Creativity Under Fire: The Effects of Competition on Creative Production

Citation

Gross, Daniel P. "Creativity Under Fire: The Effects of Competition on Creative Production." Harvard Business School Working Paper, No. 16-109, March 2016. (Revised September 2016.)

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:26211016>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Creativity Under Fire: The Effects of Competition on Creative Production

Daniel P. Gross

Working Paper 16-109

Creativity Under Fire: The Effects of Competition on Creative Production

Daniel P. Gross
Harvard Business School

Working Paper 16-109

Copyright © 2016, 2017 by Daniel P. Gross

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Creativity Under Fire: The Effects of Competition on Creative Production

Daniel P. Gross*
Harvard Business School

December 30, 2016
(first draft: November 2, 2014)

Abstract:

Though fundamental to innovation and essential to many industries and occupations, individual creativity has received limited attention as an economic behavior and has historically proven difficult to study. This paper studies the incentive effects of competition on individuals' creative production. Using a sample of commercial logo design competitions, and a novel, content-based measure of originality, I find that intensifying competition induces agents to produce original, untested ideas over tweaking their earlier work, but heavy competition drives them to stop investing altogether. The results yield lessons for the management of creative workers and for the implementation of competitive procurement mechanisms for innovation.

JEL Classification: D81, D82, D83, M52, M55, O31, O32

Keywords: Creativity; Incentives; Tournaments; Competition;
Radical vs. incremental innovation

*Address: Harvard Business School, Soldiers Field Road, Boston, MA 02163, USA; email: dgross@hbs.edu. I am grateful to Ben Handel, Steve Tadelis, and especially John Morgan, whose conversation provided much of the inspiration for this paper. I also thank Gustavo Manso and Noam Yuchtman for feedback at early stages, colleagues at the UC Berkeley Department of Economics and Haas School of Business, and seminar audiences at UC Berkeley (Economics, Haas, and Fung Institute), Carnegie Mellon, Columbia Business School, Federal Reserve Board, Federal Reserve Bank of New York, Federal Trade Commission, Harvard Business School, Indiana Kelley School of Business, Northwestern Kellogg School of Management, Microsoft Research, Rensselaer Polytechnic Institute, University of Chicago Booth School of Business, University of Toronto Mississauga/Rotman, Washington University in St. Louis Olin Business School, and the CEPR Workshop on Incentives, Management and Organizations for many useful comments. Limited short segments of the text may be similar to passages from another of my papers ("Performance Feedback in Competitive Product Development"), which uses a different dataset from the same setting. This research was supported by NSF Graduate Research Fellowship Grant No. DGE-1106400 as well as the Harvard Business School Division of Faculty and Research. All errors are my own.

The creative act is an important but under-studied phenomenon in management and economics. Millions of people in the U.S. alone work in fields where creativity is essential to job performance, including research, engineering, and professional services – industries which are the engines of innovation and growth in modern developed economies. CEO surveys also show that executives’ top concerns consistently include the creativity of their employees and pursuit of innovation within the firm. Despite its importance, the creative act itself has received limited attention as an economic behavior and has historically proven difficult to study, due to the challenge of measuring creativity and relating it to variation in incentives.

This paper studies the incentive effects of competition on individuals’ creative output, exploiting a unique field setting where creative activity and competition can be precisely measured and related: tournaments for the design of commercial logos and branding. Using image comparison tools to measure originality, I show that intensifying competition both creates and destroys incentives for innovation. While some competition is necessary to induce high-performing agents to develop original, untested designs over tweaking their existing work, heavy competition discourages effort of either kind. Theory suggests that these patterns are driven by the risk-return tradeoffs inherent to innovation. In the data, agents are most likely to produce original designs when in a horserace against exactly one other competitor of similar quality.

It is useful to begin with a definition: creativity is the act of producing ideas that are novel and appropriate to the goal at hand (Amabile 1996, Sternberg 2008). The paper opens with a simple model that provides a framework for understanding the economics of creative activity, particularly in a tournament setting, which both guides and rationalizes the empirical analysis that follows.¹ In this model, a principal seeks a new product design and solicits candidates from a pool of workers via a tournament, awarding a prize to the best entry. Workers enter designs in turns, and once entered, each submission’s quality is public knowledge. At each turn, workers must choose between developing an original design or tweaking a previous entry. The model suggests that competition increases workers’ incentives to produce original designs over tweaks – but it also shows that heavy competition depresses incentives to do either. Though intuitive, the model is useful in framing and interpreting empirical results throughout the rest of the paper.

The paper then turns to an empirical study of logo design competitions, drawing on a sample of contests from a popular online platform.² In these contests, a firm (“sponsor”) solicits custom designs from freelance designers (“players”), who compete for a winner-take-all prize. The contests in the sample offer prizes of a few hundred dollars and on average attract around 35 players and 100 designs. An important feature of this setting is that the sponsor can provide real-time feedback on players’ designs in the form of 1- to 5-star ratings. These ratings allow players to gauge the quality of their own work and the intensity of competition while the contest is underway. Most importantly, the dataset also includes the designs themselves, which

¹The model in this paper is related to Taylor (1995), Che and Gale (2003), Fullerton and McAfee (1999), and Terwiesch and Xu (2008) but differs in that it injects an explore-exploit dilemma into the agents’ choice set: whereas existing work models competing agents who must choose how much effort to exert, the agents in this paper must choose whether to build off of an old idea or try a new one, much like a choice between incremental versus radical innovation. The framework also has ties to recent work on tournaments with feedback (e.g., Ederer 2010), bandit problems in single-agent settings (Manso 2011), and models of competing firms’ choice over R&D project risk (Cabral 2003, Anderson and Cabral 2007).

²The empirical setting is conceptually similar to coding competitions studied by Boudreau et al. (2011), Boudreau et al. (2016), and Boudreau and Lakhani (2015), though the opportunity to measure originality is unique. Wooten and Ulrich (2013, 2014) have also studied graphic design competitions, focusing on the effects of visibility and feedback.

makes it possible to study creative choices over the course of a contest: I use image comparison algorithms similar to those used by commercial content-based image retrieval software (e.g., Google Image Search) to calculate similarity scores between pairs of images in a contest, which I then use to quantify the originality of each design relative to prior submissions by the same player and her competitors.

This setting presents a unique opportunity to directly observe creative production in the field. Though commercial advertising is important in its own right, the product development process observed here is similar to that in other domains where prototypes are created, tested, and refined. The nature of the setting enables a more detailed empirical study of this process, and its interaction with incentives, than is typically possible. The tournament format is especially germane: although the website advertises itself as a crowdsourcing platform, the contracting environment is fundamentally a request for proposals (RFP), a competitive mechanism widely used by firms and government agencies to procure new products or technologies – often over multiple rounds, with interim scoring, and typically with only the top bid rewarded.

The sponsors' ratings are critical in this paper as a source of variation in the information that both I and the players have about the state of the competition. Using these ratings, I am able to directly estimate a player's probability of winning, and the results establish that ratings are meaningful: the highest-rated design in a contest may not always win, but a five-star design increases a player's win probability as much as 10 four-star designs, 100 three-star designs, and nearly 2,000 one-star designs. Data on the time at which designs are entered by players and rated by sponsors makes it possible to establish what every participant knows at each point in time – and what they have yet to find out. The empirical strategy exploits naturally-occurring, quasi-random variation in the timing of sponsors' ratings and compares players' responses to information they observe at the time of design against that which is absent or not yet provided.

I find that competition has large effects on the content of players' submissions. Absent competition, positive feedback causes players to cut back sharply on originality: players with the top rating produce designs more than twice as similar to their previous entries than those with only low ratings. The effect is strongest when a player receives her first five-star rating – her next design will be a near replica of the highly-rated design – and attenuates at each rung down the ratings ladder. However, these effects are reversed by half or more when high-quality competition is present: competitive pressure counteracts this positive feedback, inducing players to produce more original designs. A battery of supporting analysis establishes that this result is econometrically identified and robust to alternative measures of the key variables.

Taken alone, these results suggest that competition unambiguously motivates creativity, but the analysis, and conclusion, presumes no outside option. In practice, players have a third option: they can stop bidding. Whether and when the participation constraint binds is its own question. Consistent with previous research (e.g., Baik 1994 or Brown 2011), I find that heavy competition discourages further investment. Empirically, high performers' tendency to produce original work is greatest when facing roughly 50-50 odds of winning – in other words, when neck-and-neck against one similar-quality competitor.

These findings contribute to a developing but mixed literature on the effects of competition on individual creative output: economists argue that competition can motivate the kind of risk-taking that is characteristic

of inventive activity (e.g. Cabral 2003, Anderson and Cabral 2007), yet many psychologists argue that high-powered incentives and other extrinsic pressures stifle creativity by crowding out intrinsic motivation (see Amabile and Hennessey 2010 for a review) or by causing agents to choke (Ariely et al. 2009). Lab-based studies of creativity are as mixed as the theory (e.g., Eisenberger and Rhoades 2001, Charness and Grieco 2014, Erat and Gneezy 2016, Bradler et al. 2016), in part due to differences in measurement and experimental design. Missing from the creativity literature is the added nuance that competition is not strictly a binary condition but rather can vary in intensity across treatments – and as this paper shows, the effects hinge crucially on the intensity of competition, as well as the existence of an outside option.

The evidence that creativity can be elicited with balanced competition has substantive implications for managers in creative industries and for the procurement practices of all organizations. Many practitioners appear to subscribe to the aforementioned intrinsic motivation theory of creativity endorsed by social psychologists, which holds that extrinsic motivators are counterproductive and is regularly communicated in the Harvard Business Review (e.g., Florida and Goodnight 2005, Amabile and Khair 2008, Amabile and Kramer 2012) and other business press. While intrinsic motivation is valuable, the results of this paper demonstrate that high-powered incentives can be effective at motivating creativity, if properly managed. The results also provide lessons for organizers of innovation prize competitions and other competitive procurement mechanisms for innovation (e.g., RFPs) on managing the intensity of competition.

The paper also makes a methodological contribution to the innovation literature. Due to data constraints, empirical research has historically measured innovation in terms of inputs (such as R&D spending) or outputs (patents), when innovation is at heart about the individual acts of discovery and invention that take place between. As a result, there is relatively little systematic, empirical evidence on the process of idea production. This paper is an effort to fill this gap, invoking new tools for content-based measurement of innovation and using them to study how ideas are developed and refined in response to incentives.

The paper is organized follows. Section 1 discusses related literatures in economics and social psychology and presents the model. Section 2 introduces the empirical setting and describes the identification strategy. Section 3 estimates the effects of competition on submissions’ originality. Section 4 presents the countervailing effects on participation. Section 5 discusses the implications of these results for management, policy, and future research on creativity and innovation and concludes.

1 Background: Creativity and Incentives

1.1 Existing Literature

Research on individual creativity has historically belonged to the realm of social psychology. The question of whether incentives enhance or impair creativity is itself the focus of a contentious, decades-old debate led by two schools of thought: one camp argues that incentives impair creativity by crowding out intrinsic motivation (Amabile 1996, Hennessey and Amabile 2010), whereas the other argues that incentives bolster

creativity, provided that creativity is explicitly what is being rewarded (Eisenberger and Cameron 1996). Scholars in each of these camps have written public rejoinders to the other (e.g., Eisenberger and Cameron 1998, Hennessey and Amabile 1998), while others have sought to develop and test more nuanced theories in an attempt to reconcile these arguments (e.g., Shalley and Oldham 1997).

The empirical literature on which these arguments are based in most cases invokes high-powered incentives (tournaments) in its experimental design. Despite dozens of experiments, the empirical evidence has been unable to clarify which of these positions is valid (Shalley et al. 2004). Different papers include different-sized rewards (which may or may not be valuable enough to overcome motivational crowd-out, to the extent it occurs), different subject pools (college students versus grade-school children), and inconsistencies in how performance is evaluated and what features of performance are rewarded: studies cited by the pro-incentives camp reward subjects for creativity, whereas studies cited by the anti-incentives camp evaluate creativity but often reward the best ideas. Experiments on both sides rely heavily on judges' assessments of creativity, which they are typically asked to score according to their own definitions.

Experimental economists have recently entered the literature, though often subject to the same limitations. Erat and Gneezy (2016) evaluate subjects' creativity in a puzzle-making task under piece-rate and competitive incentives and find that competition reduces creativity relative to an incentive-free baseline. Charness and Grieco (2014) in contrast find that high-powered incentives increase creativity in closed-ended creative tasks and have no effect on creativity in open-ended tasks. In both studies, creativity is scored by judges without guidance or a standardized definition, which leads to low inter-rater reliability. Rather than relying on subjective assessments, Bradler et al. (2016) study the effects of tournament incentives and gift exchange on creative output with an unusual uses task – where subjects are asked to think of productive uses for a common household object (e.g., a tin can), and creativity is measured by the statistical infrequency of each answer. In this case, the authors find that tournaments increase creative output relative to both gift exchange and an incentive-free baseline, though the empirical methodology makes it hard to distinguish an increase in originality (novel uses) from an increase in output alone (total uses).

This paper makes several important departures from this body of research. The logo design competitions studied here provide a field setting in which real creative professionals are competing for prizes of significantly greater value than observed in the existing lab-based studies. They also provide a setting where originality can be objectively measured with content-based assessment. Additionally, in contrast to much of the literature, it is not creativity per se that is being rewarded, but rather product quality: as in most product development settings, creativity here is a means towards an end, rather than an end in and of itself. Most importantly, however, this paper studies competition as a continuously-varying rather than binary treatment. In practice, competition is not a uniform condition, and the fact that the implementation of competitive incentives varies across the previously-cited studies might perhaps even explain their divergence.

At the heart of this paper is a set of empirical results on the originality of submissions into the sampled tournaments. To the extent that being creative is risky and its outcome uncertain, as the model below will propose, the paper is also connected to the economics literature on choices over risk in competition. Cabral

(2003) and Anderson and Cabral (2007) show that in theory, laggards will take actions with higher-variance outcomes, consistent with the intuition of needing a big hit to catch up or, in the extreme, of having nothing to lose. Similar behavior has been observed among investment fund managers, who increase fund volatility after a mid-year review which reveals them trailing their peers' performance (Brown et al. 1996) or when trailing the market (Chevalier and Ellsion 1997). In additional related work, Genakos and Pagliero (2012) study the choice over how much weight to attempt across rounds of dynamic weightlifting tournaments, and interpret the decision as a choice over risk. The authors find that whereas moderate laggards increase risk, distant laggards reduce risk – a result at odds with the existing theoretical literature and the evidence from this paper, which indicate that more distant laggards prefer greater risk, conditional on participating. The interpretation, however, may be limited by the difficulty of empirically distinguishing a choice of risk from a commitment to a specified level of effort in the weightlifting context.

1.2 Theoretical Framework

The preceding literature explains creativity and its motives primarily through narrative psychological constructs, rather than economic forces. Yet creativity can also be interpreted as an economic behavior, insofar as it involves a choice over uncertainty. This section demonstrates how this idea can be operationalized in a relatively simple tournament model whose features resemble the empirical setting. The results both guide the empirical analysis and offer a framework for interpreting the evidence that follows.

Suppose a risk-neutral principal seeks a product design. Because R&D outcomes are uncertain and difficult to value, the principal cannot contract directly on performance. It instead sponsors a tournament to solicit prototypes from J risk-neutral players, who enter designs sequentially and receive immediate, public feedback on their quality. Each design can be either original or adapted from the blueprints of previous entries; players who choose to continue working on a given design at their next turn can re-use the blueprint to create variants, with the original version remaining in contention. At each turn, the player must decide whether to continue investing and if so, whether to produce an original design or tweak an earlier submission. At the end of the tournament, the sponsor awards a winner-take-all prize P to its favorite entry.

Let each design be characterized by latent value ν_{jt} , which only the sponsor observes:

$$\nu_{jt} = \ln(\beta_{jt}) + \varepsilon_{jt}, \quad \varepsilon_{jt} \sim \text{i.i.d. Type-I E.V.} \quad (1)$$

where j indexes players and t indexes designs. In this model, β_{jt} represents the design's quality, which is revealed by the sponsor's feedback, and the latent value is a function of revealed quality and a i.i.d. random shock, which reflects idiosyncracies in the winner selection process. To hone intuition, further suppose each player enters at most two designs. Player j 's probability of winning is then:

$$Pr(\text{player } j \text{ wins}) = \frac{\beta_{j1} + \beta_{j2}}{\beta_{j1} + \beta_{j2} + \sum_{k \neq j} (\beta_{k1} + \beta_{k2})} = \frac{\beta_{j1} + \beta_{j2}}{\beta_{j1} + \beta_{j2} + \mu_j} \quad (2)$$

where $\mu_j \equiv \sum_{k \neq j} (\beta_{k1} + \beta_{k2})$ is the competition that player j faces in the contest. This function is concave in the player's own quality and decreasing in the quality of her competition.

Every player's first design in the contest is inherently novel, and at their subsequent turn, they have three options: they can exploit (tweak, or adapt) the existing design, explore (experiment with) a radically different design, or abandon the contest altogether. To elaborate on each option:

1. **Exploitation** costs $c > 0$ and yields a design of the same quality, resulting in a second-round design with $\beta_{j2} = \beta_{j1}$ and increasing the player's probability of winning accordingly.
2. **Exploration** costs $d \geq c$ and can yield a high- or low-quality design. With probability q , exploration will yield a high-quality design with $\beta_{j2}^H = \alpha \beta_{j1}$; with probability $(1 - q)$ it will yield a low-quality design with $\beta_{j2}^L = \frac{1}{\alpha} \beta_{j1}$, where $\alpha \geq 1$ is the exogenous degree of exploration.³
3. **Abandonment** is costless: the player can abstain from further investment. Doing so leaves the player's probability of winning unchanged, as her previous work remains in contention.

In this context, feedback has three effects: it informs each player about her first design's quality, influences her second design, and reveals the level of competition she faces. Players use this information to decide (i) whether to continue participating and (ii) whether to do so by exploring a new design or re-using a previous one, which is a choice over which kind of effort to exert: creative or rote.

Conditions for Exploration

To further simplify notation, let $F(\beta_2) = F(\beta_2 | \beta_1, \mu)$ denote a player's probability of winning as a function of β_2 , given β_1 and μ (omitting the j subscript). For a player to produce an original design, she must prefer doing so over both exploiting an existing design (Eq. 3.1) and abandoning (Eq. 4.1):

$$\underbrace{[qF(\beta_2^H) + (1 - q)F(\beta_2^L)] \cdot P - d}_{E[\pi|\text{explore}]} > \underbrace{F(\beta_1) \cdot P - c}_{E[\pi|\text{exploit}]} \quad (3.1)$$

$$\underbrace{[qF(\beta_2^H) + (1 - q)F(\beta_2^L)] \cdot P - d}_{E[\pi|\text{explore}]} > \underbrace{F(0) \cdot P}_{E[\pi|\text{abandon}]} \quad (4.1)$$

These conditions can be rearranged and be written as follows:

$$qF(\beta_2^H) + (1 - q)F(\beta_2^L) - F(\beta_1) > \frac{d - c}{P} \quad (3.2)$$

$$qF(\beta_2^H) + (1 - q)F(\beta_2^L) - F(0) > \frac{d}{P} \quad (4.2)$$

³For the purposes of this illustrative model, I treat α as fixed. If α were endogenous and costless, the player's optimal α would be infinite, since the exploration upside would then be unlimited and the downside bounded at zero. A natural extension would be to endogenize α and allow exploration costs $d(\cdot)$ or the probability of a successful outcome $q(\cdot)$ to vary with it. Such a model is considerably more difficult to study and beyond the scope of this paper.

In words, the probability gains from exploration over exploitation or abandonment must exceed the difference in cost, normalized by the prize. If the difference in the cost of exploration versus exploitation is small relative to the prize, as it likely is in the data, the choice between them reduces to a question of which strategy yields the greater increase in the player’s probability of winning.

Effects of Competition

This modeling infrastructure leads directly to the focal propositions. The first proposition states that when μ_j is high, exploration has greater expected benefits than exploitation, whereas when μ_j is low, the reverse holds. The second proposition states that as μ_j grows large, the benefits of a second design decline to zero. Because effort is costly, players are thus likely to abandon when competition is severe.

Proposition 1. *Suppose $q \in \{\frac{1}{1+\alpha}, \frac{1}{2}\}$. Then, there exists a μ^* such that for all $\mu_j < \mu^*$,*

$$\underbrace{F(\beta_{j1})}_{E[\text{Pr(Win)}|\text{exploit}]} > \underbrace{[qF(\beta_{j2}^H) + (1-q)F(\beta_{j2}^L)]}_{E[\text{Pr(Win)}|\text{explore}]}$$

and for all $\mu_j > \mu^*$,

$$\underbrace{[qF(\beta_{j2}^H) + (1-q)F(\beta_{j2}^L)]}_{E[\text{Pr(Win)}|\text{explore}]} > \underbrace{F(\beta_{j1})}_{E[\text{Pr(Win)}|\text{exploit}]}$$

Proposition 2. *The returns to a player’s second design decline to zero as $\mu_j \rightarrow \infty$.*

Proofs are provided in Appendix A. The necessary condition for competition to motivate exploration is that $q \in \{\frac{1}{1+\alpha}, \frac{1}{2}\}$, which implies that the probability of successful outcomes is low enough to be improbable, but holds if and only if original submissions are in expectation higher-quality than tweaks, such that exploration is not only high-risk, but also high-return (see appendix). When this is the case, the first proposition shows that competition can provoke creativity as a strategic response, a result which reproduces the findings of Cabral (2003) and Anderson and Cabral (2007) on choices over risk, but in a structure more closely matching the empirical setting. The second proposition shows, however, that large performance differences can also discourage effort. Though intuitive, it is a reminder that a participation constraint applies: in contrast to bandit models (e.g., Manso 2011) or previous models of choices over risk in competition (e.g., Cabral 2003), agents in this setting incur costs and may abstain from effort.

2 Setting, Data, and Identification

I collect a randomly-drawn sample of 122 logo design contests from a widely-used online platform to study how creative behavior responds to competition.⁴ The platform from which the data were collected hosts

⁴The sample consists of all logo design contests with public bidding that began the week of Sept. 3-9, 2013 and every three weeks thereafter through the week of Nov. 5-11, 2013, excluding those with multiple prizes or mid-contest rule changes such as prize increases or deadline extensions. Appendix B describes the sampling procedures in greater detail.

hundreds of contests each week in several categories of commercial graphic design, including logos, business cards, t-shirts, product packaging, book/magazine covers, website/app mockups, and others. Logo design is the modal design category on this platform and is thus a natural choice for analysis. A firm’s choice of logo is also nontrivial, since it is the defining feature of its brand, which can be one of the firm’s most valuable assets and is how consumers will recognize and remember the firm for years to come.

In these contests, a firm (the sponsor; typically a small business or non-profit organization) solicits custom designs from freelance designers (players) in exchange for a fixed prize awarded to its favorite entry. The sponsor publishes a project brief which describes its business, its customers, and what it likes and seeks to communicate with its logo; specifies the prize structure; sets a deadline for submissions; and opens the contest to competition. While the contest is active, players can enter (and withdraw) as many designs as they want, at any time they want, and sponsors can provide players with private, real-time feedback on their submissions in the form of 1- to 5-star ratings and written commentary. Players see a gallery of competing designs and the distribution of ratings on these designs, but not the ratings on specific competing designs. Copyright is enforced.⁵ At the end of the contest, the sponsor picks the winning design and receives the design files and full rights to their use. The platform then transfers payment to the winner.

For each contest in the sample, I observe the project brief, which includes a project title and description, the sponsor’s industry, and any specific elements that must be included in the logo; the contest’s start and end dates; the prize amount; and whether the prize is committed (the sponsor may retain the option of not awarding the prize to any entries if none are to its liking). While multiple prizes are possible, the sample is restricted to contests with a single, winner-take-all prize. I also observe every submitted design, the identity of the designer, his or her history on the platform, the time at which the design was entered, the rating it received (if any), the time at which the rating was given, and whether it won the contest. I also observe when players withdraw designs from the competition, but I assume withdrawn entries remain in contention, as sponsors can request that any withdrawn design be reinstated. Since I do not observe written feedback, I assume the content of written commentary is fully summarized by the rating.⁶

The player identifiers allow me to track players’ activity over the course of each contest. I use the precise timing information to reconstruct the state of the contest at the time each design is submitted. For every design, I calculate the number of preceding designs in the contest of each rating. I do so both in terms of the feedback available (i.e., observed) at the time of submission as well as the feedback eventually provided. To account for the lags required to produce a design, I define preceding designs to be those entered at least one hour prior to a given design, and I similarly require that feedback be provided at least one hour prior to the given design’s submission to be considered observed at the time it is made.

⁵Though players can see competing designs, the site requires that all designs be original and enforces copyright protections. Players have numerous opportunities to report violations if they believe a design to be copied or otherwise misused. Violators are permanently banned from the site. The site also prohibits the use of stock art and has a strict policy on the submission of overused design concepts. These mechanisms appear to be effective at limiting abuses.

⁶One of the threats to identification throughout the empirical analysis is that the estimated effects of ratings may be confounded by unobserved, written feedback: what seems to be a response to a rating could be a reaction to explicit direction provided by the sponsor that I do not observe. This possibility is evaluated in more depth in Appendix D, which shows that instructive comments are (i) rare, and (ii) not disproportionately given to 5-star designs relative to 3- or 4-star designs, such that they cannot be responsible for the differences in behavior that will be found in later tables.

The dataset also includes the designs themselves. Recall that creativity bears the formal definition of the act of producing ideas that are novel and relevant to the goal at hand (Amabile 1996, Sternberg 2008). To operationalize this definition, I invoke image comparison algorithms commonly used in content-based image retrieval software (similar to Google Image’s Search by Image feature) to quantify the originality of each design entered into a contest relative to preceding designs by the same and other players. I use two mathematically distinct procedures to compute similarity scores for image pairs, one of which is a preferred measure (the “perceptual hash” score) and the other of which is reserved for robustness checks (the “difference hash” score). Appendix B explains how they work. Each one takes a pair of digital images as inputs, summarizes them in terms of a specific, structural feature, and returns a similarity index in $[0,1]$, with a value of one indicating a perfect match and a zero indicating total dissimilarity. This index effectively measures the absolute correlation of two images’ underlying structure, reflecting similarities or differences in the basic shapes, outlines, and other elements that define the image.

To make this discussion concrete, the inset below demonstrates an example application. The figure shows three designs, entered in the order shown, by the same player in a logo design competition that is similar to those in the sample, although not necessarily from the same platform.⁷ The first two logos have some features in common (they both use a circular frame and are presented against a similar backdrop), but they also have some stark differences. The perceptual hash algorithm gives them a similarity score of 0.31, and the difference hash algorithm scores them 0.51. The latter two logos are more alike, and though differences remain, they are now more subtle and mostly limited to the choice of font. The perceptual hash algorithm gives these logos a similarity score of 0.71, and the difference hash scores them 0.89.

Illustration of image comparison algorithms

Notes: Figure shows three logos entered in order by a single player in a single contest. The perceptual hash algorithm calculates a similarity score of 0.313 for logos (1) and (2) and a score of 0.711 for (2) and (3). The difference hash algorithm calculates similarity scores of 0.508 for (1) and (2) and 0.891 for (2) and (3).

For each design in a contest, I compute its maximal similarity to previous designs in the same contest by the same player. Subtracting this value from one yields an index of originality between 0 and 1, which can be interpreted as an empirical counterpart to the parameter $1/\alpha$ in the model. In the empirical analysis, I use measures of similarity to a player’s *highest-rated* previous submissions, rather than all of her prior submissions, but since players tend to re-use only their highest-rated work, these two measures are highly correlated in practice (with $\rho=0.9$ under either algorithm).

⁷To keep the platform from which the sample was collected anonymous, I omit identifying information.

Creativity can manifest in other ways. For example, players sometimes create and enter several designs at once, and when doing so they can make each one similar to or distinct from the others. To capture this phenomenon, I define “batches” of proximate designs entered into the same contest by a single player and compute the maximum intra-batch similarity as a measure of creativity in batched work. Two designs are proximate if they are entered within 15 minutes of each other, and a batch is a set of designs in which every design in the set is proximate to another in the same set. Intra-batch similarity is arguably closer to a true measure of experimentation, reflecting players’ tendency to try minor variants of the same concept versus multiple concepts over a short period of time.

These measures are not without drawbacks or immune to debate. One drawback is that algorithmic comparisons require substantial dimensionality reduction and thus provide only a coarse comparison between designs based on a select set of features. Concerns on this front are mitigated by the fact that the empirical results throughout the paper are similar in sign, significance, and magnitude under two distinct algorithms. In addition, coarse comparisons will be sufficient for detecting designs that are plainly tweaks to earlier work versus those that are not, which is the margin that matters most for this paper. One might also question how well the algorithms emulate human perception, but the example provided above assuages this concern, as do other examples in Appendix B, which discusses these issues in detail.

2.1 Characteristics of the Sample

The average contest in the data lasts eight days, offers a \$250 prize, and attracts 96 designs from 33 players (Table 1). On average, 64 percent of designs are rated; less than three receive the top rating.

[Table 1 about here]

Among rated designs, the median and modal rating is three stars (Table 2). Though fewer than four percent of rated designs receive a 5-star rating, over 40 percent of all winning designs are rated five stars, suggesting that these ratings convey substantial information about a design’s quality and odds of success.⁸ The website also provides formal guidance on the meaning of each star rating, which generates consistency in their interpretation and use across different sponsors and contests.

[Table 2 about here]

Table 3 characterizes the similarity measures. For each design in the sample, we can compute its maximal similarity to previous designs by the same player, the highest-rated previous designs by the same player, and the highest-rated previous designs by that player’s competitors (in the same contest). For every design batch, I calculate the maximal similarity of any two designs in that batch. Note that the analysis of intra-batch similarity is restricted to batches that are not missing any image files.

⁸Another 33 percent of winning designs are rated 4 stars, and 24 percent are unrated.

[Table 3 about here]

The designs themselves are available for 96 percent of submissions in the sample. The table shows that new entries are on average more similar to that player’s own designs than her competitors’ designs, and that designs in the same batch tend to be more similar to each other than to previous designs by even the same player. But these averages mask more important patterns at the extremes. At the upper decile, designs can be very similar to previous work by the same player (≈ 0.75 under the perceptual hash algorithm) or to other designs in the same batch (0.91), but even the designs most similar to competing work are not all that similar (0.27). At the lower end, designs can be original by all of these measures.

Are ratings meaningful? Estimating the empirical success function

A player’s probability of winning can be directly estimated as a function of her ratings, using the win-lose outcomes of each design in a large sample of contests from this platform, which I borrow from Gross (2016).⁹ Recall from Section 1 that a design’s latent value is a function of its rating and an i.i.d. extreme value error. In the data, there are five possible ratings. This latent value can be flexibly specified with fixed effects for each rating, and the success function then be estimated as a conditional logit model. To formalize, let R_{ijk} denote the rating on design i by player j in contest k , and (in a slight abuse of notation) let $R_{ijk} = \emptyset$ when design ijk is unrated. The value of each design, ν_{ijk} , can be written as follows:

$$\nu_{ijk} = \gamma_0 \mathbb{1}(R_{ijk} = \emptyset) + \gamma_1 \mathbb{1}(R_{ijk} = 1) + \dots + \gamma_5 \mathbb{1}(R_{ijk} = 5) + \varepsilon_{ijk} \equiv \psi_{ijk} + \varepsilon_{ijk}$$

The details of the estimation are provided in Appendix C, but here we can summarize the results. The fixed effects are monotonically increasing in the rating and precisely estimated. Only a 5-star design is on average preferred to the outside option. To produce the same increase in a player’s win probability as generated by a five-star design, a player would need 12 four-star designs, 137 three-star designs, or nearly 2,000 one-star designs. The magnitudes of these differences imply that for players with a top rating, competitive pressure effectively only comes from other top-rated designs. As a measure of fit, the odds-on favorite wins almost half of all contests in the sample. These results collectively establish that this simple model fits the data relatively well and in an intuitive way, indicating that ratings provide considerable information about a player’s probability of winning – and although players do not observe the ratings on specific competing designs, knowledge of the distribution of competitors’ ratings makes it possible for players to invoke a simple model such as the one estimated here in their decision-making.

⁹Estimating the success function requires a larger sample of winners, and thus contests, than are in the primary sample of this paper. As Appendix C shows, the sample in Gross (2016) contains >4,000 contests, is empirically comparable, and includes all of the same variables except for the images themselves – sufficient for the exercise.

2.2 Empirical Methods and Identification

I exploit variation in the level and timing of the sponsor’s ratings to estimate the effects of competition on players’ creative choices. With timestamps on all activity, I can determine exactly what a player knows at each point in time about their own and their competitors’ performance and identify the effects of what is known at the time of design. Identification is thus achieved by harnessing variation in the *information* that players possess about the state of the competition. Note that although this information is broadcast through sponsor ratings, the results fundamentally reflect the effects of market structure on players’ choices, rather than feedback per se, as players are responding to the competitive conditions. In this context, the feedback process is merely how the state of the competition becomes known.¹⁰

Concretely, the analysis compares the actions of players who have received the top rating or who know they face top-rated competition against those who do not – whether it is because no prior designs will be given the top rating, or because these ratings have simply not yet been administered. The identifying assumption is that there are no omitted factors correlated with observed feedback that also affect choices. This assumption is supported by two pieces of evidence. First, the arrival of ratings is unpredictable, such that the set of ratings observed at any point in time is effectively random: sponsors are erratic, and it is difficult to know exactly when or how often a sponsor will log onto the site to rate new entries, much less any single design. More importantly, players’ choices are uncorrelated with ratings that were unobserved at the time, including forthcoming ratings and ratings on specific competing submissions.¹¹

To establish that feedback provision is unpredictable, I explore the relationship between feedback lags and ratings. In concept, sponsors may be quicker to rate the designs they like the most, to keep these players engaged and improving their work, in which case players might be able to infer their eventual ratings from the time elapsed without any feedback. Empirical assessment of this question (Appendix D) confirms that this is not the case: whether measured in hours or as a percent of the total contest duration, the lag between when a design is entered and rated is unrelated to its rating. The probability that a rated design was rated before versus after the contest ends is similarly uncorrelated with the rating granted.

Evidence that choices are uncorrelated with unobserved feedback is presented in Section 3. Following the presentation of the focal results, I show that the relationship between the similarity measures and forthcoming ratings is indistinguishable from zero. Though not reported here, I also examine players’ tendency to imitate highly-rated competing designs and find no such patterns, likely because players simply do not know which designs are highly rated (and thus which ones to imitate). The results suggest that players respond only to information available at the time of design and that there are no omitted factors correlated with this information that would confound its estimated effects.

¹⁰To see this point more clearly, note that without performance differences, there would be no variation in ratings to explain behavior. Moreover, if discretionary feedback were eliminated and instead quality always immediately revealed (a possibility supported in the data, though not common in practice), the basic approach to estimation would not change.

¹¹Though this setting may seem like a natural opportunity for a controlled experiment, the variation of interest is in the 5-star ratings, which are sufficiently rare that a controlled intervention would require either unrealistic manipulation or an infeasibly large sample. I therefore exploit naturally-occurring variation for this study.

3 Competition and Creativity

Figure 1 provides a first-cut, visual preview of this section’s results. The figure shows distribution of designs’ maximal similarity to previous submissions by the same player, conditioning on that player’s best rating and whether top-rated competition was present at the time. Recall that low similarity scores indicate that the submission is substantively original, while high scores indicate that it is a variant on a prior entry. Submissions from players with high ratings are significantly more similar to their prior work, with the effects largest for those with a 5-star rating (bottom-right panel). However, these very same players are more likely to enter original designs when facing top-rated competition than in its absence.

[Figure 1 about here]

The richness of the field setting makes it possible to evaluate whether competition affects high-performing players’ tendency to enter more original work in a variety of ways. The estimating equation in this part of the paper is as follows, with variants estimated throughout this section:

$$\begin{aligned}
 \text{Similarity}_{ijk} = & \beta_0 + \sum_{r=2}^5 \beta_r \cdot \mathbb{1}(\bar{R}_{ijk} = r) \\
 & + \beta_{5c} \cdot \mathbb{1}(\bar{R}_{ijk} = 5) \cdot \mathbb{1}(\bar{R}_{-ijk} = 5) \\
 & + \beta_{5p} \cdot \mathbb{1}(\bar{R}_{ijk} = 5) \cdot P_k \\
 & + \beta_c \cdot \mathbb{1}(\bar{R}_{-ijk} = 5) + \beta_p \cdot P_k \\
 & + T_{ijk}\lambda + X_{ijk}\theta + \zeta_k + \varphi_j + \varepsilon_{ijk}
 \end{aligned}$$

where i indexes designs by player j in contest k . The variables are as follows: Similarity_{ijk} is the maximal similarity of design ijk to the highest-rated preceding designs by player j in contest k ; $\bar{R}_{ijk}$ is the highest rating received by player j in contest k prior to design ijk ; $\bar{R}_{-ijk}$ is the highest rating received by player j ’s competitors prior to design ijk ; P_k is the prize in contest k , in units of \$100s; T_{ijk} is fraction of the contest elapsed when design ijk is entered; X_{ijk} consists of other design-level controls, including counts of previous designs by the same player and by competing players, and the number of days remaining in the contest; and ζ_k and φ_j are contest and player fixed effects. Specifications without contest fixed effects include the prize P_k as a standalone explanatory variable to identify the interaction. Standard errors throughout are clustered by player to account for any within-player correlation in the error term, though the results are robust to (and more conservative than) clustering by contest-player or by contest.

It may be helpful to provide a roadmap to this part of the analysis in advance. In the first set of regressions, I estimate the specification above. The second set estimates a specification in first differences, replacing the dependent variable with the change in similarity to previously-rated designs and independent variables with indicators for changes in highest ratings. The third set studies within-batch similarity. The fourth set tests the identifying assumption that players do not act on forthcoming ratings. The fifth set explores whether

the effects of competition are general to players with lower ratings when no 5-star ratings have been granted in a contest. Additional robustness checks are provided in the appendix.

3.1 Similarity of new designs to a player’s previous designs

I begin by examining players’ tendency to enter original versus derivative designs. Table 4 provides estimates from regressions of the maximal similarity of each design to the highest-rated preceding designs by the same player on indicators for the highest rating received. All specifications interact the indicator for the top rating with (i) the prize (in \$100s) and (ii) a variable indicating the presence of top-rated competition, and control for the fraction of the contest elapsed. Column (1) presents a baseline with no fixed effects or other controls. Columns (2) and (3) add fixed effects for contests and players, respectively, and Column (4) includes both. Column (5) additionally controls for the number of days remaining in the contest and the number of prior submissions by the same player as well as competing players.

[Table 4 about here]

Similar patterns are observed across all specifications. In the regression with both fixed effects and controls (Column 5), we see that players with the top rating enter designs that are 0.36 points, or over one full standard deviation, more similar to their previous work than players who have only low ratings or no ratings. Roughly half of this effect is reversed by the presence of top-rated competition, with this counteracting effect significant at the one percent level. When a player’s highest rating is four stars, her new designs are on average around 0.1 points more similar to previous work. This effect further attenuates as the best observed rating declines until it is indistinguishable from zero at a best rating of two stars, with all such differences statistically significant. High-rated competition is not seen to have an effect on the originality of these lower performers, who are already unlikely to reuse their low-rated submissions.

The latter regressions in Table 4 use contest and player fixed effects to control for other factors that are either common to all players within a contest or across all contests for a given player, but they do not control for factors that are constant for a given player within specific contests, as doing so leaves too little variation to identify the focal effects. Such factors could nevertheless be confounding, such as if players who continue participating in different competitive conditions are systematically more or less likely to enter original designs in that contest. The estimates in the previous tables additionally mask potential heterogeneity in players’ reactions to competitive conditions over the course of a contest.

Table 5 addresses these concerns with a model in first differences. The dependent variable here is the *change* in designs’ similarity to the player’s best previously-rated work. This variable can take values in $[-1,1]$, where a value of 0 indicates that the given design is as similar to the player’s best preceding design as was the last one she entered; a value of 1 indicates that the player transitioned fully from innovating to recycling; and a value of -1, the converse. The independent variables are changes in indicators for the highest rating the player has received, interacting the indicator for the top rating with the prize and the presence of top-rated

competition. I estimate this model with the same configurations of contest fixed effects, player fixed effects, and controls to account for other potential reasons why players’ tendency for originality changes over time, though the results are not statistically different across these specifications.

[Table 5 about here]

The results provide even stronger evidence of how competition affects creative choices and are statistically and quantitatively identical across specifications. A player who receives her first 5-star rating will typically then enter a near replica: the similarity increases by 0.9 points, or over *three* standard deviations, relative to players with low ratings. Top-rated competition again reverses roughly half of this effect, with the difference significant at the one percent level. Given their magnitudes, these effects will be plainly visible to the naked eye (the inset in Section 2 gives an example of what they would look like in practice). The effects of a new best rating of four, three, or two stars attenuate monotonically, similar to earlier results, and high-rated competition is not seen to have an effect on low performers.¹²

Table 6 tests the effects of competition in a different place: within batches of two or more designs entered by a given player, in a given contest, at once or in rapid succession. Recall that when entering multiple designs at once, players can make them similar to or different from each other. Observations here are submission batches, and the dependent variable is the maximal similarity of any two designs in the batch, estimated as a function of the ratings observed at the time of submission.

[Table 6 about here]

The results indicate that maximal within-batch similarity declines 0.3 points, or one standard deviation, for players with the top rating rating who face top-rated competition, relative to those who do not. The effect is insensitive to inclusion of controls (Columns 2 and 4) or weighting batches by their size (Columns 3 and 4), albeit significant only at the 10 percent level. High-rated players with competition are thus more likely to produce original designs not only across batches but also within them.

The appendix provides robustness checks and supplementary analysis. To confirm that these patterns are not an artifact of the perceptual hash algorithm, Appendix E re-estimates the regressions in the preceding tables using the difference hash algorithm to calculate similarity scores. The results are both statistically and quantitatively similar. In Appendix F, I split out the effects of competition by the number of top-rated competing designs, finding no statistical differences between the effects of one versus more than one: all effects of competition on originality are achieved by one high-quality competitor.

This latter result is especially important for ruling out an information-based story. In particular, the presence of other 5-star ratings might indicate that the sponsor has diverse preferences, and that unique designs have

¹²Interestingly, these regressions also find that new recipients of the top rating can also be induced to try new designs with larger prizes. The theory suggests a natural explanation: large prizes moderate the influence of costs in players’ decision-making. If original designs are more costly (take more time or effort) than tweaks, they may only be worth doing when the prize is large. This is particularly the case for players with highly-rated work in the contest.

higher likelihood of being well-received than one might otherwise believe. If this were the case, then originality should continue to increase as 5-star competitors are revealed. That this is not the case suggests that the effect is in fact the result of variation in incentives from competition.

In unreported tests, I also look for effects of 5-star competition on players with only 4-star designs, and find attenuated effects that are negative but not significantly different from zero. I also explore the effects of prize commitment, since the sponsor’s outside option of not awarding the prize is itself a competing alternative. The effect of prize commitment is not statistically different from zero. I similarly test for effects of four-star competition on players with five-star designs, finding none. These results reinforce the earlier evidence that competition effectively comes from other designs with the top rating.

3.2 Placebo test: Similarity to a player’s not-yet-rated designs

The identifying assumptions require that players are not acting on information that correlates with feedback but is unobserved in the data. As a simple validation exercise, the regressions in Table 7 perform a placebo test of whether originality is related to impending feedback. If an omitted determinant of creative choices is correlated with ratings and biasing the results, then it would appear as if originality responds to forthcoming ratings. If the identifying assumptions hold, we should see only zeros.

[Table 7 about here]

The specification in Column (1) regresses a design’s maximal similarity to the player’s best designs that have not yet been *but will eventually be* rated on indicators for the ratings they later receive. I find no evidence that designs’ originality is related to forthcoming ratings. Because a given design’s similarity to an earlier, unrated design can be incidental if both are tweaks on a third design, Column (2) adds controls for similarity to the best already-rated design. Column (3) allows these controls to vary with the rating received. As a final check, I isolate the similarity to the unrated design that cannot be explained by similarity to the third design in the form of a residual, and in Column (4) I regress these residuals on the same independent variables. In all cases, I find no evidence that players systematically tweak designs with higher forthcoming ratings. Choices are only correlated with ratings observed in advance.¹³

3.3 Generality: 4-on-4 Competition

The model suggests that similar dynamics should arise for players with 4-star ratings facing 4-star competition when no higher ratings are granted, since only relative performance matters – though this result may only arise towards the end of a contest, when the absence of higher ratings approaches finality. Table 8 tests this

¹³For more evidence that information drives behavior, in unreported results I test whether players imitate high-rated competitors – which they should not be able to identify, because players do not observe which competing designs received which rating. I find no evidence that imitation occurs in any systematic way. These results are consistent with the summary statistics in Table 3, where we see that imitation of competitors is uncommon.

prediction, regressing designs’ similarity to the player’s prior entries on indicators for their highest rating, restricting to submissions made before any 5-star designs were in play.

[Table 8 about here]

Column (1) includes all designs in the sample that meet this condition. Column (2) restricts to submissions in the second half of a contest; Column (3), to the final quarter. The table shows similar patterns for 4-on-4 competition that strengthen over the course of a contest, though the sample in the most restrictive condition is sufficiently small that standard errors do not allow us to rule out no effect.¹⁴ In addition to extending the generality of the main findings, the results reinforce the evidence that the observed behavior is a response to incentives, rather than to information about sponsors’ preferences.

4 Effects on Participation

The analysis thus far conditions on continued participation: the unit of observation is a submission, and the outcome is its originality. However, players can also stop making submissions if they perceive the returns to effort to be low. This outside option is present in many real-world competitive settings, and it distinguishes the setting of this paper from much of the existing literature on creativity, innovation, and high-powered incentives, where agents are effectively locked in to participating.

To incorporate the outside option into the empirics, I discretize outcomes and model each submission as a choice between three options: (i) entering a tweak and remaining active (“tweak”), (ii) entering an original design and remaining active (“original”), or (iii) entering any design and refraining from further submissions (“abandon”). Although the precise moment that a player decides to stop investing effort is not observable, we can use inactivity as a proxy. The unit of observation thus remains an individual submission, but I now categorize designs as original or as a tweak on the basis of discrete cutoffs for similarity scores, and I identify designs which are its creator’s final submission into each contest.¹⁵

The multinomial choice framework is necessary because the tradeoffs between three unordered options (tweak, original, or abandon) cannot be evaluated in a linear model. For this exercise, I classify a design as a tweak if its similarity to any earlier design by the same player is 0.7 or higher and original if its maximal similarity to previous designs by that player is 0.3 or lower. Designs with intermediate similarity scores are omitted from the exercise, as the player’s intentions are ambiguous in the intermediate range. Each action a in this choice set is assumed to have latent utility u_{ijk}^a for submission i by player j in contest k . Much like the

¹⁴In the data, contests rarely reach the final quarter with no players having received a 5-star rating and only one player having received a 4-star rating. As a result, there is only a small sample off which the focal coefficients (on the 4-star rating and its interaction with 4-star competition) in Column (3) can be estimated, and the estimates are imprecise.

¹⁵Note that this measure cannot distinguish a player who stops competing immediately after their final submission from one who waits for more information but later stops competing without additional entries. Because the end result is the same, the distinction is not critical for the purposes of this paper, as both behaviors will be influenced by information available at the time of the final submission. Anecdotally, according to some designers on this platform, it is often the case that players will enter their final design knowing it is their final design and not look back.

previous specifications, I model this latent utility as a function of the player’s own ratings, interacting the indicator for the top rating with indicators for 1, 2, and 3+ top-rated competitors, the fraction of the contest transpired, the number of days remaining, and a logit error term:

$$\begin{aligned}
u_{ijk}^a = & \beta_0^a + \sum_{r=1}^5 \beta_r^a \cdot \mathbb{1}(\bar{R}_{ijk} = r) \\
& + \beta_{5,1}^a \cdot \mathbb{1}(\bar{R}_{ijk} = 5) \mathbb{1}(N_{ijk} = 1) \\
& + \beta_{5,2}^a \cdot \mathbb{1}(\bar{R}_{ijk} = 5) \mathbb{1}(N_{ijk} = 2) \\
& + \beta_{5,3}^a \cdot \mathbb{1}(\bar{R}_{ijk} = 5) \mathbb{1}(N_{ijk} \geq 3) \\
& + \gamma_1^a \cdot \mathbb{1}(N_{ijk} = 1) + \gamma_2^a \cdot \mathbb{1}(N_{ijk} = 2) + \gamma_3^a \cdot \mathbb{1}(N_{ijk} \geq 3) \\
& + T_{ijk} \lambda^a + X_{ijk} \theta^a + \varepsilon_{ijk}^a, \quad \varepsilon_{ijk}^a \sim \text{i.i.d. Type-I E.V.}
\end{aligned}$$

where N_{ijk} is the number of top-rated competing designs at the time submission ijk is made, X_{ijk} controls solely for the number of days remaining in the contest, and the other variables are defined as before. Controlling for the fraction of the contest transpired and number of days remaining is especially important here, as abandonment is mechanically more likely to be observed later in a contest.

I estimate the parameters by maximum likelihood using observed behavior. I then use the results to compute the probability that a player with the top rating takes each of the three actions near the end of a contest, and to evaluate how these probabilities vary as the number of top-rated competitors increases from zero to three or more. These probabilities are shown in Figure 2. Panel A plots the probability that the player tweaks; Panel B, that she enters an original design; and Panel C, that she does either but then quits. The bars around each point provide the associated 95 percent confidence intervals.

[Figure 2 about here]

The probability that a high-performer tweaks and remains active (Panel A) peaks at 46 percent when there are no 5-star competitors and is significantly lower with non-zero competition, with all differences significant at the one percent level. The probability that the player produces an original design (Panel B) peaks at 52 percent with one 5-star competitor and is significantly lower with zero, two, or three 5-star competitors (differences against zero and three significant at the one percent level; difference against two significant at the ten percent level). Panel C shows that the probability of abandonment increases monotonically in the level of competition, approaching 70 percent with three or more competitors.

These patterns can also be demonstrated by other means. A simple alternative is to model the same discrete choice as a function of a player’s contemporaneous probability of winning (which can be computed using the conditional logit estimates described above) rather than directly as a function of ratings. Figure 3 provides estimated choice probabilities under this specification.

[Figure 3 about here]

In Figure 3, the probability that a player tweaks (Panel A) is maximized at 70 percent when she is a strong favorite and declines monotonically to zero with her odds of winning. The probability that the player produces an original design (Panel B) follows a distinct and highly significant inverted-U pattern, peaking at approximately a one-half odds of winning. Finally, the probability that she abandons (Panel C) increases from zero to around 80 percent as her odds of winning fall to zero. Table 9 demonstrates the source of these patterns, showing the mean utility of each action at different win probabilities, with abandonment normalized to zero. Standard errors are shown in parentheses, and the table marks (with a box) the highest-utility option at each competitive condition. As players' win probability declines: tweaks, original designs, and abandonment (in sequence) provide the greatest mean utility.

[Table 9 about here]

The evidence above is an important reminder that the participation constraint can bind when competition grows severe. Taken together, the evidence reveals that incentives for creativity are greatest with balanced competition: too little, and high-performers lack incentive to develop new ideas; too much, and agents stop investing effort altogether. In the data, it appears that creative effort is most attractive to high-rated players when faced off against exactly one high-rated competitor.

5 Implications and Conclusion

Within this sample of commercial logo design contests, I thus find that high-powered incentives have nuanced, multidimensional effects on creativity: some competition is required to motivate high-performing agents to develop original, untested ideas over tweaking their earlier work, but heavy competition drives them to stop investing altogether. When the two effects are considered in tandem, the evidence indicates that the likelihood that an agent produces original work is greatest in the presence of one competitor of similar ability. The results can be rationalized by a model in which creativity is inherently risky, and in which creative effort involves a choice over risk. As such, the paper ties together literatures in the social psychology of creativity and the economics of tournament competition, and provides new evidence on how competition shapes the intensity and direction of individuals' creative production.

The results have direct implications for the use of incentives as a tool for promoting creativity. The foremost lesson for managers is that competition can motivate creativity in professional settings, provided it is balanced. In designing contracts for creative workers, managers ought thus consider incentives for high-quality work relative to that of peers or colleagues, in addition to the more traditional strategy of establishing a work environment with intrinsic motivators such as freedom, flexibility, and challenge. Note that the reward need not be pecuniary: the same intuition applies when workers value recognition or status. Another advantage of the dynamic tournament incentive structure here is that it incorporates tolerance for failure, allowing players to recover from unsuccessful exploration, which has been shown to be an important feature of contracts for motivating innovation (e.g., Manso 2011, Ederer and Manso 2013).

In practice, this ‘Goldilocks’ level of competition may be difficult to achieve, let alone determine, and finding it would likely require experimentation with the mechanism itself by a principal in another setting. In this paper, the presence of one high-quality competitor was found to be sufficient to induce another high-quality player to produce original designs. A natural conjecture for other settings may be that one competitor of similar ability is enough to elicit creativity, while the presence of more than a few such competitors would be more harmful than helpful for motivating creative production – though the precise thresholds may also depend on other features of the setting, such as the prize distribution.

The results are also relevant to innovation procurement practices, particularly as governments, private foundations, and firms increasingly contract for R&D through prizes and institutionalize prize competition.¹⁶ Yet the applications are more general than R&D prizes alone: as earlier discussion explains, the mechanism in this paper is fundamentally an RFP, a standard contracting device used by firms and government agencies to solicit designs or prototypes of new products, systems, and technologies, with a prize or production contract awarded to the preferred submission. These competitions often take place over multiple rounds, with performance scored between, much like the contests studied here.

Caution is nonetheless warranted in drawing external inference to other procurement settings, as the product being procured in this paper (a logo) is relatively simple, and proposals are heavily tailored to each client. Another potential challenge to external validity is the absence of objective evaluation criteria: the ratings and winner selection are inherently at the sponsor’s subjective discretion. Yet in many RFPs, evaluation criteria similarly leave room for subjective judgments or are otherwise opaque to participants. More importantly, the defining feature of the R&D problem is not the ambiguity or clarity of the evaluation criteria, but rather the uncertainty around how any given product design will perform until it is tested and its performance is revealed. This uncertainty is present in all competitive R&D settings.

The final contribution is more methodological in its nature: this paper introduces new tools for measuring innovation in terms of its content. Whereas most recent attempts at content-based analysis of innovation have focused on textual analysis of patents, this paper demonstrates that even unpatentable ideas can be quantified, and it exploits a data-rich setting to study how ideas are developed and refined in response to competition. Many other questions about individual creativity and the process of innovation remain open, and this paper provides an example of how this agenda can be pursued.

¹⁶For example, the U.S. federal government now operates a platform (Challenge.gov) where agencies can seek solutions to both technical and non-technical problems from the public, with hundreds of active competitions and prizes ranging from status only (non-pecuniary) to tens of million dollars. Similar platforms (e.g., Innocentive) are available to organizations outside of the public sector. See Williams (2012) for a review of the literature on R&D prizes.

References

- Amabile, Teresa M.**, *Creativity in Context*, Boulder: Westview Press, 1996.
- and **Mukti Khaire**, “Creativity and the Role of the Leader,” *Harvard Business Review*, 2008, *October*.
- and —, “Creativity and the Role of the Leader,” *Harvard Business Review*, 2012, *Published Online*.
- Anderson, Axel and Luis Cabral**, “Go for broke or play it safe? Dynamic competition with choice of variance,” *RAND Journal of Economics*, 2007, *38* (3), 593–609.
- Ariely, Dan, Uri Gneezy, George Loewenstein, and Nina Mazar**, “Large Stakes and Big Mistakes,” *Review of Economic Studies*, 2009, *76*, 451–469.
- Baik, Kyung Hwan**, “Effort Levels in Contests with Two Asymmetric Players,” *Southern Economic Journal*, 1994, pp. 367–378.
- Baye, Michael R. and Heidrun C. Hoppe**, “The Strategic Equivalence of Rent-seeking, Innovation, and Patent-race Games,” *Games and Economic Behavior*, 2003, *44* (2), 217–226.
- Boudreau, Kevin J. and Karim R. Lakhani**, “‘Open’ Disclosure of Innovations, Incentives and Follow-on Reuse: Theory on Processes of Cumulative Innovation and a Field Experiment in Computational Biology,” *Research Policy*, 2015, *44* (1).
- , —, and **Michael Menietti**, “Performance Responses To Competition Across Skill-Levels In Rank Order Tournaments: Field Evidence and Implications For Tournament Design,” *RAND Journal of Economics*, Forthcoming.
- , **Nicola Lacetera**, and **Karim R. Lakhani**, “Incentives and Problem Uncertainty in Innovation Contests: An empirical analysis,” *Management Science*, 2011, *57* (5), 843–863.
- Bradler, Christiane, Susanne Neckermann, and Arne Jonas Warnke**, “Incentivizing Creativity: a Large-Scale Experiment with Tournaments and Gifts,” 2016. Working Paper.
- Brown, Jennifer**, “Quitters Never Win: The (adverse) incentive effects of competing with superstars,” *Journal of Political Economy*, 2011, *119* (5), 982–1013.
- Brown, Keith C., W. V. Harlow, and Laura T. Starks**, “Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry,” *Journal of Finance*, 1996, *51* (1), 85–110.
- Cabral, Luis**, “R&D Competition When Firms Choose Variance,” *Journal of Economics & Management Strategy*, 2003, *12* (1), 139–150.
- Charness, Gary and Daniela Grieco**, “Creativity and Financial Incentives,” 2014. Working Paper.
- Che, Yeon-Koo and Ian Gale**, “Optimal Design of Research Contests,” *American Economic Review*, 2003, *93* (3), 646–671.
- Chevalier, Judy and Glenn Ellison**, “Risk Taking by Mutual Funds as a Response to Incentives,” *Journal of Political Economy*, 1997, *105* (6), 1167–1200.
- Ederer, Florian**, “Feedback and Motivation in Dynamic Tournaments,” *Journal of Economics & Management Strategy*, 2010, *19* (3), 733–769.

- **and Gustavo Manso**, “Is Pay-for-Performance Detrimental to Innovation?,” *Management Science*, 2013, 59 (7), 1496–1513.
- Eisenberger, Robert and Judy Cameron**, “Detrimental Effects of Reward: Reality or Myth?,” *American Psychologist*, 1996, 51 (11), 1153–1166.
- **and —**, “Reward, Intrinsic Interest, and Creativity: New Findings,” *American Psychologist*, 1998, June, 676–679.
- Erat, Sanjiv and Uri Gneezy**, “Incentives for creativity,” *Experimental Economics*, 2016, 19, 269–280.
- Florida, Richard and Jim Goodnight**, “Managing for Creativity,” *Harvard Business Review*, 2005, July-August.
- Fudenberg, Drew, Richard Gilbert, Joseph Stiglitz, and Jean Tirole**, “Preemption, Leapfrogging, and Competition in Patent Races,” *European Economic Review*, 1983, 22 (1), 3–31.
- Fullerton, Richard L. and R. Preston McAfee**, “Auctioning Entry into Tournaments,” *Journal of Political Economy*, 1999, 107 (3), 573–605.
- Genakos, Christos and Mario Pagliero**, “Interim Rank, Risk Taking, and Performance in Dynamic Tournaments,” *Journal of Political Economy*, 2012, 120 (4), 782–813.
- Gross, Daniel P.**, “Performance Feedback in Competitive Product Development,” 2016. Working Paper.
- Hennessey, Beth A. and Teresa M. Amabile**, “Creativity,” *Annual Review of Psychology*, 2010, 61, 569–598.
- **and Theresa M. Amabile**, “Reward, Intrinsic Motivation, and Creativity,” *American Psychologist*, 1998, June, 674–675.
- Manso, Gustavo**, “Motivating Innovation,” *The Journal of Finance*, 2011, 66 (5), 1823–1860.
- Shalley, Christina E. and Greg R. Oldham**, “Competition and Creative Performance: Effects of Competitor Presence and Visibility,” *Creativity Research Journal*, 1997, 10 (4), 337–345.
- **, Jing Zhou, and Greg R. Oldham**, “The Effects of Personal and Contextual Characteristics on Creativity: Where Should We Go from Here?,” *Journal of Management*, 2004, 30 (6), 933–958.
- Sternberg, Robert J.**, *Cognitive Psychology*, 5th ed., Belmont: Wadsworth, 2008.
- Taylor, Curtis R.**, “Digging for Golden Carrots: An analysis of research tournaments,” *The American Economic Review*, 1995, pp. 872–890.
- Terwiesch, Christian and Yi Xu**, “Innovation Contests, Open Innovation, and Multiagent Problem Solving,” *Management Science*, 2008, 54 (9), 1529–1543.
- Williams, Heidi**, “Innovation Inducement Prizes: Connecting Research to Policy,” *Journal of Policy Analysis and Management*, 2012, 31 (3), 752–776.
- Wooten, Joel and Karl Ulrich**, “The Impact of Visibility in Innovation Tournaments: Evidence from field experiments,” 2013. Working Paper.
- **and —**, “Idea Generation and the Role of Feedback: Evidence from field experiments with innovation tournaments,” 2014. Working Paper.

Figure 1: Distribution of originality, conditional on ratings and competition

Notes: Figure shows distribution of designs' similarity to prior entries, conditional on the player's highest rating and the presence of top-rated competition. Each observation is a design, and the plotted variable is that design's maximal similarity to any previous submission by the same player in the same contest, taking values in $[0,1]$, where a value of 1 indicates the design is identical to one of the player's earlier submissions.

Figure 2: Probability of tweaks, original designs, and abandonment as a function of 5-star competition

Notes: Figure plots the probability that a player who has at least one 5-star rating in a contest does one of the following on (and after) a given submission: tweaks and then enters more designs (Panel A), experiments and then enters more designs (Panel B), or stops investing in the contest (Panel C). The bars around each point provide the associated 95 percent confidence interval.

Figure 3: Probability of tweaks, original designs, and abandonment as a function of $\text{Pr}(\text{Win})$

Notes: Figure plots the probability that a player does one of the following on (and after) a given submission, as a function of their contemporaneous win probability: tweaks and then enters more designs (Panel A), experiments and then enters more designs (Panel B), or stops investing in the contest (Panel C). These probabilities are estimated as described in the text, and the bars around each point provide the associated 95 percent confidence interval.

Table 1: Characteristics of contests in the sample

Variable	N	Mean	SD	P25	P50	P75
Contest length (days)	122	8.52	3.20	7	7	11
Prize value (US\$)	122	247.57	84.92	200	200	225
No. of players	122	33.20	24.46	19	26	39
No. of designs	122	96.38	80.46	52	74	107
5-star designs	122	2.59	4.00	0	1	4
4-star designs	122	12.28	12.13	3	9	18
3-star designs	122	22.16	25.33	6	16	28
2-star designs	122	17.61	25.82	3	10	22
1-star designs	122	12.11	25.24	0	2	11
Unrated designs	122	29.62	31.43	7	19	40
Number rated	122	66.75	71.23	21	50	83
Fraction rated	122	0.64	0.30	0.4	0.7	0.9
Prize committed	122	0.56	0.50	0.0	1.0	1.0
Prize awarded	122	0.85	0.36	1.0	1.0	1.0

Notes: Table reports descriptive statistics for the contests. “Fraction rated” refers to the fraction of designs in each contest that gets rated. “Prize committed” indicates whether the contest prize is committed to be paid (vs. retractable). “Prize awarded” indicates whether the prize was awarded. The fraction of contests awarded subsumes the fraction committed, since committed prizes are always awarded.

Table 2: Distribution of ratings (rated designs only)

	1-star	2-star	3-star	4-star	5-star	Total
Count	1,478	2,149	2,703	1,498	316	8,144
Percent	18.15	26.39	33.19	18.39	3.88	100

Notes: Table tabulates rated designs by rating. 69.3 percent of designs in the sample are rated by sponsors on a 1-5 scale. The site provides guidance on the meaning of each rating, which introduces consistency in the interpretation of ratings across contests.

Table 3: Similarity to preceding designs by same player and competitors, and intra-batch similarity

Panel A. Using preferred algorithm: Perceptual Hash							
Variable	N	Mean	SD	P10	P50	P90	
Max. similarity to any of own preceding designs	5,075	0.32	0.27	0.05	0.22	0.77	
Max. similarity to best of own preceding designs	3,871	0.28	0.27	0.03	0.17	0.72	
Max. similarity to best of oth. preceding designs	9,709	0.14	0.1	0.04	0.13	0.27	
Maximum intra-batch similarity	1,987	0.45	0.32	0.05	0.41	0.91	
Panel B. Using alternative algorithm: Difference Hash							
Variable	N	Mean	SD	P10	P50	P90	
Max. similarity to any of own preceding designs	5,075	0.58	0.28	0.16	0.62	0.94	
Max. similarity to best of own preceding designs	3,871	0.52	0.3	0.09	0.54	0.93	
Max. similarity to best of oth. preceding designs	9,709	0.33	0.21	0.09	0.29	0.63	
Maximum intra-batch similarity	1,987	0.69	0.28	0.23	0.77	0.98	

Notes: Table reports summary statistics on designs’ similarity to previously entered designs (both own and competing). Pairwise similarity scores are calculated as described in the text and available for all designs whose digital image could be obtained (96% of entries; refer to the text for an explanation of missing images). The “best” preceding designs are those with the most positive feedback provided prior to the given design. Intra-batch similarity is calculated as the similarity of designs in a given batch to each other, where a design batch is defined to be a set of designs entered by a single player in which each design was entered within 15 minutes of another design in the set. This grouping captures players’ tendency to submit multiple designs at once, which are often similar with minor variations on a theme.

Table 4: Similarity to player's best previously-rated designs

	(1)	(2)	(3)	(4)	(5)
Player's prior best rating==5	0.440*** (0.102)	0.459*** (0.092)	0.260*** (0.097)	0.357*** (0.097)	0.362*** (0.102)
* 1+ competing 5-stars	-0.197*** (0.073)	-0.245*** (0.063)	-0.158** (0.061)	-0.206*** (0.070)	-0.208*** (0.071)
* prize value (\$100s)	-0.025 (0.025)	-0.015 (0.023)	0.005 (0.032)	-0.014 (0.031)	-0.018 (0.033)
Player's prior best rating==4	0.165*** (0.024)	0.160*** (0.022)	0.128*** (0.030)	0.121*** (0.031)	0.116*** (0.032)
Player's prior best rating==3	0.079*** (0.018)	0.077*** (0.018)	0.068** (0.028)	0.060** (0.028)	0.056** (0.028)
Player's prior best rating==2	0.044** (0.021)	0.044** (0.022)	0.023 (0.029)	0.026 (0.030)	0.024 (0.030)
One or more competing 5-stars	-0.020 (0.018)	0.009 (0.020)	-0.003 (0.022)	0.004 (0.023)	0.001 (0.024)
Prize value (\$100s)	-0.014* (0.007)		-0.010 (0.010)		
Pct. of contest elapsed	-0.030 (0.034)	-0.060* (0.032)	-0.010 (0.030)	-0.018 (0.034)	-0.103 (0.084)
Constant	0.238*** (0.039)	0.207*** (0.023)	0.235*** (0.044)	0.232*** (0.061)	0.303*** (0.093)
N	3871	3871	3871	3871	3871
R^2	0.07	0.20	0.48	0.53	0.53
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of a design's similarity to the highest-rated preceding entry by the same player, taking values in $[0,1]$, where a value of 1 indicates the design is identical to another. The mean value of this variable in the sample is 0.28 (s.d. 0.27). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table 5: Change in similarity to player's best previously-rated designs

	(1)	(2)	(3)	(4)	(5)
$\Delta(\text{Player's best rating}=5)$	0.861*** (0.162)	0.878*** (0.170)	0.928*** (0.203)	0.914*** (0.205)	0.924*** (0.205)
* 1+ competing 5-stars	-0.417*** (0.118)	-0.412*** (0.125)	-0.418*** (0.144)	-0.427*** (0.152)	-0.429*** (0.152)
* prize value (\$100s)	-0.092** (0.039)	-0.094** (0.039)	-0.115** (0.049)	-0.107** (0.047)	-0.110** (0.048)
$\Delta(\text{Player's best rating}=4)$	0.275*** (0.062)	0.282*** (0.065)	0.267*** (0.073)	0.276*** (0.079)	0.279*** (0.079)
$\Delta(\text{Player's best rating}=3)$	0.143*** (0.055)	0.151*** (0.058)	0.134** (0.065)	0.137** (0.069)	0.138** (0.069)
$\Delta(\text{Player's best rating}=2)$	0.079* (0.043)	0.082* (0.046)	0.063 (0.053)	0.059 (0.056)	0.059 (0.057)
One or more competing 5-stars	-0.003 (0.007)	-0.003 (0.015)	-0.003 (0.014)	0.004 (0.025)	0.003 (0.026)
Prize value (\$100s)	0.003 (0.002)		0.003 (0.008)		
Pct. of contest elapsed	0.015 (0.012)	0.009 (0.018)	0.017 (0.024)	0.004 (0.030)	-0.048 (0.074)
Constant	-0.029*** (0.010)	-0.017* (0.010)	-0.031 (0.029)	0.063 (0.093)	0.105 (0.108)
N	2694	2694	2694	2694	2694
R^2	0.03	0.05	0.11	0.14	0.14
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of the *change* in designs' similarity to the highest-rated preceding entry by the same player, taking values in [-1,1], where a value of 0 indicates that the player's current design is as similar to her best preceding design as was her previous design, and a value of 1 indicates that the player transitioned fully from innovating to recycling (and a value of -1, the converse). The mean value of this variable in the sample is 0.00 (s.d. 0.23). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table 6: Similarity to other designs in the same submission batch

	Unweighted		Weighted	
	(1)	(2)	(3)	(4)
Player's prior best rating==5	0.223 (0.311)	0.238 (0.304)	0.254 (0.304)	0.285 (0.296)
* 1+ competing 5-stars	-0.308* (0.163)	-0.305* (0.162)	-0.303* (0.171)	-0.295* (0.168)
* prize value (\$100s)	0.016 (0.099)	0.015 (0.097)	0.010 (0.096)	0.009 (0.093)
Player's prior best rating==4	0.054* (0.032)	0.065* (0.037)	0.064** (0.032)	0.086** (0.038)
Player's prior best rating==3	0.055 (0.035)	0.062* (0.037)	0.052 (0.035)	0.065* (0.037)
Player's prior best rating==2	0.021 (0.050)	0.027 (0.051)	0.007 (0.047)	0.018 (0.047)
One or more competing 5-stars	0.025 (0.048)	0.027 (0.049)	0.027 (0.053)	0.027 (0.054)
Pct. of contest elapsed	-0.023 (0.049)	-0.093 (0.114)	-0.010 (0.050)	-0.056 (0.111)
Constant	0.400*** (0.066)	0.507*** (0.148)	0.391*** (0.060)	0.459*** (0.146)
N	1987	1987	1987	1987
R^2	0.57	0.57	0.58	0.58
Contest FEs	Yes	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes	Yes
Other Controls	No	Yes	No	Yes

Notes: Observations are design batches, which are defined to be a set of designs by a single player entered into a contest in close proximity (15 minutes). Dependent variable is a continuous measure of intra-batch similarity, taking values in $[0,1]$, where a value of 1 indicates that two designs in the batch are identical. The mean value of this variable in the sample is 0.45 (s.d. 0.32). Columns (3) and (4) weight the regressions by batch size. Columns (2) and (4) control for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table 7: Similarity to player’s best not-yet-rated designs (placebo test)

	Similarity to forthcoming			Residual
	(1)	(2)	(3)	(4)
Player’s best forthcoming rating==5	0.007 (0.169)	-0.084 (0.136)	-0.105 (0.151)	-0.113 (0.122)
* 1+ competing 5-stars	-0.094 (0.099)	0.032 (0.056)	0.027 (0.066)	0.035 (0.062)
* prize value (\$100s)	-0.003 (0.031)	0.015 (0.025)	0.021 (0.027)	0.018 (0.025)
Player’s best forthcoming rating==4	0.039 (0.066)	0.051 (0.096)	0.049 (0.094)	0.034 (0.095)
Player’s best forthcoming rating==3	0.080 (0.052)	0.049 (0.088)	0.051 (0.088)	0.036 (0.088)
Player’s best forthcoming rating==2	0.030 (0.049)	-0.010 (0.093)	-0.007 (0.094)	-0.014 (0.095)
One or more competing 5-stars	-0.080 (0.097)	-0.013 (0.110)	-0.010 (0.117)	-0.013 (0.119)
Pct. of contest elapsed	0.016 (0.242)	-0.502 (0.478)	-0.466 (0.462)	-0.468 (0.497)
Constant	0.217 (0.212)	0.556 (0.560)	0.569 (0.543)	0.398 (0.581)
N	1147	577	577	577
R^2	0.68	0.83	0.83	0.67
Contest FEs	Yes	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes

Notes: Table provides a placebo test of the effects of future feedback on originality. Observations are designs. Dependent variable in Columns (1) to (3) is a continuous measure of a design’s similarity to the best design that the player has previously entered that has yet to *but will eventually be* rated, taking values in $[0,1]$, where a value of 1 indicates that the two designs are identical. The mean value of this variable is 0.26 (s.d. 0.25). Under the identifying assumption that future feedback is unpredictable, current choices should be unrelated to forthcoming ratings. Note that a given design’s similarity to an earlier, unrated design can be incidental if they are both tweaks on a rated third design. To account for this possibility, Column (2) controls for the given and unrated designs’ similarity to the best previously-rated design. Column (3) allows these controls to vary with the highest rating previously received. Dependent variable in Column (4) is the residual from a regression of the dependent variable in the previous columns on these controls. These residuals will be the subset of a given design’s similarity to the unrated design that is not explained by jointly-occurring similarity to a third design. All columns control for days remaining and number of previous designs by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table 8: Similarity to any of player's previous designs: 4-vs-4

	(1)	(2)	(3)
Player's prior best rating==4	0.185*** (0.069)	0.382** (0.190)	0.564* (0.299)
* 1+ competing 4- or 5-stars	-0.121** (0.051)	-0.323* (0.181)	-0.315 (0.248)
* prize value (\$100s)	-0.011 (0.020)	-0.006 (0.033)	-0.045 (0.097)
Player's prior best rating==3	0.006 (0.023)	0.042 (0.041)	-0.029 (0.088)
Player's prior best rating==2	-0.016 (0.032)	-0.046 (0.057)	0.040 (0.097)
One or more competing 4- or 5-stars	0.060** (0.028)	0.076 (0.092)	0.091 (0.169)
Pct. of contest elapsed	-0.091 (0.119)	-0.638** (0.274)	0.828 (0.755)
Constant	0.457** (0.209)	0.948*** (0.317)	-0.320 (0.638)
N	2926	1557	879
R^2	0.52	0.60	0.67
Contest FEs	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Restriction	All	2nd half	4th qtr

Notes: Table shows the effects of 4-star feedback and competition on originality when no player has a 5-star rating. Observations are designs. Dependent variable is a continuous measure of a design's maximal similarity to previous entries in the same contest by the same player, taking values in $[0,1]$, where a value of 1 indicates the design is identical to another. All columns include contest and player fixed effects and control for the number of days remaining and number of previous designs entered by the player and her competitors. Columns (2) and (3) restrict the sample to submissions in the second half or fourth quarter of a contest, when the absence of 5-star ratings may be more meaningful and is increasingly likely to be final. