A distributional and dynamic theory of pricing and preference

Peter D. Kvam
University of Florida

Jerome R. Busemeyer
Indiana University

Theories that describe how people assign prices and make choices are typically based on the idea that both of these responses are derived from a common static, deterministic function used to assign utilities to options. However, preference reversals – where prices assigned to gambles conflict with preference orders elicited through binary choices – indicate that the response processes underlying these different methods of evaluation are more intricate. We address this issue by formulating a new computational model that assumes an initial bias or anchor that depends on type of price task (buying, selling, or certainty equivalents) and a stochastic evaluation accumulation process that depends on gamble attributes. To test this new model, we investigated choices and prices for a wide range of gambles and price tasks, including pricing under time pressure. In line with model predictions, we found that price distributions possessed stark skew that depended on the type of price and the attributes of gambles being considered. Prices were also sensitive to time pressure, indicating a dynamic evaluation process underlying price generation. The model out-performed prospect theory in predicting prices, and additionally predicted the associated response times, which no prior model has accomplished. Finally, we show that the model successfully predicts out-of-sample choices and choice response times. This price accumulation model therefore provides a superior account of the distributional and dynamic properties of price, leveraging process-level mechanisms to provide a more complete account the valuation processes common across multiple methods of eliciting preference.

Keywords: price, cognitive modeling, buying, selling, preference reversal

Exchanges where sums of money are traded for a desired choice option form an important part of people's daily economic behavior. As part of these exchanges, both consumers and producers frequently have to perform the task of evaluating the price of choice options by selecting a satisfactory buying price, selling price, or certainty equivalent. For example, people often evaluate the price for selling an investment, the price for buying a product, or the price equivalent of medical treatment for insurance. This ability to estimate the monetary value of choice options is a core component of a person's capacity to form preferences and make decisions. Despite the ease with which people seem to do this, the assignment of a monetary value to a choice option is based on complex, dynamic, and stochastic mental processes that involve both cognitive and affective components.

Peter D. Kvam (pkvam@ufl.edu), Department of Psychology, University of Florida. Jerome R. Busemeyer (jbusemey@indiana.edu), Department of Psychological & Brain Sciences and Cognitive Science, Indiana University Bloomington. The work presented here was supported by a grant from the Air Force Office of Scientific Research, FA9550-15-1-0343, and National Institute of Drug Abuse grant, NIDA R01DA021421. Some parts of this work were previously presented at the 2018 Cognitive Science Society and Mathematical Psychology conferences.

Traditionally, theories of price judgments have been based on deterministic and static utility theories. The basic idea behind these theories is that prices are determined by finding the exact monetary value that makes the person indifferent between the utility of the choice option and the utility of the monetary value (see, e.g., Luce, 2000; Becker et al., 1964). However, a deterministic representation of this task fails to account for the fact that people cannot reliably assign a price to a risky prospect (Schoemaker & Hershey, 1992; Butler & Loomes, 2007). Instead, there is always some variability in the price that is assigned by a person to the same choice option on different occasions. Utility theorists typically try to avoid this response variability problem by asking for a single price or, if several replications are obtained, by computing the mean or median price given to an alternative. In foregoing the distribution of prices, this method ignores the fact that the variability in price responses changes systematically across choice options with different attributes (Bostic et al., 1990), and furthermore ignores interesting properties concerning the shape of price distributions.

A further problem with traditional utility theories is that they are static, and so they fail to describe the dynamic cognitive processes that generate a price response. Instead, static utility-based theories tend to describe prices "as if" a person is transforming the attributes of an option and adding or multiplying them to compute the overall value (Berg & Gigerenzer, 2010), but do not describe the decision processes that underlie these transformations. As a result, they fail to predict the outcomes of any process measures, such as response times or process tracing data. Pachur et al. (2018) sought to remedy this issue by pursuing a link between eye tracking data and the parameters of prospect theory, and took the first steps toward relating the static structure of past models to the dynamic structure of cognitive processes, but stopped short of constructing a generative model of value or price. As we show in this paper, such a model is critical to understanding price because the time allocated to constructing a price response can influence the distribution of prices that is eventually assigned to a choice option. Therefore, a static model will not only provide an incomplete picture of the cognitive processes underlying valuation, but it will also fail to make predictions concerning the effect of internally or externally controlled stopping time on the price that is eventually selected.

One final area where these models tend to fail is in accounting for *preference reversals*, where one choice alternative A is selected over B in a binary choice task, but B is priced higher than A when they are viewed separately and assigned prices (Lichtenstein & Slovic, 1971; Slovic & Lichtenstein, 1983). One popular explanation is that the actual utilities or probability weights assigned to options differ between binary choice and pricing (Tversky et al., 1990). However, an alternative proposal is that the underlying representations of utility are coherent across choices and prices, and instead the response processes used to generate the two measures of preference differ (J. G. Johnson & Busemeyer, 2005). The latter idea is intuitively more appealing because it allows for internally consistent mechanisms for valuation, even if different empirical measures appear to diverge.

Theoretical and empirical basis

The purpose of this article is to build and expand on this previous work by developing and empirically testing a new dynamic and stochastic model of choice and pricing. It builds on previous work that empirical investigated these various forms of judgments, much of which has focused on the differences in prices between buyers and sellers, frequently referred to as an *endowment effect* (reviewed in Morewedge & Giblin, 2015). In many cases, this effect has been explained through loss aversion, where buyers treat the failure to purchase an item as foregoing a potential gain while sellers treat the sale of an item as a potential loss. Because people appear to weight losses more heavily than gains (Kahneman & Tversky, 1979; Kahneman et al., 1991) sellers place more value on losing the item than buys do on acquiring it. This is typically instantiated as an asymmetry between losses and gains in the utility function, using a parameter that multiplies the utility of losses by a constant λ relative to gains.

However, this only brushes the surface in terms of the cognitive mechanisms underlying buyer-seller differences. Work by Pachur & Scheibehenne (2012) empirically investigated these gaps and the potential underpinnings, and showed that these differences were due at least in part to differences in information search between buyers and sellers. They observed that sellers were more likely to terminate information sampling about potential prospects after experiencing a positive outcome (indicating that the potential sale of a choice option should be worth more) while buyers were more likely to terminate search after a negative outcome. They additionally showed that the buyer-seller gap was reduced when people gathered more information prior to making a decision, suggesting that a key mechanism underlying these differences is the criterion or *threshold* for stopping search and generating a price.

Further work on the endowment effect has suggested that process of sampling different aspects of the stimulus over time also leads to differences between buyers and sellers, including the order in which they are considered (query theory, E. J. Johnson et al., 2007; Weber & Johnson, 2011). This suggests that initially considered information interacts with subsequent information search to produce the overall valuation that people place on target items. This idea is reinforced by work that investigates the processes of decision-making and pricing via eye-tracking (N. J. Ashby et al., 2012) as well as the many components of bias that are integrated into these differences such as the length of time of ownership, bargaining advantages, strategic misrepresentation, (Morewedge et al., 2009)

These developments strongly suggest that a complete model of pricing ought to incorporate process-level assumptions in order to account for differences between buying and selling prices. In particular, the multitude of evidence aligns closely with well-established model components used in perceptual and preferential decision-making (Ratcliff et al., 2016; Busemeyer et al., 2019), where choice outcomes are determined by multiple factors including (1) initial biases or starting points, which is likely influenced by factors outlined by Morewedge et al. (2009) including reference / anchor prices (Pachur & Scheibehenne, 2017); (2) an information sampling and accumulation of support for different response options determined by attention, which can incorporate attentional and biased sampling elements (N. J. Ashby et al., 2012; Pachur & Scheibehenne, 2012; E. J. Johnson et al., 2007), (3) and termination of this process and generation of response once sufficient support for one of the response options has been gathered (leading to interactions with response times; N. J. Ashby et al., 2012; Kvam, 2019a). The model we propose here includes each of these elements as cognitive mechanisms that compose the price generation process, integrating a litany of empirical work into a cognitive theory that embodies these critical assumptions.

It is worth noting that Johnson & Busemeyer 2005 created a precursor of such a model by developing the sequential value matching model, which suggested that people generate prices by sequentially searching through different possible price values, then making a response when they found a price that they found equivalent to a target item (gamble stimulus). Critically, this model is motivated by searching for an a point of indifference, and the threshold mechanism controls how much evidence is needed to move from one price to another. Diverging from this perspective, the model we propose here proposes that support for a particular price, rather than indifference, is the mechanism that triggers generation of a price response. Like the model of Johnson and Busemeyer (2005), the proposed model provides a process account of how people make choices and assign prices to gambles, predicting the joint distribution of both price responses and choices. However, it departs in several ways to provide an improved account of pricing, outlined in more detail in the discussion.

The proposed model also takes process models of pricing a step further by predicting the joint distribution of both prices and the response times taken to select them, which has not been rigorously addressed before. In the following sections of the paper, we will show that this model precisely accounts for the changes in the empirical distributions of prices given to a wide range of gambles; differences between buying, selling, and certainty equivalence price distributions; the effect of time pressure on these price distributions; the relationship between the distribution of prices given to individual gambles and binary choices between gambles; and finally the distribution of response times associated with selecting prices. Static and deterministic models are unable to adequately account for any of these empirical distribution properties, suggesting that the class of models incorporating process-level mechanisms for price provide qualitatively superior accounts of how these responses are generated.

The article is organized as follows. First we present experimental results investigating how people's choices and prices vary across a wide range of gambles. In particular, we focus on the distributions of prices and how they change according to gamble attributes, price type (buying, selling, or certainty equivalent), and how prices change when decision makers are subjected to time pressure. Second, we develop a dynamic and stochastic price accumulation model that can account for the most important features of the choice, price, and response time data. Third, we evaluate the quality of the model accounts of the data and examine exactly where static utility models, such as prospect theory, fail. Finally, we draw conclusions from the empirical findings and discuss future directions.

Methods

The empirical study was designed to evaluate the effects of time pressure and price type on the prices people assigned to different gambles. The gambles shown consisted of a payoff ($0-20) and a chance of winning (0-100%). A total of 11 Indiana University students each completed five sessions of the experiment. Each session consisted of 8 blocks of 36 trials, for a total of approximately 288 trials per session and 1440 trials in total. Participants were paid $10 for attending each session, plus a bonus of $0-20 based on the price and choice responses they made during the experiment (outlined below).

Response types. Trials of the experiment were spread evenly across eight conditions: a full factorial design with four response conditions and two time pressure conditions. The four different response conditions were:

1. Buying / willingness to pay, where participants responded to a gamble with the amount of money they would pay to play it;
2. Selling / willingness to accept, where participants gave the amount of money they would accept in order to give up the chance to play the gamble;
3. Certainty equivalent [CE], where participants responded with a price that they believed was equal in value to the gamble (perspective-neutral); and
4. Choice, where participants selected which of a pair of gambles presented on the screen they would prefer to play.

Each trial began when the participant clicked inside a fixation circle presented in the middle of the screen (upper left of Figure 1). On the choice trials, two gambles appeared – one on either side of the screen. Participants entered their response (selected the gamble they preferred to play) by clicking the left or the right mouse button to choose the gamble on the left or right side of the screen, respectively. Response times were recorded from the time the gambles appeared to the time the participant clicked one of the mouse buttons. In the results and modeling, we focus first on more traditional measures of choice and prices, and consider process measures like response times and mouse trajectories later.

On the buying, selling, and CE trials – after clicking the circular fixation – participants instead saw a single gamble appear in the middle of the screen along with a semicircular scale like the one shown in Figure 1 (middle / right panels). They gave their price response using this scale. When their mouse reached the edge of the semicircle, the price indicated by the position of the mouse was shown in parenthesis above the gamble (Figure 1, right panel). They confirmed this amount and entered their response by clicking on the scale at the desired price. Again, response times were recorded from the time the gamble appeared on-screen until the participant made their response by clicking the mouse.

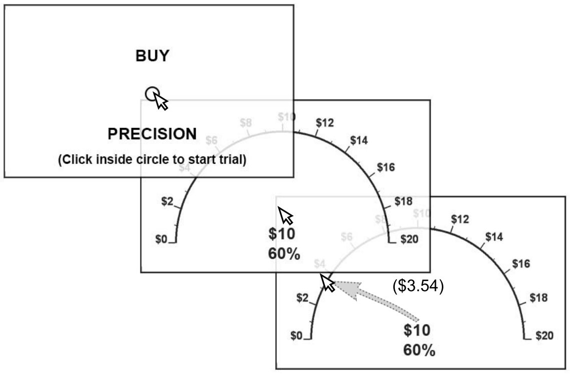

Figure 1. Diagram of the price task. Participants were reminded before each trial about what type of response they were giving (buy / sell / rate) and whether there was time pressure (speed / precision; left panel). They gave their response by clicking on a semicircular scale (middle / right panels).

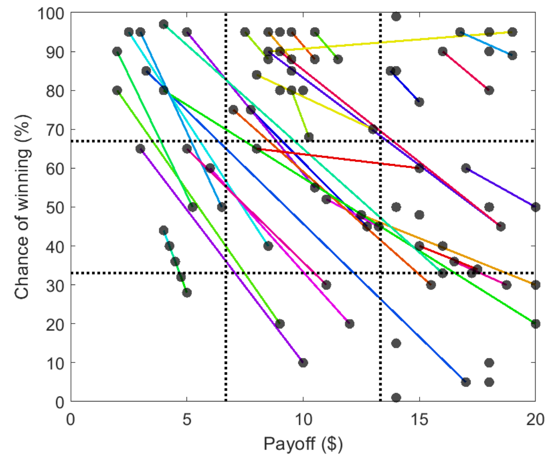

Figure 2. All priced gambles (black dots) presented in the pricing conditions of the experiment, represented in terms of their payoff (horizontal axis) and winning probability (vertical axis). Colored lines connect the pairs of gambles that were also presented in the binary choice condition.

Time pressure. In addition to the four response type conditions, we also manipulated time pressure during the task. This divided trials into two types, speed and precision trials, which were crossed with the response type manipulation for a total of eight conditions. In the speed conditions, participants had to respond in less than five seconds for the pricing conditions (buying / selling / CE) or less than two seconds for the choice condition. They were shown an error message after any trial on which they failed to respond within this time frame. In the precision condition, participants were asked to respond within ten cents of the desired price in the judgment conditions and prompted to make their response carefully in the choice condition.

The directions for the time pressure and response type conditions were given at the start of each block of trials, and participants were reminded which condition they were in by text above and below a fixation circle in the middle of the screen (Figure 1, left panel) at the start of every trial as well. Over the course of the study, participants saw 72 different individual gambles (repeated 3 times each) across pricing trials, and 36 pairs of gambles (repeated 2 times each) in the choice trials (see Figure 2 for the gambles and pairs). Trials were blocked by condition so that buying/selling/CE and speed/precision trials were not mixed together.

Mouse position during each trial was also recorded every 50 ms, allowing us to follow the (x,y) position of the mouse on the screen from trial initiation to final response. Mouse tracking and other methods for process tracing like arm or finger movements have recently been used to provide additional insight into decision processes prior to response (Koop & Johnson, 2011), including preference based choice (Chen & Fischbacher, 2016), simple perceptual decisions (Friedman et al., 2013; Dotan et al., 2018), intertemporal choice (Cheng & González-Vallejo, 2017), and recognition memory (Koop & Criss, 2016). In our case, we can convert the (x,y) position into polar coordinates (r,ϕ), which indicate the distance of the mouse from the center of the screen and the angle of the mouse relative to the scale. Respectively, these correspond to how close a person is to the response semicircle and what price their mouse position indicates at every point during the trial. The mouse data can therefore provide insight into how close a person is to making a response (how much support for a response they have collected) and what price they favor over the course of each trial.

Payment / incentives

In order to provide trial type-consistent incentives, payment was based on the response a participant gave in a randomly selected trial, where each trial was incentivized in a way that was consistent with the type of response they were prompted to give. At the end of each session of the experiment, the random trials were selected from all of the trials participants had undergone, weighted by trial type. To provide an incentive for speed and accuracy conditions, participants were awarded a flat $2 bonus if a speed trial was selected for payment (or $0 if they failed to respond within the allotted time on a speed trial). Precision trials were more likely to be selected (60% precision versus 40% speed), which provided an incentive to be more careful on these trials. Participants were informed on the details of the payment scheme prior to participating, so they were fully aware of these incentives for the entirety of the experiment.

For buying trials, participants would start with a $10 bonus (bringing them to $20 total to accommodate any bid they wished to make), then would have that amount reduced by the amount they bid to play a gamble, provided the bid was high enough. Whether or not a bid was "high enough" was determined by the price's relation to those given in a prior experiment (Kvam & Busemeyer, 2018). If the buying price was above at least 20% of selling prices given to the same or similar gambles in the previous experiment (in the experiment, they received a message indicating that they "could find a seller" for that gamble) it was accepted and the participant would forego the amount that they bid in order to play the gamble. If it was below the selling prices, the participant's bid would not be accepted and they would not get to play the gamble, instead keeping the flat bonus. In this way, they were incentivized to give low but reasonable prices as if they were buying a gamble.

For selling trials, participants would be endowed with a random gamble from the experiment that they had priced. Whether they kept the gamble or sold it depended on the selling price they assigned to it. If the selling price was below at least 20% of the buying prices for the same or similar gambles from the previous experiment, then the selling price was accepted (participants were told they "could find a buyer" for that gamble) and the participant would receive a bonus equal to the selling price they indicated. In this way, they were incentivized to give high but reasonable prices as if they were selling the gamble.

For certainty equivalent trials, participants would receive the gamble with probability p, where p is the percentile of certainty equivalent prices for that gamble – derived from the previous experiment – into which their price fell. Thus, if their CE price was unusually low, they would be more likely to receive the gamble and if their CE price was unusually high, they would be more likely to receive the fixed dollar amount (sure thing). In this way, participants were incentivized to give prices that thought were approximately equal to the value of the gamble, as this was the way to maximize their expected utility from the randomly drawn trial.

If the participant wound up with a gamble at the end of the experiment, they would then play this gamble out by rolling 1-6 dice and winning or losing based on the combination of rolls of those dice. The number and criterion number on the dice rolls was calibrated to the probability of winning the gamble; for example, if the chance of winning was 17%, they would win only if they rolled a 1 on a single die.

Results

The results focus on three main areas: (1) the distribution of prices and reliability of prices across different types of gambles, (2) the effects of time pressure on prices, (3) and preference reversals between choices and prices. Within each of these, we consider differences between types of price responses, as there are frequently interactions between the type of price elicited and the timing or distributions of price responses. Put together, they illustrate the need for a distributional and dynamic model of price, how these distributions change with time pressure, and how prices produce different preference orders when compared to binary choice.

Price distributions

The shape of the distributions of prices is perhaps the most striking element of the results of the experiment. These are shown for three example gambles in Figure 3. The first thing to note is the skew of the distributions for low and high probability gambles: when the probability of winning is low (e.g., 5%, Figure 3 top panel), responses tend to group near \$0 with a long tail out toward the right, giving a heavy positive skew. Conversely, when the probability of winning is very high (e.g., 95%, Figure 3 bottom panel), responses are grouped near the maximum payoff with a long tail toward lower price responses, yielding a strong negative skew.

We tested how the skew changed across conditions and gambles by calculating a nonparametric measure of the skew of the distributions (which is linearly related to Pearson's median skewness coefficient) as $\frac{Mean-Median}{Std.Dev.}$ for each participant and condition (Doane & Seward, 2011; Groeneveld & Meeden, 1984). These participant $\times$ condition skew measures can then be compared as a function of condition and the gamble win probability using Bayesian linear and simple difference models. In all models below, we report the mean effect and 95% highest-density interval (HDI) for the estimates of these effects. All models were fit using JAGS, MATLAB, and MATJAGS, and used 4 chains of 5000 samples each with 500 burn-in samples.

Overall, skewness became more negative as the probability of winning the gamble increased in all conditions, including the buying ($M(b_1) = -0.12$, 95% HDI = $[-.21, -.03]$), selling ($M(b_1) = -.19$, 95% HDI = $[-.28, -.11]$), and CE conditions ($M(b_1) = -0.19$ 95% HDI = $[-.28, -.10]$). As illustrated in Figure 3, this was true for almost every individual participant (bottom panels) as well as the aggregate (top), indicating that this type of effect is not driven by a select few individuals nor an effect of averaging.

These types of distributions might be expected if outcomes were simulated from a binomial distribution with probability equal to the gamble outcome probability – the unlikely outcomes (\$0 in the high-probability gamble, \$18 in the low-probability gamble) would be sampled less frequently, leading to a skewed distribution of samples. This lends support to the idea that decision makers mentally simulate payoffs when considering the value of a particular gamble, as the distributions of prices seem to mimic the distributions we might expect from this kind of mental simulation. Note that these skewed distributions will remain even with large numbers of samples – even with sample sizes near 1000, win probabilities near 1-5% will still yield substantively skewed distributions of samples, and thus price re-

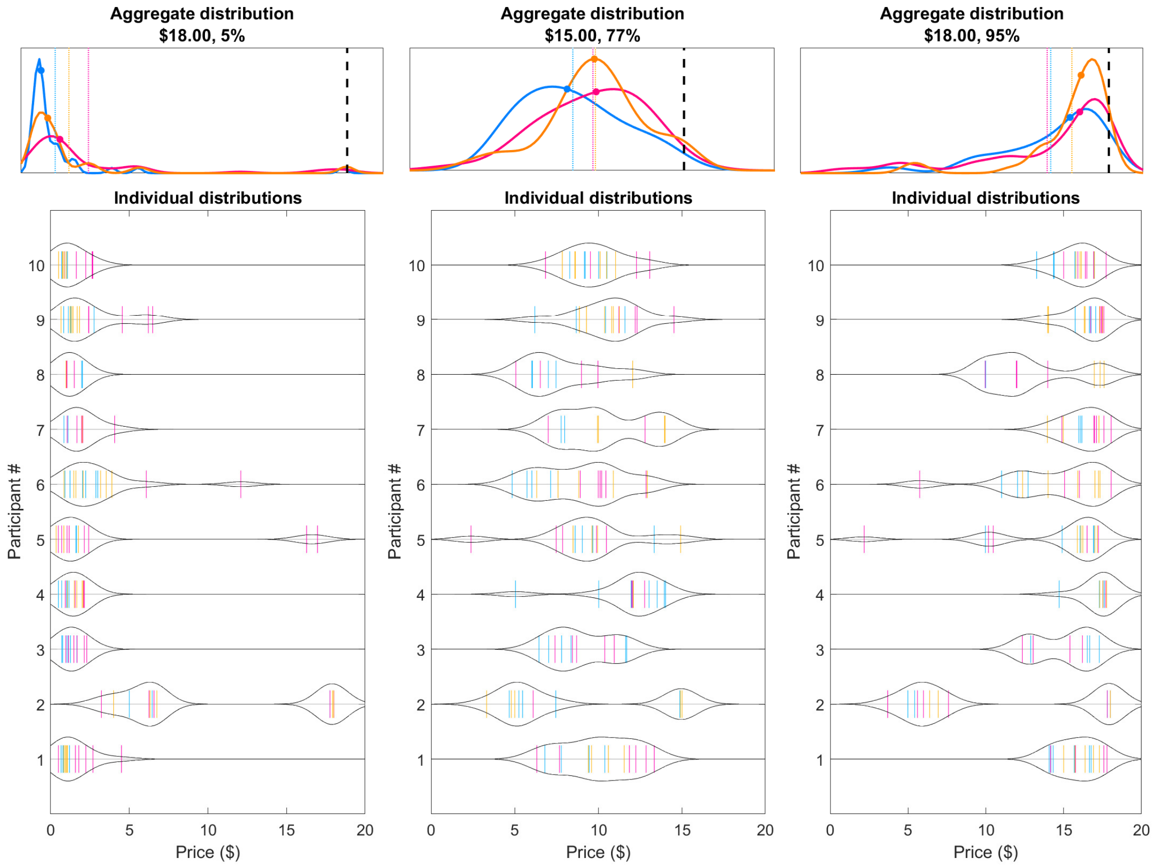

Figure 3. Distributions of buying (blue), selling (red), and certainty equivalent (orange) prices for three example gambles in the data set. In the top plots, solid curves correspond to smoothed aggregate density of responses, dotted / colored vertical lines correspond to the mean of each distribution (same color), and dots on each distribution correspond to the median of each set of responses. The vertical black line indicates the maximum payoff for each gamble. In the bottom plots, individual vertical colored lines indicate responses on separate trials (again color coded such that blue = buying, red = selling, and orange = CE), and black violin plots illustrate the smoothed overall density of these responses for each person and gamble.

sponses. It therefore does not imply undersampling of rare events like that implicated in decisions from experience (Hertwig & Erev, 2009).

The skew of prices also relates to differences between buying and selling prices. There is a consistent mean separation between buying and selling prices, which is a well-established asymmetry (Birnbaum & Stegner, 1979; Carmon & Ariely, 2000; Kahneman et al., 1990; Yechiam et al., 2017) thought to be driven by additional value conferred by ownership or *endowment* (Morewedge et al., 2009; Thaler, 1980). The results of this experiment suggest that the mean differences between buying and selling prices seem to be driven by differences in skew (mean minus median direction) between the two types of prices: buying prices are more positively skewed compared to selling prices ($M(S_{Buy-Sell}) = 0.03$, 95% HDI = $[0.00, 0.06]$). These differences are illustrated in the panels of Figure 3 , where buying and selling prices are positively skewed for low probability wins (top panel); negatively skewed for high probability wins (bottom panel); and selling prices are slightly negatively skewed and buying prices are slightly positively skewed for probabilities near .5 (middle panel).

The difference in skew between types of price may be attributable to a difference in the *starting point* of each price type. Naturally, buyers will want to start with a low price anchor and then increase the price as they consider the possibility of winning high payoffs. Conversely, sellers will want to start high and may come down as they consider the possibility of undesirable low payoffs. This difference in strategic preference distortion is advantageous to each group, and differences in response bias appear to be the most plausible mechanism for the buyer-seller gap (Pachur & Scheibehenne, 2017). The process of initial anchoring and subsequent adjustment would create a long tail to the left in selling distributions (as fewer and fewer sellers are willing to come down in price) and a long tail to the right in buying distributions (as fewer buyers are willing to come up on price). As we show in the section on price dynamics, this may also be related to dynamic properties of these prices that are revealed when decision makers are put under time pressure.

Variance and reliability. The final aspect to note about the distributions of prices is that the high-variance gambles (chance of winning around 50%) resulted in a much wider spread of responses than the low-variance (chance of winning near 0% or 100% gambles. This is illustrated well in Figure 3, where the high-variance gamble in the middle panel results in greater spread in price responses than the low-variance gambles in the top and bottom panels. It seems that a decision-maker's uncertainty about what outcome they might receive – given by the outcome probabilities of the gamble – was translated into greater uncertainty in the distributions of their price responses.

The variance in responses to gambles with high uncertainty is also manifested in the reliability of responses to these gambles. Because each participant saw each gamble a total of five times in each condition, we were able to assess the test-retest reliability of each gamble by examining the correlations between two subsequent responses to the same gamble and condition for each person. Pairs of measurements used in the test-retest correlation were taken from only the nearest successive measurements (e.g., the first and second time a gamble was presented, and the third and fourth time it was presented), which were typically within 2-3 days of one another. Overall there were approximately 200-240 (10 participants $\times$ 6 conditions $\times$ 4 repeats, minus dropped trials) test-retest values for each gamble, yielding a fairly precise picture of the reliability of each item.

The results are shown in the top panel of Figure 4. Reliability in general was poor, with most gambles producing test-retest correlations of around 0.1 to 0.6. Furthermore, reliability was consistently lowest for gambles with a probability of winning close to 50%. Gambles with very low and very high probabilities had reliability as high as 0.8, but those with greater variance (probabilities close to 50%) rarely exceeded test-retest reliability of 0.4.

Greater variance and lower reliability in responses to high-variance gambles is again characteristic of a stochastic response process where outcome uncertainty influences the spread of responses. As with the highly skewed gamble prices, this kind of result would be typical of a binomial distribution of simulated gamble outcomes. If decision makers made their responses on the basis of a relatively small set of mentally simulated outcomes of the gambles, we could expect both highly skewed responses to low-variance gambles and widely distributed responses to high-variance gambles.

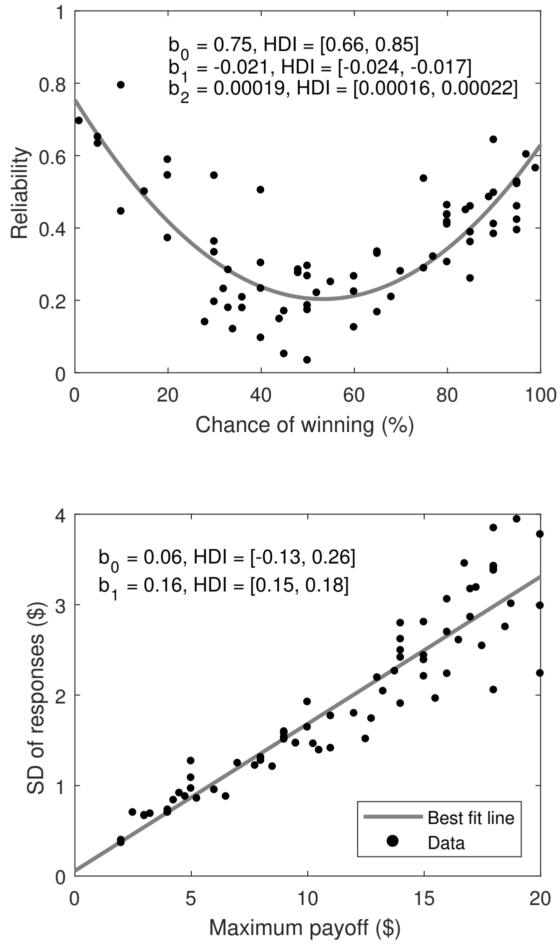

Figure 4. The relationship between chance of winning a gamble and the reliability of responses to that gamble (top), and the payoff of a gamble and the standard deviation of price responses to that gamble (bottom). Each individual dot corresponds to a single gamble, gray lines indicate best fit quadratic (top) or linear (bottom) function relating the variables.

A final note regarding the distributions of price responses is helpful mainly because it helps in building a model of the pricing process later. The bottom panel of Figure 4 shows the relation between the maximum payoff of a gamble and the variance of the distributions of responses to that gamble. Even though the maximum payoff of each gamble was slightly negatively related to probability (see Figure 2), higher payoffs were associated with greater variance in responses. The association between maximum payoff and standard deviation of responses was both strong and positive ($M_{b_0} = 0.16$, 95% HDI = $[0.15, 0.18]$ in a Bayesian regression), indicating that people gave more variable responses as the magnitude of the potential payoffs increased. Generally speaking, the variance in expected payoffs of a gamble increased with the winning payoff magnitude, which results in greater variance of prices for the gamble – the standard deviation of the expected gamble payoffs turns out to equal to

$$p \cdot (1-p) \cdot X \tag{1}$$

where p is the probability of winning, and X is the amount to win.

The skew of responses to low-variance gambles, the difference in skew between selling relative to buying prices, and the high variance and unreliability of responses to high-variance gambles all serve as indicators that there are elements missing from the deterministic valuation models that currently dominate the literature on judgment and decision making. The model we propose below incorporates a stochastic element to the pricing process that is influenced by the outcome probabilities of the gamble, allowing it to account for each of these phenomena in price distributions.

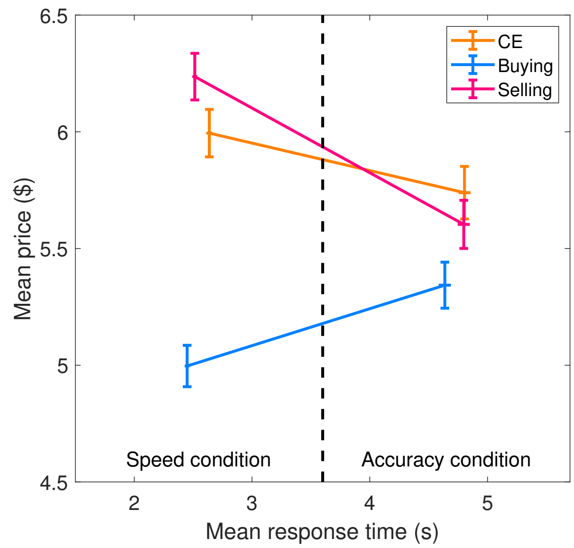

Figure 5. Mean prices for each response type (buying / selling / CE) and time pressure condition. Vertical error bars indicate ± 1 unit of standard error on the mean prices, horizontal error bars (which are all quite small) indiciate ± 1 unit of standard error on the mean response times.

Price dynamics

Another characteristic of valuation models typically used in judgment and decision making is that they do not provide any mechanism to predict differences in price as a function of how much time a decision maker takes to consider their selection. Because there is a fixed function mapping the outcomes and associated probabilities to a value, they would predict that there should be no interaction between the time taken to give a price response and the mean price response. This assumption was tested through the time pressure manipulation in the experiment which showed that price responses changed systematically over time.

Time Pressure. In order to evaluate the effects of time pressure on different types of price responses, we used a hierarchical Bayesian model to estimate the differences between conditions within and across participants. Here we report the overall differences between conditions, comparing each combination of subject and gamble between conditions using paired comparisons (e.g., we compared participant X's responses to gamble Y in the buying / speed condition against participant X's responses to gamble Y in the buying / precision condition). As before, the models were run using MATLAB and JAGS (Plummer, 2003). Code for the model can be found at osf.io/tfm4e/.

The mean pattern of prices is shown in Figure 5. In line with typical endowment effects, we found a substantial mean gap between buying and selling prices ($M_{sell} - M_{buy} = \$0.85$, 95% HDI = $[0.74, 0.97]$. However, when these prices were given under time pressure the mean difference between buying and selling was approximately a dollar ($M_{sell,speed} - M_{buy,speed} = \$1.00$, 95% HDI = $[0.42, 1.60]$). but fell to 34 cents when precision was emphasized rather than speed ($M_{sell,prec} - M_{buy,prec} = \$0.34$, 95% HDI = $[-0.36, 1.04]$). The reduced difference between conditions can be attributed to a tendency of selling prices to decrease in precision conditions relative to speed conditions ($M_{sell,prec} - Msell,speed = -\$0.25$, 95% HDI = $[-0.09, -0.43]$) and buying prices to increase when time pressure is relaxed ($M_{buy,prec} - M_{buy,speed} = 0.31$, 95% HDI = $[0.19, 0.42]]$). As a result, these prices tended to converge when participants were encouraged to consider the gambles more carefully before entering their prices. Note that the size of these increases / decreases will not perfectly match Figure 5 as they are comparing the mean differences within participant / gamble across relevant conditions rather than the difference of means across all responses and participants in a condition.

Oddly, this is at odds with the findings of N. J. Ashby et al. (2012), who found that buyer-seller differences *increased* with deliberation time. We revisit this point in the discussion, as it implicates that the model may need both starting point (anchoring & adjustment) as well as sampling process (attentional components) pieces in order to account for buy-seller differences across different types of pricing situations.

Process tracing. The dynamic nature of price is corroborated by patterns in mouse tracking data, which allows us to explore a potential source of information about the pricing process prior to when a selection was made (Freeman et al., 2011; Schulte-Mecklenbeck et al., 2011). As we suggested in the methods section, the position of the mouse on the screen was sampled at 20 Hz (every 50 ms). The average trajectory of the mouse for each price type is shown in the top panel of Figure 6. Quite clearly, prices for the selling condition tend toward higher prices than those for the buying condition, and CE condition tends to land somewhere in between. However, this only shows the spatial trajectory of these ratings rather than how mouse position changes over time. To examine the temporal properties of price exhibited in the mouse position, we must examine how the mouse position changes across time points.

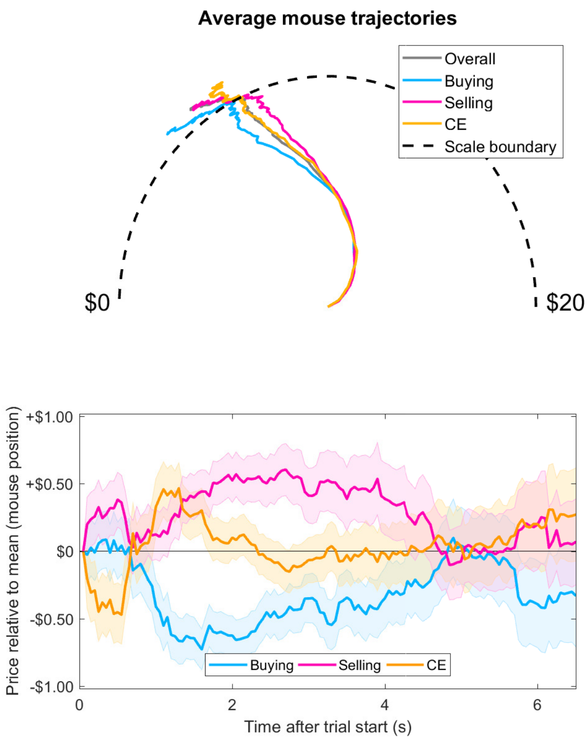

Figure 6. Average trajectory of the mouse on the screen (top panel), and broken down by radial position relative to the average trajectory (bottom panel) for each price type. The filled regions in the bottom panel indicate ± 1 unit of standard error relative to the average trajectory for the corresponding condition.

The task used a radial scale (Figure 1), which meant that the x and y positions of the mouse could be transformed into polar coordinates $r = \sqrt{x^2 + y^2}$ and $\phi = tan^{-1}(y/x)$ for convenient metrics. The coordinate r describes the distance of the mouse from its starting point in the middle of the screen, and the coordinate ϕ describes the angle of the mouse relative to direction $[1,0]$, which is where the response for \$20 is located on the screen. This makes the ϕ coordinate particularly useful, because it is most likely to track the current 'favorite' price across the course of each trial. For example, a ϕ coordinate of 0 corresponds to a favored price of \$20, a ϕ coordinate of $\pi/2$ radians (90 degrees) corresponds to a favored price of \$10, and a ϕ coordinate of π radians (180 degrees) corresponds to a favored price of \$0. While these prices may not manifest across the entire course of a trial, a participant *must* approach the appropriate ϕ angle of their desired response before entering it, simply because that is how the response is entered.

With this in mind, we can roughly track the favored price across the course of a trial as indicated by the mouse position on the screen. The bottom panel of Figure 6 shows the ϕ coordinate in terms of the price it indicates on the scale. The overall mean trajectory is removed so that we can see how buying, selling, and CE trajectories behave relative to one another. As the graph indicates, the favored price inferred from the mouse trajectory increases for selling prices and decreases for buying prices relative to the average for the first 2-3 seconds of the trial, then slowly come back together over time. This pattern suggests that early biases brought on by the price type manipulation tend to wash out over time as participants consider more information about the gamble. It therefore lines up well with the results shown in Figure 5, which also shows a convergence between price types in the conditions with later responses (precision emphasis).

The mouse tracking data alone does not guarantee that a person's true underlying preference state is evolving according to the sort of dynamics shown in Figure 6, but put together with the mean patterns shown in Figure 5 it provides strong evidence that the price type manipulation differentially impacts the dynamics of the pricing process. In particular, both sources of data suggest that buying and selling prices start with a large difference between them and that this gap diminishes over time as a person samples more information about the gamble shown on the screen. As we show in the model, this can be accounted for by treating the price type condition as a manipulation of initial bias and the price accumulation process as one which washes out this initial bias over time.

Preference reversals

Another critical phenomena that hints at richer underlying response processes is that of preference reversals. In some cases, participants will choose gamble A over gamble B when they are side by side in a binary decision, but will price gamble B as higher in value than gamble A when they see them on separate pricing trials (Lichtenstein & Slovic, 1971; Slovic & Lichtenstein, 1983; Grether & Plott, 1979). Typically, past research has found that the higher-probability gamble of the pair (p-bet) is chosen in a binary choice but the higher-payoff gamble (d-bet) is assigned a higher price when pricing the single gambles (Slovic & Lichtenstein, 1983).

A similar finding was strongly supported in our data, shown in Figure 7. For each gamble, we took the mean price

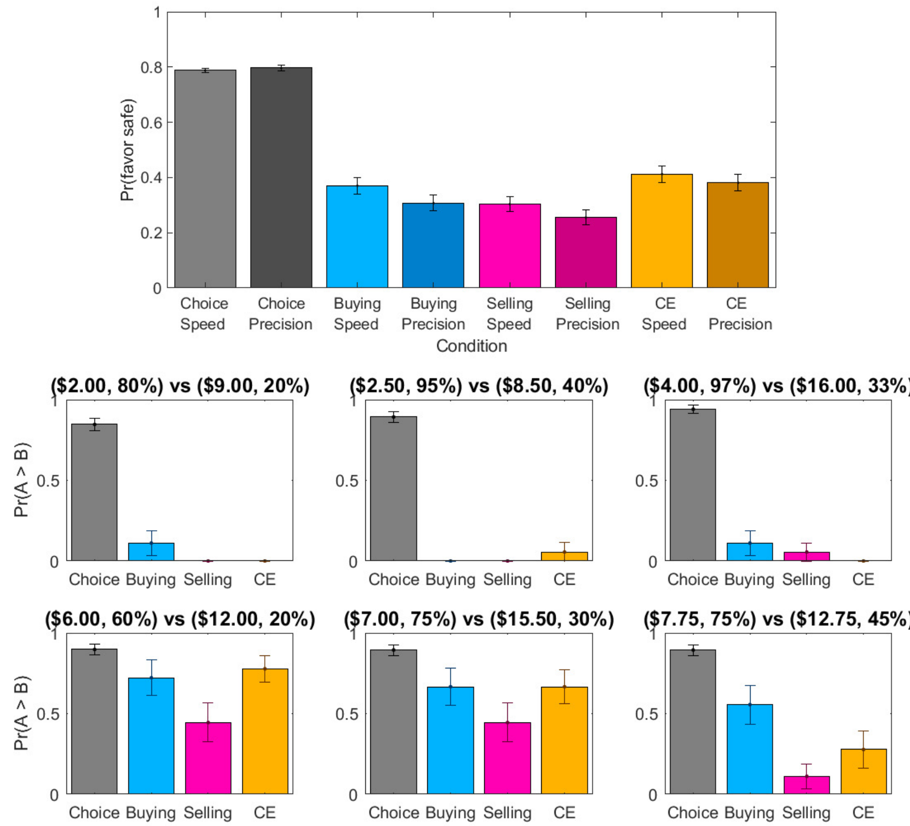

Figure 7. Proportion of participants favoring the safe (high-probability / p-bet) gamble across all conditions and gambles (top panel) and for six example gambles (smaller bottom panels). Preference between gambles is inferred from choice proportions (gray), mean buying prices (blue), selling prices (red), or certainty equivalents (orange) for each participant and then computed on aggregate. Error bars indicate ±1 unit of standard error.

[1] assigned to a gamble by a participant in a particular condition and then compared the proportion of times the mean (or median) price was higher for the p-bet versus the d-bet. As shown, the overwhelming majority of gamble pairs presented in the experiment resulted in participants choosing the p-bet more often in the choice condition when the gambles were presented side by side. However, these same participants tended to assign a higher price to the d-bet when the same gambles were presented on separate pricing trials. Out of the 30 unique, non-dominated gamble pairs presented during the experiment, 22 of them showed preference reversals between choice and pricing. This was true for buying, selling, and certainty equivalent prices – typically all three would show a reversal relative to binary choice, although buying prices tended to be most similar to the pattern obtained from the choice condition.

There was also a slight trend for precision conditions to result in prices that favored the d-bet over the p-bet more often than speed conditions (top panel of Figure 7) , but these differences did not credibly rule out zero so we did not read too far into these differences.

Although they were much more rare than preference reversals between choice and price, there were occasional reversals between buying and selling conditions. An example of this pattern is shown in the bottom right panel of Figure 7, where participants priced the p-bet ($7.75, 75%) higher than the d-bet ($12.75, 45%) more than half the time in the buying condition, but consistently priced the d-bet higher than the p-bet in the selling condition. This type of preference reversal is substantially rarer than reversals between pricing and choice conditions: a switch across 50% occurred on 4 out of 30 unique gambles, and substantial differences between buying and selling prices that did not cross 50% occurred on an additional 2 gamble pairs. These types of reversals have been found before (Mellers et al., 1992; Birnbaum & Beeghley, 1997; Birnbaum & Zimmermann, 1998), but they are even more difficult to explain than those found between pricing and binary choice. In this experiment, most reversals between buying and selling seem to be related to the differences in skew between the two price types, shown in Figure

[1] Analyzing the median prices versus the mean makes essentially no difference in the results or interpretation.

3. The model we present in the next section is able to handle buying-selling reversals by virtue of its ability to produce the different skews for each type of price, but since these were rare we do not focus on them.

Response times

Each type of response – including buying, selling, CE, and choice – exhibited the typical right-skewed distribution of response times. Each of these can be clearly seen in Figure 8. Naturally, mean response times were faster in the speed condition than in the accuracy condition, and the accuracy condition tended to exhibit a longer tail to the RT distribution. Furthermore, responses in the pricing conditions (buying, selling, CE) were substantially slower than those in the choice condition. Despite there being more information to consider in the choice condition – twice as many outcomes and probabilities – the response processes underlying pricing appear to take longer than those underlying binary choice. As we suggest in the modeling section, this may be because participants can decide based on pairwise differences in attributes between alternatives in the binary choice condition, but to come up with a price they instead weight the possibilities of winning and losing the gamble over time as they mentally simulate outcomes.

The differences in distributions of response times between price conditions were fairly small, as reflected by the x-position of the conditions in Figure 5, but were quite substantial between speed and accuracy conditions. This seems to indicate that they likely shared many properties in terms of the underlying response processes. In fact, a model with only two response thresholds – one for speed and one for accuracy – provided reasonable fits to the data. The model predictions aggregated across participants are overlaid onto the RT distributions shown in Figure 8, while the model predictions for each individual are presented in supplementary figures at osf.io/tfm4e/. In the next section, we discuss how these predictions were generated.

While the response time distributions are not especially unusual as far as decision tasks are concerned, they do constitute another major barrier for the static, deterministic models like expected utility and prospect theory. Because these theories do not provide a generative model that explains the process of how a decision maker comes up with a price – instead describing the outcome of the process *as if* the decision maker is weighing utilities and probabilities (Berg & Gigerenzer, 2010) – they are unable to provide sufficient descriptions of process-level measures such as response times or the mouse tracking data shown in Figure 6. The model we describe next provides a set of cognitive mechanisms for how these prices are generated, and is thus able to produce response times that match the real data.

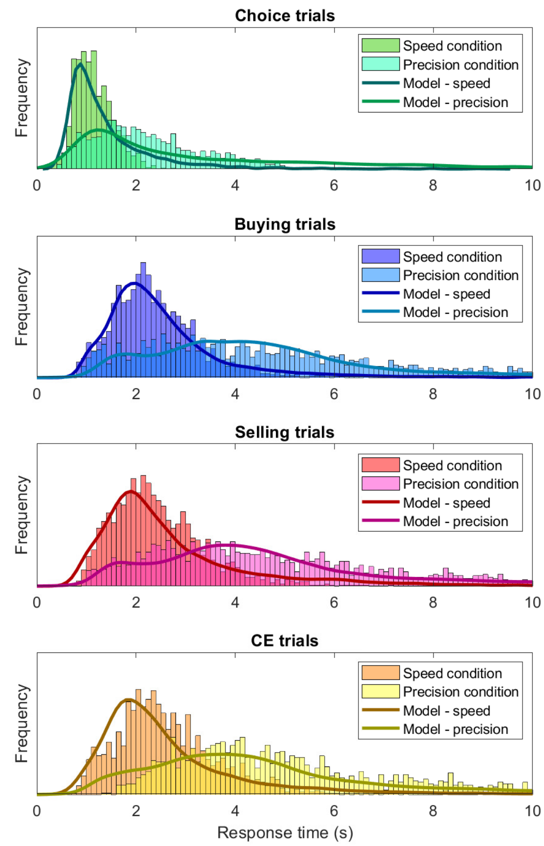

Figure 8. Distributions of response times across all eight conditions (histograms) and price accumulation / decision field theory model fits to response times for price conditions (solid lines).

Modeling

We have highlighted a number of areas where utility and prospect theory models are insufficient to capture important aspects of the empirical data. Notably, they fail to predict 1) skewed distributions of prices that differ for price type and probability of winning a gamble (Figure 3); 2) the greater variance and unreliability of prices assigned to gambles with a probability of winning near 0.5 (Figure 4); 3) the effect of time pressure on the gap between buying and selling prices (Figure 5; 4) the convergence between different types of prices over time (Figure 6); 5) preference reversals between choice and price, and buying and selling (Figure 7); or even 6) the simple response time distributions (Figure 8).

It is critical to note that these effects are individual-level ones, not simply the result of performing analyses on aggregate data, which can show patterns that are not present in any particular individual (F. G. Ashby et al., 1994; Regenwetter & Robinson, 2017). Each of these effects was observed in the majority of individual participants, which is shown in the online materials at osf.io/tfm4e/. In all of the modeling analyses presented below, we also fit individual-level data to avoid the issues associated with drawing conclusions from fits to averaged data (Estes & Maddox, 2005), although for illustrative purposes the figures aggregate data and fits across individuals.

Given the strength of the distribution-level and temporally-dependent effects in the empirical data, a viable model of pricing should be both dynamic and stochastic. To develop such a model, we examine a variant of continuous-response cognitive models that predict joint distributions of responses and response times. In addition, it is imperative to compare the proposed dynamic and stochastic model to random utility extensions of the deterministic models with respect to their ability to account for the distribution of choices and prices (ignoring response times).

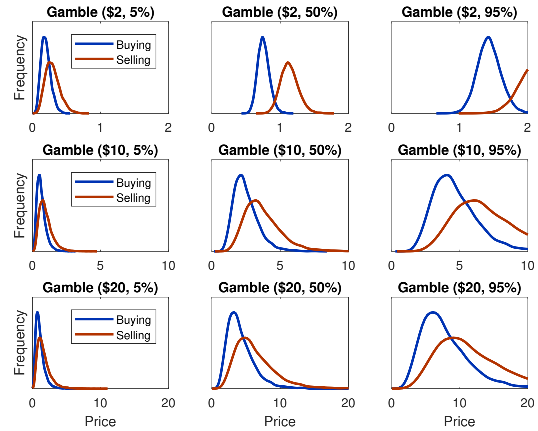

Figure 9. Predicted distributions of WTA (selling) and WTP (buying) prices generated by varying the parameters of prospect theory across trials according to a normal distribution.

Prospect theory and random utility

Thus far, we have largely dismissed prospect theory as being capable of generating the observed distributions of prices because it is inherently a deterministic theory of price. However, variability can be introduced to the model through several avenues. The two most reasonable sources of error would be random variation in the utility each participant assigns to dollar values – referred to as *random utility* models – or random error associated with the final value representation or motor response. While motor variability undoubtedly plays a part in the response processes, it offers little to no benefit in terms of capturing the skew of price distributions. The typical functional form of these error functions (normal, with mean 0 and variance σ^2) would produce a symmetric distribution of prices centered on whatever the "true" value of the gamble derived from prospect theory or expected utility. Clearly, this is insufficient to predict the probability-dependent and type-dependent distributions of prices like those shown in Figure 3.

The other reasonable possibility is to introduce cross-participant or cross-trial variability in the prospect theory parameters: utility power parameter α and probability weighting parameter γ. While we could also consider two-parameter probability weighting functions, this makes little differences in the results. It is not immediately clear how variability in these parameters should affect distributions of responses or what shape the resulting distributions would be. The most common approach would be to allow them to vary randomly according to a normal distribution. However, this is only able to create right-skewed distributions of prices, as shown in Figure 9.

Since normally distributed parameter variability seems to fail here, we tested an even more flexible implementation of prospect theory where parameters are permitted to vary according to a beta distribution. In this model, the utility parameter and probability weighting parameter were allowed to vary from trial to trial as

$$\alpha \sim Beta(A_\alpha, B_\alpha) \quad (2)$$

$$\gamma \sim Beta(A_\gamma, B_\gamma) \quad (3)$$

Furthermore, we allow these parameters to vary across price type conditions – going a step further than the traditional assumption of loss aversion / endowment affecting only the utilities and permitting price type manipulations to affect both utilities and probability weights.

We included two additional parameters in the prospect theory model in an effort to allow it the best chance at fitting price distributions. The first of these was a motor error parameter σ_{err}, which allowed responses generated based on the utility and probability weighting parameters (drawn on a particular trial) to additionally vary according to a normal distribution around the selected price, $resp \sim N(0, \sigma_{err})$. This accounted for motor processes and their effect on prices, and allowed us to separate random variation in task-relevant parameters from random variation in task-irrelevant ones (motor variation).

The second additional parameter was one that the prospect theory model shared with the price accumulation model, which was driven by a subset of responses that were made at the maximum payoff in conditions where this was not sensible. Several of the participants – particularly, participants 1, 2, and 9 – had a tendency to respond at the maximum payoff for a subset of the trials, even when the gambles featured on those trials had a very low probability of winning. This can be seen in Figure 3, where there is a small bump at the maximum payoffs in the first two panels (at $18 in the top panel, and \$15 in the middle panel). These outliers create substantial problems for the model, because they should be extremely unlikely in low- to medium-probability gambles. However, it is difficult to justify a specific rule that would be able to systematically exclude these trials, and it is entirely likely that there were some high-probability trials that contained these outliers as well. Instead, we included a contaminant process that hypothesized that participants would respond at the maximum payoff with probability p_{max}, and follow the prospect theory (or price accumulation model) prediction with probability $1 - p_{max}$. This generates a probabilistic mixture of the 'normal' response process and the contaminant process, allowing the model to capture these abnormal high price responses without sacrificing the overall quality of fits to the price data.

In total, this leaves the prospect theory model with 14 free parameters: 6 utility parameters (A_α and B_α for each of the three pricing conditions), 6 probability weighting parameters, σ_{err}, and p_{max}. This results in an extremely flexible model that is intended to giver prospect theory the best chance at accounting for the distributions of prices generated across conditions.

Price accumulation model

In response to the issues that have been identified in static and deterministic models of decision making, cognitive models incorporating process-level mechanisms have been applied to explain how preferences are formed. Many of these models take the view that preference is constructed as evaluations are accumulated over time, usually as a person samples the attributes of competing choice options (Busemeyer et al., 2019)

These models quantify the support for a particular choice option in terms of accumulators or a relative balance of support that describe how their (relative) preference for items change over time. These models provide excellent accounts of responses and response time distributions in both inferential and preferential choice (Ratcliff et al., 2016; Busemeyer et al., 2019). What is different in the present situation is that we are interested not only in discrete choice but also in continuous prices. This prevents the previous choice models from being directly applied, as they mainly predict choices between 2-3 options (or what relative preference judgments between two options should be; see Bhatia & Pleskac, 2019). This leaves the substantial task of model development for pricing scenarios like the experiment presented above, where participants can make any response between \$0 and \$20.

Here we apply an accumulation framework based in part on these preference accumulation models, which views pricing as a selection among a large number of possible responses (dollar / cent values). Recent developments in computational models of cognition have expanded theories of decision making to cases where responses can fall anywhere along a continuous range of potential responses (Kvam, 2019b; Smith, 2016; Ratcliff, 2018). Although most of the applications have been to perceptual decisions like orientation and color selection, Kvam (2019a) sets out a more general modeling framework that can be applied to preferential choices between different responses as well, such as selections along a range of prices. In this framework, each alternative is represented as a direction in a multidimensional space, where the angles between alternatives in the set correspond to similarity relations between them (as in latent semantic and cosine similarity models Bhatia, 2017; Furnas et al., 1988; Landauer & Dumais, 1997; Pothos et al., 2013). In the case of prices, this naturally forms a continuum of directions describing different possible price responses, where prices that are similar to one another (e.g., \$18 and \$19) are set closer together than those that are very different (e.g., \$18 and \$1).

Because the set of price responses composes a very simple continuum of values, they can be arranged in two dimensions as shown in Figure 10. For example, we might set a response of \$0 at 0 degrees, \$10 at 45 degrees, and \$20 at 90 degrees.[2] However, such a scale presupposes that the similarity of \$0 to \$10 is the same as the similarity of \$10 to \$20. It is well-known via studies of numerosity and number representation that human decision makers are less able to discriminate between large numbers / values relative to small ones (Feigenson et al., 2004; Longo & Lourenco, 2007), yielding a relationship between actual and perceived value that approximates a power function (Krueger, 1982). Consistent with this scaling, many applications of utility theory rely on a power function to represent utility of monetary values (e.g., Kahneman & Tversky, 1979). Naturally, we want to incorporate this diminishing sensitivity to increasing dollar values as a fundamental component of our model.

The most straightforward way to incorporate diminishing marginal sensitivity to dollar values is to build a power function into the representation of alternatives. Rather than spacing the set of alternatives uniformly across $[0, \pi/2]$, we should have larger values grouped closer together than smaller values to represent their greater representational similarity (more similar utilities). To perform this transformation, we take the initial position of a particular alternative, compute its utility according to the power function, and divide by the utility of the maximum value on the scale (\$20) to obtain its position along the scale as a proportion of the maximum. Then we multiply that value by $\pi/2$ (in radians, or 90 degrees) to obtain the angle of the alternative relative to $[1, 0]$. Thus, the angle ϕ assigned to a dollar value x is given as

[2] There is nothing that necessarily requires the maximum and minimum values to be represented orthogonally, but doing so results in natural high and low anchors that are mutually exclusive and do not provide support for one another, as the cosine between the max and min will be zero.

$$\phi = \frac{\pi}{2}\frac{x^{\alpha}}{x_{max}^{\alpha}} = \frac{\pi}{2}\left(\frac{x}{20}\right)^{\alpha} \quad (4)$$

Of course, the use of \$20 as the upper anchor / maximum value is driven by the experiment design, but we can easily substitute other values in for x_{max} as the design or set of gambles demands.

This results in a revised scale like the one shown at the bottom right of Figure 10. This transformation re-mapping the alternatives can be flexibly re-computed for different values of α, allowing it to be calculated easily as α varies across individuals or as α is estimated in the model. This transformation actually creates the effect of payoff magnitude on response variability observed in the bottom panel of Figure 4: because higher prices are grouped closer together in representation space, the same variability in cognitive state will lead to greater variability in responses that are higher on the scale relative to responses that are lower on the scale. Thus, we should observe greater variability in price responses when the expected utility of an option is higher.

Once the different price responses are represented as directions (vectors), a person's support for different prices can then be represented as a point in this 2-dimensional space. The component of their state along a vector corresponding to an alternative describes the support for that price at any given point in time. Therefore, as the preference state moves through the 2-dimensional space in which prices are represented, so too does the support for the various price responses. Once sufficient support for a particular price response is generated (the component of the state along one of the price vectors exceeds a threshold θ), that price is selected and entered. This forms a (quarter-)circular response boundary, where a person continues to consider information about the item in front of them until their preference state crosses the edge of the circle, at which point the angle of the state relative to the origin determines the response.

Formally, the model specifies a starting point for each trial, specified by two free parameters. The first parameter depends on the buying, selling, or CE condition. It is specified by s_β, which determines the initial bias in price before the gamble probabilities are considered. The initial price that a person favors is given as $s_\beta \cdot \pi/2$ (from Equation 4), and serves as a proportion of the maximum gamble payoff that a participant is initially willing to consider. It sets the angle of the starting point for the price accumulation process, such that larger values of s_β result in an initial bias toward to give greater price values and smaller values of s_β result in an initial bias to give smaller price responses. This parameter naturally varies according to the type of price a decision maker is asked to give as a fraction of the maximum payoff – for example, a decision maker may have a high value of $s_\beta = .9$ for selling prices or a low value of $s_\beta = .2$ for buying prices. This setting of the start point is motivated by the finding that response bias seems to be an important component of the endowment effect or buyer-seller gap (Pachur & Scheibehenne, 2017).

The second start point parameter specifies the strength of this initial belief as a uniform distribution $Unif(0, s_v)$, as is typical of evidence accumulation models of choice (Ratcliff, 1978; Brown & Heathcote, 2008). A higher s_v will on average result in more stubborn initial biases, making a decision maker stick with prices near their initial point even as they consider attributes of the gamble that conflict with this price. A lower s_v will allow a decision maker to be more strongly influenced by the objective properties of the gamble, allowing support for different prices to vary more flexibly as the decision maker considers different values they could assign. Each parameter has a straightforward interpretation in terms of polar coordinate: s_β determines the angle of an initial price bias relative to the origin (degree of bias), while s_v determines the radius of the starting point (strength of the bias). Participants will naturally vary in both the prices they are willing to pay / accept before considering the gamble as well as the strength of their convictions about these prices, so both starting point parameters are set as individual differences and fit as a free parameter for each person. Additionally, the starting point bias s_β is allowed to differ between buying, selling, and CE conditions as s_{buy}, s_{sell}, and s_{CE}.

Over time, the initial price a person is willing to give will be adjusted as they consider the payoffs and the probabilities of the gamble. The model suggests that a person sequentially updates their initial valuation by mentally simulating the potential outcomes of the gamble. As they think about receiving an outcome, their representation of the value of the gamble moves toward that outcome. For example, say a person is considering a gamble with a 50% chance of winning \$15. Half the time [3] they think about winning \$15, and half the time they think about winning \$0. When they think about winning \$15, their state moves in direction $\mathbf{v}_{15}$ (perhaps at 70 degrees, for example, depending on the utility representation yielded by α and Equation 4) toward high prices, or for gamble outcomes where the high payoff is lower, they will step at an angle determined by the location of that payoff given by Equation 4. When they think about receiving \$0, their state moves directly rightward toward the lower prices, stepping in direction $\mathbf{v}_0 = [1, 0]$. This model can be thought of as a dynamic variant of anchoring and adjustment models (Goldstein & Einhorn, 1987) – the initial price, impacted by the maximum payoff and the price type, is adjusted according to the potential outcomes of the gamble and their likelihoods as a person mentally simulates the gamble outcomes.

[3]For simplicity, we assume no probability weighting in the mental simulation, as this does not appear necessary for high-quality model fits. Of course, it is possible that probability weights may become important or useful in building future models of pricing so we leave the possibility open.

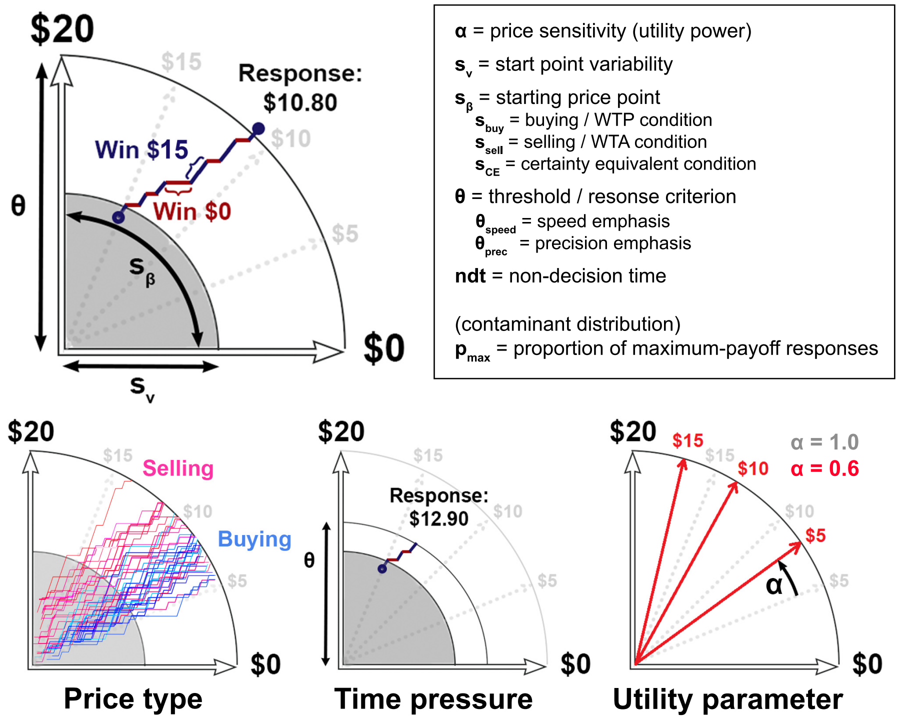

Figure 10. Diagram of the price accumulation model and the meaning of each parameter (top), showing accumulation of support for different prices that could be assigned to the gamble ($15, 55%). Differences between buying and selling (and CE) prices are influenced primarily by the start price bias s_β (bottom left), while difference between speed and precision conditions are influenced mainly by the threshold θ (bottom middle). Individuals also vary in their utility parameter α, which determines how similar the representations of low prices versus high prices are to one another (bottom right).

The model would be easily extended to situations with multiple gamble outcomes. In these cases where there are probabilities $p_1, p_2, p_3, \ldots$ of receiving outcomes $x_1, x_2, x_3, \ldots$, the probability of stepping toward a price x_i would be given by the corresponding p_i. At each step, a multinomial random variable would be drawn (parameter $n = 1$) to determine which outcome of the gamble is sampled and thus which direction the accumulation process should step.

This sequential updating process leads the state to carve a trajectory through the price representation space, as shown in Figure 10. This arrangement allows simulated outcomes to generate support for multiple prices that are consistent with that payoff – for example thinking about winning $15 might simultaneously generate support for several high-price responses like $13.50, $16, or the other surrounding values as it steps toward v_{15}. The time between these steps or the step size can be fixed at an arbitrary value in order to set the scale of the model – in our experiments, we fix the step size at 0.03 and average step time at 30 ms to provide a sufficiently fine-grained approximation for the random walk without diminishing computational efficiency to where the model took too much time to simulate.

Once support for any of the prices exceeds θ, the decision-maker responds with the corresponding price. The critical value corresponds to the amount of consideration the decision-maker puts into the incoming information before making a decision. As a result, θ impacts the amount of time it takes a person to give their prices: lower θ means they will consider the gamble attributes less, generating faster responses and giving the initial bias more sway over final prices. Higher θ means that a person will give more consideration to the gamble attributes, resulting in them taking more time to make their response and ultimately reducing the impact of their initial biases. Given its parallel role in judgment and binary decision tasks, we should expect the threshold to be higher in a precision-emphasis condition than in a speed-emphasis one. In the model, we therefore allow for two separate thresholds for each individual: one for the speed condition (θ_{speed}) and one for the precision condition (θ_{prec}).

The model uses one more parameter to describe the distributions of prices. This is the same mixture parameter p_{max}, described above in the section on prospect theory. The value of p_{max} specifies the likelihood that a participant responds with the maximum payoff for a gamble rather than going through the mental simulation and accumulation process. At this point, the price accumulation model possesses all of the parameters (6) necessary to predict distributions of prices. The parameters relevant mainly to predicting response times and dynamic properties of the model – described next – are fixed for the comparison with prospect theory.

So far, all of the parameters have described the decision processes that go into pricing, but there will naturally be some time devoted to looking at the gamble and selecting the price once the decision maker has arrived at a dollar value for their response. This is quantified by the final parameter describing non-decision time *ndt*, which quantifies the average amount of time participants take on each trial for response processes unrelated to the decision component. It is fit as a free parameter for each person and not permitted to vary across conditions or gambles. Although it is possible that there are variations in non-decision time across trials and conditions as a function of time pressure or motor preparation time (MacKenzie & Buxton, 1992; Crossman & Goodeve, 1983; Donkin et al., 2009), these can be expected to be relatively small compared to the scale of the RTs (usually 1-10 seconds) observed in the data (Figure 8).

Formally, the response is given by the angle at which the response hits the boundary and the response time is given as a random variable RT that is based on the number / length of steps it took to reach the threshold. To preserve the Markov property of the random walk, we assumed the time for each step was exponentially distributed, so that the expected time to the next step was unrelated to the time since the last step. The distribution of each t_{step} was exponential with rate parameter $\lambda = 30$ milliseconds (fixed rather than estimated, in order to set the scale of the random walk). Put together, the amount of time it takes a process to reach the threshold is given by adding up a number of these t_{step} equal to the number of steps it took to finish n_{step} (recall that the process moved 0.03 units each step). As the sum of several exponential random variables, the response time for a particular trial is therefore equivalent to a gamma random variable where $RT \sim Gamma(n_{step}, t_{step})$.

In total, this leaves us with 9 free parameters to predict the joint distributions of prices and response times across 432 combinations of gamble and price condition (72 gambles × 2 time pressure × 3 price types) for each participant in the experiment.

Model comparison and fit

The first item of business is to compare prospect theory and the price accumulation model. However, because prospect theory does not predict response times, we must use a two-step procedure to estimate the price accumulation model: a first step to apply the model to prices alone (to allow the comparison with prospect theory), and a second step to apply it also to response times. The first step of applying the model to the distributions of prices was done by ignoring *ndt* and fixing θ_{speed} and θ_{prec} to 2 and 4, respectively. The six remaining parameters – utility power / representatation parameter α, start points for each of the three conditions s_{buy}, s_{sell}, and s_{CE}, contaminant p_{max}, and the start point variability s_v – were estimated freely from the data for each participant.

This two-step procedure also allows us to compare performance of the price accumulation model on *just* the price distributions against models that cannot predict response times (which are arguably incomplete for this reason). We use this to compare the price accumulation predictions against those of a prospect theory model that includes a specialized between-trial parameter variability mechanism to attempt to produce the skew of price distributions.

Comparison

The prospect theory model contained 14 free parameters, including 12 parameters for the beta distributions of utility and probability weighting functions across conditions (3 instances of A_α, B_α, A_γ, and B_γ), plus a motor variability parameter and p_{max}. The price accumulation model contained only 6 free parameters, including a lone utility α, start point variability s_v, start point biases for the three conditions s_{buy} / s_{sell} / s_{CE}, and p_{max}. The thresholds for speed and accuracy conditions were fixed to 2 and 4, respectively, to set the scale of the price accumulation model – because they trade off with other parameters when only prices and not response times are used, there is not much to be gained from estimating them freely.

Both models were evaluated using Bayesian methodology based on a Hamiltonian Markov chain Monte Carlo sampling procedure. The procedure used 4 chains of 5000 samples each. Each participant was estimated separately, allowing us to examine individual differences between people in posterior model parameter estimates. The priors on the beta parameters for the prospect theory model were all set as uniform distributions on Unif(0,30) [4] the motor variability prior was set as an exponential distribution Exp(1), and p_{max} for both models was set as a uniform on (0,1). The priors for the price accumulation model were $\alpha \sim$ Normal(.9,.3), all s_β were uniform on (0,2) and s_v was uniform on (0,1). Note that s_v was fit as a fraction of the threshold θ to avoid situations where the start point exceeded the choice boundaries.

We used the maximum a posteriori [MAP] values of each parameter (maximum likelihoods) to generate posterior predictions and compare the models to one another. The results

[4] As shown later, the results of the model comparison are so extreme that using more constrained priors will not be enough to help the prospect theory model, and in fact this appeared to be a largely reasonable range for the priors based on the outcome of the prospect theory parameter estimates, as the posterior means lined up with the prior mean as well as those found in past work (such as Nilsson et al., 2011)

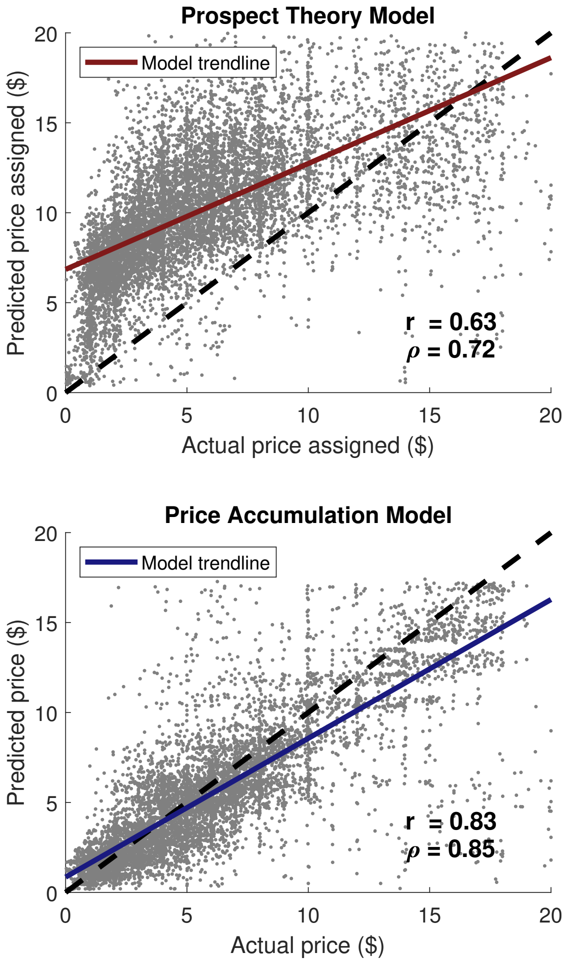

Figure 11. Observed (x) versus predicted (y) price based on the maximum posterior parameter values of the models. Linear trend lines are included to illustrate model performance relative to perfect performance (dashed black line)

are shown in Figure 11. For each response made by participants, we generated 11 predictions from each model and took the median price predicted by the model. The actual price is shown in the x-axis, while the prediction derived from the model is plotted on the y-axis.

