

## ORIGINAL ARTICLE

## The Epistemology of Popularity and Incentives

Dan Moller

University of Maryland, College Park

This paper discusses two epistemic principles that are important to buyers and sellers: the appeal to popularity and the appeal to incentive structures. I point out the various ways these principles are defeasible, and then offer some examples of them at work in the contexts of hiring, politics and the arts. Finally, I consider why these principles are generally neglected, and conclude that our neglect is unwarranted on both epistemic and moral grounds.

**Keywords** economics; epistemology; incentives; markets; popularity

DOI:10.1002/tht3.83

Social epistemologists increasingly discuss evidentiary principles in fields like law and politics, but there has been less discussion of the evidence relevant to markets of various sorts.<sup>1</sup> In this paper I describe two of the most common such principles—those pertaining to popularity and incentives—and why people often seem reluctant to employ them. Because of this reluctance, my goal here is to highlight the relevance and importance of these principles, not to offer a fully fleshed out account of how we ought to deploy them. The main point is that evidence from popularity and incentives counts in favor of shifting our beliefs, and that we tend to underutilize such evidence.

I

Assume that other things equal the popularity of a seller (e.g., a restaurant or doctor) is evidence that service will be good, and vice versa. Assume, moreover, that other things equal sellers respond to incentives, and that by extension the fact that a seller faces strong incentives to focus on the service relevant to the consumer is evidence that service will be good, and vice versa. Both kinds of evidence are of course highly defeasible, since many sellers are popular for reasons irrelevant to the buyer, and the seller may respond to irrelevant incentives that drown out those the consumer cares about. But conversely, the force of these signals can be amplified when the seller is popular *despite* providing reasons *not* to buy, or when the seller has little or no incentive to respond to any incentives but those the buyer cares about. Duly cautioned, rational consumers will reason that they generally have grounds to patronize a seller (a) when he is popular, (b) when his popularity isn't due to irrelevant factors, and (c) he is popular despite

*Correspondence to:* E-mail: dmoller@umd.edu

countervailing factors (“popularity” for short); and when (d) he has an incentive to provide good service, when (e) those incentives trump other kinds of incentives, and (f) he has few incentives besides those relevant to the buyer (“incentives” for short).

This advice about popularity and incentives may seem abstract, but it captures the guidance often given by such homely social epistemologists as food critics. Assuming we’re interested just in the quality of the food at a given price, they tell us other things equal to patronize restaurants right after opening since they will be eager to establish a reputation for good food, whereas once they actually have such a reputation, their incentives for good food diminish since their sales are in part a function of that (sticky) reputation. They also suggest avoiding restaurants in scenic locations or full of good-looking people, which reduces the incentive of restaurateurs to rely solely on quality of food and weakens the evidence due to popularity (since they will then be popular in part for food-irrelevant reasons). Instead, foodies should eat at places in obscure, ugly locations or featuring lousy atmosphere, say aggressive and alienating religious art. Once again the incentives work in the buyer’s favor, and against-the-tide popularity is an especially strong signal in this context.<sup>2</sup> Here, fashionable cynicism is supported by sound epistemology that draws on the two pieces of evidence that are among the most commonly available to discerning buyers. It’s hard to know if the cook was recently replaced or what experts say about some establishment while drifting through Barcelona; it’s easy to observe that the restaurant is on the main-square and full of attractive hipsters (portending culinary disaster).

## II

The restaurant case is a light-hearted illustration, but an important question is why this kind of information is not more widely utilized by rational consumers of services like medicine, art, and politics. Perhaps the answer is that it’s difficult to say much about how powerful this sort of evidence really is in absolute terms. But we can often bypass this worry by focusing on matched-pairs, where the potential sellers are about on par, or where we simply don’t know anything else besides popularity and incentives. Here are some examples:

**Hiring:** suppose we’re looking to hire an accountant, doctor, executive or professor strictly for their professional competence and that we face two candidates roughly on par. Much of our evidence reflects a kind of popularity: recommendations, previous hirings and many accomplishments ultimately reflect the fact that people thought highly of the candidate. How should this popularity influence us? Empirical evidence indicates that irrelevant personal attributes like beauty or warmth influence this kind of popularity to a startling degree, and accordingly we should *discount* for it.<sup>3</sup> Conversely, we should favor those swimming upstream against such attributes (point (c) above). Faced with otherwise similar candidates, then, it’s powerful evidence *against* someone’s candidacy that they are good-looking, tall, confident, warm, extroverted or well-connected, just as foodies should eat at the

same-priced, equally popular restaurant in the rotten part of town rather than the scenic square. In roughly pair-matched hiring situations it's rational to pick the stout introvert over his or her more likable rival.<sup>4</sup> After all, the rival is likely to have prospered in part due to these irrelevant factors, and the introvert succeeded despite swimming upstream, meaning their work must have been especially strong. And incentive effects suggest better work in the future from the latter, who will be less able to draw on the advantages of popularity, and will thus be forced to produce better work (point (f) above). Of course, popularity and incentives are only two signals among many, but they can be important ones about the quality of service we can expect from those we hire.

**The arts:** suppose you know of two writers, and only have time to read one of their novels. X is a well-connected schmoozer publishing his third novel; Y is a reclusive first-novelist. Both popularity (in the form of garnering the attention of a publisher and any praise by third-parties) and incentive effects suggest reading Y. The same goes for movie sequels or spin-offs that can draw on name-recognition instead of quality, or recordings that feature suspiciously attractive performers in suggestive poses on the cover (when one should hope for a rotund sexagenarian with a large wart).

**Voting:** voting is complicated because it's harder to sort out what counts as relevant; if being tall and attractive helps a politician pass legislation because of people's poor epistemic skills, then that is paradoxically a genuine reason for the voter to count such attributes as relevant. (It's reasonable to hire a good-looking litigator to persuade the unreasonable jury.) However, often we face two candidates, one of whom is tall and charismatic and the other is, say, corpulent and inarticulate, yet equally successful thus far (perhaps the two are successful governors). Other things equal, here we have evidence that the latter candidate possesses some kind of unusual talent and has the incentive to rely on it. This is only one signal among many for a voter to consider, but it's at the very least grounds for deep reflection on how the candidates became successful, and what kind of difference the talent at issue might make in office. One should expect complex, mixed feelings in reasonable voters about the charisma and height of the one candidate, and in specific cases that should count against their vote.

In these examples I have emphasized not just plumping for the popular worker, novelist or politician. Instead, following the restaurant critics' suggestion, we should look for evidence that the seller's popularity and incentive-structure are *pure*, in the sense that they are as closely aligned with our interests as buyers as possible. And to that end, we can leverage the fact that if we and others care about trait T, and there exists a bias against some other trait B, then in a pool of candidates that exhibit both traits to varying degrees, the fact that X is popular despite exhibiting B to a high degree is strong evidence

that X is popular for the right reasons, and that his incentives will be aligned with our interests, since otherwise we would expect X to have been less popular.

It's useful to compare ethnic discrimination. There is evidence that Asian students in the United States require significantly higher qualifications than others in order to be competitive at elite universities. Research suggests that Asians require a Scholastic Aptitude Test score of 1540 (out of 1600 at the time studied) in order to have the same chance of admission as a white student with a score of 1400, holding all other quantifiable qualifications constant.<sup>5</sup> Natural experiments with states that do or do not require race-blind admissions similarly display large disparities in Asian admissions, as do secular trends in the composition of the applicant pool (the Asian share rises, but admission levels plateau). This isn't conclusive evidence of bias, since Asians might be worse qualified in unquantifiable respects such as references or extracurriculars, but suppose for argument's sake there does turn out to be an Asian penalty or *de facto* quota. If you have before you Asian and non-Asian freshmen, how should you estimate their relative academic qualifications? Obviously the presumed bias against Asians means that they will tend to be more qualified than others; in fact, determining how much more group X is qualified than others in some selected-for population is a good way of *measuring* bias for or against the group. (We could thus compare bias levels against Asians and Jews or African-Americans in the present or past by noting by what proportion their qualifications had to exceed those of others in order for them to advance.) The popularity principle, at least in respect to (b) and (c), is just the same point in a different guise. Of course, publishers don't think of themselves as "discriminating" against sullen first-time authors who must produce superior work in order to succeed, but in effect they are, and that tells us that such authors who nonetheless clear the same hurdles as previously published schmoozers are likely to be special in a way the discerning book-buyer cares about.

