



17

Bargaining in Online Markets

Matthew Backus, Thomas Blake, and Steven Tadelis

The significance of bargaining and negotiations in economic life has motivated a vast theoretical and experimental literature to the end of understanding how those who bargain find profitable deals, divide gains from trade, and avoid bargaining breakdown. The welfare and policy implications of a better understanding of bargaining are vast—from avoiding breakdown at the negotiating table as nations forge compromises to addressing climate change, to the entrepreneur designing

M. Backus (✉)

Columbia University, New York, NY, USA

e-mail: matthew.backus@columbia.edu

T. Blake

Amazon, San Jose, CA, USA

e-mail: tomblake@gmail.com

S. Tadelis

University of California, Berkeley, Berkeley, CA, USA

e-mail: stadelis@berkeley.edu

markets in which parties bargain to realize gains from trade.¹ However, but for a few pioneering exceptions focusing on wage bargaining in labor negotiations Card [8], Card and Olson [9], Cramton and Tracy [11], that vast literature has forged ahead in the absence of guidance from the empirical analysis of bargaining in the field.

Much of the difficulty of studying bargaining in the field is inherent to the settings in which market participants elect to bargain instead of using other price-discovery mechanisms, e.g., fixed-price sales or auctions. Bargaining typically happens when both parties have some market power. If, alternatively, either side faces competition from a large and liquid set of substitutes, then there is little room for posturing, and so we would expect a fixed-price mechanism to dominate trade. This means that when two parties bargain over a deal, it is likely that the setting has some idiosyncrasies which are often difficult to quantify and compare across deals. These were the challenges of an early empirical literature that endeavored to study labor negotiations (see [7] for a summary). Negotiations between management and labor are particularly difficult to study because the terms are multidimensional, from wages to benefits and workplace conditions.²

Another challenge of studying bargaining in the field is that many negotiated contracts are set in the broader context of repeated interactions. A supplier may negotiate a rate with a downstream customer not merely with an eye to profits in that deal, but to the stream of profits they obtain by cultivating a long-term relationship. This continuation value is difficult to measure in the absence of a credible structural model of bargaining.

¹ Here we echo Crawford [12], who writes: “Bargaining, broadly construed, is a pervasive phenomenon in modern economies, ranging from labor negotiations to trade agreements to strategic arms limitation talks. One need only consider these examples in light of past experience to realize that the potential welfare gains from improving the efficiency of bargaining outcomes are enormous, perhaps even greater than those that would result from a better understanding of macroeconomic policy.”

² Strikes in labor negotiations provided a unique opportunity to study delay as signaling. This idea found little empirical application outside of labor until Goetz [15], which introduced delay in a model of bargaining adapted to negotiations between Comcast and internet ISPs for local content provision.

Even when negotiated contracts are one-off events, empirical work is plagued by the lack of available data, particularly at the offer level. It is perhaps not surprising that U.S. Attorney's Offices do not want the patterns of their settlement offers to be public knowledge, and that executives would not want it to be known that they frequently budge from their opening offers when making acquisitions. Perhaps more perplexing is the general unavailability of offer-level data in real-estate negotiations. This hurdle, the fact that offer-level data is generally unavailable to researchers, is in our opinion the most severe. For posted-price markets, there are innumerable sources of data on market outcomes, e.g., Nielsen data in retail. Moreover, there is increasing availability of data on negotiated prices in markets with bilateral bargaining, e.g., content providers selling to cable companies, insurance reimbursement rates for hospitals, and medical device pricing to hospitals. However, there are still precious few sources of step-by-step *offer-level* data in bargaining markets. Exceptions of which we are aware include a dataset of 780 real-estate transactions in the UK from Merlo and Ortalo-Magné [19] and the post-auction negotiations of Larsen [17].

In this chapter, we summarize our work in exploring a new data source on bargaining and negotiations: the eBay Best Offer platform. We use this platform as an example of the many distributed peer-to-peer marketplaces that have proliferated in recent years. Agents in these marketplaces negotiate over physical goods (e.g., eBay, Craigslist, Etsy, Alibaba), short-term labor contracts (Upwork, Freelancer), vacation rentals (AirBnband VRBO), and home services (e.g., Thumbtack) to name a few. The rise of e-commerce in general has been a boon to researchers: transactions that might previously have generated little more by way of record than a paper receipt now generate a vast trove of data. These data include not only information about the culminated transaction, such as what was sold, when, and at which price, but also the pre-transaction behavioral data concerning clicks, queries, and marketing that lead up to transactions. Users exchange structured offers over common contract terms such as price, service dates, and unit volumes. They exchange messages over less common terms (e.g., a vacation rental's use of air conditioning or a coder's choice of software language).

We begin in Sect. 17.1 with a summary of the Best Offer bargaining environment on the eBay platform. We also highlight some basic descriptive facts: some of them evidence in favor of rational theoretical models of bargaining, and some seemingly inconsistent. Next we turn to our own findings from the setting, which focus on understanding bargaining as negotiation, i.e., involving some kind of communication. In Sect. 17.2, we consider the role of cheap-talk communication in negotiation on the platform, taking as an example something that had been previously attributed to behavioral biases: the use of round numbers. Then in Sect. 17.3 we consider the broader question of communication and bargaining efficiency in Best Offer negotiations. While the applications we discuss are the first to which we turned, they are by no means the only questions that could be asked in this environment. To that end, Backus et al. [2] made one year of Best Offer bargaining data publicly available. The interested reader can find it available at <https://www.nber.org/data/bargaining/>. We hope that these applications will inspire more research in bargaining and that the data made available will prove fruitful to future bargaining scholars.

17.1 eBay's Best Offer Platform

The eBay online marketplace is best known for its auctions, which offers a price-discovery mechanism for sellers who may not know the optimal price of the object they wish to sell, and was complemented by a reputation system, for buyers to make transactions with sellers they might not otherwise trust. However, the use of auctions has long been on the decline on eBay, as a fraction of sales volume, as documented by Einav et al. [13]. The vast majority of sales volume on the platform are for listings with a fixed “Buy-it-Now” (BIN) price. For a subset of these fixed-price listings, sellers enable the “Best Offer” (BO) feature, which allows buyers to make an offer to the seller, and potentially negotiate a price lower than the asking, or BIN price.³

³ As documented by Backus et al. [2], sales that are bargained through the BO mechanism (that is, excluding BO listings that sold at the listed BIN price) have grown along with the




Fig. 17.1 Best offer on eBay (Notes This figure depicts a listing with the Best Offer feature enabled, which is why the “Make Offer” button appears underneath the “Buy It Now” and “Add to Cart” buttons. When a user clicks the Make Offer button, a panel appears, prompting an offer and, if desired, an accompanying message)

This process is depicted in Fig. 17.1, which depicts a fixed-price listing for a vintage box of cereal, on which the seller has enabled Best Offer as seen by the presence of the “Make Offer” button.⁴ The buyer can click the “Buy it Now” button and purchase the product for \$275, or they can click on the “Make Offer” button, which raises a pop-up screen with a numerical field for their offer. The buyer may also elect to add a brief text message to the seller, a feature to which we return in Sect. 17.3. Offers are valid for forty-eight hours, after which they expire and the buyer is no longer committed. The seller is notified of the offer and is given the option to accept—in which case the listing is closed and the buyer prompted for payment—reject, or counteroffer. If the seller makes

fixed-price portion of the marketplace, and by 2015 made up over 10% of sales volume on the U.S. website.

⁴ Buyers can also observe that a listing is Best Offer enabled on the search results page, before they get to the View Item page depicted in Fig. 17.1. There, underneath the price, the text “or Best Offer” may appear.

a counteroffer, this is valid for forty-eight hours as well, and the process is repeated in reverse.

To summarize, the process closely mirrors the sequential alternating-offers bargaining protocol in Rubinstein [23]. There are two differences: first, a buyer can make at most three offers.⁵ Second, all interactions begin with a buyer offer; sellers cannot make unsolicited offers to buyers and they cannot resurrect stalled or rejected threads, yet a buyer can (if they have not exhausted their three offers).

This marketplace offers a unique opportunity to study structured bargaining “in the wild,” with real buyers and sellers negotiating over real products. The advantages of this dataset are fourfold: first, unlike contract negotiations which are typically over many dimensions, the outcome variable is one-dimensional: price. Indeed, buyers are discouraged by the platform from negotiating over shipping, preferring instead that they load it into the agreed-upon price. Second, there are almost no repeated interactions between buyers and sellers, and so we can treat them as anonymous, one-off bargaining games. This is in contrast with many business-to-business bargaining environments, which are typically characterized by persistent relationships, e.g., licensing and supplier contracts, or union contract negotiations. Third, bargainers who bargain more than once can be linked over time by the econometrician, allowing us to include buyer and seller fixed effects. Fourth, the dataset is very large. The publicly released version includes over 100 million BO listings.

Unsurprisingly, however, there is also a particular limitation of this dataset—there is substantial unobserved heterogeneity coming from the vast array of products that buyers and sellers negotiate over. There is little need to bargain over homogeneous products for which there is a thick market, in which case the market price would be known. Therefore the product inventory in the dataset is particularly idiosyncratic and unlikely to be well-described by observable characteristics. This makes it difficult to learn very much about a bargainer’s reservation values or outside options.

⁵ This was true during the period when we collected our data. The cap has since been increased.

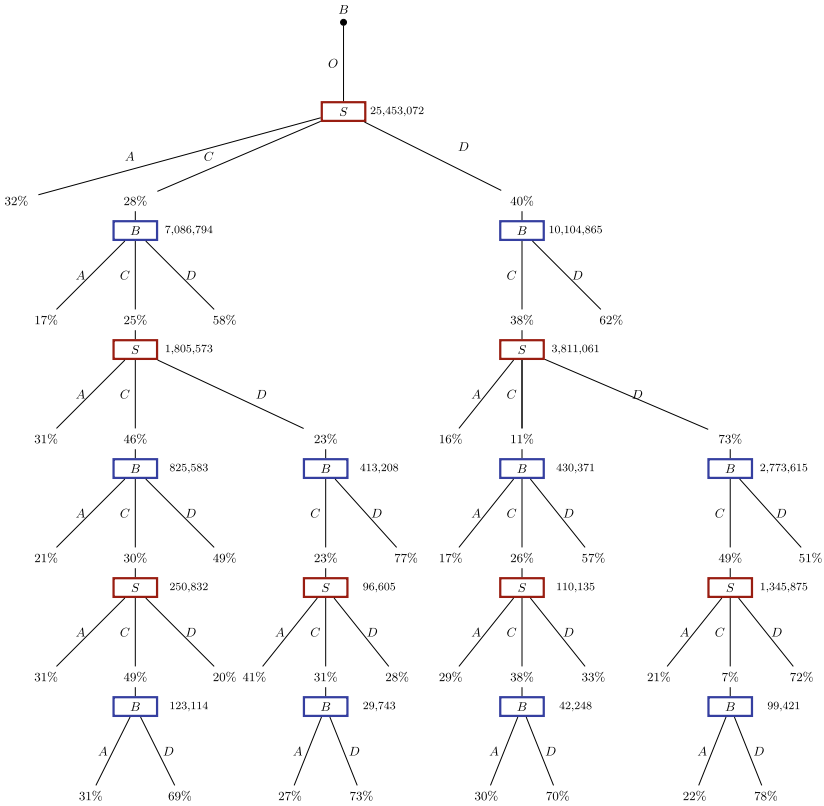


Fig. 17.2 Bargaining sequence patterns (Notes This figure summarizes the offer-level data in terms of the “game tree” of bargaining, and is borrowed from Backus et al. [2])

We begin with a few results that are immediate from the basic descriptives of the dataset. Figure 17.2, borrowed from Backus et al. [2], depicts an “empirical game tree”—the fraction of bargaining interactions that reach each node of the sequential bargaining game. Conceptually, we can think of the seller’s BIN price as an initial offer to all potentially interested buyers, the first offer in a sequence is always made by a buyer, and so the first node is “B,” after which buyers and sellers (“S”) alternate. “A” means an offer is accepted; “C” means an offer is countered, and “D” means it was declined. We report both the number and fraction

of offers that reach each node. For instance, just over 6 million cases (25%) fit the following path: an initial buyer offer is made, the seller declines (as in 40% of cases), and the buyer elects not to make another offer (as in 62% of such cases). From this figure we learn a few things. First, sellers are substantially more likely to accept an offer than buyers are. Second, we also see that frequently, if a seller declines a buyer offer (potentially allowing the offer to expire, as the dataset does not discriminate among methods of rejection), buyers re-approach the seller with a new offer (e.g., 38% of the time after a first offer is declined).

Many of the features of the dataset are broadly consistent with existing rational models of bargaining. For instance, buyers who select slower shipping methods, i.e., who may be more patient, also succeed in obtaining lower prices, consistent with the predictions of Rubinstein [23]. Of course, this is purely descriptive, and the correlation may be driven by other confounds such as a low reservation price.

Also, bargaining appears to be costly. In particular, for items that are listed at fixed prices under 50 dollars, buyers are relatively more likely to pay the seller's asking price rather than making an offer. And when the buyer does make an offer, the seller is much more likely to accept it than haggle. In particular, this supports the notion of fixed costs of bargaining. We observe that these fixed costs create qualitative differences above and below 50 dollars, which raises potential external validity concerns for the study of bargaining in laboratory settings, where stakes are low.⁶

Still other features of the data are more difficult to explain in standard theoretical models, even those that allow incomplete information and afford opportunities for bargaining breakdown. Two in particular stand out: reciprocal gradualism and split-the-difference behavior.