In general, the prospect theory model (top panel) tended to overestimate the prices that would be assigned to low-payoff gambles, resulting in over-prediction of prices on the low end of the scale. This was because of the difficulty it had in predicting the skew of the distributions – even with the trial to trial variability specified by beta parameters, prospect theory could not pick up both the left-skew in price responses for high probability gambles and right-skew for low probability gambles simultaneously. It had particular trouble with the left-skew of low-payoff, high-probability gambles (p-bets): while prospect theory can capture the right skew of distributions through variability in utility parameters (Figure 9), it appears to have trouble capturing the left-skew of high-probability gambles even when its parameters are permitted to vary according to a skewed beta distribution.

Interestingly, this over-estimation is actually in line with preference reversal phenomena. Prospect theory would not predict reversals because the same valuation process is used in both pricing and decision. The over-pricing of p-bets would allow it to capture choices better, as participants tend to choose p-bets over d-bets in binary choice. But in the case of pricing, such a tendency results in poorer performance.

Conversely, the price accumulation model does extremely well in predicting the prices assigned on each trial, shown in the bottom panel of Figure 11. It has a slight tendency to underestimate the prices assigned to high-value gambles, perhaps due to influence of the utility parameter α, but ultimately it is hard to point out any areas of clear, systematic misfit. With a value of $r = .83$, it appears to be accounting for almost 69% of the variance in prices elicited over the course of the experiment.

Table 1
Bayesian information criterion (BIC) and log likelihood (LL) for each individual based on maximum a posteriori estimates from the prospect theory model (PT) and price accumulation model (PA)

Participant #	BIC_{PT}	BIC_{PA}	LL_{PT}	LL_{PA}
1	87523	38257	-43718	-19107
2	77748	37465	-38831	-18711
3	66954	32926	-33435	-16442
4	83490	32203	-41702	-16080
5	82928	32099	-41421	-16028
6	89239	33108	-44576	-16532
7	69150	24626	-34533	-12292
8	77767	38623	-38841	-19290
9	88291	61986	-44102	-30971
10	90041	69330	-44977	-34643

This difference between the models in apparent performance is also clearly reflected in model fit statistics. The BIC difference between the models for each participant, calculated based on the maximum a posteriori parameter estimates, is given in Table 1. For each participant, this BIC difference is at least 26000, reflecting overwhelming evidence for the price accumulation model over the random utility prospect theory model. This is not simply due to the price accumulation model having fewer parameters, either: the difference in raw log likelihoods (which can be computed by summing the fourth column and fifth column for prospect theory and the price accumulation model, respectively) is

206040 favoring the price accumluation model. [5]

Overall, the model comparison strongly favors the price accumulation model. Whereas random utility prospect theory tends to overestimate the prices assigned to low-value gambles (especially p-bets), the price accumulation model provides a solid account of the variation in prices across all conditions and individuals. In the next section, we examine some of its posterior predictions, including distributions of prices, response times, and binary choices.

Price accumulation performance

We noted in the previous section that fitting of the price accumulation model as a multi-step process. For the purposes of the model comparison, only six parameters of this model were relevant. However, it is able to predict much more data than prospect theory. With the addition of three parameters, it will also predict the joint distribution of prices and response times on the task. These include the two threshold parameters that were fixed in the previous comparison (θ_{speed} and θ_{prec}) along with a non-decision time parameter ndt, which indexes the amount of time taken to encode the gamble stimulus and enter a response once it had been determined through the price accumulation process.

To fit the model to response times, we fixed the start point and sensitivity parameters at their maximum likelihood values from the price-only model and allowed ndt, θ_{speed}, and θ_{prec} to vary freely when estimating the distributions of response times. This ensured that the parameters relevant to the distributions of price were given priority over those relevant mainly to distributions of response times, although the model ultimately provided excellent fits to both outcomes (Figures 8 and 12) so this was ultimately not a substantial concern.

Because there is not a straightforward analytic solution to the likelihood of the model, we used kernel-based probability density approximation and approximate Bayesian computation to generate a likelihood from simulated data (Holmes, 2015; Turner & Sederberg, 2012; Palestro et al., 2018). For each response in the data set, we simulated 50 trials from the model under the same conditions (same gamble, time pressure, price type), and then computed the likelihood of all responses made under that condition by putting together all the simulated trials from that condition and passing an optimized kernel density estimator over the simulated trials. This procedure matched the number of simulations to the number of corresponding trials, ensuring that a representative sample of model simulations was produced on each likelihood approximation. It also apportioned the computational effort according to the importance (number of data points) of each condition and gamble.

Otherwise, the same procedure we outlined above was repeated to fit the response time distributions, with the priors of $\theta \sim Uniform(.1,5)$ and $ndt \sim Uniform(0,2)$. The resulting modal posterior (maximum likelihood) estimates are given in Table 2.

There are a number of notable properties of these parameter estimates. First, nearly all participants (with the exception of participant 4) had price sensitivity parameters α that correspond to risk-averse utilities. This indicates that most participants represented larger prices as more similar to one another than smaller prices, and results in greater variability in prices when the potential payoffs are greater. It also indicates, if these parameters translate to binary choice, that most participants would select a safe option over a risky option. In the next section, we use these parameter estimates to make an out-of-sample prediction for the choice condition to show that it results in preference reversals.

Another notable but expected property is the consistency with which initial bias was estimated as higher in the selling condition than in the buying condition ($s_{sell} > s_{buy}$). The difference in starting points in the model is the source of the buying-selling gap (Birnbaum & Zimmermann, 1998) and results in the differences in skew between the two distributions of prices. These resulting distributions are shown in Figure 12 – the model is readily able to capture the essential properties of prices, including right-skew for low-probability gambles, left-skew for high-probability gambles, greater right-skew for buying prices, and greater variability for more uncertain gambles (with outcomes near 50%).

The difference between starting points for buying and selling prices also indicates that initial prices for selling are influenced heavily by the maximum payoff, as the start point is a multiplier on the maximum payoff: in essence, in a regression framework it can be viewed as the slope of the effect of maximum payoff on initial price. Conversely, buying prices are less strongly influenced by the maximum payoff, consistent with prior work examining the asymmetry between buying and selling (Carmon & Ariely, 2000). Furthermore, the starting point for certainty equivalents is much closer to that of selling prices than buying prices, conceptually in line with the finding that selling prices are less ‘biased’ (Yechiam et al., 2017) and thus closer to what should be a perspective-neutral condition (CE).

The positions of the initial states will also determine how the mean price changes over time (i.e., under time pressure). If a starting point is especially high, the corresponding type of price will start high and come down, on average. Conversely, if a starting point is especially low the corresponding type of price will start low and come up as a decision maker considers the outcomes of the gamble. This naturally results in the converging prices shown in Figure 5. While the model

[5] We also fit a prospect theory model using the more traditional Gaussian random utility rather than beta random utility (modeled after Nilsson et al., 2011; Scheibehenne & Pachur, 2015, although not using a hierarchical implementation so as to match the structure of the other models tested). This model fared overall worse than the beta model (overall BIC = 591420, $r = .58$, $\rho = .58$)

Table 2

Modal (maximum likelihood) posterior model parameter estimates for each participant. Parameters most relevant to response distributions are presented in the left columns, while parameters most relevant to response time distributions are presented in the columns to the right of the dividers.

	Pricing (response)						Pricing (RT)			Choice (response + RT)		
Participant #	α	s_v	s_{buy}	s_{sell}	s_{CE}	p_{max}	θ_{speed}	θ_{prec}	ndt_p	$\theta_{c,prec}$	$\theta_{c,speed}$	ndt_c
1	0.91	0.39	0.33	1.00	0.85	0.04	1.71	4.10	0.51	2.90	1.62	0.84
2	0.49	0.91	0.61	0.92	0.96	0.28	0.91	4.00	0.75	1.94	1.93	0.66
3	0.77	0.46	0.46	0.93	0.98	0.00	1.15	1.51	0.24	2.34	0.97	0.45
4	1.29	0.32	0.35	0.96	0.77	0.00	2.23	4.33	0.30	3.15	1.23	0.72
5	0.77	0.86	0.54	0.82	0.89	0.00	2.17	4.80	0.74	4.05	1.31	0.72
6	0.94	0.78	0.32	1.22	0.75	0.00	1.48	3.86	1.03	3.17	1.28	0.73
7	0.77	0.74	0.55	1.08	0.98	0.00	1.77	4.60	0.99	3.77	1.17	0.59
8	0.47	0.86	0.75	1.00	0.97	0.00	1.64	3.12	1.45	4.72	0.97	0.81
9	0.90	0.64	0.47	1.16	0.87	0.08	1.71	2.80	0.48	3.24	0.85	0.69
10	0.74	0.27	0.23	1.45	1.00	0.02	1.05	2.51	0.73	2.74	1.42	0.92

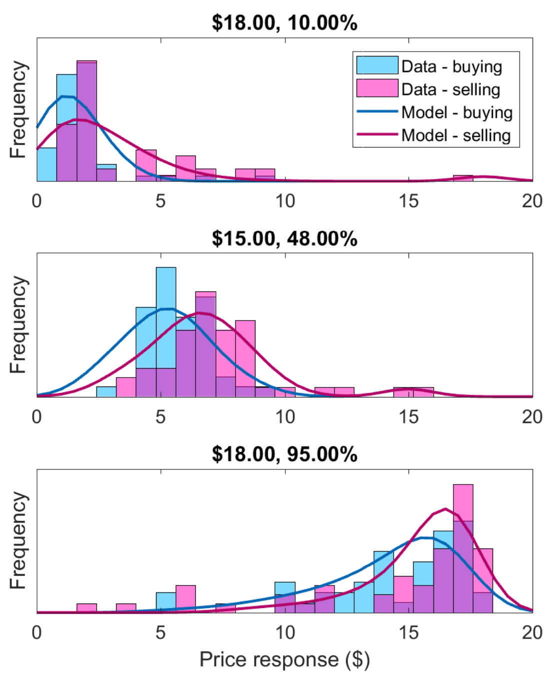

Figure 12. Model predictions (lines) of the distributions of prices (histograms) for buying (blue) and selling (pink) responses for three example gambles.

predicts that different types of prices will converge over time, it does not strictly imply that buying prices will increase or selling prices will decrease. If the mix of gambles presented to participants has an unusually low rate of success, for example, then the initial state for buying may actually be higher than the expected utility of most gambles. In these cases, we might expect both the buying and selling prices to decrease when time pressure is relaxed because the initial points are are both higher than the utility of the gamble. This is what we found in pilot data using a similar paradigm but a different mix of gambles (Kvam & Busemeyer, 2018). The model is able to account for both patterns of data via the relation of start points to gambles.

Finally, as we might expect, thresholds were substantially higher in the precision than in the speed condition. This is in line with the multitude of previous findings on the speed-accuracy trade-off (e.g., Wickelgren, 1977; Vickers & Packer, 1982; Heitz & Schall, 2012) and the difference in response times between speed and accuracy conditions interacts with the start point bias to ensure that the difference between buying and selling prices is larger in the speed condition than in the precision condition.

From these maximum posterior values, we generated predictions for response distributions and response time distributions. Fits to individual participants' response and response time data are provided on the Open Science Framework at osf.io/tfm4e. The model reproduces both of these distributions of these prices will high fidelity, which is illustrated in the aggregate fits to results shown in Figures 8 and 12.

Mouse tracking patterns. Although the process of connecting mouse tracking data to model parameters is not necessarily well-established, the model's ability to connect to mouse tracking data is potentially highly desirable. To do so, we used the angle of accumulation predicted from the model at each point in time (Figure 10) and used it to predict a mouse angle at each point in time during a trial. This was done by taking the MAP estimates from the model and using them to generate one simulated trial (with the same gamble information, participant, and trial conditions). For the simulated trials, we recorded the trajectory of the accumulation process at every 50 ms interval, yielding approximately 9500 simulated trajectories for the same number of real trials. This ensured that the simulated trajectories aligned with the number of trials and the mouse tracking sampling frequency used to analyze the experimental data. Also as with the real data, these were pruned for exceptionally long (> 10s) or short (< 300ms) response times.

For each of the simulated trajectories, we took the price that was most favored at each point in time, which was given by the angle of the state relative to the origin. This allowed us to compare it to the mouse locations in the experimental data that were given by the angle of the mouse relative to the scale 1.

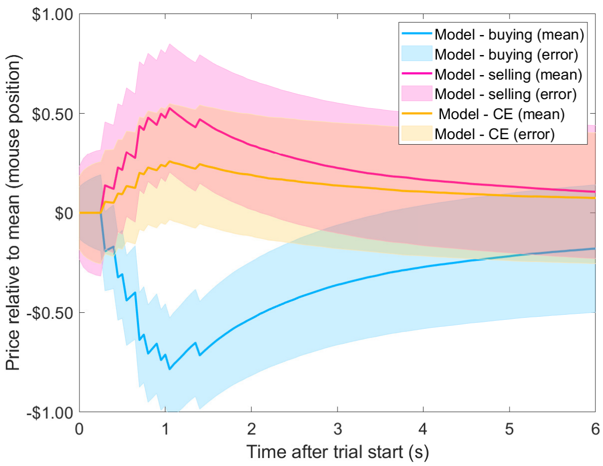

Figure 13. Model prediction for mouse trajectories in the buying (blue), selling (red), and certainty equivalent (orange) conditions. Filled area indicates ± 1 unit of standard error derived from a simulated data set from the model that matched the real data set in its size and assortment of gambles and conditions.

The results of this posterior prediction are shown in Figure 13, which can be compared to Figure 6. The mean trajectory derived from the simulated trials is shown in dark blue, red, and orange for buying, selling, and CE trials, respectively. Critically, the mouse trajectories were assumed to start after the non-decision period had elapsed, which depended on the particular participant. This meant that the first movement away from the origin occurred at 240 ms, when Participant # 3's non-decision time had elapsed. Then participant #4 was added at 300 ms, participant #9 at 480 ms, and so on up until the final participant (#8 at 1450 ms). This is why the trajectories appear more jagged toward the leading edge – as each participant joins the set of accumulating processes, their data is added altogether into the set of mouse trajectories.

While not lining up completely perfectly with the actual mouse trajectories – surely in part because not all internal cognitive processes are expressed through mouse movements – they do bear a striking resemblance to the overall trajectories observed in the experimental data. At minimum, the trajectories derived from the posterior parameter estimates follow qualitatively the pattern observed in the data. This is promising in terms of leveraging process data to inform the model – with further work, it may be possible to estimate model parameters from the mouse trajectories (or possibly eye tracking data, if it is gathered instead) in order to better inform our understanding of individual differences in the pricing process. At present, it seems sufficient to note that the model is successfully reproducing the diverging and then converging pattern of trajectories, which should be a signature of the underlying cognitive process in judgment and decision-making scenarios (Koop & Johnson, 2011; Cheng & González-Vallejo, 2017; Cox et al., 2012; Freeman et al., 2011; Schulte-Mecklenbeck et al., 2011)

Choice. Thus far we have focused almost exclusively on price, in part because the novelty of the price accumulation model is specifically related to price. In part this is because preference reversals and predictions for choice probabilities were covered thoroughly by Johnson and Busemeyer (2005), who showed that a mental simulation-based model could be implemented as a variant of decision field theory [DFT]. The price accumulation model would be identical to this DFT model for choice, creating preference reverals in a similar way, but it is still enlightening to examine how well it accounts for the empirical data from the binary choice task.

To fit the binary choice and corresponding response time data, we took the utility parameter estimates from the price accumulation model (Table 1, α) and used them to compute an expected utility for each gamble in the set. These were then fed into a diffusion model by directly calculating the drift rate using the formulas provided by the decision field theory model in J. G. Johnson & Busemeyer (2005). Before even fitting the choice response time data, we can derive directional predictions from these parameter estimates for choice – an *entirely out-of-sample* prediction made based on the results of the pricing conditions. To perform this cross-validation, the mean difference in utility is given by subtracting the expected utility of the lower-payoff, safer p-bet (S = Safe) from the expected utility of the higher-payoff, riskier d-bet (R = Risky).

$$\mu_{S>R} = EU(S) - EU(R) = p_S \cdot x_S^{\alpha} - p_R \cdot x_R^{\alpha} \qquad (5)$$

If $\mu_{S>R}$ is greater than 0 for a given pair of gambles, then a participant will on average select the safe gamble S over the risky gamble R. If it is negative, then they will select the risky gamble on average. A particular value of α will therefore give us a directional prediction for a given gamble pair about which alternative, on average, a participant should select.

When we compute $\mu_{S>R}$ for each gamble (using each individual's α estimate from the pricing data) and use it to simply predict the direction of each choice, we correctly predict on average 75.8% of the actual responses in the data set. As a out-of-sample cross-validation of the price model, this is exceptional – nearly all of the gamble pairs were approximately matched in terms of expected value, so predicting any decisions with accuracy well above 50% is nontrivial. That the model was able to make these predictions simply by estimating α from the relative distribution of prices assigned to high- and low-payoff gambles suggests that it is tapping into some true underlying representation of value that is common to both pricing and decision processes.

