

Attractors: Incidental Values That Influence Forecasts of Change

Clayton R. Critcher¹ and Emily L. Rosenzweig²

¹University of California, Berkeley,

²Ochsner Health

Author Note

Clayton R. Critcher, Haas School of Business, University of California, Berkeley; Emily L. Rosenzweig, Ochsner Health.

The research reported in this article was supported in part by U.S. National Science Foundation Award 1749608, awarded to Clayton R. Critcher. We thank Rohit Agrawal, Nathan Doiev, Noma Kahf, Philipp Loick, Tanvi Sheth, Jennifer Soliman, and Shayne Sossamon for their assistance with data collection.

All code, data, and materials are publicly available through the Open Science Framework (<https://osf.io/z9h6y>).

Correspondence concerning this article should be addressed to Clayton R. Critcher; Haas School of Business; University of California, Berkeley; 545 Student Services Bldg., #1900; Berkeley, CA 94720. Email: claytoncritcher@haas.berkeley.edu

Abstract

This article examines whether forecasts of change are influenced by *attractors*, salient values in the direction of the considered change. When an attractor is relatively distal from (vs. proximal to) the base value from which change originates, it encourages forecasts of greater change. Participants showed this pattern when predicting: which of two airfare changes was imminent (Study 1) and by how much gas prices (Study 2) or a stock's price (Study 3) would change. Attractors have this influence because they alter the way people translate even equivalent subjective interpretations of prospective changes into objective forecasts of change. In the context of a distal (vs. a proximal) attractor, forecasters thought more objective change was necessary to reflect the same subjective characterization of that change (Study 4). Having participants precommit to a subjective interpretation of an objective amount of change reduced a subsequently introduced attractor's influence on forecasting (Study 5). Following almost five decades of research showing many ways arbitrary values anchor judgments, we discuss how attractors reflect the first evidence that such values can also influence adjustment.

KEYWORDS: forecasting; trends; pricing; adjustment; anchoring

Attractors: Incidental Values That Influence Forecasts of Change

A potential traveler searches online for airfare between New York and Los Angeles. The website shows the history of fares that has been inching upward as the travel date approaches. The traveler examines this data as she considers how much prices are likely to change in the near future. A dietitian weighs a client each week and keeps records on his progress. She looks at the trajectory and considers how likely he is to lose 3 pounds again this week as she formulates a recommended regimen. With oil refineries shut down by a hurricane in the Gulf of Mexico, a driver accesses AAA's Daily Fuel Gauge to see how gas prices have been trending upward. Contemplating a quick run to the service station, he considers whether tomorrow morning will bring another spike or the beginning of a plateau.

In each of these examples, people are forecasting change. When doing so, some information is obviously relevant (e.g., the observed variability in fares over time, the general trend of a client's weight loss, previous gas spikes following hurricanes). But in this paper, we examine a more arbitrary influence on forecasts of change: salient values in the direction of change, which we call *attractors*. Values may become attractors because they are spontaneously, internally generated (e.g., a nearby round number toward which a value is trending) or because they are externally salient. For example, if one looks at a time-series graph tracking the movement of the biopharmaceutical stock Amgen, one may try to forecast by how much the stock will slide from its price of \$247.10 if the company releases lackluster results from its latest clinical trials. Most obviously, how disappointing the news is should matter. But less obviously, we posit that a salient attractor—for example, whether the next prominent axis label below the current price happens to be \$240 or 200—may matter as well.

This paper has two central goals. First, we test whether attractors—and in particular their position as relatively proximal to or distal from the value from which a change begins—influence forecasts of change. Second, we aim to explain why attractors influence forecasts of change. As we develop more fully below, we hypothesize that attractors may do so by influencing a forecaster’s interpretation of whether a potential, contemplated amount of change feels subjectively substantial or insubstantial. This can make the same potential change feel like an overestimate or underestimate, respectively. Before articulating our account in greater detail, we begin by expanding more on attractors and how they relate to anchors, a multi-faceted construct long of interest to behavioral scientists.

Attractors

There is a large literature on how arbitrary or incidental values influence judgments. Anchoring research has identified numerous ways in which numerical values or non-numerical expressions of magnitude (Oppenheimer, LeBouef, & Brewer, 2008) exert an assimilative pull on numeric judgments (Tversky & Kahneman, 1974). Anchors—whether explicitly presented in the judgment context or internally generated during the formulation of a judgment—have been shown to influence a variety of judgments: willingness to pay (Adaval & Wyer, 2011; Critcher & Gilovich, 2008; Nunes & Boatwright, 2004; Simonson & Drolet, 2004), size estimates (Wong & Kwong, 2000), trivia questions (Epley & Gilovich, 2001; Jacowitz & Kahneman, 1995), and even judicial sentencing (Englich & Mussweiler, 2001).

Forecasting change is one example of an anchoring-and-adjustment problem (Eggleton 1982). As Harvey and Reimers (2013) described, forecasters who consider how a trend line is likely to evolve “could use the last data point in the series...as their anchor and adjust away from that” (p. 589). Offering a broader characterization of the role of anchoring in the forecasting

literature, Theochari (2014) wrote, “The anchor is in most cases the last data point and adjustment is based on the patterns perceived in the data” (p. 17). And Harvey, Bolger, and McClelland (1994) showed that in forecasting change, their participants used “the previous actual value...as an anchor” (p. 215) and adjusted based on properties extracted from the observed trend. In other words, the most recent known value can be conceived of as an anchor, and the question of how much change in that value will occur is one of how much adjustment is deemed appropriate. In the context of forecasting change, anchoring-and-adjustment—given adjustment tends to be insufficient—may give rise to trend damping (Harvey, 2007). That is, forecasts of change may be inappropriately small, lying too close to the anchor (the last data point). Such underestimation of change, when present, may emerge because people adjust only far enough to reach a value that seems plausible, but little further (Epley & Gilovich, 2006).

In this paper, we do not seek to test whether forecasts of change are inappropriately anchored to the last known data point in a time series. To the contrary, we seek to understand a complementary phenomenon: how incidental values in the direction of a change—what we call *attractors*—influence forecasts of change from an anchor. In considering the novelty of this approach, it is worth appreciating that the many anchoring literatures have devoted considerable attention to different reasons why anchors restrain judgments. That is, anchoring phenomena have been tested using multiple operationalizations that invoke different psychological processes: numeric priming (e.g., Critcher & Gilovich, 2008), response scale distortion (Frederick & Mochon, 2012), the selective accessibility of anchor-consistent information (Adaval & Wyer, 2011; Strack & Mussweiler, 1997), or an effortful process of adjusting (often insufficiently) along an internal, mental number line (e.g., Epley & Gilovich, 2001). Despite this diversity, anchoring phenomena are unified by a common theme: For various reasons, anchors

restrain people’s judgments, making them unlikely to move far from their grip. We are aware of no attempts to examine a phenomenon in which an arbitrary value (here, an attractor) influences how much adjustment from an anchor seems warranted.¹

Just as Epley and Gilovich (2006) emphasized that the process of adjustment is marked by uncertainty about whether one has adjusted “far enough,” people who consider a specific change (in formulating their own forecast or considering someone else’s proposal) should possess similar uncertainty as to whether any specific change feels too small or too large. In such contexts, there are reasonable cues that people can, and no doubt do, consult (e.g., “Is this change plausible given natural constraints within the domain?”, “Is this change a big percentage shift?”, “Is this change larger or smaller than the ones in previous periods?”). We argue that attractors—and specifically, the distance between these incidental values and the starting value from which the change originates—shape subjective characterizations of the magnitude of a contemplated change. By our account, an attractor serves as a psychological benchmark or reference point that influences one’s subjective sense of whether a considered amount of change is substantial, and thus sufficient.

Our proposal is based on a relatively straightforward observation: The same objective amount of change or adjustment toward an attractor that is far away (a *distal* attractor) compared to one that is close by (a *proximal* attractor) covers proportionally less ground toward that attractor. In other words, in moving from the last known value to a forecasted value, one closes the gap between the anchor (the last known value) and the attractor (the salient value toward

¹ We concede one could choose to refer to the attractor itself as a particular kind of anchor, one that operates through a distinct mechanism from those explored in previous research. We recommend against this approach (and expand on this point in the General Discussion and the concluding Context Paragraph once the reader has the benefit of seeing our body of studies) to avoid further conceptual and taxonomic imprecision that has come from labeling essentially any influence of numbers on judgment as examples of “anchoring.” That said, we appreciate it will be up to future researchers to decide whether to subsume attractors within the umbrella concept of “anchors” or to embrace our conceptualization of attractors as a meaningfully complementary construct.

which adjustment occurs) more or less quickly depending on whether that attractor is proximal or distal from the anchor, respectively. We argue that this leads the same objective change to seem less substantial (and thus potentially not substantial enough) when considering change toward a distal instead of a proximal attractor. Given the attractor's ability to warp people's subjective interpretation of an equivalent objective amount of change, people will forecast more change when moving toward a distal (compared to a proximal) attractor. We illustrate the implications of this logic in the context of one of our opening examples, the question of how much gas will spike following a hurricane.

Imagine that gas prices have been trending up 5 or 10 cents per day. But because the gas prices were lower to begin with in South Carolina than in New York, the price in each state is now \$3.55 and \$3.85 per gallon, respectively. Drivers in both states may ask themselves how much prices are likely to rise the next day. If \$4 is a psychologically salient value in the direction in which prices are moving, then note that each additional cent increase does less to "close the gap" with the salient attractor when it is distal (for South Carolinians) than when it is proximal (for New Yorkers). To the extent that the attractor serves as a psychological benchmark, this may lead the same objective increase to feel less substantial in the context of the distal attractor. In Figure 1, this is illustrated by a possible 5-cent increase seeming "small" for South Carolinians but "medium-sized" for New Yorkers (Figure 1). Due to this way that attractors influence subjective interpretations of different contemplated objective changes, people may forecast greater objective change when moving toward distal than proximal attractors. In this example, New Yorkers and South Carolinians may even share the same subjective sense of how much gas prices are likely to change (e.g., "I hear we're likely to see a medium-sized increase in gas prices

in the next day or two”), but by our account New Yorkers would translate this into a smaller forecasted change.

Although our thesis and the contexts to which it applies may be novel foci of research, our hypotheses are made psychologically plausible by a wide range of work that examines the imperfect and inconsistent mapping between objective stimuli and a person’s subjective interpretations of those stimuli (e.g., Ostrom, 1970; Parducci, 1965). That said, we examine a new way in which the surrounding context (a salient value) changes the way that an objective value (a certain amount of adjustment) is subjectively construed and thereby alters forecasts of change.

A Complement to Previous Examinations of Evaluations or Forecasts of Change

Regardless of the specific occurrence of change under consideration, there are several common constituents of actual or forecasted change. A real or forecasted change involves the transformation of some attribute from its base or starting value to its actual (or forecasted) ending value. Change can be quantified in absolute terms (e.g., “Consumer sentiment grew by 1 point after the Democratic primary, and another 2 points after the presidential election”). But change can also be observed and considered in completely relative terms without reference to absolute numeric values (e.g., “Consumer sentiment grew twice as much after the presidential election than the Democratic primary.”)

Previous research has indeed identified influences on people’s forecasts or interpretations of change. In some cases, this research has looked at how information about historical patterns of change influences forecasts of future change (Lawrence, Goodwin, O’Connor, & Onkal 2006; Lawrence & Makridakis 1989; Reimers & Harvey 2011). For example, when people observe a time-series graph that shows a very slow rate of growth, they tend to assume growth will

accelerate, but when they observe a particularly fast rate of growth, they assume growth will decelerate (Harvey & Reimers 2013). In other cases, researchers have looked not at what influences forecasts of change, but at what properties govern evaluations or interpretations of a given change. A classic example is the *jacket-calculator* problem, which demonstrates differential sensitivity to equivalent objective changes: More than twice as many people reported being willing to travel for 20 minutes to save \$5 on a \$15 item than on a \$125 one (Tversky & Kahneman, 1981; see Thaler 1980). This reflects an extension of Fechner's (1860) law: As the objective intensity of a stimulus (here, price) grows, subjective sensitivity to equivalent change in that intensity diminishes. Also, research from the numerosity literature shows that even when the percentage change is held constant, the equivalent amount of change feels bigger when expressed in a metric that describes the change as more units in that metric (Brannon & Terrace, 1998; Pandelaere et al., 2011). For example, Wertenbroch et al. (2007) showed the same monetary difference seemed bigger when expressed in a weak currency (473.9 vs. 4,739 Chilean pesos) compared to a strong currency (1 vs. 10 U.S. dollars).

Although the specific lessons from these programs of research vary, they all highlight how properties of the change itself (e.g., the historical rate of change, the percentage change, the units in which change is expressed) influence how people evaluate or forecast it. In contrast, we propose that an incidental feature—a salient value in the direction of the forecasted change—shapes forecasts of change *because* it alters perceptions of a given, considered amount of change. In this sense, we combine these two research foci (forecasts and evaluations of change) in the present study of attractors, an influence that is independent of the change itself.

Overview of the Present Studies

Studies 1-3 test our central hypothesis that attractors influence forecasts of change. We tested whether participants thought airfare between two cities was about to experience an objectively small or large change (Study 1), how much participants believed gas prices were likely to shift (Study 2), and the amount by which a stock's closing price would move (Study 3). In all three studies, we tested how distal vs. proximal attractors—incidental values in the direction of adjustment that are relatively far from or close to the base value, respectively— influenced these forecasts. By varying features of the forecasting task, the nature by which the values' historical trajectory is presented, and the instantiation of the attractor, we could test the robustness of the basic effect. The consistent prediction is that in the presence of a distal attractor (vs. a proximal attractor), larger objective changes should seem more likely.

Studies 4 and 5 test our proposed explanation for this effect—that attractors influence forecasts of change by altering how people translate a subjective evaluation of change or adjustment into an objective amount of change or adjustment. Study 4 tested whether the same subjective amount of change (e.g., “The stock experienced a [small, medium] increase in price”) was translated into more objective change in the context of a distal (vs. a proximal) attractor. We predicted that subjectively equivalent characterization of change (“small” or “medium”) would be mapped onto objectively different forecasts of change in the context of distal and proximal attractors. To more directly show that attractors' distortion of this objective-subjective mapping is what explains the basic attractor effect on forecasting, Study 5 had some participants precommit to a certain subjective interpretation of an already-observed objective amount of change before even being exposed to the attractor. If our mechanistic logic is correct, this manipulation should interfere with attractors' ability to influence forecasts of change.

Interested readers can find more information about the creation of our experimental materials in the Supplemental Materials. Materials (as seen by participants) and data, complete with analysis scripts, can be found online: <https://osf.io/z9h6y>. Each study received ethical approval from the University of California, Berkeley, and/or Tulane University.

Study 1

When people search for flights on kayak.com, they are shown a time-series graph that illustrates how the airfare between two cities has fluctuated over the last several weeks. Although Kayak predicts whether fares are likely to rise or fall, it is up to the user to intuit how large such shifts will be. But critically, these graphs include one or more attractors as well: prominent y-axis numerical labels that have accompanying lines that cross the entire horizontal span of the graph, what are sometimes called minor axes. Inspired by this format, we showed participants trajectories of how airfare between a pair of cities had fluctuated over 9 weeks. We positioned these trajectories on graphs such that the ninth week's price was relatively proximal to or distal from an attractor. We then had participants estimate which of two possible price changes—one relatively small, one relatively large—was more likely for the next week. Our primary prediction was that participants would think objectively larger changes in the price of airfare were more likely when prices were moving toward a distal (instead of a proximal) attractor.

Method

Participants. Three hundred eighteen undergraduates from Tulane University and the University of California, Berkeley, participated in exchange for course credit. Each laboratory study in this paper was run during a different semester using a participant pool that participants are in for just one semester, meaning there was not overlap in who participated in each study.

Procedure. Participants were told that their task would be to predict fluctuations in airfare. As part of the instructions, participants saw a time-series graph that was an actual graph shown to shoppers on Kayak.com. It traced the cheapest airfare found each week between a pair of cities. Participants were told that we were testing whether their intuitions about price changes might be as good as the algorithms used by sites like Kayak.com and Bing Travel.

Participants completed sixteen trials. On each trial, participants saw four time-series graphs. Two graphs depicted recent prices on one flight route; two graphs showed recent prices for another route. We paired flight routes such that for the final price change—the price change that participants were forecasting—both routes experienced a price increase, or both routes experienced a price decrease. What participants had to forecast was whether a larger price change was imminent on one route as opposed to the other. Crucially, participants always saw one flight route (e.g., Miami to Atlanta) paired with a distal attractor; the other flight route (e.g., New York to Chicago) was paired with a proximal attractor. We counterbalanced between participants which flight route was paired with the distal as opposed to the proximal attractor. Although we describe in more detail below (and in the Supplemental Materials) how these materials were generated, we encourage readers to look to Figure 2 to see how this task was concretely experienced by participants.

Participants' task was to select which flight route (the one with the distal or the proximal attractor) was the one that experienced the large price change, and thus which flight route experienced the small price change. One such pairing was depicted on the left half of the screen; the other, on the right. This means that one half of the screen depicted price changes that matched our hypotheses (a large price change moving toward a distal attractor, a small price change moving toward a proximal attractor), while the other half depicted price changes that

mismatched our hypotheses (a small price change moving toward a distal attractor, a large price change moving toward a proximal attractor). Participants indicated which half they thought reflected the real price changes from 1 (*pretty sure Left Half is REAL*) to 6 (*pretty sure Right Half is REAL*). Of the 16 trials that participants considered, only the 8 experimental trials—those that involved price changes *toward* the attractor—are analyzed further. We included the 8 filler trials (for which the price changed *away from* the attractor) so participants would see that the attractor line was not an informative value toward which prices were converging; the lines were instead merely the only multiple of \$100 in the depicted range.

Materials. We constructed 32 time-series line graphs. Displaying data using this format is ideal when one wants to depict period-over-period changes (Hutchinson, Alba, & Eisenstein, 2010), our context of interest. Each graph supposedly reflected fluctuations in the cheapest airfare for an identified U.S. domestic route. The graphs included 10 data points, reflecting the price from “10 weeks ago” to “1 week ago.” Each graph also included an attractor, operationalized as a prominent y-axis label and accompanying horizontal line that spanned the width of the graph. For each graph, the specific attractor value was always the multiple of \$100 that was closest to the actual cheapest airfare between the two cities (as found on Kayak.com on February 27, 2014). The y-axis plotted a total range of \$70, which meant only one prominent y-axis label (and thus potential attractor) was on each graph. The full randomization procedure used to generate the graphs is detailed in the Supplemental Materials.

We created four versions of each of these 32 time-series graphs. The versions of each graph were identical except for two features. First, we varied whether the last depicted change in price (from “2 weeks ago” to “1 week ago”) was relatively small (\$3-\$5) or relatively large (\$9-\$11). Second, we shifted the entire trajectory on the graph so that the base price for this final

price change (i.e., the “2 weeks ago” price) was relatively proximal to the attractor (\$12 - \$16 away) or relatively distal from the attractor (\$24 - \$28 away, but always \$12 more than the matching proximal version). The graphs were constructed such that the base price (“2 weeks ago”) was always in the very middle of the numerical space delimited by the y-axis. In this way, the dollar amount by which airfare could adjust or change (and still remain in the area of the plot) was identical on all trials regardless of the nature of the attractor.

For half of the flight routes, we constructed graphs such that the final price change moved *toward* the attractor. For the remaining half, the final price change moved away from the attractor. These latter trials were filler trials that did not factor into analyses. Because across the stimuli, prices were just as likely to move toward as away from an attractor, this helped to demonstrate that the attractor did not provide information about the price shift, such as the airfare’s long-term average or the value most prices converged toward. Furthermore, so participants would know why that value in particular was identified on the graph, we explicitly stated that only multiples of \$100 would be labeled.

Results and Discussion

In order to test whether hypothesis-matching price changes (i.e., large—distal, small—proximal) are forecast to be relatively more likely than hypothesis-mismatching ones, we used a mixed model. We created a variable *match*, which differentiated whether on a particular trial, the hypothesis-matching price changes were on the left (-1) or right (+1) half of the screen. We also included two random effects. One was for flight route pair, thereby accounting for the fact that a large or small price change may be seen as more likely to accompany specific routes. The second was for participant, to account for the non-independence of participants’ responses across trials.

In essence, this accounts for variability in participants' tendencies to believe that graphs on the left or right half of the screen were the real ones.

Providing support for our central hypothesis, the main effect of match was significant, $B = 0.072$, $SE = 0.031$ $t(2199.24) = 2.31$, $p = .021$. As expected, participants were more likely to believe that the price changes depicted on the right side of the screen were the real ones when they depicted large price changes paired with distal attractors and small price changes paired with proximal attractors ($M = 3.33$) compared to when these were reversed to be hypothesis-inconsistent pairings ($M = 3.19$).

One question is whether these findings differ from diminishing sensitivity to change, as illustrated by the jacket-calculator problem (Tversky & Kahneman, 1981). That research shows that the same *absolute* change feels smaller (and thus, in our context, perhaps too insubstantial) when it is made from a larger base value (meaning it represents a smaller *percentage* change from the base value). Note that in the context of price decreases, the attractor effect and diminished sensitivity to change make the same prediction. That is, when the price was more distal from (vs. more proximal to) a smaller attractor, then the same downward price change may have felt subjectively smaller because it reflected a smaller percentage change from this higher initial price. If so, this offers an alternative explanation for why downward adjustment was greater in the context of a distal (vs. proximal) attractor. But this alternative explanation predicts a reversal when price changes were positive: Participants should forecast *less* change when adjusting upward toward distal attractors (vs. proximal attractors) given in such contexts the base price is lower. In other words, this alternative cannot easily account for the observed, hypothesized main effect of match. Furthermore, we did not find that our results were driven by the price-decrease trials: The effect of attractors did not differ when participants were forecasting

positive as opposed to negative price changes, $B = -0.030$, $SE = 0.033$, $t < 1$. This demonstrates both the robustness of and the distinctiveness of the attractor effect compared to what has been examined in past research.

The present results show that people believe that in the context of an uninformative distal (vs. proximal) attractor, a relatively large (vs. small) price change is more likely. Beyond testing our theoretical account, these findings have an intriguing applied implication. Because Kayak and similar travel websites are incentivized by referral fees to encourage people to purchase their airfare on the spot, Kayak could position attractors on their price graphs strategically. When Kayak predicts airfare increases, it could use graphs with distal attractors, so large price hikes seem imminent. But when Kayak predicts airfare decreases, it could use graphs with proximal attractors, which encourage the impression that the decline will be minimal.

Study 2

Whereas the first study tests how attractors changed the perceived likelihood of a possible change, Study 2 had participants forecast changes themselves. We moved to a new context in which forecasting change is relevant: predictions of the price of gas. Drivers receive frequent information about how gas prices are shifting. Their decisions about when to “fill up” can be influenced both by the quantity of gas remaining in their tanks and their estimates of how gas prices are likely to shift in the coming days.

Participants saw time-series graphs that supposedly depicted how the average price of gas fluctuated over a randomly selected 9-day period in a randomly selected U.S. state. Participants learned whether the average price increased or decreased on Day 10. They then had to adjust a “prediction bar” up or down from the Day 9 price to reach their final forecasted value

for Day 10. We predicted that participants would adjust further (thereby forecasting a larger change) when moving toward a distal compared to a proximal attractor.

Method

Participants. Two hundred eight undergraduates at Tulane University participated for extra course credit.

Materials. We generated 24 time-series trajectories, each with nine points of data that represented a gas price's value over nine days. The visible portion of the y-axis was approximately 72 cents in length. Each time-series trajectory was centered such that its value on Day 9 (the base price from which adjustment would occur) was located at the midpoint of the range delimited by the y-axis. In this way, the space for adjustment on each side of the base price was identical (36 cents) for all materials. That is, the permissible range of adjustment was the same regardless of whether a graph included a proximal or distal attractor. This feature is crucial: It means it is not the case that there is simply more room for change to occur on certain graphs compared to others. In this way, attractors are not endpoints that constrain how much change is possible, but benchmarks that may influence the subjective interpretation of considered change.

