American Sociological Review
2017, Vol. 82(2) 297–327

DOI: 10.1177/0003122417691503
journals.sagepub.com/home/asr

It's the Conventional Thought That Counts: How Third-Order Inference Produces Status Advantage

Shelley J. Correll,[a] Cecilia L. Ridgeway,[a] Ezra W. Zuckerman,[b] Sharon Jank,[a] Sara Jordan-Bloch,[a] and Sandra Nakagawa[a]

Abstract

A core claim of sociological theory is that modern institutions fall short of their meritocratic ideals, whereby rewards should be allocated based on achievement-related criteria. Instead, high-status actors often experience a "status advantage": they are rewarded disproportionately to the quality of their performance. We develop and test a theory of status advantage in meritocratic settings. The most influential model in past research derives status advantage from decision-makers' tendency to infer quality from status when quality is uncertain. The theory developed here integrates and extends this and other theories to explain the emergence of status advantage in the many meritocratic contexts where the decision-maker's personal, first-order sense of quality is less important to the decision. We argue that in such contexts, decision-makers must often coordinate with others to make the "best" decision, and thus they focus on the "third-order inference" problem of discerning who or what "most people" think is higher quality, as encoded in status beliefs. Three experiments demonstrate that under such conditions, status advantages can emerge even though (1) status information does not resolve uncertainty about quality; (2) the status belief is illegitimate; and (3) no party to the decision personally prefers the higher-status option. The theory implies that status hierarchies are resilient in the face of significant dissent but may be subject to public challenge.

Keywords

status, decision-making, meritocracy, social valuation, coordination

The tension between achievement and ascription lies at the center of modern life. Agents of the central institutions of modernity—government, science, formal organizations—justify their decisions on the basis of achievement-based criteria, whereby rewards and privileges are allocated on the basis of merit, rather than being based on "task-irrelevant but socially valued" aspects of an actor's identity (Lynn, Podolny, and Tao 2009:762; cf. Parsons 1951; Weber 1968). And yet a core claim of contemporary sociological theory is that modern institutions often fall short of satisfying the ideal of meritocratic allocation (Castilla and

[a]Stanford University
[b]MIT Sloan School of Management

Corresponding Author:
Shelley J. Correll, Department of Sociology, Stanford University, 450 Serra Mall, Stanford, CA 94305-2047
E-mail: scorrell@stanford.edu

Benard 2010). In particular, higher-status actors—those who are ranked higher than others, as enacted in public displays of honor or deference (see Goode 1978; Homans 1950; Weber 1968)—are seen as preferred targets of interaction or exchange beyond the level that can be justified on the basis of "objective" performance criteria (i.e., blind to the performer's identity). Moreover, such status advantages are important because they are at the root of cumulative advantage processes, as higher-status actors build on superior resources to develop greater capabilities over time (DiPrete and Eirich 2006; Merton [1948] 1968, 1995; Ridgeway and Correll 2004; Simcoe and Waguespack 2010).

In one type of often-studied status advantage, ostensibly meritocratic institutions fail to prevent diffuse status characteristics, such as race, gender, or motherhood, from influencing evaluations of performance (Castilla 2008; Castilla and Benard 2010; Correll and Benard 2006; Correll, Benard, and Paik 2007; Petersen and Saporta 2004; Ridgeway 2011; Turco 2010). A second type of status advantage documented in such institutions pertains to the development of status rankings, such as the ranking of firms within a given industry. Unlike students of stratification and discrimination, economic sociologists tend to regard such status advantages as a legitimate form of competitive advantage, because they emerge endogenously from market competition (Podolny 1993, 2005; see also Jensen 2006; Stuart, Hoang, and Hybels 1999). This suggests that status advantages may not always reflect a failure of meritocratic institutions to fulfill their objectives, but may be inherent in how such institutions operate.

But how and why do high-status advantages emerge in meritocratic contexts where decision-makers are supposed to justify their decisions on the basis of relative performance? The most influential approach to this question is the Socially Endogenous Inferences (SEI) model. This theory, which was extended and elaborated by Lynn and colleagues (2009; see also Gould 2002; Podolny 1993, 2005; Salganik, Dodds, and Watts 2006) and received strong empirical validation in several recent studies (Azoulay, Stuart, and Wang 2014; Kim and King 2014; Lynn et al. 2016; Simcoe and Waguespack 2010), sees status advantages as emerging when performances are difficult to assess ex ante. In these settings, decision-makers infer the quality of a performance from the choices of prior decision-makers, as encoded in a publicly observable status hierarchy. Each decision is based not on a direct assessment of quality but on others' assessments (which in turn are based on others' assessments, and so on), so one actor may rise in status relative to another actor of equal quality merely because chance factors awarded the former with an early lead in status. Thus, ironically, the endogeneity of such reasonable inferences tends to have the cumulative effect of decoupling status from quality.

The SEI model has proven useful for explaining how status advantages can emerge even in meritocratic contexts. But the fact that the processes captured by this model may be sufficient to explain status advantage does not mean that such processes are necessary to produce such advantage. Indeed, as we will discuss, status advantages seem to occur even when decision-makers are not primarily focused on inferring quality from status according to their personal standards. Our key point of departure is the recognition that decision-makers in meritocratic contexts often face a different inference problem than the one assumed by the SEI model. In meritocratic contexts, decision-makers are often *interdependent* with key resource-holders for the success of their decision. The "best" decision in such contexts is often the one that most effectively takes into account their likely reactions. Accordingly, the most salient challenge for a decision-maker is to infer how various choices are most likely to be evaluated by a third party. The decision-maker can then use these inferences to find a choice that coordinates well with the anticipated reactions of relevant resource-holders. We refer to the process of coordinating with relevant others via the anticipated criteria of most others as

"third-order inference" (Ridgeway and Correll 2006). We will argue that decision-makers generally solve these third-order inference problems by relying on status beliefs.[1]

STATUS ADVANTAGE WHERE INFERRING QUALITY FROM STATUS IS NOT CENTRAL

Our general objective is to understand the conditions under which decisions are shaped by "status beliefs"—beliefs that are discussed publicly as valid, and which rank individuals, types of people, or objects according to their expected ability to contribute to valued outcomes (Berger et al. 1977; Ridgeway et al. 2009; Ridgeway and Correll 2006). Specifically, we seek to explain the conditions under which "status advantages" will be present. Status advantages are outcomes from resource-allocation decisions where higher-status actors earn an allocation that is disproportionate to their relative quality.[2] In particular, our focus is on how status advantages emerge even when decisions are justified on a meritocratic basis—that is, based on the satisfaction of quality standards that can be articulated independently of the options available. Such quality standards could be (1) shared among multiple actors (e.g., a decision-maker and any beneficiary of the decision); (2) particular to a given decision-maker (if the context is one where differences in taste are recognized and considered legitimate); or (3) particular to the intended beneficiary of the decision (e.g., if the decision is made to satisfy a particular agent or audience). In any such meritocratic context, a higher-status actor can be said to have an advantage when the actor is not superior to lower-status actors on the relevant standard but is nonetheless favored, even when decision-makers justify their decisions based on what is objectively best.

To flesh out our approach, we begin by noting how several lines of sociological theory account for status advantages in contexts that are meritocratic but where efforts to infer a personal sense of quality from status do not appear to be important. After considering these contexts and reviewing the lines of theory developed to explain status advantage in them, we build on Ridgeway and Correll (2004, 2006) to provide an integrative framework ("third-order inference"). This framework furnishes a set of testable hypotheses that can account for status advantage in settings that fall outside the scope of the SEI model; and, as discussed below, it also applies to many contexts that are within the scope of the SEI model.

One critical line of theory motivating our approach shows how the introduction of public rankings in an industry by third-parties (i.e., neither the decision-makers nor the producers whose performances are being selected) has profound effects on demand and supply, creating competitive advantage for those who are accorded a high rank (Espeland and Sauder 2016). This phenomenon is noteworthy, in part, because researchers operating within the SEI model (e.g., Gould 2002; Lynn et al. 2009) generally work within a social networks tradition that assumes interactants are exposed to one another's deference patterns (see Zuckerman 2010). But if the introduction of public rankings has an impact even when there was already a structure of deference, this suggests that *publicity is key*. Relatedly, one might suppose that such acts of publicity merely aggregate information about relative quality implied by patterns of deference, but such rankings also introduce distinctions (e.g., overly crude distinctions between some ranks, and overly fine distinctions between others) that cannot be reduced to quality distinctions (Sauder 2006). This raises questions as to why such rankings would be heeded. One possibility is that the accuracy of rankings for measuring quality is less important than that the rankings are public and decision-makers thus expect them to be widely used.

This possibility dovetails with key observations made by economic sociologists regarding the advantages enjoyed by high-status firms, particularly in a business-to-business context. Podolny (1993, 2005) relied on the SEI model in suggesting that high-status firms are favored because they are perceived

to be higher quality, but a different logic may also be at work (see Jensen 2006; Uzzi and Lancaster 2004), whereby the appeal of a high-status firm is that it provides a surer basis for projecting accountability. In short, the decision-maker's view of actual quality may often be moot insofar as the selection of a high-status firm is more easily defended to consequential others. Perhaps the most famous expression of this mechanism is the mantra from the era of mainframe computers: "No one ever got fired for buying IBM." Or as Keynes (1936:158) put it, "Worldly wisdom teaches us that it is better to fail conventionally than succeed unconventionally." This quip is too strong because, by definition, a "successful" decision is "better" than a less successful one. The problem, however, is that the review of a decision may take place well before it can be adjudicated on the basis of objective quality standards—and so in the interim, a "conventional" option (in this case, a high-status choice) may be the superior strategy. Note that this logic may hold even though the decision-maker does not know the quality standards that are personally endorsed by those who might review the decision. Did everyone really prefer IBM?

Other research points to how interactants' personal sense of quality often seems to pale in importance relative to what is publicly endorsed in a status hierarchy. This is clearest in research on "politics of dissimulation," described by scholars of authoritarian systems (and of the ancien régime; Elias 1983), whereby public displays of loyalty to the official ideology and to leaders mask a great deal of private dissent (e.g., Jowitt 1974; Kuran 1995; Wedeen 1999). More subtle examples may be found in academia, where many citations may be "ceremonial" (Adatto and Cole 1981), reflecting less about the citer's or the intended readers' personal beliefs in the quality of the cited paper than that the citation conveys the appropriate signal to reviewers and readers. Relatedly, consider the cases noted by Centola, Willer, and Macy (2005:1010; see also Willer, Kuwabara, and Macy [2009] where experimental evidence is adduced), in which "prestigious scholars . . . are widely proclaimed as having the most brilliant new ideas" but "privately, people find the work entirely incomprehensible."

Finally, note how a similar basis for status advantage appears to be highly salient in cultural domains where the SEI model faces the most difficulty. In particular, whereas it seems relatively unproblematic for a scientist to infer that a method or theory popular among her peers would be more effective than a less popular one, it is less clear that this logic applies when it comes to cultural practices, insofar as variation in taste is expected. Accordingly, research on status advantages in this context emphasizes the importance of public discussion of quality in explaining the influence of status beliefs. For example, Bourdieu's (1984) theory of "distinction" places more emphasis on how "highbrow" culture is socially validated than on personal assessments of quality. Similarly, Goffman's (1951, 1959) dramaturgical approach to impression management invokes a complementary basis for the conferral of status, and Veblen's (1953) theory of conspicuous consumption is typically interpreted to imply that consumers favor high-status products not because they regard them as higher quality, but simply because they are socially recognized as high status (see Benjamin and Podolny 1999; Malter 2014). Relatedly, Clark, Clark, and Polborn (2006) suggest that high-status actors' consumption decisions are influential not because of their keen sense of quality, but because they tend to be socially visible or prominent, and such visible endorsement of a particular item suggests it will be popular.

The various lines of theory we reviewed generally do not engage with one another. Nor have they clearly established an alternative mechanism to the inference of quality from popularity, which is at the heart of the SEI model, but identification of such an alternative mechanism seems promising. Observations from these lines of research are highly complementary, in that they recognize that decision-makers may favor (for affiliation, exchange, or as a source of influence) those

who are publicly recognized as higher quality, independent of the personal assessments of quality held by decision-makers and the intended beneficiaries of their decisions.

THIRD-ORDER INFERENCE MODEL

We argue that the mechanism linking these lines of research is decision-makers' desire to anticipate how "most others" will assess quality. Following Ridgeway and Correll (2006), we call problems of judging what "most others" think is best third-order inference problems. By contrast, first-order inference problems involve the decision-maker deciding what she herself thinks is best, and second-order inference problems concern deciding what *a specific other* thinks is best.[3] When actors need to coordinate to make a decision about selecting one exchange partner or another, they can be expected to draw on status beliefs to solve a third-order inference problem. In all the situations described in the prior section, decision-makers have good reason to favor high-status actors independent of either personal or shared quality standards: status beliefs provide an effective means for coordinating among the relevant parties to make what will be regarded by "most others" as a good choice. We now elaborate on these ideas and develop testable hypotheses that help validate our approach.

Our motivating observation is that decision-makers in meritocratic contexts often face inference problems that require the resolution of a more other-oriented type of uncertainty than that addressed by the SEI model. Whether or not a decision-maker feels uncertain about the best choice for private use, she often faces uncertainty about which choice coordinates most effectively with the actions of consequential others.[4] Ridgeway and Correll (2004) term a "social relational context" any situation in which an actor must take the expected reactions of others into account in deciding how to act, because the others' reactions will be consequential to her interests. They argue (cf. Ridgeway 2011) that commonly known beliefs about types of people, such as those contained in gender stereotypes, serve as common codes that actors use to solve coordination problems posed in a wide variety of relational contexts where that code is salient.

More generally, public discourse tends to specify a common code for how actors and institutions are categorized and ranked in a given social domain, thereby providing a ready basis for coordination among domain-participants who are aware of the code and seek to justify their decisions on the basis of relative performance. We assume that such common codes, in this case, widely shared status beliefs, become implicitly salient for decision-makers when they must choose among alternatives that differ in rank according to these status beliefs. When salient, a widely shared status belief that ranks the relevant choices serves as a "focal point" (Schelling 1960), or a default convention that facilitates coordination by allowing the decision-maker to anticipate the likely reactions of consequential others to alternatives in order to consider these in her own choices. This happens not because the common code is widely endorsed as a personal conviction, but because it is what everyone knows that everyone knows, or what Ridgeway and Correll (2006) call a "third-order belief" (see also Adut 2008; Chwe 2001; Swidler 2001; Yamagishi and Kiyonari 2000).

To specify further, the presumption that a status belief is what "most people think most people believe" and is, in that sense, cultural common knowledge, gives it a public, *socially valid* quality, even if the decision-maker does not personally endorse it as proper. Research on legitimacy shows that the apparent social validity of something (e.g., a status belief), and the associated sense that others can be expected to act in accord with it, powerfully shapes individual behavior *independent of the individual's sense that it is right and proper* (for a review, see Zelditch 2006). Indeed, Clark and Kashima (2007) show that people spontaneously draw on common codes, such as

gender stereotypes, to communicate and facilitate coordination with strangers, even when they could communicate on the basis of stereotype inconsistent information. Thus, the quality of social validity gives a widely held status belief a kind of facticity for the decision-maker, both allowing anticipation of the reactions of consequential others and suggesting that the decision-maker consider this belief in her own behavior to make the most effective choice, even if she does not endorse the status belief.

As the above implies, the primary scope condition for the third-order inference model, and its point of contrast with the SEI model, is that a decision-maker is making a decision under a non-negligible degree of *interdependence* with consequential others whose reactions will affect the success of the decision. Accordingly, we might posit a range of decision contexts, from those where decision-makers are highly interdependent to those where the degree of interdependence is negligible. For example, when choosing a neurosurgeon to operate on you, the reactions of others might matter little. But with acts of conspicuous consumption, like buying an ostentatious house or market speculation, the success of a choice depends almost entirely on others' reactions (see Zuckerman 2012). A midpoint in this range would be hiring decisions in organizations, in which the success of a hire is substantially (see Rivera 2015) but not entirely dependent on the reactions of consequential others.

We should note that our scope condition of non-negligible interdependence does not imply a powerless decision-maker. In interdependent contexts such as groups and organizations, even leaders have non-negligible interdependence with their subordinates that affects the success of their decisions, and thus they are likely to engage in some degree of third-order inference to make the best decision. That said, within such contexts, third-order inference should be particularly determinative for the decisions of less powerful actors.

COMPARING THE TWO MODELS

So far, we have described how the SEI model provides a good account of how status beliefs help decision-makers resolve, ex ante, their first-order uncertainty about what choice would be "best" for personal consumption. Third-order inference processes, in contrast, help decision-makers resolve social uncertainty about what choice will be most effective, given interdependence. To distinguish the effects of third-order inference, it is useful to examine them after a decision-maker has made a first-order choice about quality and the distinction is especially clear where post-decision first-order judgments are found to conflict with the socially high-status choice. Status advantage is expressly not predicted under such cases by the SEI model (see Lynn et al. 2016). Yet although the SEI model is not relevant to post-decision status advantages, we do not suggest that third-order inference processes are necessarily irrelevant to the emergence of status advantage ex ante in interdependent contexts. Rather, in most such decision contexts, the processes of resolving personal uncertainty about what is "best" and resolving social uncertainty about what consequential others will accept as best are intertwined and actors may often not distinguish between them.

To illustrate this point, Figure 1 depicts how two factors—degree of interdependence and quality uncertainty—determine the contexts covered by the third-order inference model in relation to the SEI model. Status advantage should not occur in contexts lacking both interdependence and quality uncertainty (the upper-left-hand cell), but all other settings are conducive to some degree of status advantage. The SEI model uniquely explains status advantage in contexts with negligible outcome interdependence but moderate or greater first-order uncertainty about quality (the upper-right-hand cell). In addition to our earlier example of one's choice of a neurosurgeon, the now-classic "musiclab" study (Salganik et al. 2006;

		Issues of Resolving Quality Uncertainty[a]		
		Insignificant	*Moderate*	*Significant*
Decision Outcome Interdependence[b]	*Insignificant*	Neither model	SEI	SEI (e.g., Lynn et al. 2016; Salganik et al. 2006)
	Moderate	TOI	TOI & SEI	TOI & SEI
	Significant	TOI (Studies 1, 2, and 3)	TOI & SEI	TOI & SEI

Figure 1. Comparing the Scope of the Socially Endogenous Inference (SEI) and Third-Order Inference (TOI) Models
Note: Cell entries indicate whether the context is within the scope of either or both models.
[a]Status has a significant impact on the determination of the "best" choice when the decision-maker believes her first-order estimate of relative quality is imprecise and that status information can be used to increase precision.
[b]Status has a significant impact on the determination of the "best" choice when the decision-maker believes the effectiveness of a decision depends on the actions of consequential others, and that these others are known to be aware of the status hierarchy as the public code for what is "best."

Salganik and Watts 2008) provides an example of research that addresses this cell. As reviewed by Zuckerman (2012:227–28), that research demonstrated how early leads in music popularity build on themselves, even when decision-makers have no reason to care about how others will react to or coordinate with their decision. The only plausible basis for status advantage in such a context is that postulated by SEI: the inference of first-order quality from popularity.

Although our third-order inference (TOI) model is inapplicable in the upper-right-hand cell, the TOI model uniquely explains status advantage in the bottom-left cell. These are contexts in which outcome interdependence is at least moderate (e.g., the cases of hiring for compatibility mentioned earlier) or even greater (as in conspicuous consumption or speculation), but first-order uncertainty about quality is not at issue. In the context of our three studies discussed below, the absence of personal or first-order uncertainty could stem either from the fact that it has already been resolved (Studies 1 and 2) or from the fact that it has negligible importance to the success of the outcome (Study 3). Note, however, that although research conducted in the top-right or bottom-left cells shows the importance of the SEI and TOI models, respectively, most settings (e.g., organizations and markets) are in between these extremes. Thus, when we observe status advantage in such settings, we can generally assume that both SEI and TOI processes are in play, jointly shaping the emergence of status advantage from decision-makers' choices.

Three final points further clarify the logic of third-order inference and help motivate our empirical tests of the theory. First, this logic is often invoked in the context of impression and reputation management (see earlier discussion), but we argue that the logic of third-order inference can be triggered *even when a decision-maker is anonymous and is motivated by the prospect of gains rather than to avoid losses*.[5] All that is necessary is that the decision-maker believes the probability of such gains depends on successful coordination with others on the basis of what "most others" think is best. Accordingly, and as we will discuss, our first two studies create instances of such decision contexts.

Second, it is important to determine when individuals will coordinate based on status beliefs that are generally viewed as illegitimate (in the sense of not right or proper) bases for choosing among actors.[6] For example, U.S. Equal Employment Opportunity regulations emerged based on a belief that adults of all sexes and races are assumed equal in their aptitude and commitment to work. Yet, research continues to find that employers' quality judgments about workers

are affected by gender, race, and other characteristics (see, e.g., Correll et al. 2007; Ridgeway 2011; Turco 2010). *We argue that status beliefs often have force in social relations, thereby supporting status advantages, even when those beliefs are illegitimate*. In particular, this will be true insofar as the decision-maker understands the decision as one that requires coordination based on what "everyone" knows (see also Foschi 2000).

Third, we argue that status advantages are resilient even in the face of a "contrarian" code. By a contrarian code, we mean a belief about the (classification and) relative quality of actors that deviates from a community's publicly known status beliefs. Such codes may be unique to individuals or shared among a subpopulation in a community. But insofar as they are not as widely shared and publicized (i.e., not common knowledge), they may be less effective focal points for coordination (see Heath, Ho, and Berger 2006). In particular, consider the dilemma you face if you need to give a gift to someone who prefers a low-status producer, such as Pabst Blue Ribbon beer, which, despite being widely recognized as a "cheap beer," has recently developed a cult following among young, "hipster" drinkers (Gillespie 2015). It might seem straightforward that one should select a gift on the basis of the second-order belief—that is, what the gift-recipient likes (Caplow 1984). The problem, however, is that the gift-recipient is quite aware that "most people" regard the high-status option (e.g., craft beer) as best. In this situation, it is unclear whether the lower- or higher-status option (e.g., Pabst Blue Ribbon or a craft beer) will more effectively convey the intended message of the gift—that one values the gift-recipient and therefore "cares enough to give the best." Moreover, insofar as third-parties are observing and reviewing the decision, it may become even harder to convey such a message via the contrarian code. Thus, *we argue that status beliefs can have force in social relations, thereby creating status advantages, even when coordinating parties do not endorse them.* In particular, this will be true insofar as the decision-maker expects the other party will coordinate on the basis of such beliefs rather than their contrarian beliefs. This is especially likely in situations in which the contrarian code is not clearly established as common knowledge among all situational participants (Chwe 2001). As captured in our third study, the canonical context is gift-giving, where it is the *conventional* "thought" that is most clearly conveyed.

Empirical Validation

We test our theory with three experimental studies.[7] The main goal of the first study, which involves selection of high- versus low-status brands of chocolate, is to demonstrate how status advantages can emerge even in circumstances where status information does not reduce quality uncertainty. Such a case is outside the scope of the SEI model, because any favoritism toward the high-status actor cannot be rationally justified as based on the perception that it is higher quality. However, insofar as the decision-maker thinks the decision needs to be coordinated with others who are aware of the status belief (i.e., interdependence is non-negligible), we would expect her to favor the high-status producer anyway. As noted in the prior section, this experiment shows that status advantages can emerge even when (1) the decision-maker is anonymous, and (2) the motivation is to obtain gains rather than to avoid losses.

The main goal of the second study, which involves selection of a male or female police chief for U.S. towns, is to test whether decision-makers will shift to the high-status option even when such favoritism is illegitimate. This setting highlights shared standards rather than matters of individual taste. Once again, the experiment is designed in such a way that the introduction of status information is not accompanied by information about relative quality.

Finally, the third study returns to the selection of chocolate and tests whether a decision-maker will favor the high-status option even when she knows that the gift-recipient's

personal preference is for the low-status option. This test is critical because it rules out the possibility that decision-makers merely cater to the taste of those with whom they are coordinating. Rather, third-order inference has a powerful influence on the decision; the gift-giver often opts for what everyone knows is best rather than what the gift-recipient thinks is best.

STUDY 1: THE CHOCOLATE TASTE TEST

The first experiment provides an initial empirical test of our argument that coordination based on third-order inference is sufficient to produce status advantage. In the experiment, participants decide between higher- versus lower-status actors (specifically, high- or low-status chocolate producers) under conditions where the status information fails to provide information reducing first-order quality uncertainty. The study contains four conditions: two third-order inference conditions, one first-order inference condition, and one peer coordination condition. In the third-order inference conditions, participants' decisions depend on successful coordination with a large population whom they can assume is aware of the relative status of their decision options. In the first-order condition, their decision is affected solely by their personal criteria about quality. In the peer condition, participants are led to believe they will have the opportunity to discuss their choices with others who are socially similar, thereby increasing the likelihood of successful coordination based on a contrarian code. *Our main prediction for Study 1 is that status advantages will be limited or nonexistent in the first-order inference condition but will be significantly greater in the third-order inference conditions.* We also expect limited status advantage in the peer coordination condition.

Overview of Study 1

In each of two rounds, paid undergraduate volunteers rated two samples of chocolate purportedly from different chocolate producers. In actuality, both samples were of the same chocolate, thereby ensuring they were objectively of equal quality. Eighty-seven participants were randomly assigned within sex to the four experimental conditions. Sex of participant was counterbalanced within conditions. We excluded data from seven participants because of experimental error, leaving an analytic sample of 80 (40 male and 40 female) participants (20 per condition).

In the first round of all conditions, participants rated the two chocolates, with no brand names attached, and chose their favorite. This initial choice constituted their personal, or first-order, judgment about which chocolate was higher quality. In the second round, participants were told that both chocolates were of high quality, but one of the chocolates was a higher-status brand and the other was a lower-status brand. Participants also learned that they preferred the low-status brand on the first round. They then received information that manipulated their sense of the "coordination challenge" for their decision (i.e., how and whether the outcome of the decision depended on the actions of other actors). Participants then rated the chocolates again and made their final decision. Insofar as there was significant switching from the initial choice of the low-status brand to the higher-status brand, this would constitute a status advantage for the high-status chocolate producer.

Distinguishing Features of the Experimental Design

The chocolate tasting task is, by design, one for which participants are likely to have relatively little first-order uncertainty about quality judgments, defined as the chocolate they "like best." Most individuals are familiar with tasting chocolate and likely feel they have expertise as the legitimate arbiters of their personal preferences. A small pre-test clearly established that participants were able to confidently pick the chocolate they liked best.[8] In this, our experiment is intentionally different from others in which the decision-maker

could potentially worry that she does not have sufficient expertise or authority to judge the product at hand (e.g., Willer et al. 2009).

The focus of our experiment is whether subjects switch to a high-status actor (i.e., chocolate producer) even after having sampled the "performances" of each of the two actors and personally judged their relative quality. In this way, our experimental conditions are outside the scope of the SEI model, which is geared to explaining why high-status actors are favored prior to such experiential "tests" of relative quality (Lynn et al. 2016). Finally, psychological research (see Albarracin and Vargas 2010) suggests that once decision-makers make a choice, they should feel committed to that choice and be resistant to switching to an alternative choice. This set of conditions combines to create a situation in which status advantages should not emerge unless, as we argue, contexts inducing decision-makers to engage in third-order inference are sufficient to create such advantages.

Finally, our procedures create contexts where subjects are anonymous (thus eliminating impression management or conspicuous consumption motives) and they stand to gain rather than lose by favoring high-status actors. If indeed our hypothesis is validated in this context, it would demonstrate the distinctive range of the third-order inference effects we predict.

Procedures

Round 1. Participants came to the lab individually to participate in a blind taste test of two brands of premium chocolate. Participants were told that "chocolate makers have become interested in increasing the market for their products," so they are conducting taste tests "in a variety of settings such as shopping malls, college campuses, and retirement homes." Participants were also told they would be "given an opportunity to taste the samples more than once to clarify (their) preferences."

After providing consent to participate, participants were given two samples of chocolate, one wrapped in red foil and another in blue foil. Unbeknownst to the participants, both samples were of the same chocolate. Participants were instructed to close their eyes, taste one of the samples, and then rate it from poor (1) to excellent (5) on each of four dimensions of quality used in the chocolate industry—aroma, texture, taste, and aftertaste—and note their ratings on the "Chocolate Quality Survey" form provided. Participants then tasted and rated the second sample. Finally, they indicated which of the two samples—red or blue—they "like best."

Participants rated the chocolate they preferred significantly higher on all four dimensions (taste, aftertaste, texture, and aroma) compared with the chocolate they did not prefer. For the key dimension of taste, their ratings of the preferred piece of chocolate were approximately one standard deviation higher than for the chocolate not preferred, suggesting that participants formed clear, unambiguous preferences.[9]

Round 2. The second round of tasting began with a short description of the growing popularity of high-quality chocolate. The two samples were then distinguished by status with the following information:

> One of the chocolates is made by Godiva. Godiva is an industry leader with a long history of making excellent chocolate. The company was founded in Belgium in 1926 and is sold in all 50 states and throughout the world. Godiva is well known for the quality of its chocolate and has frequently been recognized with awards.
>
> The other chocolate is called Cacao Select by Safeway [a major regional supermarket chain]. This brand of chocolate was started by an expert European chocolate maker who decided to go out on his own to make great chocolate. To save costs and to more easily reach American consumers who do not know good chocolate, he decided to partner with the Safeway supermarket chain. Cacao Select by Safeway has won two blind

tasting contests by chefs when paired against long-established, prestigious brands.

Both chocolates were described as high quality, and both Godiva and Safeway are well-known, highly visible brands in the location where the study took place, but one chocolate bore the name of a high-status manufacturer and the other carried a supermarket label. Yet, participants were given reason to think that the supermarket-label chocolate might be as high quality as the high-status brand. For example, the Safeway brand was described as being made by "an expert European chocolate maker" and having won two awards from chefs. In actuality, both samples came from the same chocolate, so neither was objectively higher quality. Next, participants learned that in the first round they preferred the lower-status, supermarket chocolate. Thus, participants believed that their personal preference (i.e., their first-order belief about quality) was for a lower-status chocolate. We then manipulated the sense of the coordination challenge for their final decision, as we will describe.

First-order inference condition. In the first-order inference condition, participants were told that their ratings would be completely confidential. Thus, there were no actors (other than themselves) whose reactions would be consequential and with whom they must coordinate in making their final choice. Participants then tasted the chocolates a second time, rated them, and selected the chocolate that they personally preferred. We expected a relatively low rate of switching to the high-status option in this condition.

Peer coordination condition. In this condition, participants were told that after their second tasting, they would be invited to participate in a small focus group with people similar to them to discuss strategies for marketing chocolate to college students. They were also told that they would recommend one of the two chocolates to share with the focus group. Participants tasted the two samples a second time. Next, they rated and ranked the chocolates, and then filled out a recommendation card that contained their choice of chocolate for the focus group. We expected participants in this condition to exhibit a low to intermediate level of switching to the high-status chocolate in their final recommendation, because, although they should anticipate their peers reacting to the status of their choice, they would have an opportunity to directly explain their preference to them and thus potentially forge a contrarian code. This condition is not central to a test of our hypothesis, but we included it to further explore the circumstances under which third-order inferences influence choices independent of original quality judgments.

Third-order inference conditions. In the third and fourth experimental conditions, participants were told that they would be recommending a chocolate for placement in a gift basket marketed to one of two different categories of people—a general "all American" market segment or a "young professionals" market segment that might plausibly be more salient to undergraduate participants.[10] After hearing a description of the gift basket, they tasted the chocolates for the second time and rated and ranked them. Next, they learned more about the makers of the gift pack. To increase participants' sense that the consequences of their final recommendations depended on coordinating with the market segments to which the gift packs were intended, they were told that "you have the opportunity to receive a bonus of at least $50 (if) you recommend a chocolate that is ultimately included in the gift pack."[11] These procedures were designed to increase the interdependence of the decision-making context and place participants in contexts in which they had a chance to gain by successfully coordinating with resource-holding others while remaining interactively anonymous in the setting. Finally, participants were told that (1) the gift packs "will feature 15 chocolates of exceptional quality," (2) each will be "displayed in its distinctive brand packaging," and (3) the producers would be "making the

chocolate available for the same cost." These statements were designed to ensure that participants would understand that consumers would be able to evaluate the chocolates only on the basis of their packaging (and accordingly, the status of the brand) and there was no cost-based reason to prefer one or the other. Participants then filled out a recommendation card, selecting the chocolate they would like to include in the gift basket. In actuality, there was no bonus. Participants were told about the chance to win a bonus to increase the salience of the coordination challenge for their final choices, thereby ensuring that the decision being made required third-order inference. We expected the greatest switching to the high-status option in these two conditions.

Study wrap-up. After Round 2, participants in all conditions filled out a survey that contained background information items and manipulation check questions. They were debriefed and paid for their participation.

Dependent Variables

Our main dependent variable is *high-status advantage,* operationalized as whether participants switched their choice from Safeway (low status) in the first round of tasting to Godiva (high status) in the second round. We measure this switch in two ways. The first, which can be applied to all conditions, is based on a switch in the chocolate that the participant "liked best" from round one to round two of tasting. The second, which can be applied only to the peer coordination and the third-order inference conditions, is based on whether participants switched from the chocolate they "liked best" in the first round of tasting to the higher-status chocolate when recommending a chocolate for either the focus group or the gift packs. A switch in the recommendation is most relevant for testing our hypothesis, because what matters for status advantage is that which is enacted in public. However, by definition, there is no "recommendation" in the first-order condition, except insofar as the subject's choice of the chocolate she likes best constitutes a self-recommendation. Thus, to ensure comparability across conditions, we present results based on the switch in "liked best" as well as the recommendation.

Manipulation Checks

To ensure the manipulations of the coordination challenge were salient, participants were asked how much they took into account "what most people would like" and "my own preferences" in making their final recommendations on a scale from 1 (did not influence) to 4 (had a large influence). Confirming the effectiveness of our manipulations, participants rated the influence of "what most people would like" as significantly higher ($t = 2.82$, $p < .01$) in the third-order inference conditions ("young professionals," $m = 2.3$, s.d. = 1.03; "all American," $m = 2.2$, s.d. = 1.15) than in the first-order inference condition ($m = 1.75$, s.d. = .76). Also consistent with expectations, participants in the peer coordination condition reported an intermediate level of influence of "what most would like" ($m = 1.9$, s.d. = 1.07) that was a bit closer to that in the first-order than third-order inference conditions. Perhaps not surprisingly, all participants rated the influence of their own preferences as high (range of means = 3.45 to 3.80), and this rating did not differ by condition (F = .69, n.s.).

Results

Our main prediction is that status advantages will be greater in situations involving third-order inference. That is, when the consequences of individuals' decisions depend on coordinating their decisions with the actions of salient others, they will be more influenced by publicly held status beliefs compared with situations where their decisions are private. We also predict lower rates of status advantage when individuals have the opportunity to coordinate their decisions with specific, socially proximate others, thereby potentially establishing a contrarian code.

Our results are consistent with these predictions. Figure 2 presents the rate of switching from the low- to high-status option across the four experimental conditions. For all

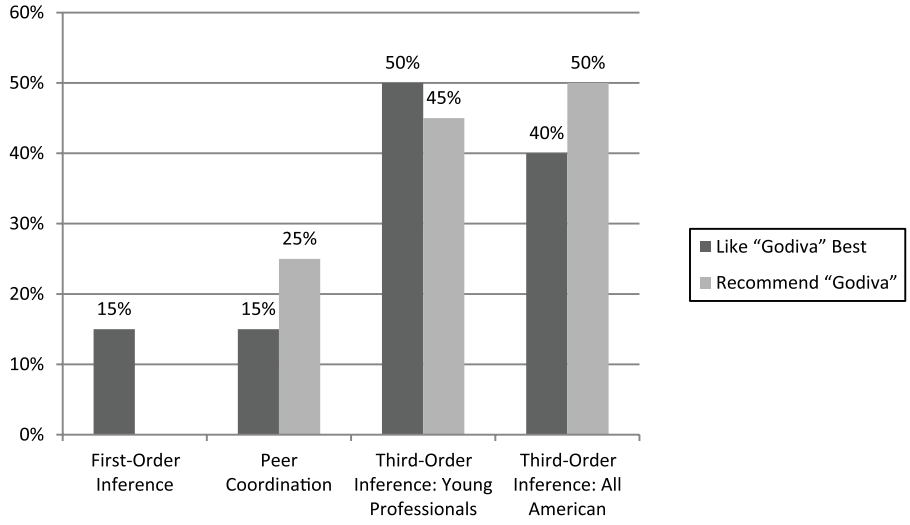

Figure 2. Percent Who Switched Their Selection to Godiva

conditions, the bar on the left indicates the percent of participants who switched which chocolate they "liked best" to the high-status brand, Godiva. The right-hand bar represents the percent who switched to the high-status chocolate (Godiva) when making a recommendation for the focus group or a gift basket (conditions 2, 3, and 4 only).

We compare both the "liked best" and "recommendation" measures for the third-order and peer coordination measures to the "liked best" measure for the first-order condition. By either of the two measures, the rate of switching was roughly three times higher in the two third-order inference conditions than in the first-order inference condition. When measured in terms of the switch in recommendation, 50 percent of participants in the "all American" third-order inference condition switched to the high-status brand, compared with only 15 percent in the first-order inference condition, a statistically significant difference ($z = 2.36$, $p < .01$, one-tailed). Similarly, 45 percent of participants in the "young professionals" third-order inference condition switched to the high-status brand, which is also significantly different from the rate of switching in the first-order inference condition ($z = 2.07$, $p = .02$, one-tailed).

When switching is measured in terms of which chocolate the subject "liked best" (left bars in Figure 2), the difference between "all American" and first-order conditions was smaller (40 versus 15 percent) but still significant in a one-tailed test ($z = 1.77$, $p = .03$). In the "young professionals" condition, 50 percent switched to the higher-status brand, which is significantly higher than the rate of switching in the first-order condition ($z = 2.36$, $p = .009$, one-tailed).

In the peer coordination condition, we expected an intermediate level of switching to the high-status brand, because participants would anticipate being able to interact with their peers, which would allow them to forge a contrarian code in favor of Safeway. As Figure 2 shows, 15 percent of participants switched to liking Godiva, and 25 percent recommended Godiva for their focus group after liking the lower-status chocolate better in the first round. The latter value represents an intermediate amount of switching in our sample, but the differences between the peer coordination condition and the other conditions were not significant.

Study 1 Discussion

Our results validate the two key aspects of our hypothesis. First, as our theoretical model predicts, a status advantage emerges when a decision-maker is induced to treat the decision

as a problem of third-order inference. Second, confirming a distinctive prediction of our theoretical model, this status advantage emerges even though the status information does not reduce the decision-maker's uncertainty about quality. On the one hand, a study of chocolate choices represents a conservative test of our theory; it seems straightforward that people should be least likely to defer to public status beliefs when they can justify their personal view by invoking "de gustibus non est disputadum"—everyone is entitled to their view in matters of taste. On the other hand, this strength of Study 1 is also a limitation. Can third-order inference also be sufficient to induce status advantage when "quality" is publically expected to be judged according to shared criteria as, for instance, in a meritocratic context such as evaluating job candidates for hire?

STUDY 2: SELECTING A POLICE CHIEF

Overview

The second experiment replicates Study 1 and extends our test of the theory to a decision context—the evaluation of job applicants—that is publically expected to be meritocratic and in which reliance on status beliefs, such as those linked to gender, is illegitimate in the contemporary United States. As with Study 1, participants again decide between higher- and lower-status actors under conditions where status information does not reduce uncertainty about quality. However, in Study 2 the high-status actor is male and the low-status actor is female. Like Study 1, participants initially learn that they preferred the lower-status option (the female applicant). However, to switch to the high-status actor in Study 2, participants must not only overcome their initial judgment of who is better, but must also knowingly favor a man over a similarly qualified woman in an ostensibly meritocratic situation, where making judgments based on characteristics like gender is largely viewed as illegitimate.

The study was conducted online via Amazon Mechanical Turk (MTurk) using paid volunteers ($1 each) restricted to those using U.S. IP addresses.[12] In each of two rounds across three conditions, 456 male and female participants rated brief résumés of two candidates for police chief for a U.S. town. Police chief is generally seen as a male-typed job (Uhlmann and Cohen 2005).[13] Evaluating candidates for a gender-typed job makes it more likely that, once participants know the gender of job candidates, gender status beliefs will become implicitly salient in the decision situation (see Wagner and Berger 2002).

In the first round, participants rated candidates without knowing their gender and made an initial recommendation about who would be the better police chief. Participants were then told more about the job qualifications and responsibilities for the job, including several qualifications designed to match stereotypically male characteristics. Next, we manipulated the coordination challenge participants faced in making the "best" decision. In two of the conditions, participants read more about the town that was selecting the police chief. In one condition, this was a "traditional" town in Kansas, evoking a coordinating audience likely to accept gender status beliefs that men are higher-quality police chiefs. In the other condition, it was a "progressive" town in Massachusetts, evoking an audience more likely to reject gender status beliefs as illegitimate. In the third condition, the nature and character of the town were left unspecified. The third condition is not critical to the test of our hypothesis, but we included it to further explore the circumstances under which third-order inferences influence choices independent of original quality judgments. Participants then saw the résumés again, this time with names on the résumés that revealed the gender of their initial recommendation. They saw that they preferred Allison Boyle, a woman, over Matthew Berk, a man. Participants were then asked to make a final recommendation. *Our main prediction for Study 2 is that status advantages (switches from the woman to the man in the final*

recommendation) will be significantly greater in the Kansas condition than in the Massachusetts condition.

Procedures

Preliminary information. Participants were "invited to participate in a research study about how people rate job applicants." They learned that the study was "a project with researchers at the Social Decision Making Center at Stanford who have been contracted to work with a national recruitment firm," and the principal investigators were "interested in understanding how age, education, and other factors influence preferences and decisions." They were further told that, while it may be difficult to assess the candidates with the information they have been provided, research indicates this is an important step employers take in making hiring decisions. Furthermore, they learned it was important to make a good match between applicants and the communities they serve. These instructions were intended to make subjects feel confident in the legitimacy of making a quick assessment of the candidates for a position that, after all, ultimately serves all U.S. residents.

To ensure that they paid attention to details of the study, participants were informed that there would be comprehension questions throughout the study and they would not be compensated if they answered these questions incorrectly. In actuality, we compensated all participants who completed the study. As with Study 1, participants were told they would have an opportunity to revise their assessment. Finally, participants learned they would be evaluating job applicants for a police chief position.

Round 1. Participants were shown two short résumés for the police chief position, identified by applicant number with no indication of gender. Participants first saw one résumé and rated the applicant on a series of seven-point scales ranging from "not at all" to "extremely" capable, warm, skilled, likeable, dominating, physically fit, and ambitious. They then saw the second résumé and rated the applicant using the same scales. The order of the two résumés was counterbalanced. After making their initial ratings, participants saw the two résumés side-by-side and were asked to rate the candidates on four dimensions—education, experience, skills, and attitude—using five-point scales ranging from "poor" to "excellent." Finally, they were asked which of the two applicants—#435 or #278—they would "recommend after (their) initial assessment." This initial recommendation represents their personal judgment of which applicant is higher quality.

Our goal was to control for the quality of applicants by designing two résumés that, while different from each other, were nonetheless judged to be equal quality. A pre-test showed that the résumés were equivalent, but results from the actual study revealed that one of the two résumés (résumé #435) was selected by two-thirds of the participants on the first round. In addition, participants who chose #435 in the first round were significantly less likely to switch their recommendation when they found out the candidate they initially chose was female. We conducted a logistic regression analysis (not shown) where a résumé dummy variable was included as a control and determined that the quality imbalance in the résumés produced no substantive differences in effects of the experimental manipulation.

Round 2. To begin the second round, participants were told that it is "important to understand a little more about what the job entails when trying to match the best applicant to the position." They then reviewed the job description for the police chief position. To strengthen the salience of gender status beliefs for decision-makers, four of the qualifications ("strength"; "(no fear) of physical confrontation"; "must exhibit bravery"; "able to behave rationally in a crisis") were stereotypically male characteristics.

Next, we manipulated the coordination challenge by varying the town that was hiring the

police chief. Participants first read that because police chiefs work closely with the community, it was important there be a good fit between the police chief and the town they serve. Participants were then randomly assigned to one of three conditions where they either (1) read a description of a town in Kansas ($N = 149$), (2) read a description of a town in Massachusetts ($N = 151$), or (3) were provided with no description of the town ($N = 156$). Participants assigned to the Kansas condition read that the town is mid-sized with traditional values, has an emphasis on law and order, and has an established business community. Massachusetts, by contrast, was described as mid-sized with innovative approaches on environmental issues, progressive politics, and an emphasis on citizen involvement.

Participants were next told they would be rating the same two applicants again, now that they had more information about the position. Before doing so, and to create a third-order inference situation, they read, "The recruitment firm we are working with has a strong history of making matches between applicants and the communities in which they will serve. As an additional incentive, you have the opportunity to receive a bonus of $.25 in addition to your payment of $1.00." They were then informed that to qualify for the $.25 bonus, they must recommend the applicant selected by the recruitment firm for this town. The possibility of a bonus was added to increase participants' sense of outcome interdependence with the audience.

Next, participants viewed the same two résumés side-by-side, but this time the résumés had names on them (Allison Boyle and Matthew Berk). The online data base "Behind the Name" (http://www.behindthename.com) lists Allison as an exclusively feminine and Matthew as an exclusively masculine (rather than unisex) name, and it reports Social Security data showing that each is among the 40 most popular names for their gender. "Allison" and "Matthew" have been successfully used to manipulate female/male gender in past studies (Bertrand and Mullainathan 2004; Correll et al 2007). At this point, participants learned that they recommended Allison in the first round. Then, using the same seven attributes from Round 1 ("capable" "warm" "skilled" "likeable" "dominating" "physically fit," "ambitious"), participants rated the applicants again, one at a time on separate screens. They then rated the two applicants side-by-side on four dimensions ("education," "experience," "skills," "attitude") and selected one of the two applicants for their final recommendation.

Study wrap-up. After Round 2, participants in all conditions were asked to rank the criteria they used in making their decision (e.g., education, experience). They also answered a series of manipulation check and demographic questions. Importantly, participants were asked how much they took various factors into account when making their decision, with one factor being fit with the town that was hiring. Participants were then sent to a site where they were debriefed and paid. All participants viewed the debriefing information before receiving compensation.

Dependent Variable

As with Study 1, the dependent variable for this study is *high-status advantage,* operationalized as whether participants switched their recommendation between Rounds 1 and 2 from "Allison Boyle" (low status) to "Matthew Berk" (high status). We predict higher rates of switching for the more traditional Kansas town than for the more progressive Massachusetts town. We do not have a strong prediction for the town-unspecified condition, but rates of switching will likely be lower than in the Kansas condition, because participants had less information about the town upon which to coordinate.

Comprehension and Manipulation Checks

We asked several attention check questions to assess whether participants correctly attended to our manipulations and other details of the study necessary for the test of our theory. The

key attention checks were whether participants could correctly recall (1) that the position being filled was police chief, (2) the gender of the candidate they initially recommended, which was always female by experimental design, and (3) the town hiring the police chief.

After their final recommendation, we asked participants what position was being filled in the study. Only three of the 452 participants who answered this question did not list "police chief," and two of these gave ambiguous rather than incorrect answers.

We also included two checks to assess whether participants recalled that their initial recommendation was a woman. This is crucial because we are measuring whether they knowingly switched their choice to the male candidate in the second round. The first of these checks occurred at the onset of Round 2, when we asked participants, "What is the name and applicant ID number of the applicant you recommended in part 1." We asked about name and ID number to avoid having participants guess our study hypothesis, but we were concerned only with whether they recalled that their initial choice was female. Of the 456 participants, 49 did not correctly recall this information. At the end of Round 2, however, after they chose a résumé ID number for their final recommendation, we again asked the name of this applicant. At this point, only four of the 456 respondents incorrectly identified the gender of the applicant, and seven gave ambiguous answers. Because 98 percent of participants did correctly identify the gender of the applicants they chose by the end of Round 2, to be conservative, we did not exclude any participants based on these gender checks from the analyses we report. However, because awareness of the applicants' gender is central to our hypotheses, we conducted sensitivity analyses of our main results with samples that excluded participants who did not correctly identify gender at the onset of Round 2. The results were robust over these analyses.

To check whether the town manipulation was successful, participants were asked after Round 2 to describe the town that was hiring the police chief. Only 13 of the 300 respondents in the Kansas and Massachusetts conditions did not correctly identify the town. To evaluate the strength of the town manipulation, during the study wrap-up, participants ranked from 1 to 5 the criteria they used to evaluate the applicants (experience, education, skills, attitude, and "match or fit with town"). Lower values indicate higher rankings. "Match with town" was significantly less important in the town-unspecified condition (4.05) than in either the Kansas (3.09; $t = 5.77$, $p < .01$) or Massachusetts (3.02; $t = 6.32$, $p < .01$) condition.

At the end of the study, respondents rated how much influence each of four factors had on their decision, using four-point scales ranging from "no influence" to "a large influence": "My own preferences," "What other people like me would like," "What most people would like," and "What most people in the town that was hiring would like." Only two of these four measures showed statistically significant differences across conditions. First, town preference had significantly more influence for participants in the Kansas condition (mean = 3.11) and the Massachusetts condition (mean = 3.11) than in the town-unspecified condition (mean of 2.11; $t = 7.28$, $p < .01$ for Massachusetts; $t = 7.49$, $p < .01$ for Kansas). Second, "own preferences" had significantly less influence in the Kansas condition than in the other two conditions (2.13 versus 2.52 for the Massachusetts condition [$t = 3.02$, $p < .01$], and 2.58 for the unspecified-town condition [$t = 3.33$, $p < .01$]). This suggests, as we might expect, that participants felt the most pressure to shift away from their initial choice of a woman in the Kansas condition. A small number of participants failed attention checks in the study, but to be conservative, we did not exclude any cases from the analyses presented here.

Results

Table 1 presents the percentage of participants who switched from their initial

Table 1. Proportion Switching to the Male Job Candidate by Subject Gender and Experimental Condition

	Unspecified	Massachusetts	Kansas
All Subjects			
78/456	20/156	20/151	**38/149**
17.1%	12.8%	13.2%	**25.5%**
Men			
35/206	*12/68*	8/70	**15/68**
17.0%	*17.7%*	11.4%	**22.1%**
Women			
41/245	*7/84*	11/80	**23/81**
16.7%	*8.3%*	13.8%	**28.4%**

Note: Italics indicate a difference between men and women at the $p < .05$ level. Bold indicates a difference between the Kansas condition and the Massachusetts condition at the $p < .05$ level.

low-status choice of the female candidate to the high-status choice of the male candidate in Round 2. As the leftmost column of Table 1 shows, people are generally reluctant to switch from a decision they have already made in such an explicitly meritocratic setting: the rate of switching was relatively low overall at 17.1 percent, and men and women evaluators were equally reluctant to switch (17.0 percent of men; 16.7 percent of women).

Our main prediction is that status advantage will be larger in the Kansas condition (see rightmost column of Table 1), where third-order beliefs about male superiority for policing are greatest. Consistent with our predictions, the rate of switching was higher for participants in the Kansas condition (25.5 percent in Kansas, versus 13.2 percent in Massachusetts [$t = 2.71$, $p < .01$] and 12.8 percent in town-unspecified [$t = 2.85$, $p < .01$]).[14] This pattern was particularly strong for female participants. Among men, 22.1 percent switched to the male applicant in the Kansas condition, compared to 11.4 percent of men in the Massachusetts ($t = 1.68$, $p < .05$, one-tailed) and 17.7 percent in the town-unspecified ($t = .64$, n.s.) conditions. For women, 28.4 percent switched in the Kansas condition, compared with 13.8 percent in the Massachusetts condition ($t = 2.30$, $p < .01$) and just 8.3 percent in the town-unspecified condition ($t = 3.44$, $p < .01$). We did not have a strong prediction about the town-unspecified condition, but note that the difference between the town-unspecified and Kansas conditions was particularly large for women. This was due to the fact that, not only were women slightly more likely than men to switch in the Kansas condition (28.4 versus 22.0 percent; $t = .88$, n.s.), the reverse was true in the town-unspecified condition, to the point that women's rate of switching was half that for men (8.3 versus 17.7 percent; $t = 1.73$, $p < .10$). The fact that men switched their choice almost twice as frequently as women when no town information was given suggests an in-group bias: male participants were more likely to switch to the man, and female participants were more likely to stay with their initial choice of the woman. However, when female participants were focused on coordinating with the citizens of the Kansas town, female in-group bias appears to have been eliminated or reversed.

Finally, our theoretical argument suggests that the more participants attended to the town's beliefs, the stronger our results should be. Our results support this prediction. In analyses not shown, we found that the effect of the town manipulation on switching recommendations was moderated by the extent to which participants took "what most people in the town that was hiring would like" into account when making their final recommendation. This effect occurred only for participants in the Kansas condition. In this

condition, the rate of switching was 45.5 percent for the 66 participants who said the town had a "large influence" on their decision, versus 9.6 percent for the 83 participants who said the town had a "moderate influence" or less ($t = 5.42$, $p < .01$). In fact, the difference in switching between Kansas and the other conditions occurred only when participants reported that the influence of the town on their decisions was large.

Study 2 Discussion

Study 2 replicates and extends the results from Study 1. As in Study 1, our results illustrate that status advantage emerges even when status information does not reduce uncertainty about quality, so long as the decision-maker is induced to think about the decision as one involving third-order inference. With decisions involving both chocolate and the hiring of police chiefs, a significant proportion of participants overrode their first-order beliefs about which choice was higher quality, but they did so only when pressured to coordinate with an audience for whom the higher-status option was conventionally believed to be best. The distinctive contribution of Study 2 is to demonstrate that third-order inference occurs even in hiring contexts, where decisions are justified based on shared rather than individual standards, and where it is understood to be illegitimate to base decisions on applicants' status characteristics, such as race and gender, rather than on quality.

However, Studies 1 and 2 cannot distinguish between two slightly different mechanisms that could underlie the third-order inference process. One possibility is that switches to the higher-status choice in third-order inference conditions is a matter of participants trying *to match their choices to the presumed (but unknown) preferences of a particular coordinating audience*. Under this scenario, the prior two studies would suggest how status advantage can be based on pluralistic ignorance, whereby many privately prefer the low-status actor, but they nevertheless favor the high-status actor when coordinating their decisions with others, because each believes that others subscribe to the conventional hierarchy.

But we argue that status advantages can even emerge when decision-makers are aware that coordinating parties subscribe to a contrarian code. In particular, the second possibility is that switches to the higher-status choice in the third-order inference conditions are a matter of participants trying *to match their choices to the presumed beliefs of a coordinating audience about what is conventionally best*—even when these audience members are not expected to personally subscribe to such conventional beliefs. In particular, switchers in Study 1 may have recognized that eventual consumers of the gift packs would focus more on the conventional message sent by the brands included in the gift packs, rather than how the chocolates actually taste. And switchers in Study 2 may have recognized that citizens of the Kansas town would be mindful of the conventional gender hierarchy regardless of whether they thought the woman was the superior candidate for their town. To ascertain whether there is more to coordination via third-order inference than just "giving them what they appear to want," we need to examine a situation in which the socially "best" choice may not simply be to match recipients' taste. Study 3 examines that situation.

STUDY 3: THE IMPRESSING THE HOSTESS EXPERIMENT

The final study is straightforward and targeted. The goal is to test whether status advantage can emerge under a situation of third-order inference even when *neither the decision-maker nor the audience* with whom she is coordinating *is thought to prefer the high-status producer*. In short, we are examining the extent to which it is the "conventional thought"—that which everyone knows is best—rather than the "solicitous thought"—that which the audience is presumed to think is best—that counts. We put the decision-maker in the role of gift-giver, giving a gift of

chocolate in a context where she seeks to make a favorable impression on the recipient of the gift and, in the critical condition, she is also informed by a third party that the target of the gift prefers a low-status chocolate producer. Thus our question is: will the gift-giver opt for the recipient's preference even when it is the lower-status choice?

In this experiment, participants again decide between a high- and a low-status choice where status information does not reduce personal uncertainty about quality. As in Study 1, participants must select between a low-status brand of chocolate (Hershey's Kisses) and a high-status brand (Godiva Truffles). Rather than selecting chocolates for a gift pack (as in Study 1), in Study 3 participants are selecting chocolate as a personal gift to the hostess of a private dinner to which they were invited. Participants were randomly assigned to one of two conditions: a "conventional taste" condition, in which there is credible information that the hostess prefers Godiva Truffles, the high-status brand; and a "contrarian taste" condition, in which the hostess prefers Hershey's Kisses, the low-status brand.[15] This study was again conducted on MTurk using 200 paid volunteers ($.50 each) and restricted to participants using U.S. IP addresses.

Procedures

Participants were informed that the purpose of the study is to "learn how people make choices in different scenarios" and that they will be asked how they would respond to a particular scenario. The scenario tells them they have been invited to "an important private dinner" and they "want to make a good impression" because "the hostess is influential in (their) line of work." They are further informed that they "are going to take a gift of chocolate to present to (their) hostess" and they "are deciding between two options, both of which cost the same amount of money."

Participants next read, "Most people think that Godiva Truffles would make a nice hostess gift." We then manipulated the hostess's chocolate preferences. For participants randomly assigned to the "contrarian taste" condition (N = 104), the scenario continues, "While you do not know the hostess personally, a friend who knows her well says that, although she has sampled a wide variety of chocolates and likes Godiva Truffles a lot, she really likes Hershey's Kisses."[16] For participants assigned to the "conventional taste" condition (N = 96), the text is the same except the friend says the hostess likes Hershey's Kisses a lot but prefers Godiva Truffles. After reading this scenario, participants are asked to choose between two options to give to the hostess at the dinner: (1) "A box of 6 Godiva Truffles, tied with a gold bow" and (2) "A large bag of Hershey's kisses, tied with a gold bow."

After participants recorded their choice of Godiva Truffles or Hershey's Kisses as a hostess gift, they then answered a series of questions designed to elicit the rationale behind their choice of gift and a series of questions rating Godiva Truffles and Hershey's Kisses as a gift. This section also included attention and comprehension check questions. Participants were asked to indicate whether the purpose of the gift was "a personal gift for the hostess," "something to be served at the dinner," or something else (with a text box where they could explain). Only five participants selected the second, incorrect, answer; five other participants selected the third option, but their written characterizations of the dinner were accurate. We do not exclude any cases from the following analyses.

Design Features and Predictions

Note several key features of the experimental scenario and the manipulation. First, the focus is on the hostess as the relevant audience for the gift, which should make the apparent default decision straightforward: the best gift is one that caters to the hostess's preferences, as attested to by "a friend who knows her well." Second, the scenario avoids the possibility that participants would choose the higher-priced item to signal their

solicitousness of the hostess, by saying that the gifts were the same price and (implicitly) that this is because the Hershey's bag was larger. To make this more credible, participants were told that the Hershey's Kisses were packaged in a bag (suggesting greater volume) rather than a box.[17]

At the same time, because the hostess and gift-giver do not know each other, the hostess's taste cannot be *common* knowledge between them. That is, the gift-giver cannot assume that the hostess knows that the gift-giver knows the hostess's taste. Indeed, first- and second-order beliefs (i.e., what a decision-maker or a specific other believes is best) provide a weaker basis for coordination precisely because they are not "what everyone knows everyone knows." Because of this essential feature of the design, in the contrarian taste condition, the gift-giver faces a dilemma in choosing the "best" gift. Is the better choice what the hostess actually prefers (Hershey's) or the high-status option (Godiva) that conventionally, and therefore reliably, signals a quality gift? This leads to our prediction for Study 3. *We assume that the majority of participants in both conditions will select the chocolate that reflects the hostess's known preference, but we expect the percentage of participants who select her preferred chocolate will differ across conditions. Specifically, we expect that a lower percentage of participants in the contrarian condition will select the hostess's preferred option than in the conventional condition.*

Results

Supporting our prediction, we find a significant difference between the two conditions in the extent to which the chocolate preferred by the hostess was selected. When the hostess's taste is conventional—Godiva Truffles—participants selected this option 94.8 percent of the time. The decision-maker faces no dilemma when she knows that the hostess prefers the high-status brand and so, not surprisingly, almost everyone chose to give her Godiva. But when the hostess's taste is contrarian—Hershey's Kisses—participants selected this option only 56.0 percent of the time ($t = 9.12$, $p < .001$). That is, almost half of the participants opted for the high-status brand even though they believed this was not the hostess's preferred chocolate.

The dilemma participants in the contrarian taste condition faced is reflected in their comments. Immediately after selection, they responded to a text box to "explain why (their) choice would make the best gift in this situation." A participant who chose Hershey's said, "Because while it might not be the more classy option, it is her favorite chocolate." By contrast, a participant who chose Godiva said, "I feel like Godiva shows that I care more, despite the fact that she may like Kisses more."[18]

Analysis of the questions asked after their initial choice provides insight into the mechanisms underlying this pattern. There was a broad consensus among participants that Godiva Truffles were the more impressive gift. Participants rated Godiva Truffles and Hershey's Kisses as to whether a gift of one or the other would "impress others at the party" along a four-point scale (1 = "Not at all"; 4 = "Very much so"). Godiva was regarded as the more impressive gift ($M_{Godiva} = 2.97$; $M_{Hershey's} = 1.92$; $t = 15.02$, $N = 200$, $p < .01$). Not only was this true in the conventional taste condition where almost all participants selected Godiva Truffles as the gift ($M_{Godiva} = 2.75$; $M_{Hershey's} = 1.89$; $t = 11.11$, $N = 96$, $p < .01$), and in the contrarian taste condition where 56 percent opted for Godiva Truffles ($M_{Godiva} = 3.16$; $M_{Hershey's} = 1.95$; $t = 10.84$, $N = 104$, $p < .01$), but it was even true for participants in this condition who chose Hershey's Kisses as the hostess gift ($M_{Godiva} = 3.28$; $M_{Hershey's} = 2.16$; $t = 7.20$, $N = 58$, $p < .01$).

Yet in this contrarian taste condition, participants were divided over which gift showed great "consideration" or "respect" for the hostess. Those who opted for the conventional taste by choosing Godiva rated it significantly *higher* than Hershey's in showing "consideration for hostess" ($Mean_{Godiva} = 3.70$; $Mean_{Hershey's} = 2.50$; $t = 7.60$, $N = 46$,

$p < .01$) and "respect for the hostess" ($Mean_{Godiva} = 3.5$; $Mean_{Hershey's} = 2.5$; $t = 7.2$, $N = 46$, $p < .01$). Participants who catered to the hostess's taste by choosing Hershey's, on the other hand, rated Godiva significantly *lower* in showing consideration ($Mean_{Godiva} = 2.97$; $Mean_{Hershey's} = 3.79$; $t = -6.34$, $N = 58$, $p < .01$) and respect ($Mean_{Godiva} = 3.18$; $Mean_{Hershey's} = 3.6$; $t = -4.20$, $N = 58$, $p < .01$) for the hostess.

The overall implication is that participants perceived two ways to make a good impression on the hostess: (1) to select the gift that caters to the hostess's taste, and thus impresses the hostess via the consideration and respect signified by catering to her unconventional taste; or (2) to select the gift that is conventionally regarded as the better gift, and thus impresses the hostess via the consideration and respect signified by choosing what is commonly known to be the best gift.

Study 3 Discussion

The evidence in Study 3 indicates that third-order inference can engender status advantage in a way that is decoupled from anyone's personal judgments of quality, whether those of the decision-maker himself or those of others, as perceived by the decision-maker. This study shows that decision-makers will often select a high-status option *even when they know that the audience believes it to be the inferior option*. The reason is that even when contrarian tastes are known to exist, they coordinate poorly (Heath et al. 2006). When one gives the higher-status option, there can be no doubt about the message that the guest is sending—one is taking care to select *the best*. Participants who chose Hershey's Kisses defended their choice as one that showed consideration and respect for the hostess by matching her tastes, even if it was not so impressive. But for many others, even this reasonable path might not make a favorable impression. The more compelling path for many was to select the conventional option and cater to the hostess in this fashion. And if this is the case when one has credible information about the audience's contrarian taste, such favoritism for the high-status option will likely be much stronger in the more usual case where one does not have such information. The path of least resistance is to opt for the conventional choice.

The gift-giving context we studied may be the canonical example of this phenomenon, but variations on our gift-giver's dilemma occur in many other meritocratic decision contexts. For instance, consider the dilemma faced by a student who must write a literature review in an exam. The student has come to appreciate a relatively obscure line of work by a less-known scholar and sees it as far superior to the work by prominent scholars. Moreover, she has it on good authority from more senior students that the professor grading the exam also prefers the obscure work. But in her exam, should she build on the obscure work and ignore the better-known research, which is cited in conventional literature reviews? At the very least, it makes sense to cite both the obscure and the prominent work, thereby conferring a status advantage on the latter, since more attention is thereby given to it than is implied by either the student's or the professor's assessment of its quality. Similarly, consider the dilemma faced by staffing agency personnel when they are asked by a hiring manager to find "out of the box" job candidates. Will such personnel recommend a candidate with brilliant on-the-job skills but only a community college education, or the candidate who has less experience but sports prestigious university credentials? By our account, many recommenders in this situation (although certainly not all) will take the safer path of recommending the conventionally high-status job candidate (exchange partner, policy decision) despite the resource-holder's preferences. Regardless of what the decision-maker in such cases recommends, the key point is that they face a dilemma they would not have faced if they were choosing strictly on the basis of the resource-holder's preferences. Because personal beliefs are often not common knowledge and status beliefs are, the high-status option is often the safest route for coordinating effectively.

GENERAL DISCUSSION: IMPLICATIONS FOR THE STUDY OF STATUS ADVANTAGE

The primary motivation for research on status advantages is the suspicion that they are pervasive in modern, meritocratic institutions, even though these institutions justify their decisions based on objective standards of achievement and quality. We argued that a powerful process by which status advantages emerge in meritocratic settings derives from the *interdependent* nature of the context in which decisions about "quality" are often made. To effectively contribute to valued outcomes in such settings (i.e., to be high quality), a decision-maker must take into account how salient others in the setting will react to her decision. Therefore, a decision-maker is often focused less on whether a decision meets her personal standards of quality (a first-order inference about quality) than on how most others are likely to judge its quality (a third-order inference). To solve this third-order inference problem and effectively coordinate with relevant others, the decision-maker draws on widely shared status beliefs: common knowledge about how "most people" judge quality. Thus, publicly recognized status beliefs act as the conventional, default standard by which decisions about "quality" are to be made in such interdependent, ostensibly meritocratic settings. Importantly, we argue, the need to coordinate induces decision-makers in such settings to rely on status beliefs in deciding quality, even when they do not personally endorse those beliefs. In some contexts, of course, a decision-maker can coordinate directly with others on the basis of unconventional beliefs. But in the many cases when this is difficult, it is beneficial for the decision-maker to engage in third-order inference, and thus opt for the high-status actor, independent of her personal evaluation of that actor. Insofar as status hierarchies are public and thus are what everyone knows that everyone knows, they provide an effective focal point for coordination (Schelling 1960), even when the parties involved do not necessarily endorse the status beliefs (Clark et al. 2006; Ridgeway and Correll 2004).

To be specific, we have considered two related mechanisms by which decision-makers might defer to conventional status beliefs despite not personally endorsing such beliefs: (1) when they believe, possibly erroneously, that coordinating audiences subscribe to the conventional status belief; or (2) when they believe their choice will be assessed in light of the conventional status belief even by audience members who do not personally subscribe to it. The former mechanism creates the potential for status advantage that is sustained by situations of pluralistic ignorance; the latter mechanism engenders status advantage on the basis of a logic of myth and ceremony that no one is presumed to endorse (see Meyer and Rowan 1977; cf. Winship 2004; Zuckerman 2004).[19] What unites these logics is that they are driven more by the need to coordinate effectively with others than a desire to select items of superior quality.

Relationship between SEI and TOI Models

The other major model of status advantage, the socially endogenous inferences (SEI) model (Gould 2002; Lynn et al. 2009; Podolny 1993, 2005), focuses on a decision-maker concerned primarily with a first-order inference about the choice that would best meet personal standards of quality when quality is uncertain. This model has appeal because it seems evident that decision-makers who choose on the basis of personal assessments of quality often do use status beliefs to make inferences about that quality. In a meritocratic context, this too can lead to a decoupling of status from quality. Our model does not contradict the SEI model, but rather suggests an additional mechanism by which status advantages can be produced over a wide range of meritocratic contexts.

The primary scope condition of our third-order inference model of status advantage, and its point of departure from the SEI model,

is that a decision-maker is making a decision under a non-negligible degree of interdependence with consequential others whose reactions will affect the success of the decision. This encompasses a wide range of meritocratic settings, including most organizational decision-making and well-known phenomena like conspicuous consumption. Because the TOI model has different scope conditions there are decision contexts (those with substantial outcome interdependence in which first-order quality is not at issue) in which status advantage is uniquely explained by the third-order inference model, just as some contexts can be uniquely explained by the SEI model (contexts with quality uncertainty but negligible outcome interdependence; see Salganik et al. 2006). Because most decision-making contexts involve some uncertainty about quality and some degree of interdependence, however, status advantage, by our account, is frequently the result of some combination of SEI and TOI processes, as illustrated in Figure 1.

The Challenge of Empirical Validation

At a theoretical level, the intertwining of third-order inference and SEI processes in common decision contexts implies that we cannot account for how status advantages emerge in ostensibly meritocratic organizations without also incorporating the effects of third-order inference, along with SEI processes. But this intertwining of the two status processes means that it is difficult, when using observational data, to isolate which mechanism is at work—the inference of quality from status or the inference of what most people think is best via status. Consequently, our studies relied on experimental methods that allowed us to create conditions that are outside the scope of SEI and where third-order inference is theorized to be uniquely driving status advantage—settings where status does not reduce quality uncertainty and where decision-makers are sensitive to their interdependence with the beneficiary of the decision. Experimental methods are uniquely well-suited for isolating factors that are confounded in daily life, but their use always raises the question of external validity. Would real-world decision-makers behave similarly to those in our experiments?

Recent ongoing work provides evidence of the third-order inference mechanism operating in a range of real-world contexts.[20] In the first paper, Kovács and Sharkey (2016) show, in an analysis of the market for books, that the sales gap between prize-winning books and their short-listed peers is greater in the December holiday season, when individuals are more likely to be purchasing books as a gift for friends and family rather than for themselves. That is, when interdependence increases (i.e., during the gift-giving season of December), books that are higher status but of similar quality (i.e., award-winning books versus their short-listed peers) are more strongly favored by decision-makers.[21]

The second paper, by Smith and Gaughan (2016), examines the stock market reaction to the appointment of a female CEO to publicly traded, U.S.-based companies. Analyzing data from all such appointments from 2000 to 2015, they find that the appointment of a female CEO results in a negative market reaction, but only when the appointment is accompanied by increased media attention. If stock market participants held first-order beliefs that women were lower-quality leaders than men and acted solely on these first-order beliefs, we would expect stock prices to drop once shareholders learned of the female appointment, regardless of media publicity. Instead, the drop in prices occurs only when the appointment is publicized by the media. This is in line with the logic of third-order inference (see also Cook and Glass 2011). With increased publicity about the female CEO, gender status beliefs become more salient as a focal point for anticipating the reactions of other shareholders, such that for any focal actor, the future value of her stock is increasingly dependent on what other shareholders (or potential shareholders) believe about the quality of the new female leader. In

the absence of specific knowledge about the beliefs of others, the rational strategy is to assume that those others hold conventional status beliefs that women are lower-quality leaders (for evidence on gender stereotypes about leaders, see Eagly and Carli 2007) and thus will sell off their stock, dropping stock prices.

Dynamics of Status Advantage

The TOI model has important implications for how status advantages may be sustained or eroded. According to the SEI model, the advantage experienced by a high-status actor is threatened only when the gap between the quality of the higher- and lower-status actors is sufficiently great that it overcomes the weight decision-makers place on an actor's status in determining their estimate of the actor's quality. The need for such a significant shock to status-based judgments of underlying quality seems like a major obstacle to change, and it helps explain the typically observed stability of status hierarchies. And insofar as higher-status actors use the superior resources they attract to develop greater capabilities over time (see DiPrete and Eirich 2006; Merton [1948] 1968, 1995; Ridgeway and Correll 2004; Simcoe and Waguespack 2010), this will make status hierarchies very stable indeed.

The TOI model also implies that status hierarchies will be very stable, and it is consistent with the observations that initial status advantages have self-fulfilling properties. Note, however, that the mechanism by which such hierarchies are undermined is somewhat different and provides a clearer basis for expecting change. In particular, what matters, from the point of view of the TOI model, is not the size of the gap in quality between high- and low-status actors, but the size of the gap that is *publicly* observed, and thus becomes the basis for conventional views. This, too, creates considerable stability. Many decision-makers might have direct experiences that contradict the status hierarchy, but as long as those experiences occur in private while they continue to favor the high-status actor when coordinating in public settings, this will have no bearing on the status hierarchy (cf. Ridgeway 2011). What everyone knows about what everyone knows will not change.

Yet, at the same time, to the extent that such private experiences are widespread, they create the potential for rapid change once the quality-gap does become public. If a public event occurs to suggest that dissent is widespread, each actor may have considerable experience to draw on in formulating a new orientation to the erstwhile high-status actor. Such dynamics are familiar from the downfall of authoritarian regimes, where public events displaying widespread dissent reveal that previously apparent widespread support for the regime had been a "politics of dissimulation" (see Jowitt 1974; Kuran 1995; Wedeen 1999). Thus, our model of status advantage helps explain not only why status hierarchies are so common and stable, but also how they may be subject to change.

APPENDIX

Table A1. Proportion Switching to the Male Job Candidate by Subject Gender and Experimental Condition: Excluding Participants Who Failed First Gender Check

	Unspecified	Massachusetts	Kansas
All Subjects			
53/407	9/134	13/136	**31/137**
13.0%	6.7%	9.6%	**22.6%**
Men			
22/180	*6/56*	6/63	**10/61**
12.2%	*10.7%*	9.5%	**16.4%**
Women			
29/223	*2/75*	6/72	**21/76**
13.0%	*2.7%*	8.3%	**27.6%**

Note: Italics indicate a difference between men and women at the $p < .05$ level. Bold indicates a difference between the Kansas condition and the control condition at the $p < .05$ level.

Authors' Note and Funding

The first three authors are equal contributors. This research was supported by collaborative National Science Foundation Grants SES-0751993 and SES-0751474 to the first three authors.

Acknowledgments

We are thankful for feedback from *ASR* reviewers, as well as participants at the Yale Sociology Department colloquium, the Columbia Business School Organizational Behavior Seminar, the MIT Sloan Marketing Seminar, the Hebrew University Department of Sociology seminar, the IESE Strategy Group seminar, the University of Toronto Rotman Strategy Group seminar, the 2009 Berkeley-Stanford Inequality Symposium, the 2010 NYU conference on Power, Status, and Influence, and especially the MIT Economic Sociology Working Group. Special thanks to Catherine Turco and Julia DiBenigno for suggestions that led to Study 3.

Notes

1. When the success of a decision depends significantly on the evaluative reactions of key others with whom one does not interact directly, one faces a coordination problem, as first analyzed by Schelling (1960). As scholars in this tradition (see especially Chwe 2001; see also Goffman 1959) have demonstrated, actors solve these coordination problems by basing their own behavior on common knowledge that they can presume others share and will use as the basis for their behavior. Status beliefs represent an example of such common knowledge.
2. When we refer to high-status "actors," we use the term broadly to mean high-status groups of people, or producers of high-status products such as Apple computers or Godiva chocolate.
3. We recognize that scholars sometimes use second, third, and so on "order" to refer to beliefs about beliefs about beliefs (Troyer and Younts 1997; Webster and Whitmeyer 1999). Our usage is based on past practice in studies of the formation of status beliefs and coordination (e.g., Ridgeway 2011; Ridgeway et al. 2009).
4. Our assumption of uncertainty about the choice that coordinates best with consequential others, an uncertainty that decision-makers use status beliefs to resolve, distinguishes our model from Whitmeyer's (2007) model of prestige conferral in non-interacting groups. Like ours, Whitmeyer's model assumes that the "effectively best" choice depends not only on first-order judgments of quality, but also on the support of others for various alternatives, but his model addresses contexts where the decision-maker has direct and accurate knowledge about that support.
5. To be sure, insofar as any opportunity forgone is in some sense a loss (relative to a counterfactual), it is difficult to distinguish conceptually between gains and losses. But following the psychological literatures on "loss aversion" (Kahneman and Tversky 1984) and "regulatory focus" (Crowe and Higgins 1997), it is useful to define a gain as something added to an endowment that one would possess without taking any action, whereas a loss is something subtracted from an endowment that one would possess without taking any action.
6. Henceforth, when we refer to status beliefs as illegitimate in a decision context, we mean they are viewed as not a morally right or proper basis for choosing among actors. Status beliefs can retain social validity, in that they are thought to be what "most people think most people believe" in a

decision context even when they are not considered a proper basis for deciding. Social validity is sometimes considered the collective dimension of legitimacy, but research shows validity can vary independently of propriety in different contexts and produce different effects (Zelditch 2006).

7. All experiments were approved by the Institutional Review Board at Stanford University. Experimental materials are available upon request.
8. In the pre-test, 23 participants were asked to taste the two samples of chocolate and describe them in terms of taste, texture, aroma, and aftertaste. The two samples were from the same bar and were the same brand of chocolate used in the actual study. Participants were then asked which sample they preferred and given the option of no preference. All 23 participants expressed a clear preference for one sample or the other. Every participant gave detailed descriptions justifying their preferences, indicating that they did develop a distinct preference.
9. These results are consistent with a long line of research in psychology showing that the placement of objects (or people) into different categories (red or blue foil chocolates) leads to the perception of exaggerated differences between them (see Brewer 2007).
10. To further understand participants' perceptions of "all Americans" and "young professionals," we conducted a small pre-test of 18 participants from the same population that participated in the actual study. Participants were asked, in a free-response question, what they thought "all Americans" or "young professionals" were like. They were then asked how confident they were in their description and asked to rate the groups on the dimensions of respect, intelligence, and capableness. Participants gave largely consistent descriptions and there was no significant difference in how confident they were in their ratings of all Americans and young professionals. They viewed young professionals as more intelligent and more capable than all Americans, but there was no significant difference in how respected they thought the two groups were. They also imagined that both groups would prefer a name-brand chocolate over a lesser-known brand.
11. By contrast, and as indicated by the manipulation checks, there was no need to increase the salience of the other parties in the first-order condition, because there was just the subject herself. Similarly, we expected (and the manipulation check found) that the anticipation of having to make a joint decision made the focus group salient in the peer coordination condition.
12. This subject pool, which has been widely used in recent years, is broadly representative of Americans who are somewhat better educated and technologically aware than the average population (Berinsky, Huber, and Lenz 2012; Buhrmester, Kwang, and Gosling 2011).
13. A job can be "male-typed" if men are overrepresented in the job or if the traits associated with the job are similar to the stereotypes of traits associated with men/masculinity. "Male-typed" does not mean that being male is a criterion for the job. Indeed, to be meritocratic is to make decisions due to achievement-based criteria rather than aspects of an actor's identity, such as gender.
14. Recall that the values in Table 1 include data from all participants, including those who failed the gender attention check. We conducted a sensitivity analysis where we dropped the cases who failed the attention check. The main difference is that the rate of switching to the male candidate, averaged over all conditions, is much higher among the 49 participants who failed this gender check (51.0 percent versus only 13.0 percent among those who did not fail the check; $t = 7.00$, $p < .01$). This suggests that this group did not pay close attention to the first round. However, when they are excluded from Table 1, the pattern of results in regard to our hypothesis is virtually identical. In particular, the rate of switching is significantly higher in the Kansas condition than in the Massachusetts condition (22.6 versus 9.6 percent; $N = 273$; $t = 2.97$, $p < .01$) and this pattern is particularly strong for female participants (27.6 versus 8.3 percent; $N = 148$; $t = 3.11$, $p < .01$). It is directionally clear for male respondents as well, but it is below conventional significance levels (16.4 versus 9.5 percent; $N = 129$; $t = 1.14$, n.s.) (see Appendix Table A1).
15. The subject's beliefs about the hostess's preferences are, by definition, a second-order belief (i.e., what a specific other thinks is best). However, the situation is still a third-order inference situation if participants coordinate based on "what most people think is best" (a third-order belief) and not on what they think the hostess specifically believes. We will test this prediction.
16. An anonymous reviewer suggested that a more conservative test would be to inform participants that the hostess does not like Godiva. This is an interesting suggestion for future research, as it would test the boundary conditions of our theory. Our goal here was simpler: we wanted to test whether status advantage *can* emerge under situations of third-order inference, even when neither the decision-maker nor the audience prefers the high-status option. We concede that this effect might not extend to situations (however rare) where the audience has a known aversion to that option.
17. This raises the possible concern that participants chose Godiva specifically because it was in a box. In fact, when asked to explain their choice of gift when the hostess was known to prefer Hershey's, only one participant who chose Godiva mentioned the box versus the bag as the sole reason for selecting Godiva. The packaging was not mentioned in 42 of the 46 other cases in which a participant chose

Godiva, and in the other three cases it was described as a minor factor.

18. An anonymous reviewer suggested that participants might be thinking that the host would be "unsettled to see that a guest knew her private tastes" and so, to "maintain proper distance," the decision-maker selects Godiva. If so, this provides another sense in which personal beliefs about quality coordinate poorly: not only are the host's preferences not common knowledge, neither are her preferences for whether others know her private tastes. As a result, decision-makers defer to a safer choice by relying on third-order beliefs when making a decision.
19. Note that whereas Kuran (1995) sees authoritarian regimes as situations of pluralistic ignorance in which the majority of people privately dissent but the veneer of public conformity fools people into overestimating the extent of public conformity, Wedeen (1999) and others (e.g., Elster 1996) contend that private dissent is more of an open secret, whereby most people assume that public conformity is fake.
20. These papers are "real world" examples, meaning they are based on observational, not experimental, data. As a result, they, like all studies relying on observational data, are subject to the criticism that some omitted independent variable is producing the key result, rather than the theoretical variables offered by the authors. Omitted variable bias is always a possibility (i.e., no study can control for all potential independent variables), but these studies do use an impressive array of controls. Our point in describing these studies here, though, is to illustrate that there is some evidence that real-world decision-makers behave in a way implied by the third-order inference model.
21. The authors use a matching procedure to ensure that the award-winning and short-listed books were similar on relevant dimensions, such as sales volume prior to the announcement of the award, price, expert ratings, and so on.

References

Adatto, Kiku, and Stephen Cole. 1981. "The Functions of Classical Theory in Sociological Research: The Case of Max Weber." *Knowledge and Society: Studies in the Sociology of Culture. Past and Present,* Vol. 3, pp. 137–62. Greenwich, CT: JAI Press.

Adut, Ari. 2008. *On Scandal: Moral Disturbances in Society, Politics, and Art*. New York: Cambridge University Press.

Albarracin, Dolores, and Patrick Vargas. 2010. "Attitudes and Persuasion." Pp. 394–427 in *Handbook of Social Psychology,* 5th ed., Vol. 1, edited by S. T. Fiske, D. T. Gillbert, and G. Lindzey. New York: Wiley & Sons.

Azoulay, Pierre, Toby Stuart, and Yanbo Wang. 2014. "Matthew: Effect or Fable?" *Management Science* 60(1):92–109.

Benjamin, Beth A., and Joel M. Podolny. 1999. "Status, Quality, and Social Order in the California Wine Industry." *Administrative Science Quarterly* 44(3):563–89.

Berger, Joseph, Hamit Fisek, Robert Norman, and Morris Zelditch. 1977. *Status Characteristics and Social Interaction*. New York: Elsevier.

Berinsky, Adam J., Gregory A. Huber, and Gabriel S. Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20(3):351–68.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94(4):991–1013.

Bourdieu, Pierre. 1984. *Distinction: A Social Critique of the Judgment of Taste*. Translated by R. Nice. London, UK: Routledge.

Brewer, Marilynn B. 2007. "The Social Psychology of Intergroup Relations: Social Categorization, Ingroup Bias, and Outgroup Prejudice." Pp. 695–715 in *Social Psychology: Handbook of Basic Principles,* 2nd ed., edited by A. W. Kruglanski and E. T. Higgins. New York: Guilford Press.

Buhrmester, Michael, Tracy Kwang, and Samuel D. Gosling. 2011. "Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?" *Perspectives on Psychological Science* 6(1):3–5.

Caplow, Theodore. 1984. "Rule Enforcement without Visible Means: Christmas Gift Giving in Middletown." *American Journal of Sociology* 89(6):1306–1323.

Castilla, Emilio J. 2008. "Gender, Race, and Meritocracy in Organizational Careers." *American Journal of Sociology* 113(6):1479–1526.

Castilla, Emilio J., and Stephen Benard. 2010. "The Paradox of Meritocracy in Organizations." *Administrative Science Quarterly* 55(4):543–76.

Centola, Damon, Robb Willer, and Michael Macy. 2005. "The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms." *American Journal of Sociology* 110(4):1009–40.

Chwe, Michael Suk-Young. 2001. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton, NJ: Princeton University Press.