We might wonder whether this point about popularity really is distinct from the one about incentives, or whether we aren't in fact double-counting. Are there really two separate reasons here for buyers to consider, one that the seller's popularity suggests quality service, and another that his or incentive-structure suggests the same? Double-counting would be a problem if popularity was merely a reflection of the seller's incentives, but it isn't. Popularity tells us something about how other *buyers* feel about the service, which is useful information for us to the extent that our interests align with theirs, whereas incentive-structures tell us something about the *seller*. The two are correlative in so far as a seller will often be popular because of how he responds to his incentives, but that doesn't mean that the two pieces of evidence aren't distinct. To see that they are, note that in unusual cases they will come apart, as when a seller is very popular because he acts contrary to his incentives (e.g., he foolishly sells too cheaply), or when he is unpopular despite acting on them. The situation is analogous to that of establishing motive and opportunity for a crime as well as providing evidence like fingerprints. The fingerprints may reflect, in a sense, the motive and opportunity of the criminal, but that doesn't impugn the independent evidentiary value of both. Jurors should raise their probability on the defendant's guilt upon hearing his fingerprints were at the scene of the crime; hearing that he had motive and opportunity they should do so again, since the conditional probability

of guilt is higher on that evidence plus the fingerprints than on the finger-print evidence alone.

### III

Now, why is it unpopular to invoke such evidence, even at the margins, even as one signal among many? Why is it rare to hear anyone talk of rejecting a candidate in part because they are warm and handsome? Faced with a pick of accountants at a firm, *sound epistemology overwhelmingly suggests barreling past attractive, polite workers and urgently seeking out the ugliest, shortest, most boorish one available*, yet this strategy is rarely even considered.

There are several objections to this kind of reasoning that might support ignoring it. To begin with, perhaps the parameters I've set seem unrealistic. How often is it irrelevant to us whether someone we hire is fun to be around? Being fun or attractive can be useful for things like teamwork and collaboration, which are in turn important dimensions of professional life. There may be some scope for this kind of consideration, but to see how much weight it really deserves, consider making things explicit: "I propose the department reject professor X's candidacy because he isn't as warm and fun as professor Y" or "I think you should choose Dr. Schmidt as your cardiologist because he's taller than Dr. Geary" or "I want someone good-looking to take care of my retirement plan." These preferences look embarrassing once we spell them out. There may be the occasional grey area generated by our concern for positive social interaction, but by and large this explicitness test doesn't seem to support dismissing popularity and incentives as signals.

Another objection is that the practical application of these principles will turn out to be at best murky. Two difficulties present themselves. On the one hand, this may all seem to depend on getting hold of empirical information that we can at best have some hunches about. Earlier, I noted the existence of at least some such information, but it's difficult to imagine knowing exactly how much good looks help political candidates, or how much less motivated movie studios are to find a good script for a sequel. If we can't get hold of empirical information of this sort, are we simply to guess? And even if we do get some information in the form, say, of wage premiums, it's hardly obvious how to incorporate that information into a function for updating our beliefs. More subtly, how can we know that we've accounted for all such influences on the popularity and incentive-structure of a buyer? And if we factor in the information we have while ignoring what we don't have, we seem to risk worsening our epistemic position rather than improving it. On the other hand, when we *are* confident about how popularity and incentives factor-in, we're likely to see a lot of conflicts, since the evidence rarely stacks up neatly in favor of one view, leaving us back where we started. One job candidate will be charming, another handsome, one governor has a disfavored ethnic background, the other comes from a disadvantaged social class, and we might wonder whether such conflicts can be resolved in a way that leaves us with any useful guidance at all.

