
Matrix (**A**) estimated from bid data. Matrix (**V**) inverse of matrix (**A**) point estimates.
See estimation section for description of matrices and estimation procedure.
Biased specification with no restrictions on contractual incompleteness for primecontracting.

Table 1.11: Cost specification comparison: Bid-skew bias

	Heavy Industries		Ancillary Tasks	
	Sub	Prime	Sub	Prime
Overrun				
Incompleteness $\left\lvert\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right\rvert \mathbf{1}(q^a_{ct} \geq q^e_{ct})$				
Most Biased	**0.556**$\uparrow\uparrow$	**0.038**$\uparrow\uparrow$	**0.115**$\uparrow\uparrow$	**0.659**$\uparrow\uparrow$
not bid-skew corrected & unrestricted	(0.196)	(0.031)	(0.076)	(0.219)
Moderate Bias	**0.461**$\uparrow$	**-0.019**$\uparrow$	**0.089**$\uparrow$	**0.619**$\uparrow$
bid-skew corrected & unrestricted	(0.190)	(0.039)	(0.090)	(0.216)
Preferred	**0.437**	**-0.037**	**0.074**	**0.535**
bid-skew corrected & restricted	(0.188)	(0.040)	(0.087)	(0.183)
Underrun				
Incompleteness $\left\lvert\frac{q^a_{ct}-q^e_{ct}}{q^e_{ct}}\right\rvert \mathbf{1}(q^a_{ct} < q^e_{ct})$				
Preferred	**0.044**	**-0.037**	**0.792**	**0.535**
bid-skew corrected & restricted	(0.281)	(0.040)	(0.268)	(0.183)
Moderate Bias	**-0.020**$\downarrow$	**-0.184**$\downarrow$	**0.738**$\downarrow$	**0.281**$\downarrow$
bid-skew corrected & unrestricted	(0.292)	(0.195)	(0.248)	(0.267)
Most Biased	**-0.113**$\downarrow\downarrow$	**-0.263**$\downarrow\downarrow$	**0.652**$\downarrow\downarrow$	**0.175**$\downarrow\downarrow$
not bid-skew corrected & unrestricted	(0.251)	(0.179)	(0.310)	(0.234)
Suppressed terms				
Fixed Effects				
Other cost controls Fringe status, Job-site distance, Task scale economies $\frac{q^e}{\bar{q}^e}$				

Regressand: Unit price bid/Blue book unit cost. This table illustrates the bias from a failure to correct for bid-skewing and the contamination bias of measurement error in the unrestricted specification. The arrows indicate the direction of bias relative to the preferred specification with restrictions and a bid skew correction. Estimates on the other cost control variables are omitted. There is no reason they would be biased by bid-skewing, and there are negligible differences in their estimates across specifications.

Chapter 2

A Dynamic Model of Housing Demand: Estimation and Policy Implications

2.1 Introduction

In this paper, we estimate and simulate a dynamic structural model of consumer demand for housing. We use this model to study how housing and non-housing demand will respond to a collapse in home prices, decrease in incomes, increases in interest rates, and a tightening of credit standards. In the early part of the recession housing was at the center of a number of important policy and macroeconomic questions, and still remains an important part of the dialogue during the recovery. Apart from the direct impact on the housing market, our model has quantitative implications for the size of the wealth effect of house price changes on nondurable consumption expenditures. The evidence can thus shed light on whether the recent decline in home prices triggered a drop in aggregate consumption demand and hence sparked the subsequent recession. Looking forward, our results can provide guidance on the possible consequences over a longer time horizon.

Over the past decade, housing has appreciated at a very fast rate compared to historical standards. Between 1997 and 2006, the nationwide Case Shiller home price index has more than doubled from 84 to 190. The rate of appreciation in certain U.S. cities was much faster than the national average. However, home prices have recently fallen and there is little evidence that there will be a sharp recovery. Between the peak in 2006 Q2 and 2009 Q2 the nationwide index fell 30%. For the most hard hit cities, such as Miami, Detroit, San Diego, Las Vegas, and Phoenix, prices have fallen between 45% and 55% as of June 2009. In addition, due to the sharp decline in economic activity and a subsequent increase in the unemployment rate household incomes have been falling significantly.

The past two decades have also seen substantial changes in mortgage markets. The traditional 30 year fixed rate mortgage was no longer the standard mortgage product. Since their introduction in the early 1980's the adjustable rate mortgage, had grown to a recent peak of a 40% share of mortgage applications in 2005.[1] In addition to adjustable rate mortgages, there has been an expansion of the subprime mortgage market and other non-conforming loans. These credit market innovations helped people with low credit quality become homeowners and also allowed households to buy larger homes. However, it is

[1]Source: Freddie Mac annual ARM (adjustable rate mortgage) survey.

now clear that the expansion of subprime credit also had a downside. Approximately 15% of subprime loans are in default as of August 2009, three times the rate in 2005.[2] As a result, non-conforming mortgages have been more difficult to obtain as lenders have tightened credit in the mortgage market.[3]

In our structural model, a household solves a life cycle dynamic programming problem. In each period, households make investment decisions in housing, choose non-housing consumption levels, and make decisions regarding mortgage borrowing and savings. Unlike a typical investment vehicle, housing provides a flow of services that enters utility along with non-housing consumption. We include additional realistic features in the model. Adjusting the stock of housing requires the consumer to incur transaction costs meant to capture realtor fees and other costs related to buying and selling a home. This gives rise to a lumpy pattern in housing investment. The model includes credit constraints in the form of minimum down payment requirements for mortgages. We also include bequest incentives by allowing end of life cycle wealth to enter into life time utility. We adopt a partial equilibrium approach by modeling the evolution of income and home prices as exogenous first order Markov processes.

We estimate the structural parameters of our model using household level data on income and housing tenure decisions from the Panel Study of Income Dynamics (PSID). Estimating a fully dynamic model of housing demand is technically challenging; solving the household's dynamic programming problem is computationally difficult because of two key types of non-convexities. First, housing demand has discontinuities arising from transaction costs to adjust housing stocks. Second, households may face credit constraints because conforming mortgages typically require a 20% down payment. As a result, it is not possible to characterize optimal decisions using Euler equations, nor are we able to use standard GMM methods to estimate the structural parameters.

Instead, we use the multistep method proposed by Bajari, Benkard, Levin (2007) (hereafter BBL) to estimate the model. The first step of BBL requires us to estimate housing decisions rules and the law of motion for the state variables. We use a multinomial logit model of housing investment decisions in

[2]Source: Mortgage Bankers Association.

[3]According to the 2009 Freddie Mac annual ARM survey, the share of ARM applications is at an all time low of 3%.

the spirit of Han (2008). In each period households either upgrade to a larger home, downgrade, or remain in their existing home. This reduced form model provides a flexible way to capture lumpy patterns in housing investment. We estimate the evolution the exogenous state variables using standard time series and panel econometric techniques. In the second step, we estimate the structural parameters in household utility. The estimator proposed by BBL solves a revealed preference problem. We assume that the policy function estimated for housing investment in the first stage is the solution to a household's dynamic programming problem. The estimator reverse engineers a period utility function that rationalizes the estimated decision rules by solving a system of revealed preference inequalities. An attractive feature of our estimator is that it allows for non-convex adjustment costs and credit constraints.

Given our parameter estimates, we use our dynamic model to simulate a typical household's response to a set of negative shocks meant to mimic the current disruptions in the U.S. housing market. The counterfactuals we consider are a housing price bust, a decrease in income, and a tightening of lending standards. Our results demonstrate that many households do not alter their housing choice in response to these shocks, but rather adjust nondurable consumption expenditures. The intuition behind this result is simple- households in the model only move two to three times before retirement (as do most households in the data). Because they are locked in, changes in housing market conditions do not influence their level of housing stock nor consumption. Instead, they experience a decline in home equity. Households respond to a decline in income by reducing expenditures on nondurable consumption good.

The paper is organized as follows. We review the theoretical and empirical literature relevant for our study in the next section. Section 3 describes the model. Section 4 presents the data and descriptive evidence. In section 5, we provide an overview of the BBL estimation technique. Section 6 and 7 describe the first and second stage of estimation. In section 8 we present our benchmark model simulation results and section 9 conducts sensitivity analysis with respect to crucial structural parameters. Section 10 concludes; additional details about computation are relegated to the appendix.

2.2 Related Literature

Following Mankiw's (1982) seminal aggregate study of consumer durables, a sizeable literature has developed that uses structural household-level models of housing demand and tenure choice to study the interaction between house prices, household consumption, and tenure decisions. The closest paper in spirit to ours is the work by Li and Yao (2007), who construct a life cycle model of housing tenure choice to study the impact of house price changes on housing choices and consumer welfare for different age groups of the population. The key distinction between their paper and ours is two-fold. First, while their paper contains a more explicit model of the life cycle and housing tenure choice, they restrict attention to a preference specification in which the elasticity of substitution between nondurables and housing services is fixed to one (that is, the aggregator is of Cobb-Douglas form). Our paper, on the other hand, uses novel empirical techniques to *estimate* this crucial parameter and finds the elasticity to be larger than one. Our model simulations document that this difference has important implications for the dynamics of housing and consumption choices.[4] Second, while their main focus lies on the impact of house price changes on consumption allocations over the life cycle and the distribution of its welfare impacts across different households, we focus more directly on the impact of house price shocks on housing demand and the demand for nondurable consumption.

Flavin and Yamashita (2002), Fernandez-Villaverde and Krueger (2002), Yao and Zhang (2005), Hintermaier and Koeniger (2009), Kiyotaki et al. (2008), Iacoviello and Pavan (2009). and Diaz and Luengo Prado (2008, 2009) use a similar life cycle model to study the impact of the presence of housing on portfolio choice, precautionary saving and the wealth distribution. The latter authors also employ their model to argue that the current user cost approach to measure price changes for housing services in the consumer price index (CPI) is biased in the presence of owner-occupied housing and household heterogeneity. In a sequence of quantitative papers Chambers, Garriga and Schlagenhauf (2007, 2009a,b) use a life cycle model with tenure choice to explore the impact of tax treatments and different mortgage designs on home

[4] Li et al. (2009) perform a structural estimation of a model very similar to Li and Yao (2007), but use different techniques than the empirical approach employed in this paper.

ownership rates and explore the reasons for the substantial increase in this rate in the U.S. in the last decade. Yang (2009) documents the role of down payment constraints and transaction costs on the life cycle profile of housing and consumption. Oralo-Magne and Rady (2006) study the interaction of financial markets conditions and home ownership rates in a structural model of housing choices. Lustig and van Nieuwerburgh (2006) and Piazessi, Schneider, and Tuzel (2007) explore the connection between house and asset prices.

The role as collateral of housing in particular in the main theme of a recent literature that focuses on the joint housing and mortgage choice. Important examples of this work include Hurst and Stafford (2004), Luengo-Prado (2006), and Chambers et al. (2009a, b). The recent increase in default on mortgages has motivated a small but growing literature on structural models of foreclosures within this context. See, for instance, Jeske and Krueger (2005) and Garriga and Schlagenhauf (2009). The same issue is analyzed empirically, among others, by Carroll and Li (2008).

On the empirical side, an attempt has been made to quantify the wealth effect on nondurable consumption from changes in house prices. Leading work in this strand of the literature include Case, Quigley, and Shiller (2003), Benjamin, Chinloy and Jud (2004) and Campbell and Cocco (2007). They document a sizable housing wealth effect and contrast their results to estimates of the wealth effects from other financial assets. Our focus is on the negative wealth effect from home price declines as well as the interaction with down payment constraints and negative income shocks.

Our paper also contributes to the literature on estimation of dynamic decision problems with non-convex adjustment costs. To the best of our knowledge, the only paper that attempts to estimate such a structural model is Hall and Rust (2003). The simulated minimum distance estimator that they propose is not computationally feasible in our application because it requires repeatedly computing the optimal policy function. Computing a single optimal policy function in our model takes one week of CPU time with an advanced workstation. Our estimator avoids the burden of repeatedly computing optimal policies. Finally, our paper is closely related to Han (2008). Our first stage is quite similar to her reduced form model of housing demand. We depart from her work because we estimate a households structural utility parameters.

This allows us to explore counterfactuals that would not otherwise be possible in a reduced form approach.

2.3 The Model

We model a typical households' consumption and housing choice as a partial equilibrium, dynamic decision problem with a finite lifetime horizon. Households live for T periods and in each period t choose consumption expenditures on nondurables, c_t,and the amount of one period-risk free financial assets to bring to the next period, a_{t+1}. We let h_t denote the size of the household's real housing stock at the beginning of the period so that h_{t+1} is the amount of housing chosen today for tomorrow. A household derives a service flow $g(h_t) = gh_t$ from the housing stock where $g > 0$ is a parameter. In our applications we shall assume that $g = 0.075$. Households value nondurable consumption and housing services according to a standard intertemporal lifetime utility function

$$U\left(\{c_t, h_t\}_{t=0}^{T}\right) = E_0\left[\sum_{t=1}^{T} \beta^{t-1} u(c_t, g(h_t)) + \gamma \beta^T \log(b_T)\right] \tag{2.1}$$

where β is the standard time discount factor and the parameter γ measures the degree of altruism that motivates households to leave bequests b_T at the end of life. Expectations E_0 are taken with respect to the stochastic processes driving labor income and house prices, to be specified now. Let p_t denote the relative price of one unit of housing, in terms of the numeraire nondurable consumption good. Housing prices $\{p_t\}_{t=0}^{T}$ follow first order stochastic Markov processes.

At time 0, agents are endowed with initial asset holdings (a_0, h_0) and one unit of time per period, which they supply inelastically to the labor market to earn labor income y_t. The labor income process is composed of two components, a deterministic mean life cycle profile ε_t (which incorporates aggregate income growth in the economy as a result of technological progress) and a stochastic component η_t that follows a first order Markov process. Households retire at an exogenous age T_r and receive a flat pension benefit b until they die. Thus labor income is given by

$$y_t = \begin{cases} \varepsilon_t \eta_t & \text{if } t < T_r \\ b & \text{if } t \geq T_r \end{cases}$$

We model the two main frictions in the housing market explicitly. First, the stock of housing is subject to nonconvex adjustment costs. Specifically, in order to purchase a home of size h_{t+1} the household has to spend

$$p_t h_{t+1} + p_t \phi(h_{t+1}, h_t)$$

where $p_t h_{t+1}$ is the purchase price of the home and $p_t \phi(h_{t+1}, h_t)$ is the transaction cost a household has to bear when adjusting the owned stock housing from h_t to h_{t+1}. We assume that the function ϕ takes the form

$$\phi(h_{t+1}, h_t) = \begin{cases} \phi * (h_{t+1} + h_t) & \text{if } h_{t+1} = h_t \\ 0 & \text{if } h_{t+1} = h_t \end{cases}$$

where the number ϕ measures the percent of the value of the house that both a seller and a buy has to bear as transaction cost. In most of our analysis we shall assume that $\phi = 0.03.$, that is, which is representative of real estate fees.[5]

The second key friction in the housing market we model explicitly is the requirement for households to acquire (and maintain) some minimal positive equity share in the house. We assume that the joint choice of financial assets and housing positions satisfies the following collateral constraint:

$$a_{t+1} \geq -(1 - \xi) p_t h_{t+1}. \tag{2.2}$$

Here $\xi \in [0, 1]$ is the fraction of the purchase price of the house that has to be paid down at purchase, i.e. $(1-\xi)$ is the fraction of the purchase price that can

[5] We do not have direct data on transaction prices. In principle, we could infer them indirectly through the use of our structural model. However, taking a direct stance on transaction costs will give a more efficient estimate of the remaining model parameters. As we shall show in our policy simulations, our model is able to reasonably match the observed moving frequencies in our data, where transaction costs play a large role. Also, the simulation results suggest that qualitatively, many of our policy conclusions will be robust to a fairly broad range of transaction costs as long as they are sizeable (and as long as the adjustment cost function has the nonconvex form we have specified).

be financed via a mortgage. In most of our experiments we shall assume that households are able to finance at most 80 percent of their housing purchases through mortgages.[6] Also note that as long as $\xi \in [0, 1]$ households can only borrow against their housing collateral; uncollateralized debt is therefore ruled out by assumption in our model.

Thus the key frictions in the housing market are summarized by the transaction cost ϕ parameter and the collateral constraint parameter ξ, with $\phi = \xi = 0$ denoting frictionless housing markets. Our simulation exercises will therefore be able to quantify the importance of frictions in the housing market by deducing optimal choices of households under various assumptions on (ϕ, ξ).

In addition to housing households can use financial assets to accumulate wealth. These assets yield a constant real interest rate r. If households borrow (subject to the collateral constraints) they face a real mortgage interest rate $r^m > r$. In most of our exercises we treat interest rates as constant, but we will analyze how the dynamic consumption and asset accumulation choices change with the level of interest rates households face.

Defining

$$r(a) = \begin{cases} r & \text{if } a \geq 0 \\ r^m & \text{if } a < 0 \end{cases}$$

the budget constraint can be written as

$$c_t + a_{t+1} + p_t h_{t+1} + p_t \phi(h_{t+1}, h_t) = y_t + (1 + r(a_t))a_t + p_t h_t \tag{2.3}$$

Finally, consumption and housing choices are constrained to be nonnegative:

$$c_t, h_{t+1} \geq 0 \tag{2.4}$$

Households maximize (2.1) subject to the constraints (2.2), (2.3) and (2.4). In the appendix we offer further details on the recursive formulation and computation of this partial equilibrium household decision problem.

[6] This is a typical down payment requirement in conforming mortgages offered by Fredie Mac and Fannie Mae. Our sample extends to the 1997 and predates the explosion in subprime mortgages. A 20 percent downpayment requirement is representative of credit constraints during this time period.

2.4 Data and Descriptive Evidence

In this section, we discuss our data which primarily comes from the Panel Survey of Income Dynamics (PSID), a national household level longitudinal survey. We then present descriptive evidence to address the behavioral implications of the model. Specifically, we document lumpy housing investment and life-cycle patterns of housing tenure decisions.

¿From the PSID, we select a panel of households between the years 1980 and 1993, including only households that own their primary residence. [7] The sample contains 1931 households for a total of 14,556 household-year observations. The key variables are income, value of primary residence, annual mortgage payment, home equity, an indicator of whether the household moved in a given year, and demographic information. We supplement our data with 30 year fixed interest mortgage rates from the Federal Housing Finance Agency.[8] We use the home price index from the Federal Housing Finance Agency which collects home price information from home sales financed with Fannie Mae and Freddie Mac conventional mortgages. Both series span the years 1975 to 2009. We also impute measures of non-housing wealth and consumption from the Consumer Expenditure Survey (CEX). The CEX data are not reliable for estimation, but are nonetheless informative for descriptive purposes.

2.4.1 Aggregate Descriptive Evidence

Table 2.1 lists descriptive summary statistics. All dollars values are deflated by the non-housing component of the consumer price index which serves as our numeraire good. The base year is 1980. Housing is a significant component of household wealth and expenditures. The average household occupies a $90,000 home valued at double its annual income. On average households have a 2/3 equity stake in their home which is by far the largest source of household wealth. Home equity averages $59,000, whereas non-housing wealth reported

[7]We only include households where the head of household is married, between the age of 20 and 65 and born between 1920 and 1959. We exclude households where the head of household changed during the sample period and any household that reports an income below $10,000 or above $150,000, a trimming of approximately the top and bottom 2% of households based on income. We also exclude a small number of households reporting negative home equity balances.

[8]Formerly the Office of Federal Housing Enterprise Oversight (OFHEO)

Table 2.1: Summary Statistics

Variable	Num Obs	Mean	Std Dev	Min	Max
Income	14,456	45,783	23,691	10,194	163,183
Home Size	14,456	80,680	64,450	363	725,986
Home Value	14,456	89,869	73,012	436	871,175
Home Equity	14,456	59,543	61,451	39	871,228
Home Value/Income	14,456	2.07	1.54	0.006	42
Home Equity/Home Value	14,456	0.660	0.029	0.0004	1
Annual Mortgage Payment	14,456	4,306	4,568	0	80,293
Consumption*	14,456	18,934	21,117	0	815,342
Non-Housing Wealth*	14,456	12,621	21,553	-38,602	139,379
Move Indicator	11,343	0.0426	0.202	0	1
Upgrade Indicator	11,343	0.0321	0.176	0	1
Downgrade Indicator	14,343	0.0105	0.102	0	1

All dollar values deflated by non-housing component of consumer price index: base year 1980. Move indicator variables reported on a truncated sample size that only includes observations with a non-missing lagged observation.
Home size represents the quantity of housing expressed in 1980 dollars. It is calculated by dividing real home values by the FHFA home price index. We do not have information on household locations and thus are not able to account for regional differences in home prices.
* Imputed from Consumer Expenditure Survey

in the CEX averages \$12,000. Most households, 77%, finance their home using a mortgage and carry a positive mortgage balance. Annual mortgage payments account for 9% of total income on average.

Adjustment costs to moving play a key role in our model. If the costs are high enough, the model predicts households should move to a different home infrequently and make lumpy adjustments to housing stock. This behavior is evident in the data. Households move in just 4.3% of the years. The size of the adjustments are large, averaging a change of \$40,000 in home size. Households typically upgrade to larger homes; 75% of moves are upgrades, with an average increase in home size of \$24,000. Table 2.2 lists frequencies for the number of moves per household during the 14 year sample period. The majority of households, 58%, never move, and it is rare for a household to move more than once. Extrapolated over a life-cycle, households move one or two times after their first home purchase. Consistent with lumpy investment behavior, the large standard deviation in the ratio of house value to income (1.54 reported in table 2.1) in part indicates households do not continuously adjust housing in response to income fluctuations.

Table 2.2: Household Moves

Number of Moves	Frequency
0	58%
1	34%
2	6.5%
3	1.4%
>4	0.4%

Sample duration 1980-1993.
Note: we may be underreporting the number of moves because of gaps in the panel.

2.4.2 Life Cycle Patterns

Table 2.3 reports summary statistics by birth cohort. Several patterns emerge about life cycle behavior.[9] Both income and housing exhibit a hump-shaped pattern. Young households move into progressively larger homes until they are middle aged and in the process accumulate home equity, both in total and as a percentage of home value. But as households approach retirement, income falls and they downgrade to smaller homes. The frequency of moves declines in age, from 6.3% for the youngest cohort to 2.7% for the oldest.

2.5 Estimation Procedure

We use the two-step method proposed by Bajari, Benkard, and Levin (2007) for estimating stochastic dynamic models. It is a burgeoning technique that is finding applications in many settings, such as the estimation of dynamic oligopoly models (Ryan (2006), Sweeting (2007), Snider (2008)). It can also be applied to dynamic single agent choice problems, and is a particularly attractive method for estimating rich models like our application. It can accommodate continuous controls and non-convexities such as adjustment costs and credit constraints. More importantly, the method is computationally feasible. In our simulation exercises it takes up to one week of CPU time to solve our model. Therefore, a nested fixed point algorithm which requires repeatedly

[9] In general a cohort table cannot distinguish life cycle effects from cohort effects; life cycle interpretation is only suggestive.

Table 2.3: Birth Cohort Sample Means

Birth Cohort	1956-59	1951-55	1946-50	1941-45	1936-40	1931-35	1926-30	1921-25
Num obs	1,244	2,690	2,973	1,907	1,321	1,458	1,851	924
Income	42,372	44,745	47,311	47,718	50,234	47,092	45,332	37,690
Home Size	69,190	78,573	88,106	81,222	80,950	84,702	81,093	71,408
Home Value	79,973	89,159	98,127	90,089	89,294	93,281	89,650	76,518
Home Equity	35,791	45,158	58,139	57,105	64,072	78,239	78,686	68,294
Home Equity/ Home Value	0.435	0.491	0.579	0.654	0.729	0.835	0.862	0.902
With a Non-missing Lagged Observation								
Num obs	945	2,209	2,349	1,512	1,032	1,148	1,466	637
Move Frequency	0.063	0.057	0.054	0.046	0.019	0.028	0.022	0.027
Upgrade Frequency	0.056	0.044	0.045	0.033	0.013	0.025	0.014	0.011
Downgrade Frequency	0.007	0.013	0.009	0.013	0.007	0.013	0.008	0.016
Conditional on Move								
Num obs	60	126	126	69	20	32	32	17
Δ Home Size	31,976	28,961	37,982	23,802	13,825	-5,682	-7,801	-10,977

All dollar values deflated by non-housing component of consumer price index: base year 1980.

calculating the full solution to the model will not be computationally possible.

The estimator is based on the principle of revealed preferences and is quite intuitive in its application. Consider a simple static setting to understand the intuition. The estimator compares utility from observed actions in the data to alternative, non-optimal actions and selects utility function parameters that best rationalize the observed choices. It only requires that the data are generated by rational agents maximizing utility. In a dynamic setting the estimator compares discounted expected utility from estimated optimal policy functions to suboptimal policy functions. Making this comparison requires an additional step of forward simulating actions to estimate expected discounted utility. We first outline this step in the context of our application and proceed to formally describe the revealed preference estimation approach.

In the model in section 3 the exogenous state variables $s = \{y, p, r\}$: income, home prices, and interest rates follow a first order Markov process.

$$s' = \pi(s) \tag{2.5}$$

where in recursive form s' is next periods value of the state variables.

A household's policy function,$\sigma(\cdot)$, inherits the Markov property;

$$(a', h') = \sigma(a, h, s) \tag{2.6}$$

where a and h denote endogenous state variables: household debt and housing stock. Using standard techniques in cross sectional and time series econometrics we form estimates $\widehat{\pi}$ and $\widehat{\sigma}$ of the Markov process for exogenous state variables and policy functions.

In the first stage we use forward simulation to estimate discounted expected utility. Consider household i at time t with current state variable $s_{i,t}$. We simulate a pseudo-random sequence of the state variable and housing and savings decisions for time period $\tau = t, t+1, t+2, ..., T$ using the following algorithm:

1. Given $(a_{i,\tau}, h_{i,\tau})$ and $s_{i,\tau}$, draw $s_{i,\tau+1} \sim \pi\left(a_{i,\tau}, h_{i,\tau}, s_{i,\tau}\right)$
2. Let $(a_{i,\tau+1}, h_{i,\tau+1}) = \hat{\sigma}(s_{i,\tau+1})$
3. If $\tau < T - 1$, return to 1.

We draw $r = 1, ..., R$ pseudo-random sequences in this manner. Denote a generic sequence as $\{a_{i,\tau,r}, h_{i,\tau,r}, s_{i,\tau,r})\}_{\tau=t}^{T}$.

Define $\hat{E}[U; \theta, \hat{\sigma}, \widehat{\pi}, s_{it}]$ as:

$$\hat{E}[U; \theta, \hat{\sigma}, \widehat{\pi}, s_{it}] = \frac{1}{R} \sum_{s=1}^{R} \sum_{\tau=t}^{T} \beta^{\tau} u(a_{i,\tau,r}, h_{i,\tau,r}; \theta) \tag{2.7}$$

We interpret $\hat{E}[U; \theta, \hat{\sigma}, \widehat{\pi}, s_{it}]$ as an estimate of agent i's expected discounted utility at time t when the state is $s_{i,t}$. Note that $\hat{E}[U; \theta, \hat{\sigma}, \widehat{\pi}, s_{it}]$ holds fixed the Markov process $\widehat{\pi}$, the (estimated) optimal policy $\hat{\sigma}$, and the parameters of the utility function θ.

Let $\sigma^a \neq \hat{\sigma}$ denote a policy function other than the estimated policy $\hat{\sigma}$. We can simulate the expected discounted utility from σ^a as:

$$\hat{E}[U; \theta, \sigma^a, \widehat{\pi}, s_{it}] = \frac{1}{R} \sum_{s=1}^{R} \sum_{\tau=t}^{T} \beta^{\tau} u(a^a_{i,\tau},_r, h^a_{i,\tau,r}; \theta) \tag{2.8}$$

where $\left\{a^a_{i,\tau,r}, h^a_{i,\tau,r})\right\}_{\tau=t}^{T}$ is draw analogously to step 1-3 above replacing $\hat{\sigma}$ with σ^a.

We perform these forward simulations for all N household-year observations and A alternative policy functions. We define our estimator as:

$$\hat{\theta} = \arg\min \frac{1}{AN} \sum_{a=1}^{A} \sum_{i=1}^{N} \max \left\{ \hat{E}[U;\theta,\sigma^a,\widehat{\pi},s_{it}] - \hat{E}[U;\theta,\hat{\sigma},\widehat{\pi},s_{it}], 0 \right\}^2 \quad (2.9)$$

If the household is rational, it must be the case that at the true utility parameter value θ_0 $E[U;\theta_0,\sigma,\pi,s_{it}] > E[U;\theta_0,\sigma^a,\pi,s_{it}]$ for any $\sigma^a \neq \sigma$. That is, the true policy yields higher utility than any alternative policy. Suppose we have a very large sample of household-year observations N and simulation draws R. Consistency of the first stage estimates guarantees that $\hat{\sigma}$ converges to σ and $\widehat{\pi}$ to π. The large value of R will then guarantee that $\hat{E}[U;\theta,\hat{\sigma},\widehat{\pi},s_{it}]$ converges to $E[U;\theta,\sigma,\pi,s_{it}]$. This implies that the sample analogue of revealed preference should hold in a limiting sense. That is, $\widehat{E}[U;\theta_0,\hat{\sigma},\widehat{\pi},s_{it}] > \widehat{E}[U;\theta_0,\sigma^a,\widehat{\pi},s_{it}]$. The loss function in the estimator penalizes those observations corresponding to large deviations from utility maximization. In this sense, the estimator selects a $\hat{\theta}$ which rationalizes $\hat{\sigma}$ as an optimal choice rule.

Bajari Benkard and Levin (2007) discuss the formal econometric theory and provide regularity conditions for consistency and asymptotic normality. In order to guarantee consistency and asymptotic normality, there are three sources of error to consider. The first is sampling error which makes our first stage policy function and exogenous state variable estimates $\hat{\sigma},\widehat{\pi}$ differ from their true value σ,π . The second is simulation error from choosing a finite number of simulations paths R. The third comes from using only a finite number of alternative policies A.

2.5.1 Utility Function

We model period utility with a familiar constant elasticity of substitution form. During the life cycle households derive utility from housing, h and non-durable consumption c, according to the period utility function,

$$u(c,h) = \log\left[(\theta c^{\tau} + (1-\theta)(\kappa h)^{\tau})^{\frac{1}{\tau}} \right]$$

where θ and τ are utility parameters to be estimated. The term, θ captures consumption shares, and the term τ captures the elasticity of substitution between housing and non-housing consumption. The elasticity parameter is particularly important as our counterfactual exercises will demonstrate in the next section. Other models of housing demand (see Li and Yao (2007)) assume this parameter takes on a value of zero (unit elasticity). We set the utility flow of housing parameter to $\kappa = 0.075$. This value is consistent with the literature on owners equivalent rent.

Households discount period utility at a rate $\beta = 0.97$ through age 70.[10] Utility from bequest motives, or—depending on interpretation—retirement wealth enters additively.[11] Lifetime utility for a household of age, $a = a_0$, is

$$U(\tilde{c}, \tilde{h}, q_{70}) = \sum_{a=a_0}^{70} \beta^{a-a_0} u(c_a, h_a) + \beta^{70-a_0} \gamma \log(q_{70})$$

where q_{70} is the amount of home equity at age 70. The log functional form on the bequest/retirement component imposes an Inada condition that home equity should be positive. The bequest parameter γ is the third parameter to be estimated.

2.6 First Stage: Reduced Form Policy Functions

In the presence of transactions cost to moving, households move infrequently and, when they do, make lumpy adjustments to their stock of housing. To capture the discreteness in the timing of moves, we model the moving decision as a discrete choice problem. There are three alternative moving decisions M_{it}: remaining in the existing home, downgrading to a smaller home, and upgrading. We use a multinomial logit model. The probability of household i making a moving choice of j at time t is given by

[10] We must fix the discount factor because they are poorly identified in dynamic models (Rust (1987)).

[11] We end the life cycle at age 70 because the support of the age distribution ends at age 65 in the data. We are only comfortable extending to age 70. Given many people live well beyond 70, the motive for have a large home equity stake could be to finance retirement expenditures.

$$p_{it,j} = P[M_{it} = j] = \frac{\exp(\mathbf{x}'_{it}\beta_j)}{\sum_{k=1}^{3} \exp(\mathbf{x}'_{it}\beta_k)}$$

Two of the primary covariates in $\mathbf{x}'_{it}$ are income and the size of the current home. We expect households that experience an increase in income to be more likely to upgrade and less likely to downgrade. Households occupying small homes would be more likely to upgrade and less likely to downgrade. To capture the impact of a down payment constraint, we include a linear spline in the home equity ratio. Specifically, we allow the effect of the home equity ratio to differ for equity ratios above and below 15%. If down payment constraints have an important effect on housing decisions we would expect the likelihood of moving—in particular, upgrading—to reduce significantly as equity falls through the lower range. We would expect the equity ratio to have much less effect in the higher home equity range. We apply the user cost method to measure the price of homeownership. It is a widely accepted measure for the cost of durable goods, such as housing. We use a simple measure: the difference in real interest rates and the expected real rate of home price inflation. This measure captures both the financing costs of mortgage interest payments and the offsetting investment component of housing capital gains. In our application, we assume households have rational expectations and thus measure expected home price inflation using the realized value of contemporaneous home price inflation. We also condition on age and age-squared.

We use a parsimonious specification to capture the lumpiness in housing stock adjustments. We model the size of housing stock adjustments as a simple average of log changes in home size:[12]

$$|\log(h_{it+1}) - \log(h_{it})| = \beta_0 + \epsilon_{it}$$

Under the assumption that the ϵ error term and u error term in the multinomial logit model are uncorrelated, we can estimate the adjustment size

[12] We also tested methods that distinguished upward and downward adjustment sizes and parameterized adjustment size through factors such as income and home size. We found that such models performed poorly when simulations veered outside the support of the distribution of the variables. Moreover, with so few observations of upgrades and downgrades there is a large loss in degrees of freedom from estimating separate adjustment sizes.

separately using the subset of households that either upgrade or downgrade. The error term in the multinomial logit model can be interpreted as capturing unobserved factors such as a job relocation. Such non-financial factors on their own would affect the likelihood of moving, but have no bearing on the adjustment size.[13]

2.6.1 Reduced Form Policy Function Results

Table 2.4 reports results for the multinomial logit model. The coefficients on the decision to remain in the current home are normalized to zero. Almost all of the coefficient signs are as expected. Households in larger homes are more likely to downgrade and less likely to upgrade. As income increases, they are more likely to upgrade and less likely to downgrade. As the price of housing increases, measured by the user cost, households are less likely to upgrade. User cost has an insignificant effect on downgrades. We find down payment constraints are quite important. In the low equity range, below 15%, a drop in equity lowers the probability of upgrading. The magnitude is very large. Somewhat surprisingly, they are also less likely to downgrade as equity falls. In fact this makes sense. For any move, even a downgrade, households would have difficulty securing a new mortgage. In the high equity range, home equity has no significant effect on upgrades, but extra equity makes it less likely for a household to downgrade. Finally, older households are less likely to make any sort of adjustment.

2.6.2 Mortgages and Savings

We do not have reliable data on household savings, mortgage choices, and bequests and cannot estimate these decisions from data. Instead we assume households used mortgage products that were common during the 1980's and early 1990's. We also assume housing is only the source of wealth. This is reasonable because our data from the CEX shows non-housing wealth is insignificant portion of households' total assets.

[13] We also experimented with an (S,s) inventory model of durable good expenditures as in Attanasio (2000) and Ryan (2006). This model imposes more structure than a flexible multinomial logit. With so few adjustments, the (S,s) model did not yield reasonable results.

Table 2.4: Upgrade/Downgrade Multinomial Logit

	Downgrades	Upgrades
Home Size	0.0057	-0.0098
	(0.0010)	(0.0014)
Income	-0.0131	0.0222
	(0.0048)	(0.0024)
Low Home Equity Ratio (< .15)	6.5557	13.3245
	(7.2891)	(4.9580)
High Home Equity Ratio (> .15)	-0.8571	-0.2155
	(0.4175)	(0.2279)
User Cost	-0.0201	-0.0442
	(0.0311)	(0.0183)
Age	-0.1479	-0.1433
	(0.1026)	(0.0661)
Age^2	0.0016	0.0009
	(0.0011)	(0.0007)
Constant	-1.6934	-0.9043
	(2.4916)	(1.5820)
Num Moves	119	364
Num Obs	11343	11343

Magnitude of Adjustment Size in Logs		
Num Obs	Mean	Std. Dev.
483	0.496	0.486

All dollar values deflated by non-housing component of consumer price index: base year 1980. There is a linear spline in the home equity to house value ratio term with a knot value at 0.15. The reported coefficients are the marginal effects in each region. Standard errors in parentheses. Magnitude of adjustment size measured as $|\log(h_{it}) - \log(h_{it-1})|$ for households that move. Sample restricted to observations with non-missing lagged observations.

We assume households finance with 30 year fixed rate mortgages. A household that moves transfers home equity from its prior residence into the new home. The remaining value of the home is financed with a new 30 year fixed rate mortgage set at the prevailing interest rate.[14]

In some instances, particularly near the end of the life cycle, a household could have excess home equity. This happens for downgrades to homes valued less than the current amount of home equity. We lock this excess wealth into an annuity with an amortization term that expires at end of the life cycle. Thus, the annuity is completely drained at the end of life cycle. Periodic payments supplement income. Consumption is simply the difference between income and mortgage payments. If a household pays off its mortgage in full, it consumes its entire income. Households do not prepay or refinance mortgages, nor can they borrow against their home equity. These assumptions restrict a household's ability to draw down home equity to smooth consumption and to save beyond the value of a home in anticipation of a bequest.

When a household moves, it pays an adjustment cost equal to six percent of the new home's value. This is the industry standard realtor fee. We interpret this cost as a conservative lower bound because there are search costs, relocation expenses, and other non-monetary adjustment costs that we cannot quantify.

We do not directly impose a down payment constraint. Instead, we let down payment constraints enter through the equity ratio term in the multinomial logit model of moving decisions. Finally, it is necessary to impose a subsistence requirement. In rare instances simulated consumption would be negative. We endow those households with $1,000 in consumption. Likewise, if home equity is negative at the end of the life cycle, we endow those households with $1,000 in equity for a bequest. Both of these assumptions can be interpreted as an insurance policy against particularly poor shocks.

We must make these assumptions because we do not have reliable data on mortgage financing and savings. Nonetheless, we feel they are reasonable, and we show that our forward simulations match key moments from the life cycle quite closely.

[14] The term is 30 years regardless of age. It may be reasonable to assume older households choose shorter terms, but we do not have data that would allow us to predict term lengths.

Table 2.5: Income Process Estimates

	Estimates: $\log(y)_{it}$
Age	0.1141
	(0.0046)
Age-squared	-0.0011
	(0.00005)
Cohort	0.01489
	(0.0012)
Constant	-28.0945
	(2.3183)
ρ	0.506
σ_ϵ	0.242
σ_ν	0.731
σ_a	0.399
Observations	14446
Households	1931
R-squared	0.104

Random Effects with AR(1) error term disturbance. Regressand: log consumption-deflated income. Cohort is the head of household's birth year. Standard errors in parentheses.

2.6.3 Exogenous State Variables: Income, Home Prices and Interest Rates

The process for income includes an age component, a cohort effect, household random effects, and an AR(1) error disturbance:

$$\begin{aligned}
\log(y_{it}) &= \beta_0 + \beta_1 age_{it} + \beta_2 age_{it}^2 + \beta_3 birthcohort_i + \alpha_i + \epsilon_{it} z_{it} \\
z_{it} &= \rho z_{it-1} + \nu_{it} \\
\epsilon_{it} &\sim N(0, \sigma_\epsilon^2) \\
\nu_{it} &\sim N(0, \sigma_\nu^2) \\
\alpha_i &\sim N(0, \sigma_\alpha^2) \qquad (2.10)
\end{aligned}$$

Estimates are reported in table 2.5 which shows that income is persistent and exhibits a hump-shaped pattern over the life-cycle.

We model the time series process for real interest rates i_t and real home price inflation π_t as a Vector Auto-Regression (VAR) with one lag,

Table 2.6: Interest Rate and Home Inflation VAR

interest rate		home inflation	
i_{t-1}	0.615	π_{t-1}	0.534
	(0.105)		(0.147)
π_{t-1}	-0.435	i_{t-1}	0.113
	(0.086)		(0.178)
constant	2.845	*constant*	0.178
	(0.641)		(1.090)
Error Covariance			
i	3.399		
π	2.543	9.823	
Num Yrs	33	1975-2009	

Standard errors in parentheses.

$$
\begin{aligned}
i_t &= \beta_{ci} + \beta_{ii} i_{t-1} + \beta_{i\pi}\pi_{t-1} + e_{it} \\
\pi_t &= \beta_{c\pi} + \beta_{\pi\pi}\pi_{t-1} + \beta_{\pi i} i_{t-1} + e_{\pi t}
\end{aligned}
$$

where the error term is distributed bivariate normal, $\mathbf{e} \sim \mathbf{N}(\mathbf{0}, \Sigma)$. We use a longer time series that spans the years 1975 to 2009. Results are presented in table 2.6. Notice the coefficient on the home inflation constant; average real home price inflation was just slightly greater than zero over this time period.

2.6.4 Goodness of Fit

We use two methods to evaluate the goodness of the fit of our forward simulations. First, we examine the entire, 40 year, simulated life cycles for the youngest cohort of households aged 30 to 34. Second, we consider the first five years of forward simulation for all cohorts. In both exercises, we average several key variables across households and simulation paths for five year intervals.

Table 2.7 reports the full life cycle simulation for the youngest cohort. As a first check, compare the cohorts' initial conditions in the data column to its first five years of simulations in the "ages 30-34" column. The fit is sensible. Beyond the first five years, the simulated life cycles show a distinct hump shaped pattern in income and consumption. There is also a hump shaped pattern in home size; in early years households upgrade rapidly, then there is a leveling off mid-life, and a slight drop at the tail end of the life cycle. This pattern is largely driven by upgrading and downgrading frequencies. In early years, there is a high rate of upgrades which gradually declines to a low level by the end of the life cycle. Downgrading frequencies are very low and constant up until age 55 at which point they start increasing. A comparison with birth cohort statistics (see table 2.3) shows that move frequencies match the data quite closely. Although the dollar-valued magnitudes cannot be directly compared, notice that for income and home size the peaks in the hump exactly match the age of the hump in the data. That is, in mid-life (ages 50-54) income peaks, and with a lag (ages 55-59) home size peaks. We also see a steady growth in house values over the life cycle which is partly due to upgrades, and also real home price inflation. In addition, home equity grows as households pay off mortgages and experience home price capital gains. As in the data, we capture an increase in the home equity ratio (home equity/home value) over the life cycle. As a minor point, we include an entry for non-house wealth. This is the size of the annuity from left over home equity for those households that downgrade to homes worth less than their home equity. The magnitude is small, on the order of at most 2% of home equity wealth, and thus the assumptions about non-house savings behavior should have a negligible effect on our final results.

Table 2.8 reports statistics for each cohort, not just the youngest. Under each cohort heading, the first column reports the statistics from the data, and the second, statistics from five years of forward simulation. Again, the simulations are reasonable; nothing wild happens. The first five years of simulation are close to the initial conditions and the trajectories of the variables match the hump-shaped patterns that we expect. Young households grow income, housing stocks, and wealth, while older households experience declines. Moving frequencies match quite closely.

These tables are useful for understanding identification of the utility function parameters. The shares of housing and consumption in the simulated data help to jointly identify the elasticity of substitution parameter, τ and consumption share parameter θ (refer to the utility function in section 5). Separately identifying these parameters is more subtle. The frequency and

Table 2.7: Forward Simulation for Youngest Cohort

Years	0	0-5	6-10	11-15	16-20	21-25	26-30	30-35	36-40
Age	*data* 30-34	30-39	35-44	40-49	45-54	50-59	55-64	60-69	65-74
Income	*42,372*	47,648	55,403	61,975	65,717	66,356	63,457	57,642	54,586
Home Size	*69,190*	74,992	85,781	95,294	102,767	107,454	109,421	108,678	107,477
House Value	*79,973*	89,650	111,681	136,486	161,189	184,227	206,344	225,041	231,180
Home Equity	*35,791*	37,579	47,063	62,420	82,252	106,761	136,716	168,169	179,835
Home Equity/ Home Value	*0.435*	0.415	0.443	0.494	0.558	0.634	0.721	0.802	0.827
Move Freq.	*0.063*	0.071	0.063	0.050	0.041	0.032	0.031	0.032	0.034
Upgrade Freq.	*0.056*	0.060	0.053	0.040	0.030	0.023	0.017	0.012	0.009
Downgrade Freq.	*0.007*	0.011	0.010	0.010	0.010	0.010	0.014	0.019	0.025
Consumption		43,876	50,608	56,305	59,361	59,588	56,694	51,591	49,118
Annuity		308	685	940	1,179	1,644	2,335	3,162	3,359

All dollar values deflated by non-housing component of consumer price index: base year 1980. Italicized entries correspond to the 1956-1959 cohort average in the data. They are aged between 30 and 34. Regular font entries correspond to forward simulations averaged across each 5 year interval.

lumpiness of housing stock adjustments pin down the elasticity of substitution parameter. The elasticity determines how long a household can tolerate a non-ideal stock of housing, and thus how often, and by how much, it adjusts its housing stock. Later, we highlight the intuition in our sensitivity analysis of the elasticity parameter. That households leave large amounts of equity at the end of the life cycle allow us to identify the bequest parameters.

2.7 Second Stage

2.7.1 Estimation Details

In this section, we describe the second stage of the estimation procedure (outlined in section 4) that estimates the primitive utility function parameters. First, we forward simulate life cycles under the optimal policy function that we estimated in the first stage. Next, to apply the revealed preference approach, we forward simulate several alternative life cycle paths that follow perturbed, non optimal policy functions.

There are three types of alternative policy functions. The first type draws uniform random perturbations of the multinomial logit model parameters. For the second type, we manually generate variation in moving frequencies and

Table 2.8: 5-year Forward Simulation: Birth Cohort

Birth Cohort	1956-59		1951-55		1946-50		1941-45	
	Data	+5yr Sim	Data	+5yr Sim	Data	+5yr Sim	Data	+5yr
Income	42,372	47,768	44,745	49,591	47,718	51,420	50,234	50,693
Home Size	69,190	75,097	78,573	83,940	88,106	92,263	81,222	83,857
House Value	79,973	89,611	89,159	99,424	98,127	106,775	90,089	96,701
Home Equity	35,791	37,481	45,158	48,898	58,139	61,934	57,105	60,562
Home Equity Home Value	0.435	0.417	0.491	0.496	0.579	0.581	0.654	0.656
Move Frequency	0.063	0.073	0.057	0.063	0.054	0.048	0.046	0.039
Upgrade Frequency	0.056	0.061	0.044	0.052	0.045	0.039	0.033	0.030
Downgrade Frequency	0.007	0.012	0.013	0.010	0.009	0.009	0.013	0.009
Consumption		43,962		45,630		47,832		47,797
Non Housing Wealth		315		261		273		317
Birth Cohort	1936-40		1931-35		1926-30		1921-25	
	Data	+5yr Sim	Data	+5yr Sim	Data	+5yr Sim	Data	+5yr
Income	50,234	52,126	47,092	47,476	45,332	44,514	37,690	36,611
Home Size	80,950	83,060	84,702	86,032	81,093	81,525	71,408	71,399
House Value	89,294	95,397	93,281	98,270	89,650	93,562	76,518	79,982
Home Equity	64,072	67,345	78,239	81,088	78,686	81,120	68,294	70,914
Home Equity& Home Value	0.729	0.729	0.835	0.831	0.0862	0.861	0.902	0.901
Move Frequency	0.019	0.032	0.028	0.026	0.022	0.025	0.027	0.023
Upgrade Frequency	0.013	0.023	0.025	0.017	0.014	0.012	0.011	0.010
Downgrade Frequency	0.007	0.009	0.007	0.010	0.008	0.013	0.016	0.013
Consumption		49,833		46,129		43,647		35,978
Non Housing Wealth		303		539		826		711

All dollar values deflated by non-housing component of consumer price index: base year 1980. "Data" columns represent cohort averages in the data. "+5yr Sim" columns correspond to forward simulations averaged across a 5 year interval. Data columns are blank for consumption and Non Housing Wealth because we do not have data on these variables.

housing shares. Specifically, households move every Nth year, where, for different alternatives, N varies between 3 and 10. In a moving year, households move to a house of size H, which varies across alternatives in a range between 88% and 112% of a household's initial housing to income ratio. This variation helps to identify housing shares, which depend on both the consumption share θ and elasticity τ parameters. It is difficult to separately identify both of these parameters because of a multiplying-a-parameter-by-a-parameter problem. The variation in moving frequencies helps to separately identify the elasticity of substitution. Intuitively, the elasticity determines the length of time that a household would tolerate living in a home that is not optimal. For the third type, we allow households to either deplete or add to home equity in the last five years of the life cycle. In each of those five years, they either convert up to 10% of home equity into consumption or add up to 10% to home equity by reducing consumption. Under the optimal policy they neither add to nor reduce home equity. These alternatives generate variation to identify the retirement/bequest parameter.

2.7.2 Second Stage Results

We estimate the model on a sub-sample of 450 households. We simulate 55 life cycle paths per household. We use 22 alternative policies with 30% manually generating moving frequency and housing share variation, 70% perturbing policy functions coefficients, and, on top of that, 30% varying end of life home equity. Despite the computational advantages of the method, this is the largest sample that our workstation can accommodate. We use a subsampling procedure to calculate standards. In particular, we estimate the parameters on 100 separate subsamples of the same size.[15]

The results are presented in Table 2.9. Notice the elasticity of substitution parameter is significantly greater than zero.

As a validation of the results, it is useful to take the estimated parameter results and calculate implied housing shares in a static model of housing demand.[16] The implied housing share of income is 17.31%, which compares

[15]This procedure does not account for first stage error in the construction of standard errors.

[16]In the static model, households maximize the period utility function subject to a budget constraint on housing rental and consumption spending. It is assumed that the housing

Table 2.9: Utility Parameter Estimates

Parameter	Estimate
Consumption Share θ	0.7653
	(0.0188)
Elasticity of Sub. τ	0.2435
	(0.0470)
Bequest γ	2.5644
	(0.2269)

quite favorably to the 14.3% share in the raw data.

2.8 Simulation Results

2.8.1 Parameterization

In the interest of clarity we briefly summarize the parameters used for the simulations of the structural model. Whenever possible and applicable we use the parameters estimated and employed in the previous sections.

We now use these parameters to simulate the response of consumption and asset accumulation choices to income and house price shocks. We proceed in three steps. First, in order to briefly explain the main mechanisms of the structural model, we display the life cycle profiles in the absence of shocks. Then we subject households to joint income and house price shocks, under the benchmark parameterization. Finally we assess the importance of the size of financial constraints, interest rates, and the elasticity of substitution between consumption and housing services in the utility function. The first two sensitivity analyses are motivated by the recent financial and macroeconomic crisis that has led to changes in the extent to which households can borrow against the value of their home and the mortgage interest rates they are able to obtain. The third exercise highlights that precise estimates of structural utility parameters are quantitatively important in assessing the model-implied consumption and wealth response to income and house price shocks.

rental rate is 7.5% of a home's value.

Table 2.10: Parameterization of Structural Model

Parameter	Value	Interpretation
	Preferences	
β	0.97	Time Discount Factor
γ	2.56	Degree of Altruism
σ	1	Intertemporal Elasticity of Subst.
g	7.5%	Service Flow from Housing Stock
θ	0.77	Consumption Share in Utility Function
τ	0.2435	Elasticity of Sub. between c, h: $\frac{1}{1-\tau}$
	Housing and Financial Markets	
ϕ	3%	Transaction Cost
ξ	20%	Downpayment Requirement
r	1%	Return on Rinancial Assets
r_m	7.24%	Mortgage Interest Rate
π_p	0.95	Persistence of House Price Shock
σ_p	0.03	Std. Dev of House Price Shock
	Labor Income Process	
π_η	0.3	Persistence of Income Shock
σ_η	0.95	Std. Dev. of Income Shock
$\varepsilon_t = \bar{\varepsilon}_t(1+g)^t$	$g = 1.9\%$ $\{\bar{\varepsilon}_t\}$ from Hansen (91)	Life Cycle Labor Income Profile
b	0.5	Pension

2.8.2 The Mechanics of the Model

Figure 2.1 displays life cycle profiles of income shocks[17], η_t, and house prices, p_t, (the exogenous stochastic driving forces of the model), consumption of the composite good, c_t, financial assets, a_{t+1}, real housing assets, h_{t+1}, and a variable we call voluntary equity, $q_{t+1} = a_{t+1} + (1-\xi)p_t h_{t+1}$. As explained in the appendix, it is helpful computationally to use this variable instead of financial wealth in the recursive formulation of the household problem. However, introducing this variable is not only useful for computation, but also helps to interpret the simulation results. Note from the financing constraint (2.2) it follows that $q_{t+1} \geq 0$, and $q_{t+1} = 0$ if the constraint binds. The variable q_{t+1} measures the equity stake in excess of the fraction, $1-\xi$, as required by the constraint. Thus $q_{t+1} = 0$ indicates that in period t households are financing-constrained whereas $q_{t+1} > 0$ indicates a non-binding constraint.

[17] Recall that total labor income is given by $y_t = \eta_t \varepsilon_t$. Thus what we plot here is income net of the deterministic life cycle component which yields profiles that are easier to interpret in the presence of income shocks.

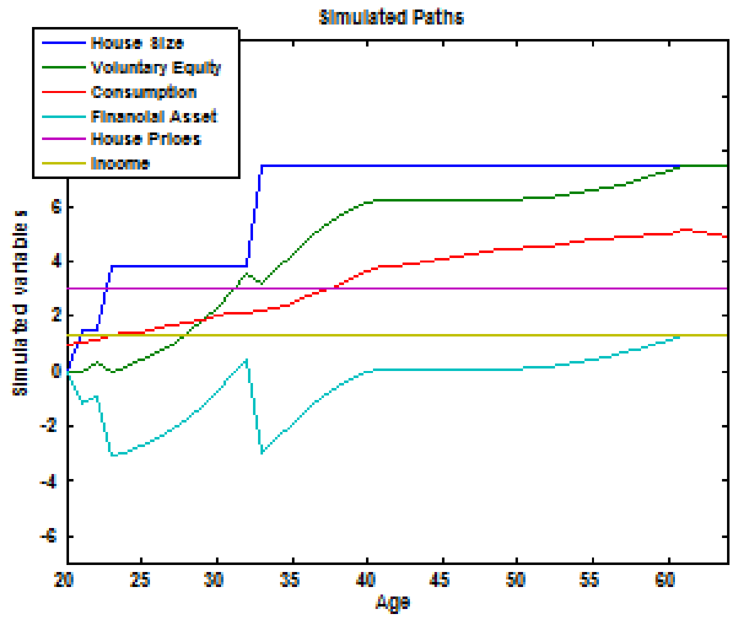

Figure 2.1: Life Cycle Profiles in Absence of Shocks

For the figure (as for all simulations to follow) households start with zero financial and minimal housing assets.[18] Underlying the life cycle paths are the decision rules in the presence of income and house price risk, but for this benchmark simulations the *realizations* of the income and house prices are constant sequences.[19]

From the figure one can clearly see the key housing market frictions in action. Due to the nonconvex adjustment cost, households alter the size of their houses infrequently: they move three times during the first 15 years of their lives, reaching their desired housing size at age 35 and then, absent further shocks to house prices and incomes, stay put. The fact that households move several times prior to reaching their unconstrained optimal size of the house is due to the presence of the financing constraint. As can be seen from

[18]We cannot start households off with zero housing since housing purchased today only generates services tomorrow and the utility function is not well-defined at $h = 0$. Instead we let $h_0 = h_{\min}$, where $h_{\min}$ is the smallest point on the housing grid.

[19]Qualitatively it makes no difference for the dynamics of consumption and housing choices whether the constant prices and incomes are equal to the low or the high realization of the corresponding Markov chains although quantitatively low house prices and high incomes generate slightly larger and more rapid adjustments in the housing stock position.

the time path of voluntary equity, after the first move the household's financing constraint is binding and voluntary equity is zero, $q_t = 0$. After three periods households start to accumulate equity in excess of what is required from the financing constraint, and another move is triggered. As with the first move, the second house the household purchases is still suboptimally small: after the purchase, voluntary equity is again zero and the financing constraint is binding once again. The households accumulate voluntary equity again until the last upward adjustment in the housing stock occurs once more.

This last adjustment is "unconstrained" in the sense that the financing constraint is not binding at this move: the households could have afforded a larger home but finds it suboptimal to choose a larger housing position. In contrast to previous moves, the equity stake exceeds the required minimum share as q_t falls following the move but remains positive.

Absent the financing constraint (but in the presence of the nonconvex adjustment cost) households would have moved only once, catapulting them to the optimal size of the house right away. Thus in order to reproduce the stylized fact documented above that households adjust the size of their home infrequently, but more than once on average over their lifetime, the combination of both frictions in the housing market is crucial. Quantitatively the model reproduces the average time in between housing adjustments and the number of moves during a households' lifetime documented above for US data rather well, at least for the early part of the life cycle.[20]

Having discussed how households behave in the absence of income and house price shocks we are now prepared to explain how these households, within the model, respond to simultaneous declines in income and house prices as observed recently for the US economy.

2.8.3 Simulating an Income and House Price Shock

The exercise we carry out is intended to mimic a sudden, unlikely, but not entirely inconceivable (from the households' perspective) decline in the housing price. At the same time the household receives a negative income shock (which

[20]To the extent that the model abstracts from housing transactions triggered by relocation shocks we would expect the model to understate the frequency with which households move. This is apparent for households in the later stages in their life cycle where the model, absent income and house price shocks, predicts no housing size adjustments at all.

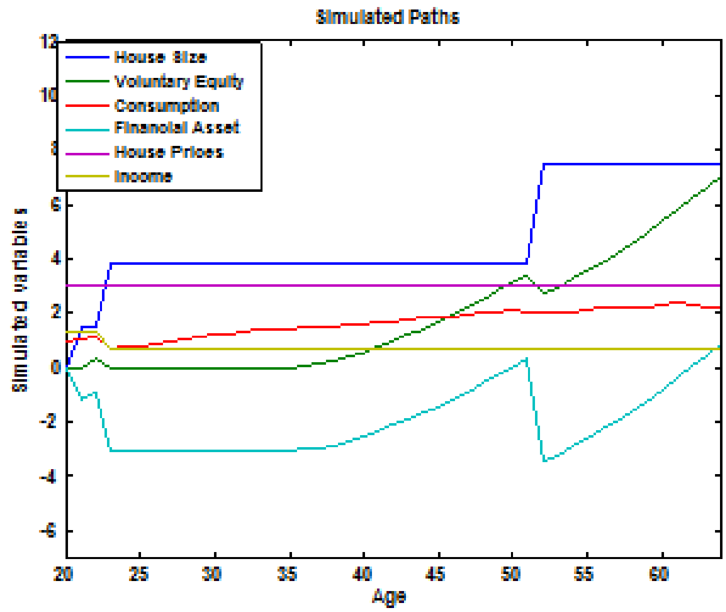

Figure 2.2: Life Cycle Profiles with Negative Income Shocks

by construction is highly, but not perfectly persistent).[21] Given the stylized nature of our structural model, which is necessitated by our desire to estimate it and provide a tight link between theory and estimation, we view our precise quantitative results as less important than to exhibit to what extent model elements and parameters (e.g. adjustment costs, the size of the down payment constraint) affect household responses.

