Deliberation during online bargaining reveals strategic information

Miruna Cotet,[1] Wenjia Zhao,[2] Ian Krajbich[1,3]

[1]The Ohio State University, Columbus, US

[2]University of Warwick, Coventry, UK

[3]University of California Los Angeles, US

Abstract

A standard assumption in game theory is that decision makers have pre-planned strategies telling them what actions to take for every contingency. In contrast, non-strategic decisions often involve an on-the-spot comparison process, with longer response times (RT) for choices between more similarly appealing options. If strategic decisions also exhibit these patterns, then RT might betray private information and alter game theory predictions. Here, we examined bargaining behavior to determine whether RT reveals private information in strategic settings. Using pre-existing and experimental data from eBay, we show that both buyers and sellers take hours longer to accept bad offers and to reject good offers. We find nearly identical patterns in the two datasets, indicating a causal effect of offer size on RT. However, this relationship is half as strong for rejections as for acceptances, reducing the amount of useful private information revealed by the sellers. Counter to our predictions, buyers are discouraged by slow rejections – they are less likely to counteroffer with slow sellers. We also show that a drift diffusion model (DDM), traditionally limited to decisions on the order of seconds, can account for decisions on the order of hours, sometimes days. The DDM reveals that more experienced sellers are less cautious and more inclined to accept offers. In summary, strategic decisions are inconsistent with pre-planned strategies. This underscores the need for game theory to incorporate RT as a strategic variable and broadens the applicability of the DDM to slow decisions.

When people interact in strategic settings, do they have a plan, or do they make decisions on the spot? Do they respond immediately, or does it take time to figure out what to do? For example, does someone selling their car know what offers they will accept or reject? And does the buyer know what counteroffers they will accept or reject?

Whether people have prepared plans is an important question, because without them, people risk revealing private information (*1-9*). Private information is a central concern in strategic settings – it is information that people have that gives them an advantage over others and that they would like to keep hidden (*10*). For example, when bargaining over a car, a seller might be willing to accept a small amount but wouldn't want to reveal that to a potential buyer, who might be willing to pay a lot more. With a prepared plan, the seller would be able to instantly accept or reject any offer from the buyer. Without a prepared plan, the seller might require time to consider the buyer's offer, inadvertently betraying how attractive they find it. A quick rejection could signal a non-competitive offer while a slow rejection could signal a close call (*11, 12*). This information could in turn be used by the buyer to make smarter follow-up offers (Fig. 1). The central question in this paper is whether bargainers' response times (RT) convey such information.

To tackle this question, we need to understand what choice process bargainers might be using, if not executing a prepared plan. Many quick decisions involve a process of accumulating and comparing evidence up to a predetermined boundary, a process which takes time and reflects strength-of-preference (*13-25*). The evidence reflects the person's evaluation of the options – a person deciding between an apple and an orange must weigh the costs and benefits of the apple against those of the orange. If these two evaluations are roughly equal, the person will struggle to decide which item to choose. On the other hand, if the person finds the orange to be much more attractive than the apple, then their choice will be quick and predictable. This relation between strength-of-preference and RT is a basic feature of evidence-accumulation or sequential-sampling models, such as the drift-diffusion model (DDM). Sequential sampling models like the DDM are typically applied to fast perceptual judgments, but in recent years they have seen increasing application to economic choice (*26-45, 23, 13*).

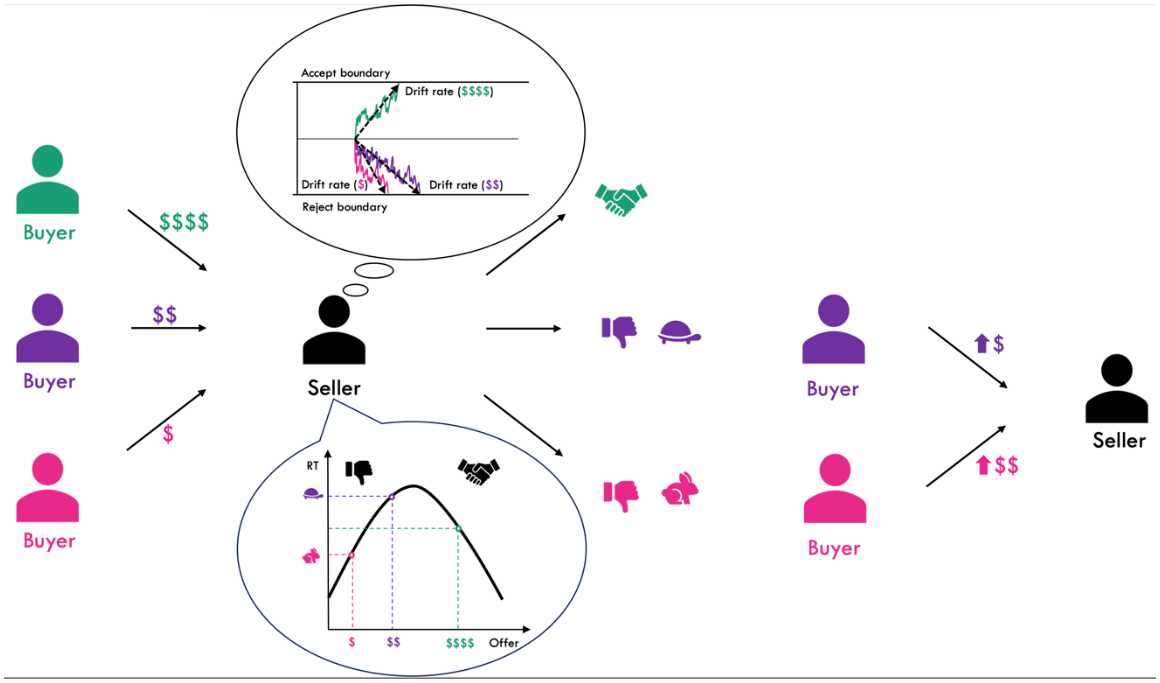

Figure 1. Bargaining setting. In a standard bargaining exchange, the proposer makes an offer to the responder, which the responder can either accept or reject. On eBay, buyers make offers to sellers. The size of the offer, in combination with the seller's private value, determines the drift rate in the DDM. Here, the high offer (in green) yields a modest positive drift rate (towards the "accept" boundary), the medium offer (in purple) yields a low negative drift rate (towards the "reject" boundary), and a low offer (in pink) yields a high negative drift rate. As a result, the seller accepts the high offer with medium speed, rejects the medium offer slowly, and rejects the low offer quickly. So, a strategic buyer will increase their next offer a lot after a quick rejection, but only a little after a slow rejection.

Despite the evidence supporting the DDM in economic choice, it is still unclear whether its predictions extend to strategic settings. One recent article by Konovalov & Krajbich (2023) (KK) examined the potential ways in which bargaining could play out with a DDM decision process. Theoretically, they found multiple possibilities: bargainers might always make instantaneous decisions (to mask their private information) or they might still take time to make their decisions. In any case, KK assumed that bargainers are forced to learn their values at the time they receive an offer. Their laboratory experiments reflect this assumption: college undergraduates bargained repeatedly over vouchers with induced values and their decisions were limited to less than 10 seconds. These results cannot tell us whether the DDM extends to experienced sellers, bargaining over familiar goods, with RT on the order of minutes, hours, or days. In fact, the DDM has not been applied to natural decisions longer than a few seconds (*46*). Thus, it is still unclear whether RT is informative during bargaining and, if so, whether DDM can account for the patterns in the data.

Here, we examine DDM predictions in bargaining (*47-50*), using field data and field experiments, with RT ranging from a few seconds to multiple days. Our field setting is eBay – one of the world's largest online marketplaces. Since 2005 eBay has allowed people to sell their products through an alternating-offer protocol where sellers post items for sale and buyers can make them offers. eBay recently released a dataset with millions of bargaining exchanges from May 2012 to June 2013 (*51*). We analyzed these data to test whether sellers' RT reflected the quality of the offers that they received. Because these data can only provide correlational results, we additionally conducted and analyzed a pre-registered field experiment on eBay where we acted as buyers, experimentally varying offer size to study sellers' resulting RT without the potential confounds of non-experimental data. To preview the results, we find that sellers' (and buyers' RT) do indeed reflect the quality of the offers they receive. The size of these effects is remarkably similar in both the pre-existing data and the field-experiment data. Moreover, a DDM, with some adjustments to account for long non-decision times, does a reasonable job of explaining sellers' choices and RT. We conclude that most sellers do not have prepared plans and instead evaluate offers as they come in. As a result, their RT convey useful information that can be used against them in strategic settings.

Results

On eBay, a seller can post an item for sale along with an accompanying list price. Buyers can then purchase the item instantly at that price. A seller can additionally enable a "make an offer" option, which allows buyers to initiate a bargaining process. The process begins with the buyer making the seller an offer. The seller can accept, reject, or counter the offer. If countered, the buyer can then accept, reject, or counter that offer. This process continues until a sale is made or until 3 offers are made by each party. There is a deadline of 48 hours to respond to each offer. After that time, the offer is automatically rejected. In the data, a single exchange is a collection of offers between a buyer and a seller for a specific item.

We examine two eBay bargaining datasets. The first dataset comes from Backus et al. The second dataset comes from a field experiment that we conducted.

The Backus et al. dataset consists of a year's worth of bargaining exchanges (between May 2012 and June 2013) on eBay (*51*). For practical reasons, we only analyze a randomly selected ~20% of that data, resulting in 1.02 million bargaining exchanges.

We excluded all exchanges where an offer arrived from a different buyer. Our model of the seller's decision is that they compare the attractiveness of the buyer's offer to the potential benefits of waiting for future offers (or perhaps keeping the item). If another offer actually arrives during this decision process, that changes the calculation, and should lead to the immediate rejection of the worse offer. So, we exclude all such cases (8.6% of all exchanges, ignoring the other restrictions).

We also exclude exchanges with items listed at more than $1000 (6.7% of all exchanges, ignoring the other restrictions), where the RT was less than 10 seconds (including automatic acceptances and rejections; see below), where the offer expired (22.1% of all exchanges, ignoring the other restrictions), where a message was included (9.8% of all exchanges, ignoring the other restrictions), and where a buyer or seller was outside of the US (15.8% of all exchanges, ignoring the other restrictions) (see Supplementary Methods Section for full details).

The second dataset was from a pre-registered field experiment that we ran from 2020 to 2023. The goal of the experiment was to address the issue of causality: does the size of the offer cause the RT, or might there be hidden variables that influence both offers and RT (e.g., unobserved seller characteristics)? In the experiment we acted as buyers on eBay, identifying a set of sellers, and making pseudo-random offers to them. We ran the experiment in two waves, the first with 50 sellers and 11 offers per seller, and the second with 150 sellers and 21 offers per seller. In both cases we preselected sellers based on their number of items for sale, identified items to bid on, and then made offers between 0.3 and 0.9 of list price for the first wave and 0.1 and 0.8 for the second wave. The items that we bid on were collectible trading cards (e.g., baseball cards, Pokemon cards) valued between ~$10-20. In total we made 3,586 offers (see Methods for additional details).

Sellers do not have prepared plans

To establish whether eBay sellers have prepared plans we look at two things. First, we examine whether they use tools that allow them to automatically accept or reject certain offers. If sellers are aware of these tools and have stable plans, using such thresholds would save them the effort of having to respond to offers. Still, sellers might not be aware of these tools or might want to adjust their plans over time. Therefore, our main analysis focuses on whether sellers' RT reflect the size of the offers that they receive. A seller who is making decisions on the spot should react to higher offers with faster acceptances and slower rejections. Conversely, a seller with a plan should show no relation between offer size and RT.

The pre-existing data from eBay indicate that most sellers (62%) do not use automatic thresholds. In the subset of data that we analyzed, 17% used only rejection thresholds, 5.4% used only acceptance thresholds, and 15.7% used both thresholds. To examine the effect of seller experience on automatic threshold use, we looked at the number of best-offer listings created by each seller (dating back to 2008), The relation between seller experience and threshold usage was quadratic – both inexperienced and highly experienced sellers used thresholds more than medium-experience sellers (Fig. S1).

While these results do not rule out that sellers have plans, the low rate of automatic thresholds suggests that sellers prefer to evaluate offers as they arrive.

Turning to our main analysis, we found that sellers' RT were highly responsive to the offers they received. We focus on buyers' initial offers and sellers' responses to those offers (excluding sellers with automatic thresholds – Methods). Over most of the offer range [20%, 100%], the median acceptance RT decreased with offer size from 2 hours down to 0.8 hours, while over most of the offer range [0, 70%], the median rejection RT increased with offer size, from 1.3 hours up to 1.8 hours (Fig. 2A). We confirmed these results using linear regressions of log(RT) on offer size (as percent of list price) over the offer range for which both mean acceptance and rejection RT were judged to be monotonic in price (range = [0.36, 0.68]; Methods). Within this range, acceptances had a negative relation between offer size and log(RT) while rejections had a positive relation between offer size and log(RT) (mixed effects with full random effects at the seller level: $b_{accept} = -0.24$, S.E. = 0.01, 95% CI = [-0.27, -0.22], $t(255{,}930) = -21.77$, $p < 10^{-16}$ vs. $b_{reject} = 0.10$, S.E. = 0.01, 95% CI = [0.07, 0.13], $t(116{,}781) = 6.86$, $p = 10^{-11}$). These results are robust to including various controls, including the offer creation hour, the price of the item, the number of views, the number of watchers, whether the item was relisted, the listing age, the number of photos of the item, buyer experience, seller experience, the number of seller listings, seller feedback, and interactions between the offer and seller experience (Table S2). Seller-level regressions revealed a similar pattern, with 69% showing negative offer-RT correlations for acceptances and 53% showing positive offer-RT correlations for rejections (of 1,127 sellers with at least 50 acceptances or rejections; Fig. S13). These effects were also largely consistent across product categories, ranging from baseball cards to vehicles, with 28/32 categories showing a positive relation between log(RT) and offer size for acceptances and 26/32 categories showing the opposite relation for rejections (Fig. 2B). Thus, in line with the DDM, sellers on eBay were slower to reject better offers and to accept worse offers.

Consistent with the DDM, we also found that large "errors" were made very quickly. By errors, we mean low-probability decisions, i.e., rejecting very high offers (> 75%) or accepting very low offers (< 25%) (Fig. 2A; Fig. S8 for breakdown by absolute offer level). A counter-intuitive prediction of the DDM is that large errors are substantially faster than decisions near indifference (*17*). Thus, the fast errors that we observed are consistent with a DDM explanation.

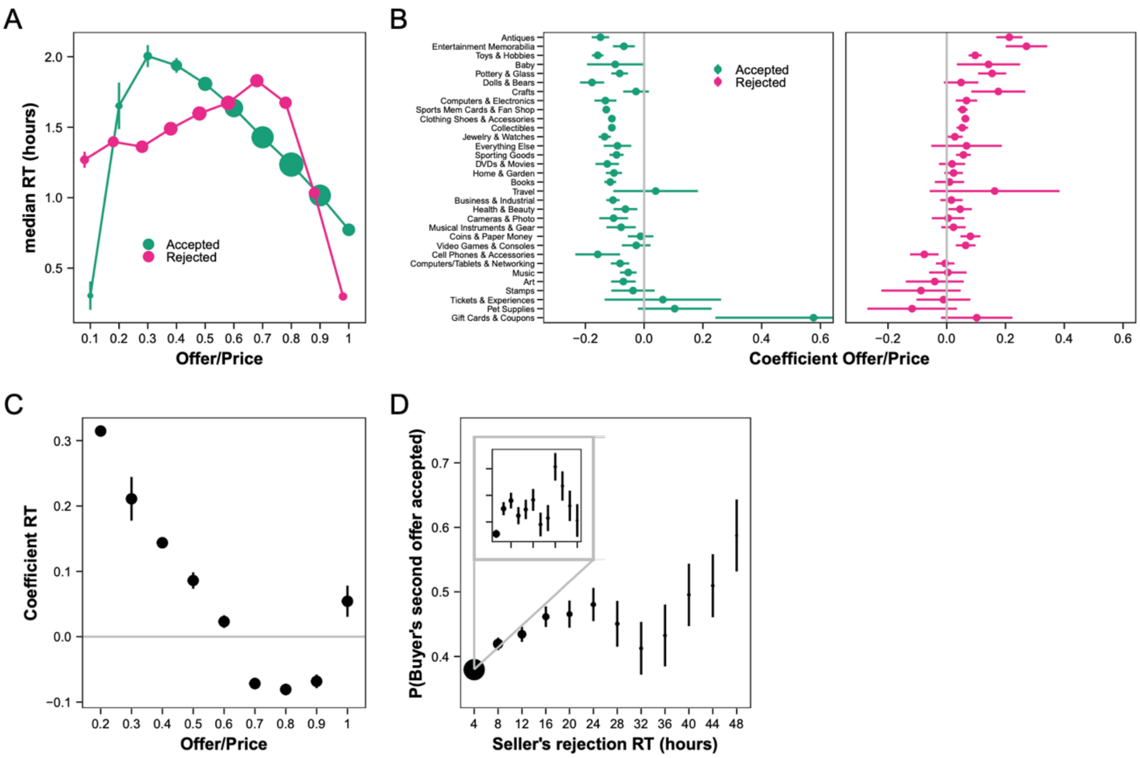

Figure 2. eBay RT reflect offer size. (A) Sellers' median RT (in hours) as a function of buyers' initial offer ratios (offer / list price), conditional on accepting or rejecting the offers. **(B)** Coefficients for the relation between seller's log(RT) and buyers' initial offer ratios, conditional on accepting or rejecting the offers and on item category. **(C)** The effect of sellers' RT on their probability of accepting the offer, conditional on the offer ratio. Plotted along the y-axis are the logistic regression coefficients for offer ratios in 10% intervals. The regression is the same as in Table S5 (column 3), but with RT (in hours) added as an explanatory variable. These regressions include random effects (clustered by seller) on the intercept. **(D)** The probability of sellers accepting buyers' first offers as a function of the sellers' rejection RT to the first offers. In the zoomed-in inset, the probability that the seller accepts the buyer's second offer ranges from [0.332, 0.418]. For **(A)** the size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors. For **(B, C)** the bars represent standard errors. For **(D)** the size of the dots indicates the relative amount of data in that bin, and the bars represent bootstrapped standard errors.

Sellers' counteroffers also displayed a significant relation between log(RT) and offer size. We had speculated that counteroffers might look more like rejections but instead they looked more like acceptances. These RT monotonically decreased with offer size over the range [10%, 100%] from 1.46 hours to 0.54 hours (Fig. S5, Table S4). Over the same offer range as before ([0.36, 0.68]), we found a significant negative relation between log(RT) and offer size (mixed-effects regression with full random effects at the seller level: $b_{counter} = -0.06$, S.E. = 0.01, 95% CI = [-0.07, -0.05], $t(187{,}735) = -9.17$, $p < 10^{-16}$). This result is robust to including the same controls as before. Interestingly, the counteroffers did not display the same fast errors as the acceptances and rejections. Indeed, fast errors are implausible with counteroffers, as sellers have to produce a new offer and not just react to the buyer's offer.

One explanation for why a seller might delay in responding to an offer is that they are waiting to see if any better offers come in. Because we only consider cases where no other offers were made on the item during the bargaining thread, additional time can only be bad news for the seller. So, an increase in RT should correspond to a higher probability of accepting the offer. Instead, for offers between 60% and 90%, longer RT corresponded to a lower probability of acceptance (Fig. 2C). Thus, the idea that sellers are waiting for better offers to arrive is not sufficient to explain their behavior.

An additional prediction that one can derive from the DDM is that sellers who reject an offer more slowly should be more likely to accept a subsequent offer from the buyer, controlling for the size of the offer, because a slower rejection signals a lower value from the seller's perspective. Indeed, sellers were more likely to accept buyers' second offers the longer they took to reject the buyers' first offers (Fig. 2D; Table S11; $b_{rejectionRT} = 0.13$, *S.E. = 0.02, 95%CI = [0.10, 0.17],* $z(23{,}432) = 7.74$, $p < 10^{-15}$) controlling for the second offer ratio. This holds also when including additional controls for the item characteristics (number of watchers, views, photos, list price and listing age) and buyer bargaining history (Table S11; $b_{rejectionRT} = 0.12$, *S.E.* $= 0.02$, *95%CI* $= [0.08, 0.15]$, $z(23,426) = 6.43$, $p < 10^{-9}$).

eBay Field Experiment

Our eBay field experiment largely confirms the results from the pre-existing dataset.

We made 3,586 offers and had 1,550 acceptances, 644 rejections, and 1392 counteroffers. Sellers were more likely to accept higher offers ($b_{Offer/Price} = 2.64$, S.E. = 0.11, 95% CI = [2.43, 2.86], $z(2,192) = 24.34$, $p < 10^{-16}$, Fig. 3A, Table S15).

In terms of RT, sellers were significantly faster to accept higher offers, and were non-significantly slower to reject higher offers (Fig. 3B; the rejection effect was significant, $p < 0.0001$, when including expired offers and assigning them the maximum RT, see Table S22). Regressions of log(RT) on offer ratio revealed nearly identical coefficients to the pre-existing eBay data ($b_{accept}= -0.24$, SE = 0.07, 95% CI = [-0.38, -0.10], $t(1,544) = -3.37$, $p = .001$, $b_{reject} = 0.11$, SE = 0.10, 95% CI = [-0.09, 0.32], $t(638) = 1.06$, $p = .29$, Table S19). Thus, the RT results from our experiment align well with those from the pre-existing data.

The RT for counteroffers were one exception. Unlike in the pre-existing eBay data, here the counteroffers displayed a positive (though not significant) rather than negative relation between log(RT) and offer ratio ($b_{counter} = 0.05$, S.E. = 0.08, 95% CI = [-0.11, 0.20], t(938) = 0.57, p = .569), though the sign was inconsistent between waves of the experiment (Table S20).

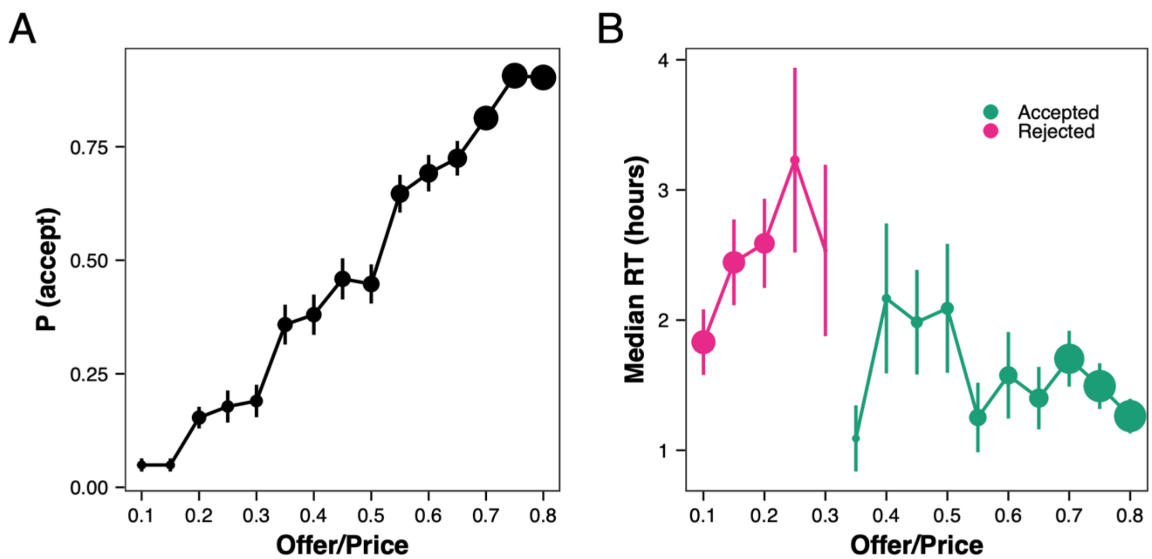

Figure 3. eBay field experiment mirrors pre-existing eBay data. (A) Sellers' probability of accepting the first offer as a function of the offer ratio (offer / list price). **(B)** Sellers' median RT (in hours) as a function of the offer ratio, conditional on acceptance or rejection. The size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent standard errors across sellers. Bins with less than 40 observations are excluded.

Buyers do not have prepared plans either

Our primary focus has been on sellers because they have at least one response in every exchange, we have more information on them, and they were the target of our field experiment. However, we can also see whether buyers' behavior is in line with DDM predictions. When sellers respond with counteroffers, the tables are turned and buyers have to decide whether to accept, reject, or counter. Like with the sellers, the DDM predicts that buyers (without plans) will respond to better counteroffers with faster acceptances and slower rejections.

To quantify the quality of a seller's counteroffer, we calculated the seller's "compromise". We defined the seller's compromise as the gap between their counteroffer and the list price as a fraction of the gap between the buyer's offer and the list price. A full compromise (=1) corresponds to the seller counteroffering with the same price that the buyer offered, while no compromise (=0) corresponds to the seller counteroffering with their list price. The bigger the compromise, the more attractive the counteroffer.

Buyers' RT did vary with the quality of sellers' counteroffers as expected (Fig. 4A Table S8). We regressed buyers' log(RT) on the sellers' compromise. Larger compromises led to slower rejections and faster acceptances ($b_{accept} = -0.13$, S.E. = 0.009, 95% CI = [-0.15, -0.11], $t(817{,}789) = -14.39$, $p < 10^{-16}$, $b_{reject} = 0.18$, S.E. = 0.008, 95% CI = [0.16, 0.19], $t(76{,}587) = 21.11$, $p < 10^{-16}$, Fig. 4A; Table S8). The difference in RT was substantial, increasing from 2.96 hours to 5.58 hours for compromises from 10% to 100%. Interestingly, buyers' RT were consistently longer for rejections than for acceptances. This may be because buyers are on eBay with the intent to purchase, so accepting is, on average, more attractive than rejecting.

Following Konovalov & Krajbich, we next attempted to analyze the size of buyers' second offers conditional on sellers' rejection RT. However, this analysis is highly contaminated by selection bias, as buyers made second offers only ~10% of the time, and those who did were surely more determined to acquire the items. In the supplementary material we report these analyses, as well as an attempt to deal with the selection bias, but given how much data is missing, we refrain from drawing any conclusions from those results.

Instead, we examined whether buyers were more likely to return with follow-up offers after being rejected more slowly. Since a fast rejection from a seller should signal to the buyer that their offer was not competitive, buyers who are rejected quickly should be discouraged from coming back with a second offer. Instead, we observed the opposite effects. Using a logistic regression, we regressed whether the buyer made a second offer on the RT of the seller's rejection. Buyers were more, not less, likely to make a second offer after a faster rejection (Fig. 4B, Table S9, $b_{rejection\ RT} = -0.098$, S.E. = 0.004, 95% CI

$= [-0.11, -0.09]$, $z(226{,}649) = -26.24$, $p < 10^{-16}$). This result was robust across all the item categories (Fig. 4C). Thus, counter to our predictions, buyers on eBay appear to be discouraged by slow rejections.

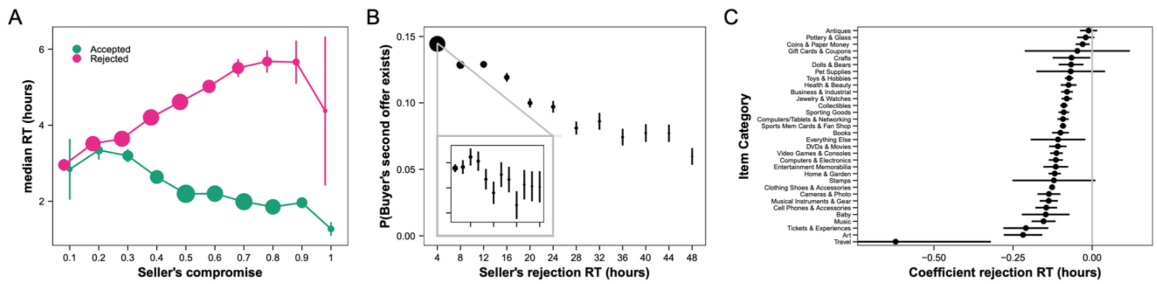

Figure 4. Buyers react adversely to slow sellers. (A) Buyers' median RT (in hours) as a function of sellers' compromise, conditional on the buyer accepting or rejecting the offers. The seller's compromise is the amount that they lowered their counteroffer, divided by the gap between the list price and the buyer's offer. A 100% compromise would be a counteroffer that matches the buyer's offer; a 0% compromise would be a counteroffer that is the list price. **(B)** Probability that buyers make second offers as a function of sellers' rejection RT to the first offers. The inset zooms in on rejection between 0 and 2 hours. In the zoomed-in inset, the y-axis ranges from [0.133, 0.152]. **(C)** Coefficients for the relation between seller's rejection RT to the first offers and probability that buyers make second offers, by item category. For **(A)** the size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors. For **(B)** the size of the dots indicated the relative amount of data in that bin, and the bars represent bootstrapped standard errors. For **(C)** the bars represent standard errors.

To understand why buyers are discouraged by slow rejections we investigated several hypotheses. One possibility is that buyers are impatient and so a slow response from a seller prompts them to move on to other sellers. A second possibility is that impatient buyers choose to simply purchase the item after a slow response in order to avoid further costly delays. A third possibility is that the effect is driven by buyers making offers on impulse goods, i.e., goods that they lose interest in over time. While our data are not ideal for testing these hypotheses, we investigate each one in turn.

The presence of multiple sellers might explain why buyers are encouraged by fast rejections on eBay, opposite to the lab (*12*). If buyers are impatient and have the option to move on to other sellers, the positive signal from a slow rejection might not be enough to compensate for the buyer's lost time (*57*). This hypothesis is difficult to test with our data because we do not know the exact identities of the items. However, we do find that buyers who were rejected more slowly by a seller were less (not more) likely to make another offer to a different seller for an item in the same category within 24 hours ($b_{rejectionRT}$ = -0.22, S.E. = 0.02, 95%CI = [-0.25, -0.19], z(226,644) = -13.15, $p < 10^{-15}$), controlling for the first offer ratio, item characteristics (number of watchers, views, photos, list price, listing age), and buyer bargaining experience. Buyers may sometimes move on to other sellers, but this is not more common after slow rejections.

A second possibility is that after a slow response, buyers simply purchase the items to avoid additional costly delays from bargaining. We find a complex relationship between seller's rejection RT and the probability of the buyer then purchasing the item. The probability of purchasing decreases with the seller's rejection RT, up to ~8 hours. The probability then rises slightly up to ~24 hours, at which point it declines again until bouncing back up at the 48-hour time limit. Thus, for most of the bargaining threads (77% of rejections are under 8 hours), these results do not support the idea that buyers are getting impatient and purchasing the items. Instead, the results support our initial hypothesis that fast rejections should be discouraging.

A third possibility is that some buyers might be making impulse offers, and delays in sellers' responses induce a 'cooling off' period after which buyers' excitement wears off and they are less likely to return with second offers. To test this hypothesis, we looked to see whether buyers were less likely to return to a slow seller for items deemed to be more impulsive or hedonic. We prompted a Large-Language Model (ChatGPT) to rank our item categories based on impulsiveness, uniqueness, and urgency, and correlated these rankings with the effect of the seller's rejection RT on the probability of the buyer making a second offer. We did not find any significant correlations between the effect of rejection RT on the probability of a second offer and the impulsiveness ranking (Spearman $r(29) = -.18$, $p = .315$), uniqueness ranking (Spearman $r(29) = -.07$, $p = .685$), or urgency ranking (Spearman $r(29) = .29$, $p = .103$), though for urgency the effect was marginal and in the expected direction. Overall, we do not find support for the effect being driven by cooling down after impulsive offers.

Modeling eBay sellers with the DDM

We have argued that bargainers do not have prepared plans, but does their decision process align with the DDM? To address this question, we adapt a value-based DDM (*12*) to the bargaining setting. Indeed, this adapted DDM provides an accurate account of sellers' behavior on eBay.

We model the drift rate in the DDM as a function of the buyer's first offer (as a fraction of the seller's list price) and of the seller's list price. In a typical value-based DDM, the drift rate is a function of the difference in subjective values between the two options. Here, we assume that the list price is a reasonable proxy for the seller's subjective value for the item. If they accept the offer, they get the buyer's cash; if they reject the offer, they get the value of the item.

We also developed a more complex formulation for non-decision time. In the DDM, non-decision time is a parameter that accounts for time that the decision maker is not evaluating or comparing the alternatives. In lab experiments, non-decision time is typically a few hundred milliseconds and modeled with a uniform random variable (*17*). It is thought to capture motor latencies (i.e., the time to press a button) and initial delays in orienting to the options. In our data there are many other factors that could contribute to non-decision time. eBay users multitask. Sellers may take hours before seeing an offer and their decision process might be interrupted by sleep, work, family, etc. Thus, we developed new models of non-decision time that use Gamma distributions (Model 2), additionally include the time of day (Model 3: sleep and work), and additionally vary with the quality of the offer (Model 4: more interruptions when sellers struggle to decide) (see Methods). Each of these models builds on the previous one (see Methods). We compared these models using WAIC, a standard method for comparing model fit while accounting for model complexity (*52*).

Relative to the model with the standard non-decision time specification (Model 1), Model 2 improved model fit for 11% of sellers, Model 3 further improved model fit for another 47% of sellers, and Model 4 even further improved model fit for 42% of sellers. Model 1 captured the choice data reasonably well but struggled to account for the RT distributions – it overpredicted the shortest RT and greatly underpredicted the longest RT.

Using the best-fitting non-decision-time model for each seller, the DDM accurately captured both choice and RT data from the eBay sellers (at least those that fit our inclusion criteria –Methods) namely the fact that sellers responded to higher offers with a higher probability of acceptance, faster acceptances, and slower rejections (Fig. 5, see Methods for additional posterior predictive checks, Figs. S25-31).

Using the best fitting model for each subject, we could capture the aggregate acceptance probability depending on offer ratio or choice difficulty (Fig. 5A, Fig. S25) and the aggregate X-shaped RT pattern (Fig. 5B, Fig. S26). The model could also account reasonably well for acceptance probability (Fig. S27) and RT quartiles (Fig. S28) across subjects. Moreover, because the models allow non-decision time to vary with the time of offer creation, they could also capture the mean RT depending on the time the offer was created (Fig. S29). The model and data regression coefficients for offer ratio on log RT (Fig. S31) and for offer ratio on choices (Fig. S30) were also highly correlated.

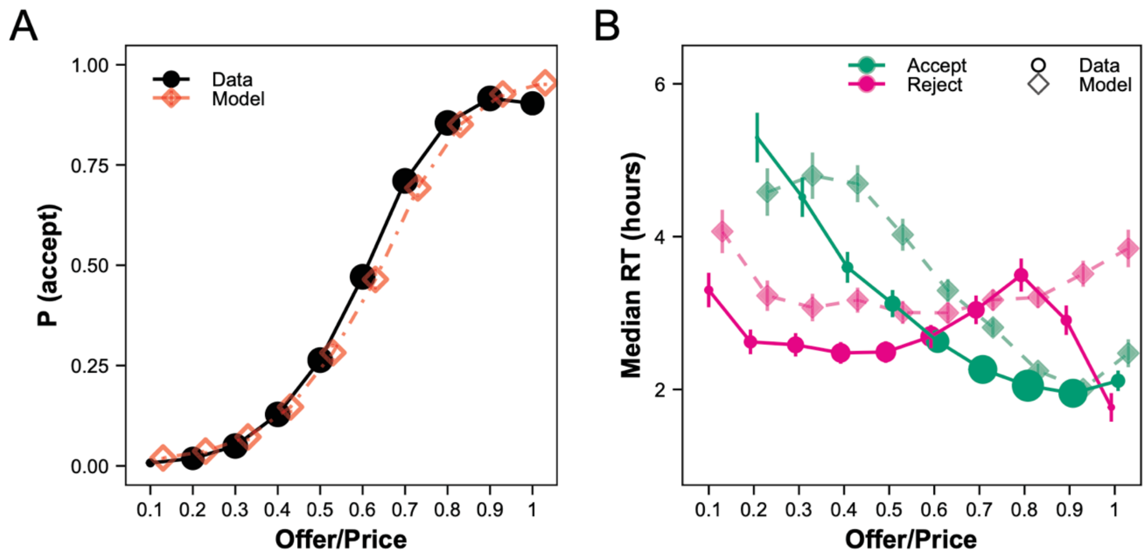

Figure 5. eBay choices and RT are well captured by the DDM. (A) Choice data and DDM fits in the pre-existing eBay data. Sellers' probability of accepting the first offer as a function of the offer ratio (offer / list price). **(B)** RT data and model fits in the pre-existing eBay data. Seller's median RT (in hours) as a function of the offer ratio. The bars represent standard errors across individuals. The dot sizes represent amount of data. Bins with less than 15 observations were excluded.