Each gas-price trajectory also included an attractor: a single value marked with a horizontal line through the y-axis. The attractor was randomly selected to be \$2 or \$3, plausible values given the depicted gas prices were said to have come from the last 10 years. To offer an explanation for why the attractor was present, the instructions explained that only the whole-dollar prices would be prominently displayed. Given the y-axis range was less than a dollar, only one attractor appeared on each graph. Short tick marks identified five-cent increments. We created two versions for eight of the 24 graphs. For the *proximal* attractor versions, the attractor was relatively close to the Day 9 value (8 to 12 cents away). For the *distal* attractor versions, the

attractor was relatively far from the Day 9 value (20 to 24 cents away, but always exactly 12 cents more than in the corresponding proximal attractor version).

For the 16 critical trials (i.e., the distal and proximal attractor versions of the eight distinct trajectories), the stated direction of adjustment was toward the attractor (see Figure 3, for an example of one of these eight key pairs). For the 16 filler trials, the direction of adjustment was always away from the attractor. Thus, as in Study 1, the direction of the attractor was not correlated with the direction of adjustment. This served to reinforce that the attractor was not meant to offer meaningful information about the Day 10 price.

Procedure. Participants were told that commodity traders are paid hefty salaries to anticipate changes in the price of oil, and thus the price of gas. We stated we were interested in whether ordinary people might be able to intuit price changes as well as professional traders can. Immediately to the right of the Day 9 value was a short red bar. Participants were asked to use the up or down arrow keys to adjust this prediction bar to their final forecasted value. Before beginning the task, participants were quizzed on their understanding of the instructions: “How many cents does each tick mark on the y-axis represent?”, “How many days’ worth of prices will you receive?”, and “How many graphs will you make predictions about?”. To correct any misunderstandings, participants were shown the answers before beginning the forecasting task.

The 32 trials (16 experimental graphs and 16 filler graphs, as described above) appeared in one of two semi-randomized orders. We placed one constraint on the randomization: The two versions of each trajectory could not appear in the same half of trials (to minimize the likelihood of detection). Each semi-randomized order differed in whether the proximal or distal version of a particular trajectory appeared in the first or second half. On each trial, participants used the up and down arrow keys to adjust the red cursor to where they thought the average gas price would

be on Day 10. Each click moved this prediction bar up or down one cent. Through iterative clicks participants adjusted to their final forecast. We excluded trials on which participants failed to adjust (2.6%) or adjusted in the opposite direction of what was instructed (5.9%).

Results and Discussion

To test whether participants adjusted more toward distal than proximal attractors, we leaned on a similar data analytic strategy to what was used in Study 1. We created the variable attractor *distance*, which differentiated those trials for which the attractor was proximal to (-1) or distal from (+1) the base price (i.e., the Day 9 value). We also defined *direction* to differentiate those trials that requested downward (-1) as opposed to upward (+1) adjustment. We permitted these two fixed effects to interact. To account for non-independence across trials, we included random effects of *participant* and gas-price *trajectory*. The predicted variable was the amount of adjustment (from Day 9 to Day 10), not the final Day 10 forecast, given the same Day 10 forecast would reflect different amounts of forecasted change depending on the attractor distance manipulation.

Confirming our central hypothesis, participants adjusted further in the direction of a distal attractor than a proximal attractor, $B = 0.33$, $SE = 0.05$, $t(2833.82) = 6.20$, $p < .001$.

Extrapolating from the model, participants adjusted an average of 6.41 cents toward distal attractors, but only 5.74 cents toward proximal attractors. In other words, participants estimated that the price change from Day 9 to Day 10 would be 12% greater in the context of a distal (vs. proximal) attractor. As in Study 1, the influence of the attractor did not differ when participants were instructed that the gas price had increased from when told that it had declined, $B = -0.07$, $SE = 0.05$, $t(2832.60) = 1.25$, $p = .210$.

Study 3

In our first two studies, participants consistently estimated greater changes when values move toward a distal compared to a proximal attractor. In combination, the two studies established the robustness of this effect by demonstrating consistent effects in perceptions of what changes were likely (Study 1) and in change forecasts participants made themselves (Study 2). That said, both studies depicted change trajectories and operationalized attractors in an identical way, on time-series graphs that were moving toward a salient minor axis label. The primary aim of Study 3 was to determine whether attractors influence forecasts of change even with novel methods.

Think back to our example from the Introduction, about gas prices in New York and South Carolina moving upward toward the salient round number of \$4. Implicit in that example is that—at least when dealing with fractional units like dollars and cents—whole number values may serve as salient attractors. After all, it is known that round numbers are salient values that serve as natural reference points in contexts as varied as SAT scores and baseball batting averages (Pope & Simonsohn, 2011). This hints at the existence of internally salient attractors—round numbers toward which change occurs but that are not made visually prominent on a graph—that may then influence forecasts of change in the way our previous studies showed.

Participants received information about how a particular stock price varied over a 9-day period. But this information was presented in a table instead of on a graph. On the critical trials, the final closing value of a stock was either relatively close to (proximal attractor) or relatively far from (distal attractor) the closest whole number value, a value toward which the stock was moving. For example, we assume that a stock that is moving upward from \$6.53 or \$6.69 are both advancing toward the naturally salient attractor of \$7. Even without the perceptual features that characterized our earlier studies' materials, we predicted that participants would continue to

forecast greater change when the value was moving toward a distal rather than a proximal attractor.

Method

Participants. Three hundred fifty-three Americans were recruited from Amazon Mechanical Turk. Thirty participants failed an attention check that asked whether participants had considered price fluctuations in stock prices (correct answer), gas prices, car prices, or the inflation rate. This left 323 participants for all analyses reported below.

Materials. We used the trajectories created for use in Study 2 and modified them in four ways for the present study. First, each 9-day trajectory was said to reflect movement in the closing value of a specific stock instead of the change in the price of gas. Second, we performed a linear transformation on the values so that the daily stock price fluctuation could range from a decline of 50 cents to an increase of 50 cents. Third, and relatedly, participants were informed of the Day 9 closing price for the stock. It had a dollar amount (accompanied by some number of cents) that was randomly sampled from the set of integers from \$1 to \$50.

Fourth, instead of displaying time-series graphs, we presented the nine days of stock prices in a table. Whereas in Study 2 we made salient the attractor (i.e., the whole number value toward which the gas price was moving) visually (by including a minor axis on a time-series graph to identify the closest integer value), the tabular presentation format lacked this feature. Instead, we assumed that the nearby whole-number value would serve as a naturally (internally) salient attractor. For the 16 critical trials (which included a proximal and a distal attractor version of the same trajectories), the stock was moving *toward* the closest integer value. For the 16 filler trials (each of which reflected a unique trajectory), the stock was said to be moving *away from* the closest integer value.

When the starting value was proximal to an attractor, it was 27, 29, or 31 cents away from the nearby integer value toward which it was moving. When the attractor was distal from the attractor, it was 43, 45, or 47 cents away. The distal and proximal attractor versions of each graph displayed identical trajectories except that they were exactly 16 cents apart from each other. For example, one stock whose value was said to experience a decline on Day 10 closed Day 9 at \$26.45 in the distal version; the entire trajectory was shifted down by \$0.16 so that the it closed at \$26.29 in the proximal version (Figure 4). For the filler trials (for which there was only one version of each trajectory), the stock was moving away from the closest integer value on Day 10. As in our earlier studies, this served to make clear that stock prices did not typically converge toward the closest round-number price.

Procedure. The basic procedure was similar to that used in Study 2, but participants learned they would be forecasting the movement of stocks instead of gas prices. Whereas the depicted range of the graph naturally constrained Study 2 participants in how much change they could forecast, the open-ended response format in Study 3 meant that constraints would have to be imposed more explicitly. (Such constraints help to avoid outlying responses that should be implausible given the depicted variability in the stock prices). In the instructions, we highlighted to participants that “we selected stocks whose value never changed by more than 50 cents, either on the days we show you or between the 9th and 10th days.” On each trial, participants were asked to “estimate the stock’s price on Day 10 (to two decimal points).” If participants attempted to enter a closing value that was in the wrong direction of change or reflected a change of more than 50 cents, participants were prompted to reenter their forecasts. The 32 trials appeared in a random order, with the constraint that two trials depicting the same trajectory (but with different attractors) could not appear in the same half.

Results and Discussion

To test whether participants still forecasted more change toward a relatively distal compared to a relatively proximal attractor, we used the same model specification used for Study 2. The predicted main effect of attractor distance emerged, $B = 0.007$, $SE = 0.001$, $t(4836) = 5.45$, $p < .001$. When considering the same trajectories, participants thought the stock would change by 20.7 cents when moving toward a distal attractor, but only 19.2 cents when moving toward a proximal one. In other words, the distal attractors prolonged forecasts of change by 8%. Analogously to the previous studies' results, this effect did not differ by whether the stock value was moving higher or lower, $B = -.001$, $SE = .001$, $t < 1$.

Although the first three studies have established that attractors influence estimates of change, they have yet to demonstrate *why* attractors have this effect. By our account, attractors change the way people construe a prospectively considered objective amount of change in subjective terms (e.g., as small vs. medium), and vice versa. This is because the same objective change closes a smaller proportion of the gap (and thus is more likely to feel implausibly small) in the context of a distal vs. a proximal attractor. The final two studies test our mechanistic account.

Study 4

If attractors alter people's forecasts of change by shaping what they think would constitute a subjectively large or small amount of change, then people should translate the same subjective description ("A *small change* in price was experienced on the last day") into different objective amounts of change in the presence of distal vs. proximal attractors. In Study 4, participants again saw graphed trendlines and had to adjust a prediction bar toward a distal or proximal attractor to indicate their forecasts. But this time, we directly told participants—in

subjective terms (“small” or “medium”)—how much change had actually occurred. Our mechanistic logic makes two predictions for how the same subjective description of change should translate into objectively different forecasts in the context of distal and proximal attractors.

First, we predicted that what constitutes either a “small” or “medium” change would be greater in the context of a distal vs. a proximal attractor. If so, this would highlight that it’s not simply that attractors influence what subjective amount of change seems likely, but instead that attractors change how much objective adjustment is required to achieve the same subjective amount of change. Second, and reflecting a more nuanced prediction, we predicted that the difference between what constituted a “small” and “medium” change would be objectively greater in the context of a distal (vs. proximal) attractor. Both predictions follow directly from the idea that each unit of adjustment in the context of a distal attractor feels subjectively smaller than each unit of adjustment in the context of a proximal attractor; furthermore, these two properties can be seen in Figure 1. Thus, in the context of a distal attractor, more adjustment would be necessary to achieve both “medium” and “small” change (the first prediction), just as more units of adjustment would be necessary to differentiate “medium” and “small” change (the second prediction). This second prediction in particular would help to address an alternative explanation that attractors do not guide the interpretation of adjustment that is moving toward it (what we suggest), but instead serve as an anchor that reflects participants’ initial guess for where the stock price will ultimately land, an initial consideration that forecasters adjust (typically backward) from. By this alternative interpretation, which sees no role of attractors guiding the subjective interpretation of change moving toward it, it is unclear why a “small” and

a “medium” change would be more differentiated in the context of a distal than a proximal attractor.

Like those in Study 3, participants in Study 4 were told they were forecasting how a stock’s price changed over time. Online investment companies like E*Trade, Scottrade, and Ameritrade have brought market investment opportunities to the masses. The brokers themselves, as well as free resources like Yahoo Money, provide information about stock price fluctuations over time that are typically presented in graphical form. In addition, these graphs often have prominent attractors. For example, intra-day stock graphs on Yahoo Money highlight the stock’s opening price with a horizontal line that covers the full length of the graph. Even though we did not want to choose an informative attractor (like an opening value), our decision to return to the graphical presentation (as opposed to tabular presentation in Study 3 that relied on the natural salience of round numbers) was rooted in the ecological validity of such presentations.

Method

Participants. One hundred thirty-eight undergraduates at the University of California, Berkeley, participated for course credit or as part of a larger session for which they received \$15.

Materials. We leaned on the same set of 32 graphs (16 experimental, 16 filler) used in Study 2, but made three modifications. First, because the trajectories were said to reflect stock-price movement, the attractors were randomly sampled from a uniform distribution of integers from \$1 to \$99. Second, instead of merely telling participants that a stock experienced an increase or decline in price from Day 9 to Day 10, participants were told before they made their forecast that each change was “small” or “medium.” For half of the trials, we said the stock experienced a “small” change. For the other half, we indicated there was a “medium” change.

Thus, across all 32 graphs (half of which were fillers), there were 8 small increases, 8 small decreases, 8 medium increases, and 8 medium decreases.

Third, in an effort to test the robustness of our effects and in keeping with previous research on forecasting trends (e.g., Harvey & Reimers, 2013), we eliminated the minor tick marks from the y-axis. Note that the decision to exclude or include the tick marks is superfluous to the internal validity of the study (i.e., it matters not whether one unit of adjustment is thought to be 1 cent, 5 cents, or some other value). The benefit of their inclusion was it enhanced the external validity of the materials. The benefit of their exclusion is that by providing the precise value for the attractor but not the base price, we could vary attractor distance without having to vary the stated base price. Given this modification, we describe adjustment in units of adjustment (1 up or down click = 1 unit) instead of the objective value of change.

Procedure. The instructions were similar to those used in Study 2, except the graphs were described as tracking changes in the closing prices of stocks traded on the NYSE over a random 9-day period. Participants were asked to adjust a prediction bar from Day 9's value to estimate the value of the stock on Day 10. After a sample trial, participants completed all 32 trials. As in Study 2, we excluded the trials (0.4%) in which participants did not adjust the bar in the requested direction.

Results and Discussion

We aimed to test two predictions for how attractors influence adjustment. First, to test whether the same subjectively characterized amount of change is translated into larger objective change in the context of a distal vs. proximal attractor, we expected to find a main effect of attractor distance while controlling for the subjective change label (small or medium). That is, even given the same subjective description of an upcoming change, people should expect that

change to be objectively larger in the context of a distal (compared to a proximal) attractor. Second, to test whether the objective gap between what constitutes a “medium” vs. “small” change is greater in the context of a distal vs. proximal attractor, we tested for an interaction between attractor distance and subjective change label.

Our statistical models were analogous to those used in our previous studies. We again created the variable attractor *distance*, which differentiated the distal (+1) and proximal (-1) attractor version of each graph. We defined the new variable *subjective change* magnitude to differentiate the trajectories that we said would experience a small (-1) versus a medium (+1) change. Beyond including the interaction of these two fixed effects, we added two random effects to our model: *participant* and *trajectory*. This accounted for the non-independence of trials completed by the same participant and any tendency for some trajectories to encourage more adjustment than others, respectively.

Providing a manipulation check of sorts, there was a main effect of subjective change, $B = 1.90$, $SE = 0.15$, $t(6.00) = 12.99$, $p < .001$. This confirmed that participants attended to the subjective change labels: They adjusted further when told the change was “medium” ($M = 7.79$ units) rather than “small” ($M = 3.99$ units). Supporting our first hypothesis, there was also a main effect of attractor distance, $B = 0.20$, $SE = 0.04$, $t(2051.89) = 5.33$, $p < .001$. This reflects that what constituted a “small” or “medium” change was different in the context of a proximal vs. a distal attractor. On average, the same subjective characterization was represented by an additional 0.41 units of objective adjustment in the context of a distal than a proximal attractor. Finally, consistent with our second hypothesis, there was also an Attractor Distance X Subjective Label interaction, $B = 0.09$, $SE = 0.04$, $t(2051.92) = 2.38$, $p = .017$. This reflected that in the context of a distal attractor, more adjustment was necessary to differentiate between two

subjective characterizations of the required adjustment: The objective gap between “small” and “medium” was wider in the context of a distal attractor (3.98 units of adjustment) than a proximal attractor (3.61 units of adjustment).

Study 4 supports our account of (one reason) why attractors influence estimates of change. Even if people are trying to forecast the same subjective amount of change (e.g., “This stock is only going to experience a *small* change tomorrow”), they will do so by translating the same subjective intent into more objective change in the context of a distal (vs. proximal) attractor. More specifically, our results showed that distal (vs. proximal) attractors elongate people’s translation of subjective perception into objective adjustment, thereby also explaining why distal attractors prompted objective differentiation between a “small” and “medium” change than did proximal attractors. Of course, this study is merely consistent with our account, but does not yet provide a strong test of whether attractors influence forecasts of change *because* they distort subjective interpretations of considered changes. Study 5 offers this more definitive test.

Study 5

By our reasoning, attractors affect forecasts of change because they shape people’s subjective representation of what constitutes different objective amounts of change that they might consider. Because each objectively equivalent unit of adjustment toward a distal attractor seems subjectively smaller, adjustment is prolonged. Study 4 supported this account by showing that attractors influence the translation between objective change and subjective interpretation, but we have yet to directly test whether it is this altered mapping that is responsible for attractors’ influence on forecasts of change. Study 5 provided this test by having some participants precommit—before being exposed to the attractor—to the subjective meaning of a certain objective amount of change.

Participants were exposed to the materials used in Study 2, time-series graphs said to depict changes in the average price of gas across days. But before being exposed to the attractor or forecasting the change in gas price, participants saw the Day 1 to Day 9 gas-price trajectory. All participants' attention was directed to the change in price from Day 4 to Day 5—i.e., a change that was merely observed instead of forecast. In the *precommitment* condition, participants were asked to subjectively characterize the magnitude of that (Day 4 to 5) price adjustment. In this way, participants were forced to precommit to a particular subjective interpretation of an objective change before the attractor had the opportunity to exert such an influence on their Day 10 forecast. In the *control* condition, participants estimated the number of eighths-of-an-inch that separated the graphical depictions of the Day 4 and Day 5 prices. In this way, all participants were focused on the same observed change in price, but only those in the precommitment condition took a position on what we expected the attractor to influence when forecasting a new price change: their subjective assessment of a specific objective amount of change. If attractors influence forecasts of change by altering forecasters' subjective interpretations of the same amount of adjustment, then the precommitment manipulation should interfere with this mechanistic pathway and thereby reduce the effect of attractors on forecasts of change. This hypothesized interaction constitutes our central prediction.

A precedent for our general empirical strategy can be found in past research that has likewise examined the latitude people have in moving between subjective labels and objective representations. For example, de Lange et al. (2011) found that the influence of the labels provided on subjective rating scales (more specifically, whether they were written in one's native tongue or a second language) was reduced when objective information (pictorial representations of emotion) was added that more precisely defined the emotional intensity meant to be

communicated by each point on the scale. In this way, the experimenters clarified the objective meaning of each scale point. By analogy, we examine whether having participants precommit to their own objective-subjective mapping will reduce the influence of attractors on their change forecasts.

Method

Participants and design. One hundred sixty-three undergraduates at the University of California, Berkeley, took part in the study. They participated as part of a longer session for which they received \$15 or course credit. Participants were randomly assigned to one of two estimation conditions: *precommitted* or *control*.

Procedure. The initial instructions described a task similar to that encountered by participants in Study 2. But participants in the present study learned that they would make two judgments for each graph. The first varied by condition. The second was the forecast that served as our dependent measure of interest.

All participants first saw a graph showing how gas prices fluctuated over the first nine days. (These were the same graphs used in Study 2 except, as in Study 4, we did not include the minor tick marks). On this initial graph, neither the attractor nor any indication that the gas price increased or decreased on the tenth day was included. Everyone's attention was drawn to the change in price from Day 4 to Day 5 (see Figure 5a).

For participants in the *precommitted* estimation condition, they were asked to indicate their subjective characterization of how much the gas price had changed in that interval. More specifically, they were asked, "What is your subjective sense of how the price of gas moved from Day 4 to Day 5?" Participants responded on a 9-point scale anchored at 1 (a small amount), 5 (a medium amount), and 9 (a large amount). This led participants to commit to a subjective

interpretation of a specific objective amount of change already present in the depicted trajectory, prior to being exposed to the attractor.

Participants in the *control* estimation condition were asked to objectively quantify the distance between the dots representing the Day 4 and Day 5 values: “What is your best estimate of how many eighths of an inch the dots representing Day 4 and Day 5 are apart?” For reference, a small line representing one eighth of an inch was provided. Control participants thus considered the same magnitude of adjustment, but did so in a way that did not prompt them to subjectively interpret that objective amount of change.

After participants had made the initial judgment, the attractor as well as a statement of whether the gas price increased or declined on Day 10 appeared. Participants then used the up and down arrows to adjust the prediction bar from the Day 9 base to their Day 10 forecast (see Figure 5b). As before, we excluded trials from further consideration if participants failed to adjust (1.6%) or adjusted in the wrong direction (2.1%).

Results and Discussion

We used a similar data analytic strategy to that from Study 2. We defined two fixed effects: attractor *distance* (distal: +1, proximal: -1) and our *precommitment* manipulation (precommitted: +1, control: -1). The crucial interaction term (Distance X Precommitment Manipulation) was entered as well. In addition, we included a random effect of participant, which accounted not only for the non-independence of observations from the same participant. Finally, a random effect of gas-price *trajectory* accounted for the fact that some trajectories invited more adjustment than did others.

Consistent with our main hypothesis, the effect of attractor distance depended on—i.e., interacted with—participants’ precommitment condition, $B = -.12$, $SE = .05$, $t(2341.20) = 2.29$, p

= .022. To interpret this interaction, we examined the effect of attractor distance separately for the precommitment and control conditions. In the control condition—in which attractors were fully unfettered in their ability to influence subjective interpretations of adjustment—participants adjusted 0.87 units more toward a distal ($M = 6.44$ units) compared to a proximal ($M = 5.57$ units) attractor, $t(2339.80) = 6.20, p < .001$. But when participants had precommitted to a subjective interpretation of a given amount of change prior to encountering the attractor, the effect of attractor was halved: Participants adjusted only 0.40 units more toward a distal attractor ($M = 5.39$ units) than a proximal attractor ($M = 4.99$ units), $t(2342.44) = 2.70, p = .007$. Stated differently, whereas a distal attractor (compared to a proximal attractor) prolonged control participants' adjustment by 16%, it only prolonged pre-committed participants' adjustment by 8%. In short, the precommitment manipulation cut the influence of attractors about in half.

These results directly build off of our previous findings, showing that attractors influence the amount of adjustment when they are able to influence the interpretation of that adjustment. When participants (in the precommitted condition) had already mapped an objective amount of change onto a subjective interpretation, the attractor exerted a diminished effect on adjustment. This supports our account that attractors influence adjustment at least in part *because* they influence one's subjective sense of just how far one has adjusted. After participants had precommitted to a subjective characterization of an objective amount of change, attractors then had a diminished ability to influence this mapping, which means the attractor's influence was predictably reduced.

One strength of these results is they more conclusively put to rest concerns that attractors influenced adjustment simply because they were mistakenly identified as providing meaningful information. In earlier studies, we described steps taken to demonstrate that attractors were not

meaningful—both by providing justifications for their inclusion and/or including filler trials that showed that attractors did not even meaningfully predict the direction of change. But still, readers may worry that participants had lay theories that the round-number values associated with the attractors served as actual restraints, reference points toward which gas prices, stock prices, or airfare naturally converged. Such alternative interpretations cannot account for the influence of the precommitment manipulation.

Three aspects of the data warrant further comment. First, it is notable that the results from the control condition are almost identical to those from Study 2. In Study 2, the graphs included minor tickmarks (which defined for participants in monetary terms the value of each unit of adjustment), whereas in Study 5, the graphs did not. The near equivalence of the results supports our assumption that this methodological difference has little bearing on the psychology underlying our effect. Second, the pre-commitment manipulation significantly reduced, but did not eliminate, the attractor's influence. This too is hardly surprising. Having indicated that the price change from Day 4 to Day 5 was, say, a 6 on a 9-point scale of subjective magnitude serves to fix somewhat a participant's understanding of what constitutes a small, moderate, or large price change. But it need not fix it entirely, which would give the attractor *some* latitude in defining the subjective magnitude of how far one has adjusted. Alternatively, it is also possible that the attractor's effect is multiply determined, and our manipulation only targeted one of the ways that attractors influence adjustment. Future research may find additional psychological mechanisms that explain attractors' influence on forecasts of change.