In figure 2.2 we display the life cycle patterns of consumption, housing, and financial wealth prior to and following a negative income shock. We add a concurrent negative house price shock in figure 2.3. The income (and house price) shock are assumed to hit early in a household's life, prior to age 35 when the household, in the absence of shocks, would have acquired its optimal housing size (see figure 2.1).

The key observations we make from both figures is that

- In response to a negative income shock, households adjust nondurable

[21]Of course, since we use the decision rule of households derived under the Markov process for income and prices, households at any period understand that the state of the Markov process can switch with certain probabilty. We merely trace out the dynamic response of households to a particular *realization* of the stochastic process.

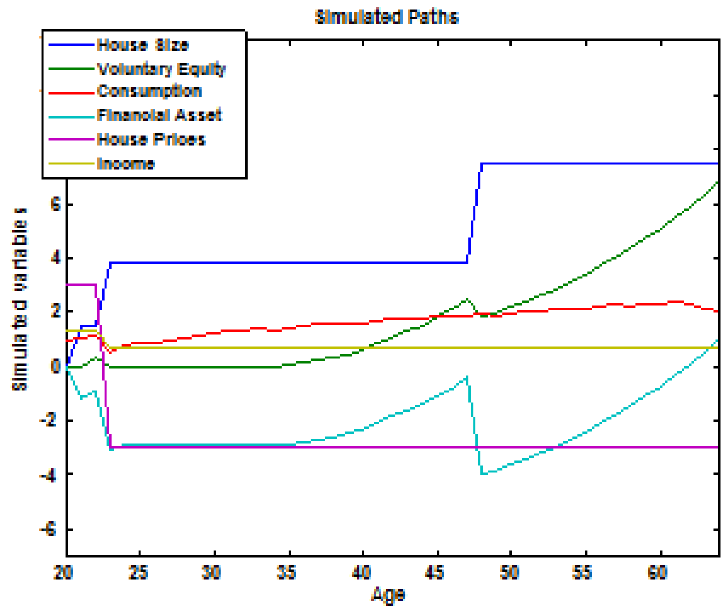

Figure 2.3: Life Cycle Profiles with Negative Income Shocks and Negative Home Price Shock

consumption but do not down size their homes. When the negative income shock hits, households are in the process of *moving up* the housing ladder and find it suboptimal to (temporarily) downsize.

- The income shock does affect the life cycle profile of housing and non-durable consumption significantly. Relative to the benchmark scenario where households reach their desired stock of housing at age 34, persistently low income delays movements up the housing ladder by a good 16 years (note that this assumes that income falls permanently).[22] The low income realizations slow down the accumulation of housing equity and thus delay the purchase of a larger house.

- While the optimal level and general life cycle profile of housing position is not affected by the fall in housing prices, the timing of the housing adjustments following the shocks is. Specifically, because houses are cheaper now households adjust to their optimal level more quickly, by

[22]Of course, given that income and house prices are driven by mean-reverting Markov chains, households perceive a positive probability that prices and income will recover.

about 5 years. This suggests the substitution effect dominates the wealth effect of a home price decline.

- The age/equity position at the time when the shock happens is important in determining the household consumption and saving response. If we subject a middle aged household (of age 35 or older) with a large home equity stake to the same shock, this household finds it optimal to keep its *entire life cycle profile* of housing unaltered and absorb the entire shock by adjusting nondurable consumption and voluntary equity. A household of this age has accumulated sufficiently many assets (relative to the loss in income over the remaining working life) to keep nondurable consumption reasonably smooth despite the unfavorable income realizations. This result confirms our first stage empirical results showing that age and housing equity are important determinants of housing choices.

- As our results without nonconvex adjustment costs below indicate, the sluggish adjustment of the housing position to shocks crucially depends on the imposition of sizeable transaction costs for adjusting the size of houses.

2.9 Sensitivity Analysis

In this section we perform a sensitivity analysis with respect to three parameters. The first set of comparative statics results is motivated by recent policy relevant changes in the US mortgage market. In particular we want to deduce the impact, within our model, of a tightening of credit lines, as observed in the current crisis of the US mortgage market.

Second we demonstrate that the model with nonconvex adjustment costs on housing gives fundamentally different predictions about household responses to income and house price changes than the frictionless benchmark model of consumer durables (as put forward by Mankiw (1982)) in which the adjustment of the stock of housing is completely costless.

Finally, we have spent considerable effort in precisely estimating the preference parameters of our model, in particular the elasticity of substitution between nondurable consumption and housing services in the utility function.

We therefore want to investigate to what extent the results of our model depend on this parameter. To this extent we repeat our simulations with a Cobb-Douglas utility specification which is commonly employed in macroeconomics (see e.g. Fernandez-Villaverde and Krueger (2002) or the discussion in Jeske (2005)).

2.9.1 Relaxing the Financing Constraint

Our model is rich enough to address the question, albeit in stylized form, of what happens to household's housing and nondurable consumption decisions as the financial sector tightens credit lines for mortgages. We now draw out the household response to a simultaneous decline of income and house prices under the assumption that households are required to hold only a $\xi = 10\%$ equity share of the value of their home as opposed to $\xi = 20\%$ as modeled so far.

Comparing figure 2.4 to our benchmark results in figure 2.1 we observe that the main consequence of a relaxed downpayment constraint is that households purchase larger homes early in life, and tend to move less frequently. This finding is due to the fact that a relaxed constraint allows households to more quickly trade up in the housing ladder since they can finance more of the purchase price of the house at young ages where they are severely constrained by the collateral constraint. The response to income and price shocks does not depend strongly on the value of ξ, since at the time adjustment after the shock is optimal the financing constraint was not binding anymore even with the tighter financing constraint $\xi = 20\%$, see figure 2.3.

2.9.2 The Role of Adjustment Costs

The second key friction in the housing market that we model, besides the downpayment constraint, is the presence of sizable transaction costs that need to be borne if (and only if) households change the size of their home. The benchmark model of consumer durables in macroeconomics abstracts from fixed adjustment costs in the market for consumer durables (see e.g. Mankiw (1982)). Our model nests this specification; by setting the adjustment cost

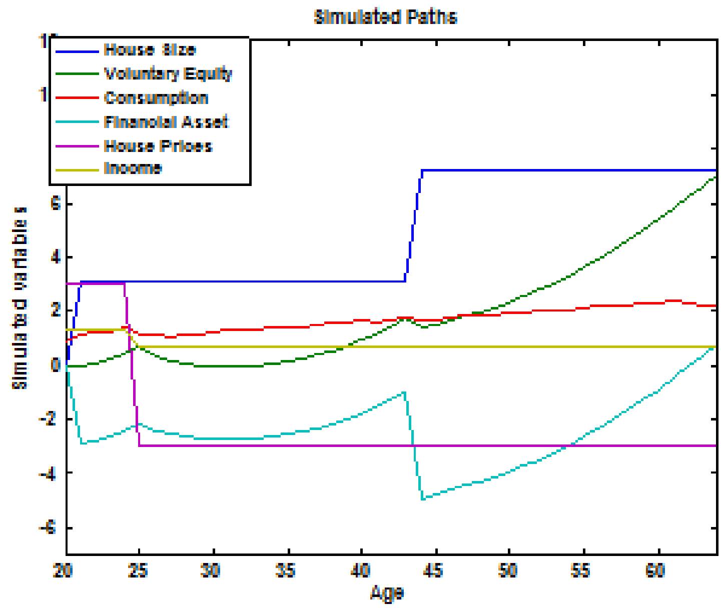

Figure 2.4: Life Cycle Profiles with Relaxed Borrowing Constraints

parameter to $\phi = 0$, we obtain the frictionless model.[23] In figure 2.5 we show, in the absence of income shocks, that this model delivers fundamentally different predictions for the life cycle profiles of consumption, housing, and financial assets.

Without adjustment costs the model loses its implications derived from lumpy housing investment decision rules. Now the build-up of consumer durables in the early stages of a household's live proceeds more gradually.[24] Note that consumption increases over time since, as long as the household is in debt, the relevant interest rate households face is the mortgage interest rate r^m and $\beta(1 + r^m) > 1$.

The model without nonconvex adjustment costs also differs substantially from our benchmark in the way households respond to income and house price shocks. See figure 2.6. Now the stock of housing reacts immediately and significantly to the persistent income decline. Nondurable consumption

[23]Mankiw's model of a representative consumer did not explicitly include downpayment constraints nor did it calibrate the household income process to micro data.

[24]We solve this version of the model with the same discrete state space dynamic programming techniques as the benchmark in order to allow for the results to be as comparable as possible.

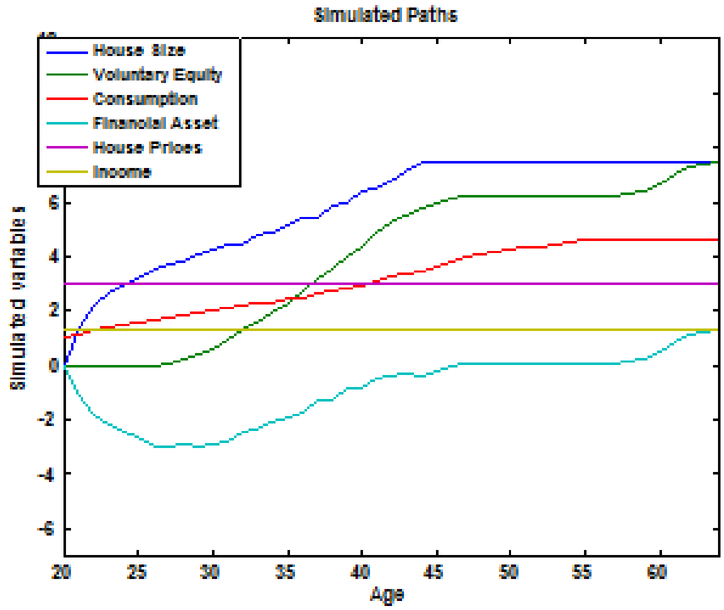

Figure 2.5: Life Cycle Profiles in Absence of Adjustment Costs: No Shocks

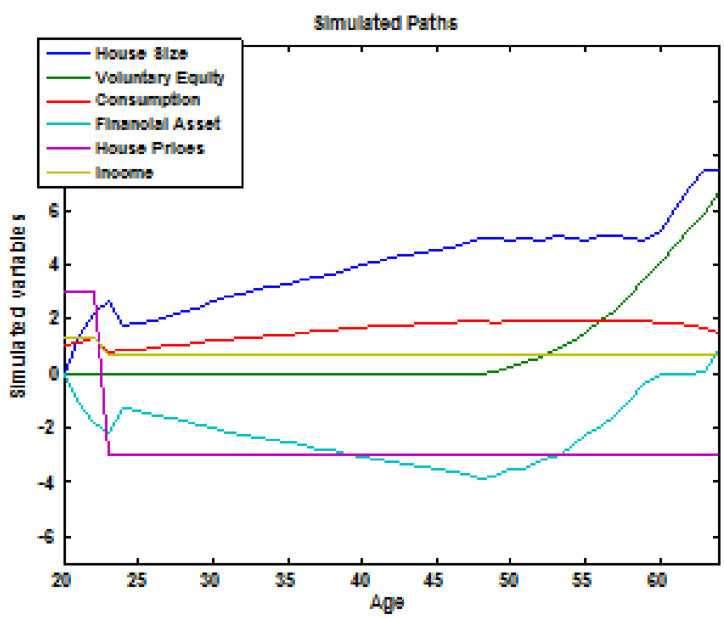

Figure 2.6: Life Cycle Profiles in Absence of Adjustment Costs: Shocks

responds as well.[25]

2.9.3 The Role of the Elasticity of Substitution

Finally we show that the dynamic consumption and asset choices crucially depend on the key preference parameters which we estimated in previous sections. We now document that the optimal choices of households depend on how substitutable nondurable consumption and housing services are in the utility function. Recall that our period utility function was specified as

$$u(c,d) = \frac{1}{\tau} \log\left[\theta c^{\tau} + (1-\theta)(\kappa h)^{\tau}\right]$$

The elasticity of substitution between nondurable consumption and housing services is given by $\epsilon = \frac{1}{1-\tau}$ and was estimated as $\epsilon = \frac{1}{1-0.2435} = 1.32$. Rather than going to the extremes of perfect or no substitutability (both of which are highly implausible given our empirical point estimate and the small standard errors around these estimates) we document how our result changes if one adopts the familiar Cobb-Douglas specification[26] with unit elasticity of substitution; that is $\tau = 0$ and thus $\epsilon = 1$. Conditional on our choice of a unit intertemporal elasticity of substitution $\sigma = 1$ this case has the additional appeal that the utility function becomes separable in nondurable consumption and housing services. That is

$$u(c,d) = \theta \log(c) + (1-\theta)\log(\kappa d).$$

Figure 2.7 shows that, while qualitatively, the life cycle pattern of consumption and housing response are similar to that in the benchmark case, timing differs. Households adjust their housing stock towards the optimal level more quickly (but to about the same extent) with the lower unit elasticity of substitution. Intuitively, the lower the elasticity of substitution, the more costly it is to sustain a suboptimally low housing/nondurables ratio.

[25] Note that the financing constraint remains present in this version of the model. When the shock hits this constraint is binding and thus a decline in the real value of houses leads to a decline in the value of collateral. Households can borrow less, and consumption of nondurables falls significantly.

[26] This parameterization was, among others, adopted by Fernandez-Villaverde and Krueger (2002)

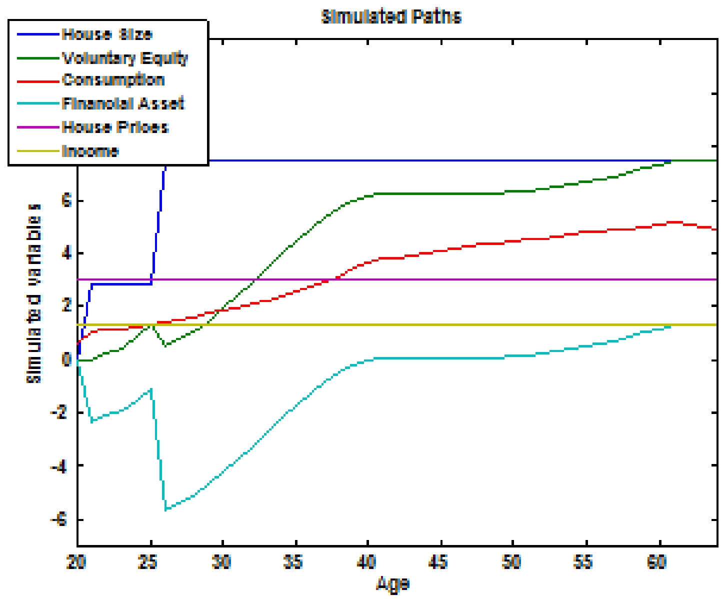

Figure 2.7: Life Cycle Profiles with Unit Elasticity of Substitution

2.10 Conclusion

In this paper, we have constructed and estimated a dynamic structural model of consumption and housing demand with a frictional housing market. We use our estimated model to simulate counterfactual household responses to a set of negative shocks to income, housing prices, and credit constraints that mimic the current calamity US housing market. In our model, we include two key frictions: down payment constraints and non-convex costs to adjust housing stocks. Using household level data from the PSID, we estimate the primitives of the model using the Bajari Benkard Levin (2007) method that is well-suited for estimating such a complicated dynamic model.

Because of the two frictions, nonconvex adjustment costs and financing constraints, households find it optimal to make infrequent housing stocks adjustments. For the benchmark parameterization, households move three times before age 35 as they climb the housing ladder, before reaching their optimal housing size. Negative income shocks at a young age slow down this accumulation process, but do not induce a downgrade to a smaller home. Adding a negative home price shock to the negative income shock causes households to move up the housing ladder more quickly because housing is cheaper. As such, any negative effect on home equity wealth is dominated by a substitution effect. The age at which the shocks hit is important. For older households, having already reached the optimal home size, the negative shocks are absorbed by a reduction in nondurable consumption and home equity. The shocks do not induce a change in housing stock.

We also document the importance of financing constraints, adjustment costs, and the precision of our utility function parameter estimates. All three have significant effects on the frequency of housing stock adjustments and the lumpiness of housing investments. At the center of this adjustment behavior is the ability of households to use home equity to smooth temporary shocks. Future research has to uncover whether introducing further frictions into the housing finance decision that make home equity lines of credit and reverse mortgages less attractive are able to overturn these results.

With regards to a housing market induced recession, our results suggest there will not be a sudden direct impact on the housing market because adjustments occur with such rarity, and household consumption will be minimally affected by the housing bust because they can use home equity to absorb the shock. Over an extended period of time, the fall in home prices has a mixed effect on welfare. Young households, climbing the housing ladder, benefit from lower home prices in that they can climb the housing ladder more quickly, but older households, already owning their optimal sized home, experience a drop in wealth. To more fully assess the housing market details of this recession, it would be interesting examine foreclosure decisions, subprime mortgage products, and the response of households at or near the borrowing constraint.

Bibliography

D. Acemoglu, P. Aghion, R. Griffith, and F. Zilibotti. Vertical Integration and Technology: Theory and Evidence. *Journal of European Economic Association(forthcoming)*, 2009.

D. Arditi and R. Chotibhongs. Issues in Subcontracting Practice. *Journal of Construction Engineering and Management*, 131(8):866–878, 2005.

J. Asker and E. Cantillon. Properties of Scoring Auctions. *RAND Journal of Economics*, 39(1):69–85, 2008.

S. Athey and J. Levin. Information and Competition in US Forest Service Timber Auctions. *Journal of Political Economy*, 109(2), 2001.

P. Bajari and G. Lewis. Procurement Contracting with Time Incentives: Theory and Evidence. *University of Minnesota and Harvard, mimeo*, 2008.

P. Bajari and S. Tadelis. Incentives versus Transaction Costs: A Theory of Procurement Contracts. *The RAND Journal of Economics*, 32(3):387–407, 2001.

P. Bajari, S. Houghton, and S. Tadelis. Bidding for Incomplete Contracts: An Empirical Analysis. *University of Minnesota: mimeo*, 2007.

P. Bajari, R.S. McMillan, and S. Tadelis. Auctions Versus Negotiations in Procurement: An Empirical Analysis. *Journal of Law Economics and Organization*, 25(2):372–399, 2009.

G. Baker and T. Hubbard. Contractibility And Asset Ownership: On-Board Computers and Governance In US Trucking. *Quarterly Journal of Economics*, 119(4):1443–1479, 2004.

G. Baker, R. Gibbons, and K. Murphy. Relational Contracts and the Theory of the Firm. *Quarterly Journal of Economics*, 117(1):39–84, 2002.

S. Bartholomew. *Construction Contracting: Business and Legal Principles.* Prentice Hall, 1998.

D. Beeston. One Statistician's View of Estimating. *Chartered Surveyor, Building and Quantity Surveying Quarterly*, 2(4), 1975.

S. Campo, E. Guerre, I. Perrigne, and Q. Vuong. Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders. 2003.

D. Cattell, P. Bowen, and A. Kaka. Review of Unbalanced Bidding Models in Construction. *Journal of Construction Engineering and Management*, 133: 562–573, 2007.

F. Ciliberto. Does Organizational Form Affect Investment Decisions? *The Journal of Industrial Economics*, 54(1):63–93, 2006.

R. Coase. The Nature of the Firm. *Economica*, 4(16):386–405, 1937.

D. De Silva, T. Dunne, A. Kankanamge, and G. Kosmopoulou. The Impact of Public Information on Bidding in Highway Procurement Auctions. *European Economic Review*, 52(1):150–181, 2008.

S. Forbes and M. Lederman. Adaptation and Vertical Integration in the Airline Industry. *American Economic Review (forthcoming)*, 2009.

S. Forbes and M. Lederman. Does Vertical Integration Affect Firm Performance? Evidence from the Airline Industry. *University of Toronto, mimeo*, 2008.

M. Gates. Aspects of Competitive Bidding. *Annual Report of Connecticut Society of Engineers*, 1959.

R. Gil. Revenue Sharing Distortions and Vertical Integration in the Movie Industry. *Journal of Law, Economics, and Organization(forthcoming)*, 2008.

R. Gil. Make-or-buy in movies: Integration and ex-post renegotiation. *International Journal of Industrial Organization*, 25(4):643–655, 2007.

R. Gil and J. Marion. The Role of Repeated Interactions, Informal Agreements and Relational [Sub]Contracting: Evidence from California Highway Procurement Auctions. *UCSB: mimeo*, 2009.

L. Hansen and K. Singleton. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica*, 50(5):1269–1286, 1982.

G. Harrison and E. Rutström. Risk Aversion in the Laboratory. *Risk Aversion in Experiments*, 2008.

O. Hart and J. Moore. Property Rights and the Nature of the Firm. *Journal of Political Economy*, 98(6):1119, 1990.

O. Hart and J. Moore. Contracts as Reference Points. *Quarterly Journal of Economics*, 123(1):1–48, 2008.

J. Heckman. Sample Selection Bias as a Specification Error. *Econometrica*, 47(1):153–161, 1979.

J. Hinze. *Construction Contracts*. McGraw-Hill, 1993.

B. Holmstrom and P. Milgrom. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics, and Organization*, 7:24, 1991.

B. Holmstrom and J. Roberts. The Boundaries of the Firm Revisited. *Journal of Economic Perspectives*, 12(4):73–94, 1998.

H. Hong and M. Shum. Increasing Competition and the Winner's Curse: Evidence from Procurement. *Review of Economic Studies*, 69(4):871–898, 2002.

T. Hubbard. Viewpoint: Empirical Research on Firms' Boundaries. *Canadian Journal of Economics/Revue canadienne d'economique*, 41(2):341–359, 2008.

M. Jofre-Bonet and M. Pesendorfer. Estimation of a Dynamic Auction Game. *Econometrica*, 71(5):1443–1489, 2003.

P. Klein. The Make-or-Buy Decision: Lessons from Empirical Studies. *Handbook of New Institutional Economics, Amsterdam: Springer*, 2005.

E. Krasnokutskaya. Identification and Estimation in Highway Procurement Auctions Under Unobserved Auction Heterogeneity. *Review of Economic Studies (forthcoming)*, 2009.

F. Lafontaine and M. Slade. Vertical Integration and Firm Boundaries: The Evidence. *Journal of Economic Literature*, 45:629–685, 2007.

J. Levin and S. Tadelis. Contracting for Government Services: Theory and Evidence from U.S. Cities. *Journal of Industrial Economics, forthcoming*, 2008.

P. Levin. *Construction Contract Claims, Changes, & Dispute Resolution*. ASCE Press, 2 edition, 1998.

J. Marion. How Costly is Affirmative Action? Government Contracting and Californias Proposition 209. *Review of Economics and Statistics: forthcoming*, 2008.

H. Markowitz. Portfolio Selection. *Journal of Finance*, 7(1):77–91, 1952.

S. Masten. The Organization of Production: Evidence from the Aerospace Industry. *Journal of Law and Economics*, 27(2):403–417, 1984.

S. Masten, J. Meehan Jr, and E. Snyder. The Costs of Organization. *Journal of Law, Economics, & Organization*, 7(1):1–25, 1991.

K. Monteverde and D.J. Teece. Supplier Switching Costs and Vertical Integration in the Automobile Industry. *Bell Journal of Economics*, 13(1):206–213, 1982.

S. Novak and S. Stern. How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry. *Management Science*, 54:12, 2008.

R. Porter and J. Zona. Detection of Bid Rigging in Procurement Auctions. *Journal of Political Economy*, 101(3):518, 1993.

A. Roy. Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3(2):135–146, 1951.

C. Semple, F. Hartman, and G. Jergeas. Construction Claims and Disputes: Causes and Cost/Time Overruns. *Journal of Construction Engineering and Management*, 120(4):785–795, 1994.

J. Sweet. *Legal Aspects of Architecture, Engineering and the Construction Process*. Brooks/Cole Publishing, 6 edition, 2004.

M. Whinston. On the Transaction Cost Determinants of Vertical Integration. *The Journal of Law, Economics, and Organization*, 19:1–23, 2003.

O. Williamson. *Markets and Hierarchies, Analysis and Antitrust Implications: a Study in the Economics of Internal Organization*. Free Press, New York, 1975.

O. Williamson. *The Economic Institutions of Capitalism*. Free Press, New York, 1985.

Attanasio, O. (2000), "Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S, s) Rules for Automobile Purchases," *Review of Economic Studies,* 67, 667-696.

Bajari, P., L. Benkard and J. Krainer (2004), "House Prices and Consumer Welfare," *NBER Working Paper* 9782.

Bajari, P., Benkard, L., Levin, J. (2007) "Estimating Dynamic Models of Imperfect Competition," *Econometrica,* 75, 1331–1370.

Benjamin, J., P. Chinloy and D. Jud (2004), "Real Estate Versus Financial Wealth in Consumption, *Journal of Real Estate Finance and Economics*, 29, 341-351.

Campbell, J. and J. Cocco (2007), "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics,* 54, 591-621.

Carroll, S. and W. Li (2008), "The Homeownership Experience of Households in Bankruptcy," *Working paper*, Federal Reserve Bank of Philadelphia.

Case, K. J. Quigley and R. Shiller (2005), "Comparing Wealth Effects: The Stock Market Versus the Housing Market," *Advances in Macroeconomics*, 5, 1-34.

Chambers, M., C. Garriga and D. Schlagenhauf (2007), "The Tax Treatment of Homeowners and Landlords," *Working paper*, St. Louis FED.

Chambers, M., C. Garriga and D. Schlagenhauf (2009a), "Accounting for Changes in the Homeownership Rate," *International Economic Review*, 50, 677-726.

Chambers, M., C. Garriga and D. Schlagenhauf (2009b) "The Loan Structure and Housing Tenure Decisions in an Equilibrium Model of Mortgage Choice," *Review of Economic Dynamics*, 12, 444-468.

Diaz, A. and M. Luengo-Prado (2008), "On the User Cost and Homeownership," *Review of Economic Dynamics*, 11, 584-613.

Diaz, A. and M. Luengo-Prado (2009), "Durable Goods and the Wealth Distribution, " forthcoming, *International Economic Review*

Flavin, M. and T. Yamashita (2002), "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review*, 92, 345-362.

Fernandez-Villaverde, J. and D. Krueger (2002), "Consumption over the Life Cycle: How Inportant are Consumer Durables?," *Working paper*, University of Pennsylvania.

Garriga, C. and D. Schlagenhauf (2009), "Home Equity, Foreclosures, and Bailouts," *Working paper*, Federal Reserve Bank of St. Louis.

Hall, G. Rust, J. (2003) "Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices" *Working paper* Brandeis University.

Han, L. (2008) "The Effects of Price Uncertainty on Housing Demand: Empirical Evidence from U.S. Markets," *Working paper*, University of Toronto.

Hintermaier, T. and W. Koeniger (2009), "Debt Portfolios," *Working paper*, Queen Mary College.

Hurst, E. and F. Stafford (2004), "Home is where the Equity is: Mortgage Refinancing and Household Consumption," *Journal of Money, Credit and Banking*, 36, 985-1014.

Iacoviello, M. and M. Pavan (2009), "Housing and Debt over the Life Cycle and over the Business Cycle," *Working paper*, Boston College.

Jeske, K. (2005), "Macroeconomic Models with Heterogeneous Agents and Housing," *Federal Reserve Bank of Atlanta Economic Review*, First Quarter, 26-34.

Jeske, K and D. Krueger (2005), "Housing and the Macroeconomy: The Role of Implicit Guarantees for Government Sponsored Enterprises," *Working paper*, Federal Reserve Bank of Atlanta

Kiyotaki, N., A. Michaelides, and K. Nikolov (2008), "Winners and Losers in the Housing Market," *Working paper*, London School of Economics.

Li, Wenli and R. Yao (2007), "The Life Cycle Effects of House Price Changes," *Journal of Money, Credit and Banking*, 39, 1375-1409.

Li, W., H. Liu and R. Yao (2009), "Housing over Time and over the Life Cycle: A Structural Estimation," *Working paper*, Federal Reserve Bank of Philadelphia (2009).

Luengo-Prado, M. (2006), "Durables, Nondurables, Down Payments and Consumption Excesses," *Journal of Monetary Economics*, 5

Lustig, H. and S. Van Nieuwerburgh (2006), "Can Housing Collateral Explain Long-Run Swings in Asset Returns?" *Working paper*, UCLA.

Mankiw, G. (1982). "Hall's Consumption Hypothesis and Durable Goods". *Journal of Monetary Economics*, 10, 417-425.

Ortalo-Magne, F. and Sven Rady (2006), "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *Review of Economic Studies*, 73, 459-485.

Rust, J. (1987) “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999-1033.

Ryan, S. (2006) “The Cost of Environmental Regulation In a Concentrated Industry,” *working paper* MIT.

Snider, C. (2008) “Predatory Incentives and Predatory Policy: The American Airlines Case”, *Working paper* UCLA.

Sweeting, A. (2007) “Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry”, *Working paper* NBER 13522.

Piazessi, M., M. Schneider and S. Tuzel (2007), “Housing, Consumption, and Asset Pricing,” *Journal of Financial Economics,* 83, 531-569.

Yang, F. (2009) Consumption Over the Life-cycle: How different is Housing?” *Review of Economic Dynamics*, 12, 423-443.

Yao, R. and H. Zhang (2005) “Optimal Consumption and Portfolio Choices with risky Housing and Borrowing Constraints,” *Review of Financial Studies*, 18, 197-239.

Appendix A

Derivation of Unit Price Bidding Equations

Omit bidder subscripts and make a change of variables. Define $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{c}$, such that positive entries for $\tilde{\mathbf{b}}$ correspond to bidding above cost (like taking a long position in an asset); negative entries, bidding below cost (a short position). According to the indeterminacy of lump sum bids, group the lump sum bids and costs as: $b_l = \sum_{t=1}^{L} b_t$ and $c_l = \sum_{t=1}^{L} c_t$. With the change of variables and CARA utility function, rewrite the first order conditions for variable quantity task t as:

$$E\left[\frac{1}{\exp((s - c_l - \mathbf{c}\cdot\mathbf{q}^e + \tilde{\mathbf{b}}\cdot(\mathbf{q}^a - \mathbf{q}^e))\gamma)}\right] q_t^e = E\left[\frac{1}{\exp((s - c_l - \mathbf{c}\cdot\mathbf{q}^e + \tilde{\mathbf{b}}\cdot(\mathbf{q}^a - \mathbf{q}^e))\gamma)} q_t^a\right] \quad \text{(A.1)}$$

The denominator term uses the equality on the constraint that bids sum up to the score to substitute out the lump sum bid, b_l. Now boldface vectors have length $T - L$. To linearize the function $u'(\pi(\tilde{\mathbf{b}}))$ take a first order Taylor series expansion around bids at cost, $\tilde{\mathbf{b}} = \mathbf{0}$

$$\begin{aligned} \frac{1}{\exp(\pi(\tilde{b})\gamma)} &\approx \frac{1}{\exp(\pi(\mathbf{0})\gamma)} - \frac{\gamma}{\exp(\pi(\mathbf{0})\gamma)}(\mathbf{q}^a - \mathbf{q}^e)'\tilde{\mathbf{b}} \\ \pi(\mathbf{0}) &= s - c_l - \mathbf{c}\cdot\mathbf{q}^e \end{aligned}$$

Substituting this linearized expression back into the first order conditions and performing further algebraic manipulation yields,

$$\frac{1}{\gamma}E[q_t^a - q_t^e] = E\left[(q_t^a - q_t^e)(\mathbf{q}^a - \mathbf{q}^e)'\right]\tilde{\mathbf{b}} \tag{A.2}$$

Then stacking by tasks and substituting $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{c}$ gives the unit price bidding equation

$$\mathbf{b}_i = \mathbf{c}_i + \frac{1}{\gamma}E\left[(\mathbf{q}^a - \mathbf{q}^e)(\mathbf{q}^a - \mathbf{q}^e)'\right]^{-1}(E[\mathbf{q}^a] - \mathbf{q}^e) \tag{A.3}$$

Appendix B

Computational Appendix

B.0.1 Recursive Formulation of the Problem

The model in recursive formulation can be written as[1]

$$\begin{aligned}
V(\eta, a, h, p, t) &= \max_{c,a',h'} \left\{ u(c, g(h)) + \beta EV(\eta', a', h', p', t+1) \right\} \\
& \quad \text{s.t.} \\
c &\geq 0 \\
h' &\geq 0 \\
a' &\geq -(1-\xi)ph' \\
c + a' + ph' + p\phi(h', h) &= \eta\varepsilon_t + (1+r)a + ph
\end{aligned}$$

B.0.2 Transformation of the State Space

One problem of the formulation of the problem above is that the constraint set for (a', h') is not rectangular, and that the constraint on h' depends on a', which is itself a choice variable. This problem can be overcome by defining a new variable, voluntary equity, q', as

$$q' = a' + (1-\xi)ph'$$

(see Diaz and Luengo-Prado (2008, 2009)). Note that this definition implies

$$q = a + (1-\xi)p_{-1}h$$

[1] This is the recursive problem, conditional on not having retired yet. The problem for retired households is similar and hence omiteed.

where p_{-1} is the price of housing in the previous period, which now becomes an additional state variable.. The recursive problem of the household with this transformation of variables now reads as

$$\begin{aligned}
v(\eta, q, h, p_{-1}, p, t) &= \max_{c,q',h'} \left\{u(c, g(d)) + \beta E v(\eta', q', h', p, p', t+1)\right\} \\
&\quad \text{s.t.} \\
c', q', h' &\geq 0 \\
c + q' + p\xi h' + p\phi(h', h) &= \eta\varepsilon_t + (1 + r(q, h, p_{-1}))q \\
&\quad + \left[(1-\delta)p - (1 + r(q, h, p_{-1}))(1-\xi)p_{-1}\right] h
\end{aligned}$$

that is, we traded off an additional state variable p_{-1} against now having a rectangular constraint set for the choice variables (c, q', h'). This is the recursive formulation of the model we compute. Clearly the consistency condition $p'_{-1} = p$ has to hold, that is, tomorrow's past housing price p'_{-1}has to equal today's price. The interest rate function now reads as

$$r(q, h, p_{-1}) = \begin{cases} r & \text{if } q - (1-\xi)p_{-1}h \geq 0 \\ r^m & \text{if } q - (1-\xi)p_{-1}h < 0 \end{cases}$$

Since we use an adjustment cost that is nonconvex, the household decision problem is not a convex programming problem and numerical approaches the require differentiability of the value function cannot be applied. Therefore we use discrete state space dynamic programming techniques to solve the problem. In particular, we discretize the state space for (q, h) into a finite rectangular grid (the income and house price process process is already a finite state Markov chain by assumption) and maximize the objective function by searching, for each (q, h) over the finite grid of admissible choices (q', h'). The consumption choice is implied by the budget constraint.

Given a terminal value function (given by the bequest function we can iterate backward in age of the household to solve for the age-dependent optimal policy (and value) functions. Once we have computed these, simulated life cycle patterns of consumption, housing and financial wealth can be generated for any sequence of house price and income shocks.