The Commonness Fallacy: Commonly Chosen Options Have Less Choice Appeal Than People Think

Emily S. Reit
Stanford University

Clayton R. Critcher
University of California, Berkeley

In predicting what others are likely to choose (e.g., vanilla ice cream or tiramisu), people can display a *commonness fallacy*—overestimating how often common (but bland) options (e.g., vanilla ice cream) will be chosen over rarer (but exciting) options (e.g., tiramisu). Given common items are often chosen merely because they are frequently offered, not because they are preferred (tiramisu is rarely offered as a dessert), commonness is not necessarily diagnostic of future choice. Studies 1a and 1b document the commonness fallacy in forecasts of single and repeated choices. Study 2 replicates it in an incentive-compatible choice context. Studies 3 and 4 uncover when and why perceived commonness is relied upon. Perceived commonness is spontaneously used as a guide when forecasting others' choices (as though people blur what has been chosen with what people will choose), but not when forecasting what others would be pleased to receive. Choice forecasters leaned upon perceived commonness over and above many other cues, including their own choices, the goods' prices, and even how much others were thought to like each option. Upon conscious reflection, choice forecasters abandon commonness and gravitate toward more normatively defensible input. Studies 5 and 6 used correlational and experimental methods, respectively, to examine antecedents of the commonness fallacy. Study 7 illustrates a literally costly consequence: A 2-part marketplace simulation study found amateur sellers' reliance on perceived commonness prompted them to systematically misprice goods.

Keywords: choice, heuristics, perspective taking, social judgment, theory of mind

Supplemental materials: http://dx.doi.org/10.1037/pspa0000172.supp

People are often tasked with predicting others' choices. A dinner host must decide how many servings of vanilla ice cream versus tiramisu to have on hand. A shoe salesman must decide how many of the new shoes to carry in blue or black. A film festival organizer must decide whether to put the indie film or the mainstream blockbuster in the larger theater.

Knowing what others will choose requires knowing their preferences, but knowing others' preferences is difficult. In understanding others, a readily accessible guide is the self (Krueger, 2000). This leads people to display a false consensus effect: For example, in estimating what percentage of college students prefer French or Italian movies, participants consulted their own prefer-

ences (Ross, Greene, & House, 1977; see also Rogers, Moore, & Norton, 2017). Although such projection can helpfully inform social knowledge (Dawes & Mulford, 1996; Krueger, 2003), it is an incomplete guide. Perfect social insight requires not only that people know that parts of their own preferences are idiosyncratic ("I have to remember not everyone thinks cilantro tastes like soap"), but also that they actually expend the effort to adjust from their own egocentric perspective (Epley, Keysar, Van Boven, & Gilovich, 2004).

For all of the difficulties in predicting others' preferences, there are additional challenges in predicting *choice*. In part, this is attributable to underappreciated factors that operate on choosers themselves—factors that lead their choices to diverge from their internal values or preferences. For example, embarrassment can exert a surprisingly strong hold that keeps people's natural inclinations in check (Bohns & Flynn, 2010). Impression management goals encourage restaurant patrons dining in groups to place different orders than if they were dining alone (Ariely & Levav, 2000; Quester & Steyer, 2010). Failing to understand these contextual or situational factors can stymie forecasters. Epley and Schroeder (2014) illustrate this point with a social choice: Even when people are open to connecting with a stranger, they frequently abstain out of a misguided assumption that such overtures are unwanted, thereby leaving unquestioned that solitude is the universal preference. States of pluralistic ignorance about the true pref-

erences of others—those that would guide their private choices—result (Prentice & Miller, 1993).

In this paper, we concur that forecasting others' choices is hard. But instead of considering how choosers themselves are influenced by many difficult-to-appreciate forces, we focus on a cause that operates on the forecasters themselves. That is, we argue that when forecasting others' choice of A or B, forecasters naturally rely on a cue that is not necessarily diagnostic of choice, thereby potentially leading to errors in judgment. Specifically, we propose that in estimating others' choice of A or B, people demonstrate a commonness fallacy—a reliance on their intuitive sense of the commonness with which people consume, use, or partake in A compared with B. For example, in the United States spaghetti and meatballs is commonly consumed for dinner, whereas chorizo-and-cheese stuffed bell peppers is more unusual. We suggest that forecasters may overestimate diners' choice of the former over the latter when diners are given the choice of the two.

The Commonness Fallacy

We argue that the commonness fallacy—much like the false consensus effect—is a broad social judgment error that stems from the overapplication of a heuristic cue. Although the false consensus effect is not typically discussed as a heuristic, when people project information about the self onto others, they essentially display a *self heuristic*. That is, they rely on what is true of the self in deciding what is true of others. We see the commonness fallacy as a complementary phenomenon, one in which forecasters lean on a *commonness heuristic* in forecasting others' choice.

A hallmark of a heuristic is that it involves attribute substitution—reliance on an imperfectly valid but readily accessible cue when making difficult, potentially intractable judgments (Kahneman, 2003; Kahneman & Frederick, 2002). But an imperfectly valid cue can often be valid. Indeed, various scholars have argued that those operating in the *heuristics* and *biases* tradition have been too eager to link those two namesake ideas (Gigerenzer, 2008; Gigerenzer & Brighton, 2009; Gigerenzer & Gaissmaier, 2011; Tversky & Kahneman, 1974). That is, heuristics do not merely lead people toward but often keep them away from bias. By understanding the imperfect but sometimes valid logic that underlies many heuristics, one can see why they are frequently useful. But also, understanding the holes in the underlying logic helps to identify the circumstances when the heuristic will lead judgments astray.

For example, when X (e.g., the size of a city) causes Y (e.g., the familiarity of a city name), Y can be taken as a valid indicator of X. Most people are more familiar with Paris, France, than Paris, Texas. Leaning on this recognition heuristic helps identify which city is bigger (Gigerenzer & Goldstein, 1996). But sometimes X is not the only cause of Y. In some cases variable Z (e.g., whether a famous historical event occurred in the city) also influences Y. We suspect the European Waterloo is more widely recognized than the Canadian one, even though the latter has a population more than three times the former's. Sometimes people spontaneously recognize when heuristic cues have been contaminated (Oppenheimer, 2004), but often they do not (Kahneman, 2003).

This logic demonstrates why perceived commonness is enticing as a mental shortcut in forecasting choice, but also why it can produce a commonness fallacy. When people's preferences lead them to seek out an item (X), then this does lead an item to become more common (Y). As an example, when we say vanilla ice cream is more common than tiramisu, we do mean people choose to eat the former more often than they choose to eat the latter. But this does not imply that people are likely to choose vanilla ice cream *over* tiramisu when given the choice between the two. After all, vanilla ice cream is not merely a more common dessert choice, it is (for a variety of reasons) more commonly offered as an option to begin with.

We argue that people lean on options' commonness (how often they have been chosen before) because it is deceptively similar to the forecast of interest (how often they will be chosen). Furthermore, people can lean on perceptions of commonness gleaned from their more general experience (People eat vanilla ice cream more often than they eat tiramisu . . .) even when they do not have information about the choice history of their specific target of judgment (. . . so I'll be sure to have a lot of vanilla ice cream on hand for my dinner guests I've never met before.) If this logic is correct, it implies that commonness may be used as a cue to estimating *choice* in particular, but less so (or maybe not at all) when considering other preference indicators (e.g., what people like, what they would be pleased to receive). It is worth noting that one remarkable thing about this prediction is its superficial inconsistency with a lesson from past research. Helzer and Dunning (2012) find that social forecasts can be superior to self forecasts because social judgments tend to give weight to past behavior, whereas self-forecasts lean on (overly-)lofty ambitions. The present research differentiates itself by arguing that a reliance on the perceived commonness of past behavior can lead people astray.

A final question we will broach is where perceptions of commonness originate. People might be guided by their own experience or familiarity with each choice option, the depth of their knowledge about what each option is, their memories of seeing others use or engage with each option, their knowledge of how much shelf space each option tends to get (e.g., how large the Miller Lite vs. Modelo display is at one's local grocery store), or even how typical of an example each option is of the choice category. In other words, to lean on perceived commonness when estimating what others will choose, people may recruit other heuristics (e.g., availability, familiarity) to inform their perceptions of commonness. We will return to this layered question with a combination of correlational and experimental evidence. But for most studies, we will defer to participants' own intuitive judgments (regardless of their origin) of stimuli's commonness.

Are People Swayed by a Heuristic? Complementary Methodological Approaches

The focus of this article is to provide evidence for a novel social heuristic, understand when and why people rely on it, and explore its downstream implications. But how empirically does one document that people lean on a heuristic? Researchers can take one of two methodological approaches, each with its own advantages and disadvantages. To simplify the presentation of our studies, in which both approaches are used, we give a name to these two methods:

The Systematic-Bias Approach

When people lean on a heuristic, it sometimes pushes them toward accuracy and sometimes toward error. However, the most

popular way to demonstrate reliance on a heuristic is to identify judgment contexts in which the heuristic will bias judgments in a predictable direction. Consider the availability heuristic (Tversky & Kahneman, 1974), that the ease of bringing to mind examples from a category gives insight into the size of that category (Schwarz et al., 1991). Early supportive evidence came from showing people mistakenly believe there are more words that start with 'r' than have 'r' as their third letter (Tversky & Kahneman, 1973). This finding was replicated with several other letters chosen precisely because they appear more often in the third than the first position: 'k,' 'l,' 'n,' and 'v' (Tversky & Kahneman, 1973). The appropriateness of these paradigms for testing the availability heuristic comes from an additional fact—that people have an easier time retrieving words by their initial than by their third letter. This is the Z described earlier, the third variable that distorts the signal value of the heuristic cue, meaning availability gives a predictably skewed indication of the dimension to be judged. By identifying a context in which Z distorts, the researchers probed for a systematic-bias that reliance on the heuristic would produce.

We call this the *systematic-bias* approach, the primary one that has been used when identifying the heuristics that guide everyday judgment. Researchers identify contexts in which heuristics should distort judgments in a particular direction and determine whether that occurs. One shortcoming of this approach is the mere observation of systematic bias means the evidence is equally consistent with any other heuristic that might make the same prediction. For example, Tversky and Kahneman's classic data is equally consistent with a heretofore unstudied *early-and-often heuristic*, that a good guide to how often letters appear in certain positions is whether that position is early (and not late) in a word. This (largely tongue-in-cheek) heuristic could also account for Tversky and Kahneman's findings. A second shortcoming of the systematic-bias approach is that reliance on a heuristic need not produce systematic biases. Imagine if Tversky and Kahneman had instead asked whether more words in the English language start with 'r' or 'k.' In this case, even if people lean on the availability heuristic ("Is it easier to think of words that start with 'r' or 'k'?"), this may not produce systematic errors in judgment. After all, in the absence of a distorting Z variable, the heuristic likely leads one to the correct answer.

Cue-Correlation Approach

Given the shortcomings of the systematic-bias approach, what is the solution to this problem? One empirical solution is to trace individual variation in judgments (i.e., the estimated proportion of words judged to start with 'r' vs. 'k') to individual variation in the heuristic cue (perhaps by using a think-aloud procedure to determine how many words starting with 'r' vs. 'k' participants spontaneously think of). We call this the *cue-correlation* approach because it tests for a correlation between the heuristic cue and the judgment the heuristic is hypothesized to inform. By tying variation in the heuristic attribute to variation in the judgment—either across participants (e.g., Critcher & Rosenzweig, 2014) or across items (e.g., Tversky & Kahneman, 1973)—one provides firmer support for reliance on a heuristic, even in cases when it keeps people from bias. The present paper ultimately uses both approaches.

Testing for the Commonness Fallacy

If people do lean on perceived commonness when forecasting others' choices, when should this cue produce systematic bias? Relying on perceived commonness should distort judgments when this cue is pit against a truly diagnostic cue, like what people would actually like to receive. When there is tension between an option's perceived commonness and liking, evidence of the commonness fallacy can be found using the systematic-bias approach. That is, people may overestimate the extent to which others will choose common but bland options (e.g., vanilla ice cream, black shoes, a mainstream blockbuster film) over rare but enticing options (e.g., tiramisu, blue shoes, an indie film).[1] But is this because forecasters lean on perceived commonness? For this question, the complementary cue-correlation approach is more useful. Individual differences in the perceived commonness of a specific option—even when controlling for the perceived liking for that option (that which should be axiomatically linked with whether people want to choose it)—should also predict forecasts of choice. Furthermore, even in the absence of a tension between perceived commonness and a truly diagnostic cue, evidence of the commonness fallacy can still be documented by way of the cue-correlation approach. We expect (and a study reported in the online supplemental materials confirms) that people lean on commonness in these cases as well, but in such cases commonness may be redundant with (instead of a distraction from) more valid cues.

Overview of the Present Studies

We present eight studies that, in combination, test whether and when people display the commonness fallacy. Guided by the systematic-bias approach, Studies 1a and 1b tested whether participants overestimate other participants' interest in choosing common items over rarer ones in both single (1a) and repeated (1b) choice contexts. Study 2 replicated these findings in an incentive-compatible context in which we went to great lengths to assure participants the experimenter would be blind to their actual choices. This study also introduces the cue-correlation approach. Studies 3–5 uncover when and why the commonness fallacy emerges. Study 3 tested whether participants leaned on perceived commonness in their choice forecasts because they truly believed it to be diagnostic, or whether it was merely an automatically accessible cue that would be set aside upon reflection. Study 4 examined whether the commonness fallacy applies to forecasts of choice in particular, or forecasts of preferences more generally. Study 5 explored a variety of possible antecedents of perceived commonness. Study 6 varied participants' exposure to a target's previously encountered choice contexts and choices to assess whether forecasters would still commit the commonness fallacy even when they had just observed choices that displayed no systematic preference for the more commonly offered and chosen item. Study 7 tested a downstream consequence of the commonness fallacy—that it may lead amateur sellers to systematically misprice goods.

---

[1] Although our goal is not to advance a formal definition or account of what makes a rare item *enticing*, we use this language merely to indicate that certain options are rare not because they are rarely available but because they simply are not particularly appealing (e.g., cockroach tacos for lunch, a beach vacation on the Arctic Ocean).

Across all studies, we wanted to maximize confidence that our tests offered sufficient statistical power to test our ideas, but we did not know a priori what the effect size was. Mindful of this common difficulty, Simmons, Nelson, and Simonsohn (2013) suggest that studies include at least 50 participants per cell or include additional justification. But in case the effects we were studying turned out to be small, we wanted to take several steps to increase our power to detect them. First, across all studies, we went well beyond Simmons et al.'s (2013) recommendations. Ignoring counterbalancing factors, we average 130 participants per condition (122 after exclusions). Second, in the majority of our studies, we did not have participants make merely one choice; instead, participants considered up to 11 choice pairs. Not only did this expand the generality of our tests, but it also permitted more statistical power for the same number of participants (Baayen, Davidson, & Bates, 2008). Third, we set *ex ante* stopping rules, but ones that would get us to at least these large sample sizes. That is, we considered practically how far we could go *beyond* a minimum threshold so that data collection could be completed using lab participants within one semester (Studies 1a–4 and 6) or would exhaust all available lab funds for online studies for that month (Studies 5 and 7).[2] All research was conducted in compliance with the rules of the relevant Institutional Review Board, and all study materials, data, and analysis scripts are publicly available on the OSF: https://osf.io/ar2e5/.

Study 1a

Study 1a tested for evidence of the commonness fallacy using the systematic-bias approach. That is, we identified 11 choice pairs that pit a common, bland item against a rarer, but relatively exciting one. Participants both indicated which option they would choose and estimated what percentage of other participants would make each choice. We predicted participants would overestimate what percentage of participants would select the relatively common (but blander) option.

Method

Participants and design. One hundred ten undergraduates at the University of California, Berkeley, completed this and other unrelated studies as part of an hour-long session. All received course credit. Seven participants failed at least one of two attention checks (see online supplemental materials). The remaining 103 participants are included in the results below.

Procedure and materials. We constructed 11 choice pairs. Each pair comprised two options—one relatively common, one relatively rare—from the same category. These materials are presented in the left half of Table 1. To make certain that the options did indeed differ on perceived commonness, we conducted a pretest on Amazon's Mechanical Turk ($N = 111$, after 9 attention check exclusions). For all 11 pairs, the common item was identified as more common for people to use or partake in than the rare item, all ts > 10.96, ps < .001. In the main study, participants completed two sets of measures; we counterbalanced which set participants answered first:

Choices. For each pair, participants were asked to consider having a choice between two options. For example, one item read, "If you had the following two options for juice to drink with breakfast tomorrow, which would you choose?" For this item, the rare option was *passion fruit juice* and the common option was *orange juice*: (In neither this nor any study did we label the options as "common" or "rare" to participants). The order of the two options was randomized, as was the order of the 11 choice pairs.

Forecasts. Participants estimated how their fellow participants responded to the choice measures. For exploratory purposes, we elicited such forecasts in one of two normatively equivalent formats (see Critcher & Dunning, 2013). Some participants were asked to answer, "The other 100 [participants] will see this question: 'If you had the following two options for juice to drink with breakfast tomorrow, which would you choose?' Predict how many of them will choose one option vs. the other." Others were instead asked to estimate how likely it was that a specific randomly selected participant would choose one option or the other (e.g., "Predict what percent chance Participant 71 will choose one option vs. the other.") Each question included two sliding scales, one for each item in the pair; the two had to add to 100%. The order of the two items within each pair was randomized, as was the order of the 11 choice pairs.

Results and Discussion

On average, participants chose the common (but bland) over the rare (but exciting) option 50.22% of the time. Did participants forecast this? Providing evidence of the commonness fallacy, participants mistakenly believed that others would gravitate toward the common item ($M = 59.12\%$, $SD = 8.57\%$)—a significant overestimation, $t(102) = 10.53$, $p < .001$, $d = 1.04$. Note that because the true answer was approximately 50%, any tendency for forecasters to regress toward 50% in their uncertainty only works against our hypotheses. The results by item are listed in Table 1. This systematic bias depended neither on the nature of the target (randomly selected individual vs. population) nor on the order with which participants indicated their own choice and forecasted others' choices, Fs < 1.[3]

Is this forecasting error large? First, looking to the effect size, we see that it is. Forecasters were off by more than a whole standard deviation. This is all the more remarkable given forecasters and choosers were the same people. But is the 8.90% systematic bias large? This is of course hard to evaluate in isolation. To offer a rough guide in interpreting our first study's results, we compared this to the size of the bias in the first study in Ross et

---

[2] We conducted post hoc tests to determine the adequacy of our sample sizes for detecting the most basic demonstration of the commonness fallacy: the extent to which participants overestimate others' interest in choosing common over rare items (the systematic-bias approach). The average effect size of this error (Studies 1a through 5) was $d = 0.69$. To have 99% power to detect such an effect, 41 participants are required. Every study recruited at least 41 participants per condition.

[3] We ran this same test concerning order effects in Studies 1b and 3. Consistent with results from Study 1a, the systematic bias documented in Studies 1b and 3 did not depend on the order with which participants indicated their own choice and forecasted others' choices, Fs < 1. These nonsignificant order effects address a concern that participants might modify their second judgments in light of their first (e.g., forecast that others will select more common items in light of one's own already-stated rare choices). Study 5 further addresses this concern by having participants make choices or forecasts with the accurate understanding they will make one but not both.

Table 1

Predicted and Actual Choice of the Common (vs. Rare) Options (Study 1a)

Category	Common option	Rare option	Predicted, common (%)	Actual, common (%)	Forecasting bias (predicted – actual)
Dinner beverage	Budweiser	Japanese imported beer	55.94	27.19	28.76***
Dinner	Pizza	Thai food	47.13	20.39	26.74***
Dessert 1	Assorted cookies	Crème brulee	38.36	19.42	18.94***
Breakfast beverage	Orange juice	Passion fruit juice	68.19	54.37	13.82***
Weeklong vacation	Hawai'i	The Galápagos	65.35	54.37	10.98***
Flowers	Daisies	Snapdragons	64.14	54.37	9.77***
Lunch	Sandwich	Curry	63.31	55.34	7.97***
Dessert 2	Vanilla ice cream	Tiramisu	50.31	46.6	3.71
Office wall paint color	White	Bright blue	70.23	69.91	0.33
Birthday celebration (with friends)	Dinner	Improv comedy	62.26	67.96	−5.70**
Concert	Classical piano	Classical harp	65.09	82.53	−17.44***
Average			59.12	50.22	8.90***

Note. Relying on the systematic-bias approach, one sees evidence consistent with the commonness fallacy when the predicted choice of the common item is greater than the actual choice. The significance level of each prediction bias is based on a one-sample t test on the predicted value compared against the actual value.
** $p < .01$. *** $p < .001$.

al.'s (1977) classic false consensus effect paper, as well as to the size of the bias produced by the false consensus effect in the present study. Reanalyzing Ross and colleagues' Study 1 results, we found that forecasters erred in the direction of their own choice by 8.60%. In the present study, forecasters tended to err in the direction of their own choice by 7.38%. Thus, the 8.90% systematic bias suggests that the commonness fallacy produces an error that is at least comparable in size to this classic social psychological phenomenon.

Study 1b

Study 1b was designed to assess the robustness of the commonness fallacy. Participants in Study 1a made a single choice for each choice pair. Perhaps the one-off nature of these decisions encouraged choosers to express more interest in unusual items than they ordinarily would. That is, perhaps forecasters do not typically overestimate others' choice of commonplace items, they simply underestimate the appeal of novel options in one-off decisions. This would suggest a different process than the reliance on commonness in forecasting choice.

To address this possibility, in Study 1b, participants considered three different choice pairs. Each had been used in Study 1a. In Study 1a's single-choice paradigm, these three pairs had shown evidence of a strong, an average, and a nonsignificant commonness fallacy. Here, instead of making each choice once, participants made each choice for every month in the upcoming year. Forecasters guessed what percentage of the time other participants chose one option or the other. If the commonness fallacy is not merely the result of estimating one-off choices, we should once again see that forecasters overestimate how often choosers will select the more commonplace options.

Method

Participants and design. Eighty undergraduates at Stanford University completed this and other unrelated studies as part of an hour-long session. All received $25 for the hour-long session. One participant failed the single attention check that asked what the

study was about (see online supplemental materials for details). The remaining 79 participants are included in the analyses below.

Procedure and materials. Like in Study 1a, participants both made choices themselves and forecasted the choices of others. But instead of making a single choice for each choice pair, participants made 12 for each. That is, participants were asked to consider making these choices once per month for the next year. We selected three choice pairs used in Study 1a that (a) could reasonably be offered on a monthly basis and (2) represented the spectrum of effect sizes observed in Study 1a. Participants completed the choices and forecasts in a counterbalanced order:

Choices. Participants were asked to select a drink to accompany a dinner (Japanese imported beer vs. Budweiser), flowers to display in one's home (snapdragons vs. daisies), and a dessert (tiramisu vs. vanilla ice cream). They began by making their selections for the then-current month (May), before considering the same choices for the next month (June), and so on until all 12 months were complete (see Table 2).

Forecasts. Given that in Study 1a it did not matter whether forecasts were elicited for all other participants or for a randomly chosen other participant, all participants in Study 1b were asked to forecast the choices of their fellow participants as a whole. Participants were told, "Every participant had to make the following choice 12 times, once for each month of the year." For each of the three choice pairs, participants estimated, "What percentage of the time do you think participants selected one item over the other?" For each pair, the two percentages had to add to 100%.

Results and Discussion

On average, participants chose the common (but bland) over the rare (but exciting) option 41.77% of the time. Illustrating the commonness fallacy, participants significantly overestimated this percentage ($M = 52.12\%$, $SD = 14.63\%$), $t(78) = 6.29$, $p < .001$, $d = 0.71$. This suggests the commonness fallacy is a mistake that characterizes more than estimates of others' one-off choices. Even for repeated choices, forecasters overestimate how much people will gravitate toward more common options.

Table 2

Predicted and Actual Choice of the Common (vs. Rare) Items (Study 1b)

Category	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Avg.	Forecast
Beverage	25.32	22.78	26.58	31.65	17.72	21.52	31.56	31.65	18.99	32.91	26.58	24.05	25.95[a]	46.82[b]
Flowers	46.84	45.57	50.63	54.43	59.49	59.49	58.23	50.63	49.37	50.63	51.90	44.30	51.79[a]	59.24[b]
Dessert	34.18	46.84	49.37	54.43	40.51	59.49	64.56	56.96	40.51	40.51	43.04	40.51	47.57[a]	50.30[a]
Average													41.77[a]	52.12[b]

Note. Mean percentages in the same row that do not share a superscript are significantly different at the $p < .05$ level.

Note that in eliciting participants' choices in Studies 1a and 1b, we were careful not to ask them which item they would buy or purchase, but instead which they would choose if offered two options. At the same time, we elected not to make explicit that (in these hypothetical choice contexts) the options would be free. We worried that calling artificial attention to this feature might create a demand, pushing people to feel like they should take advantage of this opportunity and choose what they thought would be the more expensive option. Despite these considerations, one might have concerns about the role perceived price still played in choices and forecasts. We provide additional information in the online supplemental materials about why any such influence likely worked in a hypothesis-inconsistent manner, meaning our tests are likely particularly conservative. Later studies address this issue more conclusively by: (a) using items with equivalent actual or perceived price, (b) leaning on choice contexts—both hypothetical and real—that offer natural reasons why people would not be paying for the item, and/or (c) measuring perceived price and showing a reliance on perceived commonness above and beyond the influence of perceived price.