Third, the precommitment manipulation not only diminished the influence of the attractor on adjustment, it also diminished the amount of adjustment overall. That is, there was a main effect of precommitment, $B = -0.41$, $SE = 0.12$, $t(160.50) = 3.32$, $p = .001$. Although *a priori* our

hypotheses focused on the interaction, not this main effect, the main effect may (admittedly *post hoc*) be incorporated into our account. Relatively few participants adjusted beyond the attractors, whether distal or proximal. As a result, both attractors may have been serving to reduce participants' subjective sense of the magnitude of any amount of adjustment unless the precommitment manipulation had already "fixed" the meaning of a particular amount of change. If so, distal attractors (for control participants) were simply diminishing the subjective perception of change more strongly than were proximal attractors. This would suggest that attractors—both distal and proximal—ordinarily elongate adjustment, but to different degrees. Of course, this interpretation is speculative, and inclusion of a no-attractor control condition would provide a more definitive answer.

General Discussion

We examined how and why attractors—salient values in the direction of change—influence forecasts of change. In so doing, the present findings document and explain a qualitatively new type of influence on such forecasting. Previous research has focused on how features inherent to the changing attribute itself—e.g., the percentage change, the unit in which change is expressed, the previously observed pattern of change in the attribute—influence evaluations and forecasts of change. Instead, we examined how a feature that is incidental to the change—an externally or internally salient value toward which change happens to be occurring—warps how people subjectively characterize different objective amounts of change. Distal attractors, compared to proximal ones, elongate forecasts of change. When an attractor is further away from (vs. closer to) the anchor value from which adjustment begins, forecasters think that more objective change is needed to achieve the same subjective amount of change. Suggesting that this distortion explains why distal (vs. proximal) attractors invite greater

forecasts of change, preemptively interfering with attractors' ability to shape subjective interpretations of adjustment eliminated more than half of attractors' influence on forecasts.

Our first three studies showed that attractors bias forecasts of change in predictable ways. We tested our hypotheses using complementary paradigms in three contexts: forecasting which of two price changes would be implemented on airline routes (Study 1) and estimating how much the average price of gas (Study 2) or the closing price of a stock (Study 3) would change on a given day. Although variations between the specific paradigms called for slightly different statistical tests, all three supported the hypothesis that the existence of a distal (vs. proximal) attractor makes more (vs. less) change seem likely. These results emerged regardless of whether trajectories were presented in graphical or tabular form, and whether the attractor was highlighted explicitly or was a naturally salient round number.

Our final two studies supported our account of why attractors influence forecasts of change. In particular, when people forecast change in the presence of a distal (vs. proximal) attractor, they translated the same subjective sense of how much change was said to be warranted (e.g., "The stock experienced a small loss") into a larger objective amount of forecasted adjustment (Study 4). Furthermore, the gap between what constituted a "small" and a "medium" change was objectively larger in the context of a distal attractor, supporting the hypothesis that the subjective size of each objectively equivalent unit of adjustment was smaller in the context of a distal than a proximal attractor. Of course, that it takes more objective adjustment to achieve the same subjective amount of change could simply reflect a new adjacent effect, not the mechanism underlying the effects of forecasting documented in Studies 1-3. Crucially, Study 5 provided key evidence that we had hit upon a key process explanation. In that study, we found that when participants precommitted to a certain mapping between a specific objective amount of

change (from a previous period, not the change being forecasted) and a subjective interpretation of that change, the influence of attractors on forecasts was significantly reduced. Notably, this manipulation eliminated about half of the attractor effect, which leaves open the possibility that future research may uncover additional mechanisms underlying the effect of attractors on forecasts of change.

These final two studies are not compatible with an alternative explanation that the perceived distance between the anchor and the proximal or distal attractor directly primed the concepts of “small” or “large,” and those semantically activated concepts were used as guidance in whether a small or large amount of adjustment was warranted. The alternative semantic priming account would have trouble explaining why when participants were directly told that the amount of change takes a certain form (e.g., “small” or “medium”), that this was translated into different objective amounts of change, much less why the difference between a “small” and a “medium” change was objectively larger in the context of a distal than a proximal attractor (Study 4). It also would not explain why Study 5’s precommitment manipulation worked to reduce the influence of the attractor. That manipulation served to constrain the flexibility participants had in mapping objective changes to subjective interpretations of them, but this should have no bearing on the degree to which priming the idea that a change will be “small” or “large” would then result in forecasts of less or more change. In other words, the precommitment manipulation—given it is not a precommitment to how subjectively far one plans to adjust—should have no effect on this alternative mechanism. Of course, whether such semantic priming could itself be an additional mechanism underlying our basic effect could be explored in future research.

We have referred to attractors as arbitrary values, meaning that they are not themselves informative as to the proper forecast. This feature is critical and differentiates the present focus from another literature—that on advice taking (e.g., Gino, 2008)—which examines how those who are anchored (on their own initial perspective) choose to adjust their responses based on a provided value that carries actual information. In what can reflect either an optimism in their own priors that defies rational Bayesian updating (Leong & Zaki, 2018) or a feeling of competition with advisors whom they may see as a threat to their own power (Tost, Gino, & Larrick, 2012), advisees tend to adjust insufficiently toward advice (Harvey & Fischer, 2007) and suffer in their own accuracy as a result (Lim & O'Connor, 1995). We would not conceive of advice as an attractor, for advice takes the form of a meaningful value that *should* spur more adjustment when it is further from one's initially anchored response.

We took three steps to make clear in our own studies that attractors were not meaningfully communicating advice or useful information like the target's long-term average value. First, in several studies we provided explanations as to why certain values were (and by implication, why others were not) made prominent. For example, in Study 1, we explained that only multiples of \$100 (one of which, the attractor, was in the depicted range) would be labeled on the y-axis. Second, Study 3 leaned on the natural salience of round numbers, thereby obviating the need for the explicit presentation of any attractor whose informativeness might be misinferred. Third, in all studies, change was—across the trials—just as likely to occur toward as away from the potential attractor. By including the filler trials (in which change moved away from the attractor), this demonstrated that the externally salient values were not long-term averages toward which the attribute values were converging.

Despite this, perhaps participants still believed that the attractor—explicit or implied—offered information about where the value was ultimately heading. For example, maybe people have a lay theory that values tend to ultimately converge toward round numbers. To address this worry, we conducted a follow-up study to further explore this alternative explanation.

Participants made forecasts on the same 32 graphs used in Study 5, but this time we also recorded how much participants adjusted *away from* the attractor (on filler trials). If the attractor was naturally interpreted as a long-term average, participants should adjust less when moving away from a distal (wrong-direction) attractor than a proximal one. This is because when the (wrong-direction) attractor is distal, the price has strayed especially far from what the alternative account identifies as the long-term average. If this artifactual account is right, further deviation from a wrong-direction distal attractor should be especially unlikely. Although we replicated our finding of more adjustment toward distal than proximal attractors, $t(1,743.31) = 3.56, p < .001$, we found no evidence that wrong-direction attractors influenced adjustment, $t(1,759.02) = -0.65, p = .51$.

Another worry is that even if people do not have a belief that values ultimately converge toward round numbers, they may believe that round numbers serve as “resistance levels” past which change is unlikely to go. Two details—one methodological, one empirical—speak against this possibility. Because trajectories paired with proximal attractors were, by design, closer to those attractors, the supplied historical trajectory of the price was more likely to have crossed the attractor, thereby demonstrating that the value can move on both sides of the attractor. If participants began by thinking that attractors reflected rigid resistance levels, such participants would have been more likely to have been disabused of this notion during the proximal than the distal attractor trials (see Figure 3, for one example). This would have worked against our

hypotheses. Second, the resistance level alternative cannot explain why the precommitment manipulation in Study 5 significantly reduced the attractor effect. Of course, it is also possible that people do have lay theories about how round numbers can serve as resistance levels, but attractors seem to influence forecasts of change independently of such beliefs.

That said, one open question is what constraints exist on whether a value serves as an attractor in the first place. In considering this issue, we think it useful to differentiate the properties of relevance and extremity. By relevance, we mean whether the attractor value is introduced in the context of the forecast being made. Note that a minor axis on a graph offers a value that is by this definition relevant (even if not informative): A \$50 tick mark identifies a candidate (even if not a likely) value that a currently \$42 stock could achieve. A \$50 dinner tab may make the same value salient, but the categorical irrelevance of the value may make it less likely to serve as an attractor. Even when attractors are relevant, they may be particularly implausible. Much as researchers have demonstrated that anchoring effects are just as strong when anchors are implausible because they are extreme (Strack & Mussweiler, 1997), one can ask the same about attractors. We suspect that extreme attractors are relatively unlikely to be spontaneously generated, but it is an open question to what extent they influence forecasts when they are environmentally salient. That is, when a considered amount of change does not do much to “close the gap” with an extremely distal attractor, this may not inform one’s sense that the considered change is insubstantial, but may simply reinforce the irrelevance of the arbitrary cue. Clearly, future research is necessary to examine *which* values serve as forecast-distorting attractors at all.

Two primary limitations of our studies are that: 1) participants were amateur forecasters who 2) had little valid information on which to base their forecasts beyond 9 periods of historical

returns. We suspect that both features may have contributed to effects of attractors on forecasts that would be larger than what better-informed experts would display in professional forecasting contexts. Combining across our conditions that lacked features meant to constrain or give subjective guidance about the impending amount of change, we found distal (compared to proximal) attractors prolonged change estimates by an average of 12%. Given stock market fluctuations depend on where investors think the market is going, does this really imply that the United States' 7.2% average stock market return could have been boosted to 8.1% through some strategically placed y-axis labels, leading 40-year market returns to increase 22 fold instead of 16 fold? Although such sensational speculation is tempting, we caution against it. Like many laboratory effects, these findings should perhaps best be thought of as a theoretical upper limit on what attractors can do. Future research should investigate the influence of attractors on more expert forecasters in more informationally rich contexts. We suspect attractors will have a diminished effect.

Comparisons with and Theoretical Implications for Other Literatures

Stimulus evaluation. The present work is both consistent with and distinct from research and theory in cognitive psychology that has emphasized that there is not a one-to-one correspondence between objective values (e.g., inches, years, dollars) and how they are subjectively interpreted (e.g., narrow, soon, pricey). Parducci's (1965) frequency-range theory posits that people assign different subjective labels (e.g., Likert scale values from 1[short] to 9[tall]) when describing a stimulus's property (e.g., height) depending on the range and distribution of stimuli to which they have been exposed (see also Ostrom & Upshaw, 1968). A 5'5" male may be labeled as a 9 in a room of eleven-year-old boys, but a 1 in a room of forty-year-old men. Similarly, subjective ratings are sensitive to the range of permissible objective

responses. For example, after participants committed to giving a harsh or a lenient punishment, they gave objectively longer sentences when they learned the maximum permissible sentence was 30, as opposed to 5, years (Ostrom, 1970).

Our research departs from these existing findings and frameworks in two ways. First, whereas previous research examined how people evaluate stimuli when they are given different ranges of permissible objective responses, we instead examine how people evaluate stimuli in the context of the *same* range of possible responses. That is, the attractor values in our studies changed neither: 1) the possible range into which participants could adjust, nor 2) the range of comparison stimuli (e.g., historical stock prices) to which they were exposed. To the first point, we were careful in Studies 1-2 and 4-5 to always center the base value on the y-axis, so that the number of units of potential change in each direction was held constant. In this way, we did not simply demonstrate a visual trick well-honed by those who trade in misleading bar graphs: making differences seem larger by expanding or contracting the plotted range. (And by leaning on a tabular presentation, Study 3 removed these graphs altogether.) To the second point, regardless of whether participants made estimates in the context of a proximal or distal attractor, they were exposed to the same range of previous values. More specifically, in every study, each previous price was the same number of units from the base price from which adjustment proceeded regardless of the attractor's position, meaning that the range of observed stimuli was also equivalent. Second, whereas previous research has examined how people view an individual stimulus (e.g., a square, a prison sentence) in the context of a range of other individual stimuli (e.g., a set of squares, a set of possible prison sentences), the present studies examined how a comparison between stimuli (i.e., an amount of change) was seen differently in the context of another stimulus (i.e., a distal or proximal attractor.)

