Learning from Past Negotiations: Theory and Evidence from Somalian Piracy *

Attila Ambrus† Olaf de Groot‡ Shih En Lu§ Hannes Mueller¶
Matthew D. Rablen‖ Anja Shortland**

March 2025

Abstract

We analyze the time pattern of bargaining outcomes in ransom negotiations with Somali pirates, using a unique data set with a comprehensive coverage of kidnapped ships in 2002-2012. We find that, even when controlling with ship characteristics, negotiated ransoms initially increased by a large magnitude, followed by negotiation durations sharply increasing as well. In the last years of the time period, both average ransom levels and negotiation durations seemed to stabilize. We argue that the main force behind these changes was learning by the pirates about the distribution of valuations of the buyers (ship owners). To investigate this issue theoretically, we analyze a model involving a sequence of negotiations with different buyers and sellers, in which buyers' valuations are drawn independently from the same distribution, initially unknown to the sellers. Sellers observe past negotiations and update their beliefs on the distribution accordingly. We provide conditions under which over time sellers learn the true distribution of valuations. We use our model framework for structural estimations and find that pirates' beliefs over time did move closer to the true distribution of valuations, although not all the way. We use the estimated parameters of the model for welfare analysis and investigation of counterfactual scenarios such as bargaining outcomes with pirates starting out with correct beliefs.

*This paper encompasses results from a previous working paper by de Groot, Rablen and Shortland, "Barrgh-gaining with Somali Pirates." We thank Irving Arturo De Lira Salvatierra and Erdal Karahan for very valuable research assistance.
†Department of Economics, Duke University, Durham, NC 27708. E-Mail: aa231@duke.edu
‡UN Resident Coordinator's Office in Bhutan, Thimphu, Bhutan. Email: olaf.degroot@un.org
§Department of Economics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. Email: shihenl@sfu.ca
¶IAE (CSIC), MOVE and Barcelona GSE, Campus UAB, 08193 Bellaterra (Barcelona), Spain. E-Mail: h.mueller.uni@gmail.com.
‖Department of Economics, University of Sheffield, Sheffield, S1 4DT, UK. Email: m.rablen@sheffield.ac.uk
**Department of Political Economy, Kings College London, WC2R 2LS, UK. Email: anja.shortland@kcl.ac.uk

1 Introduction

Between 2002 and 2012, kidnapping ships for ransom evolved into a widespread and lucrative illegal activity in Somalia. Over time, the pirates started employing more and more sophisticated techniques and strategies, and vastly expanded the area for hijacking activities on sea. This prompted various counter-measures by ship owners and by the international community. During this period, long hijacking durations of several hundred days imposed huge welfare losses on crew, their family and ship owners.[1] Perversely, duration seemed to increase as pirates became more sophisticated.

In this paper, we analyze data on ransom negotiations that we collected from various sources. We find that there was a marked increase in negotiated ransoms between 2008 and 2010 and a stabilization at the higher level thereafter. Overall, the increase of ransom amounts in these two years was roughly twentyfold. At the same time, negotiation lengths also increased rapidly between 2008 and 2010 (with an increased spread) and also seemed to stabilize after 2010. Average duration increased from about 100 days in 2008 to about 250 days in 2010. We argue that this joint pattern, and in particular the magnitude of the observed changes, suggests that pirates initially underestimated the distribution of valuations (hence willingness to pay by owners) for the kidnapped ships, and they have been learning about this distribution over the years from outcomes of previous negotiations. The stabilization of average ransom and negotiation duration levels in the last two years of the sample is consistent with the hypothesis that, by this time, learning had occurred, and the beliefs of the pirates regarding values got close to the true distribution.

We argue that the main driving force of the observed pattern in negotiation outcomes was learning, as there were no technological or institutional changes in our sample period that could reasonably explain the enormous changes in negotiation outcomes (such as the twenty-fold increase in prices). Moreover, we gathered data on observable characteristics of kidnapped ships (type of ship, size, age, crew size, etc.), and the above patterns hold when we control for these characteristics. While the pirates might observe additional ship characteristics influencing negotiation outcomes, we find it unlikely that this could explain the order of magnitude of ransom increases in the data, since we observe the main characteristics of ships. Finally, the observed patterns suggest that the learning mostly happened on the pirates' side, as the increasing ransom clearly benefited them, and not the ship owners. This corresponds to the pirates initially facing more uncertainty regarding the distribution of valuations than the uncertainty ship owners faced regarding the outside option of the pirates. We think this is a reasonable assumption, as the pirates had very limited options other than releasing the kidnapped ship for ransom.

As a second step, we show that controlling for proxies of pirate learning, like the time period or a rolling average of past ransoms, the length of a negotiation is negatively correlated with price, which is a basic prediction of dynamic bargaining models in which there is uncertainty regarding the buyer's valuation.

Motivated by the above observations, we theoretically investigate learning in an environment with a sequence of bargaining problems. Specifically, we consider a discrete time model in which periodically a seller and a buyer are drawn to engage in a "screening" type bargaining game (Sobel & Takahashi, 1983; Fudenberg *et al.*, 1985a; Gul *et al.*, 1986) in which all offers are made by the seller.[2] We assume that the seller's outside option is publicly known (normalized to 0), but the buyer's evaluation is drawn iid from a fixed value distribution. Moreover, the seller in our model does not know this distribution. Instead, there is a set of possible distributions and a prior belief by the seller on this set. We also assume that outcomes (agreed upon price and negotiation length) of all previous negotiations are publicly known. Hence, sellers update their prior beliefs over the set of possible distributions, based on previous outcomes.[3]

The main question we address in the theoretical investigation is whether, in the above setting, sellers learn the true distribution of valuations over time, and if their expected profits converge to the level they could achieve if they knew the true distribution. Given an equilibrium sequence of offers within a negotiation,

[1] Besley *et al.* (2015) show that a transfer of about 120 million USD to pirates caused a welfare loss in excess of 640 million USD and up to 3.7 billion USD.

[2] The assumption that all offers are made by the uninformed party makes the analysis considerably simpler. See the related discussion in Kennan & Wilson (1993).

[3] We assume that buyers do know the true distribution of valuations, but the analysis would be exactly the same if the buyers need to learn the true distribution too.

different true distributions result in different probabilities of acceptance of offers in the sequence. Thus, the number of offers made during a previous negotiation provides statistical information about the true distribution. One complication is that the offer sequence, and hence the amount of statistical information revealed, is an endogenous choice in the model. Another issue is that for a given sequence of offers, there can be different distributions implying the same probabilities of acceptance in each round of the negotiation, so that the said offer sequence does not help statistically distinguish among these distributions. A special case of this is when it becomes optimal for the seller to immediately offer a price equal to the minimal buyer valuation, which occurs when his beliefs are concentrated enough around the minimum. This offer is accepted with probability 1 no matter what the true distribution of types is, in which case learning stops, and in every future negotiation, the same offer is made and accepted.

Our main theoretical results provide sufficient conditions for full learning to occur over time, in the sense of the seller's beliefs converging to the true distribution as more and more negotiations are observed in the past. If the set of possible distributions are ranked with respect to first order stochastic dominance, then a simple condition guaranteeing that it is never optimal for the seller to start the negotiations with offering the lowest possible price (equal to the minimal buyer type) is enough to guarantee full learning.

We structurally estimate our learning model using the data on negotiations with Somali pirates and a maximum likelihood approach. We make the assumption that the possible buyer valuation distributions are (truncated) lognormal, but the sellers need to learn the mean (and possibly the variance) of the distribution. We estimate both the true distribution of valuations and the pirates' beliefs in different periods of time. The estimates show that the initial expected value in the pirates' beliefs is much smaller than the true expectation, but that over time, their expectations monotonically approached the true value, although even in the final period did not fully reach it. We observe less evidence of convergence with respect to the variance of the value distribution, suggesting that learning in that dimension is more difficult and slower.

We use our estimates for welfare analysis and find that the buyers were able to keep most of the total value in the negotiations. The pirates initially only acquired a small fraction of the total value, but over time their share of the surplus increased as they learned about the higher willingness to pay of the buyers. The welfare losses due to discounting and depreciation were initially low, but they increased over time as the negotiations became longer. In a counterfactual scenario in which the pirates started out with correct beliefs, they would have been able to keep a higher share of the surplus than actually occurred even towards the end of the actual scenario, while the buyers would have received less of the total surplus. In the counterfactual scenario, welfare losses due to discounting and depreciation would have been much higher since negotiations would have been long from the start.

2 Related Literature

In the empirical literature on bargaining, Watanabe *et al.* (2006) investigates learning during dispute resolution in medical malpractice litigation. Some important differences relative to our setting are that learning occurs within a given negotiation, and it is driven by exogenous signals publicly observed by the parties. A more similar model to the one in the current paper is estimated in Ambrus *et al.* (2018), a paper investigating negotiations between slave holders in the Corsair state of Algiers and Spanish negotiators, using historic data. However, Ambrus *et al.* (2018) assumes that the distribution of valuations is commonly known, and therefore there is no learning over time by any of the parties. On a more technical level, the estimation strategy in the current paper is different because we do not directly observe the number of negotiation rounds. Other recent papers on estimating dynamic bargaining models with asymmetric information include Keniston (2011), Backus *et al.* (2020), and Larsen (2021) although they are methodologically less related to our study, as they investigate situations with asymmetric information on both sides (and abstract away from learning).[4] Lastly, Bhattacharya & Dugar (2023) in the context of a field experiment shows that pretending to have financial constraints is a successful strategy in bargaining on the part of buyers.

[4]There is also an earlier string of literature on estimating dynamic bargaining models based on US data on wage negotiations: see Fudenberg *et al.* (1985b) and Kennan & Wilson (1989).

Outside the bargaining literature, Doraszelski *et al.* (2018) empirically investigates firms learning about the demand function in a new market.

In the theoretical literature on bargaining, learning within a negotiation is analyzed in Hörner & Vieille (2009) and Kaya & Liu (2015): a long-run seller/buyer with unit supply/demand faces an infinite sequence of short-run buyers/sellers with unit demand/supply, each making a take-it-or-leave-it offer. The version of the model having some common features with our model is when short-run agents can observe all previously rejected offers. As opposed to our model, learning in these models is about the realized value of the single good to be sold.

Lee & Liu (2013) consider learning from previous negotiation outcomes, although the main features of the model and the types of questions addressed are very different from our paper. In their model there is long-run player negotiating with a sequence of short-run players. Disagreement in each negotiation triggers an uncertain outside option, which is drawn from one of two possible distributions. The distribution is originally privately known by the long-run player, and short-run players can learn about it by observing past negotiation outcomes. The main focus of the paper is examining the reputation-building incentives of the long-run player, and the economic inefficiency it causes.

There are various other contexts in which agents learn from observing past behavior, and learning can stop if agents at some point switch to actions that do not reveal further information. The most prominent example is the literature on informational cascades and herding, started by Banerjee (1992), Bikhchandani *et al.* (1992) and Welch (1992). As opposed to our model, in the above setting agents are learning about a single parameter, as opposed to a distribution function. Moreover, in our model, sellers do not receive private information, instead in every period a new public signal (potentially uninformative) is generated by a given buyer's actions. In a different setting (learning about one's self-control) Ali (2011) investigates a dynamic decision-making situation in which the decision-maker at some point can get stuck choosing menus from which the observed choice does not reveal new information for the future, and hence learning might never occur.

There is also a large literature on learning in repeated games, both regarding other players' strategies, and regarding payoff parameters of the game (for an overview, see Fudenberg & Levine (1998)). An important insight from this literature is that when playing a Bayesian Nash equilibrium of a repeated game of incomplete information about other players' payoff matrices, players eventually end up playing a Nash equilibrium of the true game (Kalai & Lehrer, 1993).

3 Background

Between April 2005 and December 2012, Somali pirates hijacked 179 ships.[5] The vast majority of these were released by pirates following the payment of ransom. The total amount of ransom money was estimated to be between US$339 – US$413mn (Yikona, 2013).

In the early period of Somali piracy, target choice was largely based on opportunity: pirates found it easiest to board slow ships with a low free-board in the Gulf of Aden or in the vicinity of the Somali coast. As pirate groups developed and responded to naval and private sector counter-piracy efforts, pirates moved from the Gulf of Aden into the Somali Basin. Attacks occurred at an ever-increasing radius from their home ports (Shortland & Vothknecht, 2011; Percy & Shortland, 2013). Pirates also became more selective in their targets, avoiding ships that had clearly adopted BMP (travelling at high speed with barbed wire and water hoses), and possibly even researching their victims beforehand (Wired, 2009; Donohue *et al.*, 2014).

Once pirates got on board a ship, the crew were used as human shields. Navies almost always drew back at this point and pirates were allowed to direct the ships to the Somali coast for ransom negotiations to begin. Ship owners and navies preferred paying ransoms to endangering crew safety, cargoes and ships in a military rescue action (CIMSEC, 2013). By indicating their readiness to negotiate a ransom with pirates, ship owners ensured that pirates in turn would keep their property intact and the crew alive (Wired, 2009; Telegraph,

[5] There were no successful hijacks after 2012.

2013; Donohue *et al.*, 2014).[6] The actual negotiation would usually take place between a representative of the ship-owner and a "negotiator" nominated by the pirate action group and its various stake-holders.[7] Occasionally a government or groups representing crew's families would take over negotiations from ship-owners who refused to negotiate or were unable to raise a ransom. Similarly pirate groups sometimes exchanged the negotiators in which they had lost faith (Donohue *et al.*, 2014; report, 2012).

In the case of Western merchant shipping and fishing vessels, the ship-owner would have obtained insurance for travelling through the region. Insurers can only reimburse the ship owner after the ransom negotiation has concluded and the ransom been paid (Bento, 2014; Time, 2009).[8] Given their financial interest in keeping ransoms low, insurers retain the services of professional ransom negotiators, who are deployed immediately as the hijack is reported to the insurers (Telegraph, 2013; Time, 2009). The negotiators are experienced in bargaining in hostage situations and in turn coach the ship-owner's representative on how to conduct the negotiation, are present to advise throughout the ransoming process and are often involved in organising the ransom drop.[9]

The negotiations broadly took the form of demand and counter-offer, with pirates making the initial demand and the ship-owner making the final (accepted) offer. In most cases the negotiation was initiated and largely conducted on the ship's own telecommunication system. This reassured the ship owner that whoever was conducting the negotiation was actually in control of the ship. Crew members could be called in to verify this and report on the health of the crew and state of the ship. Actual calls, particularly in the early stages of the negotiation would be sporadic, with pirates imposing long periods of silence to raise pressure on the ship owner. The gaps between contact points could stretch to several weeks and even months (interviews with ransom negotiators, see also chronology of the Leopard ransoming on DR (2013)). Occasionally, when pirates seemed keen to conclude, the same tactic was employed by the ship-owner, claiming that it would be "inconvenient" to schedule the next call in the near future (interview with ransom negotiator).

Usually the first demand is a multiple of the final agreed price (Bento, 2014; Wired, 2009) – though sometimes the pirates signalled their actual reservation price very early on in the negotiation (interview with a ransom negotiator). Pirates are reported to base their assessment of a realistic ransom on a number of factors: "media attention, the country to which the ship belongs, the nationality of the crew members and many other things that arise from the context at the time of negotiations. If a country tries to re-take a ship by force, or if a pirate is killed by foreign navies, the price goes up. Before doing any calculation of ransom money, it's important to know the owners and what the ship is carrying and also the nationalities of the hostages." (pirate cited in Donohue *et al.* (2014), pp.182). Bento in addition argues that grievances over illegal and unlicensed fishing raised the ransoms for fishing vessels (Bento (2014), pp.309).

There is plenty of evidence that pirate groups shared information on ransoms with each other (Donohue *et al.* (2014), pp.182). According to one pirate negotiator this could even take the form of "workshops" to discuss strategies: "The pirates are extremely good at sharing information. We know for a fact . . . the pirates have piracy workshops. Pirates of various clans, [their] elders are getting together and they will exchange information." (NPR, 2009). Pirates also easily shared information with the media, with Somalia Report often publishing a ransom figure upon release of a ship. However, in general the ransom figures obtained from "pirate sources" tends be significantly higher than the ball-park figures published by shipping industry insiders. Overall, we can assume that pirates were well informed about the ransoms paid to their own and other pirate groups. The ship-owners, on the other hand, tended to keep actual ransoms paid confidential and typically signed confidentiality agreements with the involved parties (Time, 2009; Guardian, 2014). This is in part because of the unclear legal situation of paying ransoms and thereby abetting organised crime (Bento, 2014). Although ship-owners were not officially sharing information, in practice there are only a small number of firms which specialise in hijack for ransom cases and these would have had a fairly clear

[6] While ill-treatment of crew occurred when negotiations stalled and some crew members died of underlying medical conditions or committed suicide, there were no incidents of murder.

[7] See Yikona (2013) and Do *et al.* (2013) for details on the organisation of pirate groups.

[8] In the case of bespoke kidnap and ransom insurance the pay-out would be immediate. Where ship-owners only had mandatory crew, cargo and hull insurance pay-out only occurs once the different insurers have apportioned the losses between them (interview with a ransom negotiator).