Study 2

Study 2 replicated and extended on our earlier studies in three ways. First, we moved to an incentive-compatible choice context. Participants were given a choice between one of two candy bars—a relatively common one (an Original Milky Way) or a relatively rare one (a Midnight Dark Milky Way). Notably, these two items are equivalently priced, a property that the two hungry authors confirmed on a trip to their local drugstore. If the same forecasting error emerges in this context, we can more conclusively place fault with forecasters. By an alternative explanation, choosers misrepresent their preferences in hypothetical choice contexts by claiming more interest in rare items than they would actually display when their choices have real consequences. Given Study 1a and 1b's participants made their choices anonymously, we saw little reason for them to misstate their true choices. Still, an incentive-compatible choice context could more effectively allay this concern.

Second, we went to even greater lengths to make certain that participants knew their choices were private and anonymous. As described more fully below, we carefully constructed the choice context so that the experimenter would not know which candy bar participants selected. Crucially, we explained to participants how their choice would not be knowable by the experimenter. This would more conclusively remove any self-presentational pressures on choosers.

Third, we had each participant indicate the perceived commonness of both types of candy bars. This permitted us to supplement the systematic-bias approach with the more precise cue-correlation approach. That is, individual variation in the size of the forecasting error should be traceable to individual variation in the perceived commonness of the common versus the rare item. But also, it allowed us to probe the plausibility of an alternative explanation, that perceived commonness is influencing choices, not forecasts. By this alternative, perhaps it is not the case that commonness is used as a cue to what others choose, but as a cue to what one should avoid choosing. People may seek out variety in this way, but fail to anticipate it in others (e.g., Ratner, Kahn, & Kahneman, 1999). If so, the more that participants see the common item to be relatively prevalent, the less likely they should be to select it. Note that this would not undermine the conclusion that forecasters are committing an error, or even that perceived commonness played a role.

Before conducting our main study, we asked 92 Americans on Amazon Mechanical Turk (but excluded 10 who failed an attention check) to imagine they had the choice between two candy bars: Original Milky Way and Midnight Milky Way. Participants indicated which of the two candy bars they themselves would choose. Also, either just before or after making this choice, they estimated what percentage of others in this study would make one choice or the other. Although most participants indicated they would choose the Original Milky Way (56.10%), these same participants estimated that an even stronger majority ($M = 64.49\%$, $SD = 20.22\%$) would. Conceptually replicating Studies 1a-1b, this significant overestimate displays the commonness fallacy, $t(81) = 3.76$, $p = .001$, $d = 0.41$. Although the population from which this sample was drawn (American MTurkers) differs from the population from which our main study was drawn (undergraduates at the University of California, Berkeley), this Pilot Study can also serve as a rough benchmark for interpreting the results of our main, incentive-compatible study.

Method

Participants and design. One hundred ninety-nine undergraduates at the University of California, Berkeley, completed this and other unrelated studies as part of an hour-long session for which they received course credit. Nine participants failed the attention check. Data from the 190 remaining participants are included in analyses reported below. Two participants declined to take the chosen item with them. One such participant also failed the attention check. All analyses reported below remain statistically significant even if the other such participant is excluded (see online supplemental materials for details).

Procedure and materials. Participants made judgments about a single choice pair. They indicated which of the two options they would choose (choice), what percentage of other participants would choose one option or another (forecast), and how common it was to eat each when having a candy bar (commonness). The order of choice and forecast was counterbalanced, but both came before judgments of commonness:

Choice. When participants arrived in a private room, two items lay on a table before them. Each was wrapped in a different color paper. Participants were told (truthfully) that the experimenter did not know what was enclosed in each wrapping. But participants were told what was inside of each packaging: an Original Milky Way bar (common) or a Midnight Milky Way bar (rare). Participants were shown a picture of each option (outside of its disguising paper wrapping).

Before participants indicated which bar they wanted to take home, we included additional instructions to make sure participants knew their choice would be private: "No one will ever know which of the two you take—including the research assistants. Other undergraduates who are not affiliated with this study were responsible for wrapping each item in a specific color. . . . Additionally, your survey response of which item you choose will be kept completely anonymous." Once participants indicated which candy bar they would like to take with them, the computer informed them which color package to leave with. Before the next participant arrived, experimenters replenished whichever wrapped bar (identified by its colored wrapping) had been taken without their having to know which specific bar that was.

Forecast. Participants in Study 2 estimated what percentage of other participants would select one option or another. So that the forecasting context perfectly matched the choice context, forecasters saw exactly what choosers saw. That is, when making their forecasts, forecasters saw the exact choice prompt and the pictures of the two candy bars.

Commonness. Participants were told, "We want to get a sense for how common or uncommon you think it is for people to do each of the following. (Instead of thinking about all people in the world, think about the type of people who are participating in this study.)" Then, they were asked, "How common are each of the following options among people like those who are participating in this study?" One item read, "When having a candy bar, for it to be an Original Milky Way bar." The other, "When having a candy bar, for it to be a Midnight Milky Way bar." Ratings were made on 10-point scales anchored at 1 (*very uncommon*) and 10 (*very common*). The order of the two items was counterbalanced.

Results and Discussion

Participants chose to take home with them the more common Original Milky Way bar 55.26% of the time. But participants estimated that 63.34% (SD = 17.46%) of their peers would choose that more common option. Demonstrating the commonness fallacy, forecasters overestimated how many others would select the more common option, $t(189) = 6.38$, $p < .001$, $d = 0.46$. It is notable that this 8.08 percentage point overestimate is quite similar to the 8.39 percentage point overestimation observed in the Pilot Study (see Table 3). Although the two

studies were run using different populations, the persistence and remarkable similarity in the size of the bias does suggest there is little reason to think that self-presentational concerns or the hypothetical nature of the choice was responsible for the forecasting bias.

Our findings thus far have used the systematic-bias approach to testing for the commonness fallacy. We aimed to complement this by using the cue-correlation approach. Toward this end, we constructed a new variable *relative commonness*, the perceived commonness of the common option minus the perceived commonness of the rare option. Attesting to the validity of our common and rare operationalizations, this composite was positive ($M = 2.73$), but there was variance in these perceptions ($SD = 2.61$).

We regressed choice forecasts (of the common item) on relative commonness and participants' own choice ($+1$ = common item, -1 = rare item). Merely demonstrating the robust phenomenon of projection (Ross et al., 1977), those who selected the common option themselves estimated that more others would choose it, $\beta = .23$, $t(187) = 3.61$, $p < .001$. But providing our most direct evidence yet of perceived commonness as a heuristic cue, the perceived relative commonness of the common (vs. rare) item predicted inflated forecasts of choice of the common item, $\beta = .39$, $t(187) = 6.01$ $p < .001$. In an effort to replicate these cue-correlation findings, we returned to our Pilot Study data. And indeed, we found a similar effect of relative commonness on forecasts of choice (controlling for projection), $\beta = .32$, $t(79) = 2.94$, $p = .004$. In other words, in both the Pilot Study and the main study, forecasters—on the whole—did not merely commit a systematic forecasting error (the systematic-bias approach), but individual variation in their forecasts could be tied to individual variation in the perceived commonness of the options (the cue-correlation approach).

Instead of merely using participants' own choices as a covariate (controlling for projection), we can examine whether they are a moderator. That is, might participants who chose the common option themselves be unique in how much they lean on commonness when forecasting others' choice? We added the Self Choice × Relative Commonness interaction term to the model. Speaking against the possibility that reliance on commonness depends on participants' own selection of the common item, this interaction did not reach significance, $\beta = -.14$, $t(186) = -1.16$, $p = .249$.[4]

These results show the importance of perceived commonness to forecasts. But are the forecasting errors we observed also a function of perceived commonness guiding *choice*? If so, this would stem from choosers' perceptions that an item is rare increasing their interest in selecting it. Reducing the plausibility

---

[4] The Self Choice × Relative Commonness interaction actually did emerge in the Study 2 Pilot Study, $\beta = -.43$, $t(78) = -2.25$, $p = .028$. Given that it did not emerge in the main study, we proceeded to test for analogous effects in Studies 3 and 4 to determine which finding was anomalous. The Self Choice × Relative Commonness interaction was significant neither in Study 3, $B = -.16$, $t(1704.14) = -1.16$, $p = .248$, nor Study 4, $B = -.04$, $t(270.80) = -1.47$, $p = .142$. In summary, reliance on perceived commonness is typically independent of participants' own choices.

Table 3
Predicted and Actual Choice of Common (vs. Rare) Candy Bars, and Effect of Relative Commonness on Predicted Choice (Study 2)

	Systematic-bias approach			Cue-correlation approach
Study	Predicted choice of common candy bar, % (SD)	Actual choice of common candy bar, %	Forecasting bias, predicted − actual	Relative commonness → Predicted choice, B (SE)
Pilot	64.49 (20.22)	56.10	8.39***	2.15 (.73)**
Lab	63.34 (17.46)	55.26	8.08***	2.61 (.43)***

Note. Relying on the systematic-bias approach, one sees evidence consistent with the commonness fallacy when the forecasting bias is significantly positive. The significance level of each prediction bias is based on a test against zero, which would reflect no bias. Relying on the cue-correlation approach, one sees evidence consistent with the commonness fallacy when the relative commonness of the common item positively predicts choice of the common item (controlling for projection).
** $p < .01$.   *** $p < .001$.

of this account, those who saw the Midnight Milky Way as especially rare (compared with the Original Milky Way) were actually *less* likely to choose the rare item themselves, $\beta = -.17$, $t(188) = -2.39$, $p = .018$.[5]

Study 3

We have reasoned that in estimating what others will choose, it is natural to call to mind what people have chosen. As a result, forecasts of the choice of X over Y go beyond a consideration of whether others like X more than Y; they are also swayed by the relative commonness of X over Y. As a first test of this idea, that the commonness fallacy applies to forecasts of choice in particular (as opposed to preferences more broadly), we included a new measure. In Study 3, we measured how much participants estimated that others would *like to receive* each option. We predicted that forecasters' exaggerated estimates that others would select the commonplace option would be tied to the commonness of the options, even with perceived liking for the options statistically controlled. This is the first of two ways we test whether the reliance on perceived commonness when forecasting choice does not also affect people's understanding of others' preferences more generally, but interferes with estimates of others' *choices* in particular.

Study 3 was designed to do more than just offer this more conservative cue-correlation test of the commonness fallacy. We also wanted to understand to what extent forecasters leaned on perceived commonness over perceived liking merely because the former was a spontaneously accessible heuristic cue. Alternatively, participants might lean on perceived commonness more purposefully—choosing to give weight to perceived commonness over perceived liking as opposed to automatically relying upon it. Such intentional reliance on normatively suspect information has been identified in other judgment work (e.g., the outcome bias; Baron & Hershey, 1988).

To probe this, for some participants, we drew their attention to both perceived liking and perceived commonness *before* they made their choice forecasts. If people lean on perceived commonness because they explicitly believe that commonness is a valid cue to choice, then this manipulation should have no effect. If anything, it may increase reliance on perceived commonness, given it focuses everyone on the cue and thus makes it accessible for potential use. But if instead the commonness fallacy stems from an attribute substitution driven by commonness's accessibility and not its perceived diagnosticity in predicting choice, then calling special attention to these two cues should change what

participants lean upon when making their forecasts. More specifically, this salience manipulation should encourage people to then weight the (now accessible) information that they feel is most diagnostic for forecasting (Rosenzweig & Critcher, 2014). We expected this would be perceived liking (after all, it is almost tautological that people choose X over Y to the extent they would like to receive X over Y) at the expense of perceived commonness (see Critcher & Rosenzweig, 2014, for similar logic).

Method

Participants and design. Two hundred twenty-five undergraduates at the University of California, Berkeley, completed this and other unrelated studies as part of an hour-long session for which they received course credit. Fifteen participants failed at least one of two attention checks—one that asked for the sum of 2 and 2, one that asked what the study was about (see online supplemental materials for details). Data from the remaining 210 participants are included in all analyses reported below. Participants were randomly assigned to a *salience* or a (nonsalience) *control* condition.

Procedure and materials. As in Study 1a, participants made judgments about 11 unique choice pairs. For each, participants indicated which option they would choose (*choices*), what percentage of other participants would choose one option or another (*forecasts*), and how common each item was in people's lives (*commonness*). As a new addition, participants indicated how much other participants in the study would like to receive each option (*liking*). Participants in the salience condition completed the perceived commonness and perceived liking measures (in a counterbalanced order) *before* completing the choice and forecasting measures (also in a counterbalanced order). Those in the (nonsa-

---

[5] We ran this same test in the Study 2 Pilot Study and Studies 3 and 4. In the Study 2 Pilot Study, those who saw the rare item as especially rare (compared with the Original Milky Way) were less likely to choose the rare item themselves, $B = -.06$, $t(80) = 3.93$, $p < .001$. Proceeding to Studies 3 and 4, random-slope, random-intercept models in which we use only relative commonness to predict participants' own choices revealed a significant effect of relative commonness in Study 3, $B = -.02$, $SE = .01$, $t(280.76) = -2.75$, $p = .006$ and Study 4, $B = -.005$, $SE = .001$, $t(214.81) = -5.68$, $p < .001$. That is, as was evidenced in Study 2, across these three additional studies, those who saw the rare items as especially rare (compared with the common items) were actually *less* likely to choose the rare item themselves. This suggests (using the cue-correlation logic) that, to the extent commonness perceptions affect choosers, it is unlikely to be contributing to the systematic bias criterion for the commonness fallacy.

lience) control condition completed the commonness and liking measures *after* the choice and forecasting measures. The order of the specific choices and forecasts were randomized. Details on these measures are offered below:

Choices. These measures were equivalent to those used in Study 1a, except we substituted out two choice pairs and added in two new choice pairs. In place of the questions about wall paint color and musical concerts, we asked participants to choose either an apple or a guava for their next snack and either a traditional landscapes or abstract/modern exhibit for their next art show to attend.

Forecasts. Whereas participants in Study 2 estimated what percentage of other participants made one choice or another, those in Study 3 estimated what percent chance a specific other participant would choose one option or the other: "Predict what percentage chance Participant #X will choose one option vs. the other." For each choice pair for each participant, X was randomly sampled from the integers between 1 and 100, inclusive.

Commonness. Participants were asked to rate the commonness of having each item when partaking of or consuming an option in the relevant category. They were also told, "Instead of thinking about all people in the world, think about the type of people who are participating in this study." For example, participants indicated how common it was to have orange juice when having juice with breakfast. Ratings were made on a 1-to-10 scale, from 1 (*relatively uncommon*) to 10 (*relatively common*). The order of the 22 items was randomized.

Liking. Presumably others' choices will largely (if not entirely) be a function of whether they would like or be pleased to receive one option versus another. We aimed to measure perceptions of how much others would like to receive each item. As with the commonness item, we asked them to think about the type of people who were participating in this study. We asked participants, "If the average person learned that the following choices had been made for them, do you expect that they would feel very pleased or not at all pleased about this choice?" For example, one item asked how pleased another would be to learn "that their next juice with breakfast will be orange juice" on a 10-point scale, from 1 (*not at all pleased*) to 10 (*very pleased*). The order of the 22 items was randomized.

Results and Discussion

First, we tested whether we replicated the systematic bias found in previous studies. Overall, participants chose the common (but bland) option instead of the rare (but exciting) option 44.89% of the time. Consistent with the commonness fallacy, participants overestimated how often others would choose the common option ($M = 54.33\%$, $SD = 9.71\%$), $t(209) = 14.08$, $p < .001$, $d = 0.97$. This confirms our hypotheses using the systematic-bias approach.

Of course, this directional bias is merely consistent with, but does not offer direct evidence of, perceived commonness as the heuristic cue. To provide a more direct test, we followed the cue-correlation approach by connecting people's perceptions of commonness to their forecasts of choice. But crucially, we wanted to provide a particularly conservative test. After all, the commonness of an item may not directly inform forecasts of choice in particular, but may instead inform preferences more generally. Chocolate ice cream is more common than jalapeño ice cream in

part because people do like—and thus would likely choose—chocolate-flavored desserts. Our hypothesis is bolder—that commonness (an assessment of what has been chosen) exerts its own incremental effect on choice forecasts. That is, perceived commonness should predict forecasts of choice above and beyond such perceptions of liking.

To test these ideas, we began by defining two variables: *relative commonness* and *relative liking*. Each reflected a participant's rating of the common item minus the participant's rating of the rare item, for a particular choice pair. For example, if a participant thought liking for assorted cookies would be an 8, but for crème brulee would be a 9, the relative liking score for this participant for this choice pair would be −1. Relative commonness and relative liking were Level-1 variables that were nested within participant in a random-slope, random-intercept model predicting choice forecasts. This permitted the effects of commonness and liking to vary for each participant (random-slope), and allowed for differences between participants in how often they thought common items would be chosen (random-intercept). We also included a random effect of choice pair to account for variation between the 11 pairs in how often it was assumed others would choose the common item. Finally, given the robust phenomenon of projection, we included participants' own choice as an additional Level-1 predictor.

Unsurprisingly, there was a main effect of relative liking on choice forecasts, $B = 2.87$, $SE = 0.21$, $t(164.38) = 13.77$, $p < .001$. In other words, the more participants thought that others would like to receive the common as opposed to the rare item, the more they forecast others would make choices reflecting those preferences. But important for our purposes, there was a large residual main effect of relative commonness, $B = 0.97$, $t(266.52) = 5.74$, $p < .001$. That is, commonness clearly predicted forecasts of choice above and beyond the influence of perceived liking (a variable that includes any influence of commonness on assumed preferences more generally). This confirms our hypotheses using a more conservative version of the cue-correlation approach than that used in Study 2.

Do people intentionally lean on commonness because they embrace it as a valid cue for forecasting choice, or is it the mere accessibility of perceived commonness that prompts mindless reliance on this attribute substitution? To disentangle these possibilities, we added several terms to our model. To begin, we added *salience* (+1: commonness and liking made salient before forecasts; −1: forecasts made before commonness and liking assessed). But most centrally, we added the Salience × Relative Commonness and Salience × Relative Liking interaction terms. These interaction terms would permit us to test whether calling special attention to commonness and liking shifted how much participants relied upon them when forecasting others' choice.

We observed a negative Salience × Relative Commonness interaction, $B = -0.45$, $SE = 0.14$, $t(184.11) = -3.13$, $p = .002$. Consistent with our attribute substitution account, once special attention was called to both commonness and liking, participants reduced their reliance on commonness as a cue. But also, calling special attention to commonness and liking increased reliance on perceived liking, as reflected by a positive Salience × Relative Liking interaction, $B = 0.52$, $SE = 0.20$, $t(146.74) = 2.66$, $p = .009$. These shifted weights are depicted in Figure 1. In short, the commonness fallacy seems to emerge because commonness is a
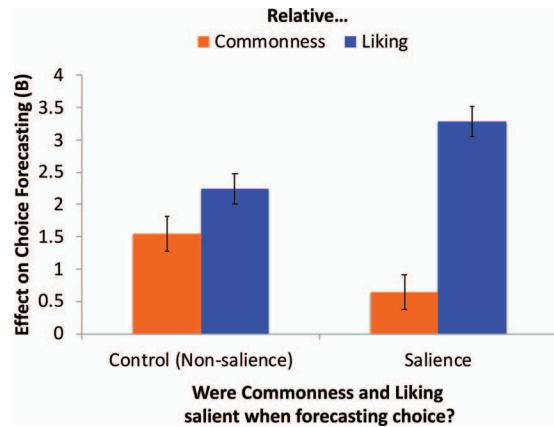
Figure 1. The independent influence of relative perceived commonness and relative perceived liking on forecasts of others' choice. The depicted values are betas (with ±1 *SE* error bars) from the model described in the text. Positive betas reflect that the more that a common item was seen to be more prevalent than or liked more than a rare item, the more it was assumed others would choose the common item. This provides support using the especially conservative cue-correlation approach that people are leaning on perceived commonness (Study 3). See the online article for the color version of this figure.

spontaneously accessible cue; reliance on commonness drops (and perceived liking increases) when both cues are made salient in forecasters' minds (see Rosenzweig & Critcher, 2014). In the online supplemental materials, we provide additional analyses that demonstrate that perceived relative liking is indeed a more diagnostic cue than perceived relative commonness; we also explore how our salience manipulation affects forecasting precision and bias.

Our goal with the salience manipulation was merely to call attention to potential inputs (i.e., commonness and liking) that might not be spontaneously accessible. This allowed us to determine how much participants were interested in leaning on the information once the spotlight had been placed on it. Might the manipulation implicitly have suggested that these were two valid cues to forecasting choices? Note that this is why it was so important that we called attention to both liking *and* commonness in the salience condition. That is, after focusing participants on both cues, they did not enhance their reliance on both, but instead shifted their use of the heuristic cue (commonness) to the more normatively defensible one (liking).

We interpret the strong residual influence of perceived commonness—above and beyond that of perceived liking—as evidence of people's direct reliance on commonness when forecasting choice. But instead of demonstrating reliance on a commonness heuristic, might the significant residual effect merely reflect measurement error in the perceived liking composite? That is, perhaps participants merely think others choose what they would like to receive (a quite reasonable approach), and a correlation between perceived commonness and perceived liking combined with measurement error in the perceived liking composite allowed perceived commonness to exert a significant residual effect. We conducted additional analyses that cast doubt on this concern.

More specifically, we first identified a variable that we expected would relate to perceived liking but not perceived commonness: participants' own choice. Projection should lead people to think that others' preferences are predicted by the self's own preferences. But projection does not suggest that the commonness of options in the world (especially controlling for what we perceive others' preferences to be) is a function of what the self would choose. As reported in more detail in the online supplemental materials, perceived liking did predict the self's own choices, but crucially perceived commonness did not once perceived liking was controlled. In other words, it is not that the perceived liking measure is so noisy that when the variable is included as a covariate, unique predictive power from perceived commonness is inevitable. This casts doubt on a measurement error artifactual explanation and reinforces our interpretation that forecasters lean directly on perceived commonness—above and beyond perceived liking—when forecasting others' choices.

Study 4

Study 4 addresses a question that typically is neglected in the heuristics and biases research tradition: What about the judgment task encourages reliance on the heuristic cue? By our reasoning, considering what others will choose is so deceptively similar to considering what has been chosen (commonness) that the heuristic attribute becomes spontaneously accessible and guides choice. But might commonness instead simply be used to understand others' preferences more generally? This possibility is called into question by Study 3's finding that perceived commonness predicts forecasts of choice over and above perceived liking, but Study 4 was designed to address this question more systematically and conclusively.

We explored the boundaries of the commonness fallacy by asking forecasters to make predictions that took one of three forms. In one condition, participants forecasted others' choice, just as participants in our previous studies have. Participants in the other two conditions were told that a computer would randomly select what another participant would receive. Unlike in previous studies, in these two conditions forecasters estimated how likely it was that another participant would be more *pleased to receive* one item or the other. The only difference between these two preference forecast conditions was that we varied whether the targets were said to know (*preference—known options* condition) or not know (*preference—unknown options* condition) what two options the computer was deciding between.

By using these three conditions, we were able to disentangle three possibilities. First, by our logic, we should find more evidence of the commonness fallacy—using both the systematic-bias and cue-correlation approaches—in the choice condition, compared with either of the preference forecast conditions. But as a second possibility, perhaps the commonness fallacy is *not* triggered by forecasts of choice in particular but is used to inform one's understanding of others' preferences more generally. If so, we should find fairly consistent reliance on perceived commonness across all three conditions. A third possibility predicts a different pattern of results. When one makes a choice between two options, one knows what options are under consideration. If the commonness fallacy is not triggered by forecasting *choice* in particular, but rather by forecasting others' *relative* preference (a circumstance in which the target is comparing one specific option to another known option), we should see more

reliance on perceived commonness in the choice and preference—known options conditions compared with the preference—unknown options condition. Although this third possibility was not as a priori compelling, we included the two preference conditions to allow us to cleanly identify forecasts of choice, not forecasts of people's satisfaction with one option instead of a known alternative, as the boundary condition.

Finally, Study 4 featured another new addition in that it permitted a detailed look at the influence of perceived price on choice forecasts by measuring it at the individual level. This allows us to be certain that participants were not leaning on the perceived cheapness of items when forecasting that others were likely to choose highly common (and often cheaper) items.

Method

Participants and design. Two hundred eleven undergraduates at the University of California, Berkeley, completed this and other unrelated studies as part of an hour-long session for which they received course credit. Participants were randomly assigned to one of three forecasting conditions: *choice, preference—known options, preference—unknown options.* Nine participants failed the attention check that asked what the study was about (see online supplemental materials for details). Data from the 202 remaining participants are reported below.

Materials and procedure. All participants began by making forecasts in one of three forms (see below). After that, participants completed measures that were nearly equivalent to those used in Study 3: (self-)choices, commonness, and liking. The only difference was that here, the commonness and liking ratings were made on a 0-to-100 sliding scale. These three sets of measures were all presented in a random order. Lastly, participants completed the perceived price measure.

Choice forecasts. These measures were nearly equivalent to the forecasts used in Study 3. Participants indicated the percentage chance that a randomly selected participant would choose one option versus the other. The instructions fleshed out one example: "Participant A will see this question: 'If you had the choice of the

following two options for lunch tomorrow, which would you choose?' Predict what percentage chance Participant A will choose one option vs. the other."

Preference—known options forecasts. Participants in the preference—known options condition were shown the same 11 choice pairs as those in the choice forecasts condition. Furthermore, they were told that a randomly chosen participant (Participant A) would see these same pairs. But it was said a computer would randomly select the option Participant A would receive. Participants' task was thus not to forecast choice, but to forecast the percentage chance that Participant A would be more pleased to receive one item or the other.

Preference—unknown options forecasts. Participants in the preference—unknown options condition were given a similar forecasting task as those in the preference—known condition, except they were explicitly told that Participant A would *not* see what the computer was selecting between. That is, a computer would randomly select the option Participant A would receive (just like in the preference—known options condition). But Participant A would not be aware of the alternative, unchosen option (unlike the preference—known options condition). This meant Participant A could not compare the choices directly, and thus could not even hope for one of the two options.

Price. Participants were asked to estimate the price of each item. They were told, "For the final task, we would like you to estimate the price of each of the following 22 items (in dollars and cents)." An example item read, "Imagine one wanted orange juice to drink with breakfast. On average, how much would someone (who lives where you do) spend on a glass of orange juice?"

Results and Discussion

When expressing their own choices, participants picked the common item only 44.46% of the time. The extent to which participants' forecasts overshot this value (as predicted by the commonness fallacy) depended on the specific forecast participants were asked to make, $F(2, 199) = 5.84$, $p = .003$, $\eta_p^2 = .06$ (see Table 4). When estimating what another participant would

Table 4

Predictions (by Forecasting Condition) and Actual Choice of the Common Options (Study 4)

Category	Common option	Rare option	Predicted choice, common (%)	Predicted preference (known), common (%)	Predicted preference (unknown), common (%)	Actual choice, common (%)
Lunch	Sandwich	Curry	59.48[a]	54.18[a]	56.31[a]	47.52[b]
Dinner beverage	Budweiser	Japanese imported beer	49.87[a]	41.54[b]	44.69[ab]	24.26[c]
Weeklong vacation	Hawai'i	The Galápagos	59.25[a]	58.14[a]	56.99[a]	55.94[a]
Dinner	Pizza	Thai food	49.25[a]	46.48[a]	50.01[a]	27.72[b]
Fruit	Apple	Guava	63.31[a]	57.96[ab]	48.99[c]	56.93[b]
Birthday celebration (with friends)	Dinner	Improv comedy	63.56[a]	53.24[b]	51.87[b]	60.89[a]
Dessert 1	Assorted cookies	Crème brulee	35.08[a]	39.16[a]	35.39[a]	21.78[b]
Flowers	Daisies	Snapdragons	57.37[ab]	53.21[b]	53.37[b]	58.91[a]
Dessert 2	Vanilla ice cream	Tiramisu	42.85[a]	41.04[a]	40.60[a]	34.16[b]
Breakfast beverage	Orange juice	Passion fruit juice	66.23[a]	54.20[b]	50.29[bc]	50.50[c]
Art exhibit	Traditional paintings	Modern/Abstract	48.92[ab]	47.14[b]	48.61[ab]	50.50[a]
Average			54.10[a]	49.66[b]	48.83[b]	44.46[c]

Note. By the systematic-bias approach, evidence consistent with the commonness fallacy is seen when the predictions are greater than the actual choice. Although forecasts in the two preference conditions did depart from actual choice, a more specific test (using the cue-correlation approach) showed that participants in those conditions did not lean on perceived commonness. Means in the same row that do not share a superscript differ at the $p < .05$ level.

choose, participants thought there was a 54.10% (SD = 10.80%) chance the common item would be chosen. This estimate was significantly higher than reality, $t(51)$ = 6.44, p < .001, d = 0.89, offering support for the commonness fallacy through the systematic-bias approach.

But by making estimates not of which item others would choose, but which they would be more pleased to receive, participants seemed to have more insight into how likely each item was to be selected. That is, compared with those in the choice forecast condition, the preference forecast participants saw less appeal in the common option regardless of whether they were asked to forecast which of two known items another would be more pleased to receive (M = 49.66%, SD = 7.84%), $t(199)$ = 2.80, p = .006, d = 0.47, or which of two unknown items another would be more pleased to receive (M = 48.83%, SD = 8.50%), $t(199)$ = 3.23, p = .001, d = 0.54. Forecasts were similar regardless of whether the recipient knew the other options or not, t < 1. Although the systematic bias was significantly reduced in both conditions, it was still significantly present: $t(79)$ = 5.93, p < .001, d = 0.66 (preference—known options); $t(69)$ = 4.30, p < .001, d = 0.51 (preference—unknown options).

Of course, merely showing that participants in the two preference conditions showed less bias in their forecasts does not directly demonstrate that they leaned less on perceived commonness. Only the cue-correlation approach can provide that convergent support. We followed the conservative procedure used in Study 3 to determine whether participants were indeed leaning on perceived commonness, over and above perceived liking, when forecasting choice. We defined three Level-1 variables—*relative commonness*, *relative liking*, and (*self-*)*choice*—and nested them within participant. The first two variables reflected the perceived commonness or liking for the common item minus the rare item. The last variable controls for the theoretically irrelevant phenomenon of projection. Of greatest interest, we included a categorical variable for the forecast condition (choice, preference—known options, preference—unknown options), as well as the Forecast × Relative Commonness and Forecast × Relative Liking interaction terms. Finally, as in our previous studies, we included a random effect of *choice pair,* which accounted for differences in forecasts among the 11 pairs related to the general popularity (regardless of the forecasting type) of the common versus rare item.

Unsurprisingly, we observed a strong main effect of relative liking, $F(1, 147.50)$ = 193.92, p < .001. That is, participants estimated that others were more likely to choose or be pleased to receive the common (vs. rare) item to the extent it was assumed others liked the common (vs. rare) item. Furthermore, this reliance on perceived liking did not vary by condition, F < 1. In contrast, we saw a more modest overall influence of perceived commonness, $F(1, 220.01)$ = 13.52, p < .001, in part because the influence of commonness varied by forecasting condition, $F(2, 160.76)$ = 8.01, p < .001.

When we tested for the influence of perceived commonness by condition, we observed a pattern of results that nicely complemented the findings from the systematic-bias approach. In a replication of our past results, those who forecast others' choice leaned on perceived commonness, B = 0.12, $t(47.18)$ = 4.19, p < .001. In contrast, those who forecast what would make others more pleased did not. This was true regardless of whether forecasts were about those who would know about the counterfactual option, B =

.00, t < 1, or those who would not know about that counterfactual option, B = .04, t < 1. In summary, even though it might seem tautological that others will select the options that they would be most pleased to receive, we found that reframing the forecasting task in this way encouraged more accurate forecasts of what people choose. This is because people leaned on (and were led astray by) perceived commonness only when predicting others' *choice*, not what they would prefer to receive. This supports our account that the commonness fallacy arises in part from people inappropriately blurring what people in general have chosen (commonness) with what an individual will choose (forecasts of choice), not because commonness is always leaned on to understand others' preferences.[6]

Is it odd that we found that those forecasting what others would be pleased to receive showed no evidence of leaning on perceived commonness (by the cue-correlation approach) but did still show some systematic bias in the direction of the commonness fallacy? One possibility—though an admittedly speculative one—is that perceived commonness may partially inform perceptions of liking across all three conditions. Consistent with this idea, there is a moderate zero-order correlation between perceived commonness and perceived liking (r = .22). But as the cue-correlation analyses showed, only those predicting others' choices lean directly on perceived commonness as well. This would explain why there is some systematic bias consistent with the commonness fallacy when predicting what others would be pleased to receive, a bias that then becomes exaggerated when predicting what others would choose.

We turn to the question of whether evidence of the commonness fallacy reflects reliance on commonness, or instead perceived price. Because of positive skew in price estimates, we first log-transformed them. We then created the variable *relative price;* akin to our relative commonness and relative liking composites, this reflected the (log-transformed) estimated price of the common item minus the (log-transformed) estimated price of the rare item. We added relative price as well as the Forecast × Relative Price interaction terms. Crucially, the focal Forecast × Relative Commonness interaction remained, $F(2, 1940.85)$ = 9.75, p < .001. In other words, the predictive power of perceived commonness did not actually reflect the unmeasured variable perceived price. We did observe a significant influence of perceived price, $F(1, 1631.60)$ = 12.57, p < .001, but the influence was *positive*, and did not vary by forecasting condition, F < 1. People assumed that others would choose items to the extent they were relatively more expensive, which is the opposite of the concern that people might overestimate others' interest in common items due to a belief that others gravitate toward inexpensive items.

Study 5

In documenting the commonness fallacy, we have used the systematic-bias and cue-correlation approaches to provide conver-

---

[6] We did find projection differed by forecasting condition, $F(2, 172.84)$ = 6.43, p = .001. But this actually reflected that people projected more when forecasting others' choice (B = 7.14) than when forecasting what others would like to receive (Bs = 4.45 and 3.16). Given projection assists with accuracy, this differential tendency actually constrained our predicted effects.

gent evidence that the perceived commonness of options guides forecasts of how often they will be chosen. But in so doing, we have relied on a mix of experimenter intuition and participants' own ratings to identify which options are more or less common. But where do such perceptions originate? No doubt, such intuitions find their origins in many sources.

For Study 5, we identified five possibilities: the frequency with which the self has previously chosen each option (self-history), how knowledgeable the self is about each option (knowledge), the ease of recalling specific instances of people selecting the option (recall ease), how much shelf or website space is devoted to each option (store space), and whether the option is a good example of the category (typicality). First, we ask which one or (likely) more of these offer unique insights into participants' perceptions of commonness. Second, we ask which of these antecedents also predict participants' forecasts of others' choice. We conclude by considering whether these new variables are plausible antecedents of the commonness fallacy—cues that guide forecasts of choice because they may shape perceptions of options' commonness.

Whereas the above tests reflect an extension of our cue-correlation approach, we made an additional methodological change that affects the systematic-bias approach. In our earlier studies, the same participants both indicated their own choices and forecasted the choices of others. In the present study, each participant either indicated their own choices or forecasted others' choices, but did not do both. Although previous studies counterbalanced the sequencing of these choice and forecast measures and did not observe order effects (see Footnote 3), having participants know they would complete only one of these two tasks obviates all concerns about how participants' stated choices might affect or be affected by their choice forecasts.

Method

Participants and design. The respondents were 349 Americans recruited from Amazon Mechanical Turk (MTurk). They participated in exchange for nominal payment. Participants were randomly assigned to one of two conditions: *choice* or *forecast*. Twenty-two participants failed the attention check, a single item that asked what the study was about (see online supplemental materials for details). The hypotheses, measures, sample size, and analysis plan were preregistered: https://aspredicted.org/7uw84 .pdf.

Materials and procedure. Participants considered 11 choice pairs, the same ones used in Study 3. Those in the *choice* condition first indicated which option they themselves would select. Those in the *forecast* condition instead estimated the percentage chance that a randomly selected other participant would choose one item or another. Next, all participants completed the commonness measure. These three measures were virtually identical to those used in Study 3, the only difference being that here, commonness was rated on a 1 (*very uncommon*) to 10 (*very common*) scale.

At that point, all participants completed in a random order a battery of measures that probed possible antecedents of perceived commonness: self-history, knowledge, recall ease, store space, and typicality.

Self-history. Participants were asked, "Please indicate whether the thing that follows is relatively common or relatively uncommon for you to do." For example, participants indicated how common it was for them, "when having juice for breakfast, for it to be orange juice." Ratings were made on a 1 to 10 scale, from 1 (*very uncommon*) to 10 (*very common*). If participants lean on their own history with the options to infer commonness, this would reflect a sort of projection (Ross et al., 1977).

Knowledge. Participants were asked the extent to which they feel knowledgeable about each option. For example, participants indicated how knowledgeable they felt about, "the vacation destination: The Galápagos." Ratings were made on a 1-to-10 scale, from 1 (*not at all knowledgeable*) to 10 (*very knowledgeable*). Participants might infer that because they know more about an option, it must be more common.

Recall ease. Participants were asked how easy it was to recall specific instances of people partaking of or consuming an option in the relevant category (Tversky & Kahneman, 1973). For example, participants indicated how easy it was to recall instances of people "eating curry for lunch." Ratings were made on a 1-to-10 scale, from 1 (*extremely difficult*) to 10 (*extremely easy*). Leaning on recall ease to infer commonness would reflect reliance on the availability heuristic (Schwarz et al., 1991; Tversky & Kahneman, 1973).

Store space. Participants were asked, "In a physical store or on an online store, how much space (e.g., shelf space, website space) is typically devoted to each option." One example read, "The dessert option: vanilla ice cream." Ratings were made on a 10-point scale, from 1 (*almost none*) to 10 (*a lot*). When stores supply more of a good, this availability in itself may lead people to infer it is more commonly chosen.

Typicality. Participants were asked, "Is _____ a good example of _____"? One example read, "daisies . . . a flower to display in one's home." Ratings were made on a 1-to-10 scale, from 1 (*not at all*), to 10 (*very much so*). Such a measure has been used to determine whether exemplars are representative or typical members of a category (Dieciuc & Folstein, 2019; Rosch & Mervis, 1975). When an option feels more like a prototypical example, it may seem like it would be more commonly consumed.

Results and Discussion

Systematic-bias approach. Choice participants chose the common (but bland) option instead of the rare (but exciting) option 59.98% of the time. Forecast participants overestimated the frequency with which others would choose such common options ($M = 63.67\%$, $SD = 10.49\%$), $t(157) = 4.42$, $p < .001$, $d = 0.35$. Given that choosers and forecasters were different participants, this systematic-bias support for the commonness fallacy cannot be attributed to interference between these two measures.

Cue-correlation approach. We began by testing whether perceptions of commonness explained participants' forecasts. As before, we created two sets of difference scores for each choice pair—how much more common the common item was perceived to be than the rare item, and how much more likely others were forecast to select the common instead of the rare item. The same random-slope, random-intercept model used in Studies 3 and 4 provided support using the cue-correlation approach. That is, the more that one option was seen to be more common than the other, the more others were forecast to choose it over the other, $t(176.86) = 14.08$, $p < .001$.

Examining where perceptions of commonness may originate involves extending the cue-correlation approach one step back-

ward. We began by creating difference scores (common option—rare option) for our five new predictors: self-history, knowledge, recall ease, store space, and typicality. We then modified our random-slope, random-intercept model by nesting these five (difference-score) predictors within participant to predict the options' relative commonness. As indicated in Table 5, all five measures offered unique predictive power, $ts > 2.49$, $ps < .014$. Although these results are fundamentally correlational, they are consistent with the possibility that commonness perceptions have their origin in all five of our examined sources.

But did each of these commonness antecedents provide incremental validity in understanding participants' forecasts? We returned to our initial cue-correlation model, but used the five new difference scores—instead of perceived commonness—to predict choice forecasts. In this case, four of the five difference scores—all except for knowledge—predicted choice forecasts, $ts > 3.10$, $ps < .007$. When we added commonness to this model, we saw that perceived commonness continued to strongly predict choice forecasts, $t(163.75) = 8.74$, $p < .001$. Three of the four predictors continued to predict choice forecasts, although less clearly than did commonness (all $ts < 3.25$). This suggests that self-history, recall ease, and typicality judgments may contribute to the commonness fallacy in part because they speak to perceived commonness. One predictor, store space, was no longer significant with commonness included, $t(76.08) = 1.14$, $p = .259$. This is consistent with an option's amount of store space contributing to the commonness fallacy only because it speaks to perceived commonness. The next study builds on these correlational findings by providing an experimental test of the commonness fallacy and its origins.

Study 6

In documenting evidence of the commonness fallacy, we have presented findings that are essentially correlational. That is, those options that are common (because they have been frequently chosen before) are also those items that are then assumed to be more popular than they actually are in new choice sets. But even this characterization of what makes items common—that they are frequently chosen—has not yet been directly examined. Instead, we have leaned on participants' own characterizations of commonness—either through pretest ratings or measured at the individual level in our studies themselves.

Table 5
Regression Output for Possible Antecedents of Commonness Predicting Commonness and Choice Forecasts (Study 5)

Predictors	DV: Relative commonness	DV: Predicted choice	DV: Predicted choice
Self-history	.09 (.01)***	.94 (.26)***	.69 (.21)**
Knowledge	.04 (.02)*	.30 (.24)	—
Recall ease	.13 (.02)***	.99 (.26)***	.55 (.20)**
Store space	.24 (.02)***	.82 (.26)***	.26 (.23)
Typicality	.15 (.02)***	1.06 (.28)***	.71 (.23)**
Commonness	—	—	2.23 (.26)***

Note. Predictive power (unstandardized betas) of five possible antecedents of commonness on relative commonness and choice forecasts. Commonness remains a significant predictor of forecasts when included in the model with the five antecedents.
* $p < .05$.   ** $p < .01$.   *** $p < .001$.

Mindful of these limitations, we designed Study 6 to extend on this past evidence in three ways. First, we offered a causal test of our hypotheses. Second, we strengthened our ability to make a normative claim by testing whether commonness influenced forecasts even when it could be identified as a non-normative cue. Third, we directly document an origin of the fallacy by tracing participants' exposure to another's choice process to their misguided reliance on that information. Recall that in Study 5 the only predictor of the commonness fallacy entirely explained by perceived commonness was store space—the amount of (physical or virtual) shelf space that the option occupied. In Study 6, we experimentally varied how frequently different options appeared on store shelves. This allowed us to determine whether options that were commonly chosen because they were commonly offered, not because they were typically chosen *when* offered, would be seen as especially likely to be chosen in the future. In other words, we examined whether an experimental manipulation of commonness would warp forecasts of choice in a way that Study 5's correlational evidence anticipates.

Participants were told they would have to forecast the likelihood that a person named John would select each of three candy bars. To aid with this forecast, participants were allowed to observe an extensive choice history, 96 separate occasions that John selected a candy bar from two that were offered. This choice history revealed that John was perfectly indifferent between the three options: He was equally likely to select each candy bar that was offered. But because the most common candy bar was offered three times as often as the least common candy bar, John chose the former three times as often as the latter. If participants commit the commonness fallacy, they should not think that John will be indifferent when offered the choice of all three candy bars. Instead, they should think it most likely he will select the most common bar and least likely that he will select the least common bar.

Method

Participants and design. One hundred thirteen Americans from Amazon Mechanical Turk (MTurk) participated in exchange for nominal payment. No participant failed the attention check, a single item that asked what the study was about (see online supplemental materials for details). Participants were randomly assigned to one of three *commonness* conditions—designed to vary which branded candy bar was the one chosen most, second most, and least frequently. Given this was merely a counterbalancing factor, all analyses collapse across these differences. The hypotheses, measures, sample size, and analysis plan were preregistered: https://aspredicted.org/pm4y6.pdf.

Materials and procedure. Participants learned they would have to make a behavioral forecast for a target named John. To provide relevant background, we showed participants 96 previous choices that John faced and made. On each of the 96 trials, participants were told that "John walks into a store without knowing what candy bar options it offers." In this way, information about John's preferences would be revealed only by his candy bar choices, not by which stores he chose to patronize. Participants were then shown which two candy bars the store stocked. After two seconds, the option John chose was circled (see Figure 2).

Across trials, participants saw three unique candy bars: Boost, Crunchie, and Twirl. Each is an actual British candy bar. This
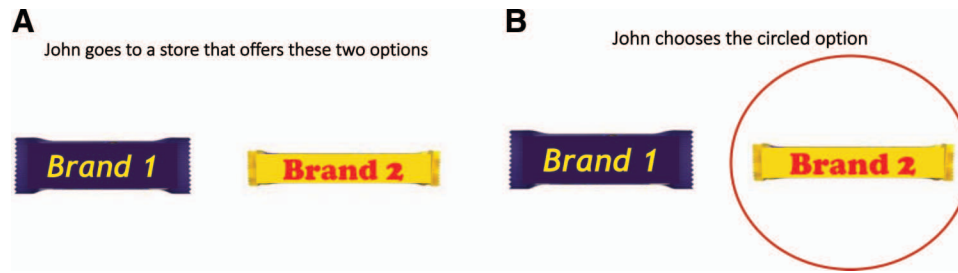
Figure 2. Mock-ups of stimuli used in Study 6. Although participants saw the original, professionally designed packaging with the commercially available brand names, those details have been altered here for copyright reasons. Screen A would appear for 2 s, followed by screen B for 2 s. This was repeated 96 times for different choice pairs, as reported in the Method section of Study 6. See the online article for the color version of this figure.

meant our American participants were unlikely to have much familiarity or prior attitudes about them. At the same time, this permitted us to use realistic stimuli, each represented by professionally designed packaging.

To modify the commonness of each choice, but to avoid John's exhibiting a relative preference for one candy bar over another, we manipulated how frequently each candy bar was offered. More specifically, the most common candy bar was present in all 96 choice contexts. John saw the second most common candy bar in two thirds (64) of the contexts. The least common candy bar was in the remaining third (32) of contexts.

In order that John not reveal any predilection for one candy bar over any other, John was equally likely to choose each candy bar on each trial. This means John selected the most common candy bar on half (48) of the 96 trials for which it was offered. Even though this candy bar was chosen most frequently, it was also offered most frequently. Similarly, John selected the second most common option on half (32) of the trials it was offered. Although John was also behaviorally indifferent between the most and least common option, this meant he chose the least common bar only 16 times. Which candy bar was most, second most, and least common was roughly equivalent across participants.

Finally, participants read, "If John had the choice between the following three candy bars, estimate the percent chance that he would choose each one." Participants saw a picture of each of the three candy bars seen previously. They offered three percentages that had to sum to 100%. We recoded these values to reflect the perceived likelihood that John would choose the most common, second most common, and least common option.

Results and Discussion

Although John's choices revealed perfect indifference among the three candy bars, we proceeded to test whether commonness guided forecasts of choice. We began by conducting two orthogonal linear contrasts. One contrast ordered the forecasts in terms of their targets' commonness: most common ($+1$), second most common (0), least common (-1). The second orthogonal contrast—which we predicted to be nonsignificant—would allow us to assess the goodness of fit of the first: most common (-1), second most common ($+2$), least common (-1).

As predicted, the first contrast was significant, $t(112) = 13.45$, $p < .001$. Suggesting this contrast fit the data well, the second

contrast was not, $t < 1$. More specifically, participants thought it more likely that John would select again the most common candy bar ($M = 46.03\%$, $SD = 14.90\%$) than the second most common candy bar ($M = 33.14\%$, $SD = 14.68\%$), paired $t(112) = 4.99$, $p < .001$, $d = 0.47$. And in turn, this latter candy bar was forecasted as more likely to be chosen than the least common one ($M = 20.83\%$, $SD = 11.11\%$), paired $t(112) = 6.13$, $p < .001$, $d = 0.58$. It is notable that these three percentages were highly similar to the percentage of the time that John chose the most common (50%), the second most common (33%), and the least common (17%) candy bar. The commonness fallacy emerges even when people were made aware of the choice contexts that led to certain items being commonly chosen.

This demonstration has some surface-level overlap with past research on people's tendency to see contingencies when none are actually present. For example, even when the direction participants pushed a joystick did not actually affect whether an image of the Lochness Monster would appear, participants misperceived the contingency—especially when the monster appeared more often (Allan & Jenkins, 1980). But note that in the present study, participants were not determining whether the presence of a candy bar correlated with its being chosen. (After all, it was not possible to have a candy bar not be present but be chosen.) In this way, we are not testing another bias in contingency detection. Instead, the present demonstration may be more related to demonstrations of denominator neglect (Mikušková, 2015; Reyna & Brainerd, 2008)—a tendency to judge proportions by the magnitude of their numerators instead of their denominators. By inferring preferences from the number of times an option was chosen while not fully adjusting for the times an option was offered, people who show the commonness fallacy make a similar error.

Study 7

Having repeatedly documented the commonness fallacy, explored its origin, and examined the conditions under which it emerges, we now move to consider how this judgment error can have costly downstream consequences. With the growth of online marketplaces and shopping communities, it is no longer only at garage sales that ordinary individuals play the role of sellers in economic exchange. Although some such marketplaces operate as auctions, in others sellers list specific prices for their goods. We developed a two-part marketplace simula-

tion to test whether the commonness fallacy would encourage sellers to misprice their goods by being more likely to suboptimally raise (vs. lower) prices for common (vs. rare) goods.

Method

Participants and design. One hundred eight-two Americans were recruited on Amazon Mechanical Turk (MTurk) to complete the main study for nominal compensation. (An additional 100 people were recruited from the same population for a pretest, described below, used to determine the optimal, profit-maximizing price for every good.) Twelve participants in the main study failed the attention check that asked what the study was about (see online supplemental materials for details). Fifteen other participants failed to follow instructions (e.g., by indicating they wished to both raise and lower the price of the same good). Data from the remaining 155 participants are reported below.

Materials and procedure. Participants were told they were taking part in a market simulation, one that would have them play the role of a Mechanical Turk shopkeeper. As part of this simulation, participants' task was to determine how they wanted to price goods they might sell to fellow Turkers. Participants were given the goal of maximizing profits. As we pointed out to them, raising prices typically reduces the number of items sold, but increases the profit margin on each sale. Lowering prices has the reverse effect: increasing sales, but reducing margins.

Participants completed 11 such simulations, corresponding to the 11 choice pairs used in Studies 3 and 4. In each, they were told to imagine they sold two items in a particular category to fellow Turkers. We asked participants to envision Turkers coming into their store and choosing whether to buy one of their items (e.g., orange juice or passion fruit juice) or nothing at all (if all the prices were too high). We then showed participants the current price of each item. Unbeknownst to participants, such default prices were the optimal, profit-maximizing prices. For each item, we asked participants, "Do you think it would be smart to *raise* the price, to *lower* the price, or to *not change* the price?" Participants indicated their responses by clicking an up arrow, a down arrow, or the price itself to indicate that they felt it would be smart to raise, low, or not change the price, respectively.

Finally, participants made forecasts of what a randomly selected other would choose if selecting between the two items. This permitted us to make certain that pricing errors could be tied to people's (known-to-be-inflated) estimates that others would choose the common item. That is, it would permit us to tie any mistakes to the commonness fallacy.

