http://pubsonline.informs.org/journal/orsc

ORGANIZATION SCIENCE
Articles in Advance, pp. 1–20
ISSN 1047-7039 (print), ISSN 1526-5455 (online)

The Road Not Taken: Technological Uncertainty and the Evaluation of Innovations

David Tan[a]

[a] Foster School of Business, University of Washington, Seattle, Washington 98195

Contact: davidtan@uw.edu, https://orcid.org/0000-0003-0466-2960 (DT)

Received: May 11, 2020
Revised: March 2, 2021; September 27, 2021
Accepted: November 10, 2021
Published Online in Articles in Advance: February 11, 2022

https://doi.org/10.1287/orsc.2021.1567

Abstract. When venturing into unfamiliar areas of technology, inventors face ex ante technological uncertainty, that is many possible alternative technological paths going forward and limited guidance from existing technological knowledge for predicting the likelihood that a given path will successfully result in an invention. I theorize, however, that this ex ante technological uncertainty becomes less apparent when evaluating inventions in hindsight. When one knows that a given technological path turned out to be successful ex post, it may be difficult to appreciate the ex ante plausibility of reasons to prefer alternative paths. As a result, inventions may seem more obvious to those evaluating inventions with the benefit of hindsight. My theory yields a counterintuitive implication; when inventors venture into less familiar areas of technology, there is a greater risk of evaluators overestimating obviousness due to hindsight bias. Empirical evidence comes from novel data on accepted and rejected patent applications, including hand-collected data from the text of applicant objections to obviousness rejections and examiners' subsequent reversals of rejections in response to applicant objections.

Keywords: innovation • technological change

Introduction

Evaluating innovations is an important process in a variety of organizational contexts, including but not limited to making technology adoption decisions, allocating resources for research and development, and responding to new technologies by competitors (Henderson and Clark 1990, Orlikowski and Gash 1994, Tripsas and Gavetti, 2000, Kaplan et al. 2003, Kaplan and Tripsas 2008, Tripsas 2009, Eggers 2012, Toh and Kim 2013, Wry et al. 2014, Ferguson and Carnabuci 2017, Polidoro 2020). Given the importance of such organizational processes, much research has sought to understand the theoretical mechanisms that determine how members of organizations evaluate—and misevaluate—innovations. A prevailing focus in this literature has been on the evaluation of an innovation's *value* (Rogers 1962, Rindova and Petkova 2007, Mueller et al. 2012, Mueller et al. 2014, Berg 2016, Boudreau et al. 2016, Criscuolo et al. 2017, Mueller et al. 2017). For instance, a common theme is that innovations are often characterized by unfamiliarity and uncertainty, and this affects perceptions about usefulness, practicality, or other aspects of the economic value of innovations.

However, a phenomenon that prior research has not examined is the evaluation of *obviousness*. I use the term obviousness to refer to the extent to which an innovation is a predictable outcome of applying prior knowledge. Although the term obviousness is associated with the context of patent examination, the underlying phenomenon is a fundamental concept in the literatures on creativity and innovation. Creativity research has argued that nonobviousness is fundamental to the definition of creativity (Simonton 2012). To be considered creative, an idea must be not only novel and useful but also surprising (Boden 2004). In the widely used "novel and useful" definition of creativity, there is a requirement that one should not be able to arrive at an idea simply by following an algorithmic process (Amabile 1983). In a similar vein, a fundamental concept in the innovation literature is that innovations vary in the extent to which they follow or deviate from prior technological trajectories (Dosi 1982, Abernathy and Clark 1985, Tushman and Anderson 1986, Anderson and Tushman 1990). Innovations that closely follow prior technological trajectories represent more predictable outcomes of applying prior knowledge, for instance, by following technological paradigms (Dosi 1982). Such innovations are more obvious. In contrast, innovations that deviate from prior technological trajectories represent greater leaps in the face of technological uncertainty (Fleming 2001). Such innovations are less obvious.

Given the fundamental nature of obviousness as a dimension of innovation, understanding the process through which organization members evaluate an innovation's obviousness is important. A core component of organizational learning and evolution is how organizations respond to stimuli that violate expectations based on prior knowledge (March and Simon 1958, Cyert and March 1963). Innovations can serve as triggers for organizational learning when they appear "unexpected" (Greve and Taylor 2000). By highlighting limitations in the predictive power of prior knowledge, innovations can motivate organizations to either strategically avoid contexts where their prior knowledge is no longer appropriate (Theeke et al. 2018) or update their knowledge. For this to occur, however, organization members must be able to accurately recognize the extent to which an innovation violates or follows expectations based on prior knowledge, that is, an innovation's obviousness.

In this study, I contribute several novel insights about the evaluation of obviousness. First, I propose that the evaluation of obviousness involves distinctive theoretical processes from those that have been previously studied in the literature on the evaluation of innovations. The prior literature has emphasized how the evaluation of an innovation's value involves a forward-looking cognitive process; individuals attempt to predict as-yet unrealized future outcomes, such as how well an innovation will work, how many customers will be interested in it, etc. (see, e.g., Åstebro and Elhedhli 2006, Rindova and Petkova 2007, Berg 2016, Boudreau et al. 2016, Criscuolo et al. 2017). A theme in the prior literature is the role of familiarity; individuals can assess an innovation's value with more certainty in areas of technology where there have been many similar innovations than in areas where there have been few.

In contrast, the evaluation of an innovation's obviousness involves a backward-looking cognitive process. Instead of attempting to predict as-yet unrealized future outcomes, individuals attempt to determine whether an already realized outcome—a focal innovation—could have been predicted using prior knowledge. A distinctive theoretical insight about this process is that, unlike the process of reasoning forward toward unrealized future possibilities, reasoning backward from a realized outcome does not naturally alert individuals to the plausibility of counterfactuals, ex ante reasons to prefer alternative technological paths besides the one that resulted in a focal innovation. If individuals overlook plausible counterfactuals, then an innovation will tend to seem more obvious in hindsight.

This distinctive theoretical insight yields a counterintuitive prediction about the role of familiarity. In a forward-looking process, unfamiliarity makes unrealized future outcomes seem less predictable. In contrast, I propose that in a backward-looking process, unfamiliarity increases the risk of hindsight bias, which makes realized outcomes seem more predictable. The logic is that, in more familiar areas, there tend to be clearer priors—such as in the form of technological paradigms—about what technological paths should and, just as importantly, should not be successful (Dosi 1982). Whereas prior research has emphasized how paradigmatic expectations instill priors against taking unorthodox approaches when it comes to forward-looking processes (see, e.g., Garud and Rappa 1994), I propose that such strong priors also provide stronger psychological guards against hindsight bias when it comes to the backward-looking process of evaluating obviousness. In more familiar areas of technology, stronger priors help to remind individuals of the plausibility of counterfactuals, ex ante reasons to prefer alternative technological paths.

To test my theory, I use novel data on the evaluation of obviousness of accepted and rejected U.S. patent applications, including (1) hand-collected data from the text of applicant objections to obviousness rejections and (2) examiners' subsequent reversals of rejections in response to applicant objections. I find that inventions in less familiar areas of technology are more likely to receive obviousness rejections in which examiners overlook counterfactuals, that is, ex ante reasons to prefer alternative technological paths. I find that such obviousness rejections are more likely to be later reversed.

Theory

Hindsight Bias and the Evaluation of Technological Innovations

Technological innovation has long been depicted as a process of search under uncertainty (Nelson and Winter 1982). For instance, inventors are thought to search across the space of technological choices in hopes of discovering combinations that successfully achieve new functionality (Fleming 2001). Under search-based depictions of technological innovation, the key intuition is that it is generally "not possible to ex ante determine the success or failure of any particular technological path" (Garud and Rappa 1994; emphasis in original). It is not simply the case that all potential future technologies can be analytically predicted based on prior knowledge. Therefore, technological innovation involves risky leaps in the face of uncertainty.

While this uncertainty is apparent ex ante to those engaged in technological search, it may not necessarily be apparent ex post to others who evaluate the outcomes of successful search. Research on hindsight bias suggests that individuals tend to perceive causal processes as being more predictable when reasoning backward from a known outcome than when reasoning forward to unknown possible outcomes (Fischhoff 1975, Hawkins and Hastie 1990, Wasserman et al. 1991, Dawes 1993, Roese and Vohs 2012). For instance, after seeing the results of scientific experiments, evaluators tend to believe that they "knew all along" how the results would turn out (Slovic and Fischhoff 1977). They tend to believe that they could have predicted the results by reasoning from prior scientific knowledge and, therefore, that the results were obvious.

I propose that a similar mechanism affects the perceived obviousness of technological innovations. In general, reasoning backward involves looking for antecedents to an observed outcome, whereas reasoning forward involves beginning from antecedents and thinking about possible outcomes. The key distinction is that backward reasoning does not naturally alert evaluators to counterfactuals, ex ante reasons to expect different outcomes from that which was actually observed.

When reasoning forward, individuals are confronted with a wide range of possible outcomes and many-to-many interdependencies among potential causal factors (Dawes 1993). Faced with this uncertainty, individuals feel cognitive pressure to consider causally complex mental models when trying to make predictions. They cannot feel confident in making simplifying leaps of logic in favor of any one specific outcome because they cannot be certain that omitted steps in their reasoning can be easily filled in to support that prediction.

In hindsight, however, there is less uncertainty to motivate consideration of more causally complex explanations. In contrast to the many-to-many interdependencies that individuals face when reasoning forward, when reasoning backward, individuals face only a single known outcome that they need only connect to some set of seemingly sufficient antecedents (Dawes 1993).

As a consequence, the cognitive process at work in the evaluation of innovations may be very different from the cognitive process behind the creation of innovations. When evaluators encounter an innovation, the thought process of reasoning backward naturally helps them to identify lines of reasoning through which prior knowledge favors the technological path that successfully resulted in the innovation. However, reasoning backward does not naturally lead evaluators to think about lines of reasoning that plausibly favored alternative technological paths. Figure 1 provides a visual depiction of this intuition.

When evaluators know that a technological path successfully resulted in an innovation, they become vulnerable to overlooking conflicting guidance from prior technological knowledge. Hindsight allows for false confidence in oversimplified understandings of how inventors arrived at successful innovative outcomes. As a consequence, aspects of an innovative process that were complex and nonobvious to inventors may seem simple and obvious to evaluators.

Technological Familiarity and Counterfactual Reasoning

The possibility of hindsight bias in the evaluation of innovations suggests that factors that make the innovative process more uncertain from the perspective of inventors may not necessarily make innovations seem more uncertain from the perspective of evaluators.