[9] As the presence of a professional ransom negotiator would indicate the financial power of an insurance company behind the shipping company, the professional negotiators usually kept a low profile.

idea of ransoms paid so far.

4 Data

Piracy incidents are reported by the International Maritime Bureau (IMB). IMB Reports are published annually and reach back to 1992. The reports deliver detailed accounts of the kind of attack, location, time and affected vessel. Attack reports contain the name of the ship, date of attack and location. The 2010 report contains, for example, the following description: *"On 05 March 2010, a Marshall Islands flagged Chemical Tanker MT UBT Ocean was attacked and hijacked by armed pirates while underway in position Latitude 09:35 South and Longitude 044:18 East about 680 NM south of Mogadishu, Somalia at approximately 0535 UTC. The pirates managed to successfully board the ship and took hostage all 21 crew members and sailed the ship to Somalia. It is believed a ransom was paid for the release of the crew and ship."*

We also went through additional sources from the UN, World Bank, Newspapers and other online sources. We used these sources to collect data on: attack and release date, ransom amount in USD, ship size (tonnage), ship type (yacht, bulk carrier, fishing vessel...) crew size and composition in nationalities. Overall there are 138 cases between 2002 and 2012 for which we have information on all of these variables. We also have the ship age in 126 cases. Summary statistics are reported in appendix Table A1.

Ransom amounts are not always publicly disclosed, making reliable estimates difficult to obtain. However, using a combination of direct contacts with ransom negotiators, primary sources within Somalia and open sources, we were able to determine ransoms for a large majority of cases. However, for a number of ships, we did not find a consistent estimate across the informal sources. In these cases, we used an average of the lower and upper bounds of the ransom reports.

Still, we lose about 30 cases due to lack of ransom data. This can have several reasons. First, there never was a ransom payment, for example, because the crew got rescued. Second, there were no credible ransom estimates. Thirdly, ship and crew were separated so that the negotiation falls outside of our model.

5 A First Look at the Data

In this section, we present our first empirical results. As a first step, we describe the raw data. We then present a simple empirical model that allows us to estimate the path of ransom levels controlling with ship and crew characteristics. Finally, we provide non-parametric estimates. Section 7 will present a full structural model of pirate negotiations.

5.1 Ransoms and Duration of Hijackings

Figure 1 provides a first look at our data. It shows the ln of ransoms by time of attack. The first thing to note is that the number of attacks increased considerably in 2005 and then again in 2008. The ln of ransoms also increased dramatically over the years from about 12 in 2006 to about 15 in 2012, which implies that ransoms had increased about twentyfold.

In Figure 2, we show that this increase in ransoms went hand-in-hand with an increase in the duration of captivity. While ships were rarely held for longer than 200 days until 2008 this subsequently changed with significant share of ships being held for a year or longer afterwards. Two observations are important here. First, we have ordered hijackings by time of attack so that this pattern cannot be explained by later releases. Secondly, there were no ships being held after 2013 so that the data is not truncated.

Finally, in Figure 3, we turn towards the relationship between duration and ln(ransoms). There is a positive association between ransoms and duration which is statistically not significant. However, as we know from Figure 1 and Figure 2, both ransoms and duration increased in later period. Our idea is that beliefs held in later period are responsible for both the increase in duration and ransom. In other words, time of attack is an omitted variable in Figure 3 and is imposing an upward bias on the observed relationship. In Figure 4, we therefore show the same relationship controlling for time of attack. The result is a striking

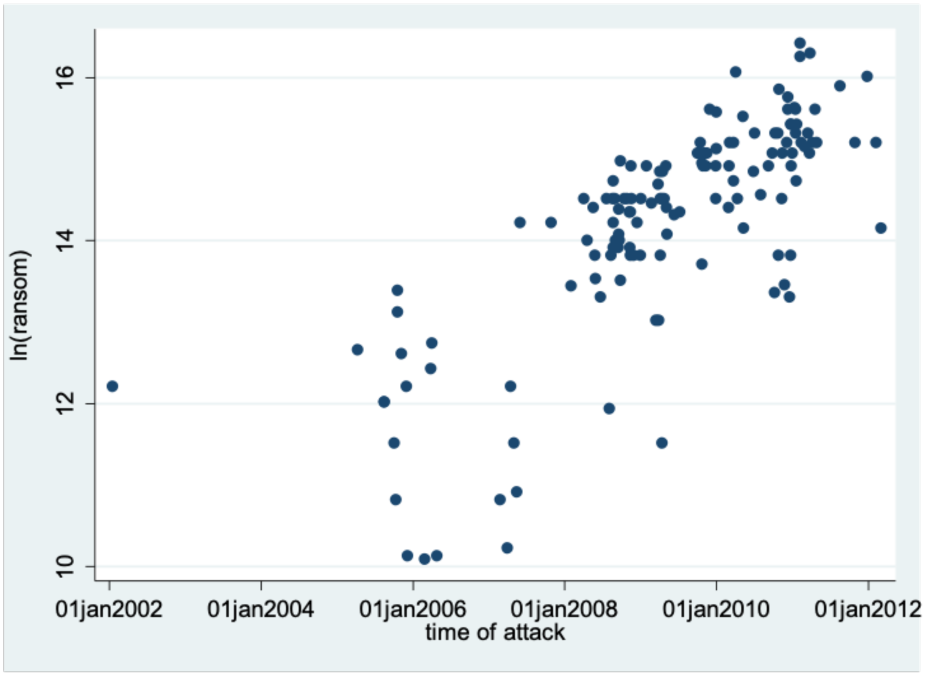

Figure 1: Ransoms and Time of Attack

turnaround of the relationship. Duration and ransoms are now negatively correlated and this relationship is statistically significant.

These patterns can be explained by a model in which pirates were updating their beliefs about the distribution of valuations throughout the years and adjusted their asking price upwards. This led to both longer negotiations on average and higher ransoms. Holding beliefs fixed, a negative relationship between ransoms and duration arises, as predicted by dynamic bargaining theory.

The main concern with this view from an empirical perspective is that ship and crew characteristics changed over the years. Put differently, the increase in ransoms could be explained with larger or more valuable ships being captured in 2010 than in 2006. Indeed, the news shock surrounding the capture of the Sirius Star in November 2008, for example, suggests that the attack on a huge oil tanker was regarded an anomaly. In what follows we therefore provide an empirical strategy to estimate how the share of the surplus the pirates were able to extract was changing over time, controlling with how ship and crew characteristics influence the surplus.

5.2 Separating Ship Value and Negotiation Outcomes

Assume that equilibrium ransoms for ship and crew i take the form

$$P_i = e^{\kappa_i} \times e^{\alpha X_i + \varepsilon_i} \tag{1}$$

where e^{κ_i} is the share of ship value that is captured by pirates in the negotiation, X_i are ship and crew characteristics and ε_i is an *iid* normal error. If we take ln on both sides we get

$$\ln P_i = \kappa_i + \alpha X_i + \varepsilon_i \tag{2}$$

so that under the assumption $E[\varepsilon_i] = 0$ we can derive a proxy for the pirate share, κ_i, from

$$\hat{\kappa}_i = \ln P_i - \hat{\alpha} X_i, \tag{3}$$

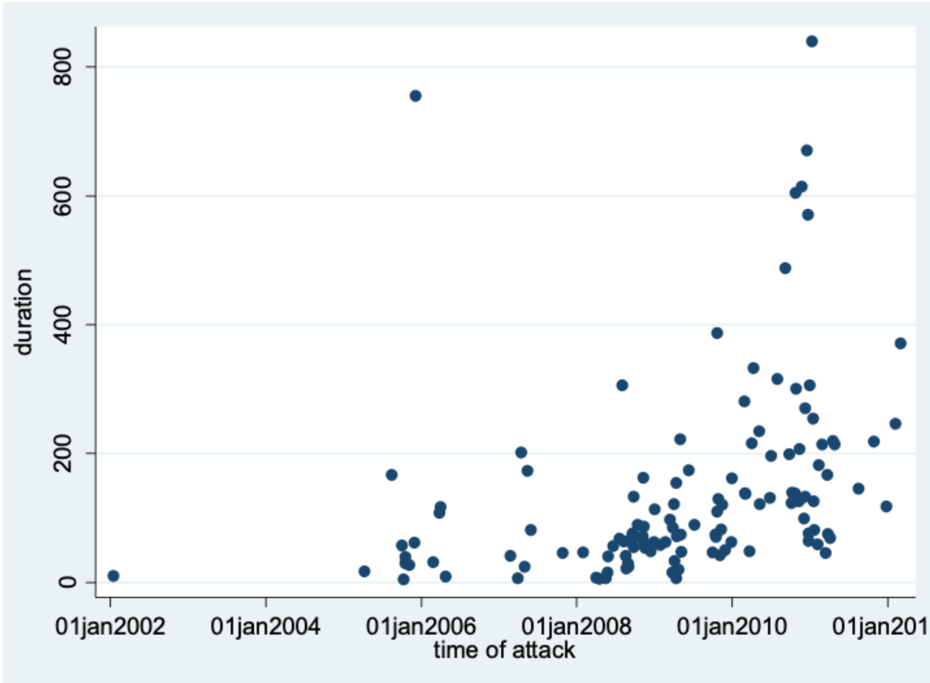

Figure 2: Captivity Duration and Time of Attack

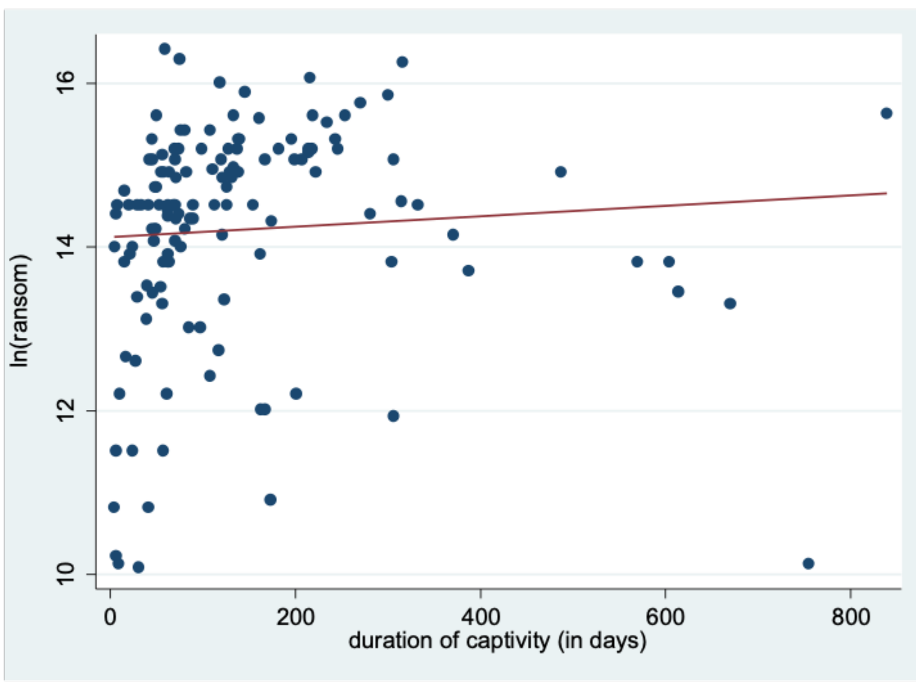

Figure 3: Ransoms and Duration Raw Data

i.e. to use the residuals of an OLS regression of $\log P_i$ on characteristics X_i. This will, of course, not be correct for each observation but it will allow us to study κ_i indirectly, for example, by looking at moving

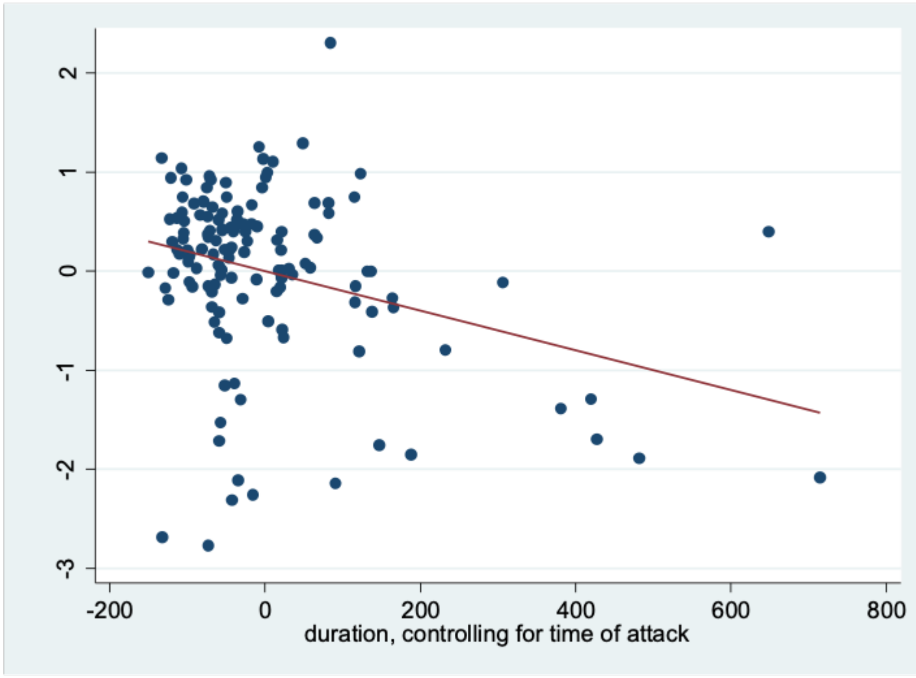

Figure 4: Ransoms and Duration (controlling for time of attack)

averages of $\hat{\kappa}_i$.

As a first step, we regress $\ln P_i$ on X_i, several specifications are reported in Table 1. The results indicate that the single most important determinant of ransoms is ship size (in tonnage). In column 2 we introduce 10 ship type dummies. We find very stark difference across ship types. Bulk carriers, chemical tankers and yachts command about 3 times the ransom than fishing vessels, for example. Crew size matters critically as well. For every crew member ransoms rise by about 3 percent. In addition, we find that different nationalities command different ransom amounts. Western European crew, for example, raise the ransom amounts by 11 percent more than other crew.

Our favorite specification is a saturated model in column 4 with a full set of ship type and crew nationality dummies. This model explains about one half of the variation in ln(ransoms). In column 5, we show that ship age, which we have for only 126 observations, is not significant when added to this model.[10] In what follows, we use the residuals $\hat{\kappa}_i$ from column 4 of Table 1. These are plotted in Figure 5. The striking feature of rising ln(ransoms) translates directly to a rising share captured by pirates, $\hat{\kappa}_i$. In other words, while there is some evidence that pirates attacked bigger and more valuable ships this alone cannot explain the drastic rise in ransoms over this period.

In Figure 6, we show rolling averages of $\hat{\kappa}_i$ and duration. This reveals an interesting comovement of these two variables that can be roughly divided in three periods. In the first period until 2008 both duration and ransoms were fairly stable. In 2008 the average ransom increases, first drastically and than at a slower (ln) rate. Duration stay stable first but then increases dramatically after January 2010 and levels off in 2011. These patterns suggest dividing the time horizon into three periods. Period 1 lasted until 2008, period 2 started January 2008 and lasted until December 2009. Period 3 started January 2010 and lasted until the last release.

The increase of $\hat{\kappa}_i$ in period 1 indicates that pirates captured an increasing share of ship and crew values. At first this did not increase duration significantly. However, in period 2 average duration increased. This is consistent with the idea that pirates had changed their strategy and increased their initial demand. This

[10] We also experimented with functional form assumptions regarding ship size and crew size. Results are robust to this.

Table 1: Explaining Ship Values

VARIABLES	(1) ln(ransom)	(2) ln(ransom)	(3) ln(ransom)	(4) ln(ransom)	(5) ln(ransom)
shipsize	1.46e-05***	8.98e-06***	5.98e-06**	8.41e-06**	8.04e-06**
	(3.31e-06)	(2.57e-06)	(2.77e-06)	(3.23e-06)	(3.55e-06)
crewsize	0.0452***	0.0336**	0.0320**	0.0229	0.0207
	(0.0136)	(0.0164)	(0.0152)	(0.0255)	(0.0298)
western european crew			0.110***		
			(0.0301)		
shipage					-0.0114
					(0.0137)
ship type dummies	no	yes	yes	yes	yes
crew count by nationalities	no	no	no	yes	yes
Observations	138	138	138	138	126
R-squared	0.205	0.397	0.425	0.458	0.342

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. "Western european" are crew from the UK, Ireland, Denmark, Spain, Germany, Italy, Greece, Spain, Norway and the Netherlands.