One question that we can address with the DDM is how sellers change with experience. Here, we measure experience by the number of bargaining exchanges sellers had participated in before the current offer. We find that more experienced eBay sellers exerted less response caution and evaluated offers more positively. When examining RT as a function of offer size for more experienced sellers, we observe a speeding up for acceptances, but not for rejections ($b_{experience}$ = -0.14, S.E. = 0.02, 95% CI = [-0.18, -0.10], t(255,916) = -6.98, p = 10^{-12} for acceptances, $b_{experience}$ = -0.013, S.E. = 0.03, 95% CI = [-0.07, 0.04], t(116,767) = -0.45, p = 0.66 for rejections, Table S2). Looking at correlations between DDM parameters and experience (Fig. S32), we find that seller experience was negatively correlated with boundary separation (Spearman r(510) = -0.2, p = 10^{-5}) and positively correlated with the intercept term in the drift-rate function (favoring acceptance) (Spearman r(510) = 0.28, p = 10^{-10}). Some sellers had a starting point bias towards acceptance or rejection (372 had no starting point bias, 17 had an acceptance starting point bias and 123 had a rejection starting point bias), but this did not correlate with their experience (Spearman r(510) = -0.07, p = .136).

To incorporate the possibility that sellers might be learning about their items from the market, we fit an additional model in which the drift rate depends not only on the offer ratio and list price, but also on the number of watchers, the number of views, and how many days the item had been for sale (Time of Day and Offer Ratio and Item Characteristics Gamma DDM - Model 5). We also fit a similar model where we excluded the offer ratio from the drift regression (Time of Day and Item Characteristics Gamma DDM excluding Offer Ratio from Drift Rate - Model 6) to see whether item characteristics had a larger effect than the offer ratio in determining sellers' RT.

When comparing all models using WAIC (*52*), we find that 3% of the sellers are better fit by Model 2, 10% of sellers by Model 3, 3% by Model 4, 84% by Model 5 and none by Model 6. If we only compare Model 4, Model 5 and Model 6, we find that 90% of sellers are better fit by Model 5, while 9% of sellers are better fit by Model 4. No sellers are better fit by Model 6 which excludes the offer ratio from the drift rate.

Model 5 reveals that sellers' decisions are indeed informed by the market. We found that the number of views and the number of watchers for an item pushed the seller's drift rate towards rejection, while the listing age pushed the seller's drift rate towards acceptance (95% credible interval for number of views: positive for 0%, negative for 42% of sellers; for number of watchers: positive for 0.6%, negative for 33% of sellers; for listing age in days: positive for 58%, negative for 0.9% of sellers; Fig. S33). Even for the model including item characteristics in the drift rate our previous results still hold (Fig. S34), namely a negative correlation between seller experience and boundary separation (Spearman $r(514) = -.19, p < 10^{-4}$) and a positive correlation between seller experience and drift bias (Spearman $r(514) = .28, p < 10^{-10}$).

Discussion

In summary, our results indicate that, counter to implicit assumptions in economics and game theory, many people do not have deterministic threshold strategies, but instead make decisions on the spot, in a way that reflects their strength-of-preference. Buyers and sellers on eBay take hours longer to accept unattractive offers and to reject attractive offers. They may have plans, but those plans involve noisy and lengthy evaluations of the received offers. As a result, their RT reveal private information, namely their evaluation of received offers, which opposing agents can use to gain an advantage in the bargaining process. Sellers' decisions are well explained by a computational model (DDM) where evidence in favor of accepting accumulates at a rate proportional to the size of the buyer's offer. The DDM also reveals that sellers with more experience use narrower decision boundaries and evaluate offers more positively. Surprisingly, eBay buyers do not respond to fast rejections as expected – rather than being discouraged, they are more likely to come back with a new offer.

There are a couple of key takeaways from our study. First, RT is an important strategic variable, one that has so far been ignored in game theory and until recently in economics (*9, 44, 53-61*). Practitioners, experimenters, and theorists alike must consider the consequences of people inadvertently revealing private information through RT. In the labor market, job applicants can observe how quickly they receive a job offer after an interview. A fast job offer may signal that the employer has a strong preference for the candidate. This information can affect the applicant's decision to accept the offer (*62*). In political campaigns, the speed of endorsements for a candidate can provide information about the strength of the endorser's preference. Slow endorsements may indicate weak support, while early endorsements can be more convincing (*63*). In academia, RT of referees and journal editors are correlated with the quality of the paper (*64*). RT can also be used in financial markets. It has been shown that the timing of stock analyst recommendations can impact subsequent analysts' recommendations (*65*). The speed of trades in asset market transactions can reveal information about the existence of insiders (*66, 67*). In both online and live auctions, the timing of bids can be informative. Bidding frenzies can be influenced by the frequency of bids and this tends to increase product valuations (*68*). In negotiations, when proposers' first offers are immediately accepted, they are more likely to generate counterfactual thoughts about how they could have done better and are therefore less likely to be satisfied with the agreement than are proposers whose offers are not accepted immediately (*69*). On the flip side, offers accepted after a delay may lead to greater satisfaction for the parties involved (*69*).

Second, the scope of sequential sampling models like the DDM extends beyond simple, quick decisions. The DDM has been a workhorse in perception (*15-17*), but its use has been limited to decisions on the order of seconds. However, there is nothing in the DDM that restricts it to that timescale. Our results indicate that the DDM is as relevant to buying a car as it is to choosing what to eat for lunch. Rather than making decisions algorithmically or heuristically, we argue that people typically rely on a noisy evidence accumulation and comparison process that approximates an optimal, but constrained, decision rule (*70*).

Some might wonder whether bargaining delays reflect sellers waiting for better offers. If this was the case, we would expect a systematic effect where more deliberation time (without any new offers) leads to a higher probability of offer acceptance. In other words, time ticking by with no new offers can only be bad news, pushing the seller to accept the offer they have in hand. Instead, we observe the opposite for high offers – as more time passes, sellers are more likely to reject the offers (Fig. 2C). Also, if eBay sellers were waiting for better offers then we would expect many acceptances just before the response deadline. Instead, we find that only 0.13% of acceptances occur in the final hour. The experiment from Konovalov & Krajbich (2023) also rules out this explanation in the lab because there is only one buyer and one seller.

So what is that sellers are thinking about during these long deliberation times? The consequences of accepting the offer are straightforward – receiving cash and parting with the item. The consequences of rejecting the offer are not. The seller must estimate the size and timing of future offers, as well as their own evaluation of those offers, and evaluate whether the potential improvement in price is enough to justify the expected delay. The seller must also account for the possibility that they won't get an acceptable offer on their item and will be stuck with it. Indeed, we found that sellers' decisions are responsive to the market interest in their goods – their evaluation of an offer is negatively affected by the number of views and watchers for the item and positively affected by the number of days the item has gone unsold. It is these considerations (and potentially others) that sellers deliberate over, sometimes extensively.

One interesting finding from our paper is that large errors – accepting very low offers or rejecting very high low offers – are made very quickly. While perhaps surprising, this is actually predicted by the DDM. However, it is possible that there are other reasons for these very quick responses. It could be that these cases are legitimate responses to certain kinds of products, e.g., gift cards or scam products (*72*). Because we don't have details about the specific items from the pre-existing eBay dataset, we can only speculate.

Using the DDM we also found that more experienced sellers exhibit less caution in their choices and are more likely to accept a given offer. More experienced sellers are perhaps less cautious because they are involved in more bargaining threads and so cannot afford to spend as much time on any one offer. We also suspect that they are more likely to accept a given offer because they are less attached to their goods. Indeed, professional sellers are less likely to exhibit the endowment effect (*73*).

More research is needed to understand the surprising effect that buyers are less likely to return to a seller after a slow rejection. Because a slow rejection signals a competitive offer, buyers should be encouraged and follow up with a slightly higher offer. We failed to find support for several possible explanations for this finding – buyers are impatient and move on to other sellers, buyers are impatient and purchase the item, buyers cool off after making offers on impulse goods. However, our data are not ideal for testing these hypotheses because we do not know the exact identities of the items being bargained over. An additional possibility is that a buyer may respond more favorably to a quick rejection if they are very uncertain about the value of the item. In that case, a quick rejection could signal a high value to the buyer, which could outweigh the presumably higher price. These hypotheses should be explored in future work.

Another major issue that should be addressed in future work is how bargainers choose to counteroffer rather than accept or reject. Counteroffers are challenging to model in the DDM framework because they involve a continuous decision of how much to counter with. Even ignoring that detail, counteroffers turn the decision into a three-alternative choice problem in which countering is a compromise between the other two options. We are not aware of any sequential sampling models that are appropriate for modeling counteroffers. Instead, as a robustness check, we pooled counteroffers with either acceptances or rejections, and verified that our main results still hold (see Supplementary Information; Fig. S25-31, Fig. S35-36, Table S12).

One implication of our results is that equilibrium concepts from game theory may not be well suited to describing behavior. Our results suggest that many agents don't have pre-planned strategies and so may not be fully reasoning through their decisions (*74, 59*). This is in line with research in behavioral game theory (*75*) such as work on Level-K (*75*) and Cognitive Hierarchy models (*77, 78*) where people use limited steps of thinking, Experience Weighted Attraction (EWA) where people learn strategies based on predispositions and payoff experience (*79, 80*), and Quantal Response Equilibrium (QRE) where people are noisy and believe that others are noisy as well (*81*). Other more recent work has begun to examine strategic decision-making from process-tracing and sequential sampling model perspectives (*82-86*). Our research builds on that work, using RT to demonstrate that the amount of thinking depends on the difficulty of the decision in extensive form (i.e., sequential) games.

In conclusion, we've demonstrated that people make strategic decisions, sometimes ones that take hours or days, in a way that is consistent with noisy evidence accumulation and comparison, i.e., sequential sampling or DDM. Rather than having prepared plans, they make decisions on the spot, revealing private information about their preferences. This means that RT is an important strategic variable to be incorporated into game theory, and that the range of applications of sequential sampling models like the DDM is greater than previously thought. We hope that future work will continue to push the bounds of how this framework can help us to understand complex, social, and strategic decision-making.

Materials and Methods

eBay observational data

The Backus et al. dataset consists of about a year's worth of bargaining exchanges (between May 1 2012 and June 1 2013) on eBay (*51*). We created a subset of the data by randomly selecting 20% of the sellers and then collecting all of their bargaining exchanges. The dataset contains several useful measures, including unique identifiers for items, buyers, and sellers; seller experience; item characteristics such as list price, sale price, category, list date, number of views, number of watchers, whether the item was relisted, and whether there are acceptance or rejection thresholds; and offer details such as the buyer's first offer, any subsequent offers, whether any messages were included, and most importantly for us, timestamps on all offers and replies.

Given the non-monotonicity in the response time (RT) data, it would be inappropriate to model the entire range of offers with linear regressions. Instead, we employed two complementary modeling approaches. First, we modeled RT with generalized additive model (GAM) regressions, which incorporate smooth functions of the independent variables (*87*) (Fig. S6). Second, we used the GAMs to identify the range of offers for which RT exhibited monotonic trends, based on when the first derivative of each GAM became significantly different from zero beyond their peaks (calculated using finite differences). We then used mixed-effects linear regressions on the resulting monotonic range (Table S2).

Item category features ranking

We prompted a Large-Language Model (ChatGPT-4o) to rank our item categories based on impulsiveness, uniqueness, and urgency to test whether delays in the sellers' responses induce a 'cooling off' period after which buyers' excitement wears off and they are less likely to return with second offers. We expected a positive correlation, with higher item category ranking (from 1 - most impulsive, most unique, most urgent to 32 - least impulsive, least unique, least urgent) and the coefficient for rejection RT predicting buyers' second-offer probability (higher coefficient corresponding to higher probability that the buyer comes back with a second offer after a slower rejection).

We defined impulsiveness as the "likelihood for a spontaneous purchase and instant gratification potential". We defined uniqueness as the "distinctiveness or rarity, novelty, custom or limited-edition nature, personal appeal or niche taste". We defined urgency as the "time-sensitive demand, scarcity, limited-time offer and immediacy of need".

The prompt we used is as follows:

*"Please rank these item categories in terms of **impulsiveness**, **uniqueness** and **urgency**. Provide a ranking for each term. Consider the following for definitions for each term:*

Impulsiveness:
***Likelihood of a spontaneous purchase**: How quickly a buyer might make an unplanned purchase based on emotional appeal or temptation. Emotional appeal means the degree to which an item attracts attention and triggers an immediate "must-have" feeling without much rational consideration. **Instant gratification potential**: How easily a purchase can provide immediate satisfaction or joy, leading to a quicker decision.*

Uniqueness:
***Distinctiveness or rarity**: How rare or hard-to-find the item is, making it stand out from similar items. **Novelty**: How fresh or different the item is compared to typical market offerings, creating a sense of originality. **Custom or limited-edition nature**: The degree to which the item is custom-made, handmade, or part of a limited run, making it more*

special. ***Personal appeal****: How much the item reflects individual taste or caters to niche interests, which might increase its perceived value.*

Urgency:

Time-sensitive demand*: How much pressure a buyer might feel to make a quick decision due to factors like limited availability or impending expiration of the offer.*

Scarcity*: How soon the item may run out of stock or become unavailable, creating a sense of "buy now or miss out."* ***Limited-time offer****: How much the item is tied to a temporary promotion or exclusive deal that pushes the buyer to act quickly.* ***Immediacy of need****: The extent to which the item fulfills an immediate, practical need, prompting the buyer to prioritize purchasing it quickly.”*

eBay field experiments

Both experiments were pre-registered. The preregistration for eBay field experiment 1 is available at AsPredicted (#44248): https://aspredicted.org/ZXD_DG4. The preregistration for eBay field experiment 2 is available at OSF: https://osf.io/evp2k/?view_only=bc7890e87e7f4a508d1c7af2af22af5e.

Experiment 1

Accounts and technical details

The study received approval from the OSU Institutional Review Board (2020B0152), which also granted a waiver of consent process. We used three different eBay accounts with different names to make offers to sellers. The accounts had no previous history on eBay. We also created a developer account that used the eBay API (Application Programming Interface) to send requests and return data from the eBay database. We used the eBay APIs to find items, find available information for items and sellers, and retrieve data about the offers we made and their RT. Unfortunately, the eBay API does not directly record the time when the seller responds to an offer, only when the offer is made. Therefore, to measure RT we used the sale timestamp for acceptances, the next offer creation time for counteroffers, and we manually recorded the rejection timestamp from the eBay website.

Item selection criteria

We selected items on eBay that had the following characteristics: single baseball cards that were in good, very good, like new, or brand new condition, best offer feature enabled, seller ships from the US, price between $10 - 20, minimum seller feedback of 99%, free shipping, no expiration within 3 days, listing not older than 30 days (Table S1). We decided on baseball cards because they are easy to ship and store, and they are frequently sold on eBay.

Offer criteria

We used a within-subject design in which we targeted specific sellers and made different offers to them on several different items. To do so, we targeted sellers with many items for sale. We selected the first 10,000 items from the baseball card category and selected the first 50 sellers that had more than 15 and less than 100 items for sale. For each seller, we first made an offer at 0.60 of list price. We then replaced 13 sellers that had an automatic acceptance or automatic rejection for this offer.

We then proceeded to make 10 offers to each seller, each for a different posted item, and in a random order. For each seller we made three offers each at 0.3 and 0.9, and two offers each at 0.45 and 0.75, in addition to the first offer at 0.6 (Fig. S20A). If after those 11 offers we had only one rejection for a particular seller, we made one more offer for another item at 0.25 of the list price. If after 11 offers we had only one acceptance for a particular seller, we made one more offer for another item at 0.95 of the list price. If we still had only one acceptance or one rejection for a seller, we excluded that seller and added a new seller. We replaced 6 sellers in this way.

There was a 12-hour expiration time on all offers. A seller received only one offer per day and we spaced the offers so that consecutive offers did not overlap. We made approximately 50 offers per day between 9am and 12pm (ET) on weekdays. We avoided making offers to sellers who indicated that they were away. We kept making offers to the selected sellers for approximately two months between July and August 2020.

Data

In one instance we couldn't retrieve the RT for a rejection so we used the rejection message timestamp instead. In total, we had 549 offers, with 3 auto-acceptances, 34 auto-rejections, 20 expirations, 242 acceptances, 79 rejections and 171 countered offers (Table S13).

In some cases, we couldn't retrieve the covariate information. This happened in 11 cases (2%) for seller's feedback, registration year on eBay, and number of listings available at the start of data collection, and 3 cases for the number of previous offers for an item (0.5%).

Experiment 2

Accounts and technical details

The study received approval from the OSU Institutional Review Board (2020B0152). We used three different eBay accounts with different names and addresses to make offers to sellers. The accounts had no previous history on eBay. One account was only used to make first offers to the sellers. We used the remaining two accounts to make the rest of the offers. We used the eBay API to find items, find available information for items and sellers, and retrieve data about the offers made and their RT. To measure RT we used the sale timestamp for acceptances, the next offer creation time for counteroffers, and a customized script for rejections. The script accessed our accounts every ~5 minutes to check whether any offers had changed from pending to rejected. In rare cases the script failed, in which case we instead used the message timestamps for rejections. These messages were sent by eBay but with some delay. We estimated this delay as being uniformly distributed between 0 and 1 hour.

Item selection criteria

We selected items on eBay that had the following characteristics: single trading cards, best offer feature enabled, seller ships from the US, price between $10-17 dollars, minimum seller feedback of 99%, and no free shipping. We chose items from the following item categories: sports cards (baseball, soccer, football, basketball, hockey) in excellent, near mint or very good condition, and collectible card games (magic: the gathering, pokemon, yu-gi-oh) in near mint or better, lightly played (excellent) or moderately played (very good) conditions. We switched to selecting only items that did not offer free shipping so that the seller's response would be solely based on their value for the item.

Offer criteria

We again used a within-subject design in which we targeted specific sellers and made different offers to them on several different items. We selected the first 10,000 items from the sports category and the first 10,000 items from the collectible card games category and then selected 150 sellers with more than 21 and less than 500 items. For each seller, we first made an offer at 0.50 of list price. We then replaced 39 sellers that had an automatic acceptance or automatic rejection for this offer. We initially set the item price range to be between $7 -17, but realized that we could not make our lowest offers for items below $10 since all offers on eBay have to be at least $0.99. Therefore, we later replaced 38 sellers who did not have 20 viable items in this new price range.

We then proceeded to make 20 offers to each of the selected sellers in. We made the following offers to each seller in a random order: two offers at 0.10, 0.15, 0.20, one offer at 0.25, 0.30, 0.35, 0.40, 0.45, 0.60, 0.65, two offers at 0.70, 0.75, 0.80 for a total of 21 offers per seller (including the first offer of 0.5) (Fig. S20B). We chose to oversample at the extreme offer levels to make it more likely to see some acceptances for low offers and some rejections for high offers.

There was a 24-hour expiration time on all offers. A seller received only one offer per day and we spaced the offers so that consecutive offers did not overlap. We made approximately 100 offers per day between 9am and 1pm (ET) on weekdays. We avoided making offers to sellers who indicated that they were away, which led to some delays and incomplete number of offers for some sellers. Some sellers also deleted their listings during the experiment which also led to incomplete number of offers. If a seller made a counteroffer, we always let it expire. We kept making offers to the selected 150 sellers for about two months between February and April 2023.

Data

In a few cases we couldn't retrieve the RT. This happened in 18 cases for acceptances (1.3% of acceptances), in 7 cases for rejections (1.2% of rejections) and 6 cases for counteroffers (0.8% of counteroffers). We excluded these observations. In total, we had 3,037 offers, with 11 auto-acceptances, 182 auto-rejections, 200 expired, 1308 acceptances, 565 rejections and 771 counteroffers (Table S14).

In some cases, we couldn't retrieve covariate information. This happened in 21 cases (0.7%) for seller's feedback, registration year on eBay, and number of listings available at the start of data collection and 54 cases for the number of previous offers for an item (1.8%).

Source of RTs for Rejections

For rejections, 24% of RT come from message notifications and 76% come from the eBay API. Based on 436 observations, the mean delay in messages timestamps compared to our script timestamp was 26.6 minutes (Mean = 26.6 min, SD = 17.6 min, Median = 26.7 min). We also ran the regression of log(RT) on offer ratio with a dummy variable for the source of the RT, but the results did not change significantly (Table S21).

eBay DDM

Data

We fitted the DDM at the seller level for the observational eBay data used in the behavioral analyses (Supplementary Methods). We imposed the same restrictions on the data as before, except that some restrictions (see restrictions 14-22) were imposed only on the first offer, not the entire exchange. Because the model requires many trials, we selected all the sellers from the dataset that had at least 50 rejections and 50 acceptances, leaving us with 534 sellers, 240,620 listings, and 263,215 bargaining threads. We didn't consider counteroffers in the modeling.

DDM specifications

The DDM assumes that people integrate evidence in favor of accepting or rejecting an offer over time, with a drift rate (v), specified by a drift intercept (b0) and drift slopes (b1 & b2), which represents the attractiveness of accepting the offer. Decisions begin at the starting point (z) and are made when the accumulated evidence reaches either the acceptance (a) or rejection boundaries (-a). The response time is composed of the amount of time needed for the accumulated evidence to reach a threshold and a non-decision time (t) that represents time that the decision-maker is not evaluating the offer.

We started by fitting a standard DDM (Model 1) to each seller's choice and RT data. For each seller we estimated the following parameters: starting point, boundary separation, non-decision time, drift rate intercept and drift rate slope for ratio of first offer to list price and drift rate slope for list price. Although this model captured the choice data reasonably well (Fig. S27A), it had trouble accounting for RT (Fig. S28A). Although the correlation between the data and the model was high and significant, the model overpredicted the fastest RTs (25th quantile), slightly underpredicted the median RTs (50^{th} quantile), and greatly underpredicted the slowest RTs (75th quantile). We first sought to improve the model before examining how its parameters relate to bargaining exchange characteristics.

On eBay, sellers might not be regularly checking their accounts, or they might be otherwise occupied (e.g., asleep or at work). Therefore, we considered non-decision times that depend on the time of day when the offer was made. Fig. S12B shows a plot of daily activity on eBay. In these data, sellers mostly responded to offers between 7am and 9pm (PT). There was also a slight bimodality in the activity levels with higher activity during the morning and evenings.

The standard DDM treats non-decision time as a stationary distribution. We next considered a version that treats non-decision time as a gamma distribution (Model 2). The gamma distribution is a waiting time distribution that is characterized by shape and rate parameters (*88*). If events occur according to a Poisson process with rate λ, then the waiting time until the occurrence of the nth event follows a gamma distribution with parameters $(n, 1/\lambda)$.

$$t \sim \mathcal{G}(\alpha, \beta) \tag{1}$$

We also considered a gamma-distribution model that further allows the shape parameter to depend on the time of day when the offer was created (Model 3). We let the shape parameter be a function of 3 other parameters (h_1, h_2, h_3). We chose this sinusoidal function because it can produce longer non-decision times at night as well as two peaks of activity during the day.

$$t \sim \mathcal{G}(\alpha_H, \beta) \tag{2}$$

$$\alpha_H = \exp\left(h_0 + h_1 \cos\frac{2\pi H}{24} + h_2 \cos^2\frac{2\pi H}{24}\right) \tag{3}$$

We also let non-decision time depend on the difficulty of the decision, namely a quadratic function of the offer ratio (Model 4). If difficult decisions take more time, they are more likely to be interrupted, resulting in additional non-decision time. So, we let the shape parameter of the gamma distribution additionally depend on linear and quadratic effects of offer ratio
(d_1, d_2).

$$t \sim \mathcal{G}(\alpha_H, \beta) \tag{4}$$

$$\alpha_H = \exp\left(h_0 + h_1 \cos\frac{2\pi H}{24} + h_2 \cos^2\frac{2\pi H}{24} + d_1\left(\frac{p_1}{p_2}\right) + d_2\left(\frac{p_1}{p_2}\right)^2\right) \tag{5}$$

We use the following weakly informative prior for the parameters.

$$\alpha \sim \mathcal{N}(2,1), a \geq 0 \tag{6}$$

$$z \sim \mathcal{N}(0.5,0.5), 0 \leq z \leq 1 \tag{7}$$

$$\beta_0 \sim \mathcal{N}(0,3) \tag{8}$$

$$\beta_1 \sim \mathcal{N}(0,3) \tag{9}$$

$$\beta_2 \sim \mathcal{N}(0,3) \tag{10}$$

$$h_0 \sim \mathcal{N}(0,2) \tag{11}$$

$$h_1 \sim \mathcal{N}(0,2) \tag{12}$$

$$h_2 \sim \mathcal{N}(0,2) \tag{13}$$

$$\beta \sim \mathcal{N}(0,2), \beta \geq 0 \tag{14}$$

$$d_1 \sim \mathcal{N}(0,3) \tag{15}$$

$$d_2 \sim \mathcal{N}(0,3) \tag{16}$$

For the standard DDM we used the same priors as above with the following exception:

$$\beta_0 \sim \mathcal{N}(0,0.3) \tag{17}$$

$$\beta_1 \sim \mathcal{N}(0,0.3) \tag{18}$$

$$\beta_2 \sim \mathcal{N}(0,0.3) \tag{19}$$

We used the Ohio Supercomputer Center to fit the data. We ran 3 separate chains for each

subject. Each chain consisted of 3,200 samples out of which 700 were warm-up samples. In cases where the models did not converge, we used 10,000 samples out of which 5,000 were warm-up samples. We computed $\hat{R}$ of all parameters to assess model convergence.

There were 22 unique sellers for which the model did not converge (7 for Standard DDM – Model 1, 5 for Gamma DDM – Model 2, 7 for Time of Day Gamma DDM – Model 3, 13 for Time of Day and Offer Ratio Gamma DDM – Model 4). For the following analyses we excluded these sellers and were left with 512 total sellers. The maximum $\hat{R}$ was less than 1.05 for all other subjects indicating the models converged successfully (*89*).

For Models 5 and 6, there were 44 unique sellers for whom the models did not converge ($\hat{R} > 1.1$; 18 sellers for Time of Day and Offer Ratio and Item Characteristics Gamma DDM - Model 5 and 33 sellers for Time of Day and Item Characteristics Gamma DDM excluding Offer Ratio from Drift Rate - Model 6). For the analyses including these models we excluded these sellers and were left with 489 total sellers.

In order to test the robustness of the results to excluding counteroffers, we also fit Models 1-4 first treating counteroffers the same as rejections and second treating counteroffers the same as acceptances.

We used the same fitting procedure as before. For the model fitting when pooling counteroffers with rejections, there were 26 unique sellers for whom the models did not converge (5 for Standard DDM – Model 1, 9 for Gamma DDM – Model 2, 10 for Time of Day Gamma DDM – Model 3, 12 for Time of Day and Offer Ratio Gamma DDM – Model 4). For the analyses using these models we excluded these sellers and were left with 507 total sellers.

For the model fitting when pooling counteroffers with rejections, there were 69 unique sellers for whom the models did not converge (17 for Standard DDM – Model 1, 31 for Gamma DDM – Model 2, 30 for Time of Day Gamma DDM – Model 3, 34 for Time of Day and Offer Ratio Gamma DDM – Model 4). For the analyses using these models we excluded these sellers and were left with 463 total sellers.

DDM posterior predictive checks

In order to check how well the models fit the data, we generated simulations of choices and RT using the mean parameter values for the best fitting model for each subject. We generated 10 simulations for each trial. Each trial's drift rate and non-decision time was determined using the trial's offer creation time, offer ratio and list price. We used the *RWiener* package to generate RT and choices (*90*). We also generated simulations for each model separately to compare how well each model fit the data.

We checked the model fit using the following criteria: how well the model matched probability of acceptance across subjects, how well the model matched RT quartiles for acceptances and rejections across subjects, how well the model could predict the aggregate relationship between offer ratio and RT and the aggregate relationship between offer ratio and probability of acceptance, how well the model matched the mean RT for each hour in the day, and how well the model matched the regression coefficients for offer ratio on RT and on choices.

Acknowledgments
We thank Fangzhi Li, Astin Lin, Vidhur Narayan, Katherine Scantling, Jocelyn Spanbauer, Selina Yu for helping with data collection for the eBay field experiments and the Castel family, Alexandrea Howard and Josh Hascher for helping with the experimental materials for the eBay field experiments. We thank Kota Saito for suggesting using eBay data. We thank the Ohio Supercomputing Center for computational resources. I.K. gratefully acknowledges support from the NSF award #2333979.

References

1. V. Gates, F. Callaway, M. K. Ho, T. L. Griffiths, A rational model of people's inferences about others' preferences based on response times. *Cognition* **217**, 1–21 (2021).
2. A. M. Evans, P. P. F. M. van de Calseyde, The effects of observed decision time on expectations of extremity and cooperation. *Journal of Experimental Social Psychology* **68**, 50–59 (2017).
3. C. R. Critcher, Y. Inbar, D. A. Pizarro, How quick decisions illuminate moral character. *Social Psychological and Personality Science* **4**, 308–315 (2013).
4. A. Konovalov, I. Krajbich, Revealed strength of preference: Inference from response times. *Judgment and Decision Making* **14**, 381–394 (2019).
5. J. A. Clithero, Response times in economics: Looking through the lens of sequential sampling models. *Journal of Economic Psychology* **69**, 61–86 (2018).
6. A. Alekseev, Using response times to measure ability on a cognitive task. *J Econ Sci Assoc* **5**, 65–75 (2019).
7. A. N. Tump, T. J. Pleskac, R. H. J. M. Kurvers, Wise or mad crowds? The cognitive mechanisms underlying information cascades. *Science Advances* **6**, 1–11 (2020).
8. C. Frydman, I. Krajbich, Using Response Times to Infer Others' Private Information: An Application to Information Cascades. *Management Science* **68**, 2970–2986 (2022).
9. F. Echenique, K. Saito, Response time and utility. *Journal of Economic Behavior & Organization* **139**, 49–59 (2017).
10. T. Groves, J. O. Ledyard, "Incentive Compatibility Since 1972" in *Information, Incentives, and Economics Mechanisms*, T. Groves, R. Radner, S. Reiter, Eds. (University of Minnesota Press, NED-New edition., 1987) *Essays in Honor of Leonid Hurwicz*, pp. 48–112.
11. P. P. F. M. Van de Calseyde, G. Keren, M. Zeelenberg, Decision time as information in judgment and choice. *Organizational Behavior and Human Decision Processes* **125**, 113–122 (2014).
12. A. Konovalov, I. Krajbich, Decision Times Reveal Private Information in Strategic Settings: Evidence from Bargaining Experiments. *The Economic Journal*, uead055 1– 12 (2023).
13. A. Diederich, Dynamic Stochastic Models for Decision Making under Time Constraints. *Journal of Mathematical Psychology* **41**, 260–274 (1997).
14. R. M. Roe, J. R. Busemeyer, J. T. Townsend, Multialternative decision field theory: A dynamic connectionst model of decision making. *Psychological Review* **108**, 370–392 (2001).
15. R. Bogacz, E.-J. Wagenmakers, B. U. Forstmann, S. Nieuwenhuis, The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences* **33**, 10–16 (2010).
16. M. N. Shadlen, D. Shohamy, Decision Making and Sequential Sampling from Memory. *Neuron* **90**, 927–939 (2016).
17. R. Ratcliff, G. McKoon, The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation* **20**, 873–922 (2008).
18. M. Usher, J. L. McClelland, The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review* **108**, 550–592 (2001).

19. J. S. Trueblood, S. D. Brown, A. Heathcote, The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychol Rev* **121**, 179–205 (2014).
20. R. Webb, The (Neural) Dynamics of Stochastic Choice. *Management Science* **65**, 230–255 (2019).
21. D. Fudenberg, P. Strack, T. Strzalecki, Speed, Accuracy, and the Optimal Timing of Choices. *American Economic Review* **108**, 3651–3684 (2018).
22. S. U. Maier, A. Raja Beharelle, R. Polanía, C. C. Ruff, T. A. Hare, Dissociable mechanisms govern when and how strongly reward attributes affect decisions. *Nat Hum Behav* **4**, 949–963 (2020).
23. I. Krajbich, C. Armel, A. Rangel, Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* **13**, 1292–1298 (2010).
24. X. Yang, I. Krajbich, A dynamic computational model of gaze and choice in multi-attribute decisions. *Psychological Review* **130**, 52–70 (2023).
25. G. Fisher, A multiattribute attentional drift diffusion model. *Organizational Behavior and Human Decision Processes* **165**, 167–182 (2021).
26. J. R. Busemeyer, Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **11**, 538–564 (1985).
27. L. T. Hunt, N. Kolling, A. Soltani, M. W. Woolrich, M. F. S. Rushworth, T. E. J. Behrens, Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* **15**, 470–476 (2012).
28. M. Glickman, O. Sharoni, D. J. Levy, E. Niebur, V. Stuphorn, M. Usher, The formation of preference in risky choice. *PLOS Computational Biology* **15**, e1007201 1– 25 (2019).
29. V. Zilker, T. Pachur, Nonlinear probability weighting can reflect attentional biases in sequential sampling. *Psychological Review* **129**, 949–975 (2022).
30. W. J. Zhao, L. Walasek, S. Bhatia, Psychological mechanisms of loss aversion: A drift-diffusion decomposition. *Cognitive Psychology* **123** 1– 22 (2020).
31. J. Dai, J. R. Busemeyer, A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *Journal of Experimental Psychology: General* **143**, 1489–1514 (2014).
32. W. J. Zhao, A. Diederich, J. S. Trueblood, S. Bhatia, Automatic biases in intertemporal choice. *Psychon Bull Rev* **26**, 661–668 (2019).
33. C. A. Hutcherson, B. Bushong, A. Rangel, A Neurocomputational Model of Altruistic Choice and Its Implications. *Neuron* **87**, 451–462 (2015).
34. M. Milosavljevic, J. Malmaud, A. Huth, C. Koch, A. Rangel, The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making* **5**, 437–449 (2010).
35. B. De Martino, S. M. Fleming, N. Garrett, R. J. Dolan, Confidence in value-based choice. *Nat Neurosci* **16**, 105–110 (2013).
36. K. X. Chiong, M. Shum, R. Webb, R. Chen, Combining Choice and Response Time Data: A Drift-Diffusion Model of Mobile Advertisements. *Management Science*, doi: 10.1287/mnsc.2023.4738 (2023).
37. M. G. Philiastides, R. Ratcliff, Influence of Branding on Preference-Based Decision Making. *Psychol Sci* **24**, 1208–1215 (2013).

38. B. M. Turner, C. A. Rodriguez, T. M. Norcia, S. M. McClure, M. Steyvers, Why more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage* **128**, 96–115 (2016).
39. D. R. Amasino, N. J. Sullivan, R. E. Kranton, S. A. Huettel, Amount and time exert independent influences on intertemporal choice. *Nat Hum Behav* **3**, 383–392 (2019).
40. S. Gluth, J. Rieskamp, C. Büchel, Deciding When to Decide: Time-Variant Sequential Sampling Models Explain the Emergence of Value-Based Decisions in the Human Brain. *J. Neurosci.* **32**, 10686–10698 (2012).
41. F. Sheng, A. Ramakrishnan, D. Seok, W. J. Zhao, S. Thelaus, P. Cen, M. L. Platt, Decomposing loss aversion from gaze allocation and pupil dilation. *Proceedings of the National Academy of Sciences* **117**, 11356–11363 (2020).
42. R. Polanía, I. Krajbich, M. Grueschow, C. C. Ruff, Neural Oscillations and Synchronization Differentially Support Evidence Accumulation in Perceptual and Value-Based Decision Making. *Neuron* **82**, 709–720 (2014).
43. R. Bhui, Testing Optimal Timing in Value-Linked Decision Making. *Comput Brain Behav* **2**, 85–94 (2019).
44. C. Alós-Ferrer, E. Fehr, N. Netzer, Time Will Tell: Recovering Preferences When Choices Are Noisy. *Journal of Political Economy* **129**, 1828–1877 (2021).
45. C. Frydman, G. Nave, Extrapolative Beliefs in Perceptual and Economic Decisions: Evidence of a Common Mechanism. *Management Science* **63**, 2340–2352 (2017).
46. V. Lerche, A. Voss, Experimental validation of the diffusion model based on a slow response time paradigm. *Psychological Research* **83**, 1194–1209 (2019).
47. P. C. Cramton, Strategic Delay in Bargaining with Two-Sided Uncertainty. *The Review of Economic Studies* **59**, 205–225 (1992).
48. J. Sobel, I. Takahashi, A Multistage Model of Bargaining. *The Review of Economic Studies* **50**, 411–426 (1983).
49. P.-H. Lin, A. L. Brown, T. Imai, J. T. Wang, S. W. Wang, C. F. Camerer, Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments. *Nat Hum Behav* **4**, 917–927 (2020).
50. C. F. Camerer, G. Nave, A. Smith, Dynamic Unstructured Bargaining with Private Information: Theory, Experiment, and Outcome Prediction via Machine Learning. *Management Science* **65**, 1867–1890 (2019).
51. M. Backus, T. Blake, B. Larsen, S. Tadelis, Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions*. *The Quarterly Journal of Economics* **135**, 1319–1361 (2020).
52. S. Watanabe, Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. *Journal of Machine Learning Research* **11**, 3571–3594 (2010).
53. C. F. Chabris, D. Laibson, C. L. Morris, J. P. Schuldt, D. Taubinsky, The Allocation of Time in Decision-Making. *Journal of the European Economic Association* **7**, 628–637 (2009).
54. A. Schotter, I. Trevino, Is response time predictive of choice? An experimental study of threshold strategies. *Experimental Economics* **24**, 87–117 (2021).
55. L. Alaoui, A. Penta, Cost-Benefit Analysis in Reasoning. *Journal of Political Economy*, doi: 10.1086/718378 881– 925 (2022).

56. J. Srivastava, S. Oza, Effect of Response Time on Perceptions of Bargaining Outcomes. *Journal of Consumer Research* **33**, 266–272 (2006).
57. J. A. Clithero, Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization* **148**, 344–375 (2018).
58. L. Spiliopoulos, A. Ortmann, The BCD of response time analysis in experimental economics. *Exp Econ* **21**, 383–433 (2018).
59. D. Gill, V. Prowse, Strategic Complexity and the Value of Thinking. *The Economic Journal* **133**, 761–786 (2023).
60. O. Gossner, J. Steiner, C. Stewart, Attention Please! *Econometrica* **89**, 1717–1751 (2021).
61. B. Hébert, M. Woodford, Rational inattention when decisions take time. *Journal of Economic Theory* **208**, 1717–1751 (2023).
62. W. J. Becker, T. Connolly, J. E. Slaughter, The effect of job offer timing on offer acceptance, performance, and turnover. *Personnel Psychology* **63**, 223–241 (2010).
63. M. Cohen, D. Karol, H. Noel, J. Zaller, *The Party Decides: Presidential Nominations Before and After Reform* (University of Chicago Press, Chicago, IL, 2008) *Chicago Studies in American Politics*.
64. D. Card, S. DellaVigna, Fast or Slow? Uncovering Expert Views from Decision Time.
65. I. Welch, Herding among security analysts. *Journal of Financial Economics* **58**, 369–396 (2000).
66. C. Camerer, K. Weigelt, Information Mirages in Experimental Asset Markets. *The Journal of Business* **64**, 463–93 (1991).
67. P. Bossaerts, C. Frydman, J. Ledyard, The Speed of Information Revelation and Eventual Price Quality in Markets with Insiders: Comparing Two Theories. *Review of Finance* **18**, 1–22 (2014).
68. G. Häubl, P. T. L. Popkowski Leszczyc, Bidding Frenzy: Speed of Competitor Reaction and Willingness to Pay in Auctions. 3296104 [Preprint] (2018). https://papers.ssrn.com/abstract=3296104.
69. A. D. Galinsky, V. L. Seiden, P. H. Kim, V. H. Medvec, The dissatisfaction of having your first offer accepted: The role of counterfactual thinking in negotiations. *Personality and Social Psychology Bulletin* **28**, 271–283 (2002).
70. A. Wald, Sequential Tests of Statistical Hypotheses. *The Annals of Mathematical Statistics* **16**, 117–186 (1945).
71. C. Terwiesch, S. Savin, I.-H. Hann, Online Haggling at a Name-Your-Own-Price Retailer: Theory and Application. *Management Science* **51**, 339–351 (2005).
72. E. Fay, Virtual Artifacts: eBay, Antiquities, and Authenticity. *Journal of Contemporary Criminal Justice* **27**, 449–464 (2011).
73. J. A. List, Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace. *Econometrica* **72**, 615–625 (2004).
74. A. Avoyan, M. Ribeiro, A. Schotter, E. R. Schotter, M. Vaziri, M. Zou, Planned vs. Actual Attention. *Management Science*, doi: 10.1287/mnsc.2023.4834 1– 22 (2023).
75. C. Camerer, *Behavioral Game Theory Experiment in Strategic Interaction* (Princeton University Press, 2003).

76. V. P. Crawford, N. Iriberri, Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? *Econometrica* **75**, 1721–1770 (2007).
77. C. F. Camerer, T.-H. Ho, J.-K. Chong, A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics* **119**, 861–898 (2004).
78. G. Coricelli, R. Nagel, Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences* **106**, 9163–9168 (2009).
79. C. Camerer, T.-H. Ho, Experienced-Weighted Attraction Learning in Normal Form Games. *Econometrica* **67**, 827–874 (1999).
80. L. Zhu, K. E. Mathewson, M. Hsu, Dissociable neural representations of reinforcement and belief prediction errors underlie strategic learning. *PNAS Proceedings of the National Academy of Sciences of the United States of America* **109**, 1419–1424 (2012).
81. R. D. McKelvey, T. R. Palfrey, Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* **10**, 6–38 (1995).
82. L. Polonio, S. Di Guida, G. Coricelli, Strategic sophistication and attention in games: An eye-tracking study. *Games and Economic Behavior* **94**, 80–96 (2015).
83. N. Stewart, S. Gächter, T. Noguchi, T. L. Mullett, Eye movements in strategic choice. *Journal of Behavioral Decision Making* **29**, 137–156 (2016).
84. R. Golman, S. Bhatia, P. B. Kane, The dual accumulator model of strategic deliberation and decision making. *Psychological Review* **127**, 477–504 (2020).
85. J. T. Wang, M. Spezio, C. F. Camerer, Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games. *The American Economic Review* **100**, 984–1007 (2010).
86. E. J. Johnson, C. Camerer, S. Sen, T. Rymon, Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory* **104**, 16–47 (2002).
87. S. N. Wood, *Generalized Additive Models: An Introduction with R, Second Edition* (Chapman and Hall/CRC, Boca Raton, ed. 2, 2017).
88. A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, D. B. Rubin, *Bayesian Data Analysis, Third Edition* (CRC Press, 2013).
89. A. Gelman, D. B. Rubin, Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science* **7**, 457–472 (1992).
90. D. Wabersich, J. Vandekerckhove, The RWiener Package: an R Package Providing Distribution Functions for the Wiener Diffusion Model. *The R Journal* **6**, 49 (2014).
91. J. J. Heckman, Sample Selection Bias as a Specification Error. *Econometrica* **47**, 153–161 (1979).
92. O. Toomet, A. Henningsen, Sample Selection Models in *R* : Package **sampleSelection**. *J. Stat. Soft.* **27** (2008).

PNAS

Supporting Information for

Deliberation During Online Bargaining Reveals Strategic Information

Miruna Cotet, Wenjia Zhao, Ian Krajbich

Ian Krajbich
Email: krajbich@ucla.edu

This PDF file includes:

Supporting text
Figures S1 to S38
Tables S1 to S22
SI References

Supplementary Text

1. eBay observational data

1.1. Data cleaning

We imposed some restrictions on the sample. All restrictions were imposed at the exchange level; even if only one offer did not meet the restrictions, the entire exchange was eliminated. For all the analyses, the following restrictions were applied (unless otherwise noted), along with the percentage of the data excluded with the restriction (ignoring the other restrictions):

Restrictions for errors:

(1) If the item is sold, the sale price is above the listing price. (0.2%)
(2) An offer is above the list price. (1.8%)
(3) Either the buyer or the seller make more than three offers. (0.2%)
(4) For an offer with a reply of “countered”, a counteroffer does not exist in the exchange. (0.1%)
(5) After an accepted offer, there is a subsequent offer in the exchange. (0.4%)
(6) There is a change in the list price during the exchange. (0.5%)
(7) There is a change in the list price during the time the item was listed. (29.7%)
(8) Buyer’s second offer is below their first offer and the third offer below their second. (2.3%)
(9) Seller’s second offer is above their first offer and third above their second. (1.7%)
(10) Seller’s offer is lower than buyer’s previous offer and buyer’s offer is higher than seller’s previous offer. (0.2%)
(11) Response time is missing or above 48 hours (0%).

After applying restrictions just for errors, we were left with 2,754,383 listings, 3,872,080 threads, 159,782 sellers and 2,002,866 buyers which included all 3 buyer offers. We have 3,872,037 first buyer offers, 846,625 second buyer offers and 266,661 third buyer offers.

Additional restrictions:

(12) List price is higher than 1000 USD. (6.7%)
(13) The seller has an accept or reject price threshold set. (36%)
(14) There is an automatic rejection. (20.4%)
(15) There is an automatic acceptance. (3.3%)
(16) There is an expired offer. (22.1%)
(17) Response time is above 47 hours and 59 minutes. (18.8%)
(18) Response time is below 10 seconds. (23.4%)
(19) Offer is rejected because the seller received a better offer. (1%)

(20) Offer is rejected because the buyer received a better offer. (1%)

(21) There is a competing offers i.e. another offer arriving before the seller decided on the initial offer. (8.6%)

(22) Message is included with the offer. (9.8%)

(23) Buyer or seller is located outside of the US. (15.8%)

The final dataset of first buyer's offers after applying all the restrictions consisted of 1,128,999 bargaining exchanges spread across 1,043,320 listings with 74,443 sellers and 743,343 buyers. Without counteroffers, we had 63,989 sellers, 602,792 buyers, 820,035 listings and 869,395 bargaining exchanges.

1.2. Results: Effects of item characteristics on RT

Looking at other influences on RT, the more desirable the items, the more cautious sellers become when responding to offers. This is illustrated by the fact that higher number of watchers led to longer RT for both rejections ($\beta_{\text{Watchers}} = 0.01, SE = 0.005, 95\%CI = [0.00, 0.02], t(116, 767) = 2.65, p = .008$) and acceptances ($\beta_{\text{Watchers}} = 0.15, SE = 0.01, 95\%CI = [0.14, 0.17], t(255, 916) = 22.96, p < 10^{-16}$), while higher number of views led to longer RT for rejections ($\beta_{\text{Views}} = 0.01, SE = 0.005, 95\%CI = [0.00, 0.02], t(116, 767) = 2.83, p = .005$). Higher list price led to slower acceptances ($\beta_{\text{Price}} = 0.03, SE = 0.01, 95\%CI = [0.02, 0.04], t(255, 916) = 5.18, p < 10^{-16}$) and faster rejections ($\beta_{\text{Price}} = -0.06, SE = 0.01, 95\%CI = [-0.07, -0.05], t(116, 767) = 10.07, p < 10^{-16}$) (Table S2 (3), (6)).

1.3. Results: RT effects of second and third buyer offers

We found similar effects of offer ratio on RT of the seller for second and third buyers offers, namely a positive relationship for rejections and negative relationship for acceptances for most of the offer ranges (Fig. S10, Tables S6, S7). We only included cases where the previous buyer's offers were rejected, not countered.

1.4. Results: RT effects with thresholds

For the main analyses, we excluded exchanges for which an accept or reject threshold was present. However, if the buyer's first offer is between these thresholds, the seller's RT can still be informative. We therefore analyze these cases as well. For these analyses, we used cases where either an accept or reject threshold was present. We regressed logRT on buyers' initial offers as a percent of list price, conditional on the seller accepting or rejecting the offers. We only used offer ratios larger than 0.36 and smaller than 0.68 similar to the analyses of observations without thresholds. As in the main analyses, we found similar effects of buyers' first offer ratio on sellers' RT, namely a positive relationship for rejections and negative relationship for acceptances (Fig. S11).

1.5. Results: RT effects with expired offers

We find that the effect of offer ratio on rejection RT is still positive when including expired offers with maximum RT ($\beta_{\text{Offer/Price}} = 0.21, SE = 0.01, 95\%CI = [0.19, 0.24], t(183,068) = 19.64, p < 10^{-15}$) (Table S3 (1)). The maximum RT was 48 hours.

1.6. Buyer's second offer and strategic use of RT

In order to investigate whether there was selection bias when analyzing buyers' second offers, we first looked at the probability of the buyer making a second offer as a function of the first offer (Fig. S16). Buyers were most likely to make a second offer if their first offer was around half of the list price. For other first offers there was a sharp decrease in the likelihood of making a second offer (Fig. S16, Table S9). To account for this selection bias, we used the Heckman correction (*1*) (Table S9-10). The Heckman correction is a two-step statistical approach that corrects for non-randomly selected samples. In the first stage, we modeled the probability of the buyer making a second offer conditional on the first offer being rejected (Table S9 (1), (2)). This model helps identify the factors that influence whether the buyer makes a second offer or the selection process and calculates the inverse Mills ratio. The inverse Mills ratio is a measure of the probability of an observation being included in the sample, given its characteristics and the estimated parameters from the selection equation. It provides a way to quantify the selection bias that arises when the sample is not randomly selected from the population. In the second stage, we correct for self-selection by incorporating the inverse Mills ratio as an additional explanatory variable (Table S10 (1), (2)). This adjusted regression accounts for the selection bias by incorporating the correction factor from the selection equation. We estimated this model using the R package sampleSelection (*2*).

Buyers were more, not less, likely to make a second offer after a faster rejection ($\beta_{\text{Rejection RT}} = -0.098, SE = 0.004, 95\%CI = [-0.11, -0.09], z(226,649) = -26.24, p < 10^{-16}$) (Fig. 4B, Table S9). Then, using a linear regression, we regressed the size of the buyer's second offer on the RT of the seller's rejection. Conditional on making a second offer, buyers offered less, not more, after faster rejections (($\beta_{\text{Rejection RT}} = 0.0026, SE = 0.0005, 95\%CI = [0.002, 0.004], t(226,645) = 5.01, p < 10^{-7}$)) (Fig. S17, Table S10). The latter result may be due to a selection bias – only the most interested buyers return after a slow rejection, and they are the ones likely to increase their offers the most. In any case, counter to the lab data and to our predictions, buyers on eBay appear to be encouraged by fast rejections.

2. eBay field experiments

2.1. Procedure: Experiment 1

We also collected additional data for a separate project that is not analyzed here. For the additional data, if the first offer was declined, we made a second offer that varied in terms of how much it increased compared to the first offer. We incremented the first offer in order to achieve one of three levels of the second offer (as a fraction of list price): {0.5, 0.7, 0.9}. These second offers were determined ahead of time. Each first offer had a uniform distribution over the higher second offers. For example, first offers of 0.6 had 50% second offers of 0.7 and 50% second offers of 0.9. First offers of 0.75 and 0.9 had no second offers.

2.2. Procedure: Experiment 2

In order to determine the appropriate sample size we performed a power analysis using the experiment 1 data. We performed the power analysis for acceptances and rejections separately. We selected sellers randomly and for each seller we selected offers randomly with replacement. We ran a mixed effects regression of log RT on offer ratio, with random intercepts at the seller level. We performed this 1000 times for each number of sellers and each number of offers per seller for acceptances and rejections and counted the number of times the coefficient on offer ratio was significant at 5% level in the expected direction. In order to reach a power level of 0.80 we needed at least 100 sellers with 10 acceptances and 5 rejections per seller. We therefore decided to collect 150 sellers and 21 offers for each seller given that we expected that some offers will be automatically accepted, countered or expire, and we might lose some sellers due to insufficient items or delays in making offers over the course of the experiment.

2.3. Results: Experiment 1

2.3.1. Summary statistics

The higher the offer was, the more likely it is to be accepted ($\beta_{\text{Offer/Price}} = 3.55, SE = 0.44, 95\% CI = [2.79, 4.54], z(319) = 8.07, p = 10^{-15}$) (Table S15 (1)). Similar to the eBay observational data, low offers were more likely to be declined while mid-range offers were more likely to be countered (Fig. S21A).

2.3.2. Hypothesis 1: main result

To test the hypothesis, a mixed-effects linear regression was performed with the log of seller's RT as the dependent variable. As independent variables the seller's response, dummy coded as 0 for acceptances and 1 for rejections, and the first offer as a fraction of list price were used, as well as their interaction with the type of response (accept or reject). Full seller random effects were included. In cases where the models did not converge or had singular boundary issues, a model comparison was performed between models with simpler random effects structure and the best model according to the AIC criterion was chosen. A significantly negative coefficient on first offer ratio was expected, indicating that higher offers were accepted more quickly. A significantly positive coefficient on the interaction between first offer ratio and reject was expected, indicating that for rejections, higher offers were responded to relatively more slowly than acceptances.

We did not find a significant effect of offer ratio on the seller's RT for acceptances ($\beta_{\text{Offer/Price}} = -0.20, SE = 0.21, 95\%CI = [-0.61, 0.22], t(315) = -0.94, p = .348$). However, we found a significantly positive coefficient on the interaction between offer ratio and rejection ($\beta_{\text{Offer/Price:Rejected}} = 1.11, SE = 0.50, 95\%CI = [0.12, 2.10], t(315) = 2.20, p = .029$), indicating that for rejections, higher offers were responded to relatively more slowly compared to acceptances (Table S16 (1)). This was not the case for counteroffers ($\beta_{\text{Offer/Price}} = -0.06, SE = 0.19, 95\%CI = [-0.42, 0.31], t(167) = -0.30, p = .766$) (Table S20 (1)). We also found a significant main effect for rejections on RTs, with rejection decisions being overall slower than acceptance decisions ($\beta_{\text{Rejected}} = 1.34, SE = 0.60, 95\%CI = [0.16, 2.52], t(315) = 2.23, p = .027$) (Table S16 (1)).

In addition to the interaction model, we also wanted to test whether rejections of higher offers are made more slowly; the previous model only evaluates rejection speed relative to acceptance speed. To do so we ran the following two mixed-effects linear regressions, analyzing accepted and rejected offers separately. Although the coefficient on offer ratio on acceptance RT is negative, it is not significant ($\beta_{\text{Offer/Price}} = -0.19, SE = 0.21, 95\%CI = [-0.61, 0.24], t(238) = -0.87, p = .384$). However, we found a marginally significant positive coefficient on ratio offer for the rejections ($\beta_{\text{Offer/Price}} = 0.84, SE = 0.45, 95\%CI = [-0.06, 1.74], t(75) = 1.87, p = .066$) (Table S17 (1), (4)).

2.3.3. Hypotheses 2 and 3: seller experience and item characteristics

We added to the models for acceptance and rejections logRT the following covariates for item characteristics: number of best offers the item received, whether the item was relisted. For this experiment, after 30 days the item is automatically relisted by eBay. The seller can also chose to relist the item manually. We didn't find any significant effects of item characteristics on RTs (Table S17 (2), (3), (5), (6)).

We also added to the models for acceptance and rejections logRT the following covariates for seller characteristics: number of feedbacks received, number of viable items considered for the experiment, years since registration. However, these covariates didn't have an effect on acceptance or rejection RTs (Table S17 (2), (3), (5), (6)).

2.4. Results: Experiment 2

2.4.1. Summary statistics

The higher the offer was, the more likely it is to be accepted ($\beta_{\text{Offer/Price}} = 2.63, SE = 0.12, 95\%CI = [2.41, 2.87], z(1,871) = 22.70, p < 10^{-16}$) (Table S15 (4)). As in the eBay observational data, low offers were more likely to be declined while mid-range offers were more likely to be countered (Fig. S21B). There was also a tendency to let very low offers expire (Fig. S21B). Even though we try to exclude sellers that had automatic thresholds set, we still encountered cases where our offers were automatically declined. As expected, this was more likely to happen for very low offers (Fig. S21B).

2.4.2. Hypothesis 1: main result

Similar to what we observed in the eBay observational data, we expected that sellers' acceptance times would decrease with increasing offer size and sellers' rejection time would increase with increasing offer size. We did not have a clear hypothesis about counteroffers but would naively expected them to look similar to rejections.

To test the hypothesis, a mixed-effects linear regression was performed with the log of seller's RT as the dependent variable. As independent variables the seller's response, dummy coded as 0 for acceptances and 1 for rejections, and the first offer as a fraction of list price were used, as well as their interaction with the type of response (accept or reject). Full seller random effects were included. In cases where the models did not converge or had singular boundary issues, we performed a model comparison between models with simpler random effects structure and chose the best model according to the AIC criterion. As expected, we found a significantly negative coefficient on offer ratio for acceptance ($\beta_{\text{Offer/Price}} = -0.25, SE = 0.07, 95\%CI = [-0.39, -0.12], t(1, 862) = -3.77, p = .0003$), indicating that higher offers are accepted more quickly. We also found an expected and significantly positive coefficient on the interaction between offer ratio and rejection ($\beta_{\text{Offer/Price:Rejected}} = 0.28, SE = 0.14, 95\%CI = [0.02, 0.55], t(1, 862) = 2.09, p = .039$), indicating that for rejections, higher offers are responded to relatively more slowly compared to acceptances (Table S16 (2)). There was also a positive main effect of rejections on RT ($\beta_{\text{Rejected}} = 0.76, SE = 0.17, 95\%CI = [0.43, 1.09], t(1, 862) = 4.50, p = 10^{-4}$), indicating that rejections were slower than acceptances.

In addition to the interaction model, we also wanted to test whether rejections of higher offers are made more slowly; the previous model only evaluates rejection speed relative to acceptance speed. To do so we ran two mixed-effects linear regressions, analyzing accepted and rejected offers separately. We found the expected negative coefficient on offer ratio for the acceptances ($\beta_{\text{Offer/Price}} = -0.26, SE = 0.07, 95\%CI = [-0.40, -0.12], t(1, 302) = -3.59, p = .0006$), and a positive, but not significant coefficient on offer ratio for the rejections ($\beta_{\text{Offer/Price}} = 0.12, SE = 0.11, 95\%CI = [-0.09, 0.33], t(559) = 1.11, p = .278$) (Table S18 (1), (4)). We did not find a significant coefficient for countered offers on offer ratio ($\beta_{\text{Offer/Price}} = 0.07, SE = 0.09, 95\%CI = [-0.11, 0.24], t(767) = 0.76, p = .448$) (Table S20 (4)).

2.4.3. Hypotheses 2 and 3: seller experience and item characteristics

We expected that sellers with less desirable items such as items with older listing dates would be faster to accept offers and slower to reject offers. To test the effect of item characteristics on RT, we added measures of item characteristics to the previous model, including number of previous offers for the item, whether the item was relisted. We did not find any significant effects of item characteristics on RTs (Table S18 (2), (3), (5), (6)).

We also expected seller experience to decrease both acceptance and rejection RT. In addition, we expected that higher seller experience would make their rejection times less responsive to offer ratio. To test the effect of seller experience on RT, we added measures of seller experience to the previous model, such as seller's number of feedback rating received, seller's registration year on eBay and number of listings available at the start of data collection.

Only one of the measures of seller experience was significant and only on acceptance RT: years since registration. Contrary to expectations, the main effect of years of experience increased acceptance RT ($\beta_{\text{Years}} = 0.31, SE = 0.14, 95\%CI = [0.05, 0.58], t(1, 272) = 2.30, p = .023$) and the interaction effect was significantly negative ($\beta_{\text{Years:Offer/Price}} = -0.18, SE = 0.08, 95\%CI = [-0.34, -0.03], t(1, 272) = -2.31, p = .023$) (Table S18 (3)).

3. eBay DDM

3.1. Parameter Recovery

We performed parameter recovery to test whether our model parameters could be successfully recovered. We generated simulated datasets and then fit the models to the simulated datasets. The range of parameters values used for each model was based on the parameter values for the fitted models on the data. For each parameter, we used a parameter range that was within 5th and 95th quantile of the fitted parameter values. We used Latin square cube sampling to generate parameter combinations for each model. We performed sampling once, generating 200 parameter combinations for each model. For each parameter combination, we selected a random seller's data to simulate a dataset using the parameter combination. The seller was selected randomly from all sellers whose calculated drift rates were within the 99th and 1st quantiles of all the seller's drift rates distribution. We did this in order to avoid cases with too low or too high drift rate that could cause issues for the RWiener function (*90*). We then fit the models to the simulated datasets.

We used the Ohio Supercomputer Center to fit the data. We ran 3 separate chains for each subject. Each chain consisted of 10,000 samples out of which 5,000 were warm-up samples. We computed $\hat{R}$ of all parameters to assess model convergence.

We had 17 parameter combination for which a model did not converge (none for Standard DDM, 6 for Gamma DDM, 3 for Time of Day Gamma DDM, 8 for Time of Day and Offer Ratio Gamma DDM). The maximum $\hat{R}$ was less than 1.05 for all other subjects indicating the models converged successfully (*89*). We plot the correlation between generated and recovered parameter values for each model. The results are summarized in Figs. S37-38, which show that parameters can be successfully recovered.

3.2. Results: DDM with item characteristics

When looking at the mean posterior values for the item-characteristic parameters of the drift rate of Model 5 we find that higher number of views and higher number of watchers make the seller more likely to reject the offer while higher number of days since the listing was posted make the seller more likely to accept the offer (Fig. S33).

Number of views: The 95% credible interval is above 0 for 0 sellers, below 0 for 215 sellers (42%), and overlaps 0 for 301 sellers (58%). Number of watchers: The 95% credible interval is above 0 for 3 sellers (0.6%), below 0 for 173 sellers (33%), and overlaps 0 for 340 sellers (66%). Listing age (days): The 95% credible interval is above 0 for 300 sellers (58%), below 0 for 5 sellers (0.9%), and overlaps 0 for 211 sellers (41%).

Even for the model including item characteristics in the drift rate our previous results still hold, namely a negative correlation between seller experience and boundary separation (Spearman $r(514) = -.19, p < 10^{-4}$) and a positive correlation between seller experience and drift bias (Spearman $r(514) = .28, p < 10^{-10}$) (Fig. S34).

3.3. Results: Pooling DDM

3.3.1. Pooling Counteroffers with Rejections

We pooled counteroffers with rejections and refitted all the models in order to show that our conclusions do not change.

7% of sellers were best fit by Model 2, 42% of sellers were best fit by Model 3 and 50% of sellers were best fit by Model 4.

Using the best-fitting non-decision-time model for each seller, the DDM still captured both choice and RT data from the eBay sellers, namely the fact that sellers responded to higher offers with a higher probability of acceptance, faster acceptances, and slower rejections (Fig. S25 E, Fig. S26 E). The model overestimates the RT for acceptances compared to the data and slightly overestimates the RT for rejections compared to the data.

We find that more experienced eBay sellers exerted less response caution and evaluated offers more positively. Looking at correlations between DDM parameters and experience (Fig. S35), we find that seller experience as measured by the number of previous best offer exchanges the seller had participated in was negatively correlated with boundary separation (Spearman $r(505) = -0.18, p = 10^{-4}$). Although seller experience as measured by number of previous best offer exchanges seller had participated in was not significantly positively correlated with the intercept term in the drift-rate function (favoring acceptance) (Spearman $r(505) = 0.28, p = .639$), other measures of seller experience were positively correlated with the intercept in the drift-rate function (Number of listings created dating back to 2008: Spearman $r(505) = 0.11p = 0.014$; Number of previous feedbacks received: Spearman $r(505) = 0.15, p = 10^{-3}$).

Using the best fitting model for each subject, we could capture the aggregate acceptance probability depending on offer ratio or choice difficulty (Fig. S25 E) and the aggregate X- shaped RT pattern (Fig. S26 E). The model could also account reasonably well for acceptance probability (Fig. S27 E) and RT quartiles (Fig. S28 E) across subjects. Moreover, because the models allow non-decision time to vary with the time of offer creation, they could also capture the mean RT depending on the time the offer was created (Fig. S29 E). The model and data regression coefficients for offer ratio on log RT (Fig. S31 E) and for offer ratio on choices (Fig. S30 E) were also highly correlated.

3.3.2. Pooling Counteroffers with Acceptances

We also pooled counteroffers with acceptances and refitted all the models in order to show that our conclusions do not change.

8% of sellers were best fit by Model 2, 46% of sellers were best fit by Model 3 and 46% of sellers were best fit by Model 4.

Using the best-fitting non-decision-time model for each seller, the DDM still captured both choice and RT data from the eBay sellers, namely the fact that sellers responded to higher offers with a higher probability of acceptance, faster acceptances, and slower rejections (Fig. S25 F, Fig. S26 F). The model slightly overestimates the RT for both acceptances and rejections compared to the data.

Similar to our original results, we find that more experienced eBay sellers exerted less response caution and evaluated offers more positively. Looking at correlations between DDM parameters and experience (Fig. S36), we find that seller experience as measured by the number of previous best offer exchanges the seller had participated in was negatively correlated with boundary separation (Spearman $r(461) = -0.18, p = 10^{-4}$) and was positively correlated with the intercept term in the drift rate function (favoring acceptances) (Spearman $r(461) = 0.29, p = 10^{-9}$).

Using the best fitting model for each subject, we could capture the aggregate acceptance probability depending on offer ratio or choice difficulty (Fig. S25 F) and the aggregate X- shaped RT pattern (Fig. S26 F). The model could also account reasonably well for acceptance probability (Fig. 27 F) and RT quartiles (Fig. S28 F) across subjects. Moreover, because the models allow non-decision time to vary with the time of offer creation, they could also capture the mean RT depending on the time the offer was created (Fig. S29 F). The model and data regression coefficients for offer ratio on log RT (Fig. S31 F) and for offer ratio on choices (Fig. S30 F) were also highly correlated.

Supplementary Figures

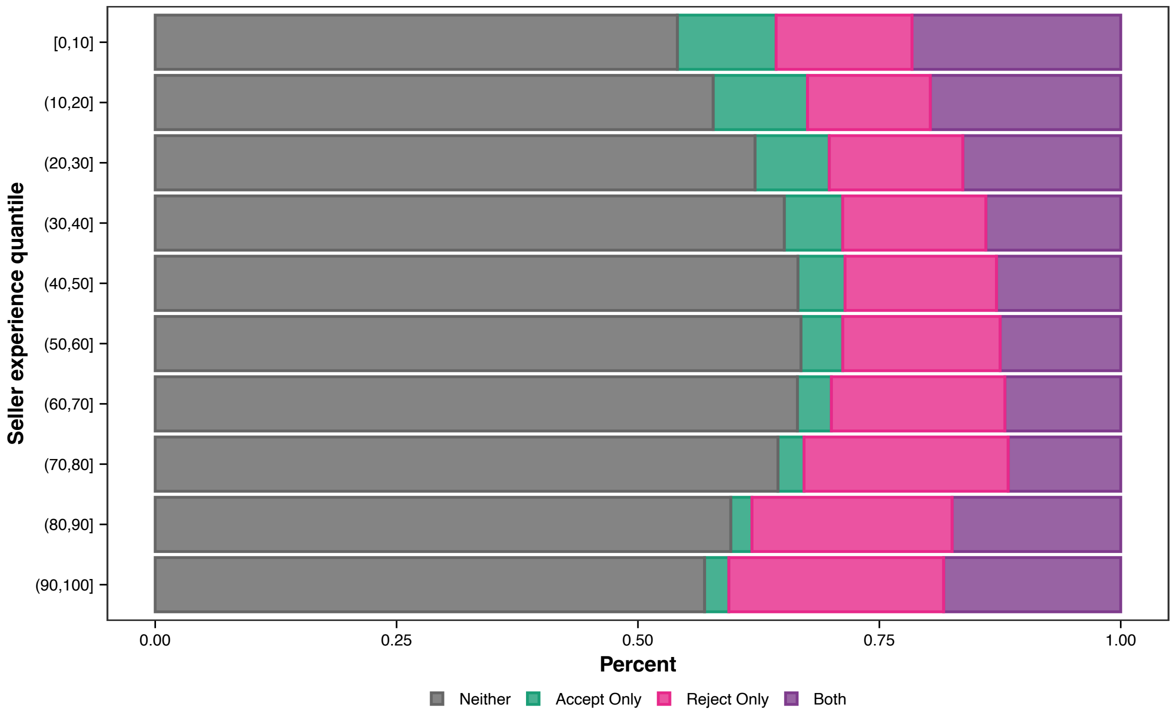

Figure S1. Acceptance and rejection threshold usage by seller experience for eBay observational data. Probability of using an accept or reject threshold, both thresholds or neither threshold, as a function of seller experience quantile. Seller experience represents the number of previous bargaining exchanges the seller has participated in. All bargaining exchanges were sorted by seller experience into ten equal sized bins. In other words, each observation was a single exchange. We did it this way, rather than sorting at the seller level, because a single seller might have some exchanges with thresholds and others without.

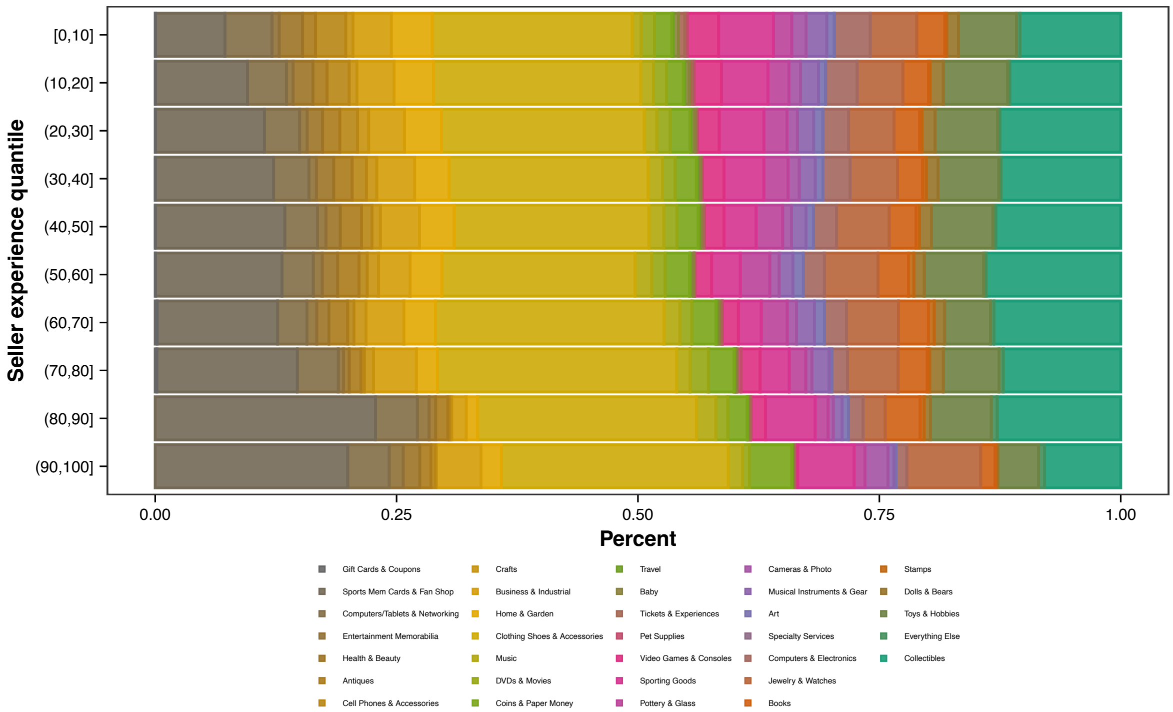

Figure S2. Item category usage by seller experience for eBay observational data. Proportion of each item category sold as a function of seller experience quantile. Seller experience represents the number of previous bargaining exchanges the seller has participated in. All bargaining exchanges were sorted by seller experience into ten equal sized bins; each observation was a single exchange.

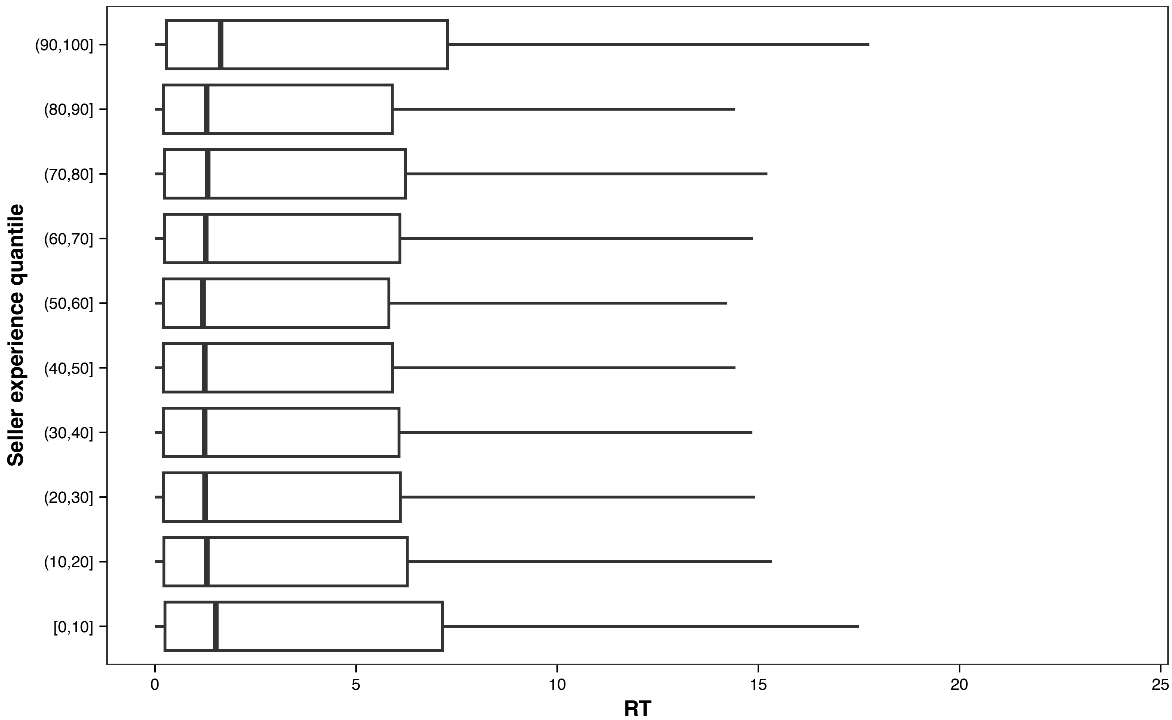

Figure S3. RT by seller experience for eBay observational data. RT boxplot as a function of seller experience quantile. Seller experience represents the number of previous bargaining exchanges the seller has participated in. All bargaining exchanges were sorted by seller experience into ten equal sized bins; each observation was a single exchange.

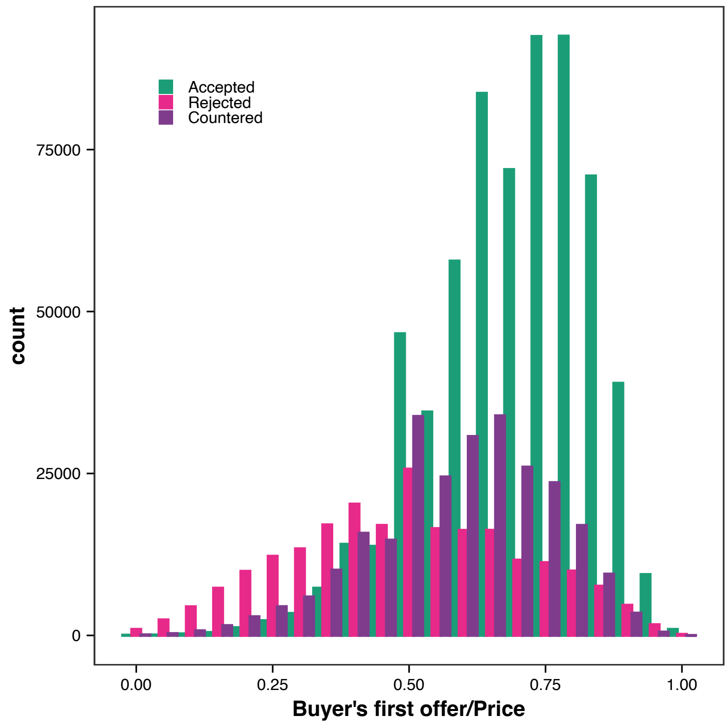

Figure S4. Distribution of buyers' first offer conditional on sellers' response for eBay observational data. Histograms of buyers' initial offers as a fraction of the sellers' list prices, conditional on the seller accepting, rejecting, or countering the offers.

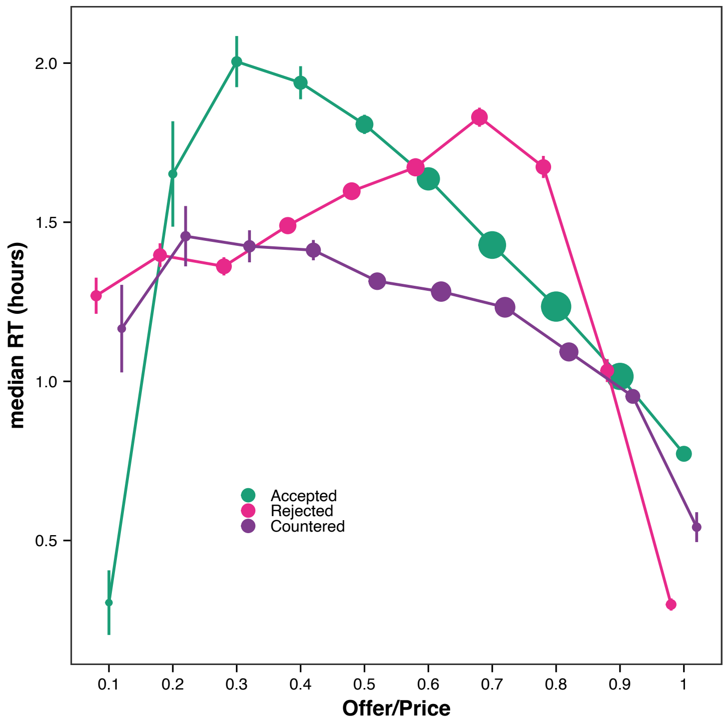

Figure S5. Sellers' median RT for buyers' initial offers by sellers' response for eBay observational data. Sellers' median RT (in hours) as a function of buyers' initial offers as a fraction of the sellers' list prices, conditional on the seller accepting or rejecting or countering the offers. The size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors.

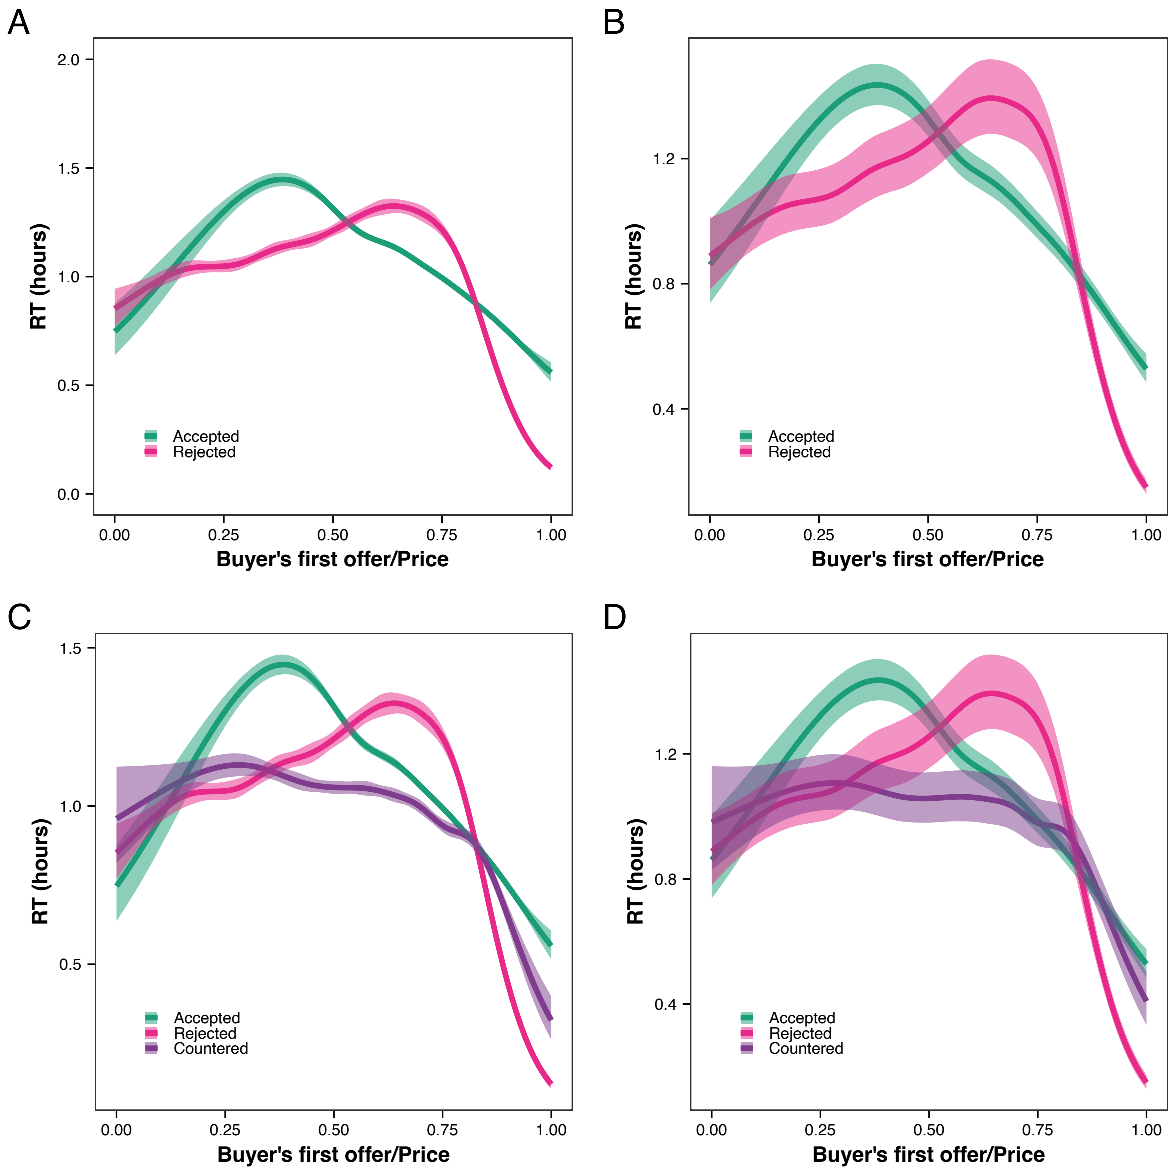

Figure S6. GAM regression predictions for sellers' RT by buyers' initial offer by type of seller response for eBay observational data. (A) GAM regression model fits without covariates. (B) GAM regression model fits with covariates. (C) GAM regression model fits with counteroffers without covariates. (D) GAM regression model fits with counteroffers with covariates. Shaded regions represent 95% confidence intervals. Covariates include item characteristics (list price, number of watchers, number of views, photo count, the list age in days, and whether the item was relisted) or not and buyer and seller characteristics (number of previous best offer exchanges buyer has participated in, number of previous best offer exchanges seller has participated in, seller's number of previous feedbacks received at the time of the offer, number of listings created by the seller dating back to 2008, number of

Best Offer-listings created by the seller dating back to 2008).

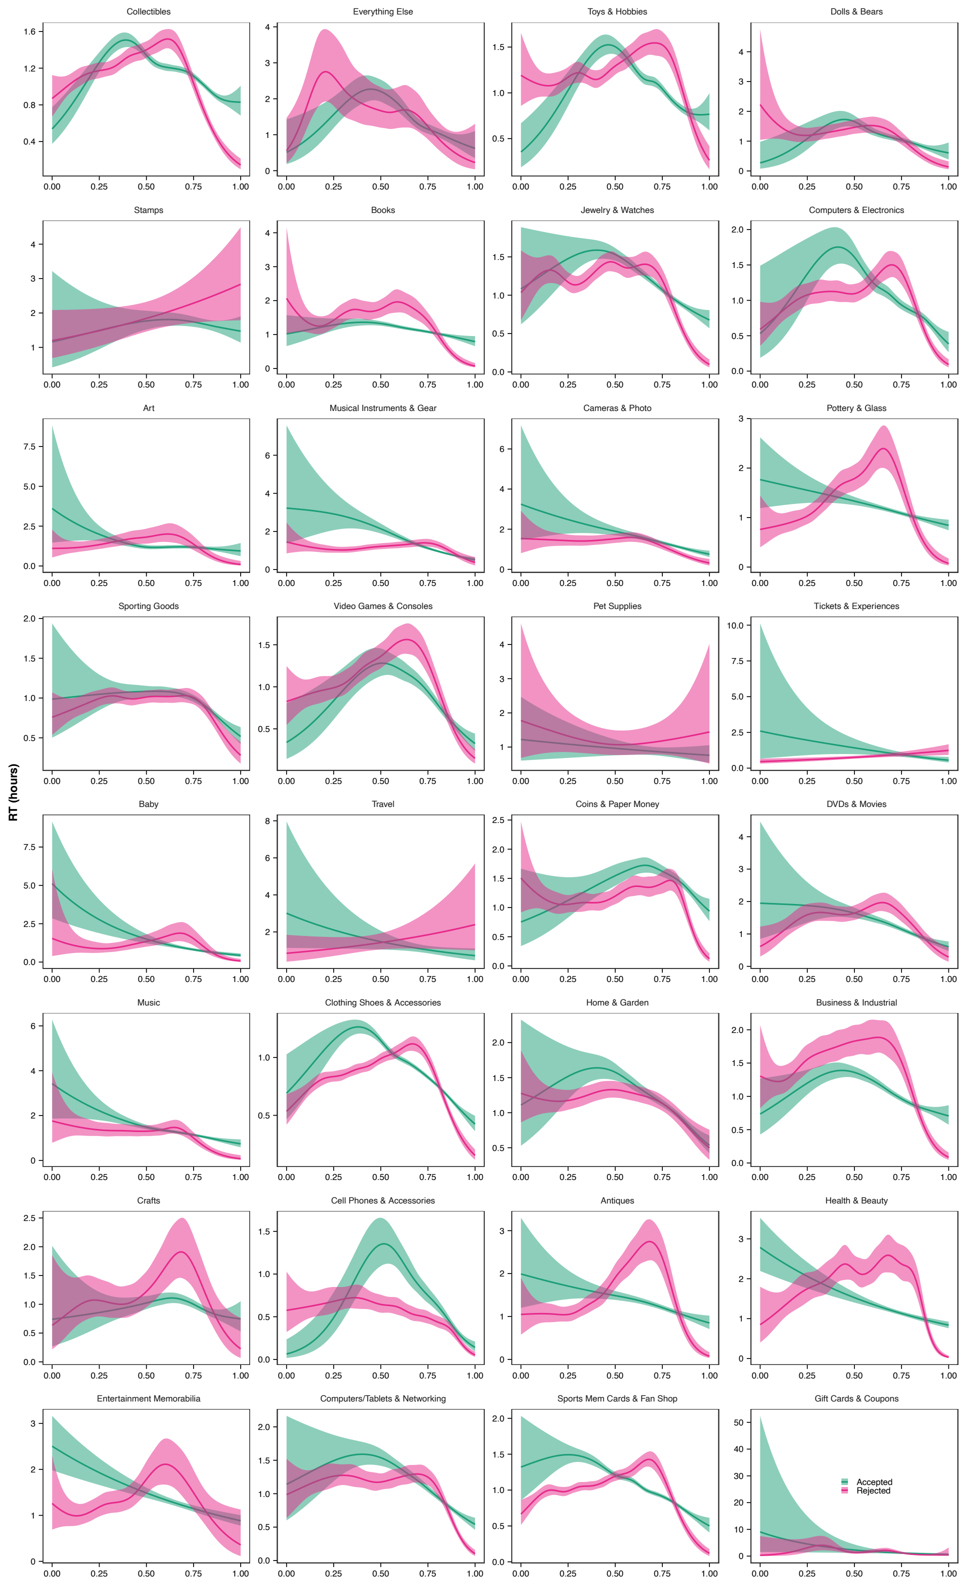
Collectibles
Everything Else
Toys & Hobbies
Dolls & Bears
Stamps
Books
Jewelry & Watches
Computers & Electronics
Art
Musical Instruments & Gear
Cameras & Photo
Pottery & Glass
Sporting Goods
Video Games & Consoles
Pet Supplies
Tickets & Experiences
Baby
Travel
Coins & Paper Money
DVDs & Movies
Music
Clothing Shoes & Accessories
Home & Garden
Business & Industrial
Crafts
Cell Phones & Accessories
Antiques
Health & Beauty
Entertainment Memorabilia
Computers/Tablets & Networking
Sports Mem Cards & Fan Shop
Gift Cards & Coupons
Accepted
Rejected
RT (hours)
Buyer's first offer/Price

Figure S7. GAM regression predictions for each item category for eBay observational data. GAM regression model predictions to sellers' RT (in hours) for buyers' initial offers as a fraction of the sellers' list prices, conditional on the seller accepting or rejecting the offers, broken down at the item category level. Shaded regions represent 95% confidence intervals.

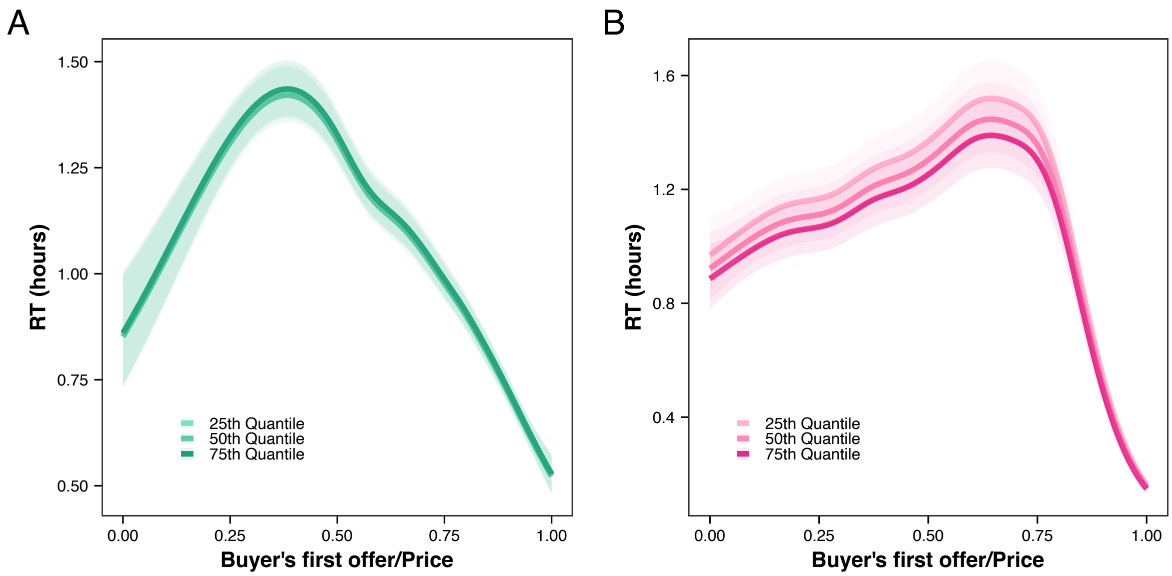

Figure S8. GAM regression predictions for different levels of list price for eBay observational data. (A) Acceptance RT. (B) Rejection RT. The shaded regions represent 95% confidence intervals. The levels of list price used for the predictions are: the 25th quantile = 17 dollars, 50th quantile = 39.95 dollars, 75th quantile = 99.99 dollars.

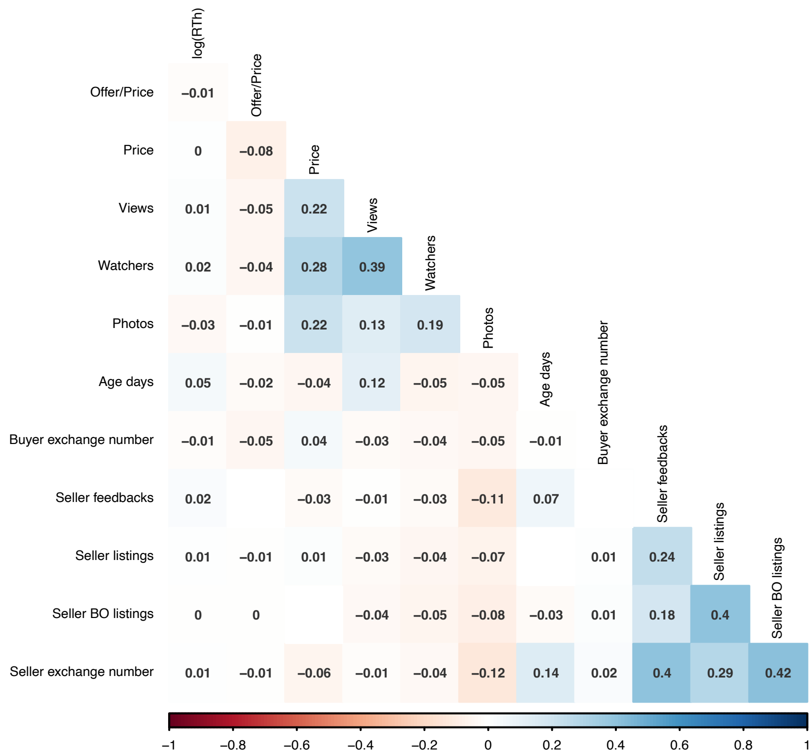

Figure S9. Correlations between variables used in mixed effects regression models for eBay observational data. Offer characteristics variables: Offer/Price - buyer's first offer as a percent of list price; log(RTh) - RT of the seller to the offer in hours. Item characteristics variables: Price - list price; Views - number of views; Watchers - number of watchers; Photos - number of photos; Age days - the listing age in days; Seller and buyer characteristics variables: Buyer exchange number - number of previous best offer exchanges buyer has participated in; Seller feedbacks - seller's number of previous feedbacks received at the time of the offer; Seller listings - number of listings created by the seller dating back to 2008; Seller BO listings - number of Best-Offer-listings created by the seller dating back to 2008; Seller exchange number - number of previous best offer exchanges seller has participated in. Only significant correlations ($p < 0.01$) are displayed. $r = 0$ indicates correlation is smaller than 0.005.

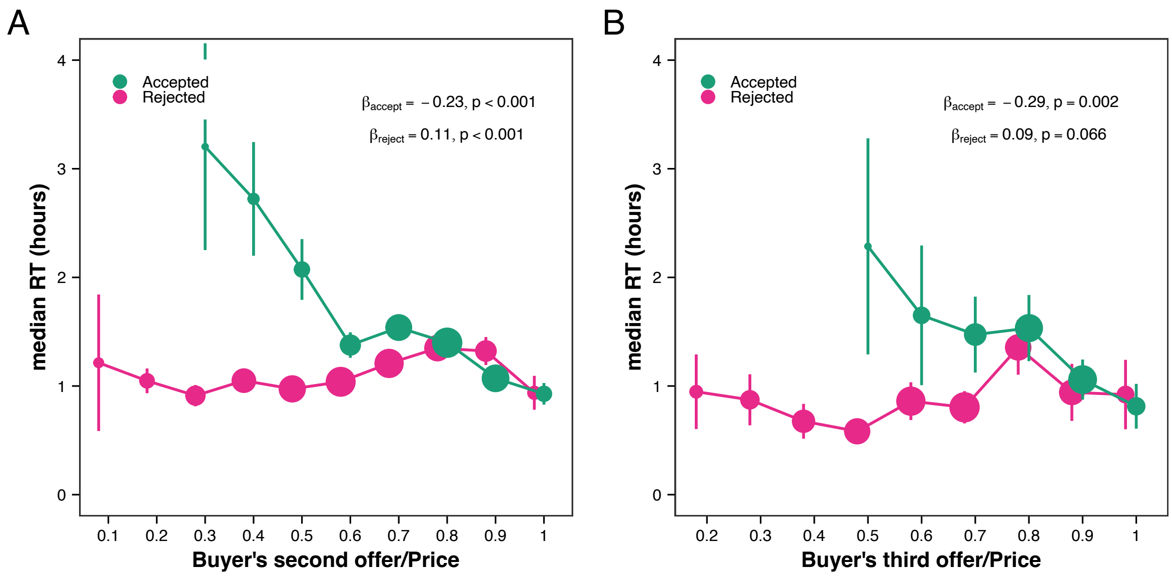

Figure S10. Sellers' median RT as a function of buyers' additional offers for eBay observational data. (A) Sellers' median RT (in hours) for buyers' second offers as a fraction of the sellers' list prices, conditional on the seller accepting or rejecting the offers. (B) Sellers' median RT (in hours) for buyers' third offers as a fraction of the sellers' list prices, conditional on the seller accepting or rejecting the offers. The size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors. Bins with less than 25 observations were excluded.

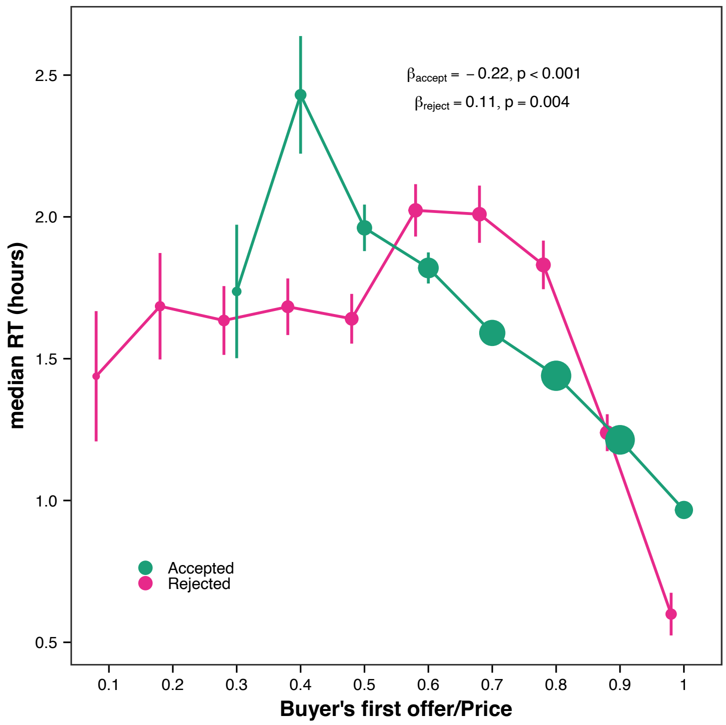

Figure S11. Sellers' median RT for buyers' initial offers by sellers' response for exchanges with thresholds for eBay observational data. Sellers' median RT (in hours) as a function of buyers' initial offers as a percent of list price (p_1/p_0), conditional on the seller accepting or rejecting the offers for bargaining exchanges with either rejection or acceptance threshold and offers that were between these thresholds. The size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors. Bins with less than 100 observations were excluded. The coefficients are from a linear regression of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller accepting, or rejecting the offers using a restricted offer range ($p_1/p_0 = [0.36, 0.68]$) similar to the analyses of observations without thresholds. The regressions also includes random effects (clustered by seller) on the intercept and first buyer offer ratio.

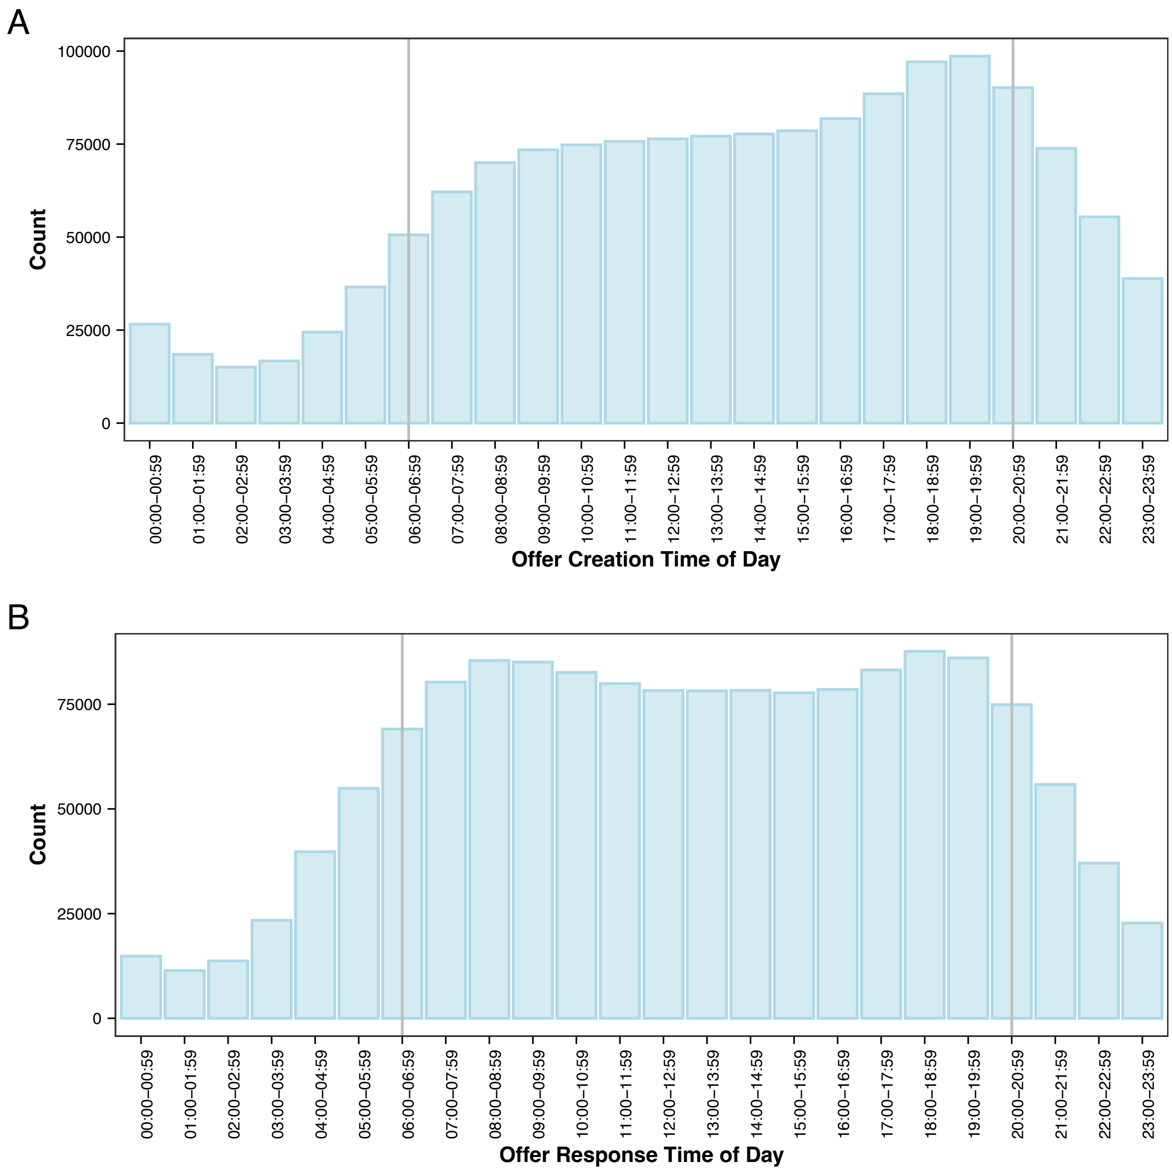

Figure S12. Activity levels by time of day for eBay observational data. (A) Histogram with number of buyer's first offer creation as a function of hour in the day (PT). (B) Histogram with number of seller's first offer responses as a function of hour in the day (PT).

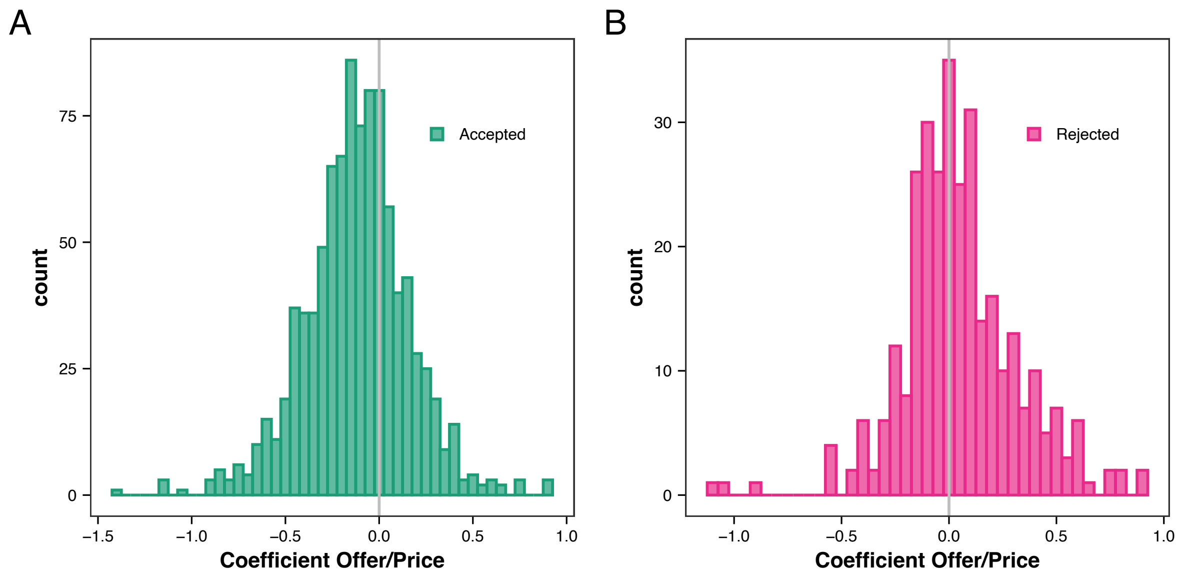

Figure S13. Seller level RT effects for eBay observational data. (A) Histogram of coefficients of buyer's first offer ratio as a fraction of the sellers' list prices at the seller level for acceptance log RT in hours. (B) Histogram of coefficients of buyer's first offer ratio as a fraction of the sellers' list prices at the seller level for rejection log RT in hours. Only sellers with more than 50 acceptances or more than 50 rejections were included.