Figure 1. Ex Ante Uncertainty About Which Technological Path Will Yield an Innovation

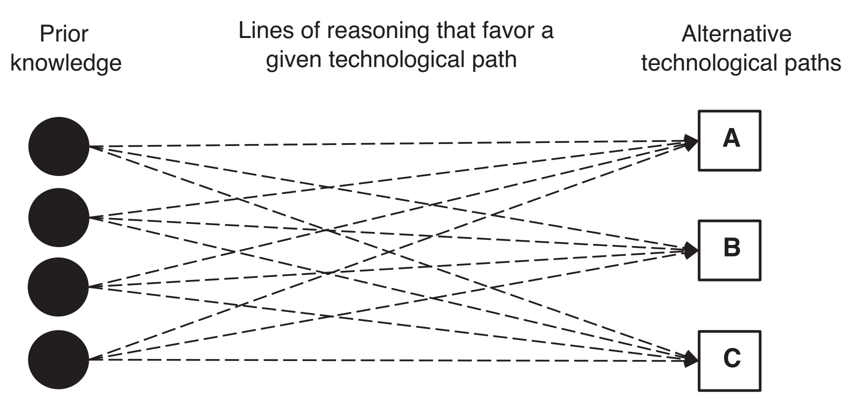

Notes. Circles represent pieces of prior knowledge. Squares represent alternative technological paths. A line between a circle and square represents a line of reasoning, whereby a given piece of prior knowledge suggests an ex ante reason to favor a given technological path, i.e., to expect that a given technological path will successfully yield an innovation. The many-to-many lines of reasoning between prior knowledge and alternative technological paths indicate ex ante technological uncertainty; prior knowledge does not allow one to perfectly predict which particular technological path will successfully yield an innovation. Even if one can think of a line of reasoning based on a given piece of prior knowledge that favors one technological path, there may be other pieces of prior knowledge or other lines of reasoning from the same piece of prior knowledge that favor alternative paths. When reasoning forward toward unrealized future outcomes, this ex ante technological certainty is more apparent. If one begins from a given piece of prior knowledge, moving from left to right will lead one down lines of reasoning that point toward a variety of alternative technological paths. However, when reasoning backward with the benefit of hindsight, ex ante technological uncertainty is less apparent. If one knows ex post that a particular technological path did in fact successfully yield an innovation, moving from right to left will only lead one backward along lines of reasoning that point toward that particular path,rather than lines of reasoning that point toward alternative paths.

Even if there was conflicting guidance from the technological knowledge of the time—that is, ex ante reasons to prefer alternative technological paths—an innovation can still seem obvious if evaluators overlook these reasons afterward. Therefore, an important part of understanding the perception of obviousness is an understanding of what makes these ex ante reasons more apparent to evaluators ex post.

A common thread between the psychological literature on hindsight bias and Kuhn's (1962) model of scientific discovery is that in order to detect anomalies, individuals must be aware of ex ante reasons why an outcome should not have been expected. In Kuhn's (1962) model of scientific discovery, this awareness stems from paradigms. Paradigms reflect the coalescence of prior scientific findings into increasingly clear and sometimes dogmatic beliefs about what should and should not be expected. For instance, Kuhn (1962, p. 60) attributes Joseph Priestley's discovery of oxygen to information-gathering processes that precluded the existence of oxygen: "His commitment to the original test procedure—a procedure sanctioned by much previous experience—had been simultaneously a commitment to the nonexistence of gases that could behave as oxygen did." Similarly, Wilhem Röntgen's discovery of x-rays "commenced with the recognition that his screen glowed when it should not" (Kuhn, 1962, p. 57).

In the context of technological innovation, paradigms can likewise emerge when findings from inventors' searches coalesce into increasingly clear expectations about what paths of technological search should and should not be successful (Dosi 1982). Dosi (1982) observed that paradigms have an "exclusion effect," whereby a "technological paradigm (or research program) embodies strong prescriptions on the directions of technical change to pursue and those to neglect." In the context of cochlear implants, for instance, Garud and Rappa (1994) showed how beliefs about what is and is not feasible give rise to this exclusion effect.

Although this aspect of paradigms may instill priors against taking unorthodox approaches, it also provides stronger psychological guards against hindsight bias when evaluating innovations. Hindsight bias arises when evaluators see that an invention works successfully and inadvertently rely on backward reasoning to see only confirmatory reasons why this should be expected. If paradigmatic expectations instill strong priors about what technological approaches should and should not work, then this actually helps guard against hindsight bias. Strong priors in an area of technology mean that there are clear ex ante beliefs about what not to expect. In areas of technology where such priors are stronger, evaluators are more capable of accurately recalling ex ante beliefs that were contrary to the technological path taken by an innovation.

This insight yields several testable predictions about the evaluation of innovations. First, it suggests, paradoxically, that evaluators will be more vulnerable to hindsight bias in unfamiliar areas of technology, where innovation is actually less predictable but where there are also less clear priors. In familiar areas of technology, where there has been more accumulation of recent inventive activity, the trial-and-error findings from prior inventors' searches are better able to crystallize into wisdom about what to do and what not to do (Fleming 2001). This provides greater depth of knowledge for making predictions about future inventions, that is, engaging in forward reasoning. On the one hand, this means that, in familiar areas of technology, there will be fewer remaining opportunities that require inventors to attempt inventive leaps in the face of uncertainty. On the other hand, from the perspective of evaluators, this provides a guard against hindsight bias. When a major inventive leap does happen, evaluators are more capable of recognizing it.

In contrast, in unfamiliar areas of technology, where there has not been as much inventive activity, innovation is less predictable. Where prior knowledge provides less guidance, inventors must take risky leaps in the face of greater uncertainty (Fleming 2001). Instead of providing clear guidance about one obvious technological direction, prior knowledge may suggest ex ante reasons both for and against many possible technological paths. This creates both greater room and greater necessity for inventors to pioneer new technological paths to address technological challenges. However, precisely because there is less wisdom to facilitate forward reasoning by inventors, there is also less guidance for evaluators. There are less likely to be strong priors about what should and should not work. This makes evaluators more vulnerable to hindsight bias. When evaluators see that an invention works successfully, they are more vulnerable to backward reasoning. Instead of seeing a balanced picture of ex ante reasons both for and against the technological path that resulted in a focal invention, evaluators may see only confirmatory reasons that favor the technological path and overlook reasons that favor alternative paths. This suggests that, in less familiar areas of technology, evaluators will be more likely to view an invention as obvious as a result of overlooking counterfactuals, ex ante reasons to prefer alternative technological paths besides the one that resulted in the invention.

Hypothesis 1. *In less familiar areas of technology, evaluators are more likely to perceive inventions as being obvious as a result of overlooking ex ante reasons that favor alternative technological paths.*

The second testable prediction is that, if inventions seem obvious because evaluators initially overlook counterfactuals, then evaluators' perceptions should change if they are subsequently made aware of these counterfactuals. Hindsight bias is a form of cognitive misestimation, not normative bias against inventions. Evaluators overestimate the predictability of a realized outcome as a result of not considering the full range of plausible alternative outcomes. The intuition is analogous to the role of comparison sets in evaluation processes (Bowers 2014). When evaluating a focal invention, an important comparison set is the range of alternative, unrealized technological paths besides the one that resulted in the focal invention. If evaluators overlook ways in which prior knowledge favored these alternative technological paths, they cannot fully appreciate the nonobviousness of a focal invention's technological path.

If this is the mechanism at work, then in principle, it should operate in the opposite direction as well. When evaluators are made more aware of plausible alternative outcomes, then holding constant the same invention, evaluators' estimates of the ex ante predictability of the observed outcome should decrease. The same invention will subsequently seem less obvious to the same evaluators when evaluators are made more aware of ex ante reasons to prefer alternative technological paths. Different evaluation contexts will vary in the extent to which inventors are able to communicate with evaluators to convey such reasons. But holding constant a given evaluation context with a given level of communication, initial perceptions of obviousness are more likely to be later reversed when these perceptions are the result of initially overlooking counterfactuals.

Hypothesis 2. *When evaluators perceive inventions as being obvious as a result of overlooking ex ante reasons that favor alternative technological paths, these perceptions are more likely to be later reversed.*

Method

Empirical Setting

To test the preceding hypotheses, I take advantage of a natural field setting: the examination of patent applications at the U.S. Patent and Trademark Office (USPTO). Each year the USPTO receives more than 500,000 applications for patent claims on technological innovations (Hegde 2012). As a field setting, patent examination provides an opportunity to study the evaluation of a wide range of real-world innovations on a large scale, a sort of *Drosophila* fly platform for the evaluation of innovations. Patent examination also offers several features that are typically only possible under artificial laboratory conditions. Unlike organizations that evaluate technologies for the purposes of making adoption or investment decisions, the USPTO does not become involved in using or commercializing the technologies being evaluated. This means that patent examiners have no economic or social stakes in defending certain existing technologies or promoting certain new ones. Their evaluations are not affected by the relevance or attractiveness of a technology with respect to existing markets or capabilities. This helps to mitigate the role of economic interests or normative values, which often shape the judgments of evaluators and gatekeepers in other contexts.

Additionally, the USPTO cannot decline to evaluate an invention because it seems difficult to understand: "Every application, no matter how peculiar or confusing, must be assigned somewhere for examination" (MPEP §903.08(d)). This means that for a technology being evaluated, some examiner must ultimately make an attempt to understand it. This helps to rule out the possibility of difficult-to-understand technologies being excluded from evaluation (Zuckerman 1999).

Finally, instead of assessing inventions based on an overall judgment of value or quality, examiners are required to assess inventions based solely on a specific and discrete set of criteria.[1] When submitting patent applications, applicants provide a list of claims describing key technological choices that are thought to distinguish an invention from prior technologies. Examiners then decide whether to allow or reject each of these claims. For each rejected claim, examiners must indicate the criteria under which the claim was rejected. For each criterion for rejection, examiners must provide a written justification, supported by explicit, affirmative reasoning along with relevant references to prior art. Examiners cannot reject patent applications for any reason beyond these specific criteria.

My analysis focuses on the evaluation of obviousness. For an invention to be patentable, it is not enough that an invention has not previously been patented, described, or used (35 U.S.C. §102). It must also be nonobvious. Nonobviousness means that an inventor should not have been able to predictably arrive at an invention by simply following prior technological knowledge (35 U.S.C. §103). To be nonobvious, an invention must move beyond the guidance of prior technological knowledge. This captures the intuition that innovations are more innovative if they break from prior technological trajectories and take risky leaps in the face of uncertainty (Dosi 1982, Fleming 2001). More broadly, in the creativity literature, Simonton (2012) has argued that the concept of obviousness in patent examination captures an important dimension of creativity, that creative ideas must be not just novel and valuable but also surprising.

For the purposes of studying hindsight bias, patent examination is also useful because it is a field setting in which the phenomenon of hindsight bias has face validity. Courts have observed that examiners sometimes "break an invention into its component parts (A + B + C), then find a prior art reference containing A, another containing B, and another containing C, and on that basis alone declare the invention obvious" (Ruiz v. A.B. Chance Co., 357 F.3d 1270, 2004). This kind of reasoning has raised concern because it acts to "discount the value of combining various existing features or principles in a new way to achieve a new result—often the very definition of invention" (Ruiz v. A.B. Chance Co., 357 F.3d 1270, 2004).

In particular, courts have pointed out that obviousness rejections of this kind essentially reflect backward reasoning with the benefit of hindsight; it "simply takes the inventor's disclosure as a blueprint for piecing together the prior art to defeat patentability—the essence of hindsight" (In re Dembiczak, 175 F.3d 994, 1999). As with backward reasoning in other contexts, backward reasoning in patent examination makes examiners vulnerable to overlooking counterfactuals, ex ante reasons to prefer alternative technological paths. In patent examination terms, examiners may overlook ways in which prior technological knowledge "teaches away" from the technological path that resulted in a focal invention. The consequence of backward reasoning and overlooking counterfactuals is that examiners may overestimate the obviousness of inventions.

Sample and Variables

For my analysis, I use data on claim rejections for a sample of 38,067 patent applications filed with the USPTO. This sample includes applications that ultimately resulted in a granted patent (73%) as well as applications that did not (27%).[2] The sample consists of those applications numbering from 10/000,001 to 10/099,999 for which the image file wrapper was available. These applications were filed between 2001 and 2002.

Claim-level data are coded from the text of rejection and allowance notices mailed by examiners as part of the application process. These can be found in the image file wrapper for each application, available through the Patent Application Information Retrieval (PAIR) system. Claim-level data permit analysis at a finer-grained level than the allowance or rejection of an application as a whole. I am able to distinguish between rejections based on obviousness (35 U.S.C. §103) from rejections for other reasons.

Additionally, the data reveal not only which claims were ultimately allowed but also which claims were allowed after being initially rejected. Final allowance rates will understate the rate at which claims were rejected by examiners during the application process by masking rejections that examiners subsequently reversed in light of applicants' responses. In my sample, 54% of claims were ultimately allowed. If one were to infer rejections based on ultimate allowance rate alone, then the implied rejection rate would be only 46%. In reality, the average rejection rate of claims is significantly higher: 71%. This can be seen in Figure 2. The most common basis for rejection is obviousness.

An applicant has the opportunity to respond to and challenge the basis for an examiner's rejections. Based

Figure 2. (Color online) Claim-Level Examination Outcomes

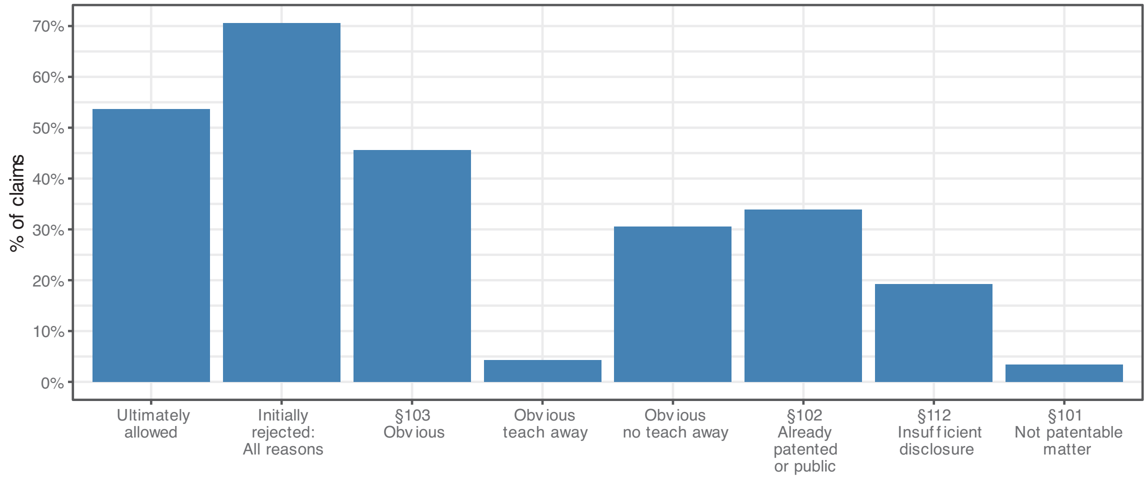

Notes. Allowance and rejection rates for a sample of 1,102,236 claims from 38,067 applications. The column "ultimately allowed" includes claims that were initially allowed as well as claims that were initially rejected but eventually allowed. The other columns refer to the percentages of claims that were initially rejected. These include claims that were initially rejected and not ultimately allowed (i.e., abandoned by the applicant) as well as claims that were initially rejected but ultimately allowed. The four bases for rejection are that a claim is not for patentable subject matter (35 U.S.C. §101), has been previously patented, publicly described, or publicly available (35 U.S.C. §102), is obvious (35 U.S.C. §103), or does not clearly describe the focal invention or disclose enough information to enable others to practice the invention (35 U.S.C. §112). In the patent context, the phrase "teach away" means that prior art actually advised against the technological direction in a focal invention. Prior art may describe reasons to prefer alternative technological paths. The column "teach away" refers to claims that were rejected for obviousness based on an examiner citation to a U.S. patent that the applicant argued teaches away from the invention. The column "no teach away" refers to claims that were rejected for obviousness based on an examiner citation to a U.S. patent that the applicant did not argue teaches away from the invention.

on the applicant's response, an examiner may maintain these rejections or decide to allow some or all of the previously rejected claims.[3] Reversals of obviousness rejections provide a window on cases when examiners may have initially overestimated obviousness. Therefore, to test my hypotheses, it is important to consider rates of all rejections, including those that examiners later reversed.

As described in more detail below, the data also provide insights about cases when examiners may have overlooked counterfactuals. Through manual reading of the text of applicant response letters, I am able to identify cases when the prior art citations that examiners used to support obviousness rejections actually contain reasons to prefer alternative technological paths.

Claims Rejected for Obviousness Because of Overlooked Counterfactuals. To test Hypothesis 1, the dependent variable is the likelihood that a claim is rejected for obviousness based on prior art citations that "teach away" from the inventions in question. In the patent context, the phrase "teach away" means that prior art actually advised against the technological direction in a focal invention (W.L. Gore & Associates, Inc., v. Garlock, Inc., 721 F.2d 1540, 1983; In re Grasselli, 713 F.2d 731, 1983). Prior art may describe reasons why alternative technological paths should be preferred. If patent examiners engage in backward reasoning with the benefit of hindsight, then they become vulnerable to selectively reading only those aspects of prior art that are consistent with a focal invention and overlooking those aspects that "teach away" from the invention. The intuition is that if examiners had read the same prior art without the bias of having first seen the invention, then they could not have predicted the invention based purely on forward reasoning alone.

Identifying such cases required three time-intensive steps. First, I used the text of claim rejection letters to determine which prior art citations were used by examiners specifically for the purposes of supporting obviousness rejections. Prior art citations are used by examiners for a variety of reasons, not all of which have to do with claim rejections. The text of rejection letters allowed me to determine definitively which prior art citations were used specifically to make obviousness rejections.

Second, I matched each individual claim to examiner citations to U.S. patents that were used to reject it for obviousness. I was able to do this for 76% of claims rejected for obviousness. The remaining 24% of claims rejected for obviousness could not be matched to a U.S. patent.[4] I verified that the independent variable of interest, *Unfamiliarity*, has no statistical effect on the probability of a claim not being matched to a U.S. patent.

Finally, for claims that could be matched to a U.S. patent, I used the text of applicant response letters to identify instances when the prior art used by an examiner to support an obviousness rejection actually teaches away from a focal invention. This was done by searching through the text of applicant response letters for variants of the phrase "teach away," including "teach against," "teach the opposite," and "contrary to the teaching of." For each instance of this phrase, the context in which the phrase appeared was manually read to identify the prior art citation that was said to teach away from the focal invention. Table 1 provides excerpts from applicant response letters that contain these phrases.

For the purpose of studying hindsight bias, these "teach away" instances are useful because they come from the same prior art chosen by examiners themselves. Examiners presumably chose the prior art that they did because it contains reasoning that supports the technological direction taken by an invention. However, if it turns out that the same prior art actually contains both reasons for and reasons against the technological direction taken by an invention but examiners only noticed the reasons for, then examiners may have inadvertently read this prior art through the lens of hindsight. They may have selectively read knowledge from the past based on filters reflecting outcomes from the future.

For a similar reason, "teach away" instances are useful for studying reversals of obviousness rejections because they are less subject to gaming and manipulation than other strategies that applicants may use. It is certainly plausible that applicants could always dig hard enough and find different prior art to serve as counterexamples to examiners' citations. But to invoke the "teach away" objection to examiners' own citations, applicants must show that the same prior art cited by examiners actually contains prescriptions against the technological paths that resulted in a focal invention. Applicants cannot freely invoke this objection as a response to every rejection if examiner citations do not in fact contain such prescriptions.

Using examiner citations to prior art that teaches away from the invention, I classify claims rejected for obviousness into two categories: "teach away" and "no teach away." The "teach away" category indicates that a claim was rejected for obviousness based on an examiner citation to a U.S. patent and the applicant argued that this patent teaches away from the invention. The "no teach away" category indicates that a claim was rejected for obviousness based on an examiner citation to a U.S. patent and the applicant did not argue that this patent teaches away from the invention.

Reversals of Obviousness Rejections. To test Hypothesis 2, the dependent variable is the likelihood

Table 1. Examples of Examiner Citations That Teach Away

Example number	Application number	Text from applicant response to examiner rejection (emphases added)
Example 1	10/024,633	Rather, the very fact that Saitoh is an ink jet printing paper requires it to be water absorbing; if this were not the case, you could not ink jet print with the inks as disclosed in that patent. Furthermore, **Saitoh does not suggest or motivate one to employ such a weatherproof coating and, in fact, teaches away from the same since such a coating would render the ink jet printing paper of Saitoh inoperative for its intended purpose**, that is, to absorb the aqueous-based ink applied thereto.
Example 2	10/027,417	Moreover, **Miki et al. teach away from using a heat curing adhesive** at column 21, lines 33-39: "When a heat curing adhesive agent is used, the catheter shaft and balloon are inevitably exposed to heat during the heating required for curing. As a result, it is entirely possible that the balloon diameter will shrink, the balloon bursting pressure will decrease, and the catheter shaft will undergo thermal degradation, and these all **lead to diminished balloon catheter performance**, so the use of a heat curing adhesive agent is not advised."
Example 3	10/043,392	Third, the examiner ignored the disclosure of Borgen et al., which clearly **teaches away from** a combination with the Jackson reference and the present application … As noted above, injection molding and press-fitting processes **operate on much different principles, which are not compatible with one another** … If the socket housing of Borgen et al. was cooler than the socket liner, then it would be inoperable for its intended purpose. Similarly, if the socket liner of Borgen et al. was hotter than the socket housing, then it would be inoperable for its intended purpose. Accordingly, one of ordinary skill in the art would be discouraged from combining the injection molding process of Jackson with the press-fitting process of Borgen et al.
Example 4	10/046,568	Moreover, **Graves et al. actually teaches away from upsetting its careful chemical balance because of the adverse bubbling effects of certain additives, like film formers, polymers, and solvents.** See column 1, lines 39 – 67, discussing how the incorporation of a particular copolymer into the nail formulations therein caused such adverse effects that "the product was not commercially viable or useful." … In light of this teaching, one of ordinary skill would not have been motivated to incorporate at least one first polymer, as presently claimed, into the nail enamel compositions of Graves et al., for fear that it may render them unsatisfactory for their intended purpose.
Example 5	10/053,085	A review of Cable and Isenberg makes clear that **the cited references teach away from combination with one another**. The Background section of Cable '903 expressly criticizes the electrodes described by Isenberg in U.S. Patent 4,582,766, noting that in devices made according to Isenberg, the electrolyte is bound to the electrode, which bonding **results in undesirable mechanical and structural complications** (see Cable '903 at col. 2, lines 38–41) … The cited references likewise teach away from the claimed invention. Whereas the claimed invention recites that the anode and electrolyte are secured to one another, the cited Cable '903 reference expressly criticizes such bonding and instead advocates the disposition of so-called "microslip zones" between the electrolyte and the electrode components (Cable '903 at column 5, lines 53–68).
Example 6	10/058,495	**Devanathan actually teaches away from the proposed combination** with Li … Unlike UHMWPE, which remains workable above its melting temperature, PMMA turns to a liquid (with the approximate viscosity of honey) above its melting temperature. As such, postirradiation quenching **would completely distort, and effectively destroy, Devanathan's acetabular cup**.
Example 7	10/083,205	In addition, **the Ward et al. and Beaupre references teach away from perforating the aluminum vapor barrier layer** of Beaupre. **Ward et al. teaches that any holes in the vapor barrier material are undesirable.** Ward et al. states that "vapor can penetrate [pin holes] to wet the insulation or condense on the underside of the superimposed metal decking. As is well known, wetted insulation has significantly reduced resistance to thermal conductivity than does the same insulation when dry." Ward et al. column 1, lines 43–52. **Beaupre teaches that it is undesirable to deteriorate the vapor barrier** characteristics of a foil vapor barrier, as would certainly be the case if the foil layer were perforated.
Example 8	10/090,293	No such suggestion or teaching is provided by Reiley or Ferree. Rather, Ferree does not mention facet joint replacement or facet joint ailments, and Reiley teaches against the use of artificial discs where facet joint ailments are present … Indeed, **not only Reiley, but the general state of the prior art teaches away from the combination of facet joint replacement with use of an artificial disc**. As expressed in Applicants'

Table 1. (Continued)

Example number	Application number	Text from applicant response to examiner rejection (emphases added)
		specification, "**contraindications** for artificial discs include arthritic facet joints, absent facet joints, severe facet joint tropism, or otherwise deformed facet joints."
Example 9	10/090,358	**Substituting the invention of Scherson into the invention of Henley would defeat this purpose**, creation of a sustained, oxygen-enriched environment through oxygen generation from ambient air under hyperbaric pressure. Indeed, the negative pressure supplied by Henley to drain wound surfaces would entirely teach away from the hyperbaric pressure teaching of Scherson for maintenance of an oxygen-rich environment. Accordingly, **one of skill in the art would not think to modify Henley by combining it with the device of Scherson, because, at the least, the two inventions teach away from each other**.
Example 10	10/097,257	However, those skilled in the art know that there are a multitude of reasons why the 1-to 15-nm particle size regime is generally not accessible to traditional emulsion polymerization or dispersion polymerization techniques … For the invention described in Chandler, the practitioner would want the maximum effective concentration range to be as large as possible and thus the particle size to be as large as possible or practical. The applicants have, surprisingly, found that the use of very low-particle size polymeric nanoparticles actually reduces the viscosity of the dispersion (paragraph 0015), which is **counterintuitive from the known art because smaller particle size usually results in higher viscosity** … Given that the fields of invention are so diverse, **Chandler teaches away from our invention**, and that one would anticipate increased viscosity by moving to lower particle size, anyone of ordinary skill in the art would not consider Chandler to be prior art for the inventive step of very low-particle size (1–15 nm) cross-linked polymeric nanoparticles with attached chromophores.

Note. These examples come from the text of applicant responses to examiners' rejection letters; they point out ways in which the prior art cited by examiners to make obviousness rejections actually teaches away from the focal invention.

that a claim is ultimately allowed. If examiners initially made obviousness rejections based on selective readings of prior art—that is, overlooking those aspects of prior art that "teach away" from the inventions in question—then examiners should be more likely to subsequently reverse these rejections. The logic is that after applicants point out the counterfactuals that examiners initially overlooked, examiners are better able to see ways in which these inventions reflected nonobvious paths that were contrary to the prior art.

Unfamiliarity. The primary independent variable in my analysis is the degree to which an invention is in an unfamiliar area of technology. To capture this, I use Fleming's (2001) measure of component familiarity. The intuition behind this measure is that when there is more recent and frequent innovative activity concentrated in an area of technology, there is greater accumulation of technological knowledge to guide subsequent search. Fleming (2001) showed evidence that greater familiarity is associated with greater certainty.[5]

The measure uses the USPTO's classification of inventions under technology subclasses as proxies for technological components. For each subclass in which an invention is classified, familiarity is measured based on the number and recency of other patents that have been classified under that subclass.

The measure is constructed in two steps, exactly as described in Fleming (2001).[6] First, for each subclass j in invention i, component familiarity is computed as

$$I_{ij} = \sum_k 1\{patent\ k\ uses\ subclass\ j\} \times e^{-\left(\frac{application\ date\ of\ invention\ i - grant\ date\ of\ patent\ k}{time\ constant\ of\ knowledge\ loss}\right)},$$

where the time constant of knowledge loss is set at five years (i.e., 5×365 days), and k indexes all patents granted prior to the application date of the focal invention i.[7]

Second, as described in Fleming (2001), I compute average familiarity across all subclasses j in invention i and then take the square root of this value:

$$component\ familiarity_i = \sqrt{\frac{\sum_j I_{ij}}{\sum_j 1}}$$

To arrive at a measure of unfamiliarity, I simply reverse the scale by taking the negative of Fleming's (2001) measure:

$$unfamiliarity_i = -component\ familiarity_i$$

Higher values of this measure can be interpreted as indicating unfamiliarity. In the analysis, I divide this measure by 100 in order to remove leading zeros in coefficient estimates.

Controls. To account for differences between examiners, all multivariate models are based on within-examiner variation. In other words, examiner fixed effects are included but not estimated as parameters so as to conserve degrees of freedom. All multivariate models also cluster errors at the level of the examiner. Additionally, I include controls for time-varying aspects of examiner experience. As a measure of total prior experience, I include a count of cumulative number of patents examined prior to the application date of a focal invention. As a measure of breadth of experience, I include a count of cumulative number of distinct technology classes examined prior to the application date of a focal invention.

To account for potential differences between applications, I include controls for inventors, prior art citations, and areas of technology. With respect to inventors, I include the number of inventors on an application and the average number of patents granted to these inventors prior to the application date of the focal invention. With respect to applicant prior citations, I include the number of patents cited by the applicant's information disclosure and the number of distinct technology subclasses to which these patents belong.

Finally, I include a set of controls to account for potential differences between law firms. I include the cumulative number of patents filed by a law firm prior to the application date of the focal invention. I also include a set of dummy variables indicating law firm status. *U.S News & World Report* provides a ranking of "Best Law Firms for Patent Law." Firms are ranked into three tiers, "National Tier 1," "National Tier 2," and "National Tier 3." For each tier, I include a dummy variable that takes a value of 1 if a law firm is ranked in that tier. Table 2 provides descriptive statistics.

Empirical Logic

Besides the hypothesized mechanism of hindsight bias, other mechanisms may affect outcomes in the patent examination process and in related ways. Examples include negative biases by examiners, random noise in the examination process, and variation in the underlying quality of patent claims. For instance, one possible mechanism is examiners' relative incentives for making false positive versus false negative errors in initial decisions. If examiners make false negative errors, applicants have strong incentives to petition for the allowance of erroneously rejected claims. However, if examiners make false positive errors, applicants have no incentives to petition for the rejection of erroneously allowed claims. This might create strategic incentives to err on the side of rejection in order to reduce the risk of false positive errors (Sah and Stiglitz 1986, Csaszar 2012). Another possible mechanism is random error. In some areas of technology, examiners may be more prone to errors in both directions, allowing more claims that are below the bar and rejecting more claims that are above the bar. Finally, in some areas of technology, applicants may have lower-quality thresholds for submitting patent applications, resulting in more applications that are below the bar for patentability.

In order to differentiate among these mechanisms empirically, it helps to specify the mathematical relationships between their unobservable data-generating processes and observable patent examination outcomes. Figure 3 illustrates the combination of observed effects that differentiates hindsight bias from negative bias, increased noise, and lower-quality claims.

The top panel of Figure 3 provides a visual representation of the initial decision to allow or reject

Table 2. Descriptive Statistics

	Mean	SD	Minimum	Maximum
Ultimately allowed	0.537	0.499	0	1
Rejected for obviousness, teach away	0.043	0.202	0	1
Rejected for obviousness, no teach away	0.306	0.461	0	1
Rejected under 102	0.339	0.473	0	1
Rejected under 112	0.192	0.394	0	1
Rejected under 101	0.034	0.180	0	1
Ln. (no. of claims)	3.593	0.679	0.000	5.602
Ln. (no. of law firm's prior patents)	4.733	3.694	0.000	10.548
Tier 1 patent law firm (1/0)	0.132	0.339	0	1
Tier 2 patent law firm (1/0)	0.083	0.276	0	1
Tier 3 patent law firm (1/0)	0.044	0.205	0	1
Ln. (no. of inventors)	0.798	0.646	0.000	4.331
Ln. (average no. of inventor prior patents)	1.181	1.140	0.000	6.723
Ln. (no. of applicant citations)	1.326	1.338	0.000	6.378
Ln. (no. of technology subclasses in applicant citations)	1.200	1.191	0.000	5.852
Ln. (no. of examiner's prior patents)	3.158	2.692	0	8
Ln. (no. of technology classes in examiner's prior patents)	2.105	1.776	0	6
Unfamiliarity	−0.012	0.009	−0.073	0.000

Note. This table shows descriptive statistics for a sample of 1,102,236 claims from 38,067 applications; the unit of observation is the claim.

Figure 3. (Color online) Mathematical Relationship Between Categories of Unobserved Data-Generating Processes and Observed Examination Outcomes

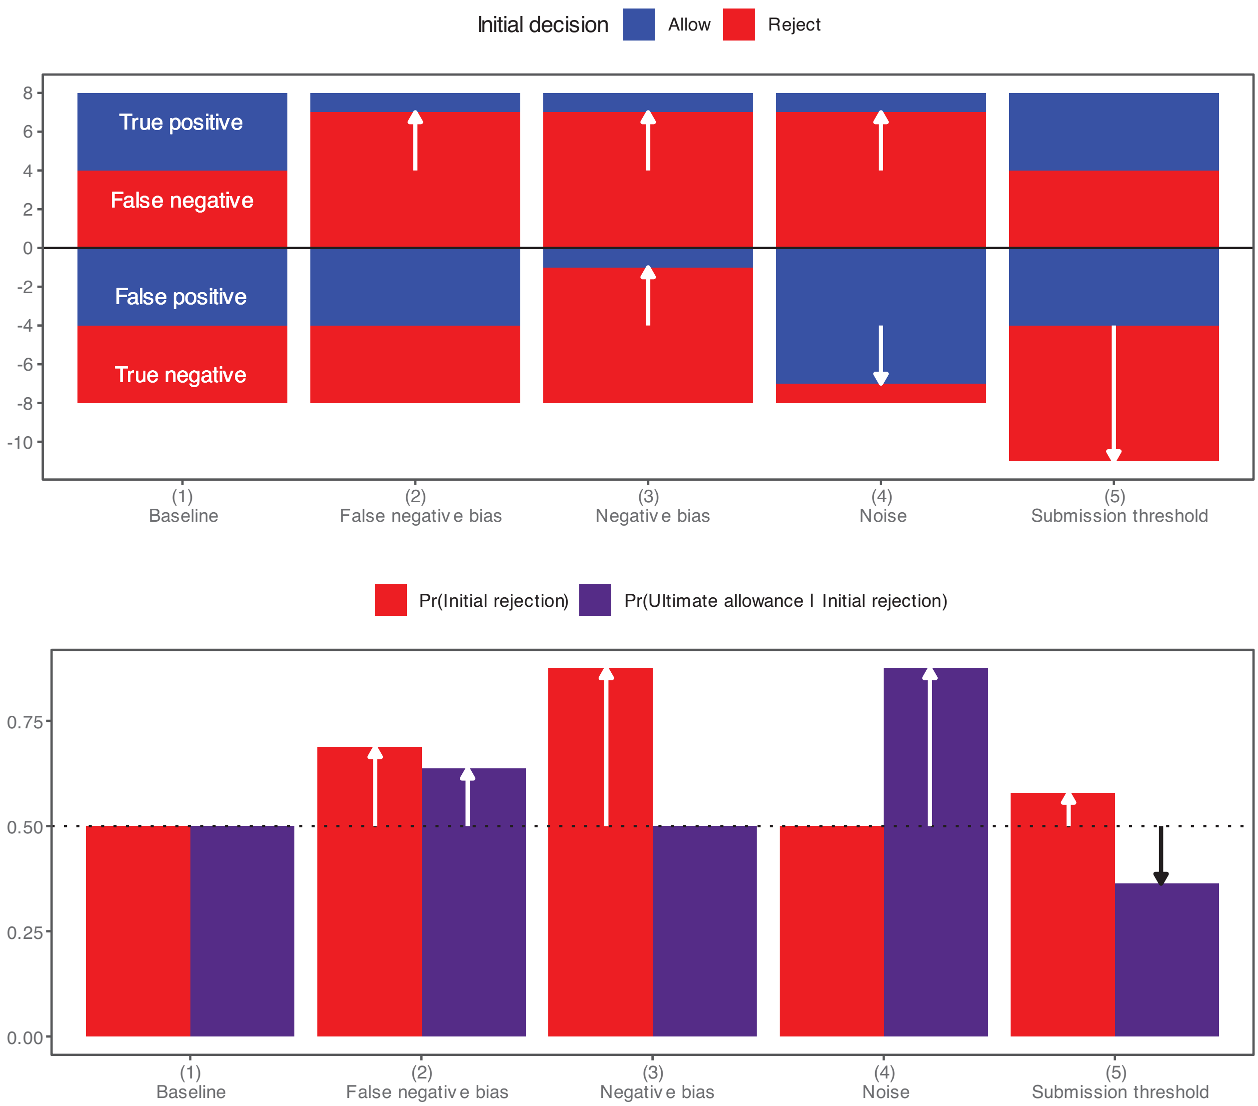

Notes. Comparison of hindsight bias to other categories of mechanisms that may affect patent examination outcomes. Here, hindsight bias falls under category (2) "false negative bias". The top panel illustrates the initial decision to allow or reject patent claims. Patent claims are drawn from a uniform distribution of patentability, ranging from −8 to +8, as represented by the vertical bars. The examination process seeks to allow claims that are above zero and reject claims that are below zero. Blue zones represent claims that are initially allowed. Red zones represent claims that are initially rejected. True positive means initially allowing a claim that is above zero. True negative means initially rejecting a claim that is below zero. False negative means initially rejecting a claim that is above zero. False positive means initially allowing a claim that is below zero. The bottom panel illustrates ways in which the different categories of unobserved data-generating processes can be distinguished based on combinations of their effects on two observed examination outcomes: 1) the probability of initial rejection and 2) the probability of ultimate allowance, conditional on initial rejection. The underlying assumption is that, if an examiner makes a genuine false negative error, an applicant will have a basis for clarifying the error and petitioning for a rejected claim to be allowed. The probability of initial rejection is visually represented as the combined area of all red zones (all initial rejections = false negatives + true negatives) divided by the area of the entire distribution (all claims). The probability of ultimate allowance conditional on initial rejection is visually represented as the area of the red zone above zero (false negatives) divided by the combined area of all red zones (false negatives + true negatives).

patent claims. Patent claims are drawn from a uniform distribution of patentability, ranging from −8 to +8, as represented by the vertical bars.[8] The examination process seeks to correctly classify claims as above the bar or below the bar of patentability. Claims that are above zero should be allowed. Claims that are below zero should be rejected. Blue zones represent claims that are initially allowed. Red zones represent claims that are initially rejected.

Column (1), labeled "baseline," illustrates four possible outcomes of the initial decision. True positive means initially allowing a claim that is above zero.

True negative means initially rejecting a claim that is below zero. False negative means initially rejecting a claim that is above zero. False positive means initially allowing a claim that is below zero. The assumption represented by these outcomes is that patent examiners make noisy evaluations of patentability. Examiners tend to be correct about claims that are either far above or far below the bar of patentability. Examiners are more vulnerable to errors for marginal claims that are just above or just below the bar.

This visual grammar provides a way to represent the effects of different categories of data-generating processes: (2) "false negative bias," (3) "negative bias," (4) "noise," and (5) "submission threshold." In this figure, hindsight bias falls under the category of (2) "false negative bias." False negative bias refers to an increased likelihood of rejecting claims that are above the bar.

This is mathematically different from the other three types of mechanisms. If examiners have strategic incentives to err on the side of making false negative errors in order to avoid false positive errors, then this would fall under the category of (3) "negative bias." Negative bias refers to an increased likelihood of rejecting claims that are above the bar as well as a decreased likelihood of allowing claims that are below the bar. Random error by examiners would fall under the category of (4) "noise." Noise refers to an increased likelihood of both rejecting claims that are above the bar and allowing claims that are below the bar. Finally, if there is a higher submission rate of applications that are below the bar for patentability, this would fall under the category of (5) "submission threshold."

The challenge of studying patent examination is that the quantities in the top panel of Figure 3 are unobserved. It is generally not possible for researchers to independently and definitively classify examiners' initial decisions as correct or incorrect. Therefore, the bottom panel of Figure 3 illustrates ways in which the different categories of unobserved data-generating processes can be distinguished based on combinations of their effects on two observed examination outcomes: (1) the probability of initial rejection and (2) the probability of ultimate allowance, conditional on initial rejection. The underlying assumption is that if an examiner makes a genuine false negative error, an applicant will have a basis for clarifying the error and petitioning for a rejected claim to be allowed. Under this assumption, differences in unobserved classification outcomes—for example, true versus false positives, true vs. false negatives—have mathematically determinate combinations of effects on the two observed outcomes.

In the bottom panel of Figure 3, the probability of initial rejection is visually represented as the combined area of all red zones (all initial rejections = false negatives + true negatives) divided by the area of the entire distribution (all claims). The probability of ultimate allowance conditional on initial rejection is visually represented as the area of the red zone above zero (false negatives) divided by the combined area of all red zones (false negatives + true negatives). To illustrate this, in the "baseline" category, the combined area of all red zones is 8 (0 to 4 and −8 to -4). The area of the entire distribution is 16 (−8 to 8). The area of the red zone above zero is 4 (0 to 4). Therefore, the probability of initial rejection is 0.5 (8 divided by 16). The probability of ultimate allowance conditional on initial rejection is 0.5 (4 divided by 8).

The bottom panel of the figure illustrates how (2) "false negative bias" has a distinctive combination of effects as compared with those of other data-generating processes. False negative bias results in increased probabilities of both initial rejection and ultimate allowance conditional on initial rejection.

In contrast, (3) "negative bias" results in an increased probability of initial rejection but no change in the probability of ultimate allowance conditional on initial rejection. The reason is that negative bias causes an increased likelihood of true negatives, which translates into more initially rejected claims that do not ultimately get allowed.

In contrast, (4) increased "noise" results in no change in the probability of initial rejection but an increased probability of ultimate allowance conditional on initial rejection. The reason is that an increase in random, unbiased error results in more initial errors in both directions.

Finally, (5) a lower "submission threshold" results in an increased probability of initial rejection but a decreased probability of ultimate allowance conditional on initial rejection. The reason is that an increased number of claims below the bar increases the number of true negatives.

The key takeaway from Figure 3 is that, whereas theoretical quantities of interest, such as probability of false negatives, are unobserved, hindsight bias should have a mathematically distinct combination of observed effects—increased probability of initial rejection and increased probability of allowance of initially rejected claims—as compared with negative bias, increased noise, or lower-quality claims.

Results

Obviousness Rejections Due to Overlooked Counterfactuals

Table 3 shows results from linear probability models. All results are based on within-examiner variation (i.e., examiner fixed-effects), and errors are clustered at the examiner level.

Table 3. Linear Probability Models of Claim Outcomes

	Allowed (1)	Rejected for obviousness, teach away (2)	Rejected for obviousness, no teach away (3)
ln(# of claims)	−0.055***	0.001	−0.024***
	(0.004)	(0.002)	(0.004)
ln(# of law firm's prior patents)	0.001	−0.001**	−0.000
	(0.001)	(0.000)	(0.001)
Tier 1 patent law firm (1/0)	−0.004	−0.002	0.015*
	(0.007)	(0.004)	(0.007)
Tier 2 patent law firm (1/0)	−0.010	−0.000	0.013+
	(0.008)	(0.004)	(0.008)
Tier 3 patent law firm (1/0)	0.040***	0.015*	0.003
	(0.011)	(0.007)	(0.011)
ln(# of inventors)	0.013***	−0.001	−0.005
	(0.004)	(0.002)	(0.004)
ln(avg # of inventor prior patents)	0.009***	−0.002	−0.002
	(0.002)	(0.001)	(0.002)
ln(# of applicant citations)	0.033**	0.013+	−0.006
	(0.012)	(0.007)	(0.013)
ln(# of technology subclasses in applicant citations)	−0.021	−0.006	0.013
	(0.014)	(0.008)	(0.014)
ln(# of examiner's prior patents)	0.020	0.008	−0.014
	(0.024)	(0.011)	(0.022)
ln(# of technology classes in examiner's prior patents)	0.006	−0.021	0.014
	(0.033)	(0.017)	(0.033)
Unfamiliarity	1.260**	0.324*	0.072
	(0.398)	(0.164)	(0.359)
Fixed effect	Examiner	Examiner	Examiner
Cluster	Examiner	Examiner	Examiner
Observations	1,102,236	1,102,236	1,102,236
Adjusted R^2	0.211	0.130	0.196

Notes. This table reports results from linear probability models of binary variables; the unit of observation is the claim. In column (1), the dependent variable indicates that a claim was ultimately allowed. In column (2), the dependent variable indicates that a claim was rejected for obviousness based on prior art that "teaches away" from the invention. In column (3), the dependent variable indicates that a claim was rejected for obviousness but not based on prior art that "teaches away" from the invention. All models are based on within-examiner variation (i.e., examiner-fixed effects), and errors are clustered at the examiner level.

$+p < 0.1$; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

As a first step, I begin by testing a baseline expectation about the *Unfamiliarity* measure. Fleming (2001) suggested that inventors face greater uncertainty when working with unfamiliar technological components. Such inventions should on average be more patentable than inventions discovered with the benefit of more guidance from recent inventive activity.

As a test of this baseline expectation, in model (1), the dependent variable indicates whether a claim was ultimately allowed. As described earlier, my sample includes applications that resulted in granted patents as well as applications that did not, and within applications, it includes claims that were ultimately allowed as well as claims that ultimately were not. In this model, *Unfamiliarity* has a significant positive effect. This suggests that inventions using unfamiliar technologies are more likely to result in patentable claims than inventions using familiar technologies.

Hypothesis 1 suggests that, in unfamiliar areas of technology, examiners are more likely to reject claims for obviousness as a result of overlooking counterfactuals. Empirically, this is captured based on a claim being rejected for obviousness based on prior art that "teaches away" from the inventions in question.

In model (2) of Table 3, the dependent variable indicates that a claim was rejected for obviousness based on prior art citations that teach away from the invention. In this model, *Unfamiliarity* has a significant positive effect. To facilitate practical interpretation of effect size, Figure 4 provides a visual representation of the estimated coefficient.

This result suggests that inventions using unfamiliar technologies are more likely to be initially rejected for being obvious as a result of examiners overlooking counterfactuals. Recall Fleming's (2001) evidence that there is less ex ante guidance and greater ex ante uncertainty for inventors working with unfamiliar

Figure 4. (Color online) Estimated Coefficients

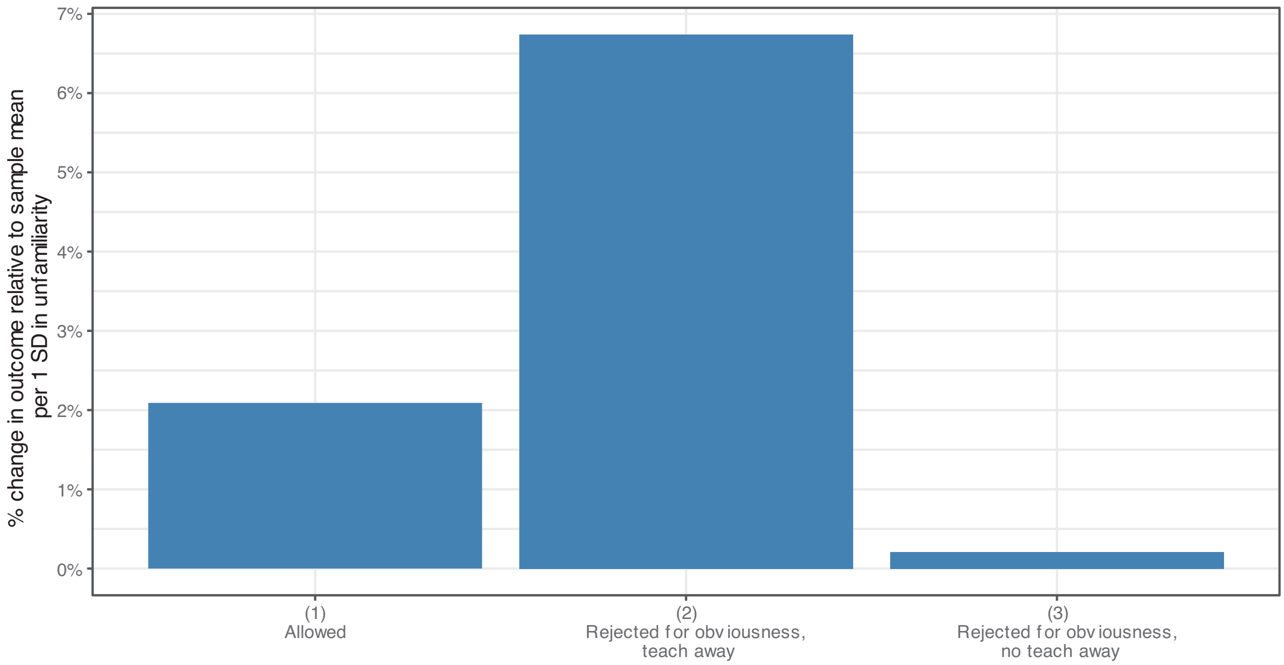

Notes. Visual representation of the estimated coefficients from Table 3. Column numbers correspond to column numbers in Table 3. Plotted coefficients are for the *Unfamiliarity* variable. To facilitate practical interpretation of effect sizes, coefficients for the *Unfamiliarity* variable are multiplied by one standard deviation in the *Unfamiliarity* variable and divided by the sample mean of the dependent variable. Therefore, the magnitudes can be interpreted as the percentage change in the dependent variable relative to the sample mean per one standard deviation increase in the *Unfamiliarity* variable.

technologies. The positive effect of Fleming's (2001) *Unfamiliarity* variable on likelihood of obviousness rejections in this model would be difficult to reconcile without the theorized mechanism. Instead of seeing ex ante uncertainty, that is, a balance of ex ante reasons both for and against a particular direction, examiners overlook ex ante reasons against a technological direction that has been shown to be successful in a focal invention. As a result, examiners are more likely to view the technological direction in an invention as being obvious.

Note that the theoretical mechanism behind Hypothesis 1 is specific to one particular source of obviousness rejections, obviousness rejections that occur because examiners overlooked counterfactuals. However, obviousness rejections can in general occur for many other reasons that do not relate to this mechanism and for which the effect of *Unfamiliarity* is theoretically indeterminate. Hence, it is important to differentiate between obviousness rejections that do and do not occur because examiners overlooked counterfactuals. Obviousness rejections due to overlooked counterfactuals are important for providing affirmative evidence of the theorized mechanism. Obviousness rejections that occur for other reasons are also important for ruling out affirmative evidence of alternative mechanisms. Even if the *Unfamiliarity* variable has the predicted effect on the dependent variable for which it is theoretically relevant, there would be concern about alternative mechanisms if it also has the same effect on other dependent variables for which it is not theoretically relevant.

As a falsification test, in model (3) of Table 3, the dependent variable indicates that a claim was rejected for obviousness but not based on prior art that teaches away from the invention. Recall that to "teach away" means that the prior art cited by examiners actually contains prescriptions against the technological paths that resulted in a focal invention. Applicants cannot freely invoke this objection as a response to every rejection if examiner citations do not in fact contain such prescriptions. The dependent variable in this model indicates that a claim was rejected for obviousness, but the applicant did not argue that the examiner citation teaches away from the invention. Contrary to the predicted positive effect in Hypothesis 1, in this model, *Unfamiliarity* does not have a significant positive effect.

As a further falsification test, Table 4 shows that the *Unfamiliarity* variable does not have positive effects on the likelihood of rejections for other reasons besides obviousness: 35 U.S.C. §101, §102, or §112. Together, these falsification tests provide confidence that the evidence for Hypothesis 1 is not a spurious byproduct of a more generalized mechanism through

Table 4. Linear Probability Models of Other Types of Rejections

	Rejected 102 (1)	Rejected 112 (2)	Rejected 101 (3)
Ln.(no. of claims)	−0.038***	−0.031***	−0.003
	(0.004)	(0.003)	(0.002)
Ln. (no. of law firm's prior patents)	−0.000	0.001*	−0.000
	(0.001)	(0.001)	(0.000)
Tier 1 patent law firm (1/0)	0.009	−0.014*	0.003
	(0.007)	(0.006)	(0.003)
Tier 2 patent law firm (1/0)	−0.014+	−0.006	0.000
	(0.008)	(0.007)	(0.004)
Tier 3 patent law firm (1/0)	−0.006	−0.011	−0.003
	(0.010)	(0.009)	(0.004)
Ln. (no. of inventors)	−0.003	0.001	0.008***
	(0.003)	(0.003)	(0.002)
Ln. (Average no. of inventor prior patents)	−0.005*	−0.002	−0.001
	(0.002)	(0.002)	(0.001)
Ln. (no. of applicant citations)	0.001	0.010	−0.003
	(0.012)	(0.011)	(0.006)
Ln. (no. of technology subclasses in applicant citations)	0.006	−0.015	0.002
	(0.013)	(0.012)	(0.007)
Ln. (no. of examiner's prior patents)	−0.017	−0.004	−0.004
	(0.021)	(0.020)	(0.009)
Ln. (no. of technology classes in examiner's prior patents)	0.004	0.013	0.003
	(0.029)	(0.026)	(0.013)
Unfamiliarity	−0.768*	0.040	−1.003***
	(0.340)	(0.355)	(0.224)
Fixed effect	Examiner	Examiner	Examiner
Cluster	Examiner	Examiner	Examiner
Observations	1,102,236	1,102,236	1,102,236
Adjusted R^2	0.167	0.218	0.203

Notes. This table reports results from linear probability models of binary variables representing other types of rejections besides obviousness (35 U.S.C. §103). The other three bases for rejection are that a claim is not for patentable subject matter (35 U.S.C. §101), has been previously patented, publicly described, or publicly available (35 U.S.C. §102), or does not clearly describe the focal invention or disclose enough information to enable others to practice the invention (35 U.S.C. §112). The unit of observation is the claim. In column (1), the dependent variable indicates whether a claim was rejected under §102. In column (2), the dependent variable indicates whether a claim was rejected under §112. In column (3), the dependent variable indicates whether a claim was rejected under §101. All models are based on within-examiner variation (i.e., examiner-fixed effects), and errors are clustered at the examiner level.

$+p < 0.1$; $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

which *Unfamiliarity* has a positive effect on all forms of rejections. The results suggest that the *Unfamiliarity* variable is operating via the hypothesized mechanism, examiners selectively reading prior art in ways that overlook counterfactuals.

Reversals of Obviousness Rejections

Hypothesis 2 suggests that obviousness rejections due to examiners overlooking counterfactuals are more likely to be subsequently reversed. As a window on this process, I look at the likelihood of claims being ultimately allowed after initially being rejected for obviousness.

Table 5 shows results from a model of claim allowance. The model includes indicator variables for different types of rejections. Being initially rejected for obviousness under §103 reduces a claim's likelihood of ultimate allowance. However, the indicator for "Rejected for obviousness, teach away" has a significant positive effect. When an obviousness rejection is based on prior art that teaches away, the claim is 6.3% more likely to be ultimately allowed as compared with all claims rejected for obviousness. This is practically significant, given that the average probability of allowance for claims rejected for obviousness is 50.9. The estimated coefficient for "Rejected for obviousness, teach away" implies a 12.3% increase (6.3/50.9) relative to the average probability of allowance for obviousness rejections in general.

Discussion

Although much research has sought to understand the theoretical mechanisms that determine how members of organizations evaluate—and misevaluate—innovations, the prevailing focus has been on the evaluation of an innovation's *value* (Rindova and Petkova 2007, Berg 2016, Boudreau et al. 2016, Criscuolo et al. 2017, Mueller et al. 2012, Mueller et al. 2014, Mueller et al.

Table 5. Linear Probability Model of Claim Allowance

	Allowed
Ln. (no. of claims)	−0.056***
	(0.004)
Ln. (no. of law firm's prior patents)	0.001
	(0.001)
Tier 1 patent law firm (1/0)	−0.003
	(0.007)
Tier 2 patent law firm (1/0)	−0.010
	(0.008)
Tier 3 patent law firm (1/0)	0.039***
	(0.011)
Ln. (no. of inventors)	0.013***
	(0.004)
Ln. (average no. of inventor prior patents)	0.009***
	(0.002)
Ln. (no. of applicant citations)	0.032*
	(0.012)
Ln. (no. of technology subclasses in applicant citations)	−0.019
	(0.014)
Ln. (no. of examiner's prior patents)	0.018
	(0.024)
Ln. (no. of technology classes in examiner's prior patents)	0.007
	(0.033)
Unfamiliarity	1.229**
	(0.400)
Rejected 101	0.027*
	(0.013)
Rejected 112	0.035***
	(0.006)
Rejected 102	−0.056***
	(0.004)
Rejected 103	−0.017***
	(0.005)
Rejected for obviousness, teach away	0.063***
	(0.011)
Fixed effect	Examiner
Cluster	Examiner
Observations	1,102,236
Adjusted R^2	0.214

Notes. Results from a linear probability model of allowance; the unit of observation is the claim. The dependent variable indicates that a claim was ultimately allowed. The model is based on within-examiner variation (i.e., examiner-fixed effects), and errors are clustered at the examiner level.
$+p < 0.1$; $^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

2017). In contrast, this study contributes novel insights about an understudied phenomenon: the evaluation of *obviousness*.

First, I show that the evaluation of obviousness involves distinctive theoretical processes that have not previously been studied. Whereas the prior literature has emphasized how the evaluation of an innovation's value involves a forward-looking cognitive process—individuals attempt to predict as-yet unrealized future outcomes, such as how well an innovation will work, how many customers will be interested in it, etc. (see, e.g., Åstebro and Elhedhli 2006, Rindova and Petkova 2007, Berg 2016, Boudreau et al. 2016, Criscuolo et al. 2017)—I propose that the evaluation of an innovation's obviousness involves a backward-looking cognitive process.

This insight yields a distinctive theoretical intuition. Unlike the process of reasoning forward toward unrealized future possibilities, reasoning backward from a realized outcome does not naturally alert individuals to the plausibility of counterfactuals, ex ante reasons to prefer alternative technological paths. As a consequence, if individuals overlook plausible counterfactuals, then an innovation will tend to seem more obvious in hindsight. This distinctive theoretical intuition, in turn, yields a counterintuitive empirical prediction about the role of familiarity. Whereas in a forward-looking process, unfamiliarity makes unrealized future outcomes seem less predictable, I propose that in a backward-looking process, unfamiliarity increases the risk of hindsight bias, which makes realized outcomes seem more predictable. The logic is that, in more familiar areas of technology, stronger priors, such as from technological paradigms (see, e.g., Dosi 1982 and Garud and Rappa 1994), help to remind individuals of the plausibility of counterfactuals, ex ante reasons to prefer alternative technological paths. I find evidence of this mechanism using novel data on the evaluation of obviousness of accepted and rejected U.S. patent applications, including (1) hand-collected data from the text of applicant objections to obviousness rejections and (2) examiners' subsequent reversals of rejections in response to applicant objections.

This study sheds light on a novel way in which innovations may be misevaluated. A prevailing theme in the prior literature is that innovations are often characterized by unfamiliarity and uncertainty, and this diminishes the perceived usefulness, practicality, or other aspects of the economic value of innovations (see, e.g., Rogers 1962). In the literature on technological paradigms, there is a similar view that innovations are discouraged when innovations are perceived as deviating from prior technological trajectories (Dosi 1982, Garud and Rappa 1994). In contrast, this study shows that innovations may be misevaluated for a very different reason. They may be perceived as being too consistent with prior technological trajectories and, therefore, not very innovative. Even though unfamiliar areas of technology present inventors with less clear guidance from prior knowledge, this also makes evaluators more vulnerable to hindsight bias, selectively seeing only those aspects of prior knowledge that seem in hindsight to be clear while overlooking those aspects of prior knowledge that presented conflicting guidance. Therefore, efforts to pioneer new technological paths might be discounted by evaluators as a result of these paths seeming predictable in hindsight. Importantly, this study suggests that such false pictures of predictability are more likely to occur in less familiar areas of technology that are actually less predictable.

At the same time, this study suggests a potential counterintuitive benefit of innovating in familiar areas. As with scientific progress, technological progress depends in part not just on discovering path-breaking ideas but also on recognizing when an idea is in fact path-breaking (Kaplan and Vakili 2015). In both science and technology, increased concentration of innovative activity within a given area can give rise to increasingly dogmatic expectations about what should and should not work and increasingly incremental innovation in service of these expectations (Kuhn 1962, Dosi 1982). Although this may seem antithetical to unorthodox thinking, Kuhn (1962) argued that the sharpening of expectations in well-trodden areas also has a flip side. It can facilitate the recognition of anomalies. The findings in this study suggest that until enough innovative activity has coalesced in an area of technology, there might not be a sufficiently clear baseline of expectations against which innovations can be contrasted. Paradoxically, increased dogmatism in beliefs and expectations in an area of technology—even if this discourages unorthodox paths—may help facilitate recognition that such paths are in fact innovative.

Theoretically, this study answers a question that is critical to Kuhn's (1962) argument but missing from his case study of anomalies: Why do anomalies go unnoticed? There is often rich historical information on cases when anomalies do get noticed. These cases go down in history as heroic stories of scientific discovery. But little is known about cases when anomalies go unnoticed. These go unrecorded precisely because they did not seem remarkable at the time. For instance, Kuhn (1962, p. 59) speculates that, although Wilhem Röntgen was the first to recognize that no prior scientific principle could predict the glow of x-rays in his experiments, "a number of other experimentalists must for some time have been producing those rays without knowing it." The novel hand-collected data in this study provide a rare window on a large sample of unnoticed anomalies. Moreover, the data provide a rare window into what evaluators were thinking when they dismissed these anomalies as unremarkable. The findings suggest that anomalies go unnoticed because evaluators selectively see only those aspects of prior knowledge that seem to predict these outcomes while overlooking reasons why these outcomes should not be expected.

The insights from this study have implications for entities involved in developing, evaluating, and competing against technological innovations. Beyond patent examination, two other examples of contexts in which the evaluation of obviousness may be important include academic science and startup financing. In academic science, there are rewards for originality (Merton 1957, Gaston 1973, Hagstrom 1974). For example, one of the criteria for National Institutes of Health grants is, "Does the application challenge and seek to shift current research or clinical practice paradigms?" (National Institutes of Health 2016). Research on the evaluation of scientific grant applications suggests that reviewers will reject applications that they perceive to be "derivative" or merely "gap-filling" (Guetzkow et al. 2004). From the perspective of scientists applying for grants or submitting discoveries to journals, it is undesirable if ideas seem obvious in hindsight. In the context of startup financing, investors may have disincentives against investing in ideas that will reinforce rather than break from existing technological trajectories in which incumbents have the advantage (Tushman and Anderson 1986). Like patent examiners, investors are often inundated with pitches for innovations. But unlike patent examiners, investors are not obligated to explain their decisions to entrepreneurs or give entrepreneurs the opportunity to respond to unfavorable decisions. Investors may just quietly screen out ideas based on "gut feel" (Huang and Pearce 2015). Under these circumstances, false negatives in initial impressions have less potential to get corrected through subsequent communication and have greater potential to shape final decisions.

In some contexts, the perception of obviousness may matter to potential adopters of a technology. For instance, in Orlikowski and Gash's (1994) study of the adoption of Lotus Notes, different individuals within the same organization had different perceptions about how Notes related to prior technologies. Some individuals adopted a "simplistic understanding" of Lotus Notes as a continuation of prior communications technology: "it is not a radical change"; "Notes will do to fax what fax did to telex" (Orlikowski and Gash 1994). Others viewed the technology as a major departure: "I knew in an hour that it was a breakthrough product, a revolution"; "We realized Notes was a transformation technology" (Orlikowski and Gash 1994). As these quotes suggest, some customers may be more interested in adopting an innovation if they perceive it as being a departure from rather than continuation of prior technology. One type of situation in which this can arise is when business customers hope to gain a competitive advantage through technology adoption. For instance, business customers may be interested in adopting an innovation if they believe that competitors using existing technologies will not see it as an obvious next step.

One final set of entities that may be affected by hindsight bias consists of organizations that compete against technological innovations. Innovations can potentially threaten the survival of organizations using prior technologies (Tushman and Anderson 1986). In order to respond appropriately, organizations must first understand how an innovation relates to prior technologies. If an innovation follows the same trajectory as prior technologies, then organizations know that their existing knowledge remains relevant (Tushman & Anderson 1986). In contrast, if an innovation follows different underlying principles, then organizations know that they will need to either invest in new knowledge in order to remain competitive or strategically avoid contexts where their prior knowledge is no longer appropriate (Greenstein 2017, Theeke et al. 2018). From this perspective, organizations can suffer as a result of false negatives, perceiving that an innovation does not require new knowledge when in fact it does. Henderson and Clark's (1990) study of the photolithographic alignment industry provides a well-known example. Kasper Instruments suffered poor product performance as a result of assuming that the new generation of proximity aligners could be understood using the same technological knowledge as prior contact aligners (Henderson and Clark 1990). More generally, across several generations of new technology, producers of photolithographic alignment equipment suffered poor returns to research and development spending as a result of relying on knowledge from prior generations of technology (Henderson 1993).

From an organizational learning perspective, recognizing that an innovation deviates from prior knowledge is a critical step. For competitors' innovations to catalyze learning by an organization, these innovations must seem "unexpected" (Greve and Taylor 2000). If organizations are unable to perceive a limitation in the predictive power of prior knowledge, they will not be motivated to search for new knowledge (March and Simon 1958, Cyert and March 1963).

Unfortunately, most settings do not have a natural corrective mechanism. The patent examination context is somewhat unique in that examiners are required to explain their decisions and give applicants the opportunity to respond. This is useful for the purpose of studying hindsight bias in a field setting. However, in most other evaluation settings, there is no analogous corrective mechanism. If an organization erroneously views a competitor's technology as being obvious, the competitor has no mechanism for correcting this misunderstanding. In fact, in the context of an organization responding to a competitor's innovation, the external entity that is best positioned to correct an organization's misunderstanding—the competitor—may have no incentive to do so. In the Kasper Instruments case, for instance, it was to Canon's advantage that Kasper viewed Canon's new technology as obvious and did not invest in new knowledge to remain competitive (Henderson and Clark 1990).

Developing practical interventions for mitigating this tendency in organizations is a challenging but fruitful avenue for future research. This study provides a starting point by identifying a novel theoretical pathway and showing why it matters. If there are structures and processes that help individuals within organizations think about why innovations should not be expected based on prior knowledge, then in theory, these should make organizations more capable of recognizing when innovations embody new principles and require new knowledge. In a manner akin to paradigms in Kuhn's (1962) model of scientific discovery, some structures and processes may help highlight what prior knowledge cannot explain and, in doing so, help alert organizations to the need for new knowledge.

Acknowledgments

This paper benefited from extremely thoughtful guidance from Senior Editor PK Toh and two anonymous reviewers. The author grateful for helpful comments on earlier versions from Chris Rider, Ming Leung, Deepak Hegde, Ezra Zuckerman, Giacomo Negro, Matt Marx, Oliver Alexy, Francisco Polidoro, Callen Anthony, Aseem Kaul, and seminar participants at the University of Washington, Washington University in St. Louis, Santa Clara University, the Strategy Research Forum, Stanford GSB, the INSEAD Doriot Entrepreneurship Conference, the West Coast Research Symposium, the National University of Singapore, and the University of Minnesota.

Endnotes

[1] The four patentability requirements are that an application: must be for patentable subject matter (35 U.S.C. §101), must not have been previously patented, publicly described, or publicly available (35 U.S.C. §102), must not be obvious (35 U.S.C. §103), and must clearly describe the focal invention and disclose enough information to enable others to practice the invention (35 U.S.C. §112). As an example, a literary work or musical composition would be rejected under §101 because it is not considered patentable subject matter. An application would be rejected under §102 if there is prior art that can be legally considered to be for the same invention. An application would be rejected under §112 if it does not clearly describe what the focal invention is.

[2] This is consistent with Lemley and Sampat's (2008) estimate of 75 percent.

[3] The reject-response process may go on for several rounds. The process ends with a granted patent when 1) all of the claims are allowed or 2) some claims are allowed and the applicant decides to abandon the remaining rejected claims. It ends with abandonment of the application when all claims are rejected and the applicant decides not to further pursue any of the claims.

[4] This could be for several reasons: 1) An examiner citation may be to prior art that is not a non-patent publication or a non-U.S. patent. Non-U.S. patent prior art is significantly more difficult to systematically recognize and clean in this context because the source files are scanned images of printed paper documents, and the documents are letters written in prose. Attempts at automated processing introduced significant errors, given the widely varying structures of non-U.S. patent prior art. 2) The source files are scanned images of printed paper documents, and optical character recognition is not always able to produce readable text. 3) Because the source files are examiner letters written in prose, there is variation in writing style, and some letters may not clearly state both the claim being rejected and the prior art used to support the rejection. Whereas the majority of letters follow a common structure in which the examiner clearly states both the claim being rejected and the prior art used to support the rejection, some letters use shorthand to refer to prior art, and it was not possible to accurately link shorthand phrases to the prior art without manually reading the context of entire letters, which was not feasible on a scale of a million claims.

[5] Fleming (2001) originally used this to measure the relative familiarity of different areas of technology from the perspective of inventors. In using this measure in a study of patent examination, I am making an assumption that the relative familiarity of different areas of technology from the perspective of inventors is correlated with familiarity from the perspective of examiners.

[6] The only modification is that I use the grant dates of prior patents instead of application dates. Fleming (2001) interpolated application dates of patents granted prior to 1975. For simplicity and internal consistency and because the interpolation formula is not known, I use grant dates.

[7] Because this measure is a composite of frequency and recency, a potential concern is that the discount factor suppresses the effect of technology areas that were well-trodden but from a long time ago. A simple way to address this is to remove the discount rate entirely (implying zero knowledge loss over time). In that case, the measure would only reflect the frequency of prior patents in a technology area, regardless of how long ago. Results are essentially identical under no discounting, Fleming's original 5 year constant, and any discount rate in between.

[8] A uniform distribution is used for mathematical simplicity.

References

Abernathy WJ, Clark KB (1985) Innovation: Mapping the winds of creative destruction. *Res. Policy* 14(1):3–22.

Amabile TM (1983) *The Social Psychology of Creativity* (Springer New York, New York).

Anderson P, Tushman ML (1990) Technological discontinuities and dominant designs: A cyclical model of technological change. *Admin. Sci. Quart.* 35(4):604–633.

Åstebro T, Elhedhli S (2006) The Effectiveness of simple decision heuristics: Forecasting commercial success for early-stage ventures. *Management Sci.* 52(3):395–409.

Berg JM (2016) Balancing on the creative highwire: Forecasting the success of novel ideas in organizations. *Admin. Sci. Quart.* 61(3): 433–468.

Boden MA (2004) *The Creative Mind* (Routledge, New York).

Boudreau KJ, Guinan EC, Lakhani KR, Riedl C (2016) Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Sci.* 62(10):2765–2783.

Bowers A (2014) Relative comparison and category membership: The case of equity analysts. *Organ. Sci.* 26(2):571–583.

Criscuolo P, Dahlander L, Grohsjean T, Salter A (2017) Evaluating novelty: The role of panels in the selection of R&D projects. *AMJ* 60(2):433–460.

Csaszar FA (2012) Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management J.* 33(6):611–632.

Cyert RM, March JG (1963) *A Behavioral Theory of the Firm* (Prentice-Hall, Englewood Cliffs, NJ).

Dawes RM (1993) Prediction of the future vs. an understanding of the past: A basic asymmetry. *Amer. J. Psych.* 106(1):1–24.

Dosi G (1982) Technological paradigms and technological trajectories. *Res. Policy* 11(3):147–162.

Eggers JP (2012) Falling flat: Failed technologies and investmentunder uncertainty. *Admin. Sci. Quart.* 57(1):47–80.

Ferguson JP, Carnabuci G (2017) Risky recombinations: Institutional gatekeeping in the innovation process. *Organ. Sci.* 28(1):133–151.

Fischhoff B (1975) Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *J. Exp. Psych. Hum. Percept. Perform.* 1(3):288–299.

Fleming L (2001) Recombinant uncertainty in technological search. *Management Sci.* 47(1):117–132.

Garud R, Rappa MA (1994) A socio-cognitive model of technology evolution: The case of cochlear implants. *Organ. Sci.* 5(3): 344–362.

Gaston JC (1973) *Originality and Competition in Science: a Study of the British High Energy Physics Community* (University of Chicago Press, Chicago).

Greenstein S (2017) The reference wars: *Encyclopaedia Britannica's* decline and Encarta's emergence: *Encyclopaedia Britannica's* Decline and Encarta's Emergence. *Strategic Management J.* 38(5):995–1017.

Greve HR, Taylor A (2000) Innovations as catalysts for organizational change: Shifts in organizational cognition and search. *Admin. Sci. Quart.* 45(1):54–80.

Guetzkow J, Lamont M, Mallard G (2004) What is originality in the humanities and the social sciences? *Amer. Sociol. Rev.* 69(2):190–212.

Hagstrom WO (1974) Competition in science. *Amer. Sociol. Rev.* 39(1):1.

Hawkins SA, Hastie R (1990) Hindsight: Biased judgments of past events after the outcomes are known. *Psych. Bull.* 107(3): 311–327.

Hegde D (2012) Funding and performance at the US Patent and Trademark Office. *Nat Biotechnol* 30(2):148–150.

Henderson R (1993) Underinvestment and Incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *RAND J. Econ.* 24(2):248–270.

Henderson RM, Clark KB (1990) Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* 35(1):9–30.

Huang L, Pearce JL (2015) Managing the unknowable: The effectiveness of early-stage investor gut feel in entrepreneurial investment decisions. *Admin. Sci. Quart.* 60(4):634–670.

Kaplan S, Tripsas M (2008) Thinking about technology: Applying a cognitive lens to technical change. *Res. Policy* 37(5):790–805.

Kaplan S, Vakili K (2015) The double-edged sword of recombination in breakthrough innovation: The double-edged sword of recombination. *Strategic Management J.* 36(10):1435–1457.

Kaplan S, Murray F, Henderson R (2003) Discontinuities and senior management: assessing the role of recognition in pharmaceutical firm response to biotechnology. *Indust. Corp. Change.* 12(2):203–233.

Kuhn TS (1962) *The Structure of Scientific Revolutions* (University of Chicago Press, Chicago).

Lemley MA, Sampat B (2008) Is the patent office a rubber stamp? *Emory Law J.* 58(1):181–206.

March JG, Simon HA (1958) *Organizations* (Wiley, New York).

Merton RK (1957) Priorities in scientific discovery: A chapter in the sociology of science. *Amer. Sociol. Rev.* 22(6):635–659.

Mueller JS, Melwani S, Goncalo JA (2012) The bias against creativity: Why people desire but reject creative ideas. *Psych. Sci.* 23(1): 13–17.

Mueller JS, Wakslak CJ, Krishnan V (2014) Construing creativity: The how and why of recognizing creative ideas. *J. Exp. Soc. Psych.* 51:81–87.

Mueller J, Melwani S, Loewenstein J, Deal JJ (2017) Reframing the decision-makers' dilemma: Toward a social context model of creative idea recognition. *Acad. Management J.* 61(1):94–110.

National Institutes of Health (2016) Definitions of criteria and considerations for research project grant (RPG/R01/R03/R15/R21/R34) Critiques | grants.nih.gov. Retrieved (April 21, 2020), https://grants.nih.gov/grants/peer/critiques/rpg.htm.

Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change* (Harvard University Press, Cambridge, MA).

Orlikowski WJ, Gash DC (1994) Technological frames: Making sense of information technology in organizations. *ACM Trans. Inf. Syst.* 12(2):174–207.

Polidoro F (2020) Knowledge, routines, and cognitive effects in nonmarket selection environments: An examination of the regulatory review of innovations. *Strategic Management J.* 41(13):2400–2435.

Rindova VP, Petkova AP (2007) When is a new thing a good thing? Technological change, product form design, and perceptions of value for product innovations. *Organ. Sci.* 18(2):217–232.

Roese NJ, Vohs KD (2012) Hindsight bias. *Perspect. Psych. Sci.* 7(5): 411–426.

Rogers EM (1962) *Diffusion of Innovations* (Free Press of Glencoe, Macmillan, New York, London).

Sah RK, Stiglitz JE (1986) The architecture of economic systems: Hierarchies and polyarchies. *Amer. Econom. Rev.* 76(4):716–727.

Simonton DK (2012) Taking the U.S. patent office criteria seriously: A quantitative three-criterion creativity definition and its implications. *Creativity Res. J* 24(2–3):97–106.

Slovic P, Fischhoff B (1977) On the psychology of experimental surprises. *J. Exp. Psych. Hum. Percept. Perform.* 3(4):544–551.

Theeke M, Polidoro F, Fredrickson JW (2018) Path-dependent Routines in the evaluation of novelty: The effects of innovators' new knowledge use on brokerage firms' coverage. *Admin. Sci. Quart.* 63(4):910–942.

Toh PK, Kim T (2013) Why put all your eggs in one basket? A competition-based view of how technological uncertainty affects a firm's technological specialization. *Organ. Sci.* 24(4):1214–1236.

Tripsas M (2009) Technology, identity, and inertia through the lens of "The Digital Photography Company.". *Organ. Sci.* 20(2):441–460.

Tripsas M, Gavetti G (2000) Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management J.* 21(10–11): 1147–1161.

Tushman ML, Anderson P (1986) Technological discontinuities and organizational environments. *Admin. Sci. Quart.* 31(3):439–465.

Wasserman D, Lempert RO, Hastie R (1991) Hindsight and causality. *Pers. Soc. Psych. Bull.* 17(1):30–35.

Wry T, Lounsbury M, Jennings PD (2014) Hybrid vigor: Securing venture capital by spanning categories in nanotechnology. *Acad. Management J.* 57(5):1309–1333.

Zuckerman EW (1999) The categorical imperative: Securities analysts and the illegitimacy discount. *Amer. J. Sociol.* 104(5):1398–1438.

David Tan ("The road not taken: technological uncertainty and the evaluation of innovations") is associate professor of management and organization at the Foster School of Business, University of Washington. He received his PhD from Emory University. His research interests include organizational status, labor markets, legal/regulatory environments, and evaluation processes.