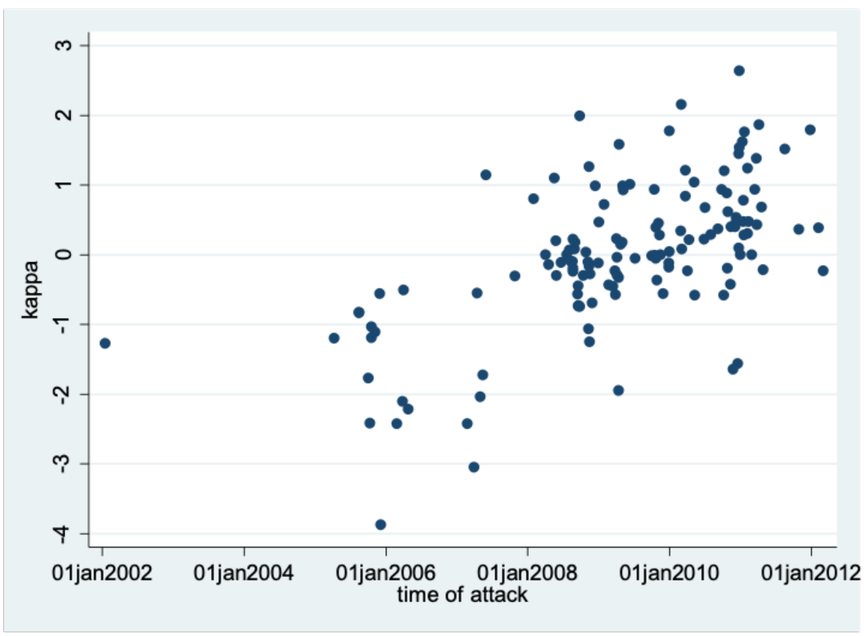

Figure 5: Estimate of Pirate Share and Time of Attack

change in strategy led to higher ransom amounts and higher duration. In 2011, asking prices stabilized and with it duration. This suggests a stabilization of pirate strategy/beliefs in this period.

5.3 Separating Out Beliefs and Negotiation Rounds

We now use the previous analysis and move one step into controlling for pirate beliefs. We can do this in two different ways. First, we can use past values of $\hat{\kappa}_i$ as a proxy for the belief environment. Second, given

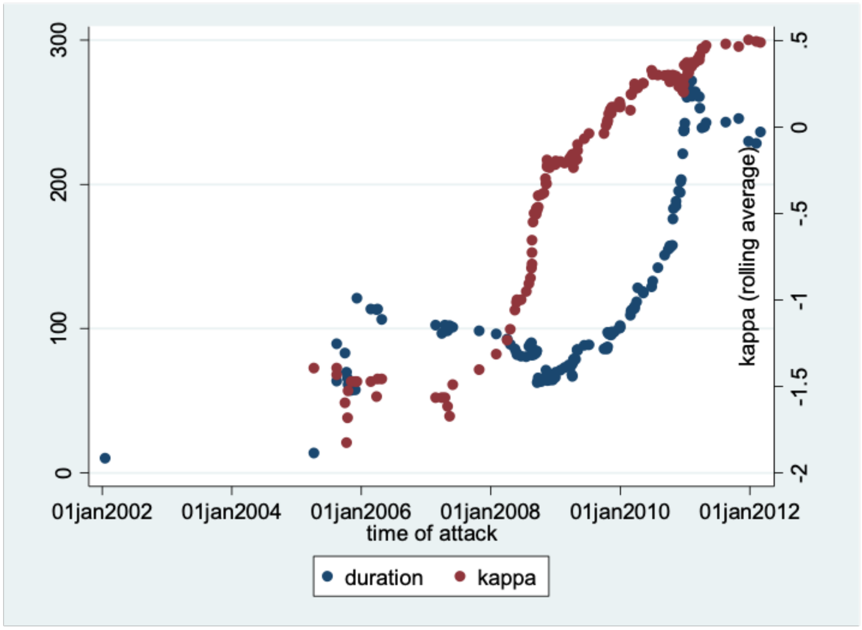

Figure 6: Rolling Averages of Duration and Kappa

the clear patterns across time, we can focus on sub-periods and assume that beliefs were constant within the period. This allows us to focus on the relationship between duration and ransom holding beliefs constant.

We first proxy beliefs by calculating rolling averages of $\hat{\kappa}_i$ and plugging these back into the regression of (2). We do this by using a rolling average of 30 attacks of $\hat{\kappa}_i$ in a regression similar to Table 1. The structure of the data complicated this step slightly. Note that capture often lasted longer than 100 days so that other ships were released during the negotiation process. This means we need to make an assumption regarding the influence of the new information contained in releases. As a default, we assume that pirates only use the information available at the time of capture, i.e. we calculate the rolling average of releases and use the average that prevailed when pirates captured a ship. An alternative is to use the rolling average at the time of release. We report results using the rolling average at time of capture in Table 2 and report the alternative in appendix Table A2.

In Table 2, Column (1) shows the same regression as Table 1 column (4) but adds the rolling mean of $\hat{\kappa}_i$ at the time of attack. Our proxy of beliefs enters positively, is significant and economically important. Ransoms increase by more than 1 percent if the average $\hat{\kappa}_i$ increases by 1 percent. In column (2) we control for the Baltic Dry Index. This is a common measure of the shipping demand which enters positively but does not affect the role of the mean significantly. Interestingly, the relationship between ransoms and the past mean also does not change if we control for the two most recent releases $\hat{\kappa}_i$ in column (3). This indicates that it is not a short term movements in $\hat{\kappa}_i$ that are driving ransoms but the mid-term movements. This is consistent with the idea that it is slow-changing pirate beliefs which are driving the share that pirates can extract in the ransom negotiation. In column (4) we control for the maximum of $\hat{\kappa}_i$ at the time of attack. This enters positively and affects the coefficient of the mean somewhat. This could imply that pirates build beliefs not only with the mean but put special attention to outliers towards the top of the distribution. However, in appendix Table A2 we show that the max does not have a robust impact in that case. Columns (4) and (5) show that duration is also positively correlated with our proxy for beliefs, irrespective of whether we control for ship characteristics or not.

As an alternative of controlling for beliefs we assume that beliefs and strategies indeed followed the three stages suggested by the data (before 2008, 2008-2009, after 2009) each with its own set of pirate beliefs Ψ_i which we then hold constant within the period. This allows us to focus on the relationship between

Table 2: The Role of Pirate Beliefs

VARIABLES	(1) ln(ransom)	(2) ln(ransom)	(3) ln(ransom)	(3) ln(ransom)	(4) duration	(5) duration
mean kappa at time of capture	1.359*** (0.137)	1.478*** (0.152)	1.174*** (0.208)	0.545** (0.225)	82.64*** (20.18)	94.54*** (30.00)
baltic dry index at time of capture		6.99e-05* (3.91e-05)				
kappa of last release			0.0160 (0.120)			
kappa two releases ago			0.168 (0.131)			
max kappa at time of attack				0.601*** (0.156)		
ship value controls	yes	yes	yes	yes	no	yes
crew value controls	yes	yes	yes	yes	no	yes
Observations	138	138	136	138	138	138
R-squared	0.751	0.760	0.753	0.779	0.133	0.315

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. "mean kappa at the time of capture" is the rolling average over 30 releases prior to the time of attack.

duration and ransom under constant pirate beliefs. Figure 7 and Figure 8 highlight these three periods in the scatterplots of $\hat{\kappa}_i$ and duration.[11]

Call the three periods discussed in the previous section $\Psi_i \in \{\Psi_1, \Psi_2, \Psi_3\}$. As an additional simplification we divide duration into brackets and assume that each of these duration brackets represents one round of negotiations n. We assume 4 brackets: 0-150 days, 151-300 days, 301-450 days and 450+ days. This allows us to estimate a regression of the form

$$\log P_i = \kappa_i (n_i, \Psi_i) + \alpha X_i + \varepsilon_i$$

where $\kappa_i (n_i, \Psi_i)$ is represented by a separate dummy for each combination of negotiation duration brackets n_i and periods Ψ_i. Since there are no observations with 3 or 4 rounds in the first period and no observations with 4 rounds in the second period we only estimate two dummies for the first period and three for the second.[12] The results are reported in Table 3. We exclude the dummy for $(n_i = 1, \Psi_3)$. All other coefficients are negative. This means that the highest share extracted by pirates was in the period after 2009 and in negotiations that lasted less than 150 days. The general pattern from these estimates is clear. Across periods the pirate share $\kappa_i (n_i, \Psi_i)$ increases. Negotiations in the first period that also lasted less than 150 days paid about $1/17th$ of what they paid in the third period ($\ln(1/17) = -2.83$). Even in the second period negotiations that lasted less than 150 days demanded significantly less ransom, about 1/3, than they did in the third period.

[11] In Figure A1 we show that this is roughly consistent with a trend-break analysis conducted on the data in Figure 4. The graph shows a positive trend break for the period around in 2008 and a negative trend break around 2010.

[12] The exception is an extreme outlier in the first period with $n = 4$ which we recode to $n = 2$.

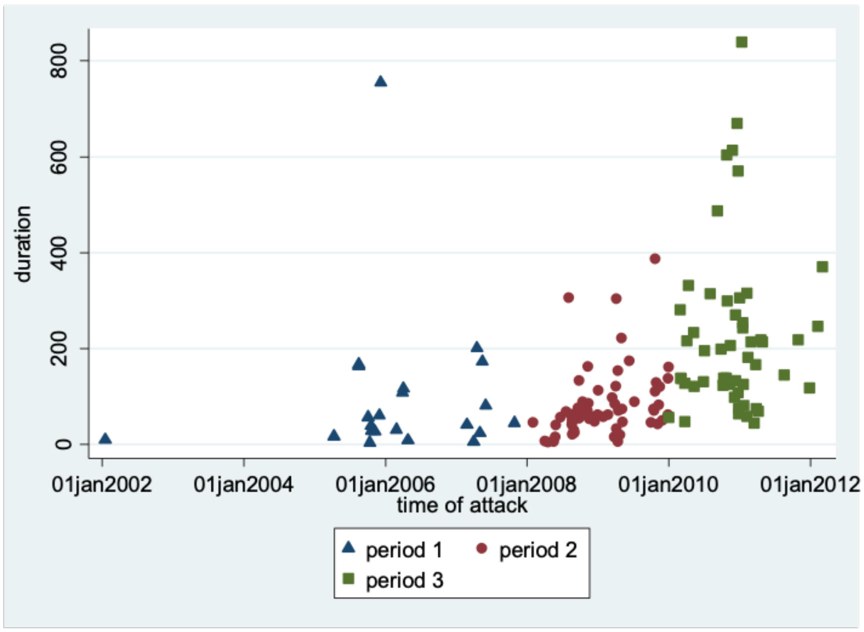

Figure 7: Duration Three Periods

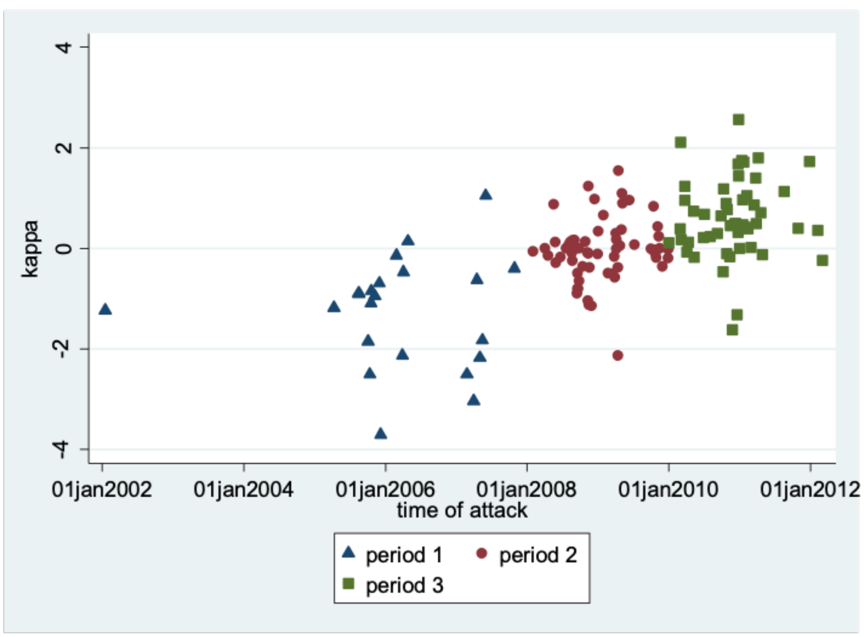

Figure 8: Kappa Three Periods

Table 3: Estimate of Pirate Share by Duration and Period

VARIABLES	(1) ln(ransom)	number of observations
n=1, Psi1	-2.828***	17
	(0.313)	
n=2, Psi1	-4.129***	5
	(0.532)	
n=1, Psi2	-1.062***	56
	(0.184)	
n=2, Psi2	-0.944***	5
	(0.301)	
n=3, Psi2	-1.266**	3
	(0.553)	
n=1, Psi3	(omitted)	24
n=2, Psi3	-0.154	17
	(0.258)	
n=3, Psi3	-0.315	5
	(0.535)	
n=4, Psi3	-1.303	6
	(0.940)	
ship value controls	yes	
crew value controls	yes	
Observations	138	
R-squared	0.831	

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

In addition, there is some evidence for a falling pirate share with a rise in n. While this pattern is not statistically significant within periods it is present in the point estimates for all periods. This is at least suggestive of a screening method used by the pirates.

Table 3 also reports the distributions of n_i for the three periods Ψ_1, Ψ_2 and Ψ_3. A clear pattern is that the fraction of negotiations with multiple rounds increased with the initial prices asked by the pirates.

6 Theoretical Investigation

6.1 A Model of Sequential Negotiations

Below, we introduce a model of sequential negotiations with the feature that negotiations are publicly observed by the participants of subsequent negotiations. Each negotiation is a dynamic bargaining game with one-sided asymmetric information, in which only the uninformed player is making offers.

Consider discrete time periods $t = 0, 1, 2, ..$, and assume that there is $K \in \mathbb{Z}_{++}$ such that at each period nK $(n = 0, 1, ...)$ a buyer and a seller are drawn to be engaged in dynamic bargaining game (negotiation) n. To simplify the analysis, we assume the following:

- K is large enough such that, in equilibrium, each negotiation ends before the next game starts (the existence of such K is guaranteed in our game), and
- no buyer or seller participates in more than one negotiation, so that players do not care about how the outcome of their negotiation impacts future negotiations.[13]

[13]In our application, this approximates the situation where for any given ship owner, the probability of being hijacked multiple

Each negotiation n is for a unit of a non-divisible good, for which the seller's valuation is commonly known to be 0, and the buyer's valuation $X_n \overset{i.i.d.}{\sim} F_0$ is the buyer's private information. The distribution F_0 is itself drawn before $t = 0$ from a collection of distributions $\mathcal{F}$. We assume that there is a commonly known prior distribution over possible valuation functions in $\mathcal{F}$, and the realized F_0 is not observed by sellers.[14] We assume that the distributions in $\mathcal{F}$:

- have a common support $[\underline{v}, \overline{v}]$, with $0 < \underline{v} < \overline{v}$, and
- exhibit continuous density functions taking values in $[\underline{d}, \overline{d}]$, with $\underline{d} > 0$.

Starting at period nK, the seller in negotiation n makes a price offer to the buyer at every period until an offer is accepted. Rejection takes the game to the next period. Parties are risk-neutral and discount future payoffs using common discount factor $\delta \in (0, 1)$. Each negotiation in our model is then equivalent to the dynamic monopoly game analyzed by Gul *et al.* (1986) if the probability that the good is sold in each period in our model is interpreted as the fraction of consumers served in that period in the dynamic monopoly game.

All actions taken in each negotiation are publicly observed. This is equivalent to only observing the price agreed in each negotiation if the sequence of offers is common knowledge, which is generically the case in equilibrium as Gul *et al.* (1986) shows that the sequence of offers is unique and deterministic.

6.2 Learning the True Distribution of Valuations

6.2.1 Finite Set of Ordered Distributions

In this section, we assume that $\mathcal{F}$ is a finite set and that the distributions in $\mathcal{F}$ are ordered according to strict first-order stochastic dominance ("FOSD ordering"): for any $F_k, F_{k'} \in \mathcal{F}$, either $F_k(v) > F_{k'}(v)$ for all $v \in (\underline{v}, \overline{v})$, or $F_k(v) < F_{k'}(v)$ for all $v \in (\underline{v}, \overline{v})$. We provide a sufficient condition under which sellers eventually learn F_0, in the sense that as the number of observed negotiations goes to infinity, the probability sellers assign to F_0 converges to 1 in probability. All proofs are provided in the Appendix.

Our main learning result is that, as long as the seller's first offer in each negotiation is bounded away from $\underline{v}$, pirates' beliefs will converge in probability to the true distribution F_0. Let f_d and F_d denote the density and c.d.f. of the distribution in $\mathcal{F}$ dominated by all other distributions in $\mathcal{F}$, and let $p_n(G)$ denote the probability placed by sellers on distribution G at the start of negotiation n.

Proposition 1. *Suppose $\mathcal{F}$ satisfies FOSD ordering. If $\max_{x \in [\underline{v}, \overline{v}]}\{[1 - F_d(\frac{x-\delta \underline{v}}{1-\delta})]x + \delta F_d(\frac{x-\delta \underline{v}}{1-\delta})\underline{v}\} > \underline{v}$, which is guaranteed if $\underline{v} f_d(\underline{v}) < 1$, then as $n \to \infty$, $\frac{p_n(F_k)}{p_n(F_0)} \xrightarrow{P} 0$ for all $F_k \in \mathcal{F}\backslash\{F_0\}$.*

We prove Proposition 1 by examining the expected log likelihood ratio between distributions F_k and F_0 for $F_k \in \mathcal{F}\backslash\{F_0\}$, computed using the outcome of a negotiation, and where the expectation is taken using F_0. First, we show that if, for any negotiation, there is a negative upper bound on this expectation, then Proposition 1 is established. We then observe that, given the strict FOSD assumption, such a bound is guaranteed by the existence of $\varepsilon > 0$ such that, for every negotiation n, there exists a threshold $\theta_n \in [\underline{v} + \varepsilon, \overline{v} - \varepsilon]$ such that a buyer with $X_n > \theta_n$ would accept a different offer than a buyer with $X_n < \theta_n$ in negotiation n. Finally, we show that this is satisfied as long as $\underline{v} f_d(\underline{v}) < 1$: intuitively, if $\underline{v}$ and/or $f(\underline{v})$ is large, then the seller may find it optimal to guarantee a first-period sale by demanding only $\underline{v}$, in which case no learning can occur.

Let s^* be the seller's optimal sequence of offers given the true distribution F_0. The next result shows that the sequence of seller offers converges in probability to s^* when the latter is unique, which is generically the case (Gul *et al.*, 1986).

times is low, while for pirates, the informational value of a single negotiation for learning about future hijacked ships' values is low.

[14]Whether buyers observe F does not impact our results.

Proposition 2. *Generically, s^* is unique. In that case, if the condition for Proposition 1 is satisfied, then the seller's equilibrium sequence of offers in negotiation n, denoted s_n, converges in probability to s^* as $n \to \infty$.*

Given the assumption that valuation distributions have well-defined densities, the seller's expected profit is continuous in the sequence of offers. Thus, Proposition 2 implies the following.

Corollary 1. *Generically, if the condition for Proposition 1 is satisfied, the seller's expected profit in negotiation n converges in probability to the expected profit she would attain if she knew F_0 as $n \to \infty$.*

6.2.2 Infinite Set of Ordered Distributions

Now suppose that the set of distributions $\mathcal{F}$ is infinite. Proposition 3, combined with the strict FOSD assumption, provides conditions under which the sellers' belief becomes arbitrary concentrated around the true distribution F_0 in probability.

Given $\eta > 0$ and a closed interval $A \subset [\underline{v}, \overline{v}]$, let $G_{\eta,A} = \{F_k : \max_{x \in A} |F_k(x) - F_0(x)| \leq \eta\}$ be the set of distributions whose c.d.f. is within η of that of the true distribution F_0 in A. As in the previous section, let f_d denote the density of the distribution in $\mathcal{F}$ dominated by all other distributions in $\mathcal{F}$.

Proposition 3. *Suppose $\mathcal{F}$ satisfies FOSD ordering, $\max_{x \in [\underline{v}, \overline{v}]} \{[1 - F_d(\frac{x - \delta \underline{v}}{1-\delta})]x + \delta F_d(\frac{x - \delta \underline{v}}{1-\delta})\underline{v}\} > \underline{v}$, and that, for any $\eta > 0$, the prior places positive probability on the set of distributions $G_{\eta,[\underline{v},\overline{v}]}$. Then there exists a closed interval $A \subset [\underline{v}, \overline{v}]$ such that as $n \to \infty$, for any $\eta > 0$, the posterior probability on $G \backslash G_{\eta,A}$ after n negotiations converges to 0 in probability.*

6.2.3 Finite Set of Non-Ordered Distributions

We return to the assumption that $\mathcal{F}$ is finite, but relax the assumption that the distributions are ordered according to FOSD. Instead, we assume:

1. the c.d.f.'s of any two distributions in $\mathcal{F}$ do not meet at more than one point in the open interval $(\underline{v}, \overline{v})$, and
2. for every distribution F with density f in $\mathcal{F}$, $v\frac{f(v)}{1-F(v)}$ is strictly increasing in v.

Condition 1 guarantees that, in a negotiation, some learning occurs in expectation as long as the seller's optimal offer sequence includes at least two offers above $\underline{v}$. Condition 2 allows us to derive a tractable condition under which this is guaranteed; it is satisfied by any distribution with increasing hazard ratio.

Proposition 4. *Suppose $\mathcal{F}$ satisfies conditions 1 and 2, and that the density f of each distribution in $\mathcal{F}$ satisfies $\underline{v}f(\underline{v}) < 1$. Let M be the highest value of $\underline{v}f(\underline{v})$ among all distributions in $\mathcal{F}$ and let v^* be such that $F(v^*) \geq M$ for all $F \in \mathcal{F}$. If $v^* \frac{f(v^*)}{1-F(v^*)} < 1$ for every distribution in $\mathcal{F}$, then as $n \to \infty$, $\frac{p_n(F_k)}{p_n(F_0)} \xrightarrow{P} 0$ for all $F_k \in \mathcal{F} \backslash \{F_0\}$.*

Therefore, as long as the hazard rate is sufficiently low for low v, the seller makes at least two offers before dropping to $\underline{v}$, which ensures continued learning.

In our structural estimation (next section), we assume that the value of ships is lognormally distributed, conditional on observable ship characteristics. We show that truncated lognormal distributions satisfy conditions 1 and 2.

Proposition 5. *Suppose every distribution in $\mathcal{F}$ is a lognormal distribution truncated at $\underline{v}$ and $\overline{v}$. Then $\mathcal{F}$ satisfies conditions 1 and 2.*

7 Structural Estimation

7.1 Econometric Model

Here we use the model framework from the previous section for structural estimations to assess the amount of learning. The model we use is similar to the one in Ambrus *et al.* (2018) (from now on ACS), with some differences described below.

We assume that the buyer's valuation of ship i with time in captivity t has the following form

$$v_{it}^b(t) = e^{-rt_i - xt_i + \boldsymbol{\alpha X_i}} e^{Z_i}$$

where Z_i is truncated normal $N(\mu, \sigma^2)$[15], $\boldsymbol{X_i}$ is a vector of observed characteristics of the ship[16], t_i denotes the length of negotiations (the time between the capture of the ship and the ransom payment), x is the depreciation rate of the ships, and r is the interest rate used for discounting. The interest rate for each period is calculated by taking the average of the LIBOR during that period. The truncation level Z_{min} determines the minimal buyer's valuation, the normalized value of which at $t = 0$ is:

$$v_{min} = e^{Z_{min}}$$

The log-likelihood function then is:

$$\mathcal{L} = -\frac{N}{2} \log \left(\frac{1}{N} \sum_i \hat{\varepsilon}_i\right)^2 + \sum_i \log \text{Prob}[n_i] \tag{4}$$

where $\hat{\varepsilon}_i$ is defined as

$$\hat{\varepsilon}_i = \log P(i, n_i, t_i) - \boldsymbol{\alpha X_i} + xt_i - \log p_{n_i}$$

Here p_{n_i} is the model predicted normalized equilibrium price, and $\text{Prob}[n_i]$ is the predicted probability that negotiations end at round n_i. Variable $P(i, n_i, t_i)$ is the observed ransom of ship i. Parameters are estimated via Maximum Likelihood. For a given value of (μ, σ), we calculate the model predicted price p_{n_i} and then estimate α by regressing $\log P(i, n_i, t_i) + xt_i - \log p_{n_i}$ on X. Depreciation, x, is fixed exogenously to 0.20 to be in line with estimates of depreciation ranging from $13\% - 30\%$.[17] Lastly, we assume that bargaining opportunities arise from a Poisson process with an average time between periods of 60 days. Thus, the Poisson process has parameter $\lambda = \frac{365}{60}$.

In order to simplify computations, we divide the data into three periods, within which beliefs about the parameters do not change, denoted μ_i for $i = 1, 2, 3$. Instead, updating occurs across periods. Lastly, for our first specification we restrict σ to have correct beliefs. Intuitively, this means the standard deviation in log prices is known by the pirates.

Fixing σ to be known by pirates does not meaningfully change our results, however it has a few benefits. First, interpretability of parameter estimates is much cleaner with only one parameter being updated. Since moments of the lognormal distribution depend on both μ and σ, even a non-monotonic trend in beliefs about each parameter could yield monotonic convergence in the moments of the distribution, while this is impossible with a fixed σ. Secondly, equilibria of the model are more stable with respect to perturbations in lognormal parameters when one is fixed. Nevertheless, we also present results allowing σ to vary, noting similar patterns in our results.

[15] We utilize the lognormal distribution due to its long tails and computational flexibility. The gamma distribution is an alternative, but more computationally burdensome than the lognormal. Additionally, this is a commonly used distribution in the auctions literature for valuations of bidders (see for example Hong & Shum (2003).

[16] X includes: ship size, number of crew members, dummies for various ship types, and dummies for having a crew member of various nationalities

[17] These values are obtained from tax authorities. These are likely to be lower bounds, as they are estimates for ships held by their owners, rather than by pirates, who may not have the same incentive or know-how to perform necessary upkeep.

7.2 Details of the Estimation

To estimate this model, we formulate the seller as solving a finite period pricing problem. The number of periods T is set equal to the maximum number of observed periods in the data. As in the Theory section, we assume the seller chooses a valuation x of buyers to make indifferent each period between buying now or waiting one more period, given the buyer types who are still considered possible (implying that all buyer types with valuation above x still considered possible in this period buy the good at this point). In the empirical model, this can be described as by the valuation which was indifferent in the prior period, which we denote y. We denote the payoff to the seller from making type x indifferent in period t when y was indifferent at $t-1$ to be $v(x, y, t)$.

Given this, the value function for the seller from having valuations y and below remaining in round t is

$$V(y,t) := \max_{0 \le x \le y} v(x,y,t)$$
$$x^*(y,t) := \arg\max_{0 \le x \le y} v(x,y,t)$$

In period T, v consists solely of the current period's expected revenue. In prior periods, we must account for the fact that the seller has a nonzero continuation value if the buyer does not purchase, $V(x, t+1)$. Additionally, the probability of purchase is $F(x)/F(y)$ where F is the cdf of buyers' valuations the seller uses to form beliefs.

Then, the interim value function takes the form

$$v(x,y,t) := \begin{cases} \left(1 - \frac{F(x)}{F(y)}\right) x, & t = T \\ \left(1 - \frac{F(x)}{F(y)}\right) P(x,t) + \beta \left(\frac{F(x)}{F(y)}\right) V(x,t+1) & 1 < t < T \\ (1 - F(x))P(x,t) + \beta F(x) V(x,y) & t = 1 \end{cases}$$

where $P(x, t)$ is the price which will make type x indifferent in period t and β is the seller's discount factor. It satisfies the equality

$$P(y,t) = (1-\delta)y + \delta P(x^*(y,t+1), t+1) \tag{5}$$

with δ being the buyer's discount factor.[18]

7.3 Results of the Structural Estimation

Parameter estimates from maximizing (4) are shown in Table 4.[19]

Table 4: Parameter Estimates with 60 Day Period

Parameter	λ	$\hat{\mu}_1$	$\hat{\mu}_2$	$\hat{\mu}_3$	$\hat{\mu}_0$	$\hat{\sigma}$	x
Estimate	6.08	10.69	12.92	14.12	16.10	1.23	0.20
	(-)	(0.19)	(0.22)	(0.23)	(0.23)	(0.02)	(-)

[18]In practice, we approximate the V terms using cubic splines and the Computational Economics toolbox for Matlab.
The likelihood function is minimized using the GlobalSearch algorithm in the Optimization Toolbox for Matlab. This tries a variety of starting points for quasi-Newton optimization algorithms. We compared results to more sophisticated algorithms, such as Simulated Annealing and Genetic Algorithms, finding similar results.

[19]Standard errors are calculated using asymptotic variance of MLE and are shown in parentheses. (-) indicate the value was calculated outside the MLE routine.

Two clear trends jump out. First, beliefs about the μ terms converge monotonically to the true values over time, suggesting pirates do indeed learn. Secondly, learning appears to be a relatively slow phenomenon. Even after many years, the beliefs of pirates differs significantly from the truth at 14.12 vs. a true value of 16.10.

We now turn our investigation to how surplus is split in this game. There are four values of interest here. First, there are the shares of surplus which go to the buyer and seller, respectively. However, we also want to know what portion of surplus is lost to depreciation and discounting. In this setting, ships physically degrade over time, lowering the total surplus accrued when an agreement is reached. In addition, there is some loss of surplus due to agreement not being reached immediately.

Table 5 shows very intuitive results. The "Predicted" lines illustrate the surplus split given pirates' estimated beliefs. In Stage 1, the buyers receive a large share of the surplus. Because sellers vastly underestimate μ_0, they price very low and sell the items very early in the bargaining game. This leads to very little loss of depreciation, while the buyers are happy to accept prices far below their willingness-to-pay. Over time, as sellers beliefs converge to the truth, we see that they achieve a larger and larger share of the surplus.

Additionally, the "Benchmark" lines show what the split of surplus would have been had pirates had correct beliefs of (μ_0, σ). As we can see, the pirates miss out on a substantial share of surplus by having incorrect beliefs. However, a large percentage of total valuation is lost to depreciation when sellers have correct beliefs, as they accurately price discriminate across bargaining rounds, leading to a substantially longer time to agreement.

Lastly, we consider the counterfactual case of what would have occurred had buyers behaved as though they had valuations (μ_1, σ). This could occur because there is a liquidity constraint on the part of buyers, restricting them from paying their full valuation. We do so by assuming that buyers agree to the sale price in each bargaining round with probability equal to the seller's expectation.

Table 5 shows that buyers lose a substantial share of surplus in the counterfactual scenario, which might seem unintuitive, as the prices they achieve are much lower. However, because the buyers are behaving as though they have much lower valuations than they do, they purchase with much lower probability in early rounds of bargaining. This means that gains from price decreases are lost to depreciation in many cases, partly due to the high amount of depreciation of ships between periods. [20]

Table 5: Split of Surplus with 60 Day Period

		Seller	Buyer	Discount	Depreciation
Stage 1	Predicted	0.15	19.95	0.16	0.66
	Benchmark	3.08	10.22	0.84	3.74
	Counterfactual	0.01	13.29	1.01	3.40
Stage 2	Predicted	1.03	18.88	0.15	0.79
	Benchmark	3.08	10.28	0.66	3.83
	Counterfactual	0.13	13.08	0.78	3.45
Stage 3	Predicted	2.08	16.86	0.10	1.46
	Benchmark	3.07	10.08	0.26	4.03
	Counterfactual	0.42	12.24	0.29	3.55

Lastly, we examine the fit of our model. For all this, a few important definitions are listed below.

- Predicted: The values the model predicts
- Observed: What we see in the data

[20] The first 5 periods are displayed here because they make up the vast majority of the data. That said, some observations do extend past this. Figures displaying the comparisons are found in the appendix.

- Benchmark: What the model would predict if pirates held beliefs (μ_0, σ)

Table 6 looks at normalized prices (that is, normalizing for attributes specific to the ship). Two facts jump out. First, the predicted prices are pretty good overall. There are some particular rounds where we do relatively poorly, but they generally have very few observations. Second, the benchmark values are very high, which is consistent with the increase in μ between pirates' beliefs and the true values.

Table 6: Predicted Offers and Observed Offers with 60 Day Period

Offer Price (in 10^6)		Offer(n=1)	Offer(n=2)	Offer(n=3)	Offer(n=4)	Offer(n=5)
Stage 1	Predicted	0.15	0.10	0.09	0.08	0.07
	Observed	0.18	0.40	0.07	0.10	-
	Benchmark	34.65	23.45	20.17	17.65	15.63
Stage 2	Predicted	1.43	0.99	0.86	0.75	0.67
	Observed	1.18	1.16	1.67	1.60	-
	Benchmark	34.39	23.74	20.58	18.12	16.13
Stage 3	Predicted	4.65	3.33	2.93	2.61	2.35
	Observed	3.98	3.96	3.18	3.02	2.97
	Benchmark	33.71	24.13	21.22	18.92	17.04

Table 7 shows that there does not appear to be a systematic error. Some purchase probabilities are overestimated, while others are underestimated, but there's no clear pattern to the errors and they are reasonably small. Again, the benchmark values are very different, due to the different pricing path.

Average Ransoms and Durations are in Table 8 and Table 9. They exhibit similar characteristics to the offers and purchase probabilities. Average duration is perhaps the worst fitting of all measures but does not have any systematic error. This suggests prices are more influential in the likelihood function than purchase probabilities.

7.4 Allowing Beliefs About σ to Vary

Here, we estimate the parameters of a model where we allow pirates' beliefs about σ to vary by period. Formally, we still divide the data into three separate periods, $i = 1, 2, 3$ and pirates have beliefs (μ_i, σ_i) about the true parameters (μ_0, σ_0). Again we assume that there is no updating within a period, and all updating occurs between periods.

The estimates from this specification are in Table 10.

The first observation here is that there is no clear pattern in the convergence of the σ terms. However, the μ terms converge monotonically over time. Additionally, the sellers' beliefs about σ in the final period are closest to the true values. This pattern could be rationalized as the sellers overestimating the variation in buyer valuations after the first period, but correcting these beliefs after the second period. This suggests that learning this feature of the value distribution is complicated and does not always move beliefs closer to the truth in any given step, but over a longer time horizon, there is progression towards the truth.

As in the case with one σ value, we examine the model fit, which appears to be quite good. Comparing the Predicted and Observed lines in Table 11 – Table 14 we see that the two line up reasonably well. Additionally, comparing to the Benchmark case, where sellers have correct beliefs, it is clear there are substantial financial losses to the sellers by underestimating the buyers' valuation distributions, though purchase would occur much more slowly.

Table 7: Accumulated Purchase Probabilities with 60 Day Period

Accumulated Agreement		Prob(n≤1)	Prob(n≤2)	Prob(n≤3)	Prob(n≤4)	Prob(n≤5)
Stage 1	Predicted	0.97	0.99	0.99	1.00	1.00
	Observed	0.59	0.77	0.91	0.95	0.95
	Benchmark	0.01	0.02	0.03	0.04	0.06
	Counterfactual	0.01	0.03	0.05	0.07	0.10
Stage 2	Predicted	0.54	0.68	0.76	0.81	0.85
	Observed	0.43	0.83	0.93	0.95	0.95
	Benchmark	0.01	0.02	0.03	0.04	0.06
	Counterfactual	0.01	0.03	0.05	0.07	0.10
Stage 3	Predicted	0.20	0.31	0.40	0.47	0.53
	Observed	0.09	0.25	0.52	0.70	0.80
	Benchmark	0.01	0.02	0.03	0.04	0.06
	Counterfactual	0.01	0.03	0.05	0.08	0.10

Table 8: Average Ransoms with 60 Day Period

	Average Normalized Random (in $\$10^6$)		
	Stage 1	Stage 2	Stage 3
Predicted	0.15	1.09	2.34
Observed	0.19	1.20	3.18
Benchmark	4.31	4.26	4.15

Table 9: Average Duration with 60 Day Period

	Average Duration (Days)		
	Stage 1	Stage 2	Stage 3
Predicted	63.11	171.66	367.50
Observed	128.18	115.00	252.86
Benchmark	408.98	400.01	379.74

Table 10: Parameter Estimates with 60 Day Period, Varying σ

Parameter	λ	$\hat{\mu}_1$	$\hat{\mu}_2$	$\hat{\mu}_3$	$\hat{\mu}_0$	$\hat{\sigma_1}$	$\hat{\sigma_2}$	$\hat{\sigma_3}$	$\hat{\sigma_0}$	x
Estimate	6.08	11.10	12.13	14.87	16.76	1.40	1.59	1.21	1.06	0.20
	(-)	(0.20)	(0.16)	(0.24)	(0.22)	(0.03)	(0.03)	(0.04)	(0.10)	(-)

Comparing average duration of bargaining games and the average ransoms paid yields similar comparisons as purchase probabilities and offers. The predictions of the model line up well with the observed distribution of the data, both in terms of normalized prices and in terms of average durations. Again, examining the benchmark case shows that sellers could achieve much higher revenue if they had correct beliefs, but would do so at the expense of taking extra bargaining rounds in expectation.

Table 11: Predicted Offers and Observed Offers with 60 Day Period, Varying σ

Offer Price (in $\$10^6$)		Offer(n=1)	Offer(n=2)	Offer(n=3)	Offer(n=4)	Offer(n=5)
Stage 1	Predicted	0.42	0.27	0.23	0.19	0.16
	Observed	0.35	0.75	0.14	0.20	-
	Benchmark	39.18	28.51	25.27	22.71	20.61
Stage 2	Predicted	2.30	1.52	1.26	1.05	0.87
	Observed	2.11	2.07	2.53	3.06	-
	Benchmark	39.49	29.07	25.88	23.35	21.27
Stage 3	Predicted	9.37	6.73	5.93	5.30	4.78
	Observed	8.36	7.97	5.80	5.67	5.51
	Benchmark	40.61	30.32	27.21	24.72	22.66

Table 12: Accumulated Purchase Probabilities with 60 Day Period, Varying σ

Accumulated Agreement		Prob(n≤1)	Prob(n≤2)	Prob(n≤3)	Prob(n≤4)	Prob(n≤5)
Stage 1	Predicted	0.97	0.99	0.99	1.00	1.00
	Observed	0.59	0.77	0.91	0.95	0.95
	Benchmark	0.01	0.03	0.05	0.08	0.10
	Counterfactual	0.07	0.12	0.17	0.22	0.26
Stage 2	Predicted	0.58	0.74	0.81	0.86	0.89
	Observed	0.43	0.83	0.93	0.95	0.95
	Benchmark	0.01	0.03	0.05	0.07	0.10
	Counterfactual	0.07	0.12	0.17	0.21	0.26
Stage 3	Predicted	0.15	0.27	0.36	0.44	0.50
	Observed	0.09	0.25	0.52	0.70	0.80
	Benchmark	0.01	0.03	0.05	0.07	0.09
	Counterfactual	0.06	0.12	0.16	0.21	0.25

As suggested in the earlier tables and similar to the case with a single σ value, Table 15 shows that buyers are gaining the largest share of the surplus in the current setting. This is due to the sellers' low beliefs about valuations, causing bargaining to end earlier and at lower prices than would occur if sellers' had correct beliefs.

Considering the counterfactual of what would happen if buyers behaved as though they had valuations drawn from (μ_1, σ_1), we can see the counterfactual surplus split in Table 15 is similar to that in the case with one σ value. Again, buyers would accrue a smaller share of the surplus as they take longer to purchase, offsetting the financial gains they would make by achieving lower prices. Given the high value of depreciation, this is rather unsurprising.

Finally, we evaluate whether allowing beliefs about σ to vary over time improves our model predictions. To do so, we measure fit using the mean squared error between model predictions and observations. The result is in Table 16.

Table 13: Average Ransoms with 60 Day Period, Varying σ

		Average Ransom (in $\$10^6$)
Stage 1	Predicted	0.42
	Observed	0.38
	Benchmark	8.03
Stage 2	Predicted	1.80
	Observed	2.12
	Benchmark	7.94
Stage 3	Predicted	4.65
	Observed	6.22
	Benchmark	7.73

Table 14: Average Duration with 60 Day Period, Varying σ

		Average Duration (Days)
Stage 1	Predicted	63.54
	Observed	128.18
	Benchmark	446.61
Stage 2	Predicted	142.09
	Observed	115.00
	Benchmark	440.78
Stage 3	Predicted	387.59
	Observed	252.86
	Benchmark	427.30

As Table Table 16 shows, there is not a strict ordering of allowing σ to vary versus not using MSE as a measure. When we allow it to vary, our purchase probabilities have a lower MSE, but prices have a higher MSE. This suggests that neither is a strictly better modeling strategy here, though this could be due to a small sample size.

8 Alternative Explanations

In this section, we discuss alternative explanations for the patterns in the data. A first alternative story would be that not just the beliefs, but the true parameters of the model were changing over time. For example, it could be possible the distribution of ship and crew values shifted upwards (and possibly spread out). It is highly unlikely that an almost twentyfold ransom increase could be explained this way given the number of controls for ship and crew value we introduce.

A second possibility is that broader technological, legal or institutional changes are responsible for the shift in ransoms. Indeed, it is true that the insurance industry made Somalia a war risk zone in May 2008 which meant that ship owners had to purchase additional insurance from that moment on onwards. Going back to the scatterplot in Figure 5 and Figure 6 is becomes clear that May 2008 is actually a very reasonable division line between periods 1 and 2. This cannot explain, however, the change from the periods 2 to 3, i.e. from 2008-2009 to 2010 onwards. Also, increases in duration are hard to explain this way.

Another alternative story would be some behavioral explanation, for example, reference-based utility for the pirates, with the reference point being the previous maximum price. Note here that, according to

Table 15: Split of Surplus with 60 Day Period, Varying σ

		Seller	Buyer	Discount	Depreciation
Stage 1	Predicted	0.40	31.66	0.26	1.05
	Benchmark	5.84	15.82	1.21	5.74
	Counterfactual	0.06	22.51	1.28	4.34
Stage 2	Predicted	1.69	30.08	0.24	1.26
	Benchmark	5.83	15.89	0.96	5.93
	Counterfactual	0.31	22.24	1.00	4.45
Stage 3	Predicted	4.05	25.41	0.18	2.82
	Benchmark	5.77	15.47	0.39	6.35
	Counterfactual	1.52	20.36	0.37	4.71

Table 16: Mean Squared Errors of Prediction with Fixed σ and Varying σ

Measure	Fixed σ	Varying σ
Price	227.34	238.34
Purchase Probabilities	1.05	0.98

Table 2, ransoms were not affected by the immediately preceding ransom outcome when controlling with the rolling average of preceding ransoms, which is evidence against a reference-point driven model.

There could also be learning on the part of ship owners. The observed patterns clearly suggest that over time pirates could retain a significantly higher share of the surplus. This suggests that learning on their part was the important contributor for the observed patterns.

References

Ali, S Nageeb. 2011. Learning self-control. *The Quarterly Journal of Economics*, **126**(2), 857–893.

Ambrus, Attila, Chaney, Eric, & Salitskiy, Igor. 2018. Pirates of the mediterranean: An empirical investigation of bargaining with asymmetric information. *Quantitative Economics*, **9**(1), 217–246.

Backus, Matthew, Blake, Thomas, Larsen, Brad, & Tadelis, Steven. 2020. Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *The Quarterly Journal of Economics*, **135**(3), 1319–1361.

Banerjee, Abhijit V. 1992. A simple model of herd behavior. *The quarterly journal of economics*, **107**(3), 797–817.

Bento, Lucas VM. 2014. Preserving negotiation whilst promoting global order: Should we bargain with salt-water devils. *Harv. Negot. L. Rev.*, **19**, 285.

Besley, Timothy, Fetzer, Thiemo, & Mueller, Hannes. 2015. The welfare cost of lawlessness: Evidence from Somali piracy. *Journal of the European Economic Association*, **13**(2), 203–239.

Bhattacharya, Haimanti, & Dugar, Subhasish. 2023. Undervaluation versus unaffordability as negotiation tactics: Evidence from a field experiment. *Journal of Economic Psychology*, **96**, 102609.

Bikhchandani, Sushil, Hirshleifer, David, & Welch, Ivo. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, **100**(5), 992–1026.

CIMSEC. 2013. *How to negotiate with pirates.*

Do, Quy-Toan, Ma, Lin, & Ruiz, Claudia. 2013. *Pirates of Somalia.* World Bank.

Donohue, William A, Pugh, Franziska, & Sabrie, Sharmaake. 2014. Somali piracy negotiations: Resolving the paradoxes of extortionate transactions. *Negotiation and Conflict Management Research*, **7**(3), 173–187.

Doraszelski, Ulrich, Lewis, Gregory, & Pakes, Ariel. 2018. Just starting out: Learning and price competition in a new market. *American Economic Review*, **108**(3), 565–615.

DR. 2013. *KRONOLOGI: Det er sket i gidselsagen.*

Fudenberg, Drew, & Levine, David K. 1998. *The theory of learning in games.* Vol. 2. MIT press.

Fudenberg, Drew, Levine, David, & Tirole, Jean. 1985a. Infinite-horizon models of bargaining with one-sided incomplete information. *Game-theoretic models of bargaining*, 73–98.

Fudenberg, Drew, Levine, David, & Ruud, Paul. 1985b. Strike activity and wage settlements. *Unpublished manuscript.*

Guardian, The. 2014. *The murky world of hostage negotiations: is the price ever right?*

Gul, Faruk, Sonnenschein, Hugo, & Wilson, Robert. 1986. Foundations of dynamic monopoly and the Coase conjecture. *Journal of economic Theory*, **39**(1), 155–190.

Hong, Han, & Shum, Matthew. 2003. Econometric models of asymmetric ascending auctions. *Journal of Econometrics*, **112**(2), 327–358.

Hörner, Johannes, & Vieille, Nicolas. 2009. Public vs. private offers in the market for lemons. *Econometrica*, **77**(1), 29–69.

Kalai, Ehud, & Lehrer, Ehud. 1993. Rational learning leads to Nash equilibrium. *Econometrica: Journal of the Econometric Society*, 1019–1045.

Kaya, Ayca, & Liu, Qingmin. 2015. Transparency and price formation. *Theoretical Economics*, **10**(2), 341–383.

Keniston, Daniel. 2011. Bargaining and welfare: A dynamic structural analysis. *Unpublished manuscript.*

Kennan, John, & Wilson, Robert. 1989. Strategic bargaining models and interpretation of strike data. *Journal of Applied Econometrics*, **4**(S1), S87–S130.

Kennan, John, & Wilson, Robert. 1993. Bargaining with private information. *Journal of Economic Literature*, **31**(1), 45–104.

Larsen, Bradley J. 2021. The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *The Review of Economic Studies*, **88**(2), 851–882.

Lee, Jihong, & Liu, Qingmin. 2013. Gambling reputation: Repeated bargaining with outside options. *Econometrica*, **81**(4), 1601–1672.

NPR. 2009. *After A Pirate Negotiation, A Personal Connection.*

Percy, Sarah, & Shortland, Anja. 2013. The business of piracy in Somalia. *Journal of strategic studies*, **36**(4), 541–578.

report, Somalia. 2012. *The Negotiators: The business side of pirate ransoms.*

Shortland, Anja, & Vothknecht, Marc. 2011. Combating "maritime terrorism" off the coast of Somalia. *European Journal of Political Economy*, **27**, S133–S151.

Sobel, Joel, & Takahashi, Ichiro. 1983. A multistage model of bargaining. *The Review of Economic Studies*, **50**(3), 411–426.

Telegraph, The. 2013. *Talk the talk: negotiation with Somali pirates.*

Time. 2009. *Why the Somali Pirates Keep Getting Their Ransoms.*

Watanabe, Yasutora, *et al.* 2006. Learning and bargaining in dispute resolution: Theory and evidence from medical malpractice litigation. *Unpublished manuscript.*

Welch, Ivo. 1992. Sequential sales, learning, and cascades. *The Journal of finance*, **47**(2), 695–732.

Wired. 2009. *Exclusive Interview: Pirate on When to Negotiate, Kill Hostages.*

Yikona, Stuart. 2013. *Pirate Trails: Tracking the Illicit Financial Flows from Pirate Activities off the Horn of Africa.* World Bank Publications.

A Appendix

A.1 Table A1, Table A2 and Figure A1

Table A1: Summary Statistics

Panel A: Main Variables

Variable	Obs	Mean	Std. Dev.	Min	Max
ransom (in USD)	138	2630191	2438301	24000	1.35E+07
duration (in days)	138	139.4565	147.9072	4	839
ship tonnage	138	16471.83	27950.35	20	162252
number of crew	138	19.73913	7.558342	2	43
number of crew from western europe	138	0.891304	2.909047	0	22
baltic dry index at time of attack	138	3178.094	2383.945	676	11465
ship age	126	17.10317	12.0365	0	60

Panel B: Ship Types

ship type	Frequency	Percent
Bulk Carrier	31	22.46
Chemical and Oil Tanker	40	28.99
Container	4	2.9
Fishing Vessel	15	10.87
General Cargo	35	25.36
Yacht	6	4.35
Other	7	5.07

Panel C: Crew Nationalities with more than 10 Hijacked

Algerian Bangladeshi Bulgarian Burmese Chinese Egyptian French Filipino Georgian Indian Indonesian Iranian Kenyan Malaysians Nigerian North-Korean Pakistani Panama Romanian Russian Spanish Srilankan

Table A2: The Role of Pirate Beliefs (by release)

VARIABLES	(1) ln(ransom)	(2) ln(ransom)	(3) ln(ransom)	(3) ln(ransom)	(4) duration	(5) duration
mean kappa at time of release	1.271***	1.295***	1.147***	1.118**	88.68***	105.5***
	(0.132)	(0.131)	(0.212)	(0.459)	(17.90)	(27.71)
baltic dry index at time of release		1.60e-05				
		(5.65e-05)				
kappa of last release			0.00389			
			(0.126)			
kappa two releases ago			0.125			
			(0.139)			
max kappa at time of release				0.112		
				(0.303)		
ship value controls	yes	yes	yes	yes	no	yes
crew value controls	yes	yes	yes	yes	no	yes
Observations	138	138	136	138	138	138
R-squared	0.751	0.752	0.749	0.752	0.183	0.365

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. "mean kappa at the time of release" is the rolling average over 30 releases prior to the time of release.

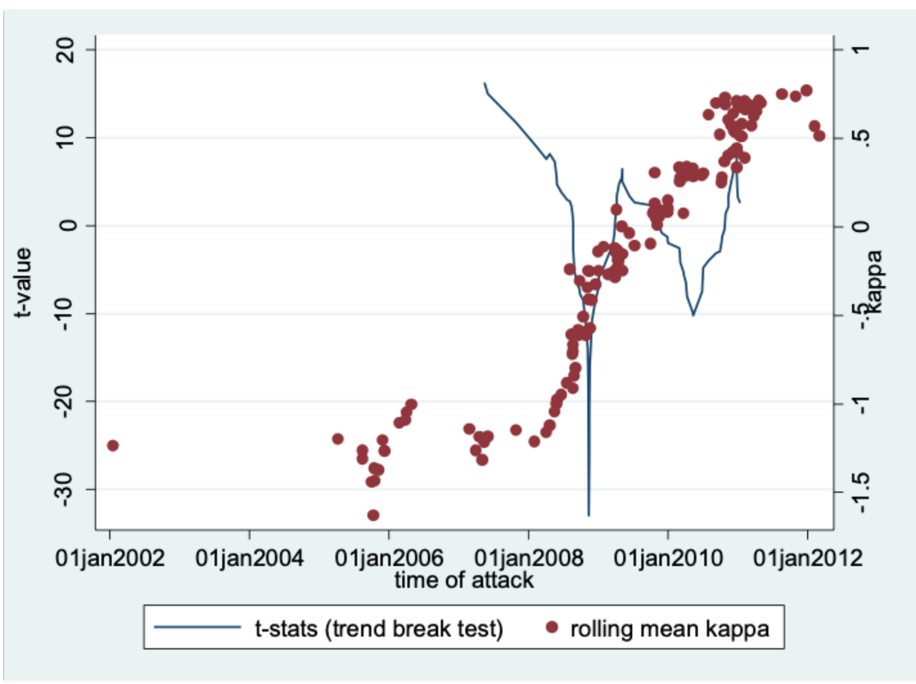

Figure A1: Clear Trend-breaks in the Data

A.2 Proofs

Proof of Proposition 1: Let P_n be the finest partition of $[\underline{v}, \overline{v}]$ such that negotiation n reveals the element of P_n that contains X_n. This is the partition of $[\underline{v}, \overline{v}]$ formed with the cutoffs corresponding to the values at which the buyer is indifferent between accepting and rejecting each offer in negotiation n.[21]

Let C_n be the element of P_n such that $X_n \in C_n$, and let $F_k(C_n)$ denote the probability that a random variable drawn according to F_k is in C_n. We start with Lemma 1.

Lemma 1: Suppose that there exists $c > 0$ such that for all $k \neq 0$ and P_n, $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)} | P_n] < -c < 0$. Then $\frac{p_n(F_k)}{p_n(F_0)} \xrightarrow{P} 0$ as $n \to \infty$.

Proof: Let $Y_n = \log \frac{F_k(C_n)}{F_0(C_n)} + (-E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)} | P_n] - c)$, so that $E[Y_n | P_n] = -c$ for all P_n. Note that $\log \frac{F_k(C_n)}{F_0(C_n)}$ is bounded by $\log \overline{d} - \log \underline{d}$, which implies that Y_n and $Var(Y_n)$ are also bounded. Claim 1 establishes that $Var(\sum_n Y_n)$ behaves as though the Y_n were independent.

Claim 1: Consider a sequence of random variables $Y_0, ..., Y_T$ with finite variances. For each $t = 0, ..., T$, let y_t be the realization of Y_t, and let $h_t = (y_0, ..., y_{t-1})$ be the history of realizations; let h_0 be the empty history. Given history h_t, Y_t is drawn according to a finite distribution $G_t^{h_t}$. For each t, assume that every $G_t^{h_t}$ has the same mean; *i.e.* $E[Y_t | h_t] = E[Y_t]$. Then:

$$Var(\sum_{t=1}^{T} Y_t) = \sum_{t=1}^{T} Var(Y_t).$$

Proof: This proof proceeds by induction. Suppose that $Var(\sum_{t=1}^{S} Y_t) = \sum_{t=1}^{S} Var(Y_t)$, which obviously holds for $S = 1$. The argument below shows that $Var(\sum_{t=1}^{S+1} Y_t) = \sum_{t=1}^{S+1} Var(Y_t)$.

We have:

$$\begin{aligned} Var(\sum_{t=1}^{S+1} Y_t) &= Var(\sum_{t=1}^{S} Y_t) + Var(Y_{S+1}) + 2Cov(\sum_{t=1}^{S} Y_t, Y_{S+1}) \\ &= \sum_{t=1}^{S+1} Var(Y_t) + 2\left(E\left[Y_{S+1} \sum_{t=1}^{S} Y_t \right] - E\left[\sum_{t=1}^{S} Y_t \right] E\left[Y_{S+1} \right] \right) \end{aligned} \tag{6}$$

Note that, letting $\Pr(h_S)$ be the *ex ante* probability of history h_S:

$$\begin{aligned} E\left[Y_{S+1} \sum_{t=1}^{S} Y_t \right] &= \sum_{h_S} \left[\Pr(h_S) \left(\sum_{t=1}^{S} y_t \right) E[Y_{S+1} | h_S] \right] \\ &= E[Y_{S+1}] \sum_{h_S} \left[\Pr(h_S) \left(\sum_{t=1}^{S} y_t \right) \right] \\ &= E[Y_{S+1}] E\left[\sum_{t=1}^{S} Y_t \right] \end{aligned} \tag{7}$$

Plugging 7 into 6 establishes the desired result. □

By Claim 1, as $n \to \infty$, the variance of $\frac{1}{n} \sum_{n'=0}^{n-1} Y_{n'}$ converges to zero. Thus, by Chebyshev's Inequality, $\frac{1}{n} \sum_{n'=0}^{n-1} Y_{n'} \xrightarrow{P} -c$.

[21] For simplicity, we assume that buyers indifferent between accepting and rejecting an offer do not mix. Note that there is a zero mass of such buyers.

Since $\log \frac{F_k(C_n)}{F_0(C_n)} < Y_n$, for any $-b > -c$, $\Pr\left(\sum_{n'=0}^{n-1} \log \frac{F_k(C_{n'})}{F_0(C_{n'})} < -bn\right) \to 1$ as $n \to \infty$. This implies $\frac{p_n(F_k)}{p_n(F_0)} = \frac{p_0(F_k)}{p_0(F_0)} \exp\left[\sum_{n'=0}^{n-1} \log \frac{F_k(C_{n'})}{F_0(C_{n'})}\right] \xrightarrow{P} 0$. □

It remains to be shown that $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)}|P_n]$ is bounded away from 0 for all possible realizations of P_n. Assume that each P_n contains at least two cells, and that each cell is a nontrivial interval. Note that $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)}|P_n] = \sum_{C \text{ is a cell of } P_n}[\log F_k(C) - \log F_0(C)]F_0(C)$. Maximizing this expression subject to $\sum_{C \text{ is a cell of } P_t} F_k(C) = 1$ yields $F_k(C) = F_0(C)$ for all C. Thus, $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)}|P_n]$ is maximized at $k = 0$ and has a maximum value of 0. Moreover, by the strict FOSD assumption, $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)}|P_n] < 0$ if $k \neq 0$.

Given the previous paragraph, it is sufficient to show that, for all P_n:
(i) the first buyer cutoff v_1 in all negotiations is bounded away from $\underline{v}$, and
(ii) the lowest buyer type $v(x)$ accepting the last price demanded x exceeding $\underline{v}$ is bounded away from $\overline{v}$.
These conditions guarantee that no sequence of partitions P_n converges to the trivial partition, which, combined with the finiteness of the family of distributions and the strict FOSD assumption, imply a uniform negative upper bound for $E_{F_0}[\log \frac{F_k(C_n)}{F_0(C_n)}|P_n]$.

Item (ii) is always satisfied: letting the last price demanded exceeding $\underline{v}$ be x, we must have $(1 - F(v(x)))x + \delta F(v(x))\underline{v} \geq \underline{v}$. Thus $F(v(x)) \leq \frac{x-\underline{v}}{x-\delta\underline{v}} \leq \frac{\overline{v}-\underline{v}}{\overline{v}-\delta\underline{v}} = 1 - \frac{(1-\delta)\underline{v}}{\overline{v}-\delta\underline{v}}$. Therefore, $v(x)$ must be at least $\frac{1}{d}\frac{(1-\delta)\underline{v}}{\overline{v}-\delta\underline{v}}$ away from $\overline{v}$.

In any negotiation, a sufficient condition for the first price demanded to exceed $\underline{v}$ is $\max_{x\in[\underline{v},\overline{v}]}\{[1 - F(\frac{x-\delta\underline{v}}{1-\delta})]x + \delta F(\frac{x-\delta\underline{v}}{1-\delta})\underline{v}\} > \underline{v}$: if the buyer has value $v > \frac{x-\delta\underline{v}}{1-\delta}$, then $v - x > \delta(v - \underline{v})$, so he will accept right away even if he believes that the next price demanded will be $\underline{v}$.
A sufficient condition for $\max_{x\in[\underline{v},\overline{v}]}\{[1 - F(\frac{x-\delta\underline{v}}{1-\delta})]x + \delta F(\frac{x-\delta\underline{v}}{1-\delta})\underline{v}\} > \underline{v}$, since the left-hand side is equal to $\underline{v}$ when $x = \underline{v}$ and the density f is continuous, is for the derivative of the left-hand side to be positive at $x = \underline{v}$, i.e. $\underline{v}f(\underline{v}) < 1$.
Let F_d be the distribution dominated by all others. If $\max_{x\in[\underline{v},\overline{v}]}\{[1 - F_d(\frac{x-\delta\underline{v}}{1-\delta})]x + \delta F_d(\frac{x-\delta\underline{v}}{1-\delta})\underline{v}\} > \underline{v}$, then if it were common knowledge that the true distribution is F_d, the seller's profit π_d would exceed $\underline{v}$. For any other distribution F that first-order stochastically dominates F_d, the seller's equilibrium profit must be at least π_d. This implies that the first price demanded, and therefore the first cutoff v_1, is bounded below by $\pi_d > \underline{v}$. Item (i) is therefore satisfied. ■

Proof of Proposition 2: By Gul *et al.* (1986), the equilibrium price path is generically unique. Suppose this holds for F_0, and suppose that the buyer best responds. Then, given a belief F about the buyer's type and a sequence of demand s, the seller's expected profit $\pi(s, F)$ has a unique maximum at s^* if $F = F_0$. Moreover, since F_0 is non-atomic, $\pi(\cdot, F_0)$ is continuous in the Euclidean metric.
Claim 2 below strengthens Gul *et al.* (1986)'s finite negotiation length result by providing a uniform bound on negotiation length in our setting.

Claim 2: In equilibrium, buyer cutoffs satisfy $v_{t+1} < v_t$ as long as $v_t > \underline{v}$, and negotiations conclude in at most $\left\lceil \frac{1}{(1-\delta)\underline{v}d} \right\rceil \equiv T$ periods.
Proof: Let z_{t+2} be the seller's expected payoff at time $t + 2$. Because, at $t + 1$, the seller can always skip demand $t + 1$, and by Gul *et al.* (1986), the continuation of on-path prices after a given on-path seller demand is unique, we have:

$$\begin{aligned}
\frac{F(v_t) - F(v_{t+1})}{F(v_t)}x_{t+1} + \delta\left(\frac{F(v_{t+1})}{F(v_t)}\right)z_{t+2} &\geq \frac{F(v_t) - F(v_{t+1})}{F(v_t)}x_{t+2} + \left(\frac{F(v_{t+1})}{F(v_t)}\right)z_{t+2} \\
[F(v_t) - F(v_{t+1})]\,(x_{t+1} - x_{t+2}) &\geq (1-\delta)F(v_{t+1})z_{t+2}
\end{aligned}$$

Note that $x_{t+1} \leq v_{t+1} \leq \underline{v} + \frac{1}{\underline{d}} F(v_{t+1})$, and $x_{t+2}, z_{t+2} \geq \underline{v}$.
Thus $[F(v_t) - F(v_{t+1})] \frac{1}{\underline{d}} F(v_{t+1}) \geq (1-\delta) F(v_{t+1}) \underline{v}$.
Therefore, if $v_{t+1} > \underline{v}$, $F(v_t) - F(v_{t+1}) \geq (1-\delta)\underline{v}\underline{d}$. $\square$

Claim 3 studies the buyer's best response to an arbitrary (*i.e.* not necessarily optimal) sequence of seller demands.

Claim 3: The buyer's best response to any planned sequence of seller demands $x_1, x_2, ...$ is given by a sequence of weakly decreasing cutoffs $v_1, v_2, ...$

Proof: The buyer wishes to maximize $\delta^t (v - x_t)$, and we need to show that if $v > v'$, then $t \leq t'$, where t is the time chosen by a buyer with value v, and t' is chosen by a buyer with value v'.
We have: $\delta^t (v - x_t) \geq \delta^{t'} (v - x_{t'})$ and $\delta^{t'} (v' - x_{t'}) \geq \delta^t (v' - x_t)$.
Adding these gives: $\delta^t v + \delta^{t'} v' \geq \delta^{t'} v + \delta^t v'$.
Thus: $v(\delta^t - \delta^{t'}) \geq v'(\delta^t - \delta^{t'})$.
Therefore $\delta^t - \delta^{t'} \geq 0$, so $t' \leq t$. $\square$

By Claim 2, for any seller belief, the seller's optimal demand sequence satisfies $s \in [\underline{v}, \overline{v}]^T$, a compact set. Therefore, for any $\kappa > 0$, $\exists \varepsilon(\kappa) > 0$ such that $\pi(s, F_0) < \pi(s^*, F_0) - \varepsilon(\kappa)$ for all s such that $|s - s^*| \geq \kappa$.

Now let two distributions F and G be within γ of each other if $|F(x) - G(x)| \leq \gamma$ for all x. Then:

- By Proposition 1, the seller's belief F converges to F_0.
- $\pi(s, .)$ is Lipschitz continuous, with constant $\overline{v}$: by Claim 3, $\pi(s, F) = \sum \delta^{t-1} x_t [F(v_{t-1}) - F(v_t)]$, where $v_0 = \overline{v}$. Note that the cutoffs v_t depend only on $s = \{x_1, ...\}$ (and not on the seller's belief F). Therefore, if $|F(x) - G(x)| \leq \gamma$ for all x, $|\pi(s, F) - \pi(s, G)| \leq \gamma(x_1 - \delta^{T-1}\underline{v}) < \gamma\overline{v}$.

It follows that when $|F - F_0| < \gamma$ and $|s - s^*| \geq \kappa$, we have: $\pi(s, F) < \pi(s, F_0) + \gamma\overline{v} < \pi(s^*, F_0) - \varepsilon(\kappa) + \gamma\overline{v} < \pi(s^*, F) - \varepsilon(\kappa) + 2\gamma\overline{v}$. Thus, if $\varepsilon(\kappa) > 2\gamma\overline{v}$, we have $\pi(s, F) < \pi(s^*, F)$, so that s cannot be optimal given F. As $\gamma \to 0$, we can take $\kappa \to 0$ and still have $\pi(s, F) < \pi(s^*, F)$ whenever $|s - s^*| \geq \kappa$. This yields Proposition 2. $\blacksquare$

Proof of Proposition 3: Fix $A = [a, b]$ such that in every negotiation, there is a buyer cutoff in A, with $a > \underline{v}$ (feasible by Proposition 1) and $b < \overline{v}$ (feasible since $\delta < 1$).

Because the distributions are ordered by strict FOSD, for any $\eta > 0$, there exists $\gamma(\eta) > 0$ such that $G \backslash G_{\eta,A} \subset \{F_k : \min_{x \in A} |F_k(x) - F_0(x)| > \gamma(\eta)\} \equiv H_{\gamma(\eta),A}$.

Given $\eta > 0$, let $F_{out} \in H_{\gamma(\eta),A}$ and $F_{in} \in G_{\gamma,A}$ for some arbitrary $\gamma > 0$. Let P be a finite interval partition of $[\underline{v}, \overline{v}]$ with at least one cell boundary within A. Then

$$\begin{aligned} & E_{F_0}\left[\log \frac{F_{out}(C)}{F_{in}(C)} | P, C \in P\right] \\ = \quad & \sum_{C \in P} [\log F_{out}(C) - \log F_0(C)] F_0(C) - \sum_{C \in P} [\log F_{in}(C) - \log F_0(C)] F_0(C) \end{aligned} \tag{8}$$

We start by seeking an upper bound for the first term $\sum_{C \in P} [\log F_{out}(C) - \log F_0(C)] F_0(C)$.

Lemma 2: $\sum_{C \in P} [\log F_{out}(C) - \log F_0(C)] F_0(C) \leq \sum_{C \in Q} [\log F_{out}(C) - \log F_0(C)] F_0(C)$, where Q is any partition that is coarser than P.

Proof of Lemma 2: Suppose that Q merges two cells of P, denoted C_1 and C_2. Let $x_k = F_{out}(C_k)$ for $k = 1, 2$, and $y_k = F_0(C_k)$ for $k = 1, 2$, and let $x = x_1 + x_2$. Then

$$\begin{aligned} & \sum_{C \in P} [\log F_{out}(C) - \log F_0(C)] F_0(C) - \sum_{C \in Q} [\log F_{out}(C) - \log F_0(C)] F_0(C) \\ = \quad & [\log x_1 - \log y_1] y_1 + [\log(x - x_1) - \log y_2] y_2 - [\log x - \log(y_1 + y_2)](y_1 + y_2) \end{aligned}$$

Maximizing this concave expression with respect to x_1 gives first-order condition

$$\frac{y_1}{x_1} = \frac{y_2}{x_2}$$

Substituting this condition into the expressives yields a value of 0, so that $\sum_{C\in P}[\log F_{out}(C)-\log F_0(C)]F_0(C) \leq \sum_{C\in Q}[\log F_{out}(C)-\log F_0(C)]F_0(C)$. Iterating the merging of cells yields the result. □

Therefore, the maximum value for $\sum_{C\in P}[\log F_{out}(C)-\log F_0(C)]F_0(C)$ is attained when P has two cells, C_1 and C_2. To find an upper bound, we maximize

$$[\log F_{out}(C_1)-\log F_0(C_1)]F_0(C_1)+[\log(1-F_{out}(C_1))-\log(1-F_0(C_1))](1-F_0(C_1))$$

subject to the constraint $|F_{out}(C_1)-F_0(C_1)|>\gamma$. The first derivative of this expression w.r.t. $F_{out}(C_1)$ is

$$\frac{F_0(C_1)}{F_{out}(C_1)}-\frac{1-F_0(C_1)}{1-F_{out}(C_1)}$$

The second derivative is

$$-\frac{F_0(C_1)}{[F_{out}(C_1)]^2}-\frac{1-F_0(C_1)}{[1-F_{out}(C_1)]^2}$$

The maximum occurs at $F_{out}(C_1)=F_0(C_1)$, where the second derivative is $-\frac{1}{F_0(C_1)}-\frac{1}{1-F_0(C_1)}\leq -4$. Therefore, for small $\gamma(\eta)$, the constraint $|F_{out}(C_1)-F_0(C_1)|>\gamma(\eta)$ implies an upper bound on the order of $\frac{1}{2}(-4)\gamma(\eta)^2=-2\gamma(\eta)^2$. For the remainder of the proof, given any $\eta>0$, we fix $\gamma(\eta)$ sufficiently small such that $G\backslash G_{\eta,A}\subset H_{\gamma(\eta),A}$ and the first term of (8) is bounded above by $-1.5\gamma(\eta)^2$.

Clearly, for any P, the second term of (8) $\sum_{C\in P}[\log F_{in}(C)-\log F_0(C)]F_0(C)\to 0$ as $\gamma\to 0$. Thus, for any $\gamma(\eta)$, there exists $\gamma>0$ such that $E_{F_0}\left[\log\frac{F_{out}(C)}{F_{in}(C)}|P,C\in P\right]<-\gamma(\eta)^2$ for any $F_{out}\in H_{\gamma(\eta),A}$, $F_{in}\in G_{\gamma,A}$ and finite interval partition P of $[\underline{v},\overline{v}]$ with at least one cell boundary within A. Since the prior places positive probability on the set of distributions $G_{\gamma,A}$, by a similar argument as in Lemma 1, the posterior on $G\backslash G_{\eta,A}\subset H_{\gamma(\eta),A}$ converges to 0 in probability as the number of observations goes to infinity. ■

Proof of Proposition 4: First, we derive two relations that will be used later. Any initial seller belief G about values is a convex combination of distributions $F\in\mathcal{F}$. Since, for any distribution $F\in\mathcal{F}$, $F(v^*)\geq \underline{v}f(\underline{v})$ for all densities f of a distribution in $\mathcal{F}$, it follows that $G(v^*)\geq \underline{v}g(\underline{v})$, where $g=G'$. Moreover, since $v^*\frac{f(v^*)}{1-F(v^*)}<1$ and $v\frac{f(v)}{1-F(v)}$ is increasing, we have $v\frac{f(v)}{1-F(v)}<1$ and thus $vf(v)<1-F(v)$ for all $v\leq v^*$. Therefore, we obtain $vg(v)<1-G(v)$ for all $v\leq v^*$.

Let the seller's stream of offers be $x_0,x_1,x_2,...$ A buyer with valuation v accepts x_t if and only if $v-x_t\geq\delta(v-x_{t+1})\iff v\geq(x_t-\delta x_{t+1})/(1-\delta)$, so the lowest type accepting x_t is $v_t=\frac{x_t-\delta x_{t+1}}{1-\delta}$. The Bellman equation corresponding to the seller's problem when choosing v_t (via the choice of x_t) is $V(v_{t-1})=\max_{v_t}\{(G(v_{t-1})-G(v_t))x_t+\delta V(v_t)\}$, where $G(v_{-1})\equiv 1$. The first-order condition is $g(v_t)x_t=(G(v_{t-1})-G(v_t))\frac{dx_t}{dv_t}+\delta V'(v_t)$. For $t=0$, we know from the proof of Proposition 1 that, since $\underline{v}g(\underline{v})<1$, the solution is interior, so the optimum must satisfy the first-order condition: $g(v_0)x_0=(1-G(v_0))\frac{dx_0}{dv_0}+\delta V'(v_0)$.

We now evaluate $V'(v_0)=(g(v_0)-g(v_1)\frac{dv_1}{dv_0})x_1+(G(v_0)-G(v_1))\frac{dx_1}{dv_0}+\delta V'(v_1)\frac{dv_1}{dv_0}$.

- If $\frac{dv_1}{dv_0}=0$, then $\frac{dx_2}{dv_0}=0$, so $\frac{dx_1}{dv_0}=0$ (since $v_t=\frac{x_t-\delta x_{t+1}}{1-\delta}$), in which case $V'(v_0)=g(v_0)x_1$.
- If not, then v_1 is interior, and:

$$\begin{aligned}V'(v_0) &= g(v_0)x_1-g(v_1)x_1\frac{dv_1}{dv_0}+(G(v_0)-G(v_1))\frac{dx_1}{dv_1}\frac{dv_1}{dv_0}+\delta V'(v_1)\frac{dv_1}{dv_0}\\ &= g(v_0)x_1+\frac{dv_1}{dv_0}(-g(v_1)x_1+(G(v_0)-G(v_1))\frac{dx_1}{dv_1}+\delta V'(v_1))\\ &= g(v_0)x_1\end{aligned}$$

where the last line follows from the $t=1$ first-order condition. Thus we have $V'(v_0)=g(v_0)x_1$.

Substituting this into the $t = 0$ first-order condition gives $g(v_0)x_0 = (1 - G(v_0))\frac{dx_0}{dv_0} + \delta g(v_0)x_1$. Using $x_0 = (1-\delta)v_0 + \delta x_1$ gives:

$$\begin{aligned} g(v_0)(1-\delta)v_0 &= (1-G(v_0))(1-\delta+\delta\frac{dx_1}{dv_0}) \\ v_0 g(v_0) &= (1-G(v_0))(1+\frac{\delta}{1-\delta}\frac{dx_1}{dv_0}) \end{aligned}$$

Since $g(v) < 1 - G(v)$ for all $v \le v^*$ and $\frac{dx_1}{dv_0} \ge 0$, it must be that $v_0 > v^*$.

The seller's belief if the first offer is rejected is then $G_1(v) = \frac{G(v)}{G(v_0)}$ for $v \in [\underline{v}, v_0]$. Thus, $\underline{v}g_1(\underline{v}) = \underline{v}\frac{g(v)}{G(v_0)} < \underline{v}\frac{g(v)}{G(v^*)} \le \underline{v}\frac{g(v)}{\underline{v}g(\underline{v})} = 1$. It follows, again from the proof of Proposition 1, that $v_1 > \underline{v}$.

Therefore, regardless of the initial belief in a negotiation, the sequence of offers contains two offers above $\underline{v}$, and the finite number of distributions guarantees that they are bounded away from each other and $\underline{v}$. The same argument as used to prove Proposition 1 establishes the result. ■

Proof of Proposition 5: To establish that families of truncated lognormal distributions satisfy condition 1, it is sufficient to show that families of truncated normal distributions satisfy condition 1 since for every lognormal random variable X, there is a normal random variable Z such that $X = e^Z$. Consider two normal distributions, F_1 and F_2, with means μ_1 and μ_2, and variances σ_1^2 and σ_2^2, respectively, and let G_1 and G_2 be the distributions obtained by truncating F_1 and F_2, respectively, at a and b, where $a < b$. Then the density of G_i is $g_i(z) = \frac{1}{F_i(b)-F_i(a)}\frac{1}{\sqrt{2\pi}\sigma_i}e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}}$. Thus, $\frac{g_1(z)}{g_2(z)} = \frac{F_2(b)-F_2(a)}{F_1(b)-F_1(a)}\frac{\sigma_2}{\sigma_1}e^{\frac{(x-\mu_2)^2}{2\sigma_2^2}-\frac{(x-\mu_1)^2}{2\sigma_1^2}}$, so that $\log(\frac{g_1(z)}{g_2(z)})$ is a quadratic function. It follows that the sign of $g_1(z) - g_2(z)$ changes either once, in which case G_1 and G_2 are ordered according to FOSD, or twice, in which case $G_1 = G_2$ exactly once in $(\underline{v}, \overline{v})$.

For condition 2, let G be a lognormal distribution (density g) where the underlying normal H (density h) has mean μ and variance σ^2, and let F (density f) be the truncation of G at a and b, where $a < b$. Then $f(v) = \frac{g(v)}{G(b)-G(a)}$ and $F(v) = \frac{G(v)-G(a)}{G(b)-G(a)}$. Thus

$$v\frac{f(v)}{1-F(v)} = \frac{\frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(\log(v)-\mu)^2}{2\sigma^2}}}{G(b)-G(v)}$$

Letting $x = \log(v)$, we have

$$v\frac{f(v)}{1-F(v)} = \frac{\frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(x-\mu)^2}{2\sigma^2}}}{H(\log b)-H(x)} = \frac{h(x)}{H(\log b)-H(x)}$$

Now let $r_k(x) = \frac{h(x)}{k-H(x)}$. We wish to show that $r_k'(x) > 0$ for $k \in (0,1)$ and x such that $H(x) < k$. We have

$$\begin{aligned} r_k'(x) &= \frac{h'(x)}{k-H(x)} + \left(\frac{h(x)}{k-H(x)}\right)^2 \\ &= \frac{1}{(k-H(x))^2}\left[h'(x)(k-H(x)) + (h(x))^2\right] \end{aligned}$$

Since the normal distribution has increasing hazard rate, we know that $r_1'(x) > 0$, so $h'(x)(1-H(x)) + (h(x))^2 > 0$. If $h'(x) < 0$, we have $h'(x)(k-H(x)) + (h(x))^2 > h'(x)(1-H(x)) + (h(x))^2 > 0$, so $r_k'(x) > 0$, as desired. The same is true by inspection if $h'(x) > 0$. ■

A.3 Figures for Section 7

A.3.1 Prices, Fixed Belief About σ

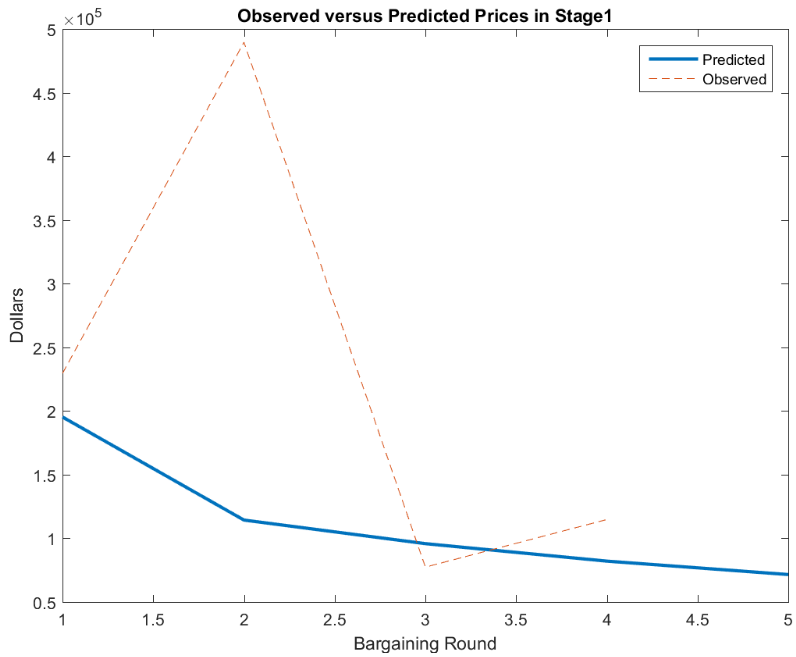

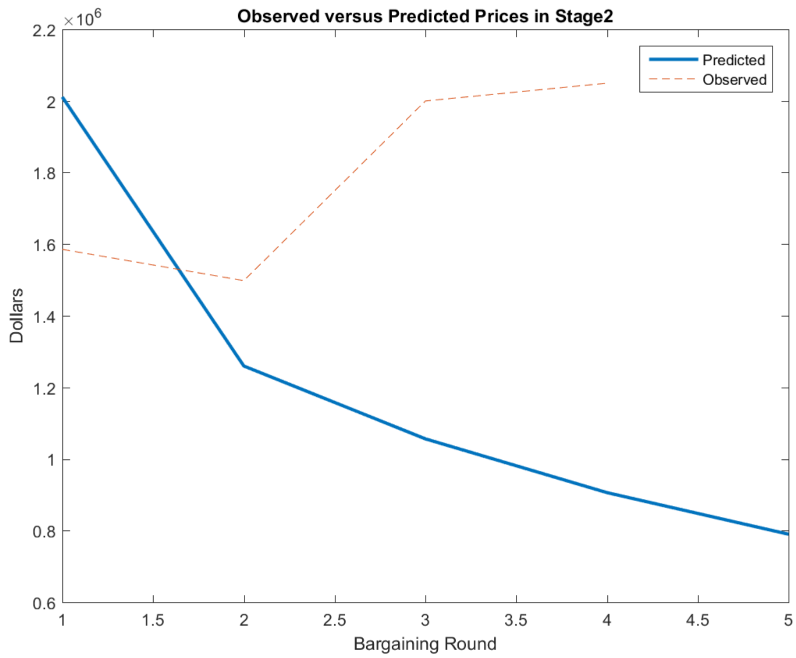

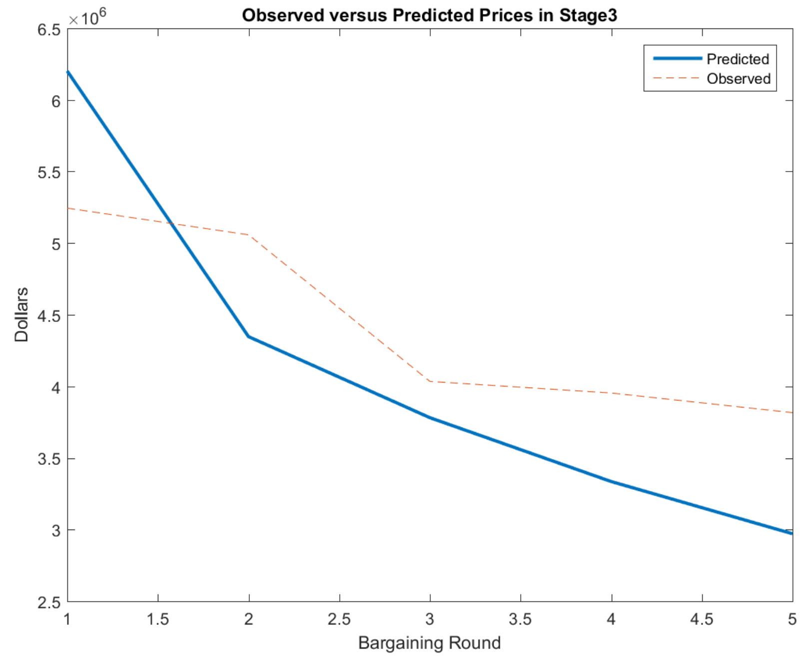

A.3.2 Purchase Probabilities, Fixed Belief About σ

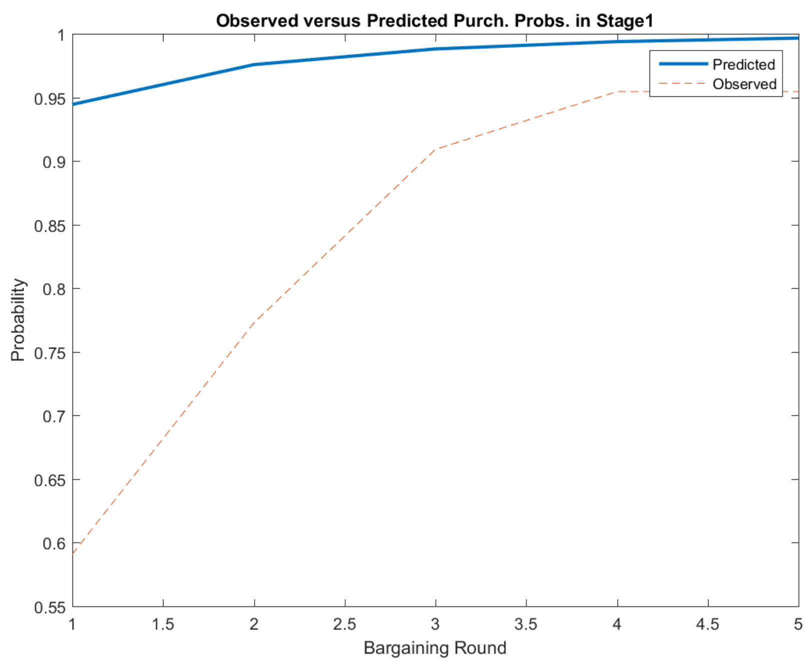

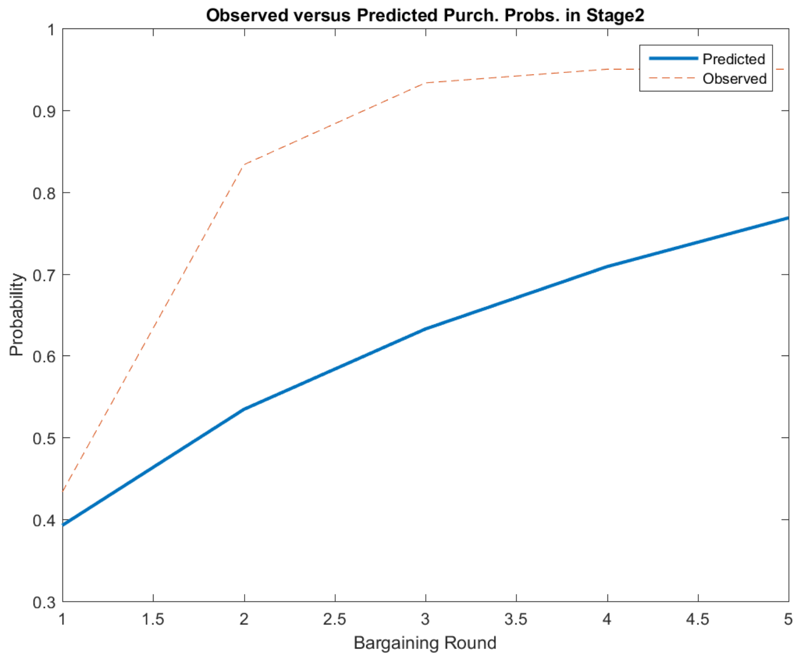
Observed versus Predicted Purch. Probs. in Stage2
Predicted
Observed
Probability
Bargaining Round
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1
1.5
2
2.5
3
3.5
4
4.5
5

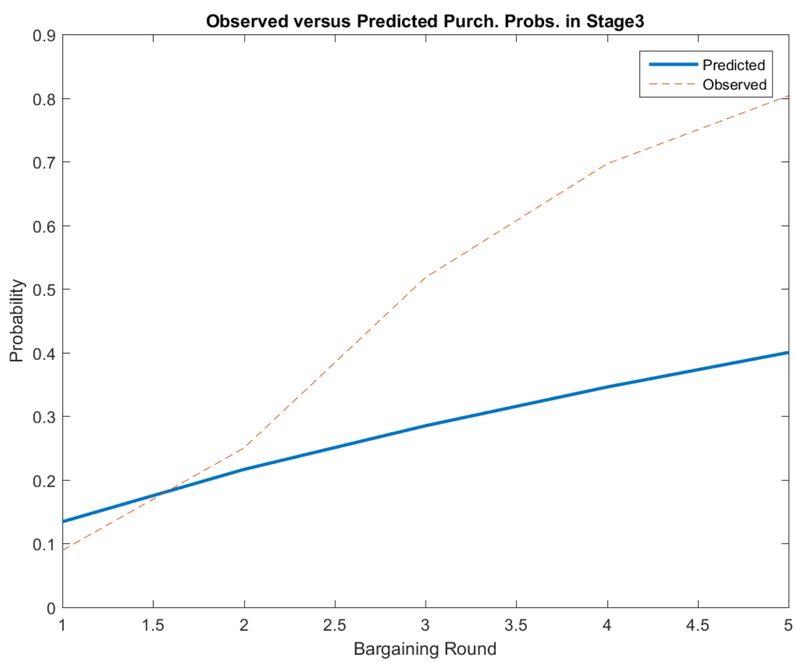
Observed versus Predicted Purch. Probs. in Stage3
Predicted
Observed
Probability
Bargaining Round
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
1.5
2
2.5
3
3.5
4
4.5
5

A.3.3 Prices, Varying Belief About σ

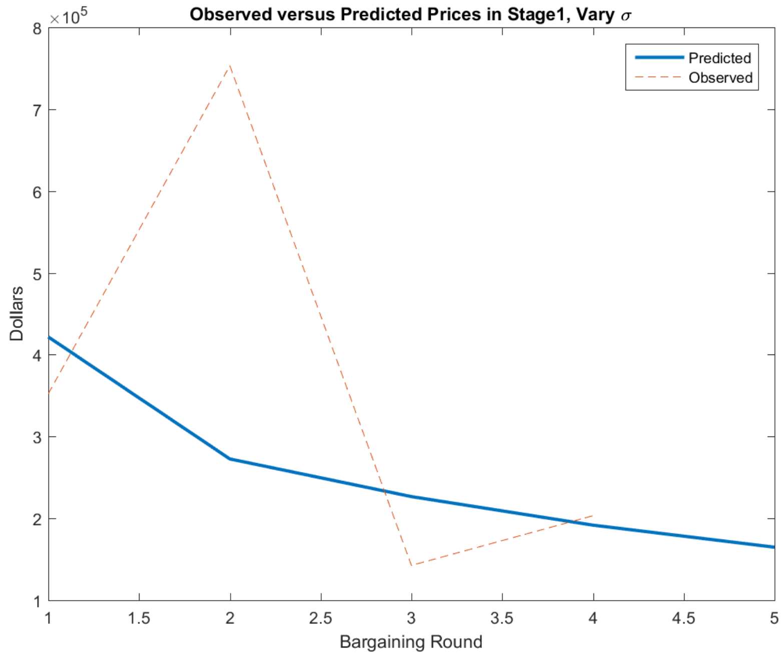

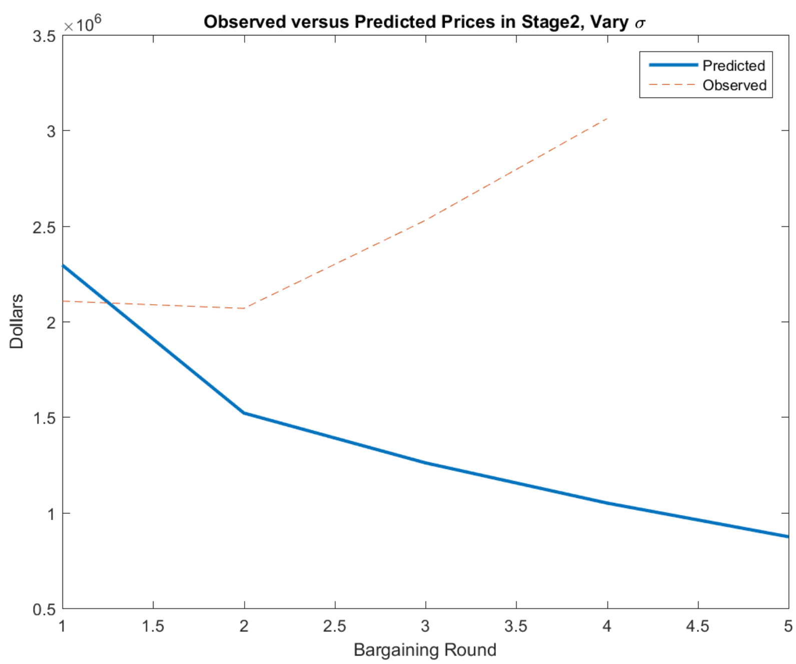

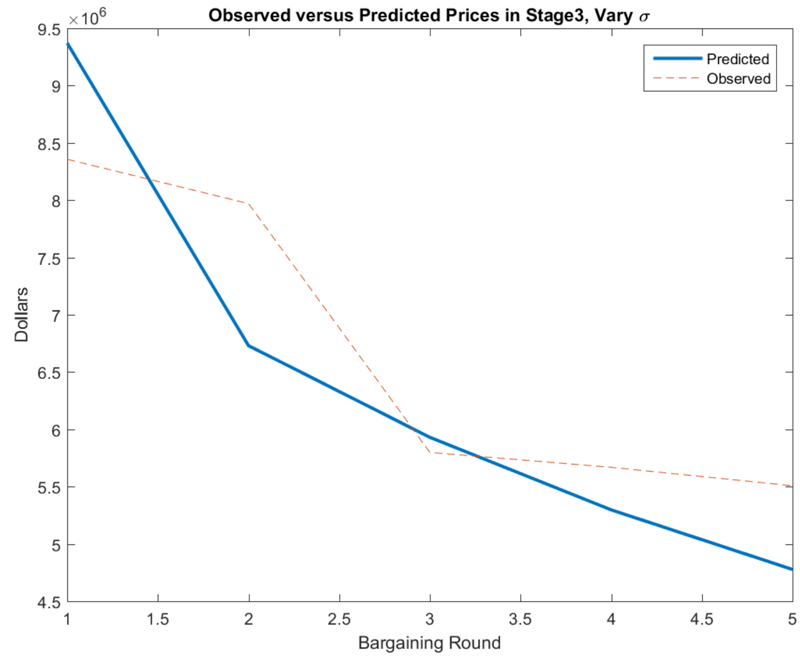

A.3.4 Purchase Probabilities, Varying Belief About σ

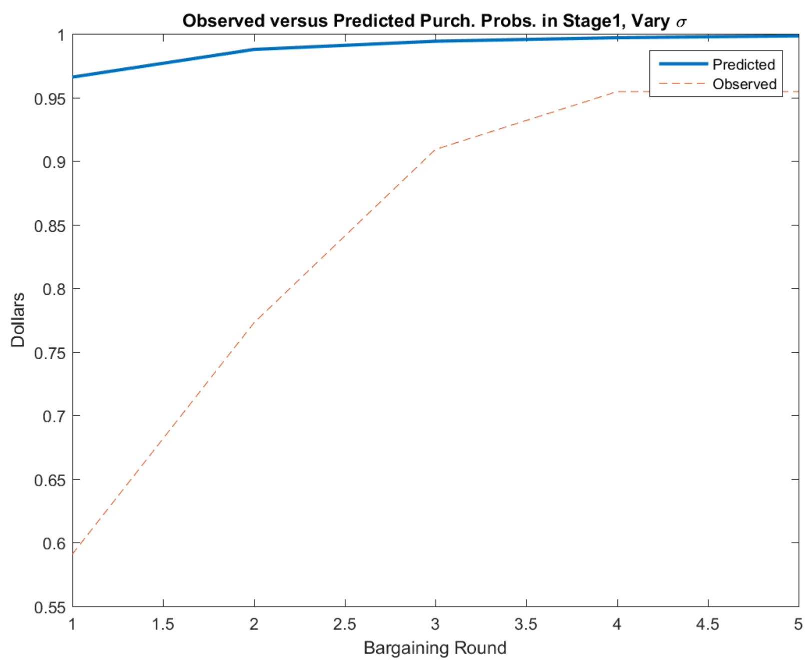

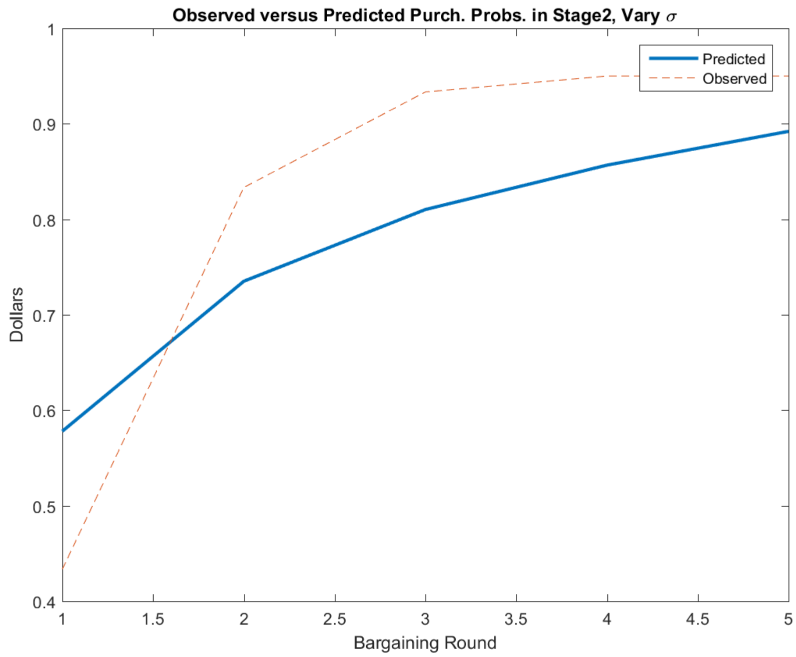

Observed versus Predicted Purch. Probs. in Stage2, Vary σ
Predicted
Observed
Dollars
Bargaining Round
1
0.9
0.8
0.7
0.6
0.5
0.4
1
1.5
2
2.5
3
3.5
4
4.5
5

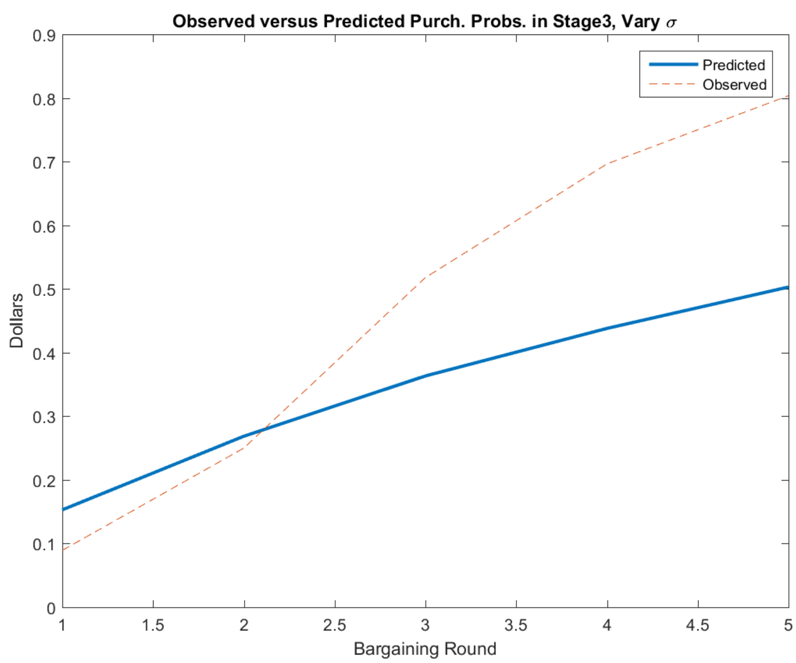

Observed versus Predicted Purch. Probs. in Stage3, Vary σ
Predicted
Observed
Dollars
Bargaining Round
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
1.5
2
2.5
3
3.5
4
4.5
5

A.4 Calculating the Split of Surplus for the Structural Estimation

To calculate the split of surplus, we perform the following steps.

First, note that

$$V = V^{buyer} + V^{seller} + C + D$$

where V is total surplus, V^i is the surplus to $i \in \{buyer, seller\}$, C is value lost to depreciation and D is value lost to discounting.

Then, we know that

$$V^{buyer} = \sum_{t=1}^{T} [E[v|v_t < v \leq v_{t-1}] - Offer_t] Pr(n = t) (\frac{\lambda}{\lambda + r + x})^t$$

$$V^{seller} = \sum_{t=1}^{T} [Offer_t] Pr(n = t) (\frac{\lambda}{\lambda + r})^t$$

$$C + D = \sum_{t=1}^{T} [(1 - (\frac{\lambda}{\lambda + r + x})^t) E[v|v_t < v \leq v_{t-1}] + (1 - (\frac{\lambda}{\lambda + r})^t) Offer_t] Pr(n = t)$$

$$C = \sum_{t=1}^{T} [(1 - (\frac{\lambda}{\lambda + r})^t) E[v|v_t < v \leq v_{t-1}] + (1 - (\frac{\lambda}{\lambda + r})^t) Offer_t] Pr(n = t)$$

Note that the difference between C and $C + D$ is that x is set to 0 when calculating the amount lost to discounting. Combined, these equations can be used to solve for the percent due to each portion of the surplus.

A.5 Empirical Distribution of the Ending Round, Structural Estimation

Table A3: Number of Observations in Each Stage-Period

Round	Stage 1	Stage 2	Stage 3
1	13	26	5
2	4	24	9
3	3	6	15
4	1	1	10
5	0	0	6
6	0	2	4
7	0	1	1
8	0	0	0
9	0	0	1
10	0	0	1
11	0	0	2
12	0	0	1
13	1	0	0
14	0	0	1