The value of $\mu_{S>R}$ only gives us the average difference in utility between the two alternatives, but does not by itself allow us to predict distributions of response times. The drift rate in decision field theory is given by considering not only the mean utility difference, but how often a decision-maker will mentally simulate contrasts between the gambles. Based on the same idea from the pricing model above, it suggests that a decision maker at each time step draws an outcome of the two gambles and uses the difference between the outcomes to shift their preference toward one gamble or the other. The average rate of motion toward the p-bet is given as a function of the difference between the value of the p-bet $V_S(t)$ and the value of the d-bet $V_R(t)$ at any given point in time (see J. G. Johnson & Busemeyer, 2005, , Equations 1-3, for the complete derivation):

$$d = \frac{E[V_S(t) - V_R(t)]}{E[V_S(t) - V_R(t) - \mu_{S>R}]^2} \tag{6}$$

The value of d therefore corresponds to the average rate of sampling information in favor of the p-bet, which provides the drift rate in a Wiener diffusion process. We were therefore able to feed the values of d into a diffusion model as a fixed drift rate for each trial, using them to estimate the overall rates of selecting the d-bet and p-bet alongside the distribution of response times associated with each decision. The choice threshold and non-decision time were then fit freely with these fixed drift rates using the *dwiener* package in JAGS (Wabersich & Vandekerckhove, 2014). The prior on both speed and accuracy condition thresholds, for each participant, was a uniform distribution on (0,5), and the prior on non-decision time was a uniform distribution on (0,1]) seconds. Bias was fixed at 0.5 so we could see how well drift alone would account for differences in choice proportions. As before, there were 4 chains of 5000 samples each, with 1000 burn-in samples. All chains were inspected visually and via $\hat{R}$ for convergence (Gelman & Shirley, 2011), and the model had little trouble converging (mainly because it was an extremely well-constrained model with only 3 free parameters per person used to predict 300-500 decisions and response times). Since the thresholds and non-decision time do not determine the *direction* of choice – instead only determining how close choice proportions will be to 50% – the α estimates are mainly responsible for the choice proportions generated by the model. Thus, the average directional predictions will be in line with those of $\mu_{S>R}$, but the model will also account for choice variability and the distributions of response times. The resulting modal posterior estimates for thresholds (speed $\theta_{c,speed}$ and accuracy $\theta_{c,prec}$ conditions) and non-decision time ndt_c in the choice task are given on the rightmost side of Table 2.

The parameter estimates for the choice model are largely unsurprising. They were by and large within reasonable ranges that we might normally expect for a drift-diffusion model: thresholds for the precision condition were consistently higher than those for the speed condition, and non-decision times were all within 400-1000ms. Non-decision time was slightly longer than for many perceptual choice tasks and in some cases longer than the non-decision time for the pricing conditions – likely due to the encoding time associated with having to inspect four rather than just two attributes before being able to make an informed choice. This indicates that the main differences in response time between choice and pricing conditions are largely the result of decision processes as opposed to non-decision ones: participants are investing more time into considering their responses in price conditions than they are in binary choice conditions. This may be a direct result of the greater number of responses available in the pricing condition, reflecting an adjustment in line with Hick's law (Hick, 1952; Usher et al., 2002) or Fitt's law (Crossman & Goodeve, 1983) where participants are more careful when they have more options to consider and select.

The predicted choice proportions from the model, generated from the modal posterior parameter estimates for each person, are shown in Figure 14. Model predictions were generated by simulating 10 iterations of each trial / choice participants made (3437 choices, for a total of 34370 simulations) and then using these to calculate the overall choice proportion for each pair of gambles.

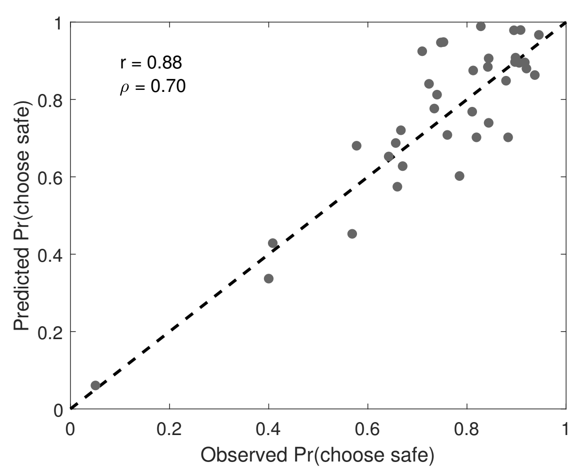

Figure 14. Observed (x) versus predicted (y) proportions of choices in favor of the safe option / p-bet.

Even though d-bets were predicted to have higher prices assigned to them, most of the choices were predicted to be in favor of the p-bet, reflecting the preference reversals shown in Figure 7. The strong correlation between observed and predicted choices (r = 0.88) suggests that the model is doing well at predicting participants' choices, and thus has no trouble predicting that p-bets should be favored in choice even when d-bets are favored in pricing.

The model also predicted response times associated with each of the choice trials participants completed. Individual-level distributions of response times and the model predictions can again be found in the online supplement at osf.io/tfm4e, but the aggregate prediction is shown in the top panel of Figure 8. As with the distributions of price response times, there are only minor deviations between the true distribution (green histogram) and model prediction (solid green lines) for the distributions of choice response times, suggesting that the model is providing an adequate account of this process-level measure as well.

A final curiosity that suggests the price and choice models are tapping into common processes is the relationship between thresholds for choice and price as well as the non-decision times for choice and price. Although these might be expected to be extremely noisy given we only have 10 participants, they are related: matched within participants by speed-accuracy condition, the thresholds across pricing and choice conditions appear to be quite strongly correlated with one another (linear $r = 0.75$, $p = .0002$; rank $\rho = .70$, $p = .0006$), and non-decision time is positively related, if not significantly so, between price and choice tasks (linear $r = .33$, $p = .34$; rank $\rho = .25$, $p = .49$). This obviously a post hoc examination of the parameters and not intended as anything more than an interesting observation, but the fact that these parameters seem to line up suggests that the model is measuring some true underlying construct that is present across conditions. Such a finding simply further demonstrates the utility of the model in describing individual differences in a meaningful way that connects behavior across tasks.

Discussion

The primary goal of this paper was to develop and test a model that could account for important characteristics of pricing by using process-level mechanisms informed by theories of cognition. It is the first to examine in-depth the distributions of prices people assign to items and how they vary as a function of manipulations of payoffs, probabilities, price types, and time pressure. Critically, the random utility extension of prospect theory missed many of the important distributional and all of the dynamic properties of these prices, and as a result was thoroughly overcome by a model that incorporated theory about the underlying processes leading to price generation.

The proposed price accumulation model was able to account for all of the important phenomena in the data through relatively few cognitive mechanisms. The gap between buying and selling prices are attributable to differences in start point, and the interactions between these start points and the mental sampling process generate the difference in skew between the two types of prices. Because mental simulation of outcomes is the mechanism by which the model allows decision makers to accumulate support for different prices, the distributions will also be sensitive to the binomial distribution of outcomes produced by the probability of winning the gambles. Specifically, this binomial will be right-skewed for low probabilities, left-skewed for high probabilities, and have large variance for probabilities near 0.5. As a result, prices given to gambles with win probability near 0 or 1 will be more tightly distributed and more consistent in addition to being more strongly skewed. The stochastic nature of these binomial distributions mean that prices assigned to gambles with high variance will have greater entropy, and thus be less reliable. This was an a prior prediction of the mental simulation process that did not depend on specific starting points or threshold parameters, and it is borne out in the empirical data as shown in Figure 4.

The price accumulation model also incorporated dynamic elements that allowed it to predict different distributions of prices over time. The dynamics of the accumulation process itself do not need to be parameterized, as the step time, length, and probabilities are all fixed or given by the stimulus. But the thresholds for accumulation help determine how the dynamics result in distributions of prices in the speed emphasis and precision emphasis conditions, allowing them to exhibit greater initial bias at earlier time points by triggering responses early in the speed condition and allowing for this bias to wash out at later time points by triggering responses later in the precision condition.

A summary of the important phenomena covered here is given in Table 3. As shown, prospect theory is only capable of produce two out of the ten important phenomena outlined in this paper. Furthermore, even another dynamic and stochastic model, the sequential value mathing model (SVM J. G. Johnson & Busemeyer, 2005) is not capable of capturing all of the important effects. In the next sections, we discuss why this is the case and why one might favor the price accumulation model over other dynamic models.

Growing or shrinking differences?. One particularly curious finding in the study reported here is that buyer-seller differences grew smaller over time, in line with predictions from an anchoring and adjustment approach to pricing (Carlson, 1990; Epley & Gilovich, 2006; Pachur & Scheibehenne, 2017). While consistent with anchoring, this result diverged from some previous work suggesting that buyer-seller differences increased with deliberation time, essentially reversing the effect found by N. J. Ashby et al. (2012, see Figure 2). In this past work, the authors emphasized attentional differ-

Table 3

Table of important phenomena reported in the paper and the price accumulation model mechanisms that allow it to account for them. We also indicate whether random utility extension of prospect theory (RUPT) can account for each phenomenon in the rightmost column.

Behavioral phenomenon (Description)	Figure	PA model mechanism	RUPT
Buyer-seller gap (Buying prices tend to be lower than selling prices for the same gamble)	3, 5, 12	s_{buy}, s_{sell}, s_{CE}	Yes
Effect of time pressure (Response times are faster, buyer-seller gap is bigger)	5	s_{buy}, s_{sell}, s_{CE}, θ_{speed}, θ_{prec}	No
Skew of buying vs selling (Buying prices are more positively skewed than selling)	3, 12	s_{buy}, s_{sell}, s_{CE}, accumulation process	No
Skew of low vs high % (Low % gambles positively skewed, high % negatively skewed)	3, 12	parameter-free	No
Unreliability vs gamble variance (Greater variance in expected payoff leads to lower reliability)	4	parameter-free	No
Variance of prices from max $ (Higher payoffs result in greater variance in prices)	4	α	Yes
Mouse trajectories (Buying and selling trajectories diverge then converge)	6, 13	All	No
Preference reversals - price & choice (Tendency to choose safe gamble over risky in choice, but price risky gamble higher than safe)	7	α, s_{buy}, s_{sell}, s_{CE}	No
Preference reversals - buying & selling (Selling prices tend to favor high-payoff gambles more often than buying prices)	7	s_{buy}, s_{sell},α, accumulation process	No
Response times (Right-skewed distribution of response times in price and choice)	8	θ_{speed}, θ_{prec}, ndt	No

ences between buyers and sellers, where buyers seemed to focus on negative aspects of the stimuli whereas sellers focused on positive ones. This difference in attention caused buying and selling prices to diverge rather than converge. Although seemingly at odds with anchoring and adjustment, the present model does in fact possess a mechanism that can handle these results. Namely, the information search process can also be biased by either over-sampling high outcomes (for sellers) or by over-sampling low outcomes (for buyers).

The presence of these dual mechanisms has been suggested in past work delving into the properties of the endowment effect and buyer-seller gap. The response bias as well as the information search and apparent stopping rules contribute to this gap in different ways. Not only do buyers and sellers appear to strategically mis-represent the prices they are willing to pay or accept (respectively), which appears as a response bias (Pachur & Scheibehenne, 2017), but they also tend to stop considering new information when they have just drawn a piece of information that aligns with their initial bias (Pachur & Scheibehenne, 2012). This work by Pachur & Scheibehenne has illustrated both of these points using decisions from experience, where participants literally draw and experience outcomes from different risky prospects, thereby allowing us to understand the external sampling process. It seems only natural that the same type of sampling process translates to internal samples, where participants might start with an initial price bias, pay attention to samples that agree with that bias, and then stop their internal sampling process when simulated a congruent outcome from the gamble. As we outline below in the the comparisons and extensions section, the price accumulation model actually integrates these disparate sources of bias, with the exception of attention bias, which could easily be integrated into the model should the data indicate it is necessary.

An attentional mechanism can also bias the information that participants search for over time or the relative weight given to different attributes when we move beyond simple gamble stimuli and into consumer choice sorts of scenarios. Early work by Birnbaum & Stegner (1979) suggested that instructions to take the point of view of a buyer or seller resulted in greater weight assigned to information sources that yielded lower or higher value estimates, respectively. This would line up with work on confirmation bias (Nickerson, 1998), where people tend to search for information that agrees with what they already believe – i.e., a buyer might believe they deserve a low price while a seller would believe they deserve a high price to give up their item and systematically look for good or bad information, respectively. The empirical results of the present experiment and in our previous work (Kvam & Busemeyer, 2018) show that buyer-seller differences are more pronounced earlier in a trial, however, which is backed up by the process tracing data (see Figures 5 and 6). This seems to suggest that biased sampling of information is not necessarily responsible for the gap between buyers and sellers in this experiment, but rather that this bias is present initially and actually washes away as the decision maker considers information about the gamble.

It may be the case that further work is needed to reconcile the two potential sources of buyer-seller asymmetries; it is entirely possible that biased sampling is at play in our experiments but that it is overall insufficient to overcome the convergence of prices in the mean. Such a reversal may be related to the scale of response times observed across different studies: while convergence is seen on the scale of 2.5 to 5 seconds in our experiments, divergence was observed on the scale of 5 to 15 seconds in the work by N. J. Ashby et al. (2012). Perhaps the most coherent solution to this problem is that biases based on seeing the maximum payoff first wash out with initial sampling (0-5 seconds to leave the start point), but are then reinstated by attentional processes that bias the sampling process in favor of positive (seller) or negative (buyer) aspects of the gambles or items. Such a finding would align with early work on cognitive dissonance (Festinger & Walster, 1964; Walster, 1964), where preference strength [in this case reflected by price] takes an initial dip before bolstering mechanisms like biased information search kick in and drive preferences back toward the bias-favored response. The price accumulation model can handle both results by building bias into the start point (0-5s result) and/or into the accumulation process (5-15s result), and so it can serve as a theoretical basis for exploring these questions further.

Common mechanisms despite preference reversals. A large body of work in economics has explored explanations for the preference reversal phenomena that appear between pricing and choice, many of which are consistent with the quantitative model that we propose here. For example, Butler & Loomes (2007) indicated that preference or response imprecision may be at play in creating differences between valuations expressed as price responses and those expressed in terms of binary choice. They suggest that anchors – such as the maximum payoff for a particular gamble – may create points of reference that draw participants (whose prices or preferences are uncertain) toward them (Loomes et al., 2009). Our empirical findings certainly reflected this proposal, as the relative poor reliability of price responses (Figure 4) and noisiness of decisions (Figure 7) both indicate that people have some intrinsic uncertainty about their valuation expressed in both pricing and choice. Furthermore, the model reflects an anchoring and adjustment process, whereby the maximum payoff creates an initial price point that is modified as a person considers the likelihood of that payoff (Carlson, 1990), thereby giving greater weight to the payoff in pricing procedures, referred to as scale incompatibility (Tversky et al., 1990).

Scale incompatibility (anchoring) and stochastic elicitation processes seem to be the most popular explanations for preference reversals in the economic and psychological literature (Tversky et al., 1990; Butler & Loomes, 2007; Cubitt et al., 2004; Loomes & Sugden, 1995; Seidl, 2002). Despite some accounts proposing only one or the other, the combination of scale incompatibility and stochastic (error-inclusive) processes seems to be necessary to account for preference reversal phenomena (Schmidt & Hey, 2004). Thus, both of these elements are an integral part of our proposed model, expressed in terms of the stochastic processes underlying both choice and pricing following the initial bias or anchor (e.g., Figure 10).

Perhaps more interestingly, the price accumulation (pricing) and decision field theory (choice) models can be described using the same parameter value (α). Our models and analyses therefore suggest that the underlying valuation process (utility) component of choice and pricing share the same assumptions about how people assign subjective value to monetary outcomes, and yet preference reversals will still arise as a consequence of the different structure of the elicitation procedures. The idea that there are common mechanisms underlying both processes is reinforced by the ability of the utility parameters derived from pricing to predict selections made in the binary choice conditions, indicating that the key "true" valuation (as opposed to perspective-specific bias) part of the price model translates to choice. This is a potentially appealing result for economic theories seeking to understand or make reference to the invariant processes underlying preference across methods of elicitation (Butler & Loomes, 2007; Cubitt et al., 2004; Loomes & Sugden, 1995), as it suggests that such an endeavor should be possible.

Comparisons and extensions. The two sources of bias – starting point and accumulation – may wind up being a point of departure between the price accumulation model and the sequential value matching model (J. G. Johnson & Busemeyer, 2005). The latter account uses similar cognitive mechanisms to account for preference reversals between price and choice, including starting point distributions and a mental simulation process. When it only uses starting point bias, the price accumulation model can in some ways be seen as an extension of the SVM model to a continuous joint distribution of prices and response times, providing the additional mechanisms necessary for it to generate the variation in distributions across conditions. However, price accumulation diverges from this model in important ways. Perhaps the most important way is in the stopping rule: the SVM triggers a price response upon reaching a level of indifference, whereas the price accumulation model triggers a response when sufficient support for one price has been accumulated. In a sense, the SVM contains a passive mechanism where the decision-maker "settles" on a particular price by failing to gather sufficient evidence to move away from it, while the price accumulation model contains an active mechanism where the decision-maker generates support for a particular price. This makes the threshold mechanism in the price accumulation model line up more closely with the empirical data suggesting that sellers tend to be more likely to stop after sampling a high outcome and buyers tend to be more likely to stop after sampling a low outcome Pachur & Scheibehenne (2012). This is at odds with the SVM model, which suggests that buyers would be more likely to hit a point of indifference after stepping up from a lower price (and thus more likely to stop after sampling positive information), and sellers would be more likely to hit a point of indifference after stepping down from a higher price (and thus more likely to terminate after sampling negative information).

Conversely, the price accumulation model suggests that a buyer, who would be toward the low/right part of the accumulation scale shown on the lower left panel of Figure 10, would be most likely to terminate search when stepping right / sampling a negative outcome. Sellers, who would be toward the top/left, would be more likely to terminate search when sampling a positive payoff, because this takes them more directly toward the threshold for high prices. Indeed, if we examine what the final step direction was for buying and selling prices for trajectories like those shown in Figure 10, buying trajectories stopped on a negative (zero) sampling outcome approximately 48% of the time, while selling trajectories stopped on a negative (zero) sampling outcome only 42% of the time (based on 10,000 simulated trials each; we would expect both to be 45% if the stopping sample was unbiased, because the gamble has a 55% chance of winning). This difference would naturally be exaggerated with information sampling biases for buyers and sellers, but certainly hints that the model lines up the the empirical findings put forth by Pachur & Scheibehenne (2012).

Another difference between the SVM model and the price accumulation model is that the present model includes a utility parameter (α) that distorts the representation of high prices relative to low prices. This corresponds to the similarity in representation of high prices insinuated by the utility function as well as work on numerosity and number representation (Feigenson et al., 2004; Longo & Lourenco, 2007; Krueger, 1982). It is a potentially important ingredient that partly accounts for the increased variability in prices assigned to high-payoff gambles shown in Figure 4. It adds to the mechanisms in the SVM model that predict response variance as a function of overall gamble variance (of which maximum payoff is naturally a part). As near as it can be interpreted, the SVM proposes a linear representation of prices, as the scale is divided into even increments (i.e., a person can step from \$1 to \$2, and \$14 to \$15, and these are functionally the same). This different function of the utility parameter in the price accumulation model therefore seems to reflect real differences in number representation that lead to differences in price distributions. However, it remains to be seen whether the greater variance in prices assigned to high payoffs is due to greater sampling variance or differences in representation because the minimum payoff was always zero, and thus greater maximum payoff was conflated with greater variance. One way to test this might be to match gambles on maximum payoff or on variance: for example, compare the variance of prices assinged to G1: 50% chance of \$20, 50% chance of \$15 vs G2: 50% chance of \$15, 50% chance of \$10 (though one would have to find some way to match the variance of the *utilities*, rather than just the dollar amounts). A natural extension of this work would test this question along with those related to negative payoffs and multi-attribute / multi-outcome gambles.

The models can also be integrated together, using mechanisms that allow one or the other to be extended in different directions. The SVM model, for example, includes mechanisms that allow it to be applied to probability equivalents (J. G. Johnson & Busemeyer, 2005). The price accumulation model can borrow these mechanisms to account for distributions of probability equivalents and their associated response times, and it could also be extended to responses like preference or favorability ratings. For example, to change the response format from price, one need simply substitute a set of ordinal or continuous ratings for the prices along the arc shown in Figure 10 and examine how the hitting points are related to the values of probability equivalent or favorability rating judgments instead. The simplicity of the price accumulation model is the key feature here. It may be the case that the SVM model could be supplemented with additional parameters and extended to generate predicts for response times, but it would provide only a discrete approximation of the continuous two-dimensional distribution of prices and response times because it uses a discrete scale. Simulating the price accumulation model and deriving its predictions is probably the simpler solution, as it removes the search and comparison layers in favor of a single support accumulation mechanism.

Emphasizing how these models can be used together, we have shown that the price accumulation model can use the relations derived in Johnson & Busemeyer (2005) to make successful predictions about preference reversals and binary choice by using the parameters estimated from pricing. In this way, they can be viewed as a complementary pair of dynamic cognitive models that predict behavior in pricing and binary choice conditions, where the price accumulation model is used for price and the SVM model is used to map its parameters onto a decision field theory model of decision-making. The success of the pricing model to make out-of-sample predictions for binary choice by using the DFT model certainly suggests that they are accounting for common cognitive mechanisms underlying valuation.

Departing from two-outcome gambles into the realm of losses, mixed outcomes, and multi-attribute / multi-outcome lotteries as well as consumer goods are perhaps the most interesting extensions of this work. Further work is surely likely to indicate that the present model is incomplete and requires additional mechanisms to account for prices and ratings assigned in these new scenarios. These may include mechanisms like rank-dependence or different weights for low outcomes, as suggested by work on configural weight models (Birnbaum et al., 1992; Birnbaum & Zimmermann, 1998). Such theories have difficulty accounting for the dynamic elements of decisions, so naturally the price accumulation model might provide an avenue for them to be built into one. The dynamic nature of the model and its connection to process tracing procedures like mouse tracking will also surely allow it to be supplemented with eye tracking data. Such an integration would permit us to more directly measure attention and feed it into the price accumulation process (as in the work of Krajbich & Rangel, 2011; Krajbich et al., 2012) as well as evaluate what parts of a gamble, product, or environment are considered in constructing price in the first place (Pachur et al., 2018; Fiedler & Glöckner, 2012; Shi et al., 2013; Vachon & Tremblay, 2014).

Conclusions

The distributions of prices, their dynamic shifts, and the additional process-level measures like response times and mouse trajectories each provide evidence against the static and deterministic frameworks usually used to model pricing tasks. Each of these sources of information point toward a dynamic and distributional account of pricing, where the properties of the gamble interact with the price type and the time it is elicited to produce unique distributions of price responses. The price accumulation model provides cognitive mechanisms that explain what changes across these conditions – initial state (bias) changes across price type, threshold changes across time pressure, and dynamics change according to gamble attributes. These mechanisms allow it to easily out-perform models like prospect theory that provide mainly descriptive accounts of pricing and choice. By providing a more complete and thorough account of the underlying processes, this theory is able to predict new phenomena in pricing as well as successfully relate price to preferential choice. Our model therefore provides insight into what psychological processes are shared across different methods of eliciting value (utility representation), which might be thought of as the invariant components of preference (Slovic, 1995).

References

Ashby, F. G., Maddox, W. T., & Lee, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science*, *5*(3), 144–151.

Ashby, N. J., Dickert, S., & Glöckner, A. (2012). Focusing on what you own: Biased information uptake due to ownership. *Judgment and Decision Making*, *7*(3), 254-263.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral science*, *9*(3), 226–232.

Berg, N., & Gigerenzer, G. (2010). As-if behavioral economics: Neoclassical economics in disguise? *History of economic ideas*, 133–165.

Bhatia, S. (2017). Associative judgment and vector space semantics. *Psychological Review*, *124*(1), 1-20.

Bhatia, S., & Pleskac, T. J. (2019). Preference accumulation as a process model of desirability ratings. *Cognitive psychology*, *109*, 47–67.

Birnbaum, M. H., & Beeghley, D. (1997). Violations of branch independence in judgments of the value of gambles. *Psychological Science*, *8*(2), 87–94.

Birnbaum, M. H., Coffey, G., Mellers, B. A., & Weiss, R. (1992). Utility measurement: Configural-weight theory and the judge's point of view. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(2), 331-346.

Birnbaum, M. H., & Stegner, S. E. (1979). Source credibility in social judgment: Bias, expertise, and the judge's point of view. *Journal of Personality and Social Psychology*, *37*(1), 48-74.

Birnbaum, M. H., & Zimmermann, J. M. (1998). Buying and selling prices of investments: Configural weight model of interactions predicts violations of joint independence. *Organizational Behavior and Human Decision Processes*, *74*(2), 145–187.

Bostic, R., Herrnstein, R. J., & Luce, R. D. (1990). The effect on the preference-reversal phenomenon of using choice indifferences. *Journal of Economic Behavior & Organization*, *13*(2), 193–212.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178.

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, *23*, 251–263.

Butler, D. J., & Loomes, G. C. (2007). Imprecision as an account of the preference reversal phenomenon. *American Economic Review*, *97*(1), 277–297.

Carlson, B. W. (1990). Anchoring and adjustment in judgments under risk. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(4), 665-676.

Carmon, Z., & Ariely, D. (2000). Focusing on the forgone: How value can appear so different to buyers and sellers. *Journal of consumer research*, *27*(3), 360–370.

Chen, F., & Fischbacher, U. (2016). Response time and click position: Cheap indicators of preferences. *Journal of the Economic Science Association*, *2*(2), 109–126.

Cheng, J., & González-Vallejo, C. (2017). Action dynamics in intertemporal choice reveal different facets of decision process. *Journal of Behavioral Decision Making*, *30*(1), 107–122.

Cox, G., Kachergis, G., & Shiffrin, R. (2012). Gaussian process regression for trajectory analysis. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *34*(34), 1440-1445.

Crossman, E., & Goodeve, P. (1983). Feedback control of hand-movement and fitts' law. *The Quarterly Journal of Experimental Psychology Section A*, *35*(2), 251–278.

Cubitt, R. P., Munro, A., & Starmer, C. (2004). Testing explanations of preference reversal. *The Economic Journal*, *114*(497), 709–726.

Doane, D. P., & Seward, L. E. (2011). Measuring skewness: a forgotten statistic? *Journal of Statistics Education*, *19*(2).

Donkin, C., Heathcote, A., Brown, S., & Andrews, S. (2009). Non-decision time effects in the lexical decision task. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society.*

Dotan, D., Meyniel, F., & Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition*, *171*, 112–121.

Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological science*, *17*(4), 311–318.

Estes, W. K., & Maddox, W. T. (2005). Risks of drawing inferences about cognitive processes from model fits to individual versus average performance. *Psychonomic Bulletin & Review*, *12*(3), 403–408.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, *8*(7), 307–314.

Festinger, L., & Walster, E. (1964). Post-decision regret and decision reversal. In L. Festinger (Ed.), *Conflict, decision, and dissonance* (pp. 100–112). Stanford, CA: Stanford University Press.

Fiedler, S., & Glöckner, A. (2012). The dynamics of decision making in risky choice: An eye-tracking analysis. *Frontiers in Psychology*, *3*, 335-.

Freeman, J., Dale, R., & Farmer, T. (2011). Hand in motion reveals mind in motion. *Frontiers in Psychology*, *2*, 59.

Friedman, J., Brown, S., & Finkbeiner, M. (2013). Linking cognitive and reaching trajectories via intermittent movement control. *Journal of Mathematical Psychology*, *57*(3-4), 140–151.

Furnas, G. W., Deerwester, S., Dumais, S. T., Landauer, T. K., Harshman, R. A., Streeter, L. A., & Lochbaum, K. E. (1988). Information retrieval using a singular value decomposition model of latent semantic structure. In *Proceedings of the 11th annual international acm sigir conference on research and development in information retrieval* (pp. 465–480).

Gelman, A., & Shirley, K. (2011). Inference from simulations and monitoring convergence. *Handbook of Markov Chain Monte Carlo*, *6*, 163–174.

Goldstein, W. M., & Einhorn, H. J. (1987). Expression theory and the preference reversal phenomena. *Psychological Review*, *94*(2), 236-254.

Grether, D. M., & Plott, C. R. (1979). Economic Theory of Choice and the Preference Reversal Phenomenon. *The American Economic Review*, *69*(4), 623–638. doi: 10.2307/1808708

Groeneveld, R. A., & Meeden, G. (1984). Measuring skewness and kurtosis. *Journal of the Royal Statistical Society: Series D (The Statistician)*, *33*(4), 391–399.

Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, *76*(3), 616–628. doi: 10.1016/j.neuron.2012.08.030

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*(1), 11–26.

Holmes, W. R. (2015). A practical guide to the probability density approximation (pda) with improved implementation and error characterization. *Journal of Mathematical Psychology*, *68*, 13–24.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: a query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(3), 461–474.

Johnson, J. G., & Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Review*, *112*(4), 841–861.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of political Economy*, *98*(6), 1325–1348.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, *5*(1), 193–206.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292.

Koop, G. J., & Criss, A. H. (2016). The response dynamics of recognition memory: Sensitivity and bias. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(5), 671.

Koop, G. J., & Johnson, J. G. (2011). Response dynamics: A new window on the decision process. *Judgment and Decision Making*, *6*(8), 750–758.

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, *3*.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857.

Krueger, L. E. (1982). Single judgments of numerosity. *Perception & Psychophysics*, *31*(2), 175–182.

Kvam, P. D. (2019a). A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *Journal of Mathematical Psychology*, *91*, 14–37.

Kvam, P. D. (2019b). Modeling accuracy, response time, and bias in continuous orientation judgments. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 301–318.

Kvam, P. D., & Busemeyer, J. R. (2018). Dynamic and distributional properties of prices. In T. T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 655–660). Austin, TX: Cognitive Science Society.

Landauer, T. K., & Dumais, S. T. (1997). A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240.

Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology*, *89*(1), 46-55.

Longo, M. R., & Lourenco, S. F. (2007). Spatial attention and the mental number line: Evidence for characteristic biases and compression. *Neuropsychologia*, *45*(7), 1400–1407. doi: 10.1016/j.neuropsychologia.2006.11.002

Loomes, G., Orr, S., & Sugden, R. (2009). Taste uncertainty and status quo effects in consumer choice. *Journal of Risk and Uncertainty*, *39*(2), 113–135.

Loomes, G., & Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review*, *39*(3-4), 641–648.

Luce, R. D. (2000). *Utility of gains and losses*. Earlbaum.

MacKenzie, I. S., & Buxton, W. (1992). Extending fitts' law to two-dimensional tasks. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 219–226).

Mellers, B. A., Chang, S.-j., Birnbaum, M. H., & Ordonez, L. D. (1992). Preferences, prices, and ratings in risky decision making. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(2), 347–361.

Morewedge, C. K., & Giblin, C. E. (2015). Explanations of the endowment effect: an integrative review. *Trends in cognitive sciences*, *19*(6), 339–348.

Morewedge, C. K., Shu, L. L., Gilbert, D. T., & Wilson, T. D. (2009). Bad riddance or good rubbish? ownership and not loss aversion causes the endowment effect. *Journal of Experimental Social Psychology*, *45*(4), 947–951.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, *2*(2), 175–220.

Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, *55*(1), 84–93.

Pachur, T., & Scheibehenne, B. (2012). Constructing preference from experience: The endowment effect reflected in external information search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(4), 1108–1116.

Pachur, T., & Scheibehenne, B. (2017). Unpacking buyer-seller differences in valuation from experience: A cognitive modeling approach. *Psychonomic Bulletin & Review*, *24*(6), 1742–1773.

Pachur, T., Schulte-Mecklenbeck, M., Murphy, R. O., & Hertwig, R. (2018). Prospect theory reflects selective allocation of attention. *Journal of Experimental Psychology: General*, *147*(2), 147–169.

Palestro, J. J., Sederberg, P. B., Osth, A. F., Van Zandt, T., & Turner, B. M. (2018). *Likelihood-free methods for cognitive science*. Springer.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003).*

Pothos, E. M., Busemeyer, J. R., & Trueblood, J. S. (2013). A quantum geometric model of similarity. *Psychological Review*, *120*(3), 679–696.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Ratcliff, R. (2018). Decision making on spatially continuous scales. *Psychological Review*, *125*(6), 888–935.

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in cognitive sciences*, *20*(4), 260–281.

Regenwetter, M., & Robinson, M. M. (2017). The construct–behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*(5), 533–550.

Scheibehenne, B., & Pachur, T. (2015). Using bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic bulletin & review*, *22*(2), 391–407.

Schmidt, U., & Hey, J. D. (2004). Are preference reversals errors? An experimental investigation. *Journal of Risk and Uncertainty*, *29*(3), 207–218.

Schoemaker, P. J., & Hershey, J. C. (1992). Utility measurement: Signal, noise, and bias. *Organizational Behavior and Human Decision Processes*, *52*(3), 397–424.

Schulte-Mecklenbeck, M., Kühberger, A., & Ranyard, R. (2011). The role of process data in the development and testing of process models of judgment and decision making. *Judgment & Decision Making*, *6*(8), 733–739.

Seidl, C. (2002). Preference reversal. *Journal of Economic Surveys*, *16*(5), 621–655.

Shi, S. W., Wedel, M., & Pieters, F. (2013). Information acquisition during online decision making: A model-based exploration using eye-tracking data. *Management Science*, *59*(5), 1009–1026.

Slovic, P. (1995). The construction of preference. *American psychologist*, *50*(5), 364–371.

Slovic, P., & Lichtenstein, S. (1983). Preference reversals: A broader perspective. *The American Economic Review*, *73*(4), 596–605.

Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, *123*(4), 425–451. doi: 10.1037/rev0000023

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, *1*(1), 39–60.

Turner, B. M., & Sederberg, P. B. (2012). Approximate Bayesian computation with differential evolution. *Journal of Mathematical Psychology*, *56*(5), 375–385.

Tversky, A., Slovic, P., & Kahneman, D. (1990). The causes of preference reversal. *The American Economic Review*, 204–217.

Usher, M., Olami, Z., & McClelland, J. L. (2002). Hick's law in a stochastic race model with speed–accuracy tradeoff. *Journal of Mathematical Psychology*, *46*(6), 704–715.

Vachon, F., & Tremblay, S. (2014). What eye tracking can reveal about dynamic decision-making. *Advances in Cognitive Engineering and Neuroergonomics*, *11*, 157–165.

Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, *50*(2), 179–197.

Wabersich, D., & Vandekerckhove, J. (2014). Extending JAGS: A tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behavior Research Methods*, *46*(1), 15–28.

Walster, E. (1964). The temporal sequence of post-decision processes. In L. Festinger (Ed.), *Conflict, decision, and dissonance* (pp. 112–128). Stanford, CA: Stanford University Press.

Weber, E. U., & Johnson, E. J. (2011). Query theory: Knowing what we want by arguing with ourselves. *Behavioral and Brain Sciences*, *34*(2), 91–92.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*(1), 67–85. doi: 10.1016/0001-6918(77)90012-9

Yechiam, E., Ashby, N. J., & Pachur, T. (2017). Who's biased? A meta-analysis of buyer–seller differences in the pricing of lotteries. *Psychological Bulletin*, *143*(5), 543–563.

Supplemental material for “A distributional and dynamic theory of pricing and preference”

Peter D. Kvam

University of Florida

Jerome R. Busemeyer

Indiana University

Supplemental material for “A distributional and dynamic theory of pricing and preference”

Price accumulation model details

The basic price accumulation model, fit to just an individual condition (ignoring differences between buying, selling, CE and speed, accuracy as they feature separate parameters), has only five parameters: utility / representation parameter α, starting point variability s_v, starting point bias s_β, threshold θ, and non-decision time ndt. In this study, we included a sixth contaminant parameter p_max, used to account for cases where participants just reported the maximum payoff as the value of a gamble. This parameter wound up only really being important for two participants (mainly participant #2, and to a lesser extent participant #9), so it may not be the case that it is a fundamentally important part of the model overall. Further studies could probably try to remove these trials or participants, avoid their occurrence by ensuring participants are following directions, or follow this same approach and model the contaminant process.

The priors for these parameters were chosen such that they restricted positive parameters to be positive (e.g., ndt, θ) and to keep simulated response times within the 0-10 second range. They were meant to be largely uninformative so as to let the data inform the model, but also to avoid cases where the accumulation process failed to finish in a reasonable time (which usually just leads to long response times, but can also cause the model to be unusable if some simulated trials take too long to generate).

Fitting procedure

The approach we took to fitting the model in this study took multiple stages, which was done for two reasons. First, prospect theory – even with a random utility component added on – does not predict response times but only predicts distributions of responses. Therefore, it seemed sensible to compare the price accumulation model to prospect theory purely on their ability to handle the prices and then later extend the model to account for response times. The second reason was that the price accumulation model, like most dynamic models of decision-making (Ratcliff et al., 2016; Busemeyer et al., 2019), can produce the same behavior from different combinations of parameter values by varying them with a scaling factor. This means that one parameter must be fixed in order set the scale of the model, or at least fixed for one condition (Donkin et al., 2009). Fixing the thresholds in order to fit the other parameters allows us to set the scale of the model for modeling the price distributions.