Determining the profit-maximizing price for each good. By setting the default prices as the profit-maximizing prices, we could classify any systematic desire by participants to raise or lower such prices as a bias toward overpricing or underpricing the items. To identify the optimal price for an MTurk store, we asked 100 participants on Amazon's Mechanical Turk to make two judgments about each of the 22 items. First, they indicated the maximum they would be willing to pay (WTP) for each item. We also had them estimate how much it would cost a store, on average, to supply each of these items (see online supplemental materials for detail).

The purpose of having all pretest participants indicate their maximum willingness to pay (WTP) is that we can understand the elasticity of demand—that is, how responsive interest in the goods would be to changes in price. Such knowledge is necessary for estimating the profit-maximizing values. We used participants' median cost estimates merely to understand what the average belief was about the cost to supply each item. This estimate is important for calculating the estimated profit margin on each item sold. (Ideal prices come from maximizing these margins, not from maximizing revenues.) For each product, we calculated the profit-maximizing price by determining what price would maximize this expression: number of units sold at a given price * [price − cost]. Of course, the true validity of this procedure depends on participants' accurately reporting their true WTP and our main study participants having beliefs about the costs of production that were similar to our pretest participants'. But even if this procedure systematically overestimates or underestimates the ideal price for each item, this should not prove problematic given our hypothesis is not that participants will show an overall tendency to raise or lower prices, but instead that pricing shifts will differ for common versus rare items.

Results and Discussion

First, we classified each participant's pricing recommendation for each product as a desire to raise (+1), lower (−1), or leave unchanged (0) each price. Thus, for each pair of products, we calculated a relative pricing strategy composite by taking the participant's preferred pricing strategy for the common item and subtracting off the preferred pricing strategy for the rare item. Higher values on this composite reflect hypothesis-consistent results—that is, a stronger tendency to want to overprice the common (and underprice the rare) items.

To test our main hypothesis, we predicted the relative pricing strategy while including only a random effect of participant. As expected, the intercept was significantly greater than zero, $B = 0.17$, $SE = .04$, $t(154.17) = 4.00$, $p < .001$. This reflects that participants were more likely to pursue an inappropriately aggressive pricing strategy when considering common (vs. rare) items. That is, participants were more likely to recommend raising the price on common (28.46%) than on rare items (23.49%), but more likely to recommend lowering the price on rare (45.16%) items than they were on common items (33.41%). In short, participants' decisions to depart from the profit-maximizing prices were anticipated by the distorting influence of the commonness fallacy. Pricing strategies by item are listed in Table 6.

But is this asymmetry actually a consequence of the commonness fallacy, or does it emerge for some other reason? Rare items' optimal price tended to be higher than common items' optimal price. Thus, if participants merely adopted a rule that the two items should be more evenly priced (for whatever reason), this could produce a pattern that misleadingly looked as if the commonness fallacy produced this error. To address this alternative, we examined whether individual variability in the extent to which participants displayed evidence of the commonness fallacy on any given pair of items—that is, the extent to which they thought others would choose the common over the rare items—would explain their pricing strategy on that pair.

Table 6
The Percentage of Participants Who Thought it Wise to Raise or Lower the Price of Each Common and Rare Item (Study 7)

Category	Rare option	Common option	Rare, raise price (%)	Common, raise price (%)	Rare, lower price (%)	Common, lower price (%)
Lunch	Curry	Sandwich	15.48[a]	23.23[a]	54.84[a]	39.35[b]
Dinner beverage	Japanese imported beer	Budweiser	19.35[a]	47.10[b]	46.45[a]	14.19[b]
Weeklong vacation	The Galápagos	Hawai'i	31.61[a]	44.52[b]	39.35[a]	27.74[b]
Dinner	Thai food	Pizza	26.45[a]	21.94[a]	39.35[a]	49.03[a]
Fruit	Guava	Apple	16.13[a]	20.00[a]	45.81[a]	30.32[b]
Birthday celebration (with friends)	Improv comedy	Dinner	25.81[a]	39.35[b]	53.55[a]	21.94[b]
Dessert 1	Crème brulee	Assorted cookies	16.77[a]	21.94[a]	54.19[a]	38.71[b]
Flowers	Snapdragons	Daisies	22.58[a]	18.71[a]	43.87[a]	50.32[a]
Dessert 2	Tiramisu	Vanilla ice cream	28.39[a]	20.65[a]	32.90[a]	38.06[a]
Breakfast beverage	Passion fruit juice	Orange juice	18.71[a]	24.52[a]	43.87[a]	16.77[b]
Art exhibit	Modern/ Abstract	Traditional paintings	35.48[a]	29.68[a]	39.35[a]	39.35[a]
Average			23.49[a]	28.46[b]	45.16[a]	33.41[b]

Note. Participants had the choice to raise, lower, or leave the price unchanged. This explains why the two rare percentages or two common percentages do not add to 100%. Means in the pairs of columns (i.e., raise rare and raise common, lower rare and lower common) in the same row that do not share a superscript are significantly different at the $p < .05$ level.

We added to our original model by including a Level-1 variable, *commonness fallacy,* that we nested within participant in a random-slope, random-intercept model. Finally, as in our previous studies, we included a random effect of *choice pair.* Consistent with hypotheses, those participants who showed the strongest evidence of the commonness fallacy for any given choice pair (because they were most confident that the common item would be chosen over the rare item) were also those who incorrectly thought it wise to more aggressively price the common (compared with the rare) item, $B = .004$, $SE = .001$, $t(144.77) = 6.18$, $p < .001$.

In summary, the commonness fallacy can underlie a costly mistake—pushing sellers to systematically misprice their goods. That said, our research simulation was better equipped to demonstrate *how* this error can occur than to estimate *how large* this effect is in the real world. The latter question requires an econometric, not a psychological, toolkit. Still, we hope this study offers an empirical illustration of how researchers can link basic psychological processes to monetarily consequential problems of broader interest.

General Discussion

Predicting others' choices is not merely a tricky task, but one on which people systematically err. Across eight studies, we found evidence that people lean on a heuristic cue to assist with such forecasts. Estimates of the likelihood that others will choose A over B are influenced by the perceived commonness of A versus B. People seem to be enticed by an intuitively appealing but ultimately imperfect logic—that people will choose what has been chosen before. The false equivalence of the contexts in which A and B were selected in the past and the context of choosing A or B in the present can make perceived commonness a source of social judgment error.

We provided evidence that perceived commonness guides social forecasts using two empirical strategies: the systematic-bias approach and the cue-correlation approach. Following the first approach, we identified choice pairs for which reliance on perceived commonness should produce systematic bias in social judgments.

That is, participants estimated that others would choose a common (but bland) over a rare (but exciting) item—whether in one-off or repeated choices—more often than choosers actually did. Following the second approach, we measured the heuristic cue (i.e., perceived commonness) directly and tied it to individual variation in the judgment of interest (i.e., choice forecasts). This provided a more direct test of our account by allowing us to localize the influence on forecasting to perceived commonness as opposed to correlated attributes (e.g., price). Also, it permitted us to observe the influence of perceived commonness above and beyond a more normatively defensible cue (i.e., perceived liking). We see the identification of two complementary approaches—systematic-bias and cue-correlation—as outlining a useful framework not only for organizing our own results but for informing future researchers considering how best to establish that people rely on a heuristic cue.

Beyond providing direct evidence of the influence of perceived commonness on choice forecasts, we delved deeper into the psychological process underlying and the origin of the fallacy in four ways. First, we determined that perceived commonness was a spontaneously accessible cue that was given less weight upon further thought. In other words, people did not lean on perceived commonness due to a faulty lay theory that commonness offered clearly valid evidence of what would be chosen. Instead, like with other heuristic-like processes, some feature of the judgment problem made it spontaneously accessible (Kahneman, 2003; Kahneman & Frederick, 2002), leading forecasters to rely on it (at the expense of perceived liking) somewhat mindlessly. Second, we identified that feature of the judgment task that cued the commonness fallacy. Namely, forecasters leaned on commonness when forecasting what others would choose, not what they would like to receive. This bolstered our account that the intuitive similarity between what has been chosen (commonness) and what another will choose is what makes the heuristic come online. These first two points identify (partial) boundary conditions on the commonness fallacy: Forecasters leaned less on commonness upon reflection and when forecasting what others would like to receive.

Third, we identified four possible antecedents of commonness that might give rise to the commonness fallacy due to their association with commonness: the ease of recalling instances of an item's use, an item's category typicality, one's own past use of an item, and how much store space the item receives. These variables predicted both perceived commonness and choice forecasts. And with all of these variables controlled, commonness remained the strongest predictor of forecasts. Fourth, we were able to trace the commonness fallacy from informational origin (observing someone making a series of choices that implied perfect indifference between choice options) to ultimate realization (predicting the person would choose the item that had been most frequently chosen—even though this was attributable to its being most frequently offered). This also reflected an experimental instead of a merely correlational demonstration of the commonness fallacy.

Although the commonness fallacy anticipates why dinner party hosts may be more likely to run out of rose geranium ice cream than vanilla ice cream, Study 7 showed how it may lead to more economically costly mistakes. Amateur sellers were misled by the commonness fallacy in setting goods' prices. Sellers were especially likely to lower prices on rarer items below what pretesting indicated would be profit maximizing. By tracing individual variation in the commonness fallacy to the pricing decision, we could identify perceived commonness as the likely culprit. In combination, this evidence establishes commonness as a cue that—much like the self's own preferences—systematically guides and sometimes misleads forecasts of others' choices.

Considering Alternative Explanations and Limitations

Does the systematic bias reside in forecasters or in choosers? Although we have characterized many of our results as reflecting that commonness biases forecasters, is it possible that the real bias comes from choosers themselves? Two concerns take this form. First, to the extent that indicating a choice of rare (but exciting) items is more socially desirable than a choice of commonplace (but bland) items, then perhaps social forecasting biases emerge only because participants are misrepresenting what they would choose. But several features of our design minimize this concern. In no study did participants make their choices public. Even when the studies were conducted in the lab, participants sat in private rooms where neither the experimenter nor other participants could observe their responses. Also, in Study 2, we went to great lengths to assure participants that the experimenter would never learn which candy bar they took home. Finally, the fact that making participants' choices incentive compatible did nothing to reduce the size of the forecasting error makes it less likely that our effects were driven by participants' costlessly misrepresenting their choices.

Second, although we have focused on the influence of perceived commonness on forecasters, we considered whether our findings might also reflect the influence of perceived commonness on choosers. That is, even if people typically prefer vanilla ice cream to tiramisu, people also like to engage in variety seeking (McAlister & Pessemier, 1982; Ratner & Kahn, 2002). If something about our experimental context—for example, the one-off nature of the choice—cued an interest in or curiosity for trying new or unusual things, then this may have pushed participants to select items to the extent they were seen as unusual, a tendency that would run counter to what forecasters (and the commonness fallacy) ex-

pected. But Study 1b showed that even when the same choice was made repeatedly (12 times), choosers selected the unusual option more often than forecasters thought. Clearly the commonness fallacy did not emerge merely due to a quirk of one-off choice. Furthermore, as first reported in Study 2 (and replicated in Studies 3 and 4, see footnote 5), participants who chose the rare instead of the common bar were actually those who saw it as relatively *more*, not less common. In combination, these findings speak to the robustness of the commonness fallacy and suggest it is driven by forecasters' overreliance on commonness, not choosers' real or feigned aversion to it.

Price. There is also a set of concerns related to the perceived price of the target items. One possibility is that people may use perceived price as a cue to their own choices ("I'd like the rare item because it looks expensive!") but fail to appreciate that others would have the same orientation. Five features of our studies and results address this concern. First, Study 2 used rare and commonplace items that were matched on actual price. Second, Study 4 measured perceived price and showed forecasters leaned on perceived commonness above and beyond perceived price. Third, we ran an additional study presented in the online supplemental materials that used a set of rare and commonplace pairs that were matched on perceived price. In that study, we outsourced stimulus selection to participants who took part in several rounds of pretesting, all in an effort to remove any experimenter bias from stimulus selection. Despite the absence of difference in perceived prices between the two items in each pair, results reveal a reliance on perceived commonness in forecasting others' choice. Fourth, in Studies 1–5, we simply asked what participants would choose, not what they would *buy*. In fact, the Supplemental Study described contexts in which such a cost-free exchange would naturally occur, and Study 2 had one unfold in reality. Fifth, even if perceived price did factor into participants' forecasts, it did so in a way that worked against our hypotheses. Study 4 found that participants assumed others would be more likely to select items to the extent they were perceived to be more expensive. But because our studies that did not match rare and common goods on price tended to use rare goods that were more expensive than common goods, such a pricing mismatch pushed forecasts closer to accuracy and thus underestimated the commonness fallacy.