Backus et al. [2] document reciprocal gradualism in their dataset, which means that larger concessions by one party appear to be met with larger concessions by the other. This feature of real-life bargaining is notoriously difficult to explain in theoretical models. Compte and Jehiel [10] obtain it in the setting of baseball negotiations, but this exploits

⁶ A similar observation, that bargaining entails fixed costs, motivates the study of negotiations in appliance sales by Jindal and Newberry [16].

the unique institutional environment of that setting, i.e., that disagreement payoffs depend on the sequence of offers. It may be possible to obtain a similar result if bargainers have reference-dependent preferences, however, this conjecture is unproven.⁷

A second robust and puzzling feature of the data is the predominance of so-called “splitting the difference.” The puzzle has two features. The first, which is perhaps less surprising, is the fact of the behavior: that bargaining parties are discontinuously more likely to make an offer that is halfway between the two prior offers. The second, however, is that these offers appear to work, and introduce a non-monotonicity in the empirical relationship between the generosity of offers and the frequency with which they are accepted. That is, offers slightly higher than 50% (e.g., 55% of the other party’s most recent ask) are less likely to be accepted even though they are more generous.⁸ What is particularly puzzling about both of these phenomena is that the reference points according to which one splits the difference are endogenous; they are simply the prior two offers, one set by each bargainer. So, anticipating such behavior, it seems one would do well to engineer extreme reference points in your favor.

These descriptive results highlight both the strengths and weaknesses of bargaining theory, and offer paths for future empirical, experimental, and theoretical work. In particular, by highlighting features of real-world bargaining, the hope is to point to research avenues that can productively engage with bargaining practitioners. This is, of course, just a first step and there are myriad alternative settings in which we could learn even more about bargaining. Even within our setting, however, we have focused exclusively so far on the offers, counteroffers, and outcomes, and neglected communication between buyers and sellers. It is to this that we turn next.

⁷ This suggestion is thanks to Philippe Jehiel.

⁸ Substantially higher offers (e.g., 70%) are still more likely to be accepted as the non-monotonicity is a local phenomenon.

17.2 Cheap-Talk Signaling and Bargaining

Previously we documented a number of patterns in the data that are seemingly at odds with rational models of bargaining. Another such pattern is documented in Fig. 17.3. On the x -axis of that plot, we have the BIN price of the BO listing. On the y -axis, we have a scale from 0 to 1, and the object of interest is the ratio of the initial buyer offer in a bargaining interaction to the BIN price. To construct the scatterplot, every point in the graph represents a group listings that all have the same BIN price *within* a unit interval according to the ceiling function. Concretely, consider an integer $z > 0$, we then take all listings with a BIN price of $x \in (z - 1, z]$ and group them together, and these listings correspond to the x -axis of the graph. Then, we calculate the ratio of

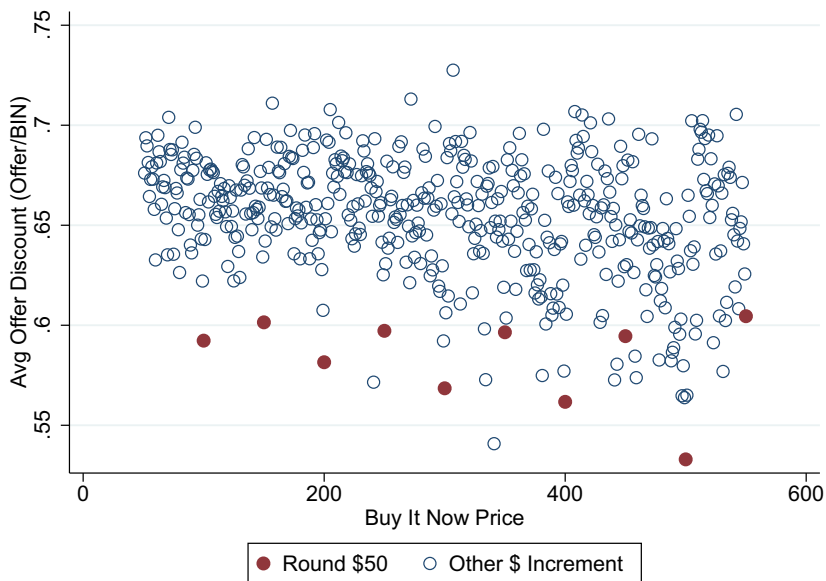


Fig. 17.3 Best offer on eBay (Notes This scatterplot presents average first offers, normalized by the BIN price to be between zero and one, grouped by unit intervals of the BIN price, defined by $(z - 1, z]$. When the BIN price is on an interval rounded to a number ending in "00" or "50," it is represented by a red circle. The figure is borrowed from Backus et al. [4])

the initial buyer offer for each of these listings and calculate the average of the ratios of all listings within such a group z , which corresponds to the y -axis of the graph. On average across all the points in the scatterplot, it appears that first offers arrive at somewhere near 65% of the BIN price. However, strikingly, if the BIN price is on a round “50” number, which we mark with a filled circle, we observe first offers between five and ten percentage points lower than other, precise number prices. A similar pattern persists for finally negotiated sale prices.

Why do round numbers fare so poorly? We were not the first to make this observation; indeed, a large literature in management and social psychology has documented it in lab experiments, see e.g., Mason et al. [18]. These papers attribute it to various perceptual biases around the use of round numbers, and conclude that bargainers should avoid the use of round numbers, favoring precise ones. However, this argument sits uneasily with economists—if the use of round numbers is a strictly dominated action then we should observe that rational bargainers avoid them. Instead, bargainers are discontinuously more likely to use round-number prices than precise number ones. Indeed, 5.3% of listings in the sample are priced at an exact multiple of \$100. Why are they doing it wrong? In Backus et al. [4], we posit an alternative hypothesis, one that rationalizes the use of round numbers by sellers: round numbers are a cheap-talk signal of eagerness to transact, analogous to rug stores “going out of business,” or forward compliments on a first date. In this light, the behavior may be entirely rational: impatient sellers (or, equivalently, sellers with worse outside options) will use round numbers to facilitate a transaction, albeit at a lower price, while patient sellers will use precise numbers and wait for higher-paying buyers.

To explore this idea in depth, we ask: what must be true in the equilibrium of a signaling game? If players are playing a separating equilibrium in which impatient sellers use round numbers and those who can hold out choose precise numbers, then (1) there must be *incentive compatibility*, a trade-off that makes some sellers prefer one outcome and others another; (2) sellers with different preferences must *sort*, i.e., if we split sellers by round- and precise-number usage, we should also see that these are fundamentally different kinds of sellers, and (3) buyers must *update*

their beliefs, that is, see the signal and form beliefs about the seller's preferences that are consistent with the sorting. The extensive data generated by online transactions make the validation of this equilibrium uniquely feasible.

Incentive compatibility becomes transparent once we construct the analogue of Figure 17.3 for the likelihood of—and time to complete—a sale. In fact, round-number sellers are 15–20% more likely to sell and, conditional on sale, do so almost 30 days earlier. In other words, we can think of the use of round-numbers as akin to moving along a demand curve, trading lower prices for higher quantities. According to basic pricing theory, we expect sellers to do this as their costs change, and the intuition in a bargaining setting is no different. Patient sellers—equivalently, those with a high marginal or opportunity cost of transacting—will prefer a higher price and a lower quantity, and impatient sellers the opposite.

Do we observe sorting of patient and impatient sellers into round- and precise-number listing prices? We can test this by looking at cases where similar offers (as a fraction of the BIN) were made to both types of sellers. If precise-number sellers are the patient sellers of the marketplace, then they should be less likely to accept an offer, other things equal. Similarly, impatient sellers should jump to accept it. Indeed, this is borne out by the data. Taking the same offer (as a fraction of the asking price) to set “other things equal,” precise number sellers are uniformly less likely to accept it, as depicted in Figure 17.4.⁹

The final diagnostic for a signaling equilibrium is that the signal is received—i.e., we need to show that buyers observe the signal and update their beliefs. This is particularly complicated because beliefs are unobservable. However, the behavioral data on buyers from the eBay platform affords an unusual opportunity. Buyers receive the “signal,” i.e., see whether there is a round or a precise asking price, when they view listings on the search results page. As it turns out, round and precise listings are equally likely to appear on the search page, yet the data show that buyers are *systematically more likely* to click through to the next step, the

⁹ Note that seller acceptance rates do not converge to one as the buyer's offer does—this may simply reflect the fact that many sellers are missing notifications that they have received an offer.




Fig. 17.4 Seller acceptance rates (Notes This figure depicts the polynomial fit of the probability of acceptance for a given offer (normalized by the BIN) on items with listing prices between \$85 and \$115, plotted separately for \$100 “Round” listings and the remaining “Precise” listings. This figure is borrowed from Backus et al. [4])

detailed item page, for round-number listings than precise number ones. And conditional on clicking through and purchasing, they are *more likely to negotiate* (rather than take the BIN price) for round-number listings. This is consistent with the hypothesis that they (correctly) anticipate negotiating lower prices with round-number sellers.

At this point, the attentive reader may rightly object—are round-number listings similar to precise-number listings on other dimensions besides price? Maybe this constellation of results is all due to unobserved heterogeneity, which was highlighted as a limitation of the dataset in Sect. 17.1. To answer this concern, Backus et al. [4] identify a natural experiment—a garbling of the signal that exposes it to some buyers and

obfuscates it to others.¹⁰ In particular, sellers from the U.K. website of eBay (eBay.co.uk) have an option to pay for their listings to be cross-listed on the U.S. site. If they do so, these listings that are originally listed in British Pounds will also appear in U.S. search results, but prices will be converted automatically to U.S. Dollars using the current exchange rate. As a result, a listing that has a round BIN price in the U.K. will appear as a precise BIN price in the U.S., holding fixed all of the other characteristics of the item (including, importantly, those not observed to the econometrician). In this setting, we showed that the price effects of round-number signaling persist after controlling for unobserved heterogeneity.

Based on all of this evidence, we concluded that cheap-talk signaling offers the simplest, most coherent explanation for the behavior of buyers and sellers concerning round numbers. And yet, it seems implausible that buyers and sellers are all knowingly engaging in this behavior. The early papers documenting this phenomenon drew on the work of psychologists who document known perceptual biases around round numbers. How do we reconcile this with the heady assumptions on rationality required to rationalize a perfect Bayesian equilibrium? The work of Thomas et al. [25] offers a way out. They document perceptual biases around round numbers, but critically, they show that these biases are mediated by experience. In other words, while the mechanical truth of decision-making may be that we have heuristics that guide perception and influence behavior, these heuristics are consistent with observed experience and are in that sense functionally equivalent to rational expectations on the equilibrium path.

¹⁰ The idea of using a garbling device to identify signaling appears also in Ambrus et al. [1]. There, the church, which serves as an intermediary in hostage negotiations between wealthy families and Tunisian pirates, must send couriers by foot to the town in which the family is located. Because the time required is unobservable to the pirates, they take this as a garbling device in the use of delay as a signaling device in negotiations and show that it predicts better bargaining outcomes, consistent with theory.

17.3 Protocol Design: Communication and Bargaining Breakdown

An important question for the bargaining literature, but also for the e-commerce platform regulating the bargaining itself, is: should buyers and sellers be allowed to communicate? From the perspective of the literature, this is an old question that relates to bargaining efficiency, and it does not help that theoretical models of communication in bargaining are divided on the efficiency implications of communication. While there may exist equilibria with credible communication in bargaining games, as in Farrell and Gibbons [14], it does not in general improve the efficiency of bargaining.

The experimental literature is rather more optimistic. In early work, Radner and Schotter [22] and Valley et al. [26] offer evidence that communication improves bargaining efficiency. Indeed, in the latter, it appears that subjects outperform the Myerson and Satterthwaite [20] upper bound on efficiency for incentive-compatible bargaining mechanisms in games with two-sided asymmetric information. The mechanism by which it does so, however, is unclear. It may be, e.g., that we over-communicate in cheap-talk games [6], we are averse to lying, or that we build amity and altruism through communication.

From the platform's perspective, however, there is an important caveat to the gains of bargaining efficiency. Platforms are particularly concerned about the risk of *disintermediation*; that is, when the buyer and the seller decide to cut out the middleman (and their fees) and instead transact independently, using the communication mechanism to foster off-platform contact. Indeed, the salience of this concern is highly predictive of the different platforms' choices regarding communication. Taobao, which does not tax transactions, allows free-form communication between bargainers with an on-site instant messenger service.¹¹ eBay, which charges fees between eight and twelve percent of the transaction price, allows brief text communication accompanying offers. And finally Amazon Marketplace, which charges fees between 15 and 25%,

¹¹ Taobao's instant messenger service was one factor in their success over the Chinese version of eBay, see Oberholzer-Gee and Wulf [21] for a discussion.

does not allow text communication on its bargaining mechanism (which is available only in a few narrow categories to begin with).