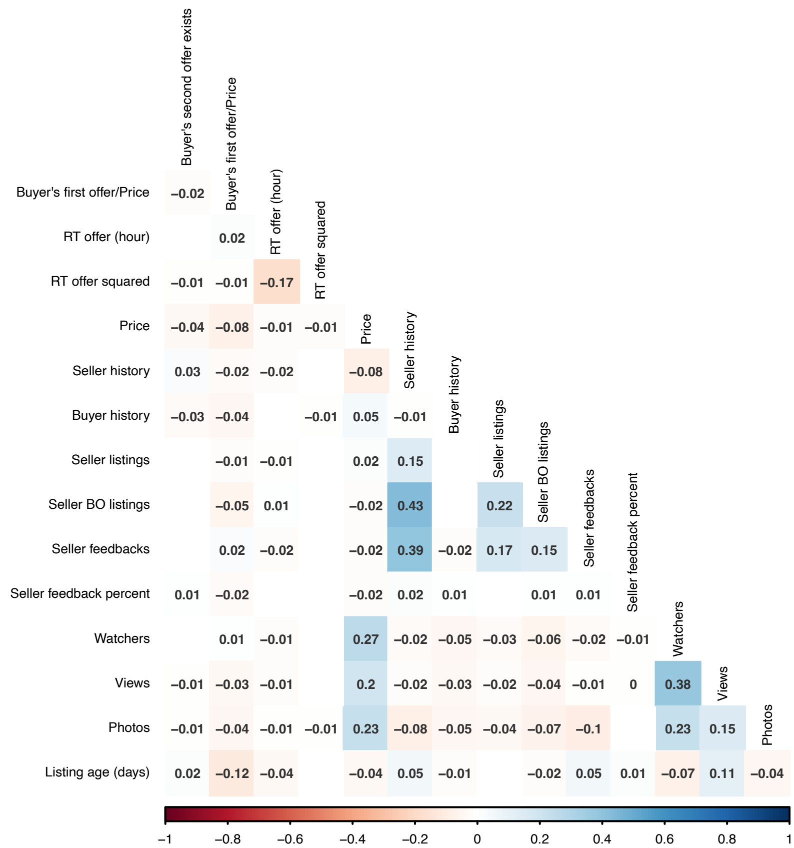

Figure S14. Correlations between existence of second buyer offer and other variables for eBay observational data. Offer characteristics variables: Buyer's second offer exists - Whether buyer's second offer exists; Buyer's first offer/Price - buyer's first offer as a percent of list price; RT offer (hours) - RT of the seller to the offer in hours; RT offer squared - squared RT of the seller to the offer in hours. Item characteristics variables: Price - list price; Views - number of views; Watchers - number of watchers; Photos - number of photos; Listing age (days) - the listing age in days; Seller and buyer characteristics variables: Buyer history - number of previous best offer exchanges buyer has participated in; Seller feedbacks - seller's number of previous feedbacks received at the time of the offer; Seller listings - number of listings created by the seller dating back to 2008; Seller BO listings - number of Best-Offer-listings created by the seller dating back to 2008; Seller history - number of previous best offer exchanges seller has participated in. Only significant correlations ($p < 0.01$) are displayed.

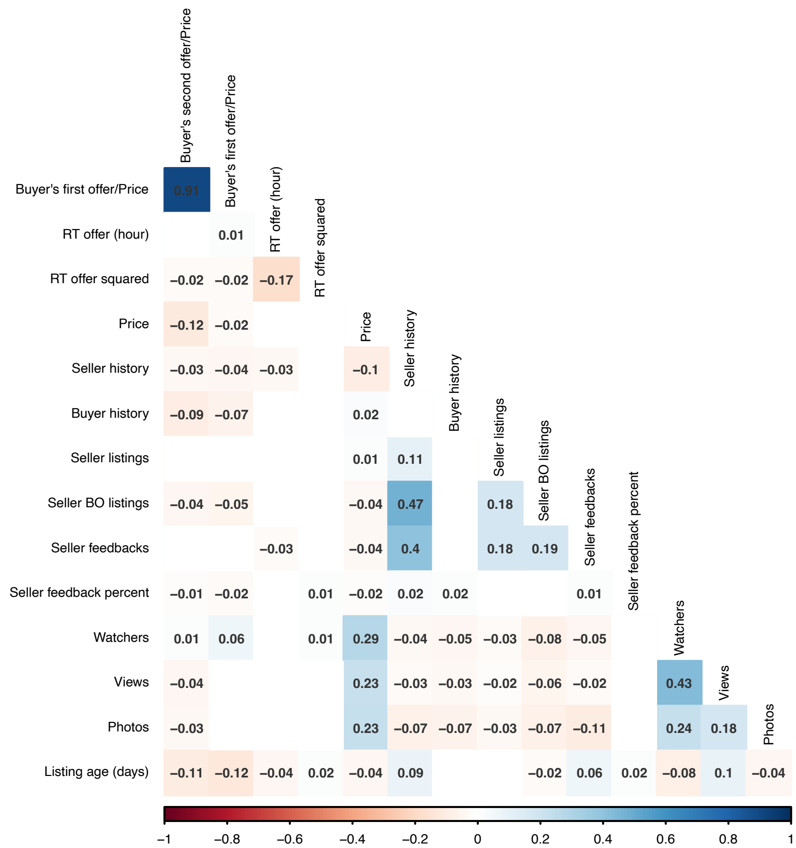

Figure S15. Correlations between second buyer offer amount and other variables for eBay observational data. Offer characteristics variables: Buyer's second offer/Price - buyer's second offer as a percent of list price; Buyer's first offer/Price - buyer's first offer as a percent of list price; RT offer (hours) - RT of the seller to the offer in hours; RT offer squared - squared RT of the seller to the offer in hours. Item characteristics variables: Price - list price; Views - number of views; Watchers - number of watchers; Photos - number of photos; Listing age (days) - the listing age in days; Seller and buyer characteristics variables: Buyer history - number of previous best offer exchanges buyer has participated in; Seller feedbacks - seller's number of previous feedbacks received at the time of the offer; Seller listings - number of listings created by the seller dating back to 2008; Seller BO listings - number of Best-Offer-listings created by the seller dating back to 2008; Seller history - number of previous best offer exchanges seller has participated in. Only significant correlations ($p < 0.01$) are displayed.

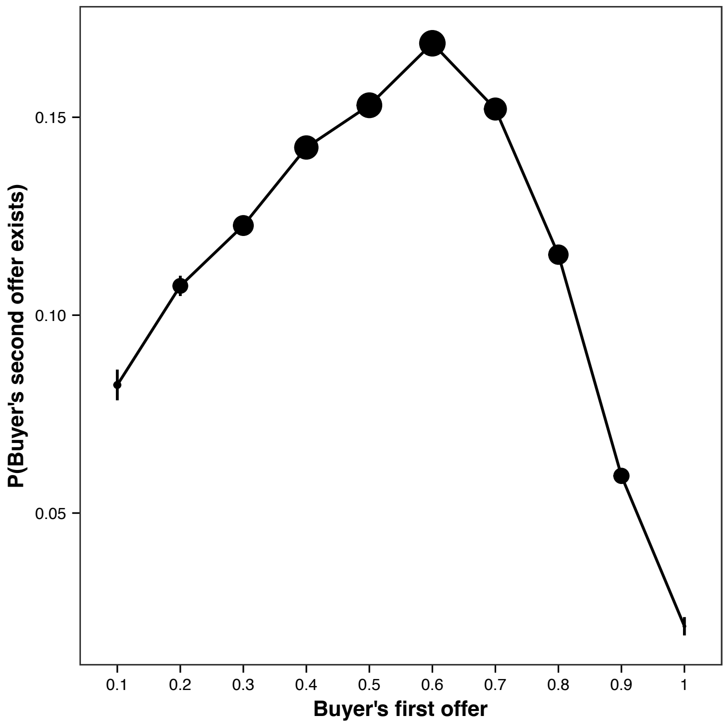

Figure S16. Second buyer offer existence probability for eBay observational data. Probability of second buyer offer existence as a function of first buyer's offer as a fraction of seller's list price.

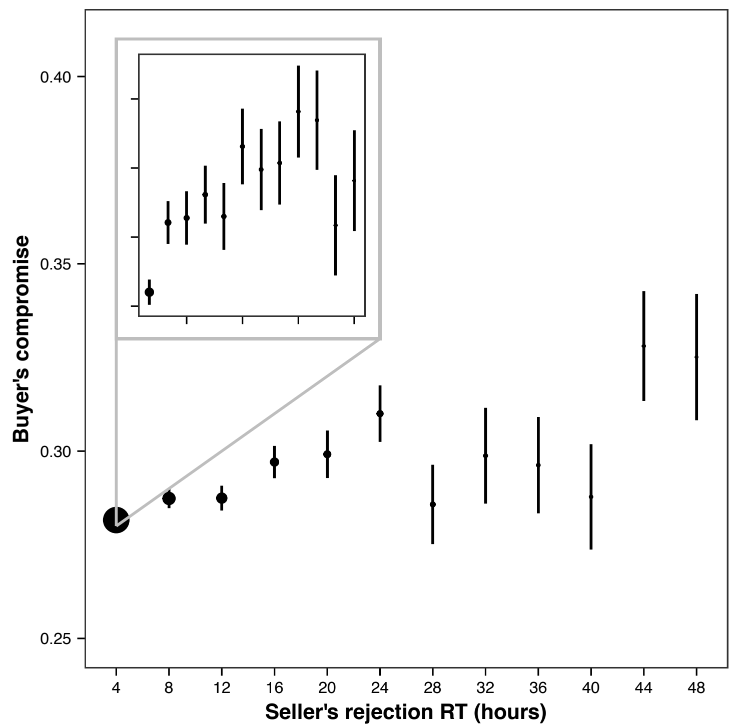

Figure S17. Buyers' compromise depending on sellers' RT. Size of buyers' compromise as a function of sellers' rejection RT to the first offers. The buyer's compromise is the amount that they raised their second offer, divided by the gap between the list price and buyer's first offer. A 100% compromise would be a second offer that is the list price; a 0% compromise would be a second offer that is the same as the first offer. In the inset zoom, the buyer's compromise ranges from $[0.272, 0.298]$ The size of the dots indicates the relative amount of data in that bin and the bars represent bootstrapped standard errors across buyers.

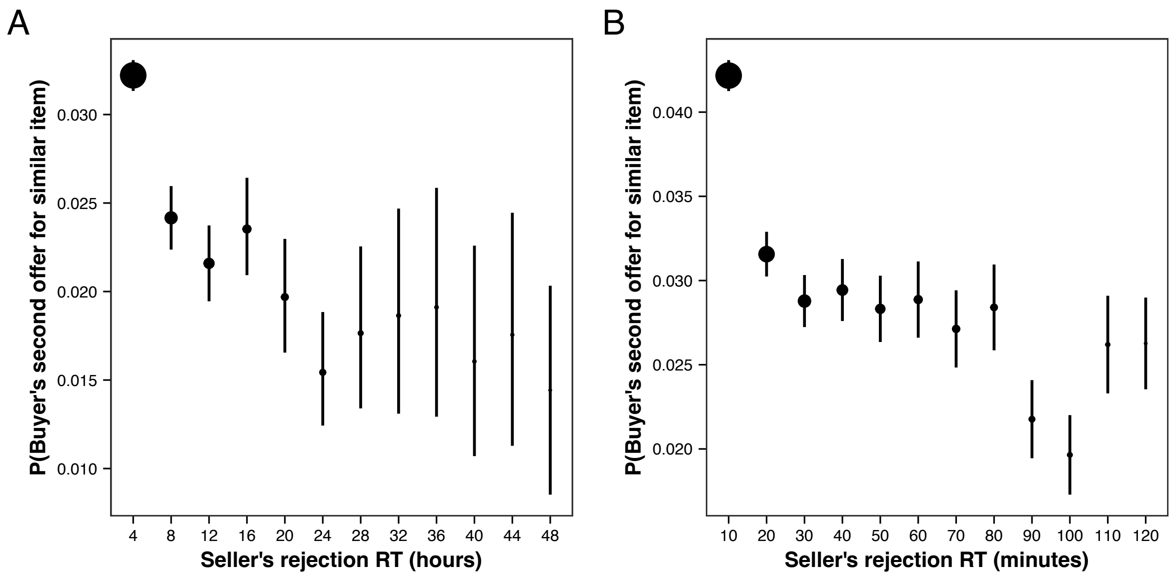

Figure S18. The probability of a buyer making an offer to a different seller on a similar item (in the same category) within 24 hours of their first offer being rejected, as a function of the seller's rejection RT. The faster the seller rejects the buyer's first offer, the more likely they are to turn to another seller. (A) Entire range of seller rejections RT (hours) (B) Zooming in on seller rejection RT (minutes) between 0-2 hours.

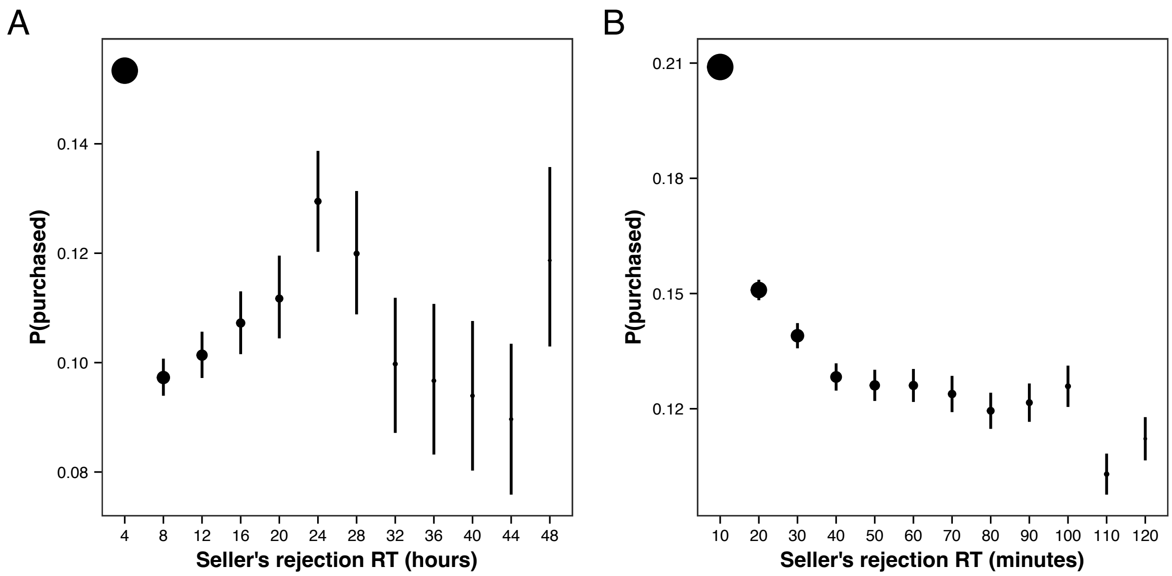

Figure S19. The probability of a buyer purchasing an item at the list price as a function of the seller's rejection RT to their first offer. (A) Entire range of seller rejections RT (hours) (B) Zooming in on the seller rejection RT (minutes) between 0-2 hours.

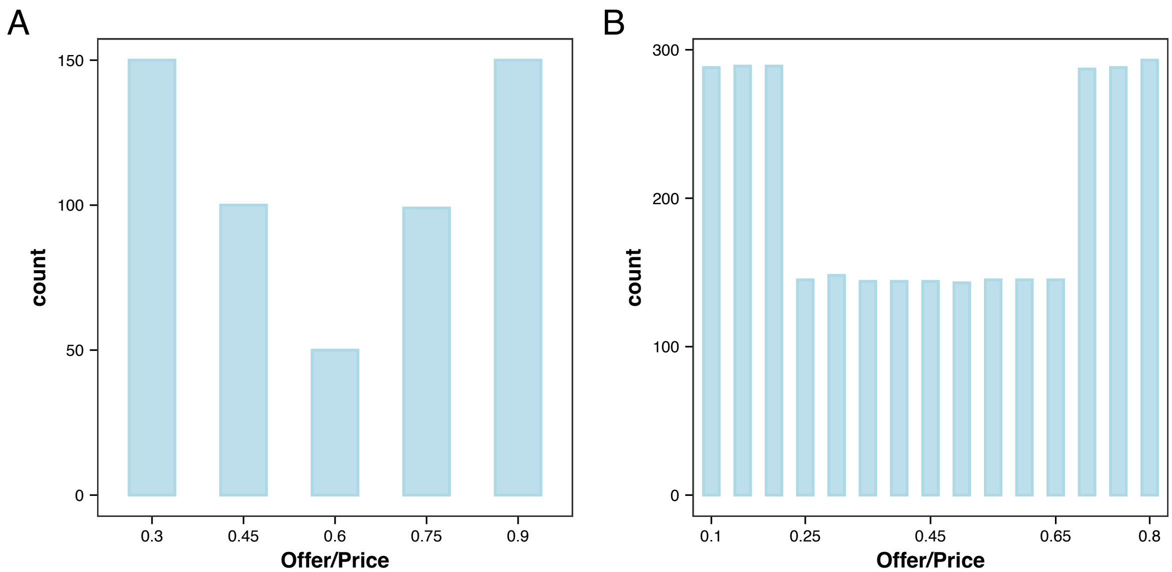

Figure S20. Histograms for offer ratios made in the eBay field experiments. First offer as a percent of seller's list price. (A) Experiment 1. (B) Experiment 2.

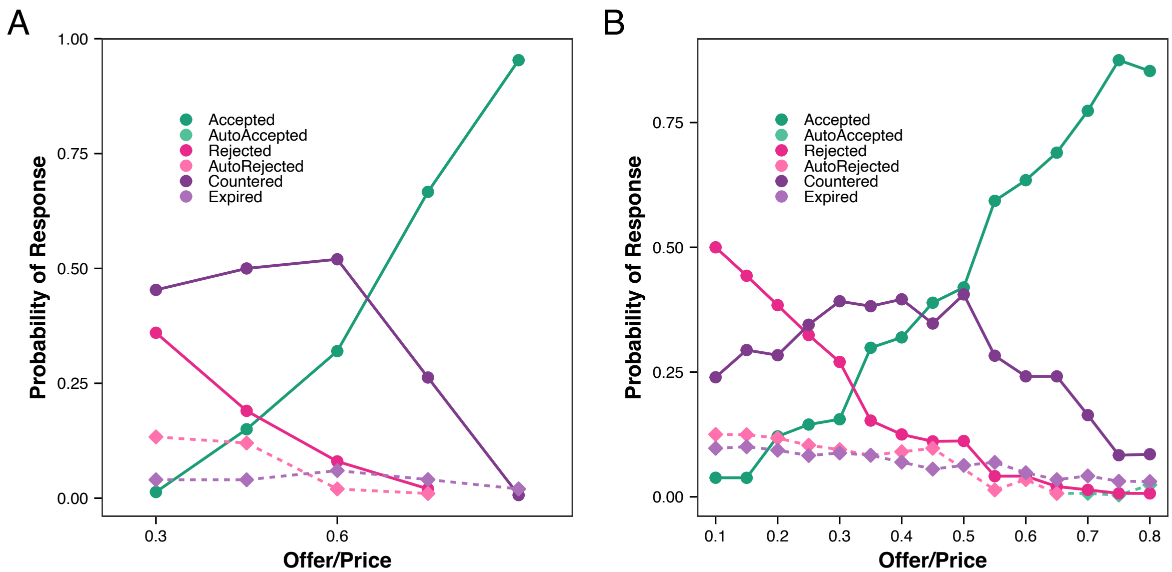

Figure S21. Probability of each type of seller response for the eBay field experiments. (A) Experiment 1. (B) Experiment 2.

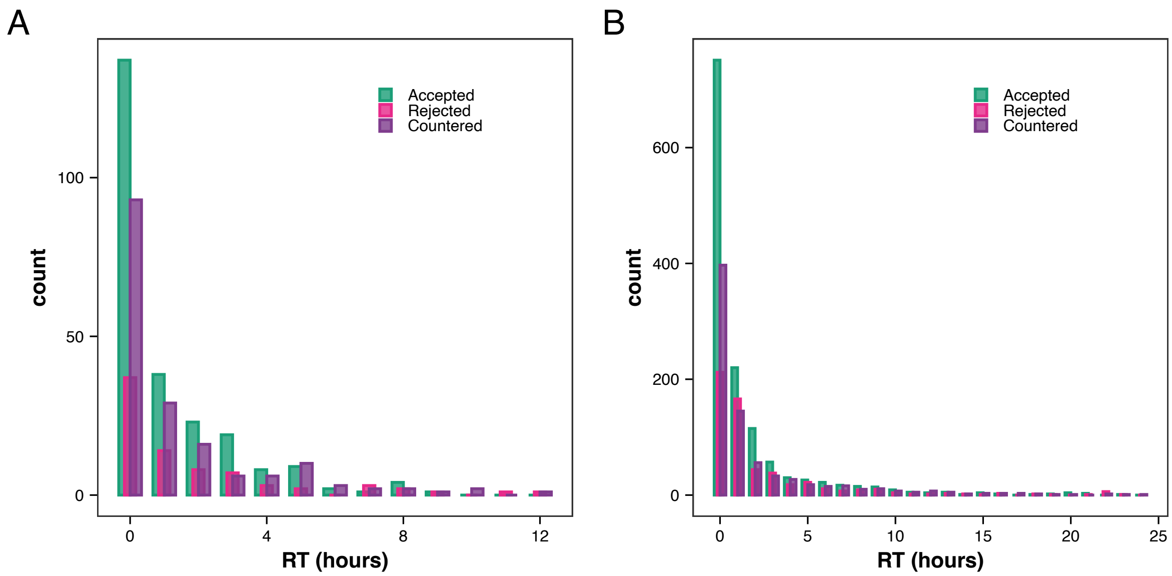

Figure S22. Response time distributions for seller's response conditional on the type of response for the eBay field experiments. (A) Experiment 1. (B) Experiment 2.

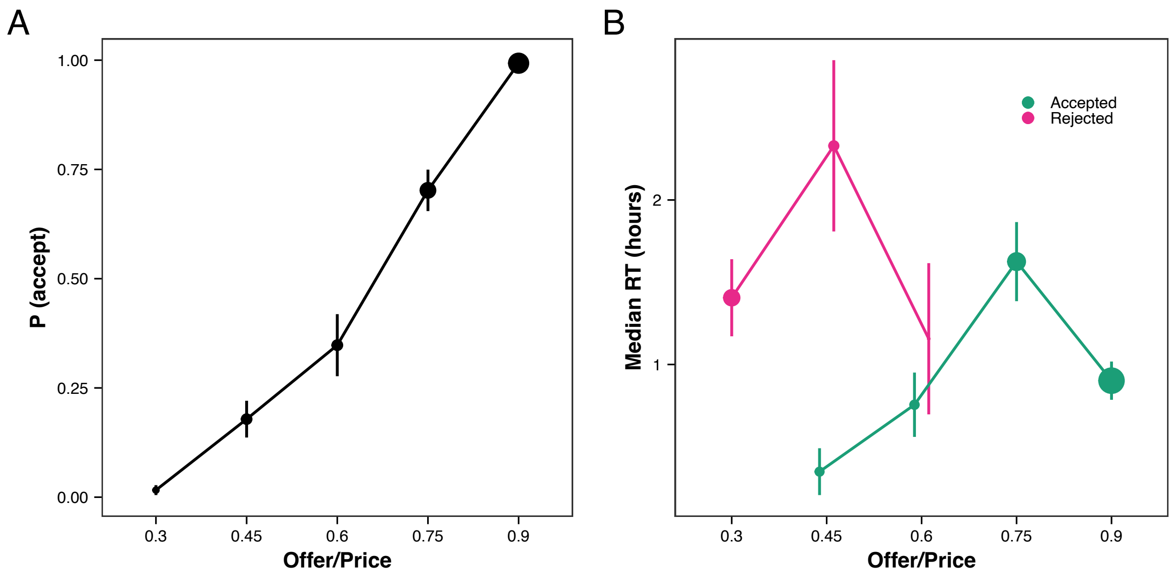

Figure S23. Choice and median RT as a function of buyer's first offer for the eBay field experiment 1. (A) Probability of seller's response by offer ratio conditional on the type of response. (B) Sellers' median RT (in hours) as a function of of buyers' initial offers, as a fraction of the sellers' list prices, conditional on the seller accepting or rejecting the offers. Only bin sizes with more than 2 data points were included. The size of the dots indicates the relative amount of data in that bin, across both curves, and the bars represent bootstrapped standard errors across sellers.

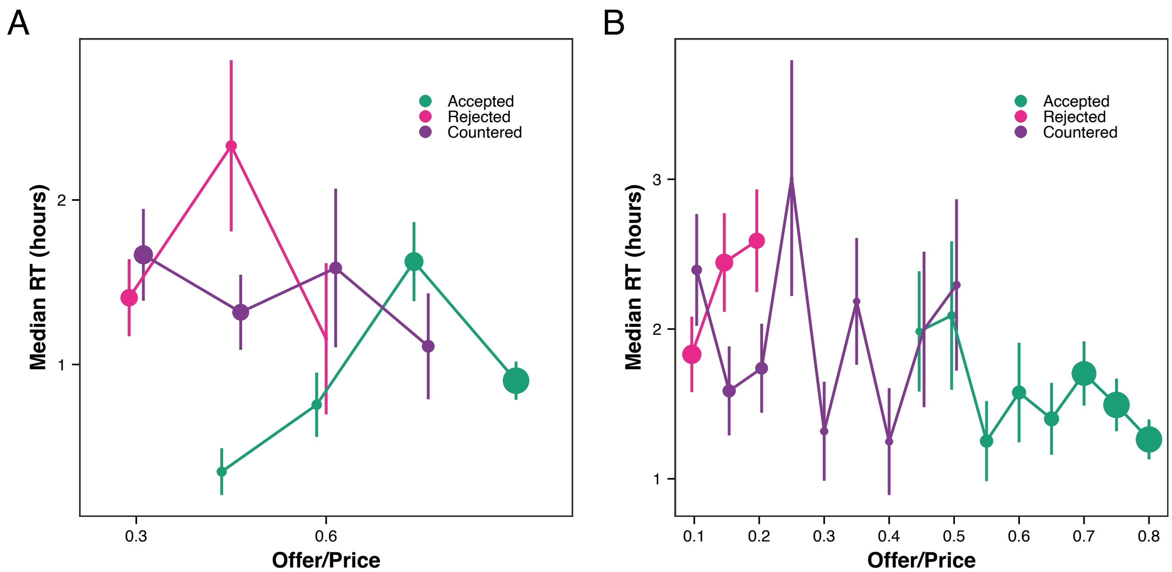

Figure S24. RT (hours) as a function of offer ratio. (A) Experiment 1. (B) Experiment 2.

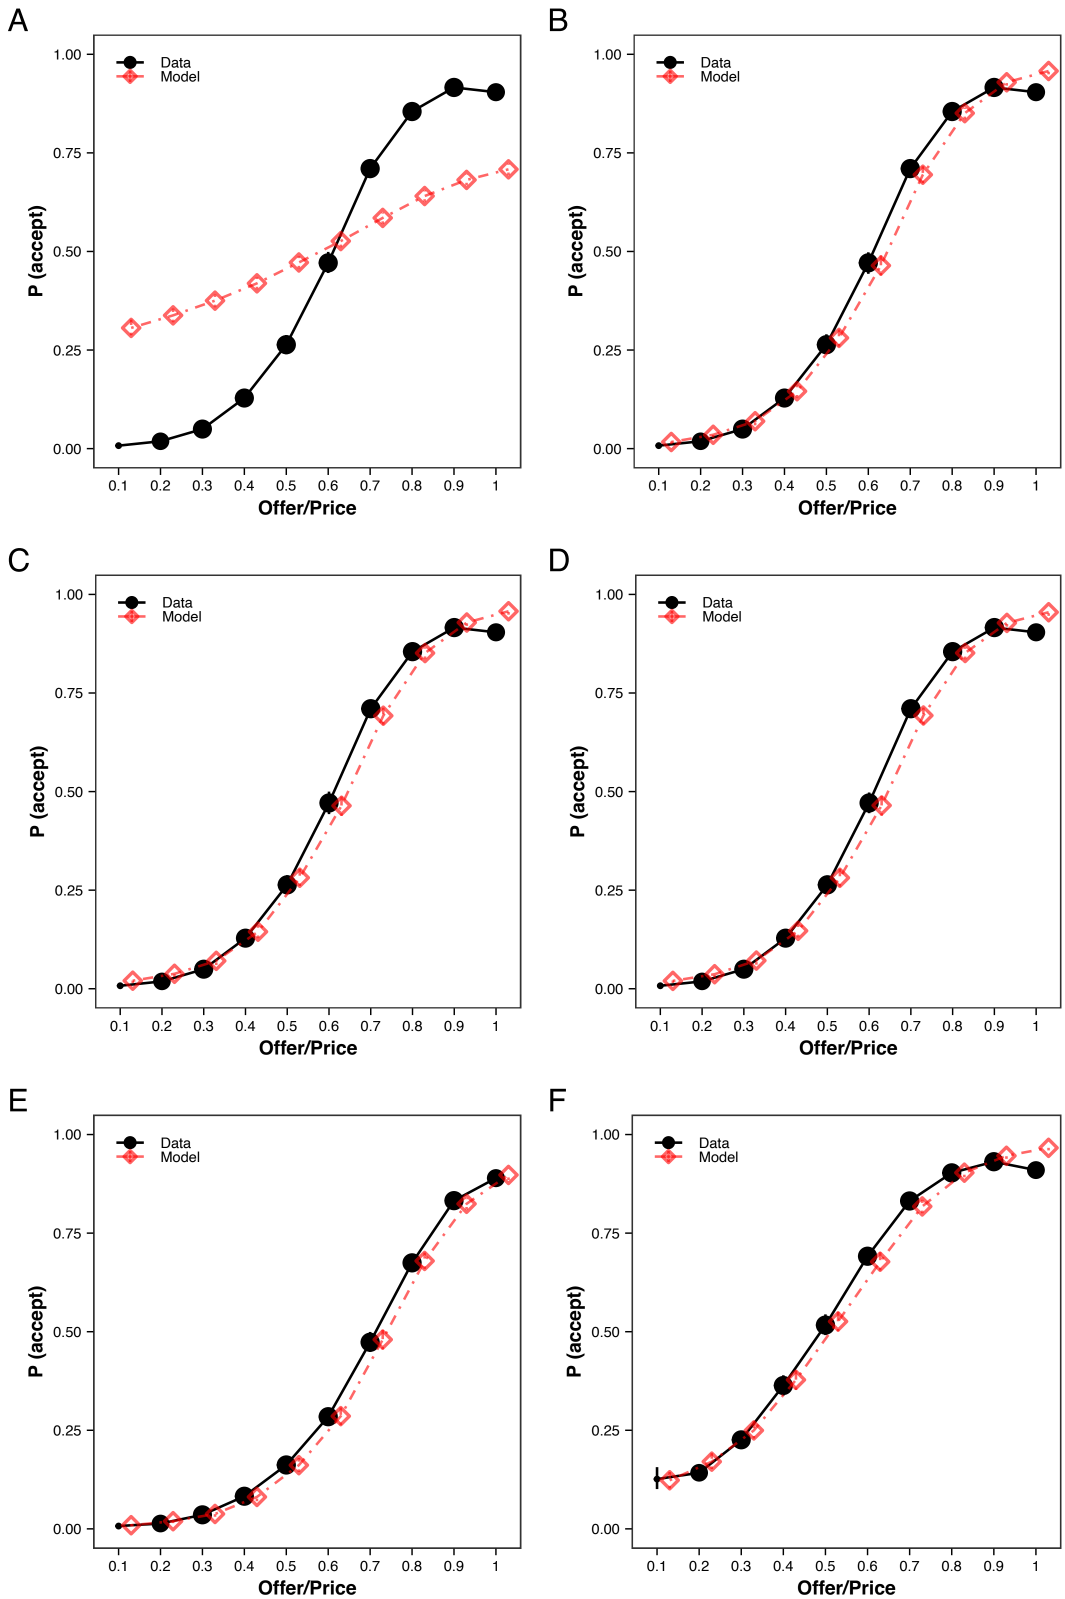
A
B
C
D
E
F
Data
Model
P (accept)
Offer/Price
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Figure S25. Probability of acceptance for data versus model predictions for eBay observational data. Sellers' probability of accepting the first offer as a function of the buyer's first offer as a fraction of seller's list price. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances.

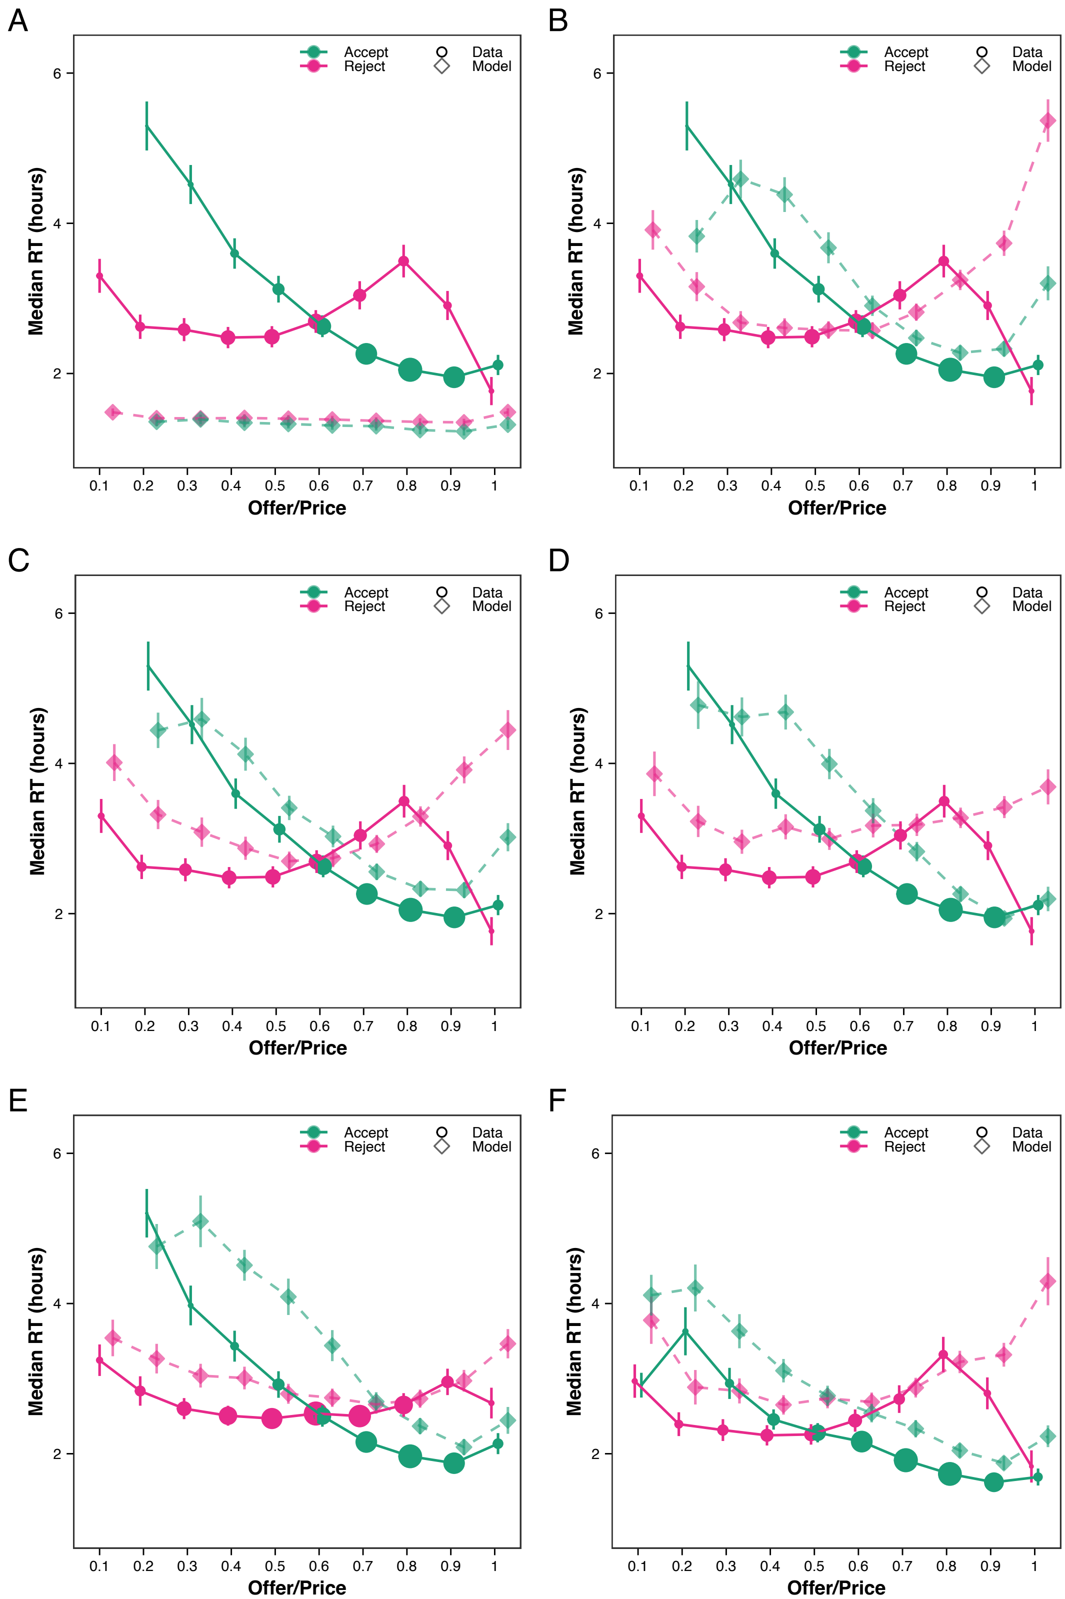

A
B
C
D
E
F
Accept
Reject
Data
Model
Median RT (hours)
Offer/Price

Figure S26. Median RT for data versus model predictions for eBay observational data. Seller's median RT (in hours) as a function of the buyer's first offer as a fraction of seller's list price, in the data and DDM fits. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. Bins with less than 12 observations were excluded. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances. Bins with less than 100 observations were excluded.

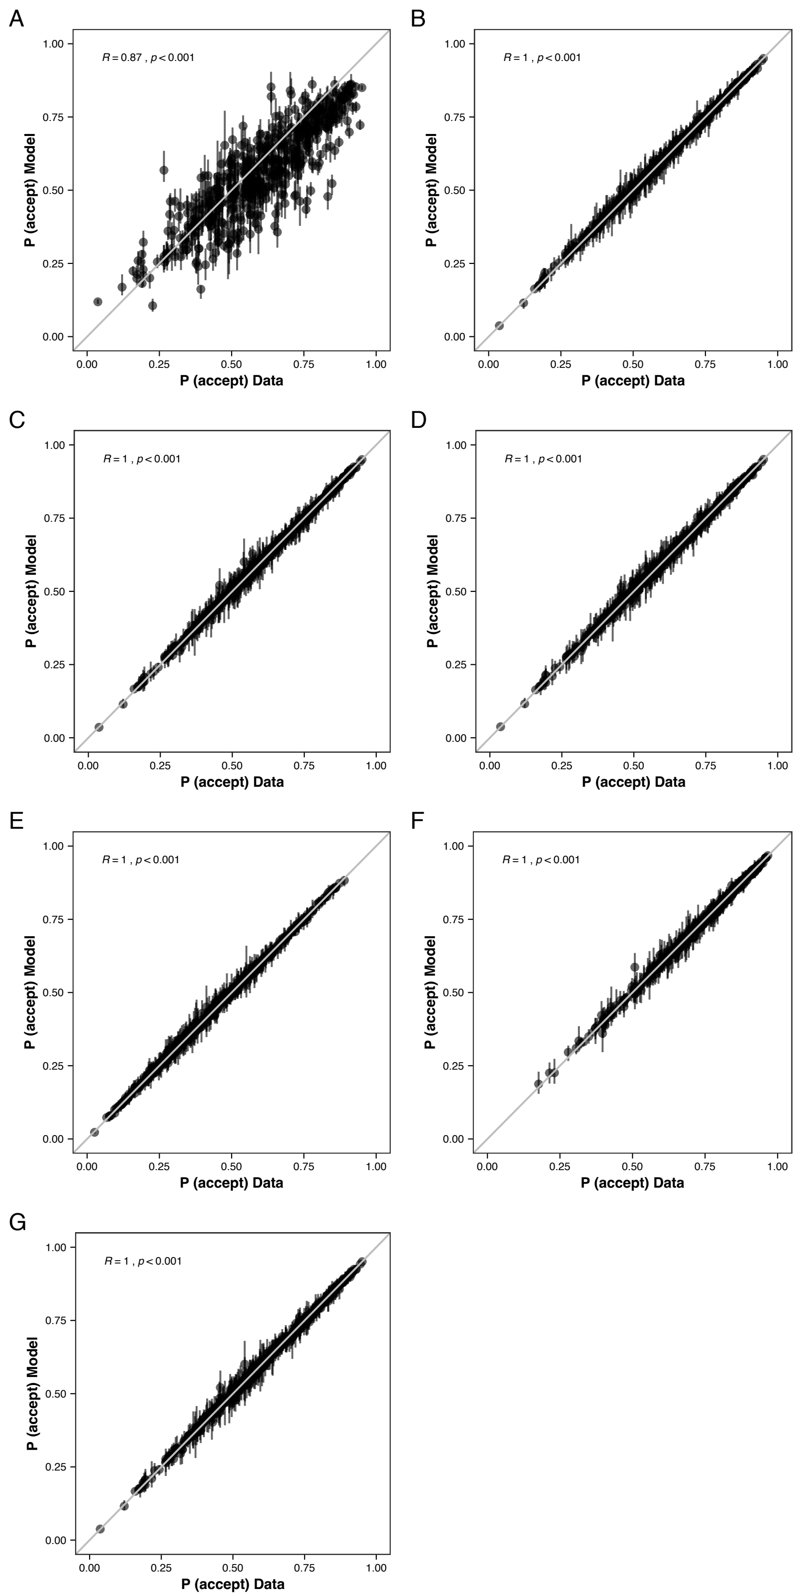
A
R = 0.87 , p < 0.001
B
R = 1 , p < 0.001
C
R = 1 , p < 0.001
D
R = 1 , p < 0.001
E
R = 1 , p < 0.001
F
R = 1 , p < 0.001
G
R = 1 , p < 0.001
P (accept) Model
P (accept) Data
0.00
0.25
0.50
0.75
1.00

Figure S27. Probability of acceptance at the subject level for data versus model predictions for eBay observational data. The bars represent 95% HDIs for 10 simulations per trial using the mean posterior of the best fitting subject level parameters. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances. (G) Best fitting model for each seller.

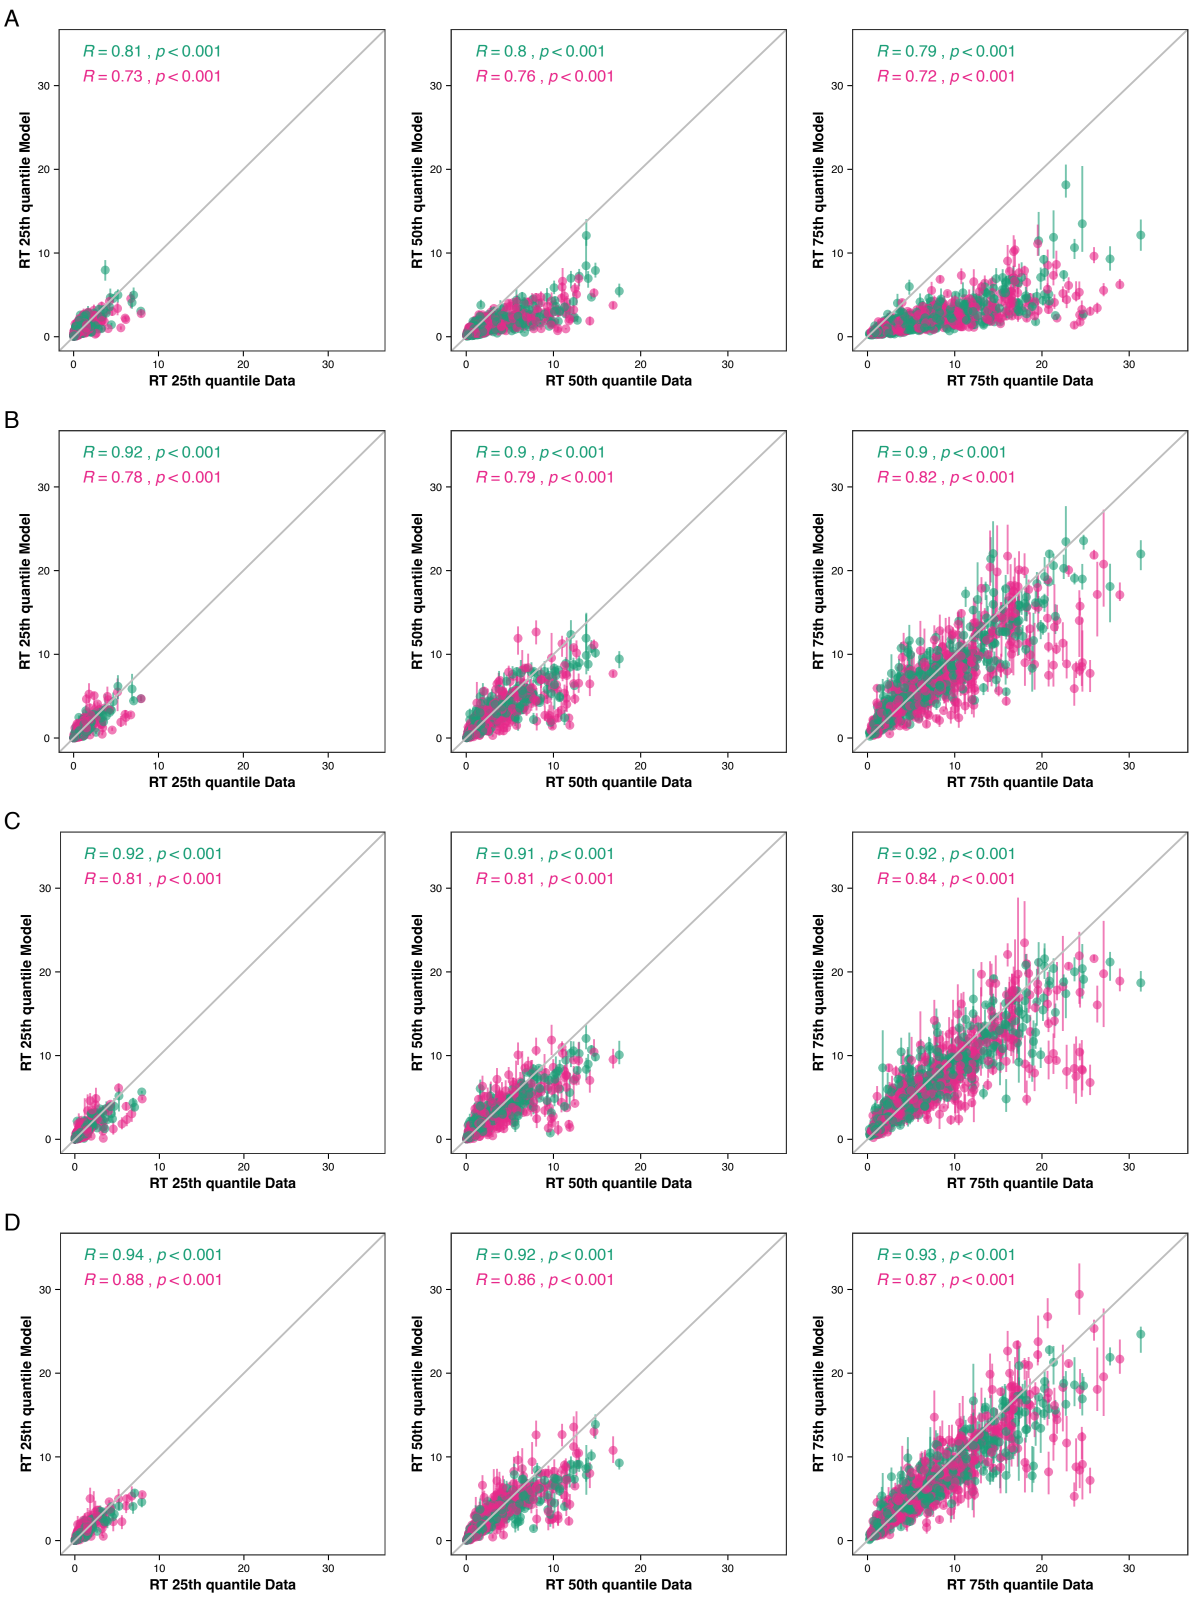
A
R = 0.81 , p < 0.001
R = 0.73 , p < 0.001
RT 25th quantile Model
RT 25th quantile Data
R = 0.8 , p < 0.001
R = 0.76 , p < 0.001
RT 50th quantile Model
RT 50th quantile Data
R = 0.79 , p < 0.001
R = 0.72 , p < 0.001
RT 75th quantile Model
RT 75th quantile Data
B
R = 0.92 , p < 0.001
R = 0.78 , p < 0.001
RT 25th quantile Model
RT 25th quantile Data
R = 0.9 , p < 0.001
R = 0.79 , p < 0.001
RT 50th quantile Model
RT 50th quantile Data
R = 0.9 , p < 0.001
R = 0.82 , p < 0.001
RT 75th quantile Model
RT 75th quantile Data
C
R = 0.92 , p < 0.001
R = 0.81 , p < 0.001
RT 25th quantile Model
RT 25th quantile Data
R = 0.91 , p < 0.001
R = 0.81 , p < 0.001
RT 50th quantile Model
RT 50th quantile Data
R = 0.92 , p < 0.001
R = 0.84 , p < 0.001
RT 75th quantile Model
RT 75th quantile Data
D
R = 0.94 , p < 0.001
R = 0.88 , p < 0.001
RT 25th quantile Model
RT 25th quantile Data
R = 0.92 , p < 0.001
R = 0.86 , p < 0.001
RT 50th quantile Model
RT 50th quantile Data
R = 0.93 , p < 0.001
R = 0.87 , p < 0.001
RT 75th quantile Model
RT 75th quantile Data

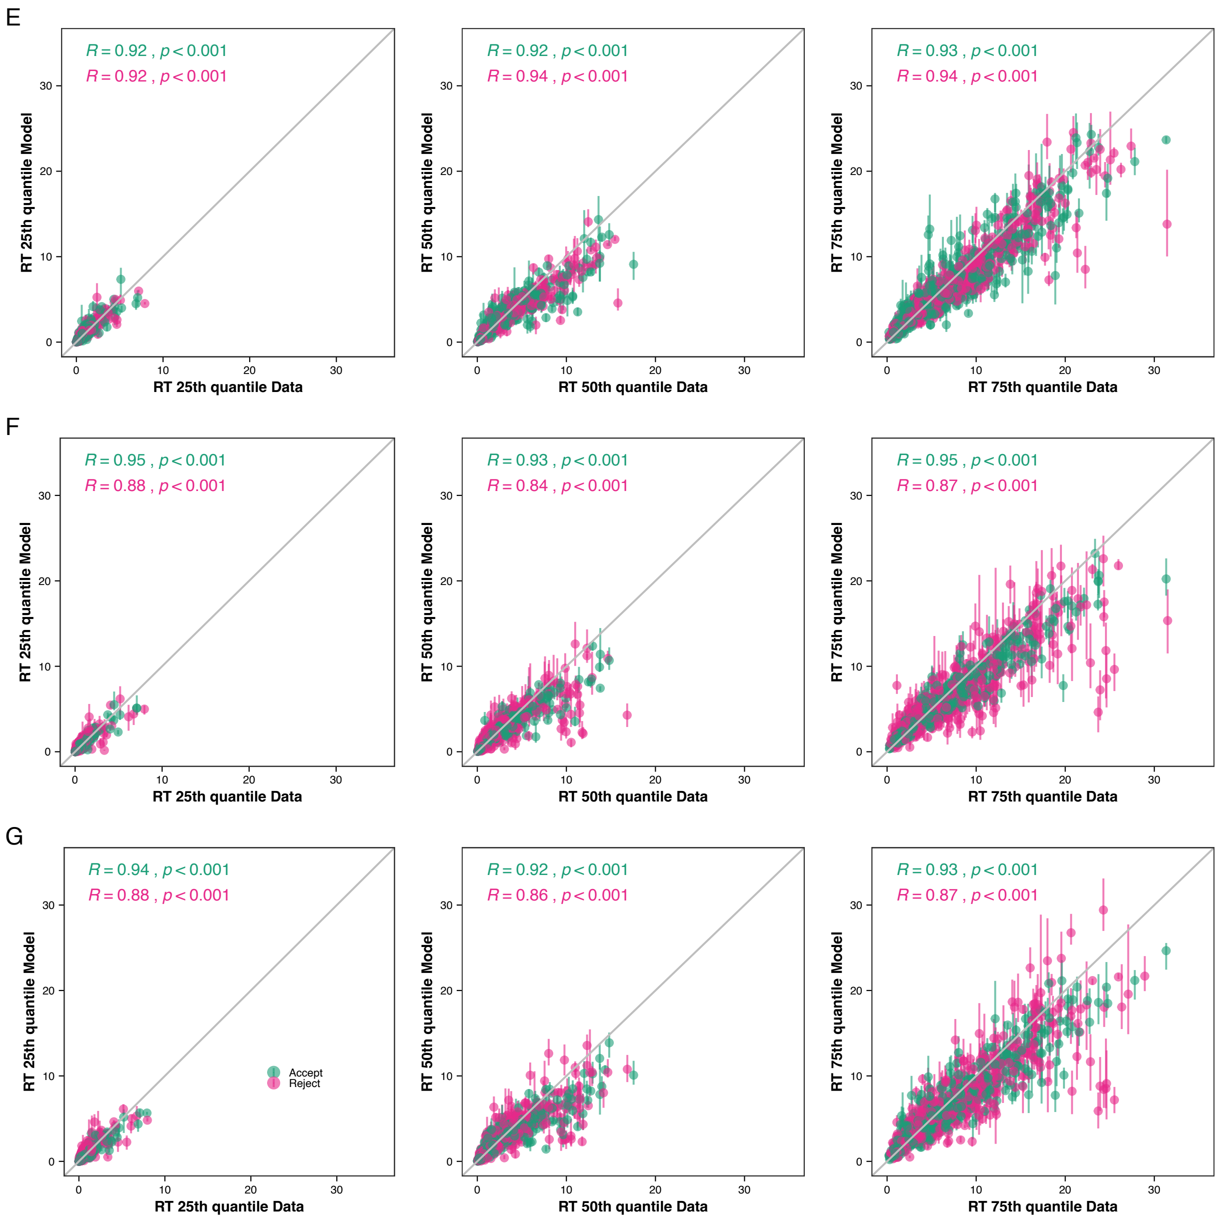

Figure S28. Response time quartiles at the subject level for data versus model predictions for acceptances and rejections for eBay observational data. The bars represent 95% HDIs for 10 simulations per trial using the mean posterior of the best fitting subject level parameters. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances. (G) Best fitting model for each seller.

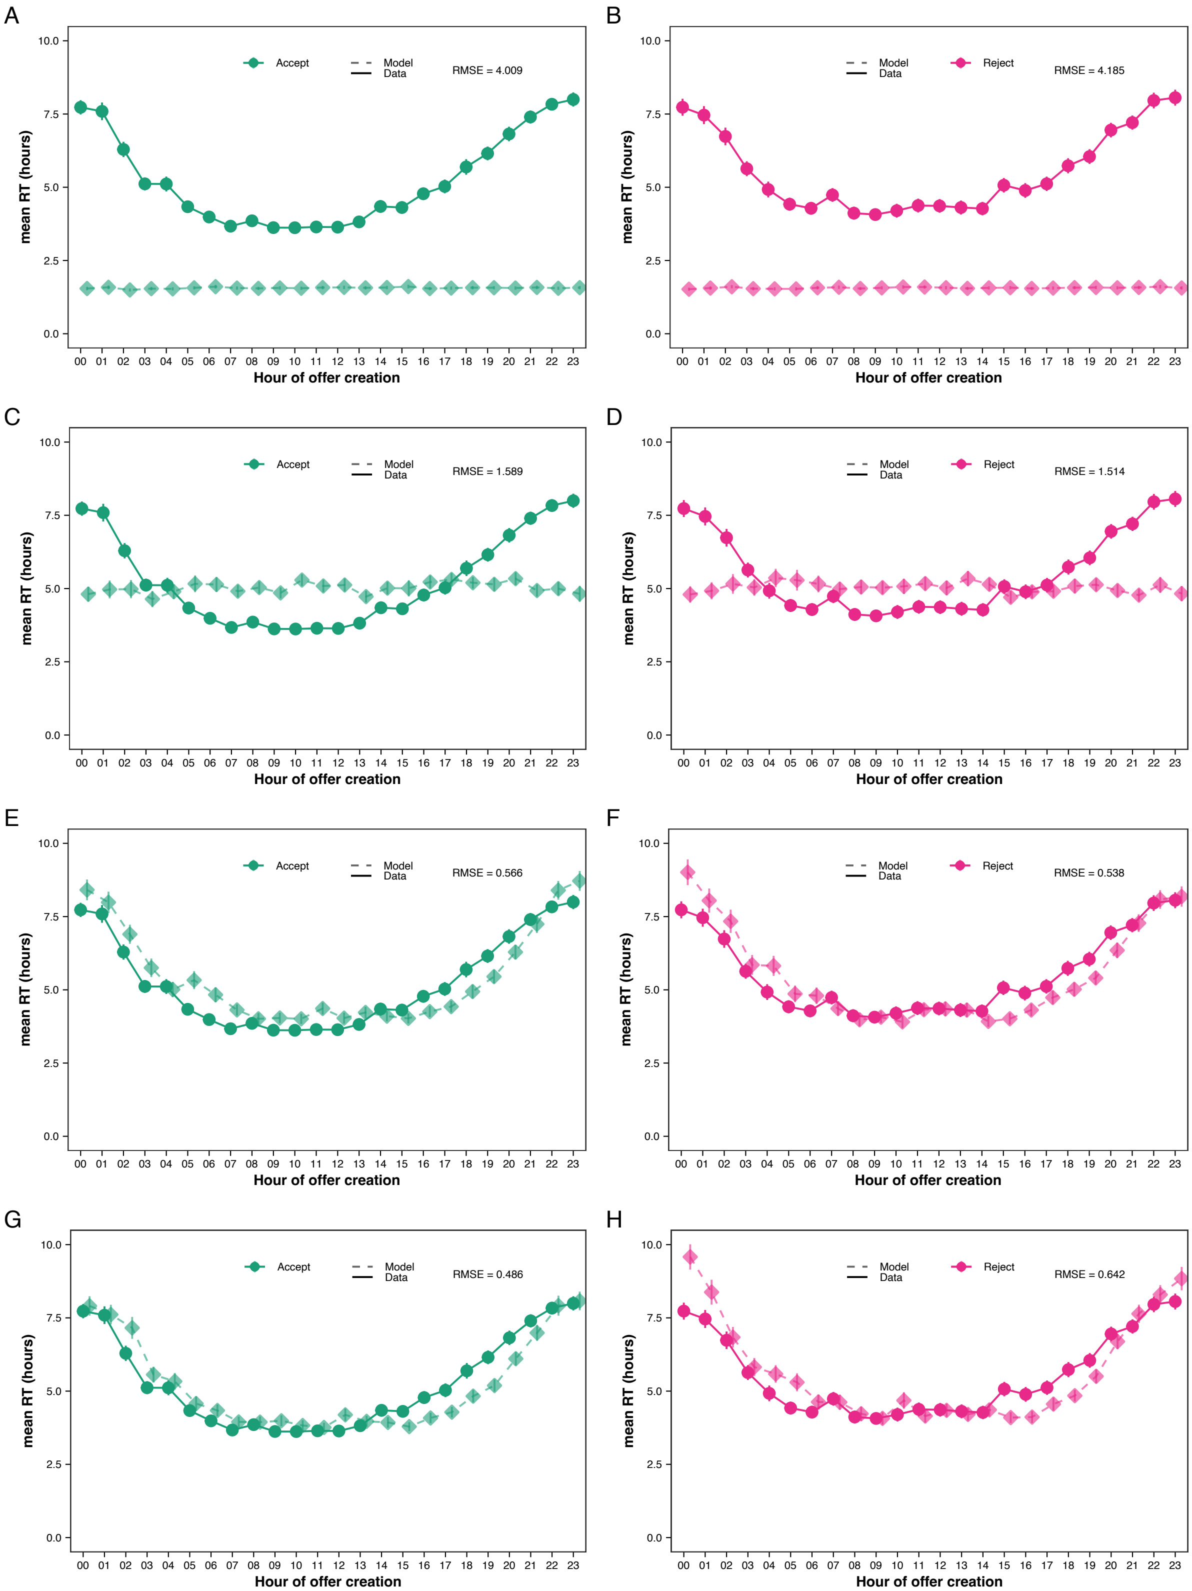
A
Accept
Model
Data
RMSE = 4.009
B
Reject
RMSE = 4.185
C
RMSE = 1.589
D
RMSE = 1.514
E
RMSE = 0.566
F
RMSE = 0.538
G
RMSE = 0.486
H
RMSE = 0.642
mean RT (hours)
Hour of offer creation
0.0
2.5
5.0
7.5
10.0
00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23

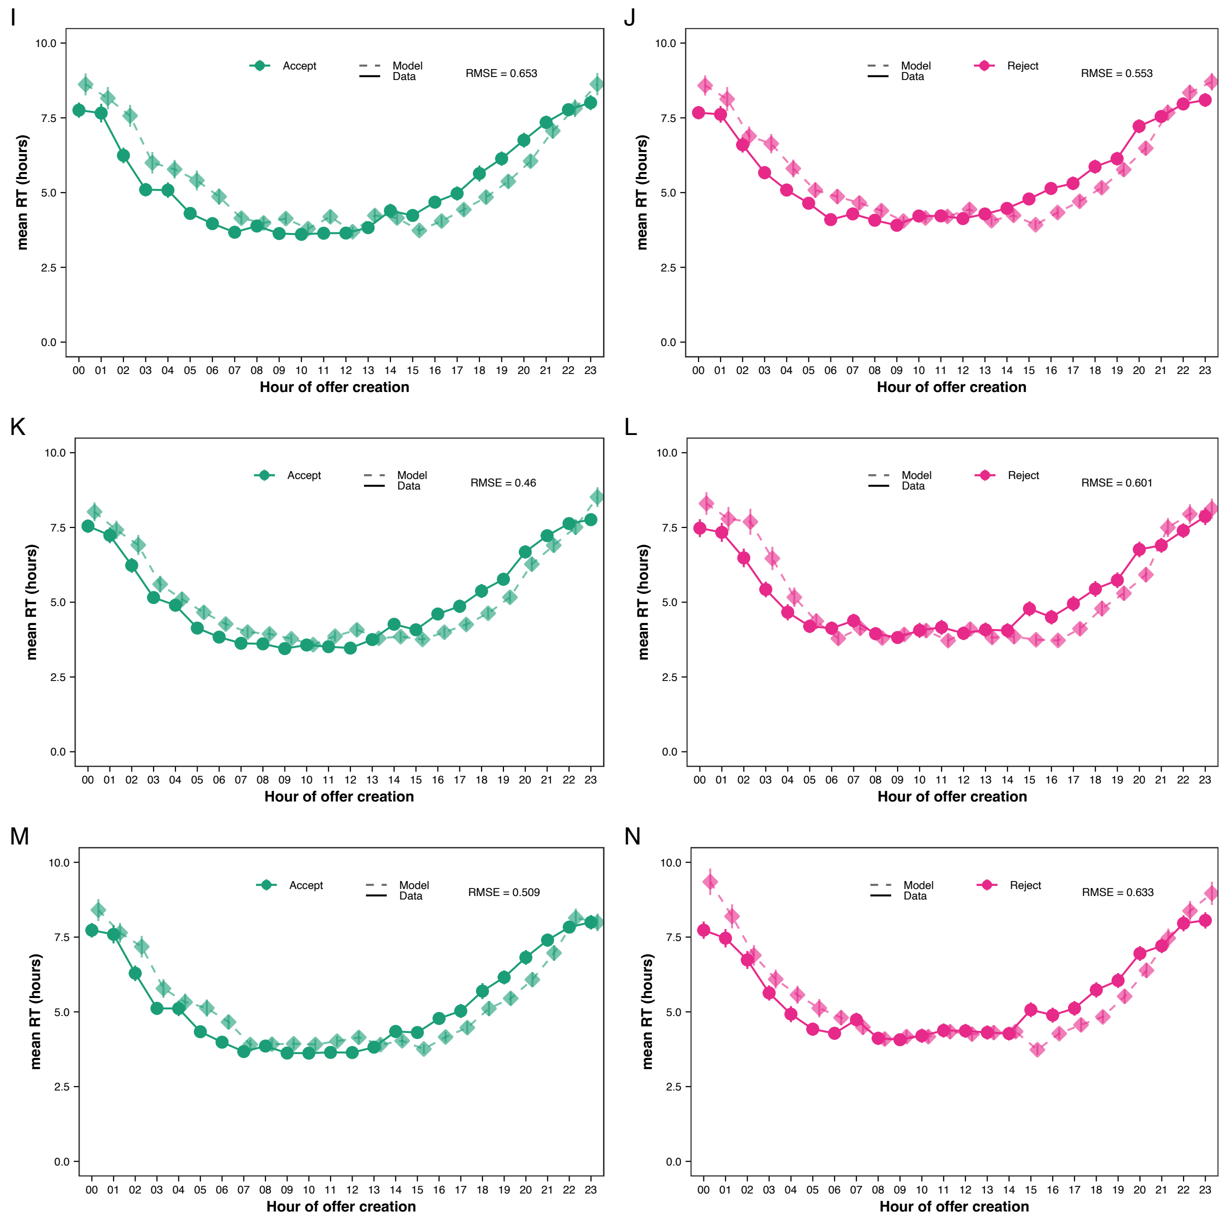

Figure S29. Mean Response time in hours as a function of hour in the day when the offer was created starting at midnight (PT) conditional on acceptance or rejection for the data and the model predictions for eBay observational data. The bars represent 95% HDIs for 10 simulations per trial using the mean posterior of the best fitting subject level parameters. (A), (B) Standard DDM. (C), (D) Gamma DDM. (E), (F) Time of Day Gamma DDM. (G), (H) Time of Day and Offer Ratio Gamma DDM. (I), (J) Best fitting model for each seller pooling counteroffers with rejections. (K), (L) Best fitting model for each seller pooling counteroffers with acceptances. (M), (N) Best fitting model for each seller.

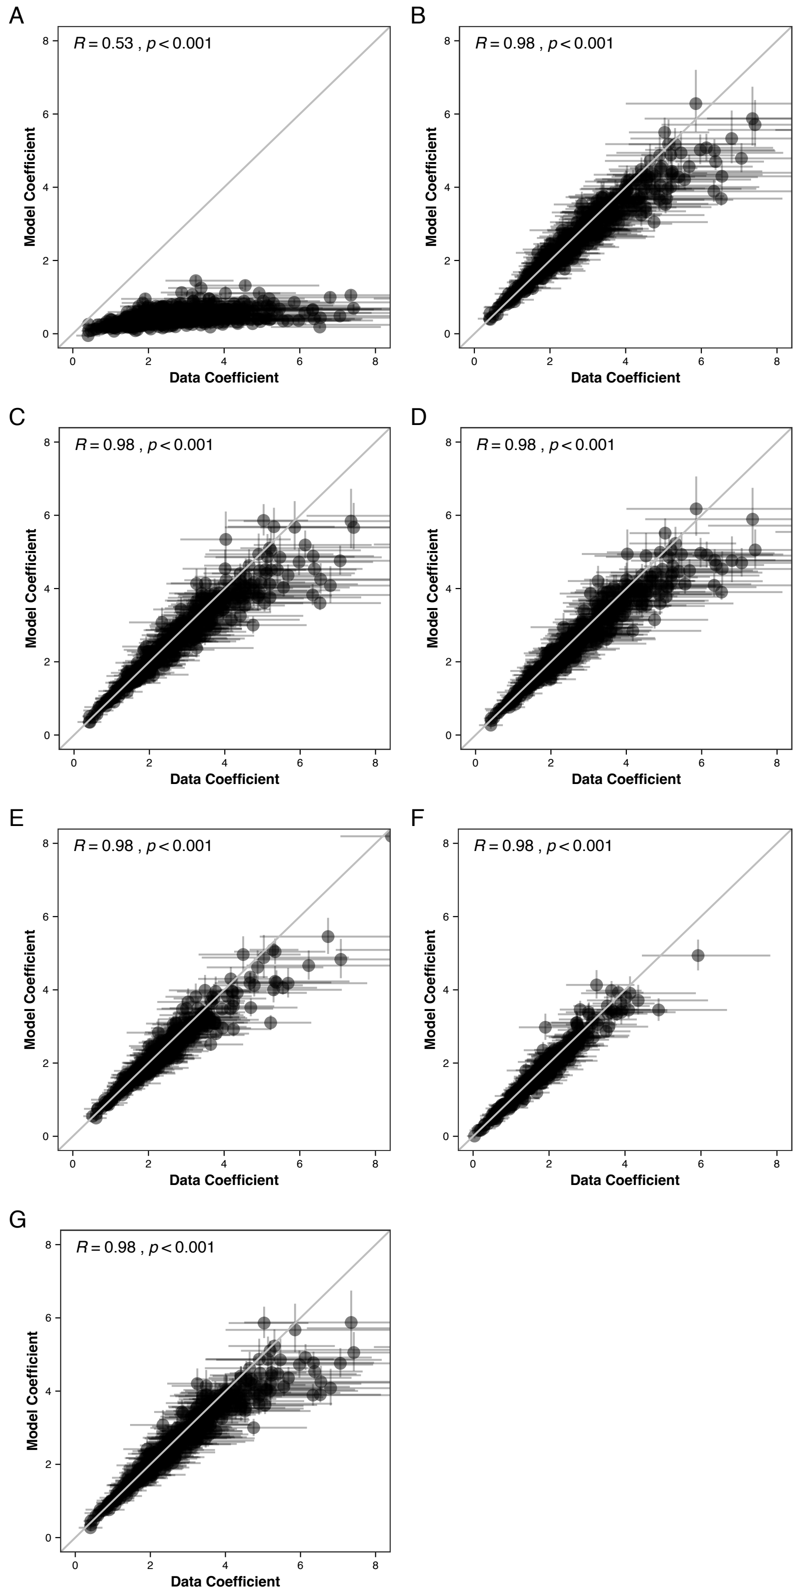
A
R = 0.53 , p < 0.001
B
R = 0.98 , p < 0.001
C
R = 0.98 , p < 0.001
D
R = 0.98 , p < 0.001
E
R = 0.98 , p < 0.001
F
R = 0.98 , p < 0.001
G
R = 0.98 , p < 0.001
Model Coefficient
Data Coefficient
0
2
4
6
8

Figure S30. Model versus data coefficients of acceptance on first offer ratio at the seller level for eBay observational data. The coefficients are from logistic regressions at the seller level using acceptance dummy as dependent variables and z-scored first offer ratio and list price as independent variables. The bars represent 95% confidence intervals. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances. (G) Best fitting model for each seller.

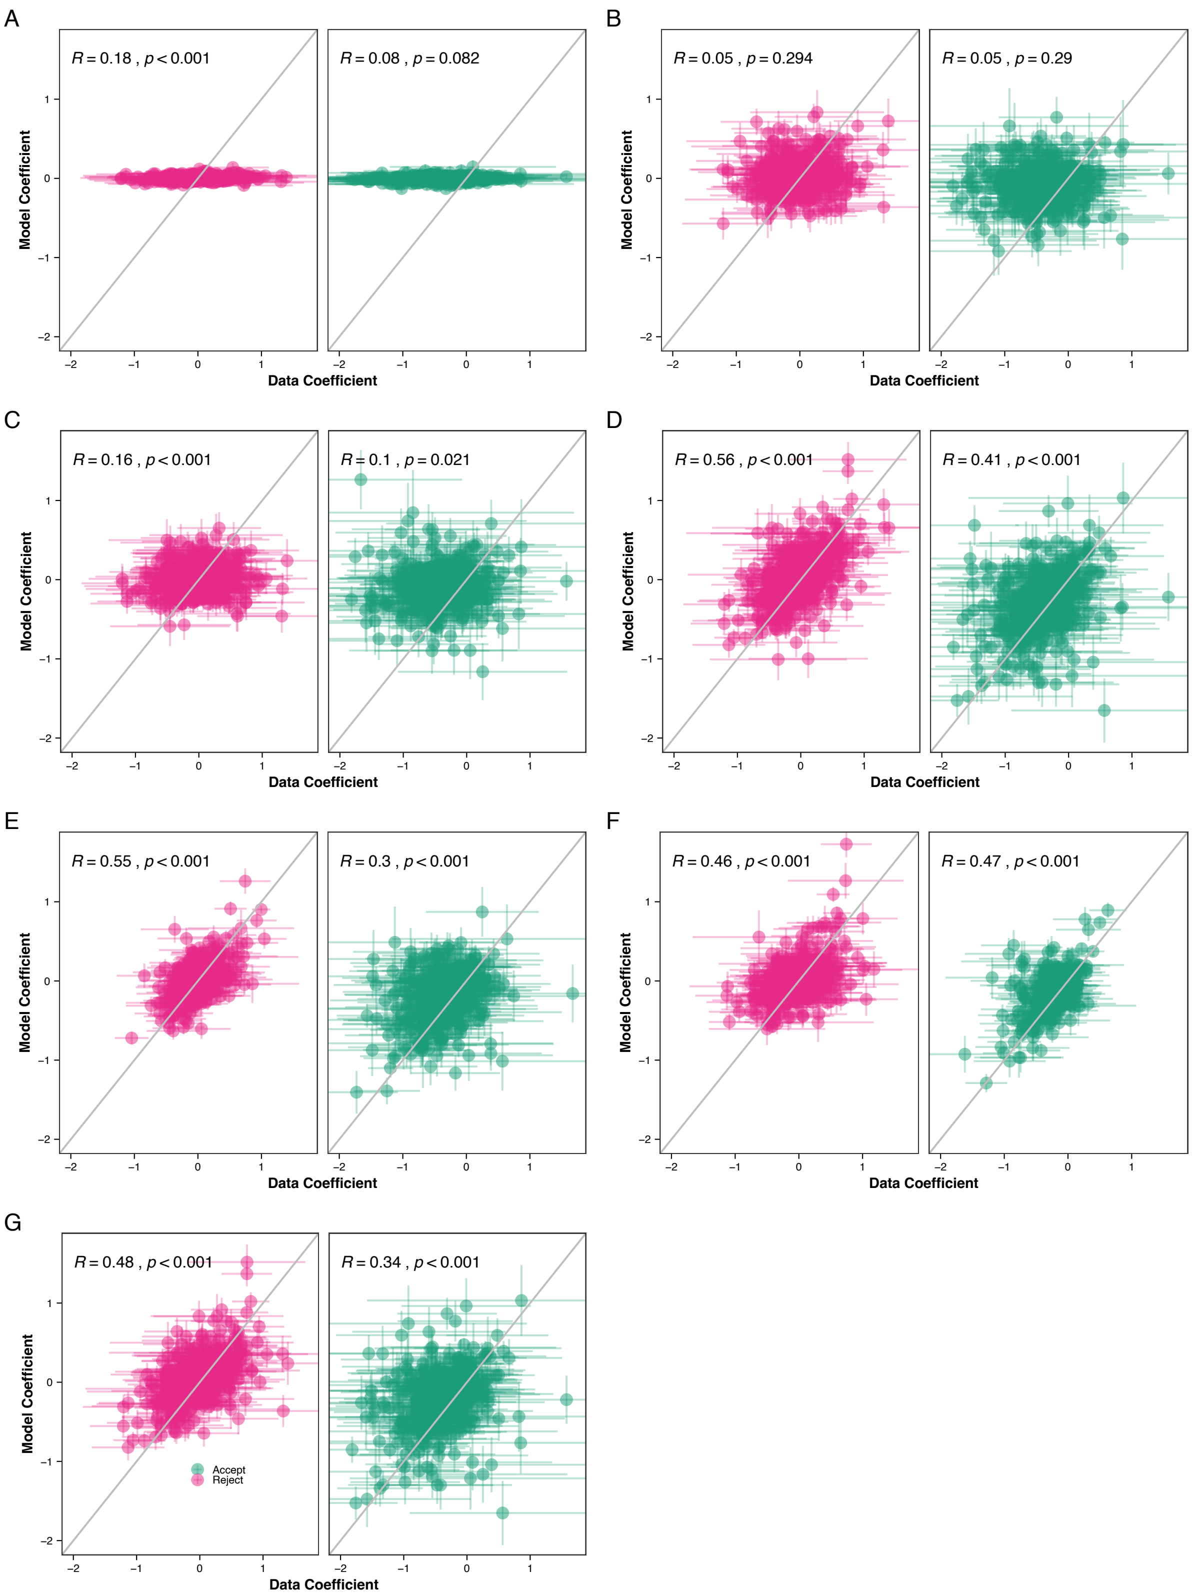
A
R = 0.18 , p < 0.001
R = 0.08 , p = 0.082
B
R = 0.05 , p = 0.294
R = 0.05 , p = 0.29
C
R = 0.16 , p < 0.001
R = 0.1 , p = 0.021
D
R = 0.56 , p < 0.001
R = 0.41 , p < 0.001
E
R = 0.55 , p < 0.001
R = 0.3 , p < 0.001
F
R = 0.46 , p < 0.001
R = 0.47 , p < 0.001
G
R = 0.48 , p < 0.001
R = 0.34 , p < 0.001
Accept
Reject
Model Coefficient
Data Coefficient

Figure S31. Model versus data coefficients of logRT on first offer ratio at the seller level for eBay observational data. The coefficients are from linear regressions at the seller level using log RT in hours as dependent variables and z-scored first offer ratio and list price as independent variables for acceptances and rejections respectively. The bars represent 95% confidence intervals. (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Best fitting model for each seller pooling counteroffers with rejections. (F) Best fitting model for each seller pooling counteroffers with acceptances. (G) Best fitting model for each seller.

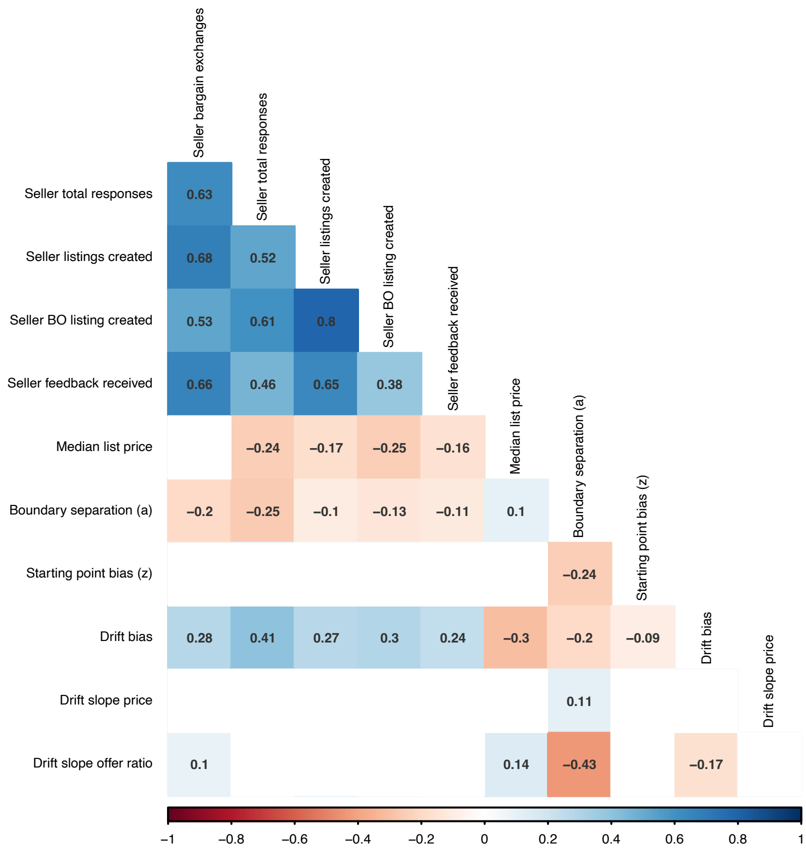

Figure S32. Correlations between mean posteriors of seller model parameters and seller experience using the best fitting model for each seller in the eBay observational data. The seller characteristics are: number of previous best offer exchanges seller has participated in, number of total seller responses (acceptances and rejections), seller's number of previous feedbacks received at the time of the offer, number of listings created by the seller dating back to 2008, number of Best-Offer-listings created by the seller dating back to 2008. Only significant correlations ($p < 0.01$) are displayed.

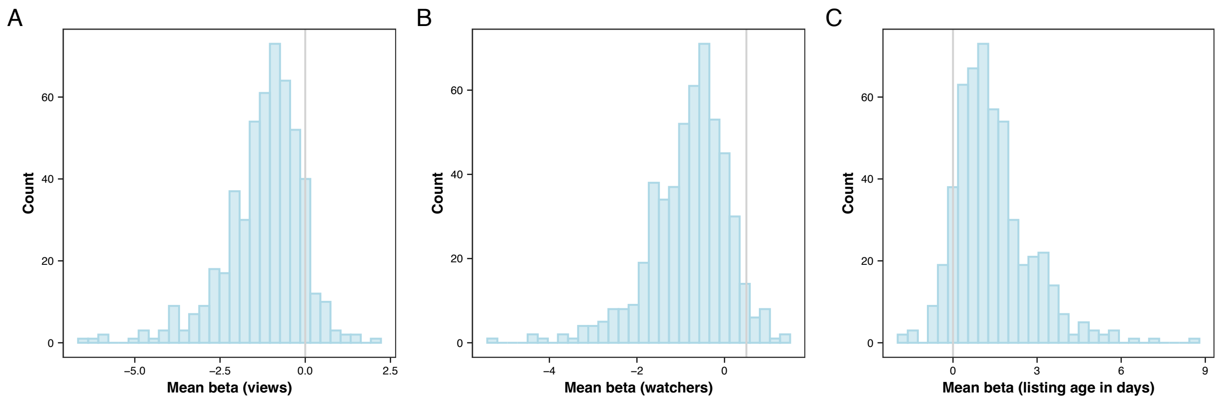

Figure S33. The effect of (A) the number of views, (B) the number of watchers, and (C) the listing age (in days), on the drift rate in Model 5. Positive effects increase the probability of accepting offers. Displayed are the distributions of mean posterior parameter estimates.

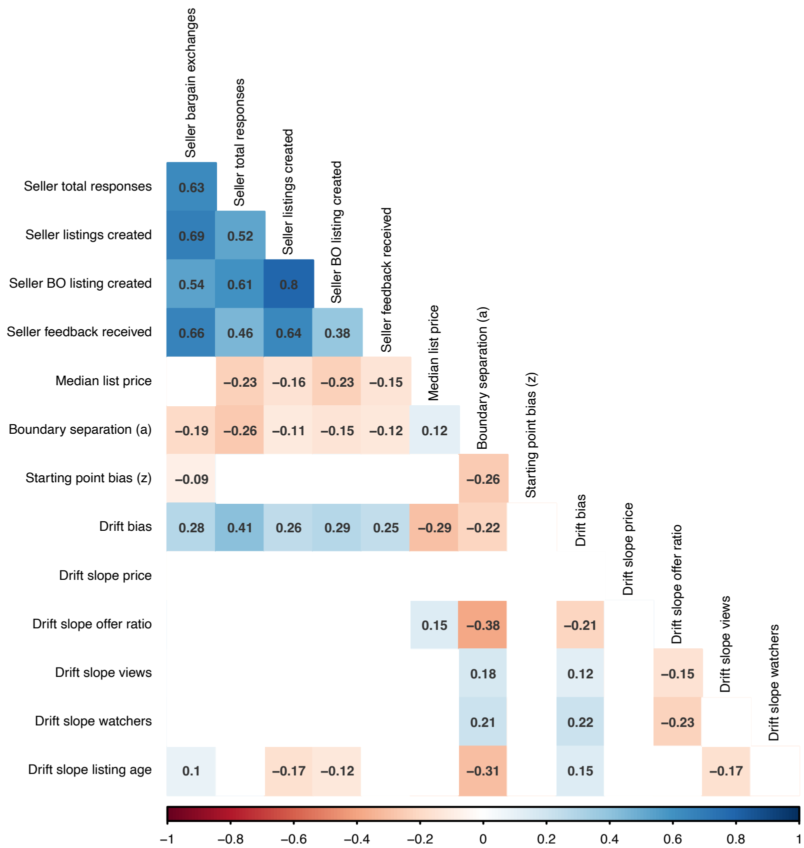

Figure S34. Correlations between mean parameter estimates from Model 5 and seller characteristics. The seller characteristics are: Seller bargain exchanges - number of previous best offer exchanges the seller has participated in; Seller total responses - seller's total number of Best-Offer responses (acceptances and rejections); Seller listings created - number of listings created by the seller dating back to 2008; Seller BO listing created - number of Best-Offer-listings created by the seller dating back to 2008; Seller feedback received - seller's amount of previous feedback received at the time of the offer. Only significant correlations ($p < 0.01$) are displayed.

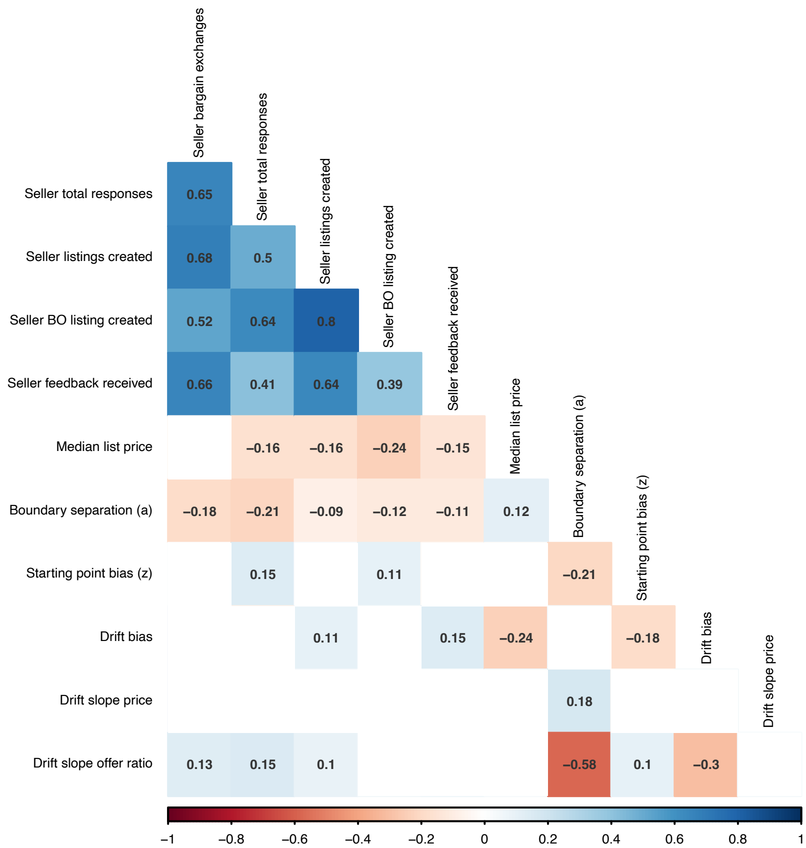

Figure S35. Correlations between mean posteriors of seller model parameters and seller experience using the best fitting model for each seller in the eBay observational data, but with counteroffers pooled with rejections. The seller characteristics are: Seller bargain exchanges - number of previous best offer exchanges the seller has participated in; Seller total responses - seller's total number of Best-Offer responses (acceptances and rejections); Seller listings created - number of listings created by the seller dating back to 2008; Seller BO listing created - number of Best-Offer-listings created by the seller dating back to 2008; Seller feedback received - seller's amount of previous feedback received at the time of the offer. Only significant correlations ($p < 0.01$) are displayed.

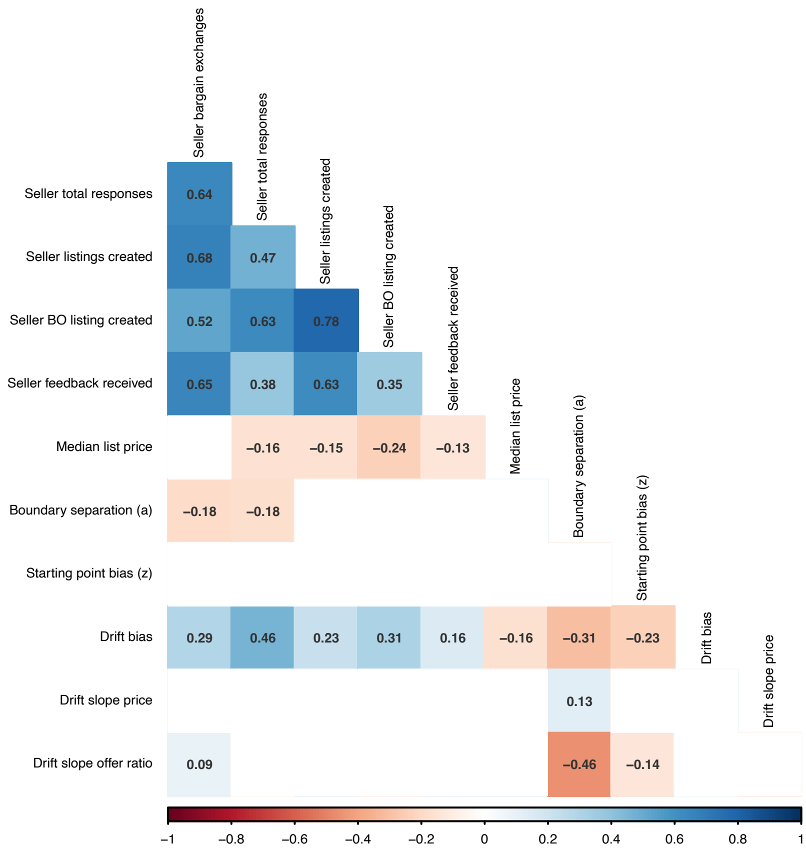

Figure S36. Correlations between mean posteriors of seller model parameters and seller experience using the best fitting model for each seller in the eBay observational data, but with counteroffers pooled with acceptances. The seller characteristics are: Seller bargain exchanges - number of previous best offer exchanges the seller has participated in; Seller total responses - seller's total number of Best-Offer responses (acceptances and rejections); Seller listings created - number of listings created by the seller dating back to 2008; Seller BO listing created - number of Best-Offer-listings created by the seller dating back to 2008; Seller feedback received - seller's amount of previous feedback received at the time of the offer. Only significant correlations ($p < 0.01$) are displayed.

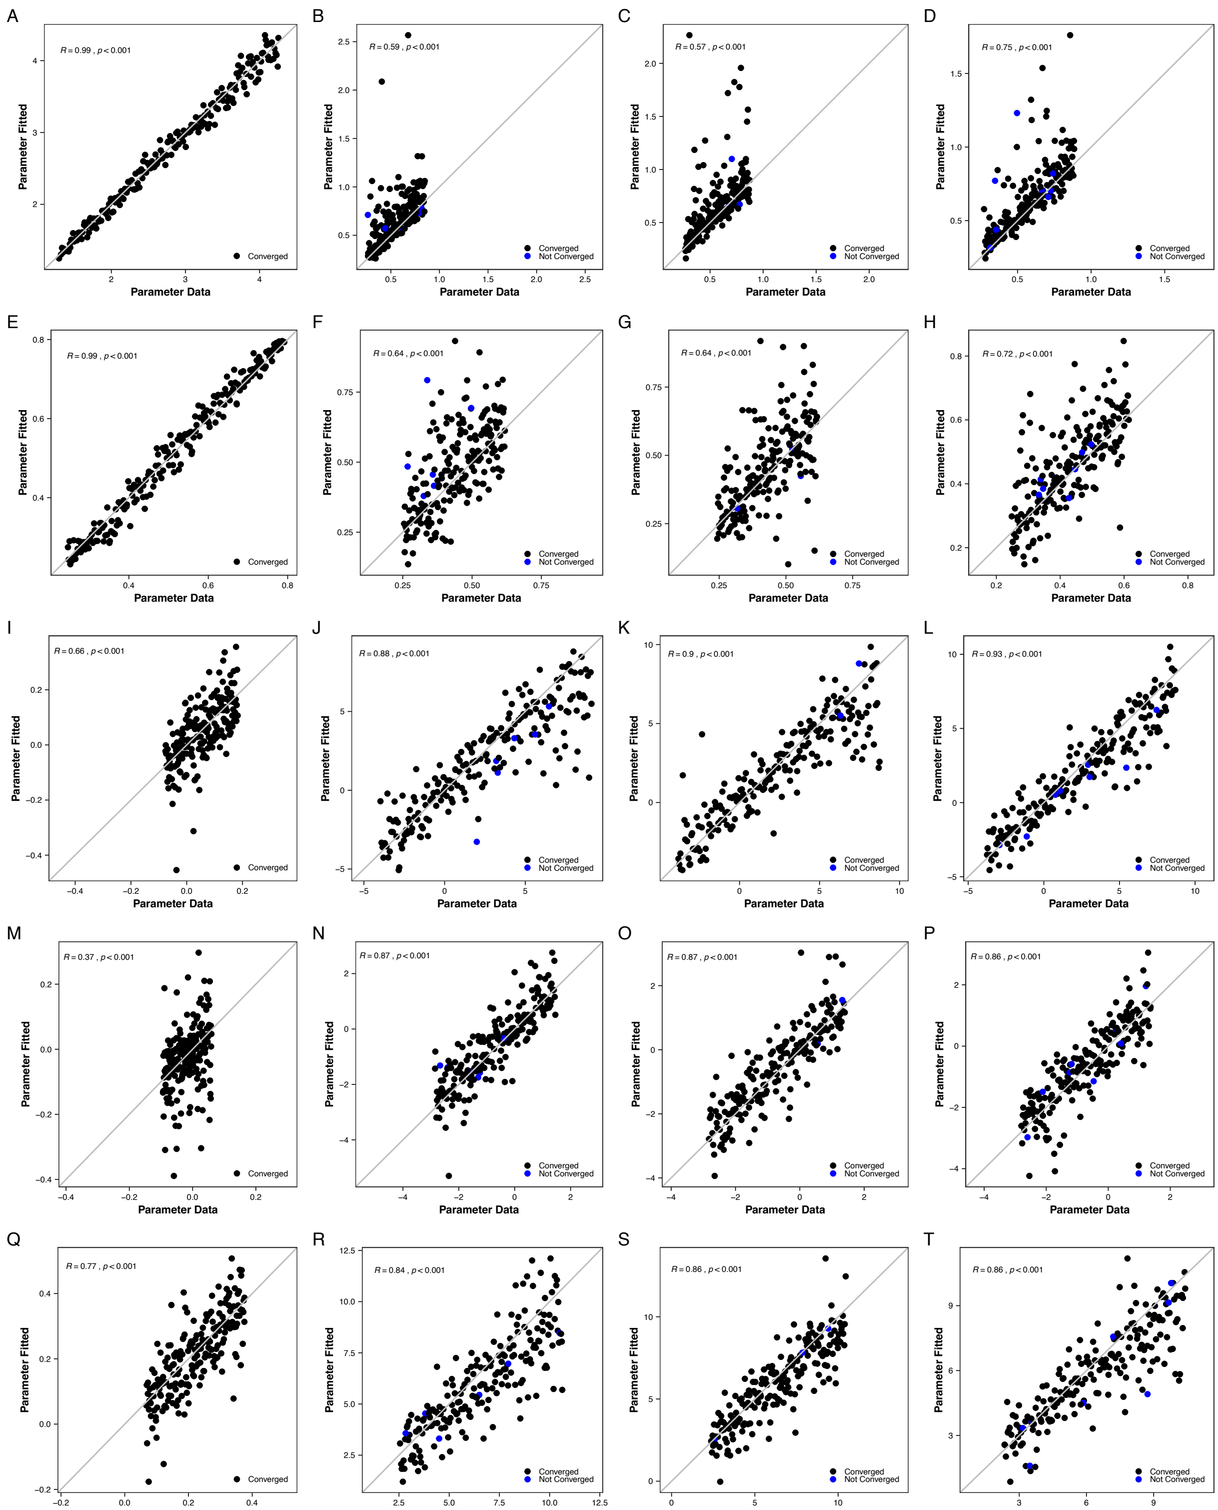

A
R = 0.99 , p < 0.001
B
R = 0.59 , p < 0.001
C
R = 0.57 , p < 0.001
D
R = 0.75 , p < 0.001
E
R = 0.99 , p < 0.001
F
R = 0.64 , p < 0.001
G
R = 0.64 , p < 0.001
H
R = 0.72 , p < 0.001
I
R = 0.66 , p < 0.001
J
R = 0.88 , p < 0.001
K
R = 0.9 , p < 0.001
L
R = 0.93 , p < 0.001
M
R = 0.37 , p < 0.001
N
R = 0.87 , p < 0.001
O
R = 0.87 , p < 0.001
P
R = 0.86 , p < 0.001
Q
R = 0.77 , p < 0.001
R
R = 0.84 , p < 0.001
S
R = 0.86 , p < 0.001
T
R = 0.86 , p < 0.001
Parameter Fitted
Parameter Data
Converged
Not Converged

Figure S37. Parameter Recovery. Correlations between simulated and fitted parameter values. (A, B, C, D) Boundary separation (a). (A) Standard DDM. (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E, F, G, H) Starting point bias (z). (E) Standard DDM. (F) Gamma DDM. (G) Time of Day Gamma DDM. (H) Time of Day and Offer Ratio Gamma DDM. (I, J, K, L) Drift intercept (β_0). (I) Standard DDM. (J) Gamma DDM. (K) Time of Day Gamma DDM. (L) Time of Day and Offer Ratio Gamma DDM. (M, N, O, P) Drift slope for list price (β_1). (M) Standard DDM. (N) Gamma DDM. (O) Time of Day Gamma DDM. (P) Time of Day and Offer Ratio Gamma DDM. (Q, R, S, T) Drift slope for offer ratio (β_2). (Q) Standard DDM. (R) Gamma DDM. (S) Time of Day Gamma DDM. (T) Time of Day and Offer Ratio Gamma DDM.

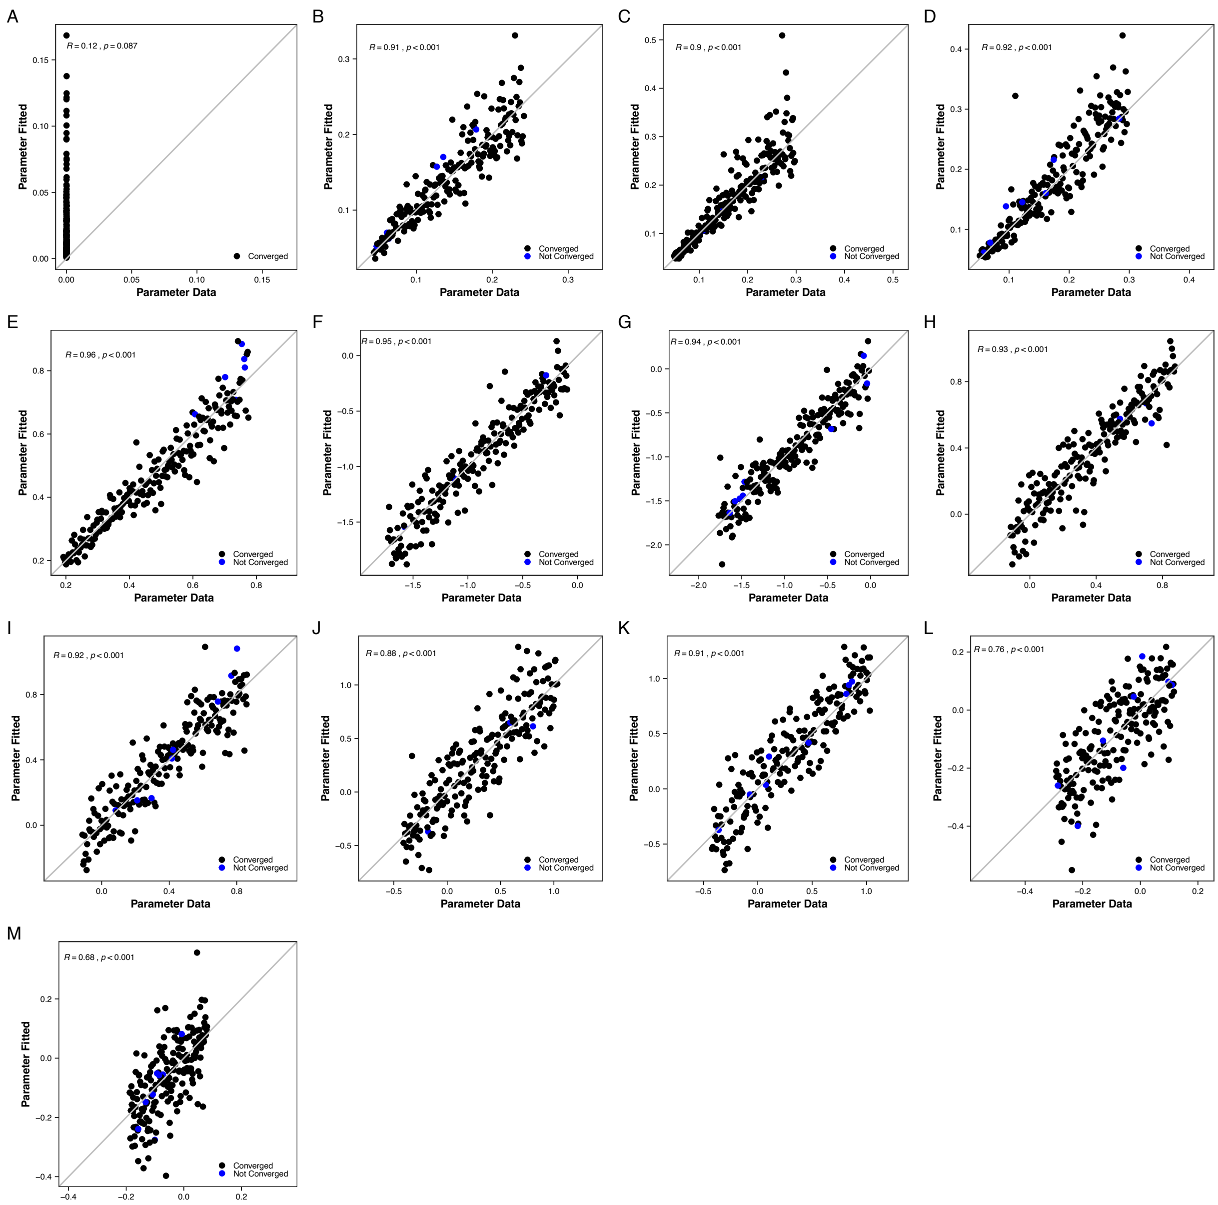

Figure S38. Parameter Recovery. Correlations between simulated and fitted parameter values for non-decision time parameters. (A) Standard DDM (t). (B, C, D) Rate parameter of Gamma distribution (β). (B) Gamma DDM. (C) Time of Day Gamma DDM. (D) Time of Day and Offer Ratio Gamma DDM. (E) Shape parameter of Gamma distribution (α) for Gamma DDM. (F, G, H, I, J, K, L, M) Parameters in sinusoidal function of the shape parameter. (F) Intercept Time of Day Gamma DDM (h_0). (G) Intercept of Time of Day and Offer Ratio Gamma DDM (h_0). (H) Linear time of day parameter of Time of Day Gamma DDM (h_1). (I) Linear time of day parameter of Time of

Day and Offer Ratio Gamma DDM (h_1). (J) Quadratic time of day parameter of Time of Day Gamma DDM (h_2). (K) Quadratic time of day parameter of Time of Day and Offer Ratio Gamma DDM (h_2). (L) Linear offer ratio parameter of Time of Day and Offer Ratio Gamma DDM (d_1). (M) Quadratic offer ratio parameter of Time of Day and Offer Ratio Gamma DDM (d_2).