Cultural generality. We observed remarkably consistent results across both American university students and Amazon Mechanical Turk workers. This provides initial evidence of the generality of the commonness fallacy. Of course, even if future work confirms that the commonness fallacy is a psychological universal, how it manifests in different contexts will be variable. That is, what is a common dinner item in one cultural context may be a rarity in another.

But more intriguing is that the validity of the commonness fallacy as a source of social insight may vary from culture to culture. In contexts where conformity pressures are stronger, perceived commonness may itself be a strong cue to choice (Bond & Smith, 1996; Kim & Drolet, 2003). In other contexts, knowledge that an item is especially common can make it less attractive (e.g., Berger & Heath, 2007). In other words, contextual variability in the extent to which perceived commonness informs preferences and ultimately choice should explain variability in the extent to which the commonness fallacy helps or hurts social accuracy.

Relation to Other Work

Base rate neglect. At first glance, reliance on perceived commonness may appear inconsistent with a well-studied phenomenon: base rate neglect. In a prototypical demonstration of that effect, judgments are disproportionately swayed by case-based information ("This rash looks remarkably similar to the one I saw in the dengue fever documentary!") People fail to correct for just how uncommon (and thus improbable) such an attribution may be (Ajzen, 1977; Bar-Hillel, 1980; Tversky & Kahneman, 1973). The commonness fallacy would seem to suggest that people are too embracing of baserates. But there are two crucial differences between the base rate neglect literature and the present research.

First, we study the commonness fallacy outside of contexts in which case-based information may have a deceptive allure. If we inserted case-based information that would seem to imply interest in the rare, but exciting item, reliance on perceived commonness might decline. Instead, in our studies participants are actually trying to *estimate* a base rate ("What percentage of people will choose A over B?") and find it all too easy to consult another base rate with which they have greater familiarity (commonness). Second, and relatedly, commonness is not the directly applicable base rate; it does not reflect people's past selections between two options (but may instead reflect the relative supply or ease of attaining each). That said, some options (e.g., Diet Coke, regular coffee) almost always co-occur with a corresponding option (e.g., Coca-Cola, decaffeinated coffee). In those more limited circumstances, commonness itself may be the valid behavioral base-rate.

Availability heuristic. As we have stated throughout the paper, the availability heuristic—that the ease of generating exemplars from a category is leaned upon when assessing the scope of that category—is particularly relevant to the commonness fallacy (Tversky & Kahneman, 1973). And as empirically tested in Study 5, the ease of recalling specific instances of people selecting the option does indeed factor into the perceived commonness of the option. Moreover, people also might use an egocentric version of availability—relying on their own experience with each item (e.g., "How often do I have pizza or Thai food for lunch?")—when assessing the broader commonness of each item. Study 5 supports this possibility as well. Note how this makes it all the more remarkable that the self leans on (and can be led astray by) this limited-diagnostic cue. That is, through the self's own experience, the self should realize that when it chooses to consume an item (e.g., vanilla ice cream), it did not necessarily choose it *over* another (e.g., tiramisu). Perhaps this is why Study 3 found that on closer inspection, the self will essentially abandon perceived commonness in forecasting choice.

Broader Implications

A substantial body of work has documented individuals' propensity to stick with the status quo over available alternatives (Kahneman, Knetsch, & Thaler, 1991; Samuelson & Zeckhauser, 1988). Owing to a tendency to default to the status quo, firms and individuals often fail to adequately explore alternative options (Schotter & Braunstein, 1981), policymakers continue to fund ineffective programs (Gilbert, Light, & Mosteller, 1975), and scientists continue with outmoded theoretical frameworks (Kuhn, 2012). These observations emphasize that a stability in what people choose can lead to undesirable intransigence. The commonness fallacy illustrates how suppliers of options end up being complicit in such stasis.

If people assume that others will choose what has always been chosen, it means there will be a reluctance to update what is offered for fear that it will not be chosen. When this keeps clearly superior options from entering the marketplace of goods or ideas, there is a welfare loss. Americans' general pessimism about the prospects of a successful third political party may reflect this process. But in other contexts, the commonness fallacy may help to preserve a desirable diversity of offerings. Various regions of the world continue to offer and specialize in distinct cuisines—in part, we suspect—because that is what people in those regions have historically chosen to eat. Of course, that is also what has historically been available. But the commonness fallacy may help discourage a convergence of the world's purveyors of products and ideas from merely catering to a median cultural and intellectual palette.

Conclusion

The present paper documents a new social judgment phenomenon, one that can distort people's forecasts of others' choices. When psychologists identify a bias, they often hope that educating the public about the mistake can help them be vigilant for the circumstances in which the distorting cue will lead them astray. But unfortunately, there are many instances in which awareness (Koehler, Brenner, & Griffin, 2002; Wilson & Brekke, 1994) and even explicit warnings (Sloman, 1996) do not insulate people from such errors. Because we identified what leads people to call upon perceived commonness, we can offer more substantive advice in helping people avoid the heuristic cue's pull. In contexts in which people consider what someone else would choose, they may do well to reframe this question as "Which would they be more pleased to receive?" Such reframing may help to free forecasters from the intuitive but often misleading allure of commonness and spare us all many a melted scoop of vanilla ice cream along the way.

References

Ajzen, I. (1977). Intuitive theories of events and the effects of base-rate information on prediction. *Journal of Personality and Social Psychology, 35,* 303–314. http://dx.doi.org/10.1037/0022-3514.35.5.303

Allan, L. G., & Jenkins, H. M. (1980). The judgment of contingency and the nature of the response alternatives. *Canadian Journal of Psychology/Revue Canadian Psychology, 34,* 1–11. http://dx.doi.org/10.1037/h0081013

Ariely, D., & Levav, J. (2000). Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of Consumer Research, 27,* 279–290. http://dx.doi.org/10.1086/317585

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language, 59,* 390–412. http://dx.doi.org/10.1016/j.jml.2007.12.005

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica, 44,* 211–233. http://dx.doi.org/10.1016/0001-6918(80)90046-3

Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology, 54,* 569–579. http://dx.doi.org/10.1037/0022-3514.54.4.569

Berger, J., & Heath, C. (2007). Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research, 34,* 121–134. http://dx.doi.org/10.1086/519142

Bohns, V., & Flynn, F. (2010). "Why didn't you just ask?" Underestimating the discomfort of help-seeking. *Journal of Experimental Social Psychology, 46,* 402–409. http://dx.doi.org/10.1016/j.jesp.2009.12.015

Bond, R., & Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) line judgment task. *Psychological Bulletin, 119,* 111–137. http://dx.doi.org/10.1037/0033-2909.119.1.111

Critcher, C. R., & Dunning, D. (2013). Predicting persons' versus a person's goodness: Behavioral forecasts diverge for individuals versus populations. *Journal of Personality and Social Psychology, 104,* 28–44. http://dx.doi.org/10.1037/a0030836

Critcher, C. R., & Rosenzweig, E. L. (2014). The performance heuristic: A misguided reliance on past success when predicting prospects for improvement. *Journal of Experimental Psychology: General, 143,* 480–485. http://dx.doi.org/10.1037/a0034129

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes, 65,* 201–211. http://dx.doi.org/10.1006/obhd.1996.0020

Dieciuc, M. A., & Folstein, J. R. (2019). Typicality: Stable structures and flexible functions. *Psychonomic Bulletin & Review, 26,* 491–505.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology, 87,* 327–339. http://dx.doi.org/10.1037/0022-3514.87.3.327

Epley, N., & Schroeder, J. (2014). Mistakenly seeking solitude. *Journal of Experimental Psychology: General, 143,* 1980–1999. http://dx.doi.org/10.1037/a0037323

Frederick, S. (2012). Overestimating others' willingness to pay. *Journal of Consumer Research, 39,* 1–21. http://dx.doi.org/10.1086/662060

Gigerenzer, G. (2008). *Rationality for mortals: How people cope with uncertainty.* New York, NY: Oxford University Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science, 1,* 107–143. http://dx.doi.org/10.1111/j.1756-8765.2008.01006.x

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology, 62,* 451–482. http://dx.doi.org/10.1146/annurev-psych-120709-145346

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103,* 650–669. http://dx.doi.org/10.1037/0033-295X.103.4.650

Gilbert, J. P., Light, R. J., & Mosteller, F. (1975). Assessing social innovations: An empirical base for policy. *Evaluation and experiment: Some critical issues in assessing social programs,* 39–193.

Helzer, E. G., & Dunning, D. (2012). Why and when peer prediction is superior to self-prediction: The weight given to future aspiration versus past achievement. *Journal of Personality and Social Psychology, 103,* 38–53. http://dx.doi.org/10.1037/a0028124

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist, 58,* 697–720. http://dx.doi.org/10.1037/0003-066X.58.9.697

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 49–81). New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511808098.004

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives, 5,* 193–206. http://dx.doi.org/10.1257/jep.5.1.193

Kim, H. S., & Drolet, A. (2003). Choice and self-expression: A cultural analysis of variety-seeking. *Journal of Personality and Social Psychology, 85,* 373–382. http://dx.doi.org/10.1037/0022-3514.85.2.373

Koehler, D. J., Brenner, L. A., & Griffin, D. (2002). The calibration of expert judgment: Heuristics and biases beyond the laboratory. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 686–715). Cambridge, UK: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511808098.041

Krueger, J. I. (2000). The projective perception of the social world: A building block of social comparison processes. In J. Suls & L. Wheeler (Eds.), *Handbook of social comparison: Theory and research* (pp. 323–351). New York, NY: Plenum Press/Kluwer. http://dx.doi.org/10.1007/978-1-4615-4237-7_16

Krueger, J. I. (2003). Return of the ego—Self-referent information as a filter for social prediction: Comment on Karniol (2003). *Psychological Review, 110,* 585–590. http://dx.doi.org/10.1037/0033-295X.110.3.585

Kuhn, T. S. (2012). *The structure of scientific revolutions: 50th anniversary edition.* Chicago, IL: University of Chicago Press. http://dx.doi.org/10.7208/chicago/9780226458144.001.0001

McAlister, L., & Pessemier, E. (1982). Variety Seeking Behavior: An Interdisciplinary Review. *Journal of Consumer Research, 9,* 311–322. http://dx.doi.org/10.1086/208926

Mikušková, E. B. (2015). The denominator neglect in decision-making. *Studia Psychologica, 57,* 255–269. http://dx.doi.org/10.21909/sp.2015.03.698

Oppenheimer, D. M. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science, 15,* 100–105. http://dx.doi.org/10.1111/j.0963-7214.2004.01502005.x

Prentice, D. A., & Miller, D. T. (1993). Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology, 64,* 243–256. http://dx.doi.org/10.1037/0022-3514.64.2.243

Quester, P., & Steyer, A. (2010). Revisiting individual choices in group settings: The long and winding (less traveled) road? *Journal of Consumer Research, 36,* 1050–1057. http://dx.doi.org/10.1086/644750

Ratner, R. K., & Kahn, B. E. (2002). The impact of private versus public consumption on variety-seeking behavior. *Journal of Consumer Research, 29,* 246–257. http://dx.doi.org/10.1086/341574

Ratner, R. K., Kahn, B. E., & Kahneman, D. (1999). Choosing less-preferred experiences for the sake of variety. *Journal of Consumer Research, 26,* 1–15. http://dx.doi.org/10.1086/209547

Reyna, V. F., & Brainerd, C. J. (2008). Numeracy, ratio bias, and denominator neglect in judgments of risk and probability. *Learning and Individual Differences, 18,* 89–107. http://dx.doi.org/10.1016/j.lindif.2007.03.011

Rogers, T., Moore, D. A., & Norton, M. I. (2017). The belief in a favorable future. *Psychological Science, 28,* 1290–1301. http://dx.doi.org/10.1177/0956797617706706

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7,* 573–605. http://dx.doi.org/10.1016/0010-0285(75)90024-9

Rosenzweig, E., & Critcher, C. R. (2014). Decomposing forecasting: The salience-assessment-weighting (SAW) model. *Current Directions in Psychological Science, 23,* 368–373. http://dx.doi.org/10.1177/0963721414540300

Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology, 13,* 279–301. http://dx.doi.org/10.1016/0022-1031(77)90049-X

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty, 1,* 7–59. http://dx.doi.org/10.1007/BF00055564

Schotter, A., & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry, 19,* 1–25. http://dx.doi.org/10.1111/j.1465-7295.1981.tb00600.x

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology, 61,* 195–202. http://dx.doi.org/10.1037/0022-3514.61.2.195

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2013). Life after P-hacking. Meeting of the Society for Personality and Social Psychology, New Orleans, LA. Retrieved from SSRN: https://ssrn.com/abstract=2205186

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3–22. http://dx.doi.org/10.1037/0033-2909.119.1.3

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5,* 207–232. http://dx.doi.org/10.1016/0010-0285(73)90033-9

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124–1131. http://dx.doi.org/10.1126/science.185.4157.1124

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116,* 117–142. http://dx.doi.org/10.1037/0033-2909.116.1.117