From a theoretical as well as a practical market design perspective then, it is important to understand how communication affects bargaining outcomes. To this end, Backus et al. [3] identify a natural experiment in the availability of communication in bargaining. Recall from Fig. 17.1 that there is an option to “add message to seller” when a buyer makes an offer. This requires an extra click, but it is available to both parties. Importantly, messages on the bargaining platform can only accompany a price offer.

Historically, this feature was unavailable on eBay.de, the German incarnation of eBay, while all other features of the Best Offer platform were equivalent to the U.S. (and other) sites. On May 23, 2016, the messaging feature was enabled on the German site, but it was only immediately available for buyers accessing the eBay.de marketplace from a computer browser, rather than a mobile app. Mobile app users, who made up approximately half of bargainers on the site, were not able to send messages.

The rollout of messaging in Germany offers a simple difference-in-differences design: before and after versus desktop and mobile. Backus et al. [3] use this natural experiment to identify the effect of the availability of communication on the likelihood that any particular bargaining interaction—inclusive of all offers between a buyer and a seller—ends in a transaction. They estimate the intent-to-treat effect to be approximately half a percentage point. This seems small, but it is important to remember that the magnitude is determined in part by the number of bargainers who actually take the feature up, i.e., “compliers,” which in this case was only 6%. Adjusting for the compliance rate, the estimated effect of actually sending a message is 7 percentage points, against a baseline success probability of 44%. This implies that among compliers, the treatment effect of messaging was to increase the odds of a successful interaction by a staggering 15%.

The estimated effects in Backus et al. [3] were not realized immediately, however. Instead, they observe that although the take-up rate is almost instantaneous, the treatment effect takes a few weeks after introduction to stabilize. The paper argues that this is evidence of learning

by bargainers who participate in multiple bargaining sequences. In an involved text analysis exercise, they document (1) that the text content of messages evolves over time, (2) that it becomes more similar, on a week-to-week basis, over time, and (3), that the changes are isolated among users who are repeat bargainers on the platform. In contrast, the text content of one-off buyers and sellers is stable over the ten weeks following the introduction of messaging.

The golden question for both academics and practitioners is, of course, what should we say when we bargain? The natural experiment in that paper does not offer an answer because it generates pseudo-experimental variation in the *availability* of communication, but not in the *content* of what is actually said. But as a second best, we might substitute an alternative question: what are more experienced sellers learning to say when they bargain? If sellers are learning from experience to bargain more effectively, then this offers some potential guidance for bargainers. At the very least, it reflects on existing work and has the potential to generate new hypotheses for what is to come.

Adopting the distributed multinomial regression framework of Taddy [24], they identify word pairs that are predictive of experience using the sample of observations in the ten weeks following the change. The exercise offers several reflections on messaging strategies. They find that inexperienced sellers are more likely to use effusive greetings, whereas more experienced sellers use polite, but restrained greetings. They also find that inexperienced sellers emphasize free shipping, which is salient in the listing, whereas experienced sellers are more likely to remind buyers of less-salient cost factors, e.g., the fact that PayPal and eBay charge them fees.

These findings are purely descriptive and may be context-specific, however, they offer a much-needed datapoint on communication in bargaining. Using natural language processing tools, exercises like this can describe what bargainers are actually doing in the field, and use that to motivate theoretical and experimental inquiry. Especially in light of the large positive effects of communication we found on eBay.de—a fourteen percent decrease in the rate of bargaining breakdown for interactions that involved a message after the change—we hope that this

approach will foster a research agenda that helps us better understand the mechanisms by which communication affects bargaining.

17.4 New Tools and Directions

The Best Offer research agenda has offered an empirical setting in which to assess the performance of existing theoretical models and shows that on many elements, economic theory holds up surprisingly well. At the same time, however, it also raises new puzzles and opportunities for future research.

We observe evidence consistent with cheap-talk signaling in the strategic choices of bargainers, evidence that bargaining ability and patience matter for outcomes, and that bargaining itself is costly in a way that affects outcomes. Then again, we also observe puzzles. In the patterns of offers, we see reciprocal gradualism and splitting the difference, neither of which yet has a satisfying theoretical motivation. And in the messages that buyers and sellers send, we see what any practitioner surely already knows: that what we say when we bargain, and the opportunities we have to say it, matters greatly for determining outcomes.

In making the data public we hope to encourage new research on this question. Indeed, answers to these puzzles and others may lie in the Best Offer data. We believe, however, that growing this research agenda will also depend on finding new large-scale bargaining datasets. For example, Bagwell et al. [5] have constructed a novel large-scale dataset on the trade negotiations behind GATT. Moreover, we conjecture that with the emergence of online real-estate agent platforms (e.g., Redfin), there may one day be a large-scale dataset of offer-level bargaining for real-estate transactions.

While we pin much of our hopes on new data sources, we should also highlight the role of new tools. Especially insofar as we endeavor to think of bargaining as a communicative act, natural language processing tools for parsing text documents may prove critical to empirical attempts to understand bargaining. Best Offer bargaining has the advantage of being structured—every message is accompanied by a numerical offer in an

alternating, sequential-offers setting. A similar advantage is shared by the GATT negotiations. While convenient, this is not generic to bargaining in the wild, and so we believe that new ML tools will have a central role in modeling unstructured bargaining.

References

1. A. Ambrus, E. Chaney, and I Salitsky. Pirates of the mediterranean: An empirical investigation of bargaining with transaction costs. *Quantitative Economics*, 9(1):217–46, 2018.
2. Matthew Backus, Thomas Blake, Bradley Larsen, and Steven Tadelis. Sequential bargaining in the field: Evidence from millions of online bargaining interactions. NBER Working Paper No. 24306, 2018.
3. Matthew Backus, Thomas Blake, Jett Pettus, and Steven Tadelis. Communication and bargaining breakdown: An empirical analysis. NBER Working Paper No. 27984, 2020.
4. Matthew Backus, Thomas Blake, and Steven Tadelis. Cheap talk, round numbers, and the economics of negotiation. NBER Working Paper No. 21285, 2015.
5. K. Bagwell, R. W. Staiger, and A. Yurukoglu. Multilateral trade bargaining: A first peek at the GATT bargaining records, 2017. Working paper, Stanford University.
6. Hongbin Cai and Joseph Tao-Yi Wang. Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56:7–36, 2006.
7. David Card. Strikes and bargaining: A survey of the recent empirical literature. *American Economic Review*, 80(2):410–415, 1990a.
8. David Card. Strikes and wages: A test of an asymmetric information model. *Quarterly Journal of Economics*, 105(3):625–659, 1990b.
9. David Card and Craig A. Olson. Bargaining power, strike durations, and wage outcomes: An analysis of strikes in the 1880s. *Journal of Labor Economics*, 13(1):32–61, 1995.
10. Olivier Compte and Philippe Jehiel. On the role of outside options in bargaining with obstinate parties. *Econometrica*, 70(4):1477–1517, 2002. ISSN 00129682.

11. P. Cramton and J.S. Tracy. Strikes and holdouts in wage bargaining: Theory and data. *American Economic Review*, 82(1):100–121, 1992.
12. Vincent P. Crawford. A theory of disagreement in bargaining. *Econometrica*, 50(3):607–637, 1982.
13. Liran Einav, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–247, 2015.
14. Joseph Farrell and Robert Gibbons. Cheap talk can matter in bargaining. *Journal of Economic Theory*, 48(1):221–237, 1989.
15. Daniel Goetz. Dynamic bargaining and scale effects in the broadband industry. Working Paper, 2019.
16. Pranav Jindal and Peter Newberry. To bargain or not to bargain - the role of fixed costs in negotiations. *Journal of Marketing Research*, 55(6):832–851, 2018. Working Paper, Pennsylvania State University.
17. B. Larsen. The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. NBER Working paper 20431, 2018.
18. Malia Mason, Alice Lee, Elizabeth Wiley, and Daniel Ames. Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *Journal of Experimental Social Psychology*, 49:759–763, 2013.
19. Antonio Merlo and Francois Ortalo-Magné. Bargaining over residential real estate: Evidence from England. *Journal of Urban Economics*, 56(2):192–216, 2004.
20. R.B. Myerson and M.A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983. ISSN 0022-0531.
21. Felix Oberholzer-Gee and Julie Wulf. Alibaba's taobao. Technical Report 9-709-456, Harvard Business School, 2009.
22. Roy Radner and Andrew Schotter. The sealed-bid mechanism: An experimental study. *Journal of Economic Theory*, 48(1):179–220, 1989.
23. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, pages 97–109, 1982. ISSN 0012-9682.
24. Matthew Taddy. Distributed multinomial regression. *The Annals of Applied Statistics*, 9(3):1394–1414, 2015.
25. Manoj Thomas, Daniel H. Simon, and Vrinda Kadiyali. The price precision effect: Evidence from laboratory and market data. *Marketing Science*, 29(1):175–190, 2010.

26. Kathleen Valley, Leigh Thompson, Robert Gibbons, and Max Bazerman. How communication improves efficiency in bargaining games. *Games and Economic Behavior*, 38(1):127–155, 2002.