These problems are real, and there are some specific things to be said to them below. But the first thing to stress is that they aren't *special* to popularity and incentives. And

insisting that everyday epistemic principles in general avoid such problems would be optimistic indeed. Consider garden-variety testimonial evidence. If eyewitnesses agree that they saw that X, then that generally counts in favor of believing that X. But of course we are beset with problems in making sense of such evidence. We rarely have access to fine-grained information about how probative various forms of testimony are. What we have is a random hodgepodge—the occasional terrifying experiment on eyewitness fallibility or selective recall, and otherwise just intuitions and life-experience.<sup>6</sup> We recognize that testimony can be more or less reliable under various conditions, but we have only a rough sense of the extent that this is so, despite drawing on testimony constantly for very important decisions. And here, too, it's not obvious how exactly to incorporate such data as is available into a belief-update function, and we cannot be sure that the testimonial evidence that we respond to is representative of all such evidence that is out there. (We can hardly assume, in a trial say, that the testimony we hear is an unbiased random sample of all that was witnessed.) And even if we did have such evidence, in everyday life testimonial evidence often conflicts in various ways that are hard to net-out.<sup>7</sup> But in the case of testimony it's obvious that these are grounds for caution, not for doubting the existence of such evidence, that is, that appropriate testimony really does count in favor of believing certain things, and that consequently we ought to pursue it as best we can, duly mindful of the pitfalls. We muddle through, informally trying to determine how much weight to accord a given piece of testimony, and sorting out conflicts by seeing where the balance lies in net, or whether such conflicts cancel the value of such evidence altogether, and so on. There's no reason to think popularity and incentives are any worse off than testimony in these respects.

Apart from this general response, notice that the specific information relevant to (a)–(f) is hardly obscure or unattainable, even if our knowledge of it remains intuitive and hard to quantify. Determining how popular a job candidate or book or politician is often presents no serious obstacle. Determining, more subtly, the extent to which buyers discriminate against irrelevant features of the seller requires simply that we estimate the extent to which buyers punish such traits as shortness or obscurity or a lack of charisma. (Sometimes it will be more natural to think of this in the equivalent terms of rewarding height or fame or charm.) This, too, seems far from hopeless. Knowing a popular restaurant is down an alley off the square is helpful; knowing that people perceive it as a *rather dangerous* alley strongly amplifies the signal, given our background information about how patrons feel about dangerous alleys. All we really need to know is how much people prize or dislike the relevant traits, and allow such feelings to influence their market-behavior.

Perhaps the hardest information to come by is the extent to which sellers are moved by incentives. Here, we need to know to what extent sellers will attempt to make up for drawbacks to their services by improving quality in other ways. Although it is difficult to have any very precise knowledge about this, all we need is a sense of how much harder, say, a writer will work to produce high quality work as against someone famous following up on several previous bestsellers. This is something we can get a sense for after

reconnoitering the relevant market a bit. To give an example close to home, in academia professors just before and just after tenure face different incentive-structures, and we might have insight into how productive such people respectively are at our institution. (As this example shows, the incentives at stake may be subtle; academics (and others) are motivated by pride, a sense of role-related responsibility, and many other concerns besides money or fear of being fired, making the question a serious and open one.) Overall, I see no reason to doubt that popularity and incentives give us one more kind of reason to add to the epistemic mix, about on par with other kinds of everyday evidence in respect to tractability.

#### IV

So far I have been rejecting various explanations of our foot-dragging when it comes to popularity and incentives. I have two positive alternatives to propose. One is simply that while it's obvious that popularity at a given price is evidence of quality, and that people respond to incentives, the point about *discounting* for irrelevant popularity, and the one about *diminished* incentives, are subtler. It requires slightly more abstraction to see the point about choosing between equally popular restaurants, one of which has attractive clientele or scenery; similarly for the point about the seller's incentive to focus on what the buyer cares about when the seller's sales are impacted by other factors as well.

Second and more interestingly, there is something disquieting about "punishing" a job-candidate for a non-job-related personal attribute, as it may seem to us. Telling someone "you lost the job because you were attractive" intuitively feels not just wrong-headed but perhaps even discriminatory and unethical. Something like this, it seems to me, explains some of the resistance to the social epistemology I have laid out. However, there are two ways in which this reaction is, I think, a mistake. First, it makes it sound as if we do have the right *epistemic* responses, but then refrain from implementing them out of ethical delicacy. As a descriptive matter, that sounds much too optimistic to my own jaded ears. What the examples given above and our reactions to them suggest is rather resistance to admitting the epistemic situation, not to admitting it but virtuously prescinding from its application. Moreover, normatively, the ethical judgment here looks questionable. The usual ground for objecting to employer discrimination is that it rests on prejudice against a group and/or has a deleterious effect on a victimized class. But neither is true here. Correcting for the misperception due to irrelevant popularity isn't a form of invidious discrimination but the removal of a distortion in judgment. And considering incentive effects is no more unfair in the labor market than when insisting on paying a painter by the job instead of by the hour. If anything, it seems unfair to the stout and sullen candidate *not* to frame one's hiring decisions in the way suggested, since he will otherwise be harmed by a competition that is judged in an unreasonable manner. Framing popularity and incentives in terms of "punishing" the tall and attractive is entirely optional; we can equivalently describe the process as "pro-actively boosting the prospects of the less attractive." The "punishing" perspective assumes that the current