One question is whether the present findings could be seen as an extension of this past work if one merely identified the attractors themselves as another observed value—a value that was salient in the judgment context, even if it was not one of the stimulus values observed in the time-series graph—that helped to determine what values were subjectively large or small. But even this would not seem to capture the novel attractor effect documented herein. Consider again our follow-up study in which we showed that attractors in the opposite direction of adjustment failed to influence how much change was expected. In other words, even by leaning on more inclusive criteria for what set of values define the frequency or range of observed values, it seems that this would not allow one to easily assimilate the present effects into such previously articulated theories. Instead, attractors seem to operate by a distinct mechanism: the tendency for salient values in the direction of adjustment to warp subjective interpretations of different objective amounts of change to which forecasters are considering committing.

Anchoring. When anchoring was first introduced, it was assumed that anchoring phenomena were the result of the psychological process of anchoring and (insufficient) adjustment. By this understanding, anchors serve as a starting point for judgment, but because attempts to adjust away from the anchor are insufficient, anchors exert an assimilative pull on responses. Anchoring and (insufficient) adjustment can influence numeric judgments (e.g., Epley & Gilovich, 2001), trend estimation (e.g., Eggleton 1982), or attempts to adopt another person's perspective (e.g., Epley, Keysar, Van Boven, & Gilovich 2004). And although there have been numerous paradigms to test, and proposed mechanisms to explain, why arbitrary numbers anchor judgments, the present paper offers the first paradigm (of what may be many more) and accompanying mechanistic account that demonstrates and explains why arbitrary values influence how much adjustment is seen to be warranted.

Still, one may ask whether it is useful to introduce the new concept of attractors, or whether they are instead better conceived of as second anchors. After all, if distal attractors prolong adjustment, don't both anchors and attractors exert assimilative pulls on judgment? The shortcoming of equating anchors and attractors is that attractors—either definitionally or as demonstrated here, empirically—do not operate like anchors. For example, one reason anchors exert an assimilative pull is adjustment from such anchors tends to be insufficient. But attractors are not starting points for adjustment. When participants were deciding by how much a stock price (Study 4) or the price of gas (Study 5) fell, this required participants to decide (and actually implement through iterative downward adjustments) how much the value declined, not how much the value increased from the base of the attractor.

If forecasters merely anchored on the attractor as a first guess of the to-be-forecasted value, it is unclear why the gap between a “small” and a “medium” amount of change would be greater in the context of a distal than a proximal attractor (Study 4). Also, and even more convincingly, it would be unclear why Study 5's precommitment manipulation would dampen the attractor's influence. That is, we can think of no reason why subjectively characterizing an objective amount of change would then weaken an anchor's tethering hold. For these reasons, we think that simply calling the attractor an anchor (or a second anchor) would gloss over the meaningfully different role played by the base value from which change originates (and thus defines one end of a forecasted change) from the arbitrary value in the direction of change (that colors the subjective interpretation of adjustment that heads toward it).

It is natural to ask whether attractors apply only to the task of forecasting a change, or whether they would apply to other anchoring-and-adjustment problems as well. Arguments could be made either way. Consider what differentiates forecasts of change from other anchoring-and-

adjustment phenomena. In answering the trivia question, “How many days does it take Mars to orbit The Sun?”, people self-generate the anchor 365 days (a readily available value known to be both relevant and wrong), and then adjust up from there (Epley & Gilovich, 2006). In this case, the end goal is merely to estimate Mars’s orbit. The amount one has adjusted does not itself have significance (i.e., we have no term for Mars’s minus Earth’s orbit). In contrast, when forecasting change, the amount of adjustment is itself meaningful; it constitutes the forecasted change. In fact, the change may be more important than the end value: A stock investor may care more about the change in a share’s value (their return), not its end valuation (the final share price.) Because we showed that attractors change whether or not a certain amount of adjustment seems subjectively substantial and thus sufficient, it may be that attractors exert a bigger effect on forecasts of change (for which one is very focused on the amount of adjustment itself, perhaps even more so than on the end value) as opposed to other anchoring-and-adjustment problems (for which adjustment is simply a means to an end value). On the other hand, we know that even in standard anchoring-and-adjustment problems, people are essentially asking themselves, “Does it feel like I have adjusted far enough?”, thereby suggesting that attractors’ influence—given their ability to change people’s subjective characterization of the same objective amount of adjustment—may generalize.

We conducted a study to begin to explore this question. Details of this Supplemental Study are reported in the Supplemental Materials. In brief, participants adjusted from a base value to a final answer either to estimate how the price of gas was likely to change in a given month (a forecast of change) or to estimate the price of gas at another gas station (a more traditional anchoring-and-adjustment paradigm, given the adjustment itself is merely a means to offering an end response, meaning it lacks an inherent meaning.) We found that attractors

influenced adjustment similarly in each case. Although this provides initial evidence that attractors can influence judgments more generally, we also encourage readers to keep in mind that we would not expect attractors to influence responses (at least by the mechanism documented in the present paper) in anchoring paradigms that do not actually involve adjustment. For example, Tversky and Kahneman's (1974) classic anchoring paradigm—at least by the selective accessibility account (Mussweiler, 2003; Strack & Mussweiler, 1997)—involves no adjustment (cf. Simmons, Lebouef, and Nelson 2010). Without adjustment, it is hard to understand how our demonstrated mechanism of distorting the subjective interpretation of adjustment could operate.

It is also worth considering how the mechanism we established is compatible with and divergent from other mechanisms accounting for why anchors restrain judgments. Frederick and Mochon (2012) presented a scale-distortion theory of anchoring that predicts, for example, that people will give a lower estimate of a giraffe's weight if they first estimate the weight of a raccoon. Committing to a low weight for a "small animal" appears to distort the meaning of the numeric weight scale in the direction of the anchor. As a result, one then uses a relatively small number to describe the weight of a "big animal." Despite our shared emphasis on the latitude inherent in mapping subjective representation onto an objective scale, the effects themselves as well as the basic mechanisms differ. Frederick and Mochon might predict that participants in our Study 4 would interpret a "medium gain" as less substantial (i.e., requiring less adjustment) after having considered a "small gain," but their account does not entail that the subjective interpretation of a medium or small gain would shift in response to an attractor's location.

In other research, Janiszewski and Uy (2008) found that the precision with which an anchor is stated can influence adjustment. People tend to adjust further from a round anchor

value (“10”) than from a more precise anchor (“9.8” or “10.2”). Janiszewski and Uy speculated that more precise anchors encourage adjustment in smaller units, which results in less distance covered after the same number of iterative adjustments. Our account likewise emphasizes how features of the judgment context can influence how numerical space is psychologically partitioned, but our research differs in terms of what features influence that partitioning (the anchor’s precision vs. the attractor’s location) and how this translates into different amounts of adjustment. Furthermore, whereas Janiszewski and Uy emphasized that adjustment may occur in shorter versus longer iterative leaps, we emphasize how attractors influence people’s holistic assessments of a given amount of adjustment, even when participants are not doing the adjusting (and thus the iterative leaps) themselves (Study 1).

Practical Implications

Although our focus has been on how and why attractors influence forecasts of change, we hope future research will explore the practical behavioral consequences of attractors’ influence. For example, might a retirement advisor create a greater sense of urgency to save more money by depicting a potential fund as trending upward toward a distal (instead of a proximal) attractor? Might public health officials help to dissuade policymakers from prematurely easing social distancing measures during a pandemic if they show infection rates declining toward a proximal versus a distal attractor? If so, attractors may serve as useful nudges that accompany otherwise equivalent presentations of information.

Context Paragraph

Anchoring-and-adjustment has been of longstanding interest to psychologists, economists, and practitioners. Over five decades, the term “anchoring” has ballooned to include almost any example in which a piece of information exerts a non-normative, assimilative pull on judgment. This liberal conceptualization has contributed to this literature’s breadth and richness. But characterizing so many distorting stimuli as “anchors”—whose biasing effects actually reflect unrelated, independent processes—reduces the single term’s usefulness. We started this project by considering how arbitrary values may not simply be initial anchors whose influence is hard to escape, but cues that shape the interpretation of adjustment from a starting value. This moves beyond conceiving of adjustment as the effort to break free of an anchor’s influence, to instead consider adjustment as the vehicle through which an arbitrary value’s influence is realized. We initiate these efforts in a domain (forecasting change) for which the judgment of interest is itself an amount of adjustment, a common and broadly applicable context to consider an arbitrary influence on adjustment. Although we appreciate the temptation to characterize the present efforts as documenting another (perhaps a fifth!) mechanism by which anchoring occurs, we hope this section provides context for why we caution against this.

References

- Adaval, R., & Wyer, R. S., Jr. (2011). Conscious and nonconscious comparisons with price anchors: Effects on willingness to pay for related and unrelated products. *Journal of Marketing Research*, 48, 355-365.
- Brannon, E. M., & Terrace, H. S. (1998). Ordering of the numerosities 1 to 9 by monkeys. *Science*, 282, 746-49.
- Critcher, C. (2020, April 3). Attractors: Incidental Values That Influence Forecasts of Change. Retrieved from osf.io/z9h6y
- Critcher, C. R., & Gilovich, T. (2008). Incidental environmental anchors. *Journal of Behavioral Decision Making*, 21, 241-251.
- De Langhe, B., Puntoni, S., Fernandes, D., & Van Osselaer, S. M. J. (2011). The anchor contraction effect in international marketing research. *Journal of Marketing Research*, 48, 366-380.
- Eggleton, I. R. C. (1982). Intuitive time-series exploration. *Journal of Accounting Research*, 20, 68-102.
- Englich, B., & Mussweiler, T. (2001). Sentencing under uncertainty: Anchoring effects in the courtroom. *Journal of Applied Social Psychology*, 31, 1535-1551.
- Epley, N., & Gilovich, T. (2001). Putting adjustment back in the anchoring and adjustment heuristic: Divergent processing of self-generated and experimenter-provided anchors. *Psychological Science*, 12, 391-396.
- Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological Science*, 17, 311-318.

- Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking and anchoring and adjustment. *Journal of Personality and Social Psychology, 87*, 327-339.
- Fechner, G. T. (1860). *Elemente der Psychophysik* [Elements of psychophysics]. Leipzig, Germany: Breitkopf und Härtel.
- Frederick, S., & Mochon, D. (2012). A scale distortion theory of anchoring. *Journal of Experimental Psychology: General, 141*, 124-133.
- Gino, F. (2008). Do we listen to advice just because we paid for it? The impact of advice cost on its use. *Organizational Behavior and Human Decision Processes, 107*, 234-245.
- Harvey, N. (1997). Use of heuristics: Insights from forecasting research. *Thinking & Reasoning, 13*, 5-24.
- Harvey, N., Bolger, F., & McClelland, A. (1994). On the nature of expectations. *British Journal of Psychology, 85*, 203-229.
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment and sharing responsibility. *Organizational Behavior and Human Decision Processes, 70*, 117-133.
- Harvey, N., & Reimers, S. (2013). Trend damping: Under-Adjustment, experimental artifact, or adaptation to features of the natural environment?" *Journal of Experimental Psychology: Learning, Memory, and Cognition, 39*, 589-607.
- Hutchinson, J. W., Alba, J. W., & Eisenstein, E. M. (2010). Heuristics and biases in data-based decision making: Effects of experience, training, and graphical data displays, *Journal of Marketing Research, 47*, 627-642.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin, 21*, 1161-1166.

- Janiszewski, C., & Uy, D. (2008). Anchor precision influences the amount of adjustment,” *Psychological Science, 19*, 121-127.
- Lawrence, M., Goodwin, P., O’Connor, M., & Onkal, D. (2006), Judgemental forecasting: A review of progress over the last 25 years. *International Journal of Forecasting, 22*, 493-518.
- Lawrence, M., & Makridakis, S.(1989). Factors affecting judgmental forecasts and confidence intervals. *Organizational Behavior and Human Decision Processes, 43*, 172-187.
- Leong, Y. C., & Zaki, J. (2018). Unrealistic optimism in advice taking: A computational account. *Journal of Experimental Psychology: General. 147*, 170-189.
- Lim, J. S., & O’Connor, M. (1995). Judgmental adjustment of initial forecasts: Its effectiveness and biases. *Journal of Behavioral Decision Making, 8*, 149-168.
- Mussweiler, T. (2003). Comparison processes in Social Judgment: Mechanisms and Consequences,” *Psychological Review, 110*, 472-489.
- Nunes, J. C., & Boatwright, P. (2004). Incidental prices and their effect on willingness-to-pay,” *Journal of Marketing Research, 14*, 457-466.
- Oppenheimer, D. M., LeBoeuf, R. A., & Brewer, N. T. (2008). Anchors aweigh: A demonstration of cross-modality anchoring and magnitude priming, *Cognition, 106*, 13-26.
- Ostrom, T. M. (1970). Perspective as a determinant of attitude change. *Journal of Experimental Social Psychology, 6*, 280-292.
- Ostrom, T. M. , & Upshaw, H. S. (1968). Psychological perspective and attitude change. In A. C. Greenwald, T. C. Brock, & T. M. Ostrom (Eds.), *Psychological Foundations of Attitude* (pp.217-241). New York, Academic Press.

- Pandelaere, M., Briers, B., & Lembregts, C. (2011). How to make a 29% increase look bigger: The unit effect in option comparisons. *Journal of Consumer Research*, *38*, 308-322.
- Parducci, A. (1965). Category judgment: A range-frequency model, *Psychological Review*, *72*, 407-418.
- Pope, D., & Simonsohn, U. (2011). Round numbers as goals: Evidence from baseball, SAT takers, and the lab. *Psychological Science*, *22*, 71-79.
- Reimers, S., & Harvey, N. (2011). Sensitivity to autocorrelation in judgmental time series forecasting. *International Journal of Forecasting*, *27*, 1196-1214.
- Simmons, J. P., LeBouef, R. A., & Nelson, L. D. (2010). The effect of accuracy motivation on anchoring and adjustment: Do people adjust from provided anchors? *Journal of Personality and Social Psychology*, *99*, 917-932.
- Simonson, I., & Drolet, A. (2004). Anchoring effects on consumers' willingness-to-pay and willingness-to-accept. *Journal of Consumer Research*, *31*, 681-690.
- Strack, F. & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, *73*, 437-446.
- Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, *1*, 39-60.
- Theochari, Z. (2014). *Judgmental forecasting from graphs and from experience*. [Doctoral dissertation, University College London].
- Tost, L. P., Gino, F., & Larrick, R. P. (2012). Power, competitiveness, and advice taking: Why the powerful don't listen. *Organizational Behavioral and Human Decision Processes*, *117*, 53-65.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases.

Science, 185, 1124-1130.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice.

Science, 211, 453-458.

Wertenbroch, K., Soman, D., & Chattopadhyay, A. (2007). On the perceived value of money:

The reference dependence of currency numerosity effects. *Journal of*

Consumer Research, 34, 1-10.

Wong, K. F., & Kwong, J. Y. (2000). Is 7300 m equal to 7.3 km? Same semantics but different

anchoring effects. *Organizational Behavior and Human Decision Processes, 82*, 314-

333.

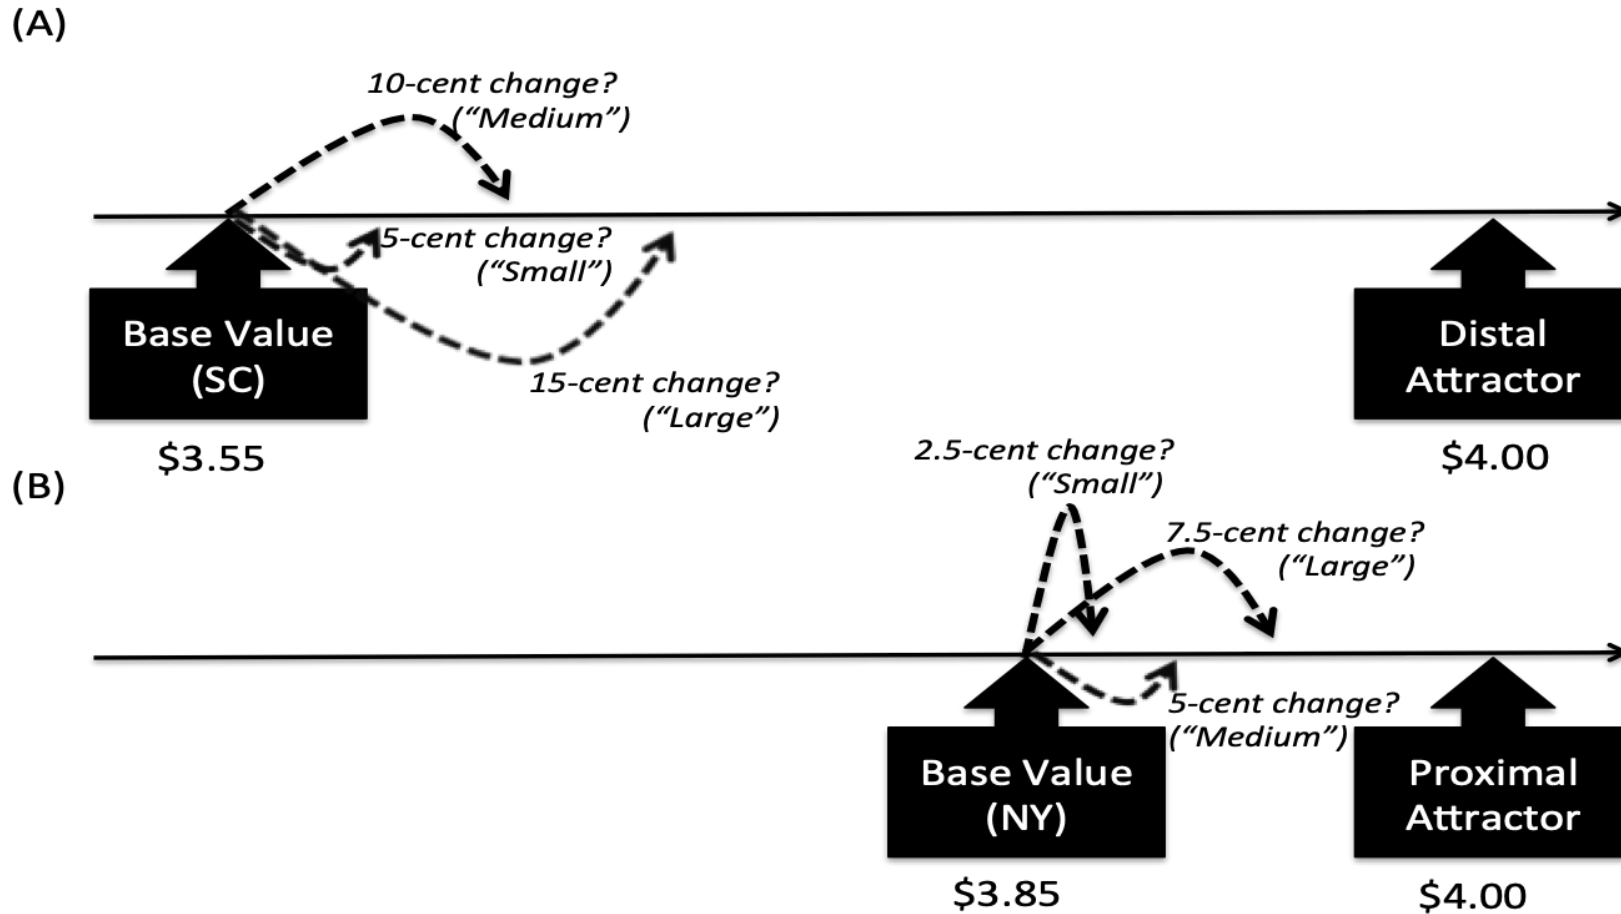


Figure 1. By our account, the same prospectively contemplated price adjustment (5 cents) will be made to feel subjectively smaller when a distal attractor is salient (A) instead of when a proximal attractor is salient (B). Note how in panel A, the same subjective labels are mapped to larger objective changes. As a result, and all else equal, people will tend to forecast more change when adjusting toward a distal attractor compared to a proximal attractor. Furthermore, the amount of objective change that differentiates two subjectively characterized changes (e.g., small and medium) should be greater in the context of a distal than a proximal attractor.

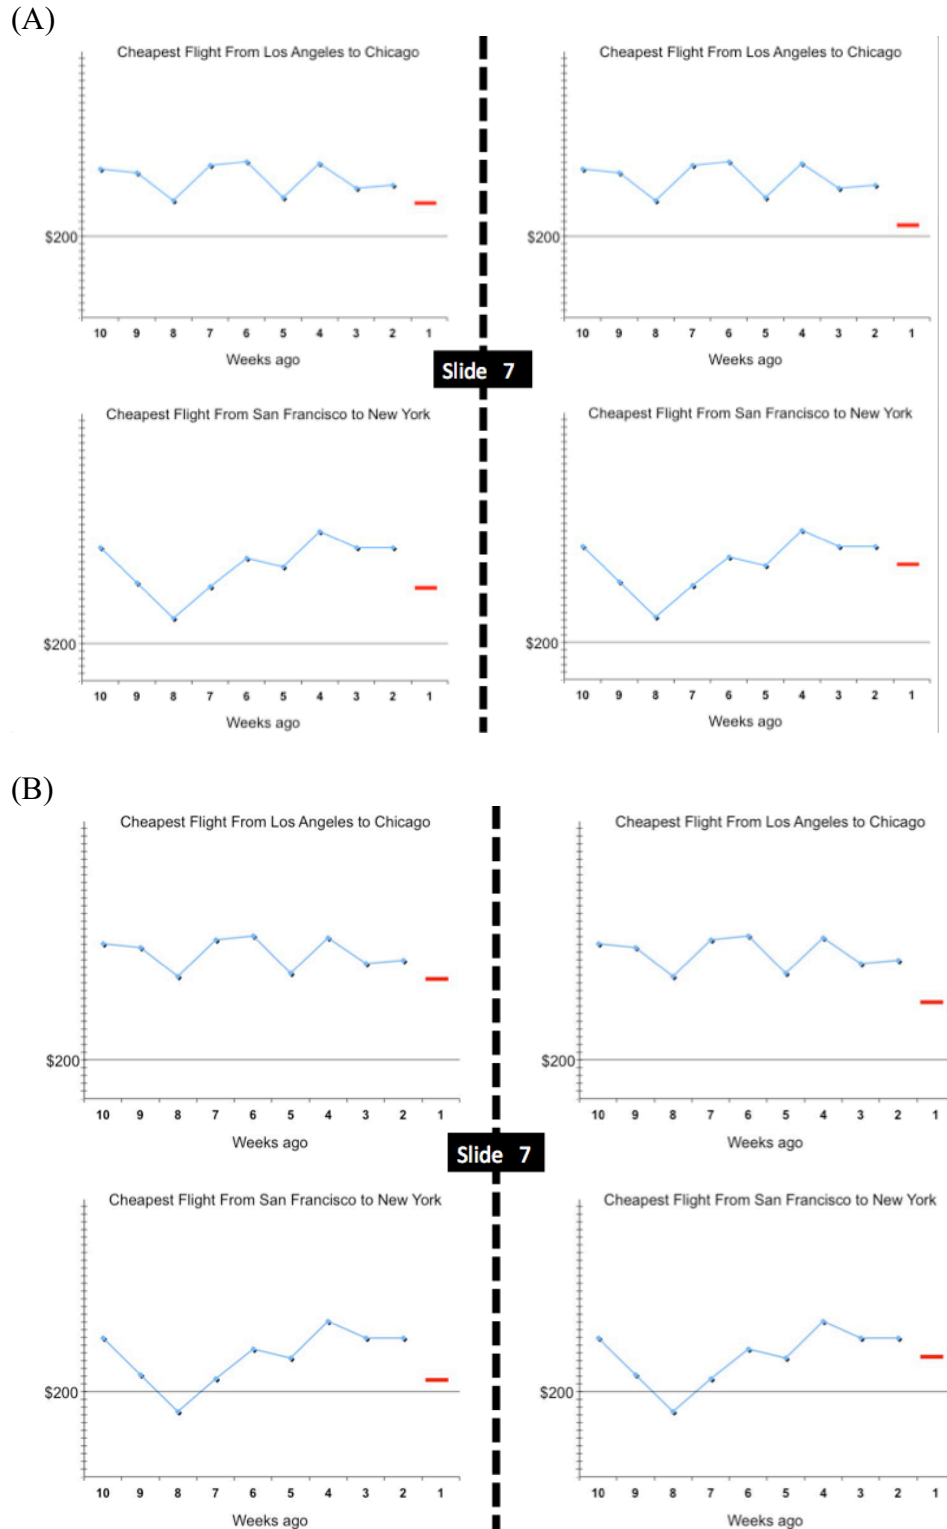


Figure 2. Each participant in Study 1 saw either (A) or (B). The hypothesis-matching half—pairing a distal [proximal] attractor with a larger [smaller] objective price change—is the left half of Panel A and the right half of Panel B. In Panel A [B], the two top [bottom] graphs include a proximal attractor; the bottom [top] graphs contain a distal attractor.

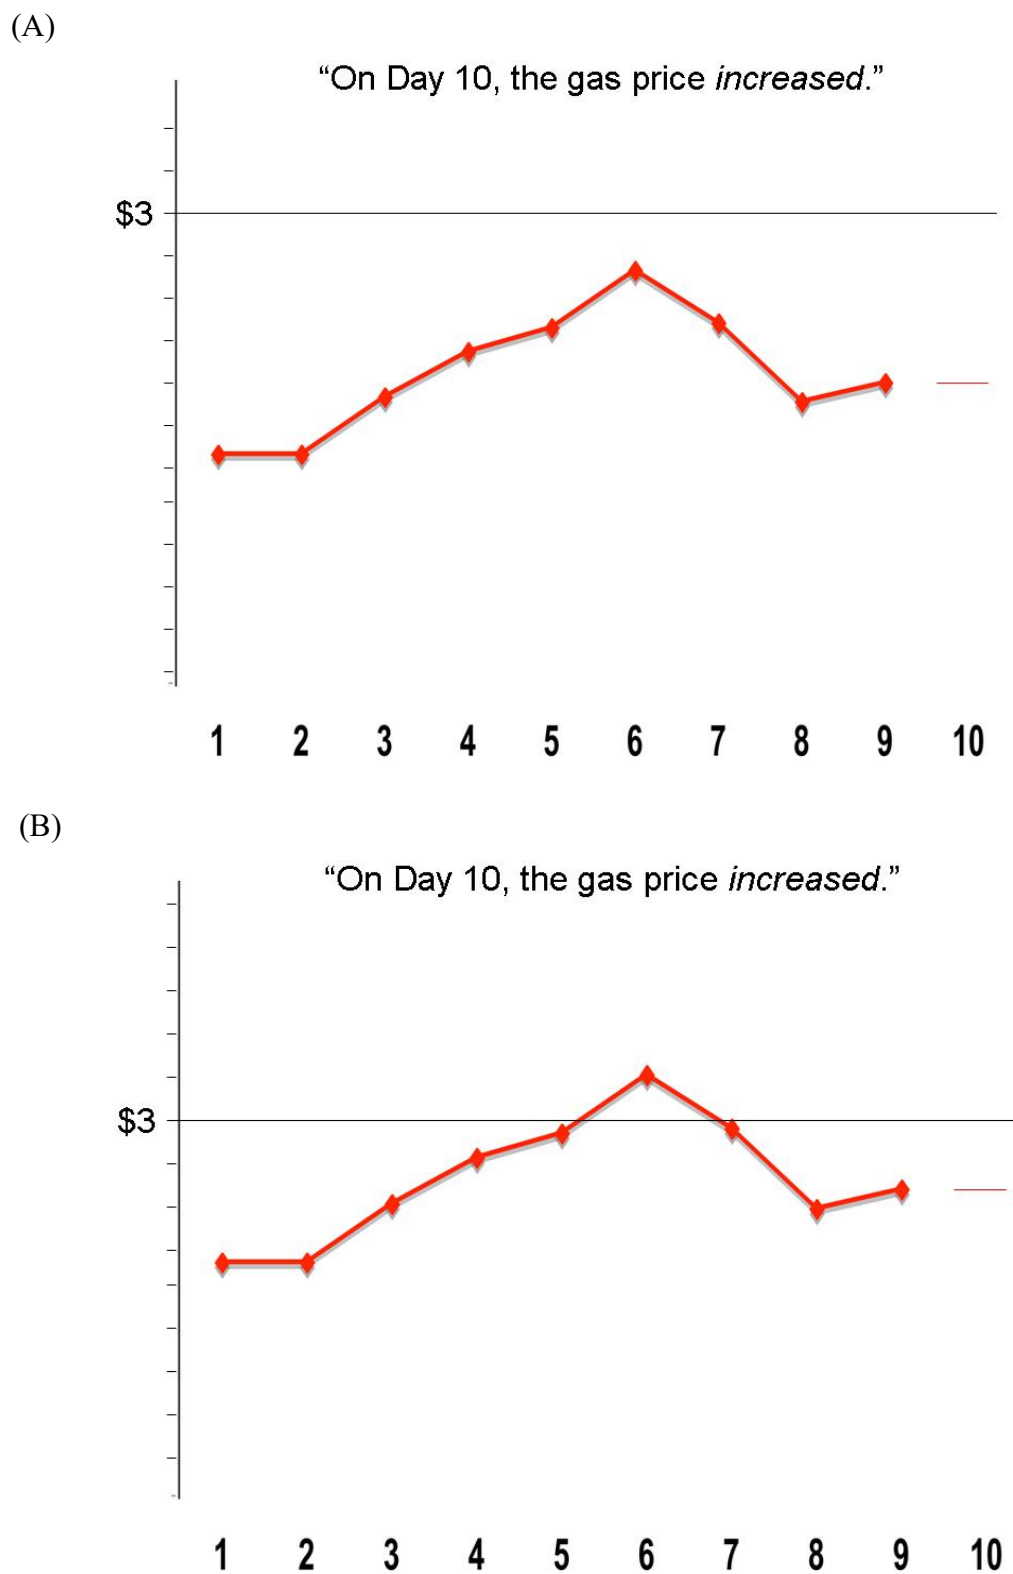


Figure 3. One of the eight key gas-price trajectories seen by Study 2 participants. Participants saw both attractor versions—distal (A) and proximal (B)—though always in separate halves of the 32 trials.

A. Distal attractor

Day	1	2	3	4	5	6	7	8	9	10
Closing Price	26.36	26.62	26.88	26.70	26.83	26.87	26.88	26.52	26.45	<i>lower</i>

B. Proximal attractor

Day	1	2	3	4	5	6	7	8	9	10
Closing Price	26.20	26.46	26.72	26.54	26.67	26.71	26.72	26.36	26.29	<i>lower</i>

Figure 4. One of the eight stock price trajectories seen by Study 3 participants. In each panel, the stock price is said to be moving toward the closest whole number value (26), but is currently relatively distal from (A) or proximal to (B) that round-number attractor.

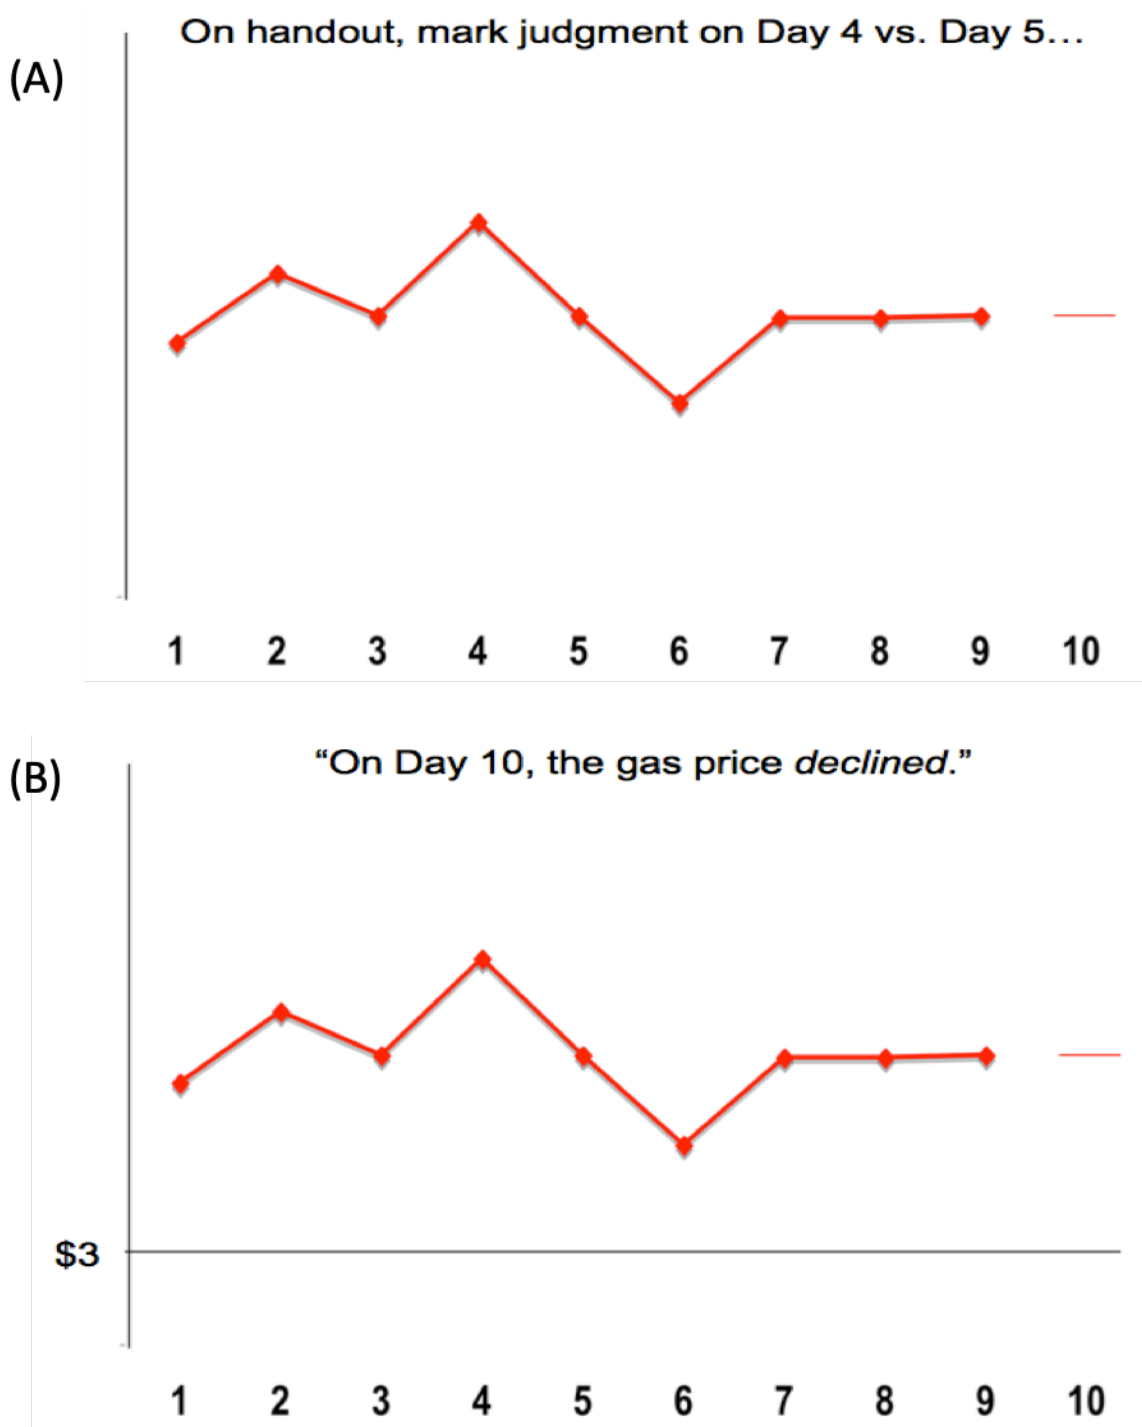


Figure 5. Participants in Study 5 were asked to characterize the Day 4 to Day 5 price change (A) in a way that asked them to precommit to a subjective interpretation of an objective amount of change or to judge the distance between the two points. Only following either judgment did the attractor and indicated direction of adjustment appear, at which point participants were to forecast how the price changed from Day 9 to Day 10 (B).