Supplementary Tables

	Mean	St.dev.	Median
Listing-level data			
List price	491	40751	50
Sold Price	142	629	32
Decline price	281	5902	35
Accept price	191	21135	10
Item relisted	0.393	0.489	0
Number of views	40	111	17
Number of watchers	2	4	1
Number of photos	3	3	2
Sold through Best Offer	0.865	0.342	1
Sold price/List price	0.77	0.16	1
Bargained price/List price	0.734	0.142	1
Accept price/List price	0.826	0.117	1
Decline price/List price	0.693	0.18	1
Accept threshold present	0.203	0.402	0
Reject threshold present	0.284	0.451	0
Offer-level data			
Buyer's first offer/List price	0.609	0.195	1
Seller's response time (hours)	7	12	1
Seller bargain exchange	2926	6872	443
Buyer bargain exchange	132	592	26
Seller feedbacks	6242	23027	1471
Seller rating (percent)	99.646	2	99.87
Seller listings	11993	111900	1367
Seller BO listings	2465	14514	222
Messages included in offer	0.073	0.261	0
Item listing age (days)	25	53	6
Offer expired	0.069	0.254	0
Offer accepted	0.289	0.453	0
Offer rejected	0.135	0.342	0
Offer countered	0.262	0.44	0
Offer autodeclined	0.205	0.403	0
Offer autoaccepted	0.03	0.171	0
Offer rejected – another offer accepted	0.01	0.098	0

Table S1. Summary statistics for eBay observational data variables.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer First Offer/Price (p1/p0)	-0.24^{***}	-0.23^{***}	-0.23^{***}	0.10^{***}	0.12^{***}	0.11^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Offer Create Hour		0.20^{***}	0.20^{***}		0.15^{***}	0.15^{***}
		(0.004)	(0.004)		(0.01)	(0.01)
Offer Create Hour Squared		0.46^{***}	0.46^{***}		0.41^{***}	0.41^{***}
		(0.003)	(0.003)		(0.01)	(0.01)
Price		0.03^{***}	0.03^{***}		-0.06^{***}	-0.06^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Views		0.01	0.01		0.01^{**}	0.01^{**}
		(0.01)	(0.01)		(0.005)	(0.005)
Watchers		0.16^{***}	0.15^{***}		0.01^{**}	0.01^{**}
		(0.01)	(0.01)		(0.005)	(0.005)
Item Relisted		-0.01	-0.01		0.08^{***}	0.08^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Listing Age in Days		0.02^{***}	0.02^{***}		0.12^{***}	0.11^{***}
		(0.004)	(0.004)		(0.01)	(0.01)
Photos		-0.02^{**}	-0.02^{**}		-0.03^{***}	-0.03^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Buyer Exchange Number			-0.003			-0.03^{***}
			(0.004)			(0.01)
Seller Exchange Number			-0.14^{***}			-0.01
			(0.02)			(0.03)
Seller Listings			0.002			0.01
			(0.05)			(0.01)
Seller BO Listings			0.002			-0.05^{**}
			(0.01)			(0.02)
Seller Feedbacks			0.05^{*}			0.07^{***}
			(0.02)			(0.02)
p1/p0:Seller Exchange Number			-0.003			-0.04
			(0.02)			(0.02)
Constant	0.30^{***}	-0.12^{***}	-0.16^{***}	0.42^{***}	0.02	0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Observations	255,936	255,936	255,936	116,787	116,787	116,787
Log Likelihood	-531,562.50	-522,356.50	-522,346.10	-248,936.60	-245,697.90	-245,692.10
Akaike Inf. Crit.	1,063,137.00	1,044,741.00	1,044,732.00	497,885.20	491,423.80	491,424.10
Bayesian Inf. Crit.	1,063,200.00	1,044,887.00	1,044,941.00	497,943.20	491,559.20	491,617.50
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S2. Linear regressions for eBay observational data of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller accepting, or rejecting the offers. These regressions are restricted to the range $p_1/p_0 = [0.36, 0.68]$. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	logRT (hours)		
	(reject)	(reject)	(reject)
	(1)	(2)	(3)
Buyer First Offer/Price (p1/p0)	0.21^{***}	0.23^{***}	0.23^{***}
	(0.01)	(0.01)	(0.01)
Offer Create Hour		0.10^{***}	0.10^{***}
		(0.005)	(0.005)
Offer Create Hour Squared		0.28^{***}	0.28^{***}
		(0.004)	(0.004)
Price		-0.10^{***}	-0.10^{***}
		(0.01)	(0.01)
Views		0.02^{***}	0.02^{***}
		(0.004)	(0.004)
Watchers		-0.02^{***}	-0.02^{***}
		(0.004)	(0.004)
Item Relisted		0.26^{***}	0.26^{***}
		(0.01)	(0.01)
Listing Age in Days		0.23^{***}	0.23^{***}
		(0.01)	(0.01)
Photos		-0.07^{***}	-0.07^{***}
		(0.01)	(0.01)
Buyer Exchange Number			-0.04^{***}
			(0.01)
Seller Exchange Number			-0.08^{***}
			(0.02)
Seller Listings			0.01
			(0.01)
Seller BO Listings			-0.01
			(0.01)
Seller Feedbacks			0.01
			(0.02)
p1/p0:Seller Exchange Number			-0.04
			(0.02)
Constant	1.74^{***}	1.45^{***}	1.43^{***}
	(0.01)	(0.01)	(0.02)
Observations	183,072	183,072	183,072
Log Likelihood	-399,243.00	-396,114.20	-395,834.60
Akaike Inf. Crit.	798,494.00	792,252.50	791,709.20
Bayesian Inf. Crit.	798,534.50	792,373.90	791,911.50
Note:			*p<0.05; **p<0.01; ***p<0.001

Table S3. Linear regressions for eBay observational data of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller rejecting or letting the offer expire. These regressions are restricted to the range $p_1/p_0 = [0.36, 0.68]$. These regressions include random effects (clustered by seller) on the intercept for (1) and (2) and intercept and offer ratio for (3). Expired offers are considered rejection with a 48 hour response time. All variables are z-scored.

	logRT (hours)		
	(counter)	(counter)	(counter)
	(1)	(2)	(3)
Buyer First Offer/Price (p1/p0)	-0.06^{***}	-0.03^{***}	-0.03^{***}
	(0.01)	(0.01)	(0.01)
Seller Counteroffer/Price	-0.10^{***}	-0.11^{***}	-0.11^{***}
	(0.01)	(0.01)	(0.01)
Offer Create Hour		0.15^{***}	0.15^{***}
		(0.005)	(0.005)
Offer Create Hour Squared		0.44^{***}	0.44^{***}
		(0.004)	(0.004)
Price		-0.03^{***}	-0.03^{***}
		(0.01)	(0.01)
Views		0.003	0.003
		(0.005)	(0.005)
Watchers		0.09^{***}	0.09^{***}
		(0.01)	(0.01)
Item Relisted		0.07^{***}	0.07^{***}
		(0.01)	(0.01)
Listing Age in Days		0.07^{***}	0.07^{***}
		(0.01)	(0.01)
Photos		-0.02^{**}	-0.02^{**}
		(0.01)	(0.01)
Buyer Exchange Number			-0.02^{***}
			(0.004)
Seller Exchange Number			-0.11^{***}
			(0.02)
Seller Listings			0.12^{**}
			(0.04)
Seller BO Listings			-0.02
			(0.01)
Seller Feedbacks			0.17^{***}
			(0.02)
p1/p0:Seller Exchange Number			-0.03^{***}
			(0.01)
Constant	0.06^{***}	-0.37^{***}	-0.36^{***}
	(0.01)	(0.01)	(0.01)
Observations	187,740	187,740	187,740
Log Likelihood	-395,095.60	-389,353.70	-389,308.40
Akaike Inf. Crit.	790,201.30	778,733.50	778,654.80
Bayesian Inf. Crit.	790,252.00	778,865.30	778,847.60
Note:			*p<0.05; **p<0.01; ***p<0.001

Table S4. Linear regressions for eBay observational data of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller countering the offer. These regressions are restricted to the range p_1/p_0 =

$[0.36, 0.68]$. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	P (accept)	P (accept)	P (accept)
	(1)	(2)	(3)
Buyer First Offer/Price (p1/p0)	2.13^{***}	2.04^{***}	1.97^{***}
	(0.01)	(0.01)	(0.01)
Price		-0.28^{***}	-0.28^{***}
		(0.01)	(0.01)
Views		-0.26^{***}	-0.26^{***}
		(0.01)	(0.01)
Watchers		-0.40^{***}	-0.40^{***}
		(0.01)	(0.01)
Item Relisted		0.29^{***}	0.29^{***}
		(0.01)	(0.01)
Listing Age in Days		0.39^{***}	0.37^{***}
		(0.01)	(0.01)
Photos		0.08^{***}	0.08^{***}
		(0.01)	(0.01)
Buyer Exchange Number			-0.05^{***}
			(0.004)
Seller Exchange Number			0.52^{***}
			(0.02)
Seller Listings			-0.01
			(0.01)
Seller BO Listings			-0.07^{***}
			(0.01)
Seller Feedbacks			-0.01
			(0.03)
p1/p0:Seller Exchange Number			-0.16^{***}
			(0.02)
Constant	1.12^{***}	1.20^{***}	1.42^{***}
	(0.01)	(0.01)	(0.02)
Observations	869,395	869,395	869,395
Log Likelihood	-272,164.40	-260,517.90	-260,010.40
Akaike Inf. Crit.	544,338.80	521,057.90	520,054.90
Bayesian Inf. Crit.	544,397.20	521,186.30	520,253.40
Note:			*p<0.05; **p<0.01; ***p<0.001

Table S5. Logistic regressions for eBay observational data of acceptance versus rejection of first buyer offer ratio (p_1/p_0) **(z-score)**. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer Second Offer/Price (p2/p0)	-0.23^{***}	-0.22^{***}	-0.22^{***}	0.11^{***}	0.12^{***}	0.12^{***}
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)
Offer Create Hour		0.12^{***}	0.12^{***}		0.01	0.01
		(0.02)	(0.02)		(0.02)	(0.02)
Offer Create Hour Squared		0.35^{***}	0.35^{***}		0.28^{***}	0.28^{***}
		(0.02)	(0.02)		(0.01)	(0.01)
Price		0.001	-0.002		-0.06^{**}	-0.05^{**}
		(0.03)	(0.03)		(0.02)	(0.02)
Views		0.05	0.05		0.01	0.01
		(0.05)	(0.05)		(0.01)	(0.01)
Watchers		0.16^{***}	0.16^{***}		-0.01	-0.01
		(0.04)	(0.04)		(0.02)	(0.02)
Item Relisted		0.11^{*}	0.11^{*}		-0.04	-0.03
		(0.04)	(0.04)		(0.04)	(0.04)
Listing Age in Days		0.07^{***}	0.07^{***}		0.11^{***}	0.11^{***}
		(0.02)	(0.02)		(0.02)	(0.02)
Photos		-0.07^{**}	-0.07^{**}		-0.06^{**}	-0.06^{**}
		(0.03)	(0.03)		(0.02)	(0.02)
Buyer Exchange Number			0.02			-0.02
			(0.06)			(0.01)
Seller Exchange Number			-0.03			-0.001
			(0.05)			(0.01)
Seller Listings			0.27			0.0004
			(0.16)			(0.02)
Seller BO Listings			-0.05			-0.08^{**}
			(0.03)			(0.03)
Seller Feedbacks			0.04			0.13^{***}
			(0.04)			(0.04)
p2/p0:Seller Exchange Number			0.04			-0.002
			(0.04)			(0.02)
Constant	0.31^{***}	-0.03	-0.03	-0.03	-0.27^{***}	-0.27^{***}
	(0.03)	(0.04)	(0.04)	(0.03)	(0.03)	(0.04)
Observations	11,667	11,667	11,667	17,346	17,346	17,346
Log Likelihood	-24,936.59	-24,732.11	-24,742.15	-38,292.71	-38,099.42	-38,107.38
Akaike Inf. Crit.	49,885.18	49,492.23	49,524.29	76,597.42	76,226.84	76,254.77
Bayesian Inf. Crit.	49,929.37	49,595.33	49,671.58	76,643.99	76,335.50	76,409.99
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S6. Linear regressions for eBay observational data of log(RT) on second buyer offer ratio (p_2/p_0) (z-score) conditional on the seller accepting, or rejecting the offers. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer Third Offer/Price (p3/p0)	-0.29^{**}	-0.27^{**}	-0.26^{**}	0.09	0.09	0.09
	(0.09)	(0.10)	(0.10)	(0.05)	(0.05)	(0.05)
Offer Create Hour		0.08	0.08		-0.01	-0.01
		(0.07)	(0.07)		(0.05)	(0.05)
Offer Create Hour Squared		0.36^{***}	0.36^{***}		0.25^{***}	0.25^{***}
		(0.06)	(0.06)		(0.04)	(0.04)
Price		0.01	0.01		0.003	-0.003
		(0.09)	(0.09)		(0.05)	(0.05)
Views		-0.01	-0.03		-0.001	0.001
		(0.15)	(0.15)		(0.04)	(0.04)
Watchers		0.03	0.04		0.03	0.03
		(0.13)	(0.13)		(0.05)	(0.05)
Item Relisted		-0.05	-0.06		-0.16	-0.16
		(0.14)	(0.14)		(0.11)	(0.11)
Listing Age in Days		0.005	0.003		0.22^{***}	0.21^{***}
		(0.07)	(0.07)		(0.05)	(0.05)
Photos		-0.12	-0.11		-0.13^{*}	-0.12^{*}
		(0.08)	(0.08)		(0.06)	(0.06)
Buyer Exchange Number			0.05			
			(0.03)			
Seller Feedbacks			0.02			0.20^{**}
			(0.14)			(0.07)
Seller Listings			0.97			-0.01
			(0.73)			(0.04)
Seller BO Listings			-0.05			0.08
			(0.08)			(0.06)
Constant	0.13	-0.22	-0.19	-0.33^{***}	-0.51^{***}	-0.47^{***}
	(0.09)	(0.12)	(0.13)	(0.06)	(0.09)	(0.09)
Observations	1,204	1,204	1,204	2,211	2,211	2,211
Log Likelihood	-2,635.86	-2,626.74	-2,629.29	-4,991.67	-4,976.35	-4,976.07
Akaike Inf. Crit.	5,279.72	5,277.48	5,290.58	9,991.34	9,976.70	9,982.14
Bayesian Inf. Crit.	5,300.10	5,338.60	5,372.08	10,014.14	10,045.12	10,067.66
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S7. Linear regressions for eBay observational data of log(RT) on third buyer offer ratio (p_3/p_0) (z-score) conditional on the seller accepting, or rejecting the offers. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Seller compromise (p0-p1s)/(p0-p1b)	-0.13^{***}	-0.12^{***}	-0.12^{***}	0.18^{***}	0.18^{***}	0.18^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Offer Create Hour		-0.04^{***}	-0.04^{***}		-0.06^{***}	-0.06^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Offer Create Hour Squared		0.31^{***}	0.31^{***}		0.25^{***}	0.25^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Price		-0.04^{***}	-0.04^{***}		-0.03^{***}	-0.03^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Views		0.04^{*}	0.04^{*}		-0.02^{*}	-0.02^{*}
		(0.02)	(0.02)		(0.01)	(0.01)
Watchers		0.25^{***}	0.25^{***}		0.08^{***}	0.08^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Item Relisted		0.15^{***}	0.15^{***}		-0.10^{***}	-0.10^{***}
		(0.02)	(0.02)		(0.02)	(0.02)
Listing Age in Days		0.10^{***}	0.10^{***}		-0.03^{***}	-0.03^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Photos		-0.02^{*}	-0.02^{*}		0.01	0.01
		(0.01)	(0.01)		(0.01)	(0.01)
Buyer Exchange Number			−0.02			−0.01
			(0.03)			(0.01)
(p0-p1s)/(p0-p1b):Exchange			0.02			−0.002
			(0.04)			(0.02)
Constant	0.57^{***}	0.25^{***}	0.25^{***}	1.06^{***}	0.83^{***}	0.83^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	81,793	81,793	81,793	76,591	76,591	76,591
Log Likelihood	-178,137.00	-176,720.40	-176,725.50	-167,897.40	-167,204.40	-167,210.70
Akaike Inf. Crit.	356,281.90	353,464.70	353,478.90	335,802.70	334,432.80	334,449.50
Bayesian Inf. Crit.	356,319.20	353,576.50	353,609.30	335,839.70	334,543.70	334,578.90
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S8. Linear regressions for eBay observational data of buyers' log(RT) on seller's compromise conditional on the buyer accepting, or rejecting the offers. The seller's compromise is the amount that they lowered their counteroffer, divided by the gap between the list price and the buyer's offer. A 100% compromise would be a counteroffer that matches the buyer's offer; a 0% compromise would be a counteroffer that is the list price. These regressions include random effects (clustered by buyer) on the intercept. All variables are z-scored.

	P (buyer second offer exists)		
	(1)	(2)	(3)
Buyer first offer/Price	-0.04^{***}	-0.04^{***}	-0.04^{***}
	(0.004)	(0.004)	(0.004)
Buyer first offer/Price Squared	-0.17^{***}	-0.16^{***}	-0.17^{***}
	(0.003)	(0.003)	(0.003)
Rejection RT (hours)			-0.10^{***}
			(0.004)
Time Offer Response Hour		0.001	
		(0.003)	
Time Offer Hour Response Squared		-0.01^{**}	
		(0.003)	
Price		-0.05^{***}	
		(0.004)	
Buyer Exchange Number		-0.08^{***}	
		(0.01)	
Seller Exchange Number		0.03^{***}	
		(0.003)	
Seller Feedbacks		-0.01^{***}	
		(0.004)	
Listing Age in Days		0.02^{***}	
		(0.003)	
Constant	-0.96^{***}	-0.96^{***}	-0.97^{***}
	(0.004)	(0.01)	(0.004)
Observations	226,653	226,653	226,653
Log Likelihood	-53,084.50	-51,874.45	-51,838.82
ρ	-0.54*** (0.01)	-0.49*** (0.02)	-0.49*** (0.02)
Note:			*p<0.05; **p<0.01; ***p<0.001

Table S9. Probit regression for eBay observational data for existence of buyers' second offer. All variables are z-scored. (1), (2) are the first stage of the Heckman correction.

	Buyer second offer/Price	
	(1)	(2)
Rejection RT (hours)	0.003***	0.003***
	(0.001)	(0.001)
Buyer first offer/Price	0.19***	0.19***
	(0.001)	(0.001)
Price		−0.02***
		(0.001)
Views		−0.003***
		(0.0005)
Watchers		−0.002***
		(0.001)
Photos		0.0004
		(0.0005)
Listing Age in Days		−0.001
		(0.0004)
Buyer Exchange Number		−0.01***
		(0.001)
Constant	0.70***	0.69***
	(0.003)	(0.003)
Observations	226,653	226,653
Log Likelihood	-53,084.50	-51,874.45
ρ	-0.54*** (0.01)	-0.49*** (0.02)
Note:	*p<0.05; **p<0.01; ***p<0.001	

Table S10. Heckman correction regression second stage for eBay observational data for buyers' second offer amount as a fraction of list price (p_2/p_0). All variables are z-scored.

	P(buyer second offer accepted)	
	(1)	(2)
Rejection RT (hours)	0.13^{***}	0.12^{***}
	(0.02)	(0.02)
Buyer Second Offer/Price (p2/p0)	1.00^{***}	1.08^{***}
	(0.02)	(0.02)
Price		-0.49^{***}
		(0.02)
Buyer Exchange Number		0.01
		(0.02)
Photos		0.17^{***}
		(0.02)
Views		-0.60^{***}
		(0.04)
Listing Age in Days		0.20^{***}
		(0.02)
Watchers		-0.55^{***}
		(0.03)
Constant	-0.51^{***}	-0.74^{***}
	(0.02)	(0.02)
Observations	23,435	23,435
Log Likelihood	-13,543.06	-12,279.99
Akaike Inf. Crit.	27,092.13	24,577.98
Note:	*p<0.05; **p<0.01; ***p<0.001	

Table S11. Logistic regression for eBay observational data for probability of seller accepting buyer's second offer. All variables are z-scored.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer First Offer/Price (p1/p0)	-0.12^{***}	-0.11^{***}	-0.11^{***}	-0.0001	0.03^{**}	0.02^{*}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Offer Create Hour		0.19^{***}	0.18^{***}		0.15^{***}	0.15^{***}
		(0.003)	(0.003)		(0.004)	(0.004)
Offer Create Hour Squared		0.45^{***}	0.45^{***}		0.43^{***}	0.43^{***}
		(0.003)	(0.003)		(0.003)	(0.003)
Price		-0.0004	0.0000		-0.04^{***}	-0.04^{***}
		(0.004)	(0.004)		(0.004)	(0.004)
Views		0.002	0.002		0.01^{**}	0.01^{**}
		(0.004)	(0.004)		(0.004)	(0.004)
Watchers		0.09^{***}	0.09^{***}		0.03^{***}	0.03^{***}
		(0.004)	(0.004)		(0.004)	(0.004)
Item Relisted		0.03^{***}	0.03^{***}		0.08^{***}	0.08^{***}
		(0.01)	(0.01)		(0.01)	(0.01)
Listing Age in Days		0.04^{***}	0.04^{***}		0.09^{***}	0.09^{***}
		(0.003)	(0.003)		(0.005)	(0.005)
Photos		-0.01^{*}	-0.01		-0.02^{***}	-0.02^{***}
		(0.004)	(0.004)		(0.01)	(0.01)
Buyer Exchange Number			-0.01^{***}			-0.03^{***}
			(0.003)			(0.004)
Seller Exchange Number			-0.13^{***}			-0.06^{***}
			(0.02)			(0.02)
Seller Listings			0.06			0.01
			(0.04)			(0.01)
Seller BO Listings			-0.01			-0.004
			(0.01)			(0.01)
Seller Feedbacks			0.08^{***}			0.09^{***}
			(0.02)			(0.02)
p1/p0:Seller Exchange Number			-0.02			-0.03^{*}
			(0.01)			(0.01)
Constant	0.25^{***}	-0.20^{***}	-0.23^{***}	0.25^{***}	-0.17^{***}	-0.17^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	413,609	413,609	413,609	274,460	274,460	274,460
Log Likelihood	-859,399.10	-845,408.80	-845,378.00	-577,747.80	-569,602.30	-569,583.20
Akaike Inf. Crit.	1,718,810.00	1,690,846.00	1,690,796.00	1,155,508.00	1,139,233.00	1,139,206.00
Bayesian Inf. Crit.	1,718,876.00	1,690,999.00	1,691,015.00	1,155,571.00	1,139,380.00	1,139,417.00
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S12. Linear regressions for eBay observational data of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller response with counteroffers pooled with acceptances or rejections. Linear regressions for eBay observational data of log(RT) on first buyer offer ratio (p_1/p_0) (z- score) conditional on the seller accepting or countering, or rejecting or countering the offers. These regressions are restricted to the range

$p_1/p_0 = [0.36, 0.68]$. These regressions include random effects (clustered by seller) on the intercept and offer ratio. All variables are z-scored.

	Mean	St.dev.	Median
Offer Ratio	0.6	0.24	1
Price (dollars)	15	4	15
Sale Price (Offer Accepted or Auto-accepted)	0.955	1	1
Bargained Price (Offer Accepted)	0.806	0.138	1
Accept Price Ratio	0.9	0	1
Decline Price Ratio	0.375	0.106	0
Response Time (hour)	1	2	0
Year since Registration	18	6	21
Viable Items	36	21	28
Seller Feedbacks	9644	25588	2746
Number of Best Offers	0.46	0.754	0
Item Relisted	0.667	0.472	1
Offer Expired	0.036	0.188	0
Offer Accepted	0.441	0.497	0
Offer Declined	0.144	0.351	0
Offer Countered	0.311	0.464	0
Offer Auto-declined	0.062	0.241	0
Offer Auto-accepted	0.005	0.074	0

Table S13. Summary statistics of variables of interest for eBay field experiment 1.

	Mean	St.dev.	Median
Offer Ratio	0.45	0.244	0
Price (dollars)	13	2	13
Sale Price (Offer Accepted or Auto-accepted)	0.734	1	1
Bargained Price (Offer Accepted)	0.629	0.172	1
Accept Price Ratio	0.764	0.055	1
Decline Price Ratio	0.243	0.133	0
Response Time (hour)	3	7	1
Year since Registration	14	8	14
Viable Items	56	67	36
Seller Feedbacks	8876	20191	3028
Number of Best Offers	0.33	0.94	0
Item Relisted	0.433	0.496	0
Offer Expired	0.066	0.248	0
Offer Accepted	0.431	0.495	0
Offer Declined	0.186	0.389	0
Offer Countered	0.254	0.435	0
Offer Auto-declined	0.06	0.237	0
Offer Auto-accepted	0.004	0.06	0

Table S14. Summary statistics of variables of interest for eBay field experiment 2.

	P(accept)	P(accept)	P(accept)	P(accept)	P(accept)	P(accept)	P(accept)	P(accept)	P(accept)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Offer/Price (p1/p0)	3.55***	4.90***	6.02***	2.63***	2.76***	2.80***	2.64***	2.83***	2.88***
	(0.44)	(1.05)	(1.58)	(0.12)	(0.13)	(0.13)	(0.11)	(0.12)	(0.13)
Number of Best Offers		−4.45***	−5.81***		−0.86***	−0.88***		−1.11***	−1.10***
		(0.94)	(1.65)		(0.12)	(0.12)		(0.12)	(0.12)
Item Relisted		1.24	2.63		−0.35*	−0.22		−0.41*	−0.21
		(1.07)	(1.61)		(0.17)	(0.18)		(0.16)	(0.17)
Feedbacks		−0.34	−0.78		0.14	0.17*		−0.02	0.02
		(0.43)	(0.61)		(0.08)	(0.09)		(0.08)	(0.08)
Years since Registration			0.41			−0.08			0.02
			(2.38)			(0.11)			(0.14)
Viable Items			−1.24			−0.03			−0.15
			(0.96)			(0.12)			(0.11)
Price			−0.99			0.03			0.12
			(0.73)			(0.17)			(0.17)
p1/p0:Feedbacks			−0.58			−0.02			0.03
			(1.89)			(0.12)			(0.15)
p1/p0:Years			−0.46			0.22			0.15
			(0.90)			(0.12)			(0.12)
p1/p0:Viable Items			−2.20			0.21			0.24
			(1.23)			(0.21)			(0.20)
Constant	1.91***	4.47***	5.39***	1.53***	1.12***	1.04***	1.53***	2.19***	2.11***
	(0.36)	(1.08)	(1.59)	(0.10)	(0.15)	(0.16)	(0.09)	(0.15)	(0.15)
Observations	321	319	308	1,873	1,858	1,844	2,194	2,177	2,152
Log Likelihood	-52.28	-20.55	-16.25	-507.16	-472.23	-465.06	-590.15	-526.22	-513.77
Akaike Inf. Crit.	108.56	51.11	54.51	1,018.32	954.46	952.12	1,184.30	1,062.45	1,049.54

Note: *p<0.05; **p<0.01; ***p<0.001

Table S15. Logistic regressions for eBay field experiments 1 and 2 of acceptance versus rejection of first buyer offer ratio (p_1/p_0) (z-score). (1), (2), (3) Experiment 1. (4), (5), (6) Experiment 2. (7), (8), (9) Experiments 1 and 2. All variables are z-scored except Number of Best Offer.

	logRT (hours)		
	(all)	(all)	(all)
	(1)	(2)	(3)
Offer/Price	-0.20	-0.25^{***}	-0.23^{***}
	(0.21)	(0.07)	(0.07)
Rejected	1.34^{*}	0.76^{***}	0.72^{***}
	(0.60)	(0.17)	(0.16)
Offer/Price:Rejected	1.11^{*}	0.28^{*}	0.25
	(0.50)	(0.14)	(0.13)
Constant	-1.10^{***}	-0.93^{***}	-0.97^{***}
	(0.27)	(0.12)	(0.11)
Observations	321	1,873	2,194
Log Likelihood	-663.72	-3,555.83	-4,231.55
Akaike Inf. Crit.	1,339.43	7,133.66	8,485.11
Bayesian Inf. Crit.	1,362.06	7,194.54	8,547.74
Note:	*p<0.05; **p<0.01; ***p<0.001		

Table S16. Linear regressions for eBay field experiments 1 and 2 for log RT (in hours) for all types of seller responses. (1) Regression includes random effects (clustered by seller) on the intercept. (2), (3) Regressions include random effects (clustered by seller) on the intercept and offer ratio and type of seller response. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. The first offer as a fraction of list price (p_1/p_0) is z-scored.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Offer/Price (p1/p0)	−0.19	−0.16	−0.21	0.84	0.79	0.90
	(0.21)	(0.22)	(0.23)	(0.45)	(0.47)	(0.92)
Number of Best Offers		0.05	0.08		−0.01	0.03
		(0.35)	(0.35)		(0.40)	(0.44)
Item Relisted		0.44	0.60		0.43	0.52
		(0.42)	(0.42)		(0.53)	(0.59)
Price		−0.07	−0.11		−0.14	−0.17
		(0.14)	(0.15)		(0.22)	(0.25)
Feedbacks			0.24			−1.10
			(0.44)			(4.92)
Years since Registration			0.30			−0.34
			(0.27)			(1.22)
Viable Items			−0.31			0.15
			(0.27)			(1.03)
p1/p0:Feedbacks			0.23			−1.09
			(0.38)			(3.72)
p1/p0:Years			−0.24			−0.16
			(0.21)			(0.97)
p1/p0:Viable Items			0.11			0.16
			(0.22)			(0.84)
Constant	-1.10^{***}	-1.43^{***}	-1.52^{***}	0.27	−0.08	−0.14
	(0.28)	(0.41)	(0.41)	(0.56)	(0.82)	(1.38)
Observations	242	242	237	79	77	71
Log Likelihood	-509.61	-510.04	-498.42	-158.83	-153.95	-138.15
Akaike Inf. Crit.	1,027.21	1,034.08	1,022.84	325.65	321.89	302.30
Bayesian Inf. Crit.	1,041.17	1,058.50	1,067.92	335.13	338.30	331.72
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S17. Linear regressions for eBay field experiment 1 of log(RT) on first buyer offer ratio (p_1/p_0) **(z-score) conditional on the seller accepting, or rejecting the offers.** These regressions include random effects (clustered by seller) on the intercept. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Offer/Price (p1/p0)	-0.26^{***}	-0.25^{***}	-0.26^{***}	0.12	0.13	0.13
	(0.07)	(0.07)	(0.07)	(0.11)	(0.11)	(0.11)
Number of Best Offers		-0.16	-0.15		-0.02	-0.02
		(0.08)	(0.08)		(0.04)	(0.04)
Item Relisted		0.10	0.03		0.25	0.25
		(0.16)	(0.16)		(0.14)	(0.15)
Feedbacks			0.02			-0.15
			(0.13)			(0.18)
Years since Registration			0.31^{*}			-0.06
			(0.14)			(0.17)
Viable Items			-0.01			0.13
			(0.16)			(0.19)
p1/p0:Feedbacks			0.05			-0.11
			(0.07)			(0.13)
p1/p0:Years			-0.18^{*}			-0.08
			(0.08)			(0.12)
p1/p0:Viable Items			0.15			0.01
			(0.11)			(0.14)
Constant	-0.91^{***}	-1.09^{***}	-1.03^{***}	-0.03	-0.12	-0.12
	(0.13)	(0.16)	(0.16)	(0.15)	(0.17)	(0.18)
Observations	1,308	1,298	1,286	565	560	558
Log Likelihood	-2,608.27	-2,589.82	-2,569.63	-927.14	-921.76	-924.37
Akaike Inf. Crit.	5,228.55	5,195.64	5,167.27	1,866.29	1,859.51	1,876.75
Bayesian Inf. Crit.	5,259.61	5,236.99	5,239.50	1,892.31	1,894.14	1,937.29
Note:			*p<0.05; **p<0.01; ***p<0.001			

Table S18. Linear regressions for eBay field experiment 2 of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller accepting, or rejecting the offers. These regressions include random effects (clustered by seller) on the intercept and offer ratio. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

	logRT (hours)					
	(accept)	(accept)	(accept)	(reject)	(reject)	(reject)
	(1)	(2)	(3)	(4)	(5)	(6)
Offer/Price (p1/p0)	-0.24^{***}	-0.24^{**}	-0.26^{***}	0.11	0.11	0.12
	(0.07)	(0.07)	(0.07)	(0.10)	(0.11)	(0.11)
Number of Best Offers		-0.15	-0.14		-0.03	-0.02
		(0.08)	(0.08)		(0.05)	(0.05)
Item Relisted		0.15	0.09		0.17	0.19
		(0.15)	(0.15)		(0.14)	(0.15)
Feedbacks			0.08			-0.02
			(0.12)			(0.18)
Years since Registration			0.29^{*}			-0.16
			(0.12)			(0.16)
Viable Items			-0.02			0.11
			(0.14)			(0.19)
p1/p0:Feedbacks			0.10			-0.07
			(0.08)			(0.14)
p1/p0:Years			-0.19^{*}			-0.09
			(0.08)			(0.12)
p1/p0:Viable Items			0.12			-0.02
			(0.11)			(0.14)
Constant	-0.95^{***}	-0.99^{***}	-0.93^{***}	-0.14	-0.20	-0.22
	(0.11)	(0.13)	(0.14)	(0.14)	(0.17)	(0.18)
Observations	1,550	1,540	1,523	644	637	629
Log Likelihood	-3,120.71	-3,102.28	-3,070.76	-1,102.36	-1,091.27	-1,082.53
Akaike Inf. Crit.	6,253.43	6,220.56	6,169.51	2,216.72	2,198.54	2,193.07
Bayesian Inf. Crit.	6,285.50	6,263.28	6,244.11	2,243.53	2,234.20	2,255.29
Note:						*p<0.05; **p<0.01; ***p<0.001

Table S19. Linear regressions for eBay field experiments 1 and 2 of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller accepting, or rejecting the offers. These regressions include random effects (clustered by seller) on the intercept and offer ratio. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

	logRT (hours)								
	(counter)	(counter)	(counter)	(counter)	(counter)	(counter)	(counter)	(counter)	(counter)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Offer/Price (p1/p0)	−0.06	−0.08	−0.01	0.07	0.08	0.07	0.05	0.06	0.05
	(0.19)	(0.19)	(0.20)	(0.09)	(0.09)	(0.09)	(0.08)	(0.08)	(0.08)
Number of Best Offers		−0.16	−0.16		−0.07	−0.06		−0.08	−0.07
		(0.13)	(0.13)		(0.06)	(0.06)		(0.06)	(0.06)
Item Relisted		0.16	0.27		−0.04	−0.06		−0.01	−0.01
		(0.47)	(0.48)		(0.20)	(0.20)		(0.18)	(0.18)
Feedbacks			0.33			−0.10			0.07
			(0.27)			(0.14)			(0.12)
Years since Registration			−0.24			0.10			−0.02
			(0.27)			(0.15)			(0.13)
Viable Items			0.04			0.20			0.17
			(0.28)			(0.12)			(0.11)
Price		0.12	0.07						
		(0.14)	(0.15)						
p1/p0:Feedbacks			−0.07			−0.04			−0.02
			(0.19)			(0.11)			(0.10)
p1/p0:Years			−0.06			0.10			0.06
			(0.21)			(0.10)			(0.09)
p1/p0:Viable Items			0.26			−0.05			−0.03
			(0.20)			(0.07)			(0.06)
Constant	-1.02^{***}	-1.03^{**}	-1.03^{*}	-0.77^{***}	-0.77^{***}	-0.77^{***}	-0.81^{***}	-0.76^{***}	-0.76^{***}
	(0.25)	(0.39)	(0.41)	(0.13)	(0.16)	(0.16)	(0.11)	(0.15)	(0.15)
Observations	171	171	171	771	756	749	942	927	920
Log Likelihood	-333.18	-334.00	-334.66	-1,563.93	-1,536.37	-1,529.92	-1,898.63	-1,870.94	-1,865.88
Akaike Inf. Crit.	674.36	682.01	695.32	3,135.86	3,084.73	3,083.84	3,805.27	3,753.89	3,755.77
Bayesian Inf. Crit.	686.93	704.00	736.16	3,154.45	3,112.50	3,139.27	3,824.66	3,782.88	3,813.66
Note:									*p<0.05; **p<0.01; ***p<0.001

Table S20. Linear regressions for eBay field experiments 1 and 2 of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller countering the offers. (1), (2), (3) Experiment 1. (4), (5), (6) Experiment 2. (7), (8), (9) Experiment 1 and 2. Regressions include random effects (clustered by seller) on the intercept. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

	logRT (hours)		
	(reject)	(reject)	(reject)
	(1)	(2)	(3)
Offer/Price (p1/p0)	0.12	0.14	0.13
	(0.11)	(0.11)	(0.11)
Number of Best Offers		−0.02	−0.02
		(0.04)	(0.04)
Item Relisted		0.26	0.26
		(0.14)	(0.15)
Feedbacks			−0.15
			(0.18)
Years since Registration			−0.06
			(0.17)
Viable Items			0.12
			(0.19)
p1/p0:Feedbacks			−0.11
			(0.13)
p1/p0:Years			−0.09
			(0.12)
p1/p0:Viable Items			0.001
			(0.15)
RT Source Message	0.08	0.09	0.09
	(0.12)	(0.12)	(0.12)
Constant	−0.04	−0.14	−0.14
	(0.16)	(0.17)	(0.18)
Observations	565	560	558
Log Likelihood	-928.15	-922.68	-925.26
Akaike Inf. Crit.	1,870.30	1,863.35	1,880.53
Bayesian Inf. Crit.	1,900.65	1,902.30	1,945.39
Note:	*p<0.05; **p<0.01; ***p<0.001		

Table S21. Linear regressions for eBay field experiment 2 of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller rejecting the offers controlling for the source of the response time. Regressions include random effects (clustered by seller) on the intercept and offer ratio. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

	logRT (hours)		
	(reject)	(reject)	(reject)
	(1)	(2)	(3)
Offer/Price (p1/p0)	0.45***	0.44***	0.44***
	(0.09)	(0.09)	(0.10)
Number of Best Offers		−0.04	−0.04
		(0.04)	(0.04)
Item Relisted		0.12	0.13
		(0.15)	(0.15)
Feedbacks			−0.04
			(0.14)
Years since Registration			−0.06
			(0.16)
Viable Items			0.09
			(0.16)
p1/p0:Feedbacks			−0.08
			(0.09)
p1/p0:Years			0.02
			(0.10)
p1/p0:Viable Items			−0.04
			(0.10)
Constant	0.79***	0.73***	0.74***
	(0.15)	(0.17)	(0.18)
Observations	864	842	834
Log Likelihood	-1,532.55	-1,491.98	-1,483.57
Akaike Inf. Crit.	3,077.09	2,999.95	2,995.14
Bayesian Inf. Crit.	3,105.66	3,037.84	3,061.31
Note:			*p<0.05; **p<0.01; ***p<0.001

Table S22. Linear regressions for eBay field experiments 1 and 2 of log(RT) on first buyer offer ratio (p_1/p_0) (z-score) conditional on the seller rejecting the offers or letting the offer expire. These regressions include random effects (clustered by seller) on the intercept and offer ratio. The random effects structure was chosen based on the model with the lowest AIC from a model comparison. All variables are z-scored except Number of Best Offer.

SI References

1. J. J. Heckman, Sample selection bias as a specification error. *Econometrica* **47**, 153–161 (1979).
2. O. Toomet, A. Henningsen, Sample selection models in R: Package sampleSelection. *J. Stat. Soft.* **27**, 1–23 (2008).