Clark, Anne E., and Yoshihisa Kashima. 2007. "Stereotypes Help People Connect with Others in the Community: A Situated Functional Analysis of the Stereotype Consistency Bias in Communication." *Journal of Personality and Social Psychology* 93(6):1028–39.

Clark, C. Robert, Samuel Clark, and Mattias K. Polborn. 2006. "Coordination and Status Influence." *Rationality and Society* 18(3):367–91.

Cook, Alison, and Christy Glass. 2011. "Leadership Change and Shareholder Value: How Markets React to the Appointments of Women." *Human Resource Management* 50(4):501–519.

Correll, Shelley J., and Stephen Benard. 2006. "Biased Estimators? Comparing Status and Statistical Theories of Gender Discrimination." *Advances in Group Processes* 23:89–116.

Correll, Shelley J., Stephen Benard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112(5):1297–1338.

Crowe, Ellen, and E. Tory Higgins. 1997. "Regulatory Focus and Strategic Inclinations: Promotion and Prevention Focus in Decision-Making." *Organizational Behavior and Human Decision Processes* 69(2):117–32.

DiPrete, Thomas A., and Gregory M. Eirich. 2006. "Cumulative Advantage as a Mechanism for Inequality: A Review of Theoretical and Empirical Developments." *Annual Review of Sociology* 32:271–97.

Eagly, Alice H., and Linda L. Carli. 2007. *Through the Labyrinth: The Truth about How Women Become Leaders*. Cambridge, MA: Harvard Business School Press.

Elias, Norbert. 1983. *The Court Society*. Translated by E. Jephcott. Oxford, UK: Basil Blackwell.

Elster, Jon. 1996. "Review of Kuran, *Private Truths, Public Lies*." *Acta Sociologica* 39(1):112–15.

Espeland, Wendy Nelson, and Michael Sauder. 2016. *Engines of Anxiety: Academic Rankings, Reputation, and Accountability*. New York: Russell Sage Foundation.

Foschi, Martha. 2000. "Double Standards for Competence: Theory and Research." *Annual Review of Sociology* 26:21–42.

Gillespie, Patrick. 2015. "Is PBR Still Cool? America's Hipster Beer Is Slowing Down." *CNN* June 30 (http://money.cnn.com/2015/06/30/investing/pbr-beer-sales-growth/).

Goffman, Erving. 1951. "Symbols of Class Status." *British Journal of Sociology* 2(4):294–304.

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life*. New York: Doubleday.

Goode, William J. 1978. *The Celebration of Heroes: Prestige as a Control System*. Berkeley: University of California Press.

Gould, Roger V. 2002. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." *American Journal of Sociology* 107(5):1143–78.

Heath, Chip, Ben Ho, and Jonah Berger. 2006. "Focal Points in Coordinated Divergence." *Journal of Economic Psychology* 27(5):635–47.

Homans, George C. 1950. *The Human Group*. New York: Harcourt, Brace.

Jensen, Michael. 2006. "Should We Stay or Should We Go? Accountability, Status Anxiety, and Client Defections." *Administrative Science Quarterly* 51(1):97–128.

Jowitt, Kenneth. 1974. "An Organizational Approach to the Study of Political Culture in Marxist-Leninist Systems." *American Political Science Review* 68(3):1171–91.

Kahneman, Daniel, and Amos Tversky. 1984. "Choices, Frames, and Values." *American Psychologist* 39(4): 341–50.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. New York: Harcourt, Brace.

Kim, Jerry W., and Brayden G. King. 2014. "Seeing Stars: Matthew Effects and Status Bias in Major League Baseball Umpiring." *Management Science* 60(11):2619–44.

Kovács, Balázs, and Amanda J. Sharkey. 2016. "The Many Gifts of Status." Unpublished manuscript, The University of Chicago Booth Graduate School of Business, Chicago, IL.

Kuran, Timur. 1995. *Private Truths, Public lies: The Social Consequences of Preference Falsification*. Cambridge, MA: Harvard University Press.

Lynn, Freda B., Joel M. Podolny, and Lin Tao. 2009. "A Sociological (De)construction of the Relationship between Status and Quality." *American Journal of Sociology* 115(3):755–804.

Lynn, Freda B., Brent Simpson, Mark H. Walker, and Colin Peterson. 2016. "Why Is the Pack Persuasive? The Effect of Choice Status on Perceptions of Quality." *Sociological Science* 3:239–63.

Malter, Daniel. 2014. "On the Causality and Cause of Returns to Organizational Status: Evidence from the Grands Crus Classés of the Médoc." *Administrative Science Quarterly* 59(2):271–300.

Merton, Robert K. [1948] 1968. "The Self-Fulfilling Prophecy." Pp. 475–90 in *Social Theory and Social Structure,* enlarged ed. New York: The Free Press.

Merton, Robert K. 1995. "The Thomas Theorem and the Matthew Effect." *Social Forces* 74(2):379–422.

Meyer, John W., and Brian Rowan. 1977. "Institutionalized Organizations: Formal Organizations as Myth and Ceremony." *American Journal of Sociology* 83(2):340–63.

Parsons, Talcott. 1951. *The Social System*. Glencoe, IL: The Free Press.

Petersen, Trond, and Ishak Saporta. 2004. "The Opportunity Structure for Discrimination." *American Journal of Sociology* 109(4):852–901.

Podolny, Joel M. 1993. "A Status-Based Model of Market Competition." *American Journal of Sociology* 98(4):829–72.

Podolny, Joel M. 2005. *Status Signals*. Princeton, NJ: Princeton University Press.

Ridgeway, Cecilia L. 2011. *Framed by Gender: How Gender Inequality Persists in the Modern World*. New York: Oxford University Press.

Ridgeway, Cecilia L., Kristen Backor, Yan E. Li, Justine E. Tinkler, and Kristan G. Erickson. 2009. "How Easily Does a Social Difference become a Status Distinction? Gender Matters." *American Sociological Review* 74(1):44–62.

Ridgeway, Cecilia, and Shelley J. Correll. 2004. "Unpacking the Gender System: A Theoretical Perspective on Cultural Beliefs and Social Relations." *Gender and Society* 18(4):510–31.

Ridgeway, Cecilia L., and Shelley J. Correll. 2006. "Consensus and the Creation of Status Beliefs." *Social Forces* 85(1):431–53.

Rivera, Lauren A. 2015. *Pedigree: How Elite Students Get Elite Jobs*. Princeton, NJ: Princeton University Press.

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science* 311(5762):854–56.

Salganik, Matthew J., and Duncan J. Watts. 2008. "Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market." *Social Psychology Quarterly* 71(4):338–55.

Sauder, Michael. 2006. "Third Parties and Status Position: How the Characteristics of Status Systems Matter." *Theory and Society* 35(3):299–321.

Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Simcoe, Timothy S., and David M. Waguespack. 2010. "Status, Quality, and Attention: What's in a (Missing) Name?" *Management Science* 57(2):274–90.

Smith, Edward Bishop, and Kevin Gaughan. 2016. "Better in the Shadows? Public Attention, Media Coverage, and Market Reactions to Female CEO Announcements." Presented at the European Group for Organizational Studies, July 7–9, 2016, Naples, Italy.

Stuart, Toby E., Ha Hoang, and Ralph C. Hybels. 1999. "Interorganizational Endorsements and the Performance of Entrepreneurial Ventures." *Administrative Science Quarterly* 44(2):315–49.

Swidler, Ann. 2001. *Talk of Love: How Culture Matters*. Chicago: University of Chicago Press.

Troyer, Lisa, and C. Wesley Younts. 1997. "Whose Expectations Matter? The Relative Power of First-Order and Second-Order Expectations in Determining Social Influence." *American Journal of Sociology* 103(3):692–732.

Turco, Catherine J. 2010. "Cultural Foundations of Tokenism: Evidence from the Leveraged Buyout Industry." *American Sociological Review* 75(6):894–913.

Uhlmann, Eric Luis, and Geoffrey L. Cohen. 2005. "Constructed Criteria: Redefining Merit to Justify Discrimination." *Psychological Science* 16(6):474–80.

Uzzi, Brian, and Ryon Lancaster. 2004. "Embeddedness and Price Formation in the Law Market." *American Sociological Review* 69(3):319–44.

Veblen, Thorstein. 1953. *The Theory of the Leisure Class*. New York: New American Library.

Wagner, David G., and Joseph Berger. 2002. "The Evolution of Expectation States Theories." Pp. 41–78 in *Contemporary Sociological Theories,* edited by M. Zelditch Jr. and J. Berger. New York: Rowman & Littlefield.

Weber, Max. 1968. *Economy and Society*. Edited by G. Roth and C. Wittich and translated by E. Frischoff et al. New York: Bedminster Press.

Webster, Murray, and Joseph M. Whitmeyer. 1999. "A Theory of Second-Order Expectations and Behavior." *Social Psychology Quarterly* 62(1):17–31.

Wedeen, Lisa. 1999. *Ambiguities of Domination: Politics, Rhetoric, and Symbols in Contemporary Syria*. Chicago: University of Chicago Press.

Whitmeyer, Joseph M. 2007. "Prestige from the Provision of Collective Goods." *Social Forces* 85(4):1765–86.

Willer, Robb, Ko Kuwabara, and Michael W. Macy. 2009. "The False Enforcement of Unpopular Norms." *American Journal of Sociology* 115(2):451–90.

Winship, Christopher. 2004. "Veneers and Underlayments: Critical Moments and Situational Redefinition." *Negotiation Journal* 20(2):297–309.

Yamagishi, Toshio, and Toko Kiyonari. 2000. "The Group as the Container of Generalized Reciprocity." *Social Psychology Quarterly* 63(2):116–32.

Zelditch, Morris. 2006. "Legitimacy Theory." Pp. 324–52 in *Contemporary Social Psychological Theories,* edited by P. Burke and J. Stets. Stanford, CA: Stanford University Press.

Zuckerman, Ezra W. 2004. "Towards the Social Reconstruction of an Interdisciplinary Turf War: Comment on Zajac and Westphal, *ASR,* June 2004." *American Sociological Review* 69(3):458–65.

Zuckerman, Ezra W. 2010. "Why Social Networks Are Overrated: Downsides of the Commensuration That Underlies Social Network Analysis." *Perspectives: Newsletter of the ASA Theory Section* 32:3–5, 15.

Zuckerman, Ezra W. 2012. "Construction, Concentration, and (Dis)Continuities in Social Valuations." *Annual Review of Sociology* 38:223–45.

Shelley J. Correll is Professor of Sociology and, by courtesy, Organizational Behavior at Stanford University where she also serves as the Barbara D. Finberg Director of the Clayman Institute for Gender Research. Her current research examines how legal mandates and organizational policies and practices can reduce the effect of status beliefs on workplace gender inequalities and how gender status beliefs produce inequalities in product markets.

Cecilia L. Ridgeway is Lucie Stern Professor of Social Sciences in the Department of Sociology at Stanford University. Her research addresses the role that social status in everyday interaction plays in stratification and inequality, especially in regard to gender. A recent book is *Framed by Gender: How Gender Inequality Persists in the Modern World* (Oxford 2011). Current projects include experimental studies of "is deference the price of being seen as reasonable" for low status members and a book project on the nature and importance of status as a form of social inequality.

Ezra W. Zuckerman Sivan is the Alvin J. Siteman Professor of Entrepreneurship and Strategy at MIT Sloan School of Management. He is also co-founder of Sloan's Economic Sociology PhD Program and is currently serving as deputy dean, with responsibility for faculty affairs. The current paper relates to his broader research interests on social valuation. Key questions include how and when social valuations depart from objective conditions (see his 2012 essay in the *Annual Review of Sociology*) and

how status hierarchies relate to hierarchies of moral worth (see his 2014 article in the *American Journal of Sociology* with Oliver Hahl).

Sharon Jank is a PhD candidate in the Department of Sociology at Stanford University specializing in social psychology, implicit bias, and diversity analytics within organizations. She is currently researching how replacing traditional résumé screening with blind auditions affects the hiring process and outcomes of underrepresented applicants.

Sara Jordan-Bloch is a sociologist at the Clayman Institute for Gender Research at Stanford University. Her research centers on the intersection of social psychology and education. Current projects include work on perceptions of mattering as a mechanism in teacher-student relationships, the relational context of cohort-based learning, and how gender inequality is experienced and understood during critical transitions for young women.

Sandra Nakagawa is a PhD candidate in the Department of Sociology at Stanford University. Her research focuses on gender differences in health and environmental behaviors and how concerns about masculinity contribute to men's lower rates of many healthy and sustainable practices.