Fixing the parameters once they have been fit to the price distributions also allows us to conduct a more stringent test of the model. Rather than allowing all nine parameters vary to fit prices and response times, we only use six for price distributions and then three free parameters for response times. This also ensures that the fit of the models to price distributions is maintained when we move to modeling response times. In turn, this second step added a dimension to the model in terms of its ability to handle multiple process-level measures, including both RT distributions for speed and accuracy conditions (Main text Figure 8) as well as predict mouse trajectories (Figures 6 and 13).

The price accumulation and prospect theory models were fit using a kernel density estimation method to turn the simulated data into a truly continuous, 2-dimensional distribution of responses and response times. This method has been effectively used to approximate the likelihoods of several types of simulation-based models (Palestro et al., 2018; Turner & Van Zandt, 2012; Turner & Sederberg, 2014), is reasonably efficient especially with the addition of signal processing methods (Holmes, 2015; Lin et al., 2019), and can be easily adapted to a two-dimensional joint distribution. For the price accumulation model, we can simulate a large number of trials from the model, use the kernel density method to generate an approximate likelihood, and then impute the likelihood of each combination of response and response time in the observed data set.

This was combined with a Markov chain Monte Carlo method for estimating the posterior distribution of parameters using a Metropolis-Hastings algorithm. For each new proposed sample, the proposed parameters of the model were used to generate a set of simulated data, pass a kernel density estimator over the simulated data, and impute the likelihood of the data by evaluating the height of the kernel density estimator at the x, y (response, response time) location of each data point. The log likelihoods of each data point were then summed across all data from a single participant to get the overall log likelihood of the participant's data given the proposed model parameters. These were added to the log likelihood of the priors for the proposed parameter values to form the posterior likelihood of the proposed parameter values. For posterior likelihoods that were greater than those of the current parameters, the MCMC process would then step to the new proposed location in the parameter space, repeat the process for the next sample, and so on.

It is critical to note that simulation-based likelihoods have some variability in the approximated likelihood function because they are generated based on simulated random draws rather than a true probability density function. Therefore, it is possible to randomly sample a likelihood for a given combination of real / experimental data and simulated data that is higher than the "true" likelihood of the data for those parameter values (Lin et al., 2019). Therefore, rather than re-using the same likelihood every time a set of parameters was used (for example, if the process did not take a step to a new location), a new set of simulated data was generated and the likelihood re-computed on every fourth step that was not taken. This prevented the sampler from getting stuck on abnormally high likelihoods that were due to noise in the simulated data and likelihood rather than truly greater likelihoods for particular sets of parameter values.

To fit each iteration of the price accumulation model and prospect theory model, we used 4 chains of 4000 samples, including 500 burn-in samples. It is possible to fit the prospect theory model using traditional MCMC samplers rather than the simulation-based model, but simulation versus analytic approaches can yield slightly different posterior likelihoods (Lin et al., 2019) and so we used the same method to compute the likelihoods for both models. Reinforcing the practical utility of the approach, we found no substantial differences in the likelihoods or parameter estimates generated for prospect theory using the simulation method compared to a JAGS implementation. For each sample, the approximate likelihood was generated by simulating 100 simulated trials for each real data point, and then calculating the likelihood of that data point using the kernel density estimator. The various chains were visually inspected and used $\hat{p}$ (Gelman & Shirley, 2011) to determine convergence. Finally, the maximum a posteriori (MAP) estimates used to generate the posterior predictive plots (Figures

Simulation and recovery

Despite its relative simplicity for a dynamic model of continuous responses, there are some relationships between its parameters, and so it is critical to ensure that the various parameters are identifiable and meaningful. The key way to establish this is by simulating data from the model and then recovering that data (Heathcote et al., 2015). Doing so ensures that a know true generating process can be accurately and precisely estimated, and allows us to understand where and how the model estimates might be biased.

For the model recovery, we simulated 360 trials (approximately the number that we might expect from one of the price conditions) across the same 72 gambles that were used in the actual study. We then fit a 5-parameter version of the model to the data, fixing the start point variability to set the scale of the model. We then estimated parameters of the model using 2 chains of 1000 samples each (200 burn-ins) and the simulation-based likelihoods as described above.

The results are shown in Figure S1. For the most part, parameters were estimated reasonably well, with α (utility) and s_β (start point bias) being slightly under-estimated. This is likely due to the two parameters trading off, as an increase in α that shifts the locations of prices on the response scale can be partly compensated for by a corresponding increase in s_β that shifts the location of the state on the price scale. Since the unique contribution of α is to increase the variance of high-payoff prices, and these prices do not compose every gamble in the study, it is reasonable that it might be estimated with somewhat lower accuracy and precision.

The other minor area of misfit seems to be in the threshold, where a few posterior estimates were about 0.2 units higher than the true generating parameter. This seems to be related to non-decision time, where there are two “islands” with thresholds low and non-decision time high, and thresholds high and non-decision time low (Figure 1 row 5, column 4). This is a consequence of the fairly commonly observed trade-off between non-decision time and thresholds, where response times can become shorter by decreasing non-decision time or threshold or become

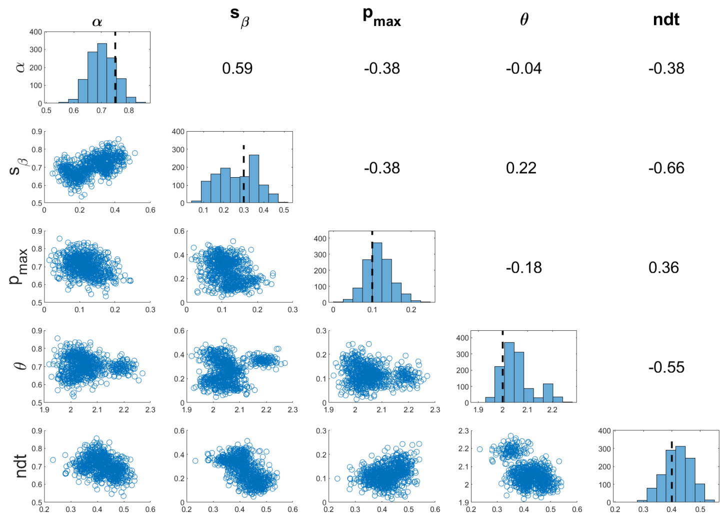

Figure 1. Model recovery for the price accumulation model. The main diagonal slides show the estimated posterior (histogram) compared against the true generating parameters (dotted black line). Bottom-left plots show the scatter plot of samples compared against one another, while the upper-right plots simply give the correlation between samples of the different parameters

longer by increasing non-decision time or threshold. This naturally results in a negative correlation between the estimates, and one of the chains seemed to sample around a local minimum along this ridge before migrating and “discovering” the true generating parameter value. Fortunately, the modal posterior estimates for these parameters till reflect the bulk of samples that were near the true generating parameter, but it might be wise for modelers to exercise caution and be vigilant in inspecting posterior distributions to catch instances of multimodality in future work using the model.

Although not entirely perfect, the model recovery was generally successful and suggests that the model is at least able to estimate parameters with reasonable precision when we know that it is the true generating process.

Selective parameter disabling

A natural question to ask regarding any model is what contribution its parameters make to the overall fit of the model. In order to ensure maximum parsimony, any parameters that do not substantially contribute to the quantitative or qualitative can potentially be removed. Here, we determine the usefulness of each of the main parameters from the model – including utility power α, separate start points for different conditions s_β, start point variability s_v, and the contaminant distribution parameter p_{max} – by disabling them one by one and examining how the model fit changes by doing so.

Utility parameter. The α parameter that controls the relative orientation of different prices in the price accumulation model can be "turned off" by setting $\alpha = 1$. In the realm of expected utility, this would be equivalent to using the expected value of computations rather than characterizing payoffs in terms of diminishing or increasing marginal returns. In the current model, it is equivalent to assuming that the representations of all numbers or prices are commensurate with their objective values – i.e., the psychological difference between \$1 and \$2 is the same as the difference between \$18 and \$19.

There are strong theoretical reasons to include a parameter like α in a model of valuation, driven by fundamental principles related to both the economic concept of utility (Friedman & Savage, 1948; Savage, 1954) and psychological concepts of number sense (Krueger, 1982). However, removing α from the model at hand has relatively subtle effects. It reduces the ability of the model to predict the variance of high-payoff gambles. This is shown in the left panel of Figure S2 – when the parameter is removed from the model, it has trouble capturing responses to high-payoff gambles like (\$17, 60%) because the variance of the predicted distributions is too small (and it also appears to bias estimates toward higher values than are observed in the data).

Removing the utility / representation parameter from the model naturally has consequences for the quantitative fit. Overall, considering just the fit to the distribution of prices, removing this parameter decreases like log likelihood of the model by 1620. The improvement in BIC conferred by dropping a parameter is not enough to overcome this difference, as the overall BIC of the

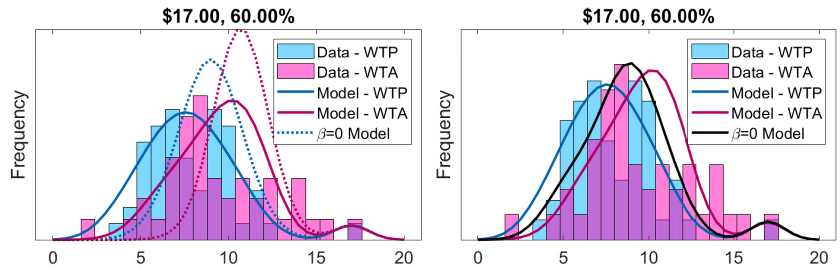

Figure 2. Performance of the model when the utility α parameter is removed (left) or when the start point bias parameters β are removed.

model with $\alpha = 1$ is 3230 higher than the model freely estimating this parameter (higher BIC indicating worse fit).

Start point. In the price accumulation model, start point bias is mainly responsible for differences in price between buying / WTP and selling / WTA conditions. Naturally, removing this parameter results in just a single distribution of prices for buying, selling, and certainty equivalent conditions. Interestingly, this is also true for fixing the start point variability parameter. Because the strength of a starting point bias is modulated by the strength of the start point (degree of start point variability), fixing $s_v = 0$ is equivalent to removing all of the start point bias as well. Doing so allows us to test both the presence of start point variability and bias as well.

Fixing this parameter to be equal to zero results in a distribution of prices shown on the right side of Figure S2. Failing to differentiate between buying and selling prices results in this model under-predicting the variance of responses in general because it misses the between-condition variance that comes from this manipulation. As a result, it misses prices on the high side in the selling condition and it misses prices on the low side in the buying condition.

As with the utility parameter, fixing the start point parameter(s) also has quantitative consequences when it comes to predicting price distributions. Fixing it results in a log likelihood drop of 3070. Since doing so effectively removes four parameters for each participant (three start point biases and a start point variability), it could be the case that the reduction in complexity justifies such a large drop in likelihood, but again the improvement in BIC conferred by having fewer parameters is insufficient to overcome this difference (BIC of fixed model is 5761 higher).

Fixing start point variability also has fairly dire consequences for the model's ability to predict response time distributions. Without this parameter, the distance from start point to threshold is fixed on every trial. As a result, the response time distributions predicted by a model lacking start point variability has much lower variance than the data (as shown in the main text in Figure 6).

Contaminant parameter. Probably the easiest parameter to remove would be the p_{max} parameter, as it simply signifies the prevalence of contaminant processes (participants giving the maximum payoff as their price response). For eight out of the ten participants, this parameter was effectively zero, and so it could easily be removed for all except two participants and have minimal effect on the overall fit.

For the other two participants, however, the effect of setting $p_{max} = 0$ is quite significant. The prevalence of responses at the maximum payoff (small humps at $17 in Figure S2) necessitates the presence of this parameter for participants 2 and 9. Thus, the model does not improve overall when it is removed for everyone, which results in a modest overall drop in log likelihood of 200. The improvement in BIC conferred by dropping 10 parameters is 94.75 (10 times ln(13072)), and thus the overall BIC obtained by dropping p_{max} is approximately 305 points worse (higher).

Threshold. There is little doubt that threshold shifts are necessary to account for differences in response times resulting from time pressure. There are no reasonable alternative options to predict the $\sim$2.5 second difference between speed-emphasis and precision-emphasis conditions. However, the thresholds also seem to contribute to the model fit in terms of the price distributions. Since prices tend to converge over time, differentiating between early and late prices appears to be added value that is conferred by the threshold parameters. Fixing the two thresholds to be the same results in a log likelihood decrease of 3190, even greater than the contribution of the start point bias and variability parameters. However, it removes fewer parameters than fixing the start point, so naturally the BIC for the fixed-threshold model is worse (6285 points higher).

Decision field theory model

The decision field theory model that we used to predict choice and response times in the choice condition was based on a simple version of the diffusion model (Ratcliff et al., 2016) and implemented in JAGS using the dweiner package (Wabersich & Vandekerckhove, 2014). In the version presented in the main text of the paper, we derived exact values for the drift rates based on the utility parameter estimated from the price distributions. These were fed directly into JAGS as fixed values, while thresholds and non-decision time were freely estimated. To fit the model, we used four chains of 5000 samples, with 500 burn-in samples per chain. These chains were inspected using the $\hat{r}$ metric for convergence, and all were within reasonable tolerance. The results of this analysis suggested that there was some degree of convergence between the price accumulation model and decision field theory.

However, it may be the case that this only occurred because the drift rates were constrained directly by the price model. It is worth double checking that the model reaches similar values for drifts (as well as thresholds and non-decision time), so we examine the results of an analysis based on freely estimated parameters for this same model. Here, thresholds were permitted to vary between speed and precision instructions but all other parameters were fixed across conditions. Table S1 shows the resulting estimates.

As shown, the estimates line up reasonably well with the results presented in the main paper (Table 2). Thresholds for the precision condition are well above those for the speed condition and generally in the same ranges, while non-decision times seem to be similar – if slightly shorter – than the estimates from the original model. The estimates of α all seem to line up well with the fixed values from the price model, showing the same trend toward risk aversion in people's decisions. The was one participant (4) who appeared to switch from risk-seeking to risk-averse, but this participant was still the most risk-seeking out of all of the participants in the study.

The introduction of bias (β) appeared to be mainly responsible for the shifts in other parameters. In this model, the "safe" (low payoff, high probability gamble) was set to be the upper boundary, whereas the "risky" (high payoff, low probability gamble) was set to be the lower

Table 1

Parameter estimates of utility α*, thresholds for speed and precision conditions* θ_s *and* θ_p*, non-decision time ndt, and predecision bias in favor of the safe gamble* β*. Estimates are generated from the freely estimated decision field theory model of responses and response times in the choice condition.*

Participant	α	θ_p	θ_s	*ndt*	β
1	0.8161	2.7406	1.7978	0.7037	0.3565
2	0.4922	2.0510	1.4181	0.5048	0.4742
3	0.7475	1.7091	1.3972	0.3162	0.4299
4	0.9373	2.8295	1.5578	0.5921	0.3733
5	0.7088	3.3444	1.7403	0.5718	0.4108
6	0.8384	2.6997	1.6560	0.6045	0.3600
7	0.7232	3.4655	1.7335	0.4203	0.3780
8	0.4632	3.1266	1.7991	0.6416	0.4372
9	0.8247	2.4949	1.2858	0.5550	0.4150
10	0.7122	2.3365	1.4867	0.7905	0.4364

boundary. An unbiased decision-maker would start at $\beta = .5$. Thus, a bias parameter above .5 would indicate an initial bias toward choosing the safe option, and one below .5 would indicate an initial bias toward choosing the risky one. Most participants appeared to have a slight initial bias toward choosing the risky option ($\beta < .5$), but then gathered information favoring the safe option over time ($\alpha < 1$).

This is perhaps not too surprising, as the first piece of information that participants would typically look at is the maximum payoff of the gamble. This would "bias" them toward choosing the risky option, first allowing them to accumulate information favoring the high-payoff, low probability alternative before looking at probability information that pushed them toward the low-payoff, high-probability (safe) option. This aligns well with work on two-stage or attention switching mechanisms (Diederich & Busemeyer, 2006; Diederich & Trueblood, 2018; Guo et al., 2017), which suggest that an initial attraction toward a risky alternative (due to so-called System 1 reasoning) is eventually overcome by preferences for the safe one (due to executive control exerted by so-called System 2 reasoning).

References

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, *23*, 251–263.

Diederich, A., & Busemeyer, J. R. (2006). Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or two-stage-processing hypothesis. *Perception & Psychophysics*, *68*(2), 194–207.

Diederich, A., & Trueblood, J. S. (2018). A dynamic dual process model of risky decision making. *Psychological Review*, *125*(2), 270-292.

Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, *16*(6), 1129–1135.

Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of political Economy*, *56*(4), 279–304.

Gelman, A., & Shirley, K. (2011). Inference from simulations and monitoring convergence. *Handbook of Markov Chain Monte Carlo*, *6*, 163–174.

Guo, L., Trueblood, J. S., & Diederich, A. (2017). Thinking fast increases framing effects in risky decision making. *Psychological science*, *28*(4), 530–543.

Heathcote, A., Brown, S. D., & Wagenmakers, E.-J. (2015). An introduction to good practices in cognitive modeling. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 25–48). Springer.

Holmes, W. R. (2015). A practical guide to the probability density approximation (pda) with improved implementation and error characterization. *Journal of Mathematical Psychology*, *68*, 13–24.

Krueger, L. E. (1982). Single judgments of numerosity. *Perception & Psychophysics*, *31*(2), 175–182.

Lin, Y.-S., Heathcote, A., & Holmes, W. R. (2019). Parallel probability density approximation. *Behavior Research Methods*, *51*(6), 2777–2799.

Palestro, J. J., Sederberg, P. B., Osth, A. F., Van Zandt, T., & Turner, B. M. (2018). *Likelihood-free methods for cognitive science*. Springer.

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in cognitive sciences*, *20*(4), 260–281.

Savage, L. J. (1954). *The foundations of statistics*. New York, NY: John Wiley & Sons, Inc.

Turner, B. M., & Sederberg, P. B. (2014). A generalized, likelihood-free method for posterior estimation. *Psychonomic bulletin & review*, *21*(2), 227–250.

Turner, B. M., & Van Zandt, T. (2012). A tutorial on approximate bayesian computation. *Journal of Mathematical Psychology*, *56*(2), 69–85.

Wabersich, D., & Vandekerckhove, J. (2014). Extending JAGS: A tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behavior Research Methods*, *46*(1), 15–28.