status quo is the right frame of reference, but the epistemology of popularity and incentives suggests that it is not.

A related concern is that focusing on the biases of others interferes with the aim of avoiding biases of our own. There is increasingly an awareness that market-participants ought to avoid prejudices influencing their decision-making, with an attendant rise in blind refereeing, orchestral auditions behind opaque screens, anonymized resumes, and so on. A renewed emphasis on the irrelevant characteristics of applicants may look like we're going backwards. And it is certainly true that there are procedural conflicts that could arise: if an orchestra auditions behind an opaque screen to avoid bias, then the judges cannot simultaneously discount for beauty. Efforts to avoid our own bias may make it harder to counteract the bias of others, and there's no reason to deny that the former may sometimes take precedence over the latter. It is also true that the strategy outlined above depends on other people succumbing to various prejudices, which may seem inconsistent with our own efforts to avoid bias. But there is a higher-level consistency at work, I think. Buyers are looking for good service, and in the process are rewarding sellers for quality. Discounting popularity for beauty is one way of connecting with high quality sellers and in the process rewarding them for their service, while avoiding one's biases in hiring is another. Both are attempts to purify the connection between buyer and seller so as to ensure that the buyer gets what he or she really wants and the seller is appropriately rewarded for their quality of service.

## Notes

- 1 See, e.g., Alvin Goldman's *Knowledge in a Social World* (Oxford: Oxford University Press, 1999). An interesting exception is Cass Sunstein's "Deliberating Groups versus Prediction Markets" in Alvin Goldman and Dennis Whitcomb (eds.) *Social Epistemology: Essential Readings* (Oxford: Oxford University Press, 2011).
- 2 Tyler Cowen gives all this advice and further references to food critic lore sprinkled across *An Economist Gets Lunch* (New York: Dutton, 2012). Economists are of course pioneers of incentive-based epistemology.
- 3 The lifetime wage premium for being above-average in attractiveness is estimated in the *hundreds of thousands* of dollars. See David Hamermesh, *Beauty Pays* (Princeton: Princeton University Press, 2011), p. 47. The effect of height alone for males is estimated at several percentage points of annual wages per inch difference in adolescence, according to Nicola Persico et al., "The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height," *The Journal of Political Economy* 112 (2004). See also Harold Kelley's classic work, "The Warm-Cold Variable in First Impressions of Persons," *Journal of Personality* 18 (1950), 431–439. A complication involves covariation. Persico et al. argue that the vehicle of the height premium is largely adolescent-social activity, not employer bias; if so, it's not so much a height premium as an adolescent-social-activity-premium, and it's at least slightly murkier whether *that* trait is irrelevant.
- 4 Here and elsewhere, when I refer to more or less likable or attractive persons on the basis of their various traits, I don't mean to assert (or deny) that these judgments are correct. The point is only that this is how people generally *judge* the relevant traits. What

I say is compatible with the view that in fact, stout introverts are truly likable and beautiful.

- 5 Thomas Espenshade and Alexandria Radford, *No Longer Separate, Not yet Equal* (Princeton: Princeton University Press, 2009), pp. 92–93.
- 6 To experience some of the terror, see Elizabeth Loftus, *Eyewitness Testimony* (Cambridge: Harvard University Press, 1996). We should presumably keep an eye out for material like this, but to say that our information is unsystematic and fragmented, even as we go about our lives crucially reliant on testimony, would be an understatement.
- 7 See CAJ Coady, *Testimony* (Oxford: Oxford University Press, 1992), pp. 172–173, for some thoughts on dealing with non-cohering testimonial evidence.