

Bargaining and Welfare : A Dynamic Structural Analysis*

Daniel E. Keniston[†]
Massachusetts Institute of Technology
January 20, 2011

Abstract

Bargaining for retail goods is common in developing countries, but rare in the developed world. The welfare implications of this difference are theoretically ambiguous—if bargaining is a low cost form of price discrimination, it may lead to greater trade and welfare and even approximate the optimal incentive compatible outcome. However, if bargaining imposes large utility costs on the participants, then a fixed price may be preferable. I develop the tools to resolve this question, specifying a mechanism design problem adapted to the context of bargaining, and developing a dynamic structural estimation technique to infer the structural parameters of the market. I then apply these techniques to the market for local autorickshaw transportation in Jaipur, India, using data I collected over 2008-2009.

*Preliminary draft: Do not quote. I am grateful for helpful conversations and suggestions from Abhijit Banerjee, Esther Dufo, Panle Jia, Michael Peters, Stephen Ryan, Simone Schaner, and Robert Townsend. This paper has also benefited from comments during the MIT Development and IO lunch seminars

[†]Department of Economics, Massachusetts Institute of Technology, E53-387, Cambridge, MA 02139. Email: keniston@mit.edu

1 Introduction

A defining feature of developing societies is the informality of their markets. Individuals, firms, and governments frequently lack the ability or the incentives to enforce contracts and commit to fixed trading rules. In the name of development, governments and aid agencies have often pushed to formalize these markets, bringing them closer to those in the developed world. Yet little is known of the actual costs or benefits of informal market institutions. This paper examines a particularly widespread case: bargaining for retail goods. It develops the theoretical and empirical tools to compare bargaining to a counterfactual formal market mechanism—fixed prices—and applies these techniques to the specific case of the market for autorickshaws in Jaipur. More broadly, the choice to promote fixed prices versus bargaining is one that is faced by governments and firms in many contexts both in and out of emerging markets. The techniques developed in this paper allow us to analyze and inform this decision, and open the path to a broader study of the shift to more formal market structures in the process of development.

Bargaining is a ubiquitous feature of markets in the developing world, yet little is known about its implications for the efficient functioning of markets and social welfare. Is bargaining a drag on the economies of developing countries, imposing high transactions costs and reducing trade? Or is it an efficient means of bilateral price discrimination, allowing gains from trade to occur between buyers and sellers who would never have participated in a fixed-cost market? Economic theory suggests that the welfare effects are a priori ambiguous: If the right equilibrium is played, bargaining has the potential to yield as much surplus as the best incentive compatible mechanism. Ausubel, Cramton, and Deneckere (2002) show that, for buyer and seller distributions with monotone hazard rates, any ex-ante efficient division of the gains from trade is implementable in an alternating offer bargaining game. However, there is no guarantee that the equilibrium actually being played will yield an efficient outcome. If bargaining imposes high costs and traders are homogeneous, then a fixed price may generate greater overall surplus. Furthermore, the selection of the equilibrium mechanism may be influenced by factors not determined by the market such as the penalties to deviation. For instance, taxis use fixed meter rates in American cities partly because policy makers have chosen and enforced a fixed price, thus making any attempt to deviate by either passengers or drivers costly in terms of transactions costs and possible fines. In order to evaluate the costs and benefits of mandating a fixed price (or any other mechanism), the policy maker must know the structural parameters of the market: the costs of bargaining, and the distribution of valuations of the players.

This paper carries out the estimation of these parameters and subsequent policy evaluation, including the first structural analysis of a two-sided incomplete information bargaining model in the literature. I use data collected on series of individual bargaining offers and bargainers' decisions to purchase a ride or walk away from the autorickshaw (local transportation) market in Jaipur, India. The data combines an audit survey with an experimental approach. Some observations were collected from bargaining sessions using surveyors who were the full residual claimants of any financial gains from bargaining and could spend as much time as they wanted searching for a good price. These observations serve to identify the range and probability of the counter-offers that bargainers anticipate after each offer. Another set of observations were collected in which surveyors posing as potential passengers gave randomized counter-offers to the drivers' offers. These observations expand the range of the observed data, and provide a useful test that drivers' counter-offers are not influenced by unobserved signals. This data collection strategy allowed for a much greater degree of control and homogeneity on unobservables than would have been possible using purely observational data.

Using the set of bargaining sequences collected in these data, I back out the player's valuations

and bargaining disutility, and consider the counterfactual policy of giving drivers and passengers the option to pay a fixed price for an autorickshaw ride instead of entering the bargaining market. This models the creation of a pre-paid taxi stand- an institution that already exists at certain airports or taxi stands in India. Choosing a per-kilometer meter rate that would set drivers ex-ante indifferent between taking passengers by bargaining or fixed rates, I find that introducing the option of the fixed price raises welfare by 28% due to high value, high bargaining disutility passengers switching to the fixed price mechanism. However, even at the optimal price for passengers (holding drivers indifferent) 63% of the buyers in my sample would remain in the bargaining market, suggesting that bargaining may still hold substantial benefits for large segments of the population.

In contrast to the auctions literature, where structural estimation of valuations is common, structural models of bargaining are rarely specified for two major reasons: First, the data required on individual bargaining offers are not frequently collected. Second, even if the data were available, the recovery of structural parameters from bargaining offers poses several challenges: Bargaining is an inherently dynamic process—for instance a low-ball early offer by the buyer will affect the course of the whole bargaining interaction. Any estimation must incorporate these dynamic considerations, and techniques for the analysis of dynamic games have become available only recently (Aguirregabiria and Mira 2010). Furthermore, the players’ valuations of the good, and their expectations of their opponent’s valuations, are crucial to the outcome of the bargaining but unobserved to the econometrician. The presence of these unobserved state variables makes many of the standard methods of estimating dynamic games inapplicable. Finally, unlike similar auctions problems, there exists no canonical closed form bargaining model that incorporates a realistic set of market features, and those that do exist and contain sufficient richness to approximate real-world bargaining often feature multiple equilibria.

To address these challenges I employ an estimation strategy motivated by Bajari, Benkard, and Levin (2007) and Pakes, Ostrovsky, and Berry (2007)’s approaches to dynamic games. Without solving for a game theoretical equilibrium of the bargaining game, I use the data on bargaining sequences to estimate players’ beliefs on their opponents’ strategies (e.g. the probability of accepting an offer, or making any of a discrete set of counter-offers, or exiting) at every bargaining round. Using these action probabilities, I can, through backwards induction, calculate a bargainer’s expected payoffs for any action conditional on the agent’s true cost or value of the good, and his or her costs of bargaining. These payoffs in turn imply the player’s best responses, and yield predictions for the actions taken at each state of the game. Finally, for each buyer (seller) observed in the data, I estimate the set of valuations (costs) and bargaining disutility that rationalize the observed series of offers or exit/accept decisions actually taken by that individual. The outcome is a fully non-parametric estimate of the players’ distributions of values (costs) and bargaining disutility.

The source of the data, and the context of the models presented in this paper is the market for intracity transportation by autorickshaw (a type of mini-taxi) in India. While the market and bargaining game are discussed in further detail in section 5, it may be useful to summarize the procedure of bargaining for an autorickshaw ride in order to provide a motivating example for the theory and empirics to follow. Potential passengers make an initial decision of whether to take an autorickshaw ride or not. If they elect to participate in the market, they stand by the side of the road and hail passing autorickshaws. The driver stops and demands where the passenger would like to go. Upon learning the destination, the driver makes an initial offer, and buyer and seller exchange offers at approximately 20 second intervals. The bargaining terminates when one party accepts the other’s last offer, or exits the bargaining¹. If either party exits from the negotiations,

¹In the data, passengers who are residual claimants unilaterally exit the bargaining only 11% of the time, perhaps because the driver can usually catch up to a passenger in order to make a final offer while the passenger is walking

the passenger waits to hail another autorickshaw and the process repeats. The trips in question are short (2-8 km.), with a cost of between 30 and 80 rupees ($\approx$ \$.75-2.00). Bargaining interactions are fairly quick—from 2 to 9 rounds, and passengers rarely encounter the same driver twice. All participants have extensive experience in bargaining.

Section 2 summarizes the various literatures upon which this paper draws for its theoretical and empirical components.

Section 3 specifies the economy in which the bargaining and trade of goods takes place. I consider a market in which buyers are randomly matched with sellers whose valuations they do not know. Whether trade occurs, and at what cost to the market participants, depends on the bilateral mechanism used to allocate the goods; I model the payoffs from bargaining or fixed prices.

Section 4 describes the estimation of the structural parameters necessary to evaluate the welfare implications of the different mechanisms.

Section 5 presents the data and results of the estimation. Reduced form specifications, although confounded by various types of selection bias, provide suggestive evidence regarding the bargaining game and the trade-offs faced by the players when considering making an offer. The structural estimation yield bounds on the distributions of valuations for buyers and sellers.

Finally, section 6 combines the results of the theory from section 3 with the parameters of section 5 and evaluates the welfare consequences of bargaining versus fixed prices. Initially I focus on the supply side of the market, estimating the fraction of drivers who would be willing to provide rides at the official, government sanctioned meter rate and at another rate proposed by some of the drivers themselves. Then, using the valuations estimated for the surveyors to represent the buyer population, I investigate the counterfactual policy of introducing the option to purchase an autorickshaw ride for a fixed price without bargaining, and derive its welfare implications.

2 Literature Review

This paper builds on a wide variety of existing literature: the theoretical framework is drawn from the dynamic mechanism design literature, combined with a specification of the extensive form and payoff functions from the game theoretical literature on bargaining, and an estimation technique adapted from the dynamic structural games literature. Here I review the sources from which I have adapted specifications and results, and finally highlight the areas in which my approach differs from the few previous structural approaches to bargaining.

Two subfields of the market design literature examine the welfare implications of alternative market mechanisms in economic environments with repeated bilateral trade. The first, typified by papers by Rubinstein and Wolinsky (1985), Fraja and Sakovics (2001), and Satterthwaite and Shneyerov (2007), focuses on the relationship between equilibrium prices and amounts of trade and the transactions costs for the traders. These papers consider an economic environment similar to mine, in which traders have imperfect information on each others' valuations, incur some utility cost from each interaction, and participate in the market over several periods. However, their focus is on the conditions under which decentralized markets converge to a Walrasian single price and often abstracts away the details of the specific bilateral mechanism. A second literature, originating in the classic paper by Myerson and Satterthwaite (1983) focuses on the efficiency, or lack thereof, of specific mechanisms for bilateral trade. In more recent work, Athey and Miller (2007) consider an economy similar to that examined in this paper, where buyers and sellers draw new valuations each period, and continue to participate in the market once they have completed a trade. These papers

away, whereas passengers cannot make a final offer to the driver after he has driven away. See table 3 for complete data on exit/accept probabilities at each bargaining round.

resemble this study in their focus on evaluating the efficiency of alternative market mechanisms, with the distinction that the mechanisms considered in this paper, bargaining and fixed prices, are chosen for their empirical relevance rather than their theoretical efficiency.

The bargaining mechanism itself is the subject of a vast literature. Starting with Rubinstein's canonical model of alternating offers bargaining with full information (Rubinstein 1982), a substantial body of work on bargaining with two-sided imperfect information has developed, most recently reviewed by Ausubel, Cramton, and Deneckere (2002). Unfortunately, this literature has not generated a canonical model comparable to Rubinstein's, perhaps due to the multiplicity of equilibria in models with sufficient detail to be realistic. Another ambiguity in the imperfect information bargaining literature is how to model the uncertainty in players' payoff functions. While the majority of research assumes that players are uncertain as to each other's valuations for the good, a separate stream of the literature (Rubinstein 1985, Bikhchandani 1992) assumes that the uncertainty pertains to the opponent's discount factor, while yet another models the possibility that the opponent might be an irrational type (Abreu and Gul 2000). Finally, although most research has focused on a multiplicative discount factor, bargaining could also be modeled as having a fixed per-round transactions cost as in Rubinstein (1982). Perry (1986) shows that this formulation will lead to immediate agreement unless players are uncertain of their opponent's fixed bargaining cost.

Despite the lack of consensus on a realistic model of incomplete information bargaining, two main strands of literature on empirical bargaining have emerged, one concentrating on bargaining experiments in the lab and another (substantially smaller) analyzing the outcomes of real-world bargaining. Lab experiments on bargaining, as surveyed by Roth (1995) set up the environment assumed by a specific model of bargaining, and test whether the outcomes of the bargaining performed by study participants match the predictions of the model. The results of studies in which players have full or partial knowledge of each other's valuations have generally rejected the quantitative and often also the qualitative implications of game theoretic models of bargaining, often in favor of fairness norms such as a 50-50 division of the gains from trade. Fairness effects should be much less important in experiments with two-sided incomplete information, however these are almost nonexistent. The only study, Yan and Lu (2008), tests the model of Abreu and Gul (2000) and finds preliminary support for it. As many authors have noted, bargaining with incomplete information is an area under-researched in the lab.

Empirical studies of bargaining are also remarkably rare. In contrast to the highly model-driven experimental literature, the majority of studies of real world bargaining compare reduced form outcomes of empirical bargaining with the broad implications of game theoretic bargaining. Several papers have examined bargaining in retail transactions. Morton, Zettelmeyer, and Silva-Risso (2004) compare survey data on buyer characteristics and price/purchase outcomes from car dealerships, and find that buyer characteristics are correlated with outcomes in ways predicted by bargaining theory, for instance that better informed, more patient buyers pay less. Ayers and Siegelman (1995) use an audit survey methodology similar to my own to test for discrimination against blacks and women in bargaining for cars. The only work to examine the actual series of bargaining offers as well as the only study of bargaining in developing countries of which I am aware is by Jaleta and Gardebroek (2007), who examine the effect of buyer and seller characteristics on the spread between buyer and seller initial offers and the amount by which the seller is willing to decrease his initial offer.

The work most similar to this study in its approach is the series of papers by Sieg (2000), Watanabe (2004), and Merlo, Ortalo-Magne, and Rust (2008) that use bargaining data to estimate structural models of medical malpractice legislation and the housing market. All three authors explicitly model the bargaining process as a game with a specific, unique equilibrium, and then use a maximum-likelihood or simulated moment-based approach to infer structural parameters of

the particular model specified. In particular, Merlo, Ortalo-Magne, and Rust (2008) consider a strategy similar to mine of using bidding functions estimated from the data, but ultimately prefer a more structural model of bidding due to endogeneity concerns. Endogeneity is less of a concern in my estimation due to the presence of data generated by randomized surveyor bids, which can be used to test for the endogeneity of driver's bids with respect to the passenger's previous offer.

This study differs from previous work in several ways. First, while my approach is structural, I do not claim to know which equilibrium of a specific game theoretic model of bargaining is being played in the market, nor do I explicitly specify the process of Bayesian updating that the players carry out during the game. Instead, I estimate the best response probabilities from whatever equilibrium is being played in the data, and use these estimated probabilities, the structural payoff functions and the extensive form of the bargaining game to infer the structural parameters of the bargaining model. Second, while other studies have examined the effect of observable signals (such as race or gender) on bargaining, I focus on a market in which, arguably, the most important aspects of preference heterogeneity are unobservable. Third, I maintain the assumption that the traders are fully rational. While there is strong evidence to the contrary from lab experiments, it is unclear to what degree this would apply to real-world bargaining², and the assumption of rationality seems to be an appropriate starting place. In future work I plan to examine various behavioral or rule-of-thumb strategies using the same data.

3 Bargaining vs. Fixed Prices: Theory

The autorickshaw transport market falls into a class of markets in which professional sellers sell multiple types of goods to casual buyers interested in purchasing only one specific item. Other examples include a shopkeeper selling many kinds of products to consumers who are each shopping for one special product, or a consultant providing different services to different clients. Each buyer (henceforth referred to using the feminine pronoun) has a permanent valuation v for the good she wants to purchase, and this valuation is known to her before she encounters a seller. Sellers, because they sell a different good to each buyer, have a buyer-specific cost c , that is only known to the seller once he meets the buyer. Thus buyers keep the same valuations across negotiations with different sellers, while sellers draw new valuations for each buyer they encounter.

Buyers and sellers have different future outcomes if trade occurs. Sellers are professional traders, and always return to the market and begin negotiations with a fresh buyer at the conclusion of a successful or unsuccessful trade. Buyers, on the other hand, are temporary participants in the market. Upon completion of a successful sale, they exit the market permanently. However, if negotiations with a given seller fail to result in trade, the buyer seeks out another seller.

At the beginning of the game, sellers decide once and for all whether to enter the market by comparing the expected returns to the trader profession with some exogenous outside option. In contrast, a new round of buyers enters the market each period to replace those that have traded in the past period and exited the market. Entering buyers first decide whether to participate or not in the market; if not they exit and are replaced next period. Buyers who do elect to participate remain in the market until they have successfully completed a trade. Due to the stationarity of the matching process outline below, once they have entered the market buyers will continue to participate until they successfully trade.

I examine the welfare implications of alternative mechanisms under the assumption that the distribution of buyers in the market has reached a steady state. That is, the cohort of buyers

²Indeed, the most commonly cited behavioral strategy, an even division of the gains from trade, is impossible in an environment with two-sided imperfect information where agents do not know the gains from trade.

who make a purchase and exit the market is replaced by an identical group of new buyers entering the market. This assumption, shared by similar work by Fraja and Sakovics (2001), abstracts away from any transitional dynamics, a issue that is outside the scope of this paper. Note that if different types of buyers have different probabilities of purchasing then the distribution of buyers participating in the market will differ from the fundamental distribution of entering buyer types, an issue I address in further detail below. Buyers who decide not to participate in the market exit instantly and are also replaced the next period.

Once entering buyers and sellers have decided to participate in the market, they engage in a process of costly search at the end of which they are randomly matched one-to-one with each other. If the number of buyers and sellers is not equal, some traders will be unmatched and will continue to the next period and have another chance at matching. Matched sellers learn their costs of providing the good, either engage in bargaining or trade at a fixed price (depending on which mechanism is in place in the market), and buyers who have purchased their desired good exit the market. Discounting occurs, new buyers enter, and the process begins again. Figure 1 depicts the market timeline in graphical form.

3.0.1 Type Distributions

The theoretical bargaining literature considers multiple dimensions of incomplete information: players' costs or valuations (c or v), their multiplicative bargaining time costs, or their per-round fixed costs of bargaining (k) may all be unknown to their opponents. In the context of autorickshaw bargaining, it is reasonable to assume that there may be more than one dimension of uncertainty, and this issue is explored in the estimation section. However, for expositional purposes, I set up the model with only unobserved heterogeneity in the costs and valuations of players. Additional dimensions of heterogeneity could be added at some cost in notation.

Sellers are assumed to be ex-ante identical, with free entry into the market. Sellers will thus enter the market until their ex-ante expected utility from market participation, $\mathbb{E}_{c,v} [U_S(c, v)]$ is equal to their outside option w :

$$\mathbb{E}_{c,v} [U_S(c, v)] = w \quad (1)$$

where expectations are taken over both the distribution of potential buyer values and sellers' own future costs since sellers' costs are only realized upon meeting with a buyer. Define this (exogenous) distribution of seller match-specific costs as.

$$c \sim g_S(\cdot) \text{ on } [\underline{c}, \bar{c}]$$

Buyers participate in the market if their interim expected utility ($\mathbb{E}_c [U_B(c, v) | v]$) is greater than their outside option y :

$$\mathbb{E}_c [U_B(c, v) | v] \geq y$$

where expected utility is conditioned on v because buyers know their own valuations before choosing to participate in the market. Let the mass B_0 of buyers who enter the market each period have a fundamental distribution of valuations

$$v \sim f_B(\cdot) \text{ on } [\underline{v}_0, \bar{v}]$$

Assuming that a buyer's expected utility under the mechanism is increasing in the valuation of the good, define $\underline{v}$ to be the buyer type or set of types indifferent between entry and remaining out of the market: $\mathbb{E}_c [U_B(c, v) | \underline{v}] = 0$. While the distribution of entering buyers' valuations is exogenous to the mechanism, the steady state distribution of buyers participating in the market at

any given time depends the market mechanism, since some trading rules might cause certain types of buyers to accumulate in the market awaiting an acceptable trading partner. Let this (endogenous) distribution of buyer costs be $g_B(\cdot)$ on $[\underline{v}, \bar{v}]$.

Matching of buyers and sellers occurs randomly, and match probabilities are uncorrelated with trader types both within and across periods. If the number of buyers in the market is B and the number of sellers is S , then buyers and sellers are matched with probabilities $\mu_B(S, B)$ and $\mu_S(S, B)$. Match probabilities for each side of the market are weakly increasing in the number of potential trading partners and weakly decreasing in the number of other buyers or sellers, i.e. $\frac{\partial \mu_B}{\partial S} \geq 0$, $\frac{\partial \mu_B}{\partial B} \leq 0$, and likewise for sellers' μ_S . Let $p(c, v)$ be the probability that trade occurs after matching between a buyer with value v and a seller with cost c . Then the mass of buyers in the market, B , is pinned down by the steady state condition:

$$B\mu_B\mathbb{E}_{c,v}[p(c, v)] = B_0(1 - F_B(\underline{v})) \quad (2)$$

which sets the number of buyers exiting the market (on the LHS) equal to those entering (on the RHS).

3.0.2 Utilities

Each interaction, players receive utility from two sources: First, they expect to trade the good with some probability $p(c, v)$. Second, in the course of the interaction, they receive (or lose) some utility $x_i(c, v)$. The $x_i(c, v)$ may represent either a payment received from or made to their trading partner, or any other utility cost incurred in the course of the interaction. The function through which these additional gains and losses from trade are related to the types of the traders depends on the mechanism used to allocate the goods in the market; below I outline the form of $\{p(c, v), x(c, v)\}$ for the specific cases of bargaining and fixed price mechanisms.

As is standard in the mechanism design literature, I assume that the traders' interim per-period utility functions are separable in utility from acquiring or selling the good and other costs or gains accrued during the trading process:

$$\begin{aligned} \text{Seller: } \mathbb{E}_v[u_S(c, v) | c] &\equiv \mathbb{E}_v[-cp(c, v) + x_S(c, v) | c] \\ \text{Buyer: } \mathbb{E}_c[u_B(c, v) | v] &\equiv \mathbb{E}_c[vp(c, v) + x_B(c, v) | v] \end{aligned}$$

where for both players expectations are taken over the distribution of their opponents' types in the market, $g_S(c)$ and $g_B(v)$.

I consider only stationary mechanisms. This constraint, combined with the assumptions of IID matching and a steady state distribution of buyers, ensures that buyers will never chose to exit the market without trading since their expected future gains from trade are unaffected by any previous failed negotiations.

Let κ_S and κ_B denote the search costs that sellers and buyers, respectively, incur at the beginning of each period and let δ denote the players' discount factor. A seller draws a new c cost each period, and remains in the market after trade. His ex-ante dynamic utility is

$$\begin{aligned} \mathbb{E}_{c,v}[U_S(c, v)] &= -\kappa_S + \mu_S\mathbb{E}_{c,v}[u_S(c, v)] + \delta\mathbb{E}_{c,v}[U_S(c, v)] \\ \mathbb{E}_{c,v}[U_S(c, v)] &= \frac{1}{1-\delta}(-\kappa_S + \mu_S\mathbb{E}_{c,v}[u_S(c, v)]) \end{aligned} \quad (3)$$

where the free entry condition ensures that $\mathbb{E}_{c,v}[U_S(c, v)] = w$. Thus the market mechanism has no effect on the welfare of sellers who, in expectation, earn their outside options. It does, however, affect the number of sellers that the market can sustain.

A buyer has the same value every period she remains in the market, and exits the game permanently whenever trade occurs. Her ex-ante expected utility is

$$\begin{aligned}\mathbb{E}_c[U_B(c, v) | v] &= -\kappa_B + \mu_B \mathbb{E}_c[u_B(c, v) | v] + (1 - \mu_B \mathbb{E}_c[p(c, v) | v]) \delta \mathbb{E}_c[U_B(c, v) | v] \\ \mathbb{E}_c[U_B(c, v) | v] &= \frac{-\kappa_B + \mu_B \mathbb{E}_c[u_B(c, v) | v]}{1 - \delta(1 - \mu_B \mathbb{E}_c[p(c, v) | v])}\end{aligned}\quad (4)$$

The total utility of the buyers who enter in period t is the sum of the market participants' and non-participants' utilities:

$$W_B^t = B_0 (F_B(\underline{v}) * y + (1 - F_B(\underline{v})) \mathbb{E}_{c, f_B(v)}[U_B(c, v) | v \geq \underline{v}])$$

Since new buyers enter every period, the total utility from the buyer side of the market is

$$\begin{aligned}W_B &= \sum_{t=1}^{\infty} \delta^t W_B^t \\ &= \frac{B_0}{1 - \delta} (1 - F_B(\underline{v})) \mathbb{E}_{f_B(v)}[\mathbb{E}_c[U_B(c, v) | v \geq \underline{v}]]\end{aligned}\quad (5)$$

Total welfare in the market is then $W = w + W_B$, which is a function of both the market mechanism $\{p(c, v), x(c, v)\}$ and the structural parameters, $\{f_B(v), g_S(c), \kappa_S, \kappa_B, \delta, B_0\}$. Given estimates of these parameters from any market, the social welfare generated by the status quo mechanism can be calculated, as can that of any other potential market design. I focus on the comparison between bargaining (the status quo mechanism) and an alternative fixed price market design.

3.1 The Fixed Price Mechanism

Fixed prices (or a fixed per-kilometer price in the case of a market for transportation) have the advantage that they are simple, satisfy ex-post IR, IC and budget-balance constraints, and are a clear policy alternative to bargaining since they have been adopted in many markets. While a fixed price is not, in general, the optimal mechanism, Athey and Miller (2007) show that in a stationary environment where all constraints must hold ex-post, the optimal mechanism approaches a fixed price under a wide range of parameter values. In the notation that follows, let η be an exogenously determined fixed price or rate, determined by a regulator or by market forces outside of the economy in question. Traders in the fixed price market pay nothing if they are not matched, and the buyer pays the seller η if trade occurs:

$$x_S(c, v) = -x_B(c, v) = \begin{cases} \eta & \text{if } c \leq \eta \leq v \\ 0 & \text{otherwise} \end{cases}$$

Buyers whose value for the good is below the fixed price will never enter, and those who do enter will trade with probability one, conditional on being matched with a seller whose cost is below η . Since matching is random, the distribution of buyers' valuations in the market $g_B(v)$ will be a truncation of the fundamental distribution of entering buyers' valuations ($f_B(v)$). Similarly, sellers will always make a sale if it is profitable for them to do so, since they gain the same continuation value regardless of whether trade occurs.³ Trade occurs with probability

$$p(c, v) = \begin{cases} 1 & \text{if } c \leq \eta \leq v \\ 0 & \text{otherwise} \end{cases}$$

³Here I assume that traders are not capacity constrained in the short-term. An alternate formulation would be to interpret the seller's cost c as including the opportunity cost of any sales that might have been lost by trading with the current buyer.

Substituting these fixed price trade probabilities and prices into equations 3 and 4 yields the sellers' ex-ante expected utility and the buyers' interim expected utility. These utility functions, together with equations 1 and 2, generate a system of non-linear equations that can be solved to yield the equilibrium masses of sellers and buyers in the market and hence the overall surplus W^η for a given fixed price η .

A weakness of fixed prices, and all mechanisms that require traders to walk away from gains from trade, is that these may not be sustainable if traders can communicate or learn some information about each other's types. For instance, buyers may encounter sellers for whom $c < v$, so that there are gains from trade, but if $c > \eta$ trade will not occur under the fixed price mechanism. If the buyer can somehow learn c , and the chance of punishment for deviation from the fixed price is low—as it is in many cases of decentralized trade—traders may prefer to make a side deal and divide the surplus of $v - c$ rather than search for new trading partners. In this case the fixed price will not be a feasible mechanism. This example highlights the difficulties in switching from a bargaining mechanism to a fixed price when traders can credibly reveal their costs and valuations, and punishment is unlikely. If expectations change slowly, immediately after a fixed price is introduced the market may still contain many agents who cannot afford to trade at the fixed cost but still can generate gains from trade. If other traders expect this broad distribution of types, then they may be willing to deviate from the fixed price to increase trade and the fixed price equilibrium will be impossible. On the other hand, in markets where fixed price equilibria are firmly established, only buyers and sellers with valuations (costs) above (below) the price will be present in the market. Because traders will expect these limited distributions, it may be impossible (or very costly) for a trader to credibly signal a valuation (cost) below (above) the price, and the fixed price equilibrium will be sustainable.

3.2 The Bargaining Mechanism

Many bargaining models exist that yield functions relating trader types with trade probabilities and expected payoffs $\{p(c, v), x(c, v)\}$, analogous to those derived above for the model of fixed prices. However, these functions will depend on the specific equilibria of the game being played, which may not be possible to determine a priori in the case of multiple equilibria. Thus instead of following the game theoretical bargaining literature and solving analytically for the full set of beliefs and strategies for each type, this section specifies only the extensive form and payoff functions of the bargaining game, without solving for a specific equilibrium. Later, I use these two elements, together with the pattern of actions observed from the specific equilibrium that is played in the data, to estimate $\{p(c, v), x(c, v)\}$, thereby avoiding the problem of multiple equilibria.

The extensive form is standard in the bargaining game: bargainers alternate taking actions, with the seller taking the first action and the buyer deciding whether to exit the bargaining, accept the seller's previous offer, or make a counteroffer from one of a set of discrete possible offer amounts⁴. Henceforth I refer to these exchanges of offers as a bargaining "rounds", to distinguish them from trade "periods" which represent a whole set of interactions between buyers and sellers in the multilateral market economy, each potentially including many bargaining rounds.

Formally, in bargaining round t player i chooses action a_{it} out of the set of possible actions A_{it} where:

$$a_{it} \in A_i(x_{-i(t-1)}, x_{ij(t-2)}) = \begin{cases} \chi & \text{exit} \\ \alpha & \text{accept player } -i\text{'s offer} \\ x_j \in X_i(x_{-i(t-1)}, x_{ij(t-2)}) & \text{counteroffer } x_j \end{cases}$$

⁴While the assumption of discrete actions is appropriate for the autorickshaw market, it is not without loss of generality on both game theoretic (Muthoo 1991) or econometric grounds.

The set of feasible actions depends on the previous offers due to the monotonicity of the bargaining game: sellers (buyers) can never rescind an earlier offer and raise (lower) their asking price.⁵ Nor can sellers (buyers) make a counteroffer below (above) their opponent's last offer, since in it would always be more profitable for them to accept that offer instead. For brevity of notation, let $A_{it} \equiv A_i(x_{-i(t-1)}, x_{i(t-2)})$ and $X_{it} \equiv X_i(x_{-i(t-1)}, x_{i(t-2)})$.

The players' per-round payoffs π_i depend on their actions and their beliefs over their opponents' types. I define these posterior distributions of beliefs as $h_S(v | \{x_\tau\}_{\tau=1}^{t-1})$ for the seller and $h_B(c | \{x_\tau\}_{\tau=1}^{t-1})$ for the buyer, where $\{x_\tau\}_{\tau=1}^{t-1}$ denotes the past history of offers.

Following the concept of Markov Perfect Equilibrium, let the state variable s_{it} to be the vector of payoff-relevant state variables for player i and time t :

$$\begin{aligned} s_{St} &= \left\{ x_{S(t-2)}, x_{B(t-1)}, h_S(v | \{x_\tau\}_{\tau=1}^{t-1}), c \right\}, \text{ for } t \text{ odd} \\ s_{Bt} &= \left\{ x_{B(t-2)}, x_{S(t-1)}, h_B(c | \{x_\tau\}_{\tau=1}^{t-1}), v \right\}, \text{ for } t \text{ even} \end{aligned}$$

Player have a fixed utility cost k of making a single bargaining offer, in contrast with κ_i the search cost of seeking out a new trading partner and entering into a fresh bargaining interaction. Since bargaining rounds are extremely short in the autorickshaw market—offers occur at roughly 20-second intervals—I do not include a multiplicative time cost of bargaining in the payoff function. For contexts in which bargaining takes longer, for instance negotiations between a firm and a union, these time costs may be more appropriate and could easily be incorporated into the model instead of or in addition to an additive cost of bargaining. Note that I do not impose any restrictions on the correlation between k and c and v .

With these variables defined, the seller's per round payoffs for each action can be written as

$$\begin{aligned} \pi_S(a_{St} = \chi | s_{St}) &= \delta \mathbb{E}_{c,v} [U_S(c, v)] \\ \pi_S(a_{St} = \alpha | s_{St}) &= x_{B(t-1)} - c + \delta \mathbb{E}_{c,v} [U_S(c, v)] \\ \pi_S(a_{St} = x_j | s_{St}) &= \Pr(a_{B(t+1)} = \chi | x_j, s_{St}) \delta \mathbb{E}_{c,v} [U_S(c, v)] \\ &\quad + \Pr(a_{B(t+1)} = \alpha | x_j, s_{St}) (x_j - c + \delta \mathbb{E}_{c,v} [U_S(c, v)]) - k_S \end{aligned} \tag{6}$$

Within-round payoffs for the buyer are similar, and reflect the fact that the buyer exits permanently after trade:

$$\begin{aligned} \pi_B(a_{Bt} = \chi | s_{Bt}) &= \delta \mathbb{E}_c [U_B(c, v) | v] \\ \pi_B(a_{Bt} = \alpha | s_{Bt}) &= v - x_{S(t-1)} \\ \pi_B(a_{Bt} = x_j | s_{Bt}) &= \Pr(a_{S(t+1)} = \chi | x_j, s_{Bt}) \delta \mathbb{E}_c [U_B(c, v) | v] \\ &\quad + \Pr(a_{S(t+1)} = \alpha | x_j, s_{Bt}) (v - x_j) - k_B \end{aligned} \tag{7}$$

Note that every seller action ultimately results in their returning to the market to seek another customer. Accordingly, I subtract the sellers' continuation value $\mathbb{E}_{c,v} [U_S(c, v)]$ from the payoff of each action, thereby normalizing the value of exiting the bargaining to zero. Buyers do not receive the continuation value of remaining in the market if trade occurs; however, when estimating buyers' valuations from only their bargaining offers, $\delta \mathbb{E}_c [U_B(c, v)]$ is not separately identified from

⁵While I impose monotonicity a priori, it can be shown that this condition will hold in virtually all models of bargaining.

the value of the value of the ride. I therefore normalize the driver's type-specific continuation value to zero and estimate the surplus net of the continuation value, or $\hat{v} = v - \delta \mathbb{E}_c [U_B(c, v) | v]$.⁶

Because players only make offers every other round, I write their payoffs from their opponent accepting or rejecting a counteroffer as part of their within-round payoffs, although these payoffs are incurred in the next round and are subject to the utility cost of bargaining. The first four rounds of the bargaining game extensive form are shown graphically in figure 2, with a particular series of offers highlighted for illustrative purposes. At each round of bargaining, the payoffs are written for accepting (α) and exiting (χ). Nodes at which the seller may take an action are colored black, and buyer's nodes are colored white.

An agent's beliefs on his opponent's type enter through the expected payoffs of making a counteroffer, $\pi_i(a_{it} = x_j | s_{it})$. For instance, when assessing the probability that his opponent will exit after an offer of x_j , the seller must account for both his uncertainty over the type of his opponent, and the potentially mixed strategies played by each type of buyer. Thus to calculate the probability that an offer of x_j results in opponent exit, he integrates his beliefs on each buyer type's exit probability over the distribution of buyers that he expects at state s_{it} :

$$\Pr(a_{B(t+1)} = \chi | x_j, s_{St}) = \int \Pr(a_{B(t+1)} = \chi | s_{B(t+1)}(x_j)) h_S(v | \{x_\tau\}_{\tau=1}^{t-1}) dv \quad (8)$$

where $s_{B(t+1)}(x_j)$ denotes the buyer state that would occur in the next round if the seller counteroffers x_j ; note that this contains v , the buyer's unobserved type over which the seller integrates. The expected opponent probability of acceptance is defined analogously.

The dynamic nature of the bargaining game enters through the opponent counteroffer probabilities, since these are, from the player's perspective, the probabilities of transitioning to each future state. I define these conditional state transition probabilities as $\Psi(s_{i(t+2)} | s_{it}, a_{it})$; note that due to the alternating offer form of the game, bargainers are concerned with their transition to a state two rounds ahead. For instance, for the seller, the probability of transitioning to state $s_{S(t+2)} = \{x_j, x_{B(t+1)}, h_S(v | \{x_\tau\}_{\tau=1}^{t+1}), c\}$ after making an offer of x_j is

$$\Psi(s_{S(t+2)} | s_{St}, a_{it} = x_j) = \Pr(a_{B(t+1)} = x_{B(t+1)} | s_{St}, x_j)$$

If the player accepts or exits $\Psi(s_{i(t+1)} | s_{it}, \alpha) = \Psi(s_{i(t+1)} | s_{it}, \chi) = 0$, and the bargaining interaction ends. I define the choice-specific utility of taking action a as

$$u_i(a | s_{it}) = \pi_i(a | s_{it}) + \int (V_i(s_{i(t+2)}) - k) d\Psi(s_{i(t+2)} | s_{it}, a)$$

where $V_i(s_{it})$, player i 's dynamic utility of state s_{it} is then equal to the value of their highest payoff action

$$V_i(s_{it}) = \max_{a \in A_{it}} \{u_i(a | s_{it})\}$$

Having defined the extensive form and payoff functions of the bargaining game, I do not specify or solve any particular game theoretical equilibrium for the players' beliefs or strategies. Instead my estimation technique will estimate players' beliefs from the cross sectional distribution of opponent actions observed in the data and take their strategies to be the actions maximizing their payoff functions conditional on these beliefs. Consistent with the dynamic games literature, I make three broad assumptions about the equilibrium :

⁶ Assuming a one-to-one correspondence between v and $\hat{v}$, the buyer's dynamic utility can be rewritten to include this normalization as: $\mathbb{E}_c [U_B(c, v) | v] = \frac{1}{1-\delta} (-\kappa_B + \mu_B \mathbb{E}_c [u_B(c, \hat{v}) | \hat{v}])$ Intuitively, one can think of reframing the passenger's utilities such that they now expect to participate in the market forever, but in each period gain only the difference between their absolute valuation and outside option.

- Single Equilibrium: All players in the market are playing the same equilibrium.
- Rational Expectations: Player's beliefs on their opponents' actions are correct along the equilibrium path.
- Optimizing Behavior: Player's strategies at each state maximize their expected payoffs subject to their beliefs.

These assumptions are similar to Fudenberg and Levine (1993)'s concept of "Self-Confirming Equilibrium". Intuitively, one can imagine that if players have participated long enough in the market they will learn about their opponent's distribution of actions on the equilibrium path. Since players only observe equilibrium play, the observed actions correspond to their beliefs, which are then "self-confirming".

4 Estimation

I employ the techniques of the dynamic games literature (Aguirregabiria and Mira 2007; Pakes, Ostrovsky, and Berry 2007) to estimate the parameters of the model outline above from data on the series of offers, counter-offers, and accept/exit decisions across multiple bargaining interactions. At every state, I calculate the players' expected utility of each potential action: accept, exit, or counteroffer, and then estimate the parameters such that, at each state, the actions most frequently taken in the data are those with the highest calculated payoffs. The algorithm follows series of distinct steps:

1. In the first stage, I estimate the state-specific probabilities of opponent actions that players would face after making each potential counteroffer.
2. Next, given a set of candidate parameters for each state I calculate the expected payoff of every action—accept, exit, or counteroffer— by backwards induction:
 - At nodes when it is the opponent's turn to act, I use the first stage estimates of the opponent action probabilities to calculate expectations over opponents' future actions.
 - At nodes when it is the player's turn to act, I use the values of the future states, together with a specification of action-specific utility shocks, to calculate the player's own future best response actions.
3. Once calculated, the action-specific utilities imply a set of probabilities that players take each action in each state. I maximize a dynamic logit likelihood function, using random coefficients to model the distribution of players' types, to fit the sequences of actions to in the data to the probability of these sequences implied by the model. Note that for each set of candidate parameters in the process of maximizing the likelihood function I must recalculate the action-specific utilities by repeating the backwards induction in step 2.

Individual-Specific Fixed Effects The most direct and intuitive approach to estimation of the players' types is to use the series of actions made by each individual bargainer to infer the parameters, or set of parameters that would rationalize each of their actions in the bargaining interaction as a best response. For instance, if a seller takes some action a_{St} , then by revealed

preference his cost and bargaining disutility must be such that his expected utility $u(a_{st}|s_{st}; c, k)$ of that action is greater than that of any alternative action a'_{st} he might have taken at that point:

$$u(a'_{st}|s_{st}; c, k) - u(a_{st}|s_{st}; c, k) \leq 0 \quad \forall a' \neq a_{st} \in A_S(s_{st})$$

Thus if player i takes a total of T_i actions, then the set of parameters that rationalize all those actions as best responses are the solution to the following minimization problem:

$$\{c_i, k_i\} = \operatorname{argmin}_{c_i, k_i} \left\{ \sum_{t=1}^{T_i} \sum_{a' \in A_{it}} \min(u(\hat{a}_{it}|s_{it}; c_i, k_i) - u(a'|s_{it}; c_i, k_i), 0)^2 \right\}$$

which imposes a quadratic loss function on all binding best-response inequalities. By repeating this estimation for the seller in each bargaining interaction, a set of $\{c_i, k_i\}$ sets can be computed (and likewise $\{v_i, k_i\}$ sets for the buyer) which provide an estimate of the distribution of costs, values and bargaining disutility.

This estimation approach has the advantage that it imposes no additional distributional or parameteric assumptions on the data beyond those implicit in the payoff functions, extensive form, and first round estimation of the players' beliefs. However, due to the shortness of the bargaining interactions (players make at most 5 offers), the individual valuations are likely to be estimated with substantial noise. As discussed in the results section, *prima facie* evidence of this noise is the fact that (as is common in the moment inequality literature) at least one best response inequality is typically binding even at the point in the parameter space that best rationalizes the data. This noise will then bias the empirical CDF of the estimated types, both expanding its support and flattening the distribution. Despite this bias, the non-parameteric distributions can deliver valuable insight into which parametric distributions might best fit the data in a random effects framework.

Random Coefficients Logit A random effects approach to estimation provides an alternative to the individual estimation of types that is consistent, but at the price of imposing a known distributional form on the buyer and seller values, costs, and bargaining disutility. Motivated by the shapes of the non-parametric distributions generated by the fixed effects estimation, I model the distribution of types as bivariate log-normal in $\{c, k\}$ (or $\{v, k\}$ for buyers), with the mean and standard deviation of c and v estimated separately for each trip distance, and the mean and standard deviation of k assumed to be the same regardless of the length of the journey.

Random effects estimation also requires more structure to be put on the data in the form of idiosyncratic shocks to the values of the bargainers' choices. As in many applications of dynamic discrete choice models, I follow Rust (1987), in assuming that the per-round utility from each action is hit by an additive extreme-value type 1 distributed shock, and that these shocks are IID across choices and across shocks:

$$\tilde{\pi}_i(a_{it}|s_{it}, \varepsilon_t) = \pi_i(a_{it}|s_{it}) + \varepsilon(a_{it})$$

These action-specific shocks are private information—conditional on the player's own actions they do not affect her opponent's actions—but they known to the agent prior to taking an action in each round. This error structure is by far the most commonly adopted in applied dynamic models (Aguirregabiria and Mira 2010), since it allows for closed form solutions to the action probabilities and generates a concave log likelihood. In the context of bargaining for autorickshaw rides, these shocks can be interpreted as a player's idiosyncratic belief that making a given offer will yield a relatively good outcome, or that the returns to accepting the opponent's previous offer or exiting the bargaining may be temporarily high, perhaps due to the passing of another autorickshaw or potential customer.

Agents' probability of taking an action a_{it} is the familiar conditional logit formula

$$\Pr(a_{it} = a | s_{it}) = \frac{\exp(u_i(a, s_{it}))}{\sum_{a' \in A_{it}} \exp(u_i(a', s_{it}))} \quad (9)$$

and the interim expected value of each state, now including the action-specific shocks, can be expressed in closed form as

$$\tilde{V}_i(s_{it}) = \log \left(\sum_{a \in A_{it}} \exp(u_i(a, s_{it})) \right)$$

Finally, let the dynamic payoff of choice a_{it} , now including the value of the shocks, be

$$\tilde{u}_i(a | s_{it}) = \tilde{\pi}_i(a_{it} | s_{it}, \varepsilon_t) + \int \left(\tilde{V}_i(s_{i(t+2)}) - k \right) d\Psi(s_{i(t+2)} | s_{it}, a)$$

First Stage: Opponent Action Probabilities The key ingredient to the backwards induction are the player's beliefs on his opponent's actions at each node. These beliefs incorporate both the posterior on the opponent's type, and each type's strategies. For instance, when making an offer of x_j , the seller estimates the probability that his offer will lead to the buyer exiting by integrating over her posterior on buyer's valuation as in equation 8. The econometrician cannot directly evaluate equation 8, since it contains the posterior distribution $h_i(\cdot)$, which is unobserved. Attempting to solve an analytic solution for $h_i(\cdot)$ would lead to all the problems of multiple equilibria discussed in section 2.

Instead, following the dynamic games literature (Aguirregabiria and Mira 2007; Bajari, Benkard, and Levin 2007), I rely on the assumption of a single equilibrium and estimate the players' beliefs on opponent types and strategies from the distribution of counteroffers observed at each state in the data. The critical feature of the bargaining game that permits this approach is that the unobserved component of the state vector, c or v , is also unobserved to the bargainer's opponent, so players know that their opponents will take actions based only on the history of offers (which is observable to the econometrician). For example, the seller's expectation of the probability that the buyer exits after he counteroffers x_j (from equation 8) can be calculated using the empirical distribution of the buyer's actions following that counteroffer:

$$\begin{aligned} \Pr(a_{B(t+1)} = \chi | x_j, s_{St}) &= \int \Pr(a_{B(t+1)} = \chi | s_{B(t+1)}(x_j)) h_S(v | \{x_\tau\}_{\tau=1}^t) dv \\ &= \Pr(a_{B(t+1)} = \chi | x_j, \{x_\tau\}_{\tau=1}^{t-1}) \end{aligned}$$

Estimating these opponent action probabilities raises an issue common to many dynamic structural papers: if the state space is large relative to the data (as it is in bargaining games) the action probabilities estimated by a simple count estimator will be very noisy and contain many values of 0 or 1. For instance, if a given state is observed only once in the data, estimating beliefs by the simple action probabilities in the data would imply that players expect their opponents to carry out the action performed at that state in data with 100% probability. To generate more continuous beliefs, I follow Aguirregabiria and Mira (2002) and estimate a first stage multinomial logit model to smooth the action probabilities:

$$\Pr(a_{-i(t+1)} = a | s_{-i(t+1)}(x_j)) = \frac{\exp(\theta_a \mathbf{q}(s_{-i(t+1)}(x_j)))}{\sum_{a' \in A_{it}} \exp(\theta_{a'} \mathbf{q}(s_{-i(t+1)}(x_j)))}$$

where $\mathbf{q}(s_{i(t+1)}(x_j))$ is a vector of linear and squared state values and interactions, and θ_a is a vector of coefficients for action a . This estimation procedure is repeated for each bootstrap sample when calculating the standard errors,

Backwards Induction While the probabilities of opponent actions are identified by the observed actions in the data, the player's own best response at a given state cannot be estimated in the same manner because it depends on the player's own type, which is unobservable to the econometrician but (unlike the opponent's type) known the player herself. Thus the observed distribution of actions does not coincide with the player's own expectations of her actions once she reaches that state. I must therefore calculate the best responses at each node in the game tree as a function of the valuation (or cost) and bargaining disutility.

I do this through the backwards induction technique outlined above. As an illustration of the algorithm, consider the example of estimating the action-specific payoffs for the seller. Assume for simplicity that all bargaining interactions last at most T rounds, and that all possible states and actions are observed in the data. Assume, furthermore, that it is the seller's turn to act in round T . Then in round T the seller faces the choice of whether to accept or exit and selects the choice with the highest value. The value of each last-round state s_{iT} is then

$$\bar{V}_i(s_{iT}) = \int_{\varepsilon} \max \{ \tilde{\pi}_i(a_{iT} = \chi | s_{iT}, \varepsilon_T), \tilde{\pi}_i(a_{iT} = \alpha | s_{iT}, \varepsilon_T) \} d\Gamma(\varepsilon)$$

Round $T - 2$ is the seller's previous action. In this round the seller considers exiting, accepting, or making one of the remaining feasible offers. Each of these offers would generate an action by the opponent, ultimately resulting in some exit, accept, or continuation value, $\bar{V}_i(\tilde{s}_{iT})$. The opponent action probabilities estimated in the first stage are then plugged into the backwards induction to calculate the expected payoffs of the counter-offers, and the value of each state in round $T - 2$ is determined by the value of the best response action. Backwards induction continues analogously in round $T - 4, \dots, 1$ until the value of each state has been calculated.

Empirical States Empirical estimation of the bargaining game necessitates some reduction of the dimension of the state space. Since players' posteriors about their opponent's type are the same after identical histories of offers, the $h(v | \{x_{\tau}\}_{\tau=1}^{t-1})$ term in the state vector can be replaced with just the history of offers itself, $\{x_{\tau}\}_{\tau=1}^{t-1}$. However, including the whole history as a state would be infeasible due to the very large number of implied states⁷. Thus I proxy the whole past history of offers with just the last two offers and the round itself, t , a simplification which is consistent with the equilibrium in many models of bargaining as a Markov perfect equilibrium. Also included in the state vector is a measure of the distance of the trip in discrete kilometers, d , and the unobserved type of the player, c or v . Thus the full empirical state vector, denoted $\tilde{s}_{it}$, is composed of the five elements

$$\begin{aligned} \tilde{s}_{St} &= \{x_{S(t-2)}, x_{B(t-1)}, t, d, c\}, \text{ for } t \text{ odd} \\ \tilde{s}_{Bt} &= \{x_{B(t-2)}, x_{S(t-1)}, t, d, v\}, \text{ for } t \text{ even} \end{aligned}$$

Final Round States The previous discussion of the backwards induction process assumed that all states are observed and that the game had a fixed terminal round T . In reality, neither of these conditions are likely to hold. As the theoretical literature recognizes, bargaining games can, in

⁷On the order of $\sum_{t=1}^T |A|^t$ if there are T periods per bargaining interaction, and $|A|$ possible actions per period.

principle, continue for infinitely many rounds. Although any finite sample of data will only contain a maximum of rounds T , at time T the players themselves could have chosen to counteroffer and thus extend the game to $T+1$, or to have taken an action at an earlier round that would have led to $> T$ rounds. Infinite horizon models are commonly estimated in the dynamic structural literature, thanks to the presence of some stationarity assumption. However, even imposing a Markov-perfect equilibrium, the bargaining game is only stationary conditional each player's posterior distribution of the other's type which is not observed.

As discussed in the context of the state variable, the past history of actions (or in this case a subset of the past actions plus a time index) may be substituted for the posterior in the state variable, but with this substitution stationarity is lost. Thus the number of potential states (where now the state depends directly on t) becomes infinite, and the analysis must address the problem of unobserved states.

However, while the unobserved states cannot be measured directly from the data, some structure can be placed on their values. Just as the theory literature often uses conjectures on off-equilibrium path actions to focus on specific equilibria, I can estimate structural parameters using both "optimistic" and "pessimistic" priors as to the type of the other player.

- The "optimistic" outcome of a counter-offer is that the opponent is revealed to be a "soft" type, and immediately accepts the counter offer. Thus (for the seller)

$$u^{\max}(x_{St}) = x_{St} - c - k$$

and analogously for the buyer,

$$u^{\max}(x_{Bt}) = v - x_{Bt} - k$$

- The "pessimistic" outcome of any counter-offer is it leads to the opponent deciding that the player is a "hard" type, and immediately exiting:

$$\text{Seller: } u^{\min}(x|\tilde{s}_{St}) = -k$$

$$\text{Buyer: } u^{\min}(x|\tilde{s}_{Bt}) = -k$$

Note the slight asymmetry between the two bounds: while the "optimistic" outcome gives the best possible payoff conditional on the action, the "pessimistic" payoff is still greater than, for example, if the player had made the offer then wasted her time making several more offers before her opponent exited. Nevertheless, the "pessimistic" scenario is consistent with the assumptions on off-equilibrium path actions used in, for instance, Chatterjee and Samuelson (1988).

In practice I apply these bounds only to the values of actions taken in the final round ($T = 9$) of the game and calculate the values of all other states (whether observed or not) by the backwards induction procedure described above.

Likelihood Function Assuming a bivariate log-normal distribution of individual costs/valuations and bargaining disutility, I solve for the parameters of the distribution by maximum likelihood. Letting μ_c^d, σ_c^d denote the mean and variance of the log normal distribution of driver costs for a trip of distance d , and $\mu_{k,S}, \sigma_{k,S}, \rho_S$ denote the (distance invariant) mean and variance of the bargaining disutility and its correlation with the cost, then the full parameter vector for the seller is $\theta_S = \{\mu_c^1, \sigma_c^1, \dots, \mu_c^D, \mu_{k,S}, \sigma_{k,S}, \rho_S\}$. The buyers' parameter vector, θ_B , is defined analogously, with valuations also estimated distance by distance.

Define N to be the number of bargaining interactions observed in the data, and T_{in} to be the number of rounds that player i may take an action in bargaining interaction n (as always, i indexes the buyer or seller). Following Judd (2006) I integrate the likelihood over the $\{c, k\}$ and $\{v, k\}$ distributions using Gauss-Hermite cubature, creating a two dimensional $M = 16$ node grid with parameter values θ_{im} at node m . Defining $\omega(\theta_{im})$ to be the appropriately rescaled Gauss-Hermite weight of node m , the full likelihood is then

$$\mathcal{L}_i(\theta) = \prod_{n=1}^N \sum_{m=1}^M \omega(\theta_{im}) \prod_{t=1}^{T_{in}} \Pr(a_{itn}|s_{itn}; \theta_{im})$$

where $\Pr(a_{itn}|s_{itn}; \theta_{im})$ is the logit action probability, as defined in equation 9. Note that the likelihood makes use of the fact that a player's type is constant throughout the T_{in} actions that he takes in the bargaining interaction.

5 Data

The data for this study comes from the market in local transportation by autorickshaw in Jaipur, India. An autorickshaw is a form of three-wheeled mini-taxi, officially capable of carrying three passengers (although often far more in practice) in a semi-enclosed back seat. Autorickshaws are the primary means of rented transportation in Jaipur, a city of approximately 3.2 million people. Although in Jaipur they are technically outfitted with meters, during the survey period of January 2008 to January 2009 the meter rates had become surpassed by inflation and no autorickshaw driver ever used the meter. There were no police or government efforts to enforce the meter, and all prices were set by negotiation.

In addition to the universal prevalence of bargaining, at least three factors make the autorickshaw market an ideal test case for the economic analysis of bargaining. First, autorickshaw rides are very homogeneous conditional on observables. Given the day, time, and physical appearance of the autorickshaw (all of which are observable), a ride from point A to B is the same (in expectation) regardless of which driver provides it. Second, the autorickshaw market is an excellent candidate for potential policy interventions since a fixed price per kilometer or other non-linear price schedule is a feasible policy alternative. In many Indian cities similar to Jaipur (e.g. Ahmedabad, or Bangalore) virtually all autorickshaw rides are priced by the meter. Many other cities have set up "pre-pay" autorickshaw stands, where rides from the stand to various destinations each have a specific, pre-determined price. This variation in market equilibria suggests that it is differences in local government policy rather than in the underlying structural parameters that determines whether a city is in a bargaining or fixed-price mechanism. Third, the price of a ride is low enough so that data can be collected on actual transactions under the control of the researcher.

In order to collect the data, surveyors followed two protocols:

- In "real" bargaining, surveyors were told to travel through a pre-assigned series of waypoints (for instance, from point A to B, to C, then back to A), and given a lump sum of money to pay for the travel. Any of the payment that they did not spend on the autorickshaw fare was theirs to keep, and once they had finished their day's assigned circuit they were free to return home. Thus the surveyors' opportunity cost of money and value of time in terms of money should have been similar to what it would have been had they been bargaining on their own.⁸ However, unlike real bargainers, the surveyors were required to complete their trips by

⁸Because of concerns that surveyors might take buses between waypoints, or negotiate with a single autorickshaw driver for the entire route supervisors were stationed at waypoints to monitor the surveyors.

autorickshaw, thus the "market entry" stage of the model is inapplicable to them, and their values of rides may not be representative of the general population who choose to travel by autorickshaw.

- In "scripted" bargaining, surveyors were assigned to stand in specific locations and were given a written bargaining protocol consisting of a destination and a sequence of pre-determined counteroffers. They then hailed passing autorickshaws, requested a ride to the destination and bargaining with the driver according to the protocol. If the driver were to accept a counteroffer, the surveyor invented an excuse not to take the ride. Although they had no personal stake in the outcomes, surveyors were instructed to act as if they were bargaining in a realistic manner so that drivers would themselves respond as naturally as possible. These scripted bargaining sequences were both cheaper and faster to collect than the real bargaining interactions. They also allowed the driver's action probabilities to be measured more accurately in states that rarely occurred in the real bargaining data.

The data from these scripted interactions can also offer a test of endogeneity of the state variable. A potential concern in estimation using only real bargaining is that drivers are somehow able to signal their types to passengers in ways not captured by their previous offers. These signals would then influence passengers' offers and introduce biases into the drivers' conditional action probabilities. For instance, if one type of driver can signal a minimum passenger counteroffer below which he commits to exit, passengers may respond by shifting their counter-offers upward or exiting themselves. These drivers will then never be forced to make good their threat, and their valuations will be misestimated. The scripted bargaining is free from this problem, since surveyors' offers cannot be correlated with any statements made by drivers. The disadvantage is that they provide no information about the passenger's bargaining habits. In all analysis that follows, the passenger's offers from the scripted bargaining interactions are dropped except insofar as they serve as the state variables for the driver's choices.

Immediately after the conclusion of the bargaining, surveyors wrote down the series of offers made by the drivers and themselves, and noted the duration of the whole interaction (in seconds). They also recorded the model and quality of the autorickshaw, as well as details about the environment such as whether other drivers had attempted to interrupt the bargaining, the weather, and the time of day.

In August 2010, approximately 18 months after the conclusion of the bargaining data collection, more data was collected on the characteristics of the drivers themselves. To avoid selection bias, surveyors hailed autorickshaw drivers from the streets and administered a short questionnaire in exchange for a small payment (10 rupees) to encourage compliance. Despite the 18 months gap between the driver survey and the bargaining, the qualitative features of the autorickshaw market remained essentially the same in terms of the market structure, timing, and bargaining process. The results of this survey, presented in table 2 below, are used to derive other parameters about the choices of autorickshaw drivers to enter the market, in particular their potential wages outside of the autorickshaw market.

The characteristics of the surveyors employed to collect the data has greater significance in this experimental setup than in many others, since their actions are essential for predicting the distribution of counteroffers and exit/accept probabilities that drivers face after each possible offer. In order to reduce the number of state variables as much as possible, surveyors were chosen to be homogenous on characteristics that could be observed by the drivers. Surveyors hired for this project were all males, between the ages of 20 and 35 and dressed in similar casual clothing. All

had finished 10th grade, and some had several years of college or had graduated. All made the same salary of rs. 200/day, in addition to whatever they earned from bargaining. Finally, all surveyors took autorickshaws routinely as part of their normal personal transportation, often on exactly the same routes as assigned those for this research. Thus the maintained hypothesis that surveyors know the distribution of driver types and counter-offers and the equilibrium being played is reasonable in this environment.

Summary statistics Table 1 presents summary statistics on bargaining for autorickshaw rides. In total, 2993 bargaining interactions were recorded, of which 2369 (79%) were conducted with the surveyor making offers from a predetermined list, and 624 (21%) in which the surveyor was the residual claimant of any gains from bargaining. Of the real bargaining interactions, 67% resulted in the surveyor actually taking the ride. The average duration of a bargaining interaction was 5 rounds, which implies two exchanges of offer/counteroffer between the players, and a final accept/exit decision by the driver. Bargaining was over quickly—interactions usually lasted less than a minute. If the interaction did not lead to trade, the average wait for another potential driver was almost 4 minutes. Some surveyors had to wait substantially longer for the next ride—the 90th percentile of the times between autorickshaws was 7:27 minutes. Trips were chosen to average about 5 kilometers, with the longest being 8.06 kms., and the shortest 1.65 kms. Finally, the data on offers suggests the role of bargaining in dividing the gains from trade between buyers and sellers: successful interactions had an average final price of rs. 41 between the driver’s initial offer of rs. 56 and the passenger’s average counter-offer of rs. 35.

The set of seller offers observed in the data is from 20 to 100 rupees, increasing at 5 rupee intervals, and the range of offers for buyers is from 15 to 65 rupees, again increasing at 5 rupee intervals. I define the set of possible actions in the backwards accordingly.

Table 2 presents the results of the interviews with drivers. On average, drivers were 35 years old and had about 10 years of experience driving an autorickshaw. The mean driver had attended school for over 5 years, although a substantial number (31%) reported having never attended school at all. 46% rented the rickshaw, and of those who rented, the average daily payment was 177 rupees. The median and mode rental rate, with 52% of the observations, was rs. 200, with substantial concentrations at rs. 150 (22%) and rs. 100 (8%). Average daily profit was rs. 267 (around \$5.80 at the exchange rates during the time the data was collected), with rs. 151 being paid for fuel costs. Interestingly, despite the substantial rental costs of an autorickshaw, drivers who rented their autorickshaws made only 88 rupees less than owners, due to the fact that their revenues were higher.

Most drivers work full-time in the autorickshaw market. The average driver worked 6.55 days per week, and the median driver works every day. Each day, the average driver works around 10 hours, in which he transports, on average, slightly over 8 passengers or groups of passengers. Very few (<1%) of drivers also have other jobs, but 8.44% share their autorickshaws with other drivers. Since these drivers do not work significantly fewer hours or days per week than those who do not share their autorickshaws, they are presumably passing the autorickshaw to other drivers who work a night shift.

Informal conversations with autorickshaw drivers suggest that there is substantial variation in drivers’ individual costs for a given ride. Many drivers operate primarily in a certain neighborhood, and return there after delivering a passenger. Thus their cost of a trip in the direction of their "home base" is substantially lower than a trip elsewhere. Similarly, drivers often have regular passengers that they transport daily—for instance taking children to school. A trip in the direction of the driver’s next scheduled pickup would have a low cost, since the driver would be headed in

that direction anyway. Finally, the 46% of drivers that rent or share their autorickshaws and must return them at specific times. If the return time is approaching, the costs to taking a long ride away from the drop off point will be idiosyncratically high.

Different passengers will also have highly variable willingness to pay for the same ride, depending on their urgency, their wealth, and other factors. Anecdotal evidence suggests that, as theory predicts, drivers acknowledge the variability of their passengers' valuations and adjust their bargaining accordingly. This may explain the experience of many foreigners who are quoted (relatively) exorbitant initial offers for retail goods by Indian merchants.

Reduced form evidence from the bargaining interactions also suggests that traders' actions shift their opponents' posterior distributions on their types. The most natural test of bargaining behavior is whether players respond to relatively high or low offers by accepting or exiting the bargaining. Figure 3 shows the probability that a player exits the bargaining conditional on the opponent's previous offer (top and bottom lines show 95% confidence interval). Higher passenger offers decrease drivers' probability of exit, and increase passenger exit probability (except in round 2, when no passenger ever exited).

Figure 4 shows the relationship between previous offers and the probability a player accepts their opponent's offer. As expected, higher offers increase the probability of drivers accepting and decrease passenger's acceptance probability. The sole exception is in the third round where the relationship between driver acceptance and passenger previous offers appears somewhat bell-shaped, although the confidence intervals are large. Perhaps drivers interpret a high second round offer as an opportunity to extract even greater surplus. Note that in both figures the results in rounds 4 and 5 are on the selected sample of drivers and passengers who did not exit in the earlier rounds.

Understanding the reduced form effect of an offer on the potential counteroffers is more challenging because the two selection effects of acceptance and exit now operate within the round as well. On the acceptance side, all driver (passenger) offers below (above) the preceding offer are mechanically unobserved, because they would lead to accepting the ride. For example, if the passenger makes a high offer to the driver, the driver's counteroffer conditional on not accepting the passenger's offer must be even higher. This can be overcome, as I do below, by examining the empirical CDFs of counteroffers with a mass point equal to the fraction of bargainers who accepted located at the value of the previous offer. The rest of the CDF shows the relative probability of offers conditional on not accepting, and by comparing these portions of the CDF conditional on different values of the opponent's previous offer we can determine whether, for instance, drivers counteroffer differently after a low offer that they rejected than after a higher offer that they would also have rejected. There is no analogous direct method of dealing with the differences in exit rates conditional on the previous offer since it is unclear (without additional structure) what offer exiting players would have made had they not exited.

Subject to this caveat, Figure 5 displays how the distribution of counter-offers depends on the preceding offer conditional on not accepting or exiting. Each panel contains three kernel weighted⁹ CDFs, each corresponding to a different cross section of the distribution of previous opponent offers. Each panel tells a similar story: conditional on accepting players respond to high offers by making high counteroffers. In each case the CDF of the counteroffers made to a previous offer in the 80th percentile is to the right of the 50th percentile previous offer CDF, which is (weakly) to the right of the 20th percentile. The vertical portions of the CDF on the low end of the driver's CDFs and high end of the passenger's CDFs are caused by the mass points of accepting drivers and passengers.

⁹A Gaussian kernel was used, with bandwidth of $\sigma \left(\frac{4}{3 \cdot N} \right)^{1/5}$ where N is the sample size and σ is the standard deviation of the variable being smoothed. For the CDFs, both the initial offers and the counteroffers were smoothed using this kernel.

Taken together, these graphs suggests that players in this market face the classic bargaining trade-off: Passengers can choose to make a high offer and get a higher probability of a favorable counter-offer at the risk of the driver leaving and then being forced to wait for a new autorickshaw, or make a low offer and get a high probability of rapid acceptance at the cost of lower surplus. Drivers face the same trade-off in reverse when it is their turn to make an offer. Interestingly, figures 3-5 suggest that these mechanisms may work slightly differently for the two sides of the market. Since they are more mobile, drivers seem to punish/reward low or high passenger offers more through the exit/accept decision, whereas passengers may be more responsive in the counter-offer dimension.

These patterns are broadly consistent with the signalling equilibria in many game theoretical models of bargaining. Drivers with high costs of providing the good or low costs of bargaining make high initial offers to signal their types (since low cost/impatient drivers would not want to take this risk), and passengers update their beliefs and modify their bargaining strategies accordingly.

6 Results

The techniques and data described above can be used to calculate both the distribution of costs of the drivers and the valuations of passengers, representing the supply and demand sides of the market. However, the costs and valuations recovered from this specific data collection strategy have different external validities: the costs are representative of the population of drivers, since drivers approached surveyors essentially at random and were unaware that their bargaining opponents were recording their offers until after the bargaining was complete. Thus the supply side is representative of the true market supply. The demand curve represented by the estimated passenger valuations, on the other hand, is specific to the context of the data collection for this paper. The surveyors represent only a small portion of the potential market for autorickshaw rides, and are thus unlikely to be representative of the general market for autorickshaw rides.

Due to these differences in external validity, I split the results section into two parts: In the first section, I analyze the distribution of drivers' cost and the supply curve implied by these costs. Here, I am able to examine the impact of policies such as enforcing the official meter rate, a policy in fact carried out by the government shortly after the study period concluded. In the second section I interpret the valuations of the surveyors literally, and derive the optimal fixed price and implied welfare outcomes if the surveyor valuations were those of the full population.

6.0.1 Estimated Parameters

The fixed effects estimation results for the drivers are displayed in figures 6, 7, and 10. Figure 6 displays histograms containing the distribution of driver costs for the four most common trip distances in the sample, 3, 4, 5 and 6 kilometers. Bootstrapped 95% confidence intervals are displayed as lines above the bars. The costs are imprecisely estimated, an outcome which may be due to the non-parametric nature of the estimation technique. Nevertheless, the distribution of costs appears reasonable—driver costs increase for higher distances, while their variance remains roughly the same. Note that the range of values between the lower and upper bounds on the identified set of driver costs is so small that it is not visible on the graph. This is due, in large part, to the fact that the moment inequalities are binding for the large majority of the sample. This suggests substantial noise in the inequalities, perhaps coming from imprecise estimates of the drivers' beliefs on their potential passenger's future actions. Since this uncertainty is incorporated into the bootstrapping (the first stage estimation of beliefs is repeated for each bootstrapped sample), the imprecision of the beliefs may also be contributing to the large confidence intervals on the value distribution.

Figures 7 and 10 display estimates of drivers' bargaining disutility k_S . While the majority of values are in the range of 0-2 rupees, the estimation suggests that over 10% of drivers have bargaining disutilities of greater than 4 rupees per offer—a relatively large amount considering that the mean price of a ride taken in the data is in the range of 35-55 rupees (depending on distance), and on average 5 offers are made per bargaining interaction. Finally, figure 10 shows the three dimensional histogram of drivers' cost and bargaining disutility for rides of 4 kilometers. The results show a strong negative correlation ($r = -.82$, $p = .00001$) between driver's costs and dislike of bargaining: drivers with lower costs (and thus more potential surplus) are estimated to have higher additive costs of bargaining. Interestingly, this result is consistent with a multiplicative discount factor, which would also generate this same type of correlation.

The results of the random effects logit estimation of the drivers' parameters are presented in table 4. As expected, mean costs reported in box A are generally increasing in the distance of the journey, although 5 kilometer rides appear to have both a very high mean and variance, and costs appear concave with respect to distance. This may be due to differences in drivers' outside options after trips to farther, more remote destinations. Estimates are reasonably precise, especially for the most common distances of 4-7 kilometers. The correlation between driver costs and bargaining disutility is displayed in the third column of table 4. As in the fixed effects estimation, it is negative but of a much smaller magnitude and very noisily estimated. This suggests that perhaps the log-normal distribution cannot capture the exact features of the underlying distribution of types, in particular the . Box B presents the estimated driver disutility of bargaining. It is substantially lower than that of the fixed effects estimation (.31 versus 1.3 rupees per offer) although imprecisely estimated.

The equivalent parameters of the distribution of passenger valuations are displayed in table 5. Strikingly, mean passenger valuations are often estimated to be lower than drivers' costs for the equivalent distance, although the large estimated standard deviations of the passenger's log-normal valuations create some overlap between the cost and valuation distributions at each distance. Passenger bargaining disutility is estimated extremely imprecisely under bivariate log-normal distribution, with several large outliers in the bootstrapping generating extremely large standard errors. Finally, passenger values do not, on average show a substantial correlation with bargaining costs, although here too the standard errors are large.

6.1 Fixed meter rates and driver utility

A policy change by the Government of Rajasthan provides a useful opportunity to test whether the estimated driver cost distributions are in fact reasonable. The month after the final bargaining data was collected (February, 2009) the Jaipur Road Transport Officer (the road safety enforcement agency) implemented a policy to force autorickshaw drivers to travel by the meter rate. Police set up checkpoints and drivers found to be travelling with meters turned off were fined and threatened with confiscation of their licenses. This initiative was met with strong resistance from the drivers, many of whom went on strike and demonstrated outside government buildings. At the time, the official meter rate was 11 rupees for the first kilometer, and 6 rupees for each additional kilometer. Spokesmen for the drivers' association demanded an increase of 4 rupees for the first kilometer, and one rupee for each subsequent kilometer.

Table 6 shows the fractions of drivers whose estimated costs are below the meter rate, for both the official February, 2009 meter rate, and that proposed by the drivers. The results suggest that under the official meter rate relatively few drivers would be willing to accept passengers on trips of distances in the 3-6 kilometer range. For instance, I estimate that only 18.4% of drivers would have idiosyncratic cost shocks low enough to be willing to take passengers on a typical 4 kilometer

journey. Although the standard errors are large, the fraction of drivers willing to accept the meter rate is bounded below 60% for distances of 3, 4, 6, and 7 kilometers. Drivers' preferred rates would, naturally, allow a substantially larger proportion of the drivers to operate profitably, and suggest that drivers would be willing to accept virtually any trip of 7 or 8 kilometers. For shorter distances, however, a substantial fraction of drivers remain might be unwilling to travel. Note that, because drivers' costs are passenger specific, this does not imply that these high cost drivers will never accept a passenger. This fraction represents those drivers whose costs are high because the proposed trip is in the opposite direction of where they are travelling at the moment, or who have other, transient reasons to turn down a passenger.

Given the estimated costs and bargaining disutility of drivers in the bargaining status quo, I can investigate the question of what alternative meter rates would in fact ensure that the drivers continue to participate in the market. In accordance with standard taxi or autorickshaw meters, I allow the price to have both a lump sum "meter down" component η_1 and variable per-kilometer rate η_2 ; to conserve notation, let the vector $\eta = \eta_1 + \eta_2 d$ for distance d . I then solve for the price that sets drivers ex-ante utility under the fixed price regime with price η equal to their expected welfare under bargaining

$$\frac{1}{1-\delta} (-\kappa_S + \mu_S \mathbb{E}_{c,d} [\eta - c]) = \mathbb{E}_{c,v,d} [U_S(c, v)]$$

In a partial equilibrium framework, where both drivers' matching probabilities μ_S and their search costs and discount rates remain unchanged by the switch to fixed prices, the price that would make drivers indifferent between the fixed price market and the bargaining status quo is implied as the solution to

$$\eta : \mathbb{E}_{c,d} [\eta - c] = \mathbb{E}_{c,k,d} [V(t=1; c, k, d)]$$

where $V(t=1; c, k, d) = \mathbb{E}_v [u_S(c, k, v) | c, k, d]$ is the drivers' dynamic utility of the first round of the bargaining game for a journey of distance d , net of the continuation value of remaining in the market after trade occurs or does not occur.

Since autorickshaw meter prices have both an intercept (first kilometer price) and slope (price for subsequent kilometers), there is a continuum of prices that satisfy this condition. Figure 11 plots these prices on the lower, dashed line, with the higher, solid line showing the set of prices that would set driver's expected fixed price profits equal to the nominal amount they received from passengers, not including any disutility costs. Both price schedules are substantially higher than the existing government rate of 11+6rs./km, and confirm that drivers' welfare would decrease from enforcing the official rate unless a substantial number of new passengers entered the market. Interestingly, the drivers' suggested rate of 15+7rs./km lies very close to the set of prices which would keep drivers indifferent, in utility terms, to the current bargaining system. While these results cannot be interpreted too strongly given the partial equilibrium nature of the analysis, they suggest that the distribution of costs recovered for the drivers may indeed reflect the true parameters.

6.2 The Pre-Paid Autorickshaw Stand

While the drivers' valuations are informative about the supply side of the market, a full comparison of the welfare implications of bargaining versus fixed prices must include the distribution of valuations and bargaining costs of the passengers as well. However, as mentioned earlier, the survey design can provide information only on the parameters of the valuations of the surveyors themselves, who are both a restricted portion of the distribution of passengers, and are taking rides because they are required to do so for their jobs.

Subject to these caveats, I consider the counterfactual policy of allowing passengers the option to purchase an autorickshaw ride at a fixed meter price instead of beginning bargaining with a driver. This option, known as the “pre-paid stand”, currently exists in some Indian airports and train stations, although primarily for the taxi transportation markets. Passengers choose to take the pre-paid autorickshaw at price η instead of bargaining if

$$v - \eta \geq \mathbb{E}_c [U_B(c, v) | v] + \kappa_B$$

where the search cost is added back to the passenger’s expected utility since the passenger could choose to immediately engage in bargaining with one of the drivers waiting at the autorickshaw stand. Under the partial equilibrium assumptions that passengers’ outside bargaining option remains the same, their choice to purchase the fixed price autorickshaw ride can be expressed in terms of the estimated valuation $\hat{v}$ as

$$\hat{v} - \eta \geq V_B(t = 1; \hat{v}, k)$$

where $V_B(t = 1; \hat{v}, k) = \mathbb{E}_c [u_B(c, v) | v, k]$ is the passenger’s first round dynamic utility from the bargaining game.¹⁰

I solve for the optimal pre-paid price, subject to the constraint that drivers are indifferent (in ex-ante utility) between going to the pre-paid stand and seeking a new bargaining interaction. This condition ensures that drivers themselves would be willing to pick up passengers at the fixed rate, while avoiding the issue of congestion and long waits at the auto stand. While a price that does not satisfy this constraint might yield higher overall surplus (for instance one in which drivers wait longer to find a passenger at the stand than to find a bargaining passenger) the lack of information about the matching function makes evaluation of this case difficult given the available data. I thus solve for

$$\begin{aligned} \eta^* &= \arg\max_{\eta} \mathbb{E}_{\hat{v}} [\max \{ \hat{v} - \eta, V_B(t = 1; \hat{v}, k) \}] \\ \text{subject to } \mathbb{E}_c [\eta - c] &= \mathbb{E}_{c,k} [V(t = 1; c, k)] \end{aligned}$$

Searching through the set of prices yields an optimal price of rs. 8.7 for each kilometer and 0 fixed meter down costs. At these prices, passengers’ mean per-period surplus would be 28% higher than their surplus without the option of using the autorickshaw stand. However even with the option of the pre-paid stand, 63% of passengers would still prefer to make remain in the bargaining market, suggesting that, at least for the sample of passengers in the data, bargaining remains a valuable option despite its costs.

Several factors suggest that these welfare estimates may be a lower bound. First, if relative to the full population of buyers, the sample buyers had low valuations of the ride and were relatively patient (as might be expected given their incentive structure) then higher valuation passengers with greater bargaining disutility would stand to gain more from the introduction of the fixed price autorickshaw stands. Second, if certain types of buyers choose the fixed price over bargaining, drivers will have better information regarding the types of buyers who choose to remain in the bargaining market. This extra information may improve the efficiency of the bargaining outside option, albeit at some transfer of surplus from passengers to the now better informed drivers.

¹⁰For passengers, who make the second offer, calculating mean expected interim utility requires first averaging over driver initial offers:

$$\mathbb{E}_c [u_B(c, v) | v, k] = \sum_{x_j \in X_{S1}} \Pr(x_{S1} = x_j) V_B(x_j, t = 2 | \hat{v}_n, k_n)$$

7 Conclusion

This paper has applied the theories of market design and bargaining and the empirical techniques of dynamic structural estimation to the comparison of markets with either bargaining or fixed prices. Unlike earlier papers on empirical bargaining that have tested the reduced-form implications of bargaining models, I have used the structure of the extensive form of the bargaining game to calculate payoffs and hence the structural parameters of the model. However, in contrast with the laboratory bargaining literature, I do not specify a particular game theoretical model of bargaining. Instead, following the literature on the estimation of dynamic games, I estimate the equilibrium strategy functions from the data and then solve for the parameters that imply that players' actions are optimal given their opponents' expected responses. These techniques, originally developed to analyze firm entry and exit, have rarely been applied outside that context and this study represents their first application to the field of bargaining.

In the specific case studied—the market for local transportation in Jaipur, India—giving traders the option to avoid bargaining and purchase a ride immediately at a fixed price is shown to increase overall welfare by allowing high value buyers to opt out and avoid incurring the disutility of bargaining. However, even with the option of a fixed price, the majority of buyers are estimated to have valuations and bargaining disutility such that they would prefer to remain in the bargaining market, suggesting that, at least for the buyers in this sample, flexible mechanisms such as bargaining retain substantial value. Generalizing these results to consider the counterfactual of a complete switch to fixed prices would require more information about buyers' and sellers' entry choices into the market and matching probabilities, and is a natural next step for future research.

More broadly, structural parameters of bargaining models are of great interest in a variety of settings, for instance in determining the cost of firm/union wage negotiations, or the transactions costs to purchasers of new homes. The choice between bargaining and fixed prices is itself relevant in many other contexts, including markets in developing countries for automobiles, or negotiations between health-care providers and insurance companies. Given data on the series of bargaining offers, the estimation performed in this paper for the autorickshaw market could be applied to any market with bargaining, and the same analysis could be undertaken. Structural analyses of this type have proven to be a robust and useful tool for the analysis of other market mechanisms, in particular auctions, and this paper has taken the first step toward applying them to the rich set of questions in the field of bargaining.

References

- Abreu, D. and F. Gul (2000, January). Bargaining and reputation. *Econometrica* 68(1), 85–117.
- Aguirregabiria, V. and P. Mira (2002, July). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica* 70(4), 1519–1543.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 70, 1519–1543.
- Aguirregabiria, V. and P. Mira (2010, May). Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156(1), 38–67.
- Athey, S. and D. A. Miller (2007). Efficiency in repeated trade with hidden valuations. *Theoretical Economics* 2(3), 299–354.
- Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002). *Bargaining with Incomplete Information*, Volume 3 of *Handbook of Game Theory*, Chapter 50, pp. 1896–1945. Elsevier Science B.V.
- Ayers, I. and P. Siegelman (1995, June). Race and gender discrimination in bargaining for a new car. *The American Economic Review* 85(3), 304–321.
- Bajari, P., L. Benkard, and J. Levin (2007, September). Estimating dynamic models of imperfect competition. *Econometrica* 5(5), 1331–1370.
- Bikhchandani, S. (1992, January). A bargaining model with incomplete information. *The Review of Economic Studies* 59(1), 187–203.
- Chatterjee, K. and L. Samuelson (1988, July-August). Bargaining under two-sided incomplete information: The unrestricted offers case. *Operations Research* 36(4), 605–618.
- Fraja, G. D. and J. Sakovics (2001). Walras retrouve: Decentralized trading mechanisms and competitive price. *Journal of Political Economy* 109(4), 842–862.
- Fudenberg, D. and D. Levine (1993, May). Self-confirming equilibrium. *Econometrica* 61(3), 523–545.
- Jaleta, M. and C. Gardebroek (2007). Farm-gate tomato price negotiations under asymmetric information. *Agricultural Economics* 36, 245–251.
- Merlo, A., F. Ortalo-Magne, and J. Rust (2008, March). The home selling problem: Theory and evidence. Working Paper, University of Maryland.
- Morton, F. S., F. Zettelmeyer, and J. Silva-Risso (2004, July). A test of bargaining theory in the auto retailing industry.
- Muthoo, A. (1991). A note on bargaining over a finite number of feasible agreements. *Economic Theory* 1, 290–292.
- Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 265–281.
- Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *Rand Journal of Economics* 38, 373–399.
- Perry, M. (1986, March). An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica* 54(2), 313–321.
- Roth, A. (1995). *Bargaining Experiments*, Chapter 4, pp. 253–348. Princeton University Press.
- Rubinstein, A. (1982, January). Perfect equilibrium in a bargaining model. *Econometrica* 50(1), 97–109.

- Rubinstein, A. (1985, September). A bargaining model with incomplete information about time preference. *Econometrica* 53(5), 1151–1172.
- Rubinstein, A. and A. Wolinsky (1985). Equilibrium in a market with sequential bargaining. *Econometrica* 53, 1133–1150.
- Rust, J. (1987, September). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica* 55(5), 999–1033.
- Satterthwaite, M. and A. Shneyerov (2007, January). Dynamic matching, two-sided incomplete information, and participation costs: existence and convergence to perfect competition. *Econometrica* 75(1), 155–200.
- Sieg, H. (2000, Oct). Estimating a bargaining model with asymmetric information: Evidence from medical malpractice disputes. *The Journal of Political Economy* 108(5), 1006–1021.
- Watanabe, Y. (2004, May). Learning and bargaining in dispute resolution: Theory and evidence from medical malpractice litigation. Mimeo, Northwestern University.
- Yan, H. and Y. Lu (2008, August). Strategic delay in bargaining with two-sided reputation: an experimental study.

Table 1: Summary Statistics of Bargaining

	Mean	Median
Number of Bargaining Interactions	2993	
Real bargaining	624	
Scripted bargaining	2369	
Percentage ending in trade:	66.51%	
Number of periods per interaction:	4.93	5.00
	(1.25)	
Total duration of an interaction: (seconds)	:55	:49
	(:36)	
Length of time between autorickshaws	3:49	2:25
	(5:14)	
Distance of trip (kms)	5.03	4.87
	(1.22)	
Driver initial offer	56.26	60.00
	(10.70)	
Passenger counteroffer	34.68	35.00
	(8.50)	
Final price if trade occurs	41.29	40.00
	(7.73)	

All details of bargaining calculated from real bargaining interactions.
Standard errors in parentheses.

Table 2: Summary Statistics of Drivers

	Mean	Median
Total number of driver interviews	678.00	
Driver age	37.44	36.00
	(10.27)	
Years spent as autorickshaw driver	10.30	8.00
	(8.51)	
Percentage renting the autorickshaw	0.35	7.50
	(8.13)	
Rental rate	166.97	186.67
	(52.42)	
Daily revenue	383.22	400.00
	(115.43)	
Daily profit (not including rental)	243.66	200.00
	(96.76)	
Percentage renting the autorickshaw	46%	
Rental rate	177.17	200.00
	(44.48)	
Daily revenue	494.31	500.00
	(130.19)	
Daily profit (not including rental)	267.15	250.00
	(103.97)	

Standard deviations in parentheses

Table 3: Players' Actions in Each Bargaining Round -
Interactions with surveyor as residual claimaint

Round	Player:	N	Action		
			Accept	Counteroffer	Exit
1	Driver	624	0%	100%	0%
2	Passenger	624	0%	100%	0%
3	Driver	617	1%	99%	0%
4	Passenger	523	10%	83%	8%
5	Driver	435	44%	34%	22%
6	Passenger	165	46%	37%	16%
7	Driver	68	65%	20%	15%
8	Passenger	15	73%	20%	7%
9	Driver	3	100%	0%	0%

Table 4: A: Estimated Driver's Parameters - Log-normal Types:

	Costs		Correlation with bargaining disutility
	Mean	Std. Deviation	
2 km	41.86 (5.48)	0.25 (6.90)	-0.01 (0.16)
3 km	42.18 (2.80)	3.03 (13.20)	-0.12 (0.15)
4 km	46.14 (2.98)	4.09 (1.96)	-0.15 (0.11)
5 km	66.14 (4.17)	13.85 (1.83)	-0.37 (0.09)
6 km	56.49 (3.18)	2.11 (2.37)	-0.06 (0.10)
7 km	49.99 (5.43)	6.71 (6.11)	-0.23 (0.13)
8 km	56.08 (4.24)	10.80 (11.86)	-0.34 (0.15)

B: Drivers' Bargaining Disutility

Mean	Std. Deviation
0.31 (0.17)	0.11 (0.19)

Table 5: A: Estimated Passengers' Parameters - Log-normal Types:

	Costs		Correlation with bargaining disutility
	Mean	Std. Deviation	
2 km	31.38 (12.31)	18.05 (8.72)	0.00 (0.21)
3 km	25.61 (8.64)	9.31 (3.95)	0.00 (0.22)
4 km	46.14 (4.57)	12.20 (4.57)	0.00 (0.13)
5 km	49.30 (4.79)	6.81 (3.84)	0.00 (0.10)
6 km	49.86 (1.74)	6.10 (1.44)	0.00 (0.07)
7 km	56.17 (5.98)	12.76 (6.99)	0.00 (0.10)
8 km	87.88 (19.84)	37.96 (21.66)	0.00 (0.25)

B: Passengers' Bargaining Disutility

Mean	Std. Deviation
0.57 (2.88E+06)	2688.75 (2.70E+21)

Table 6: Percentage of drivers with costs below meter rate

	Official meter rate			Drivers' proposed rate		
	Price	Cost- high bound	Cost- low bound	Price	Cost- high bound	Cost- low bound
2 km	rs. 23	27.78% (14.67%)	55.56% (14.98%)	rs. 29	61.11% (12.78%)	66.67% (12.94%)
3 km	rs. 29	13.66% (5.00%)	13.66% (4.97%)	rs. 36	42.44% (11.44%)	42.44% (11.40%)
4 km	rs. 35	18.60% (9.99%)	18.60% (10.04%)	rs. 43	55.14% (8.34%)	55.14% (8.37%)
5 km	rs. 41	38.75% (16.74%)	38.75% (16.74%)	rs. 50	74.79% (20.18%)	74.79% (20.13%)
6 km	rs. 47	32.60% (9.61%)	32.60% (9.62%)	rs. 57	66.89% (6.39%)	66.89% (6.38%)
7 km	rs. 53	31.78% (12.32%)	31.78% (12.42%)	rs. 64	95.33% (4.92%)	96.26% (5.02%)
8 km	rs. 59	45.24% (13.84%)	46.03% (13.65%)	rs. 71	99.21% (1.78%)	99.21% (1.78%)

Bootstrapped standard errors in parentheses

Figure 1: Market Timeline




Figure 2: Extensive Form of Bargaining Game




Figure 3: Exit Probabilities




Figure 4: Accept Probabilities




Figure 5: CDF of Counter-offers and different previous offer percentiles




Figure 6: Driver Costs




Figure 7: Driver Bargaining Disutility




Figure 8: Passenger Valuations




Figure 9: Passenger Bargaining Disutility




Figure 10: Driver Costs and Bargaining Disutility for 4 km. trips:




Figure 11: Set of Prices such that Drivers are Indifferent Between Fixed Price and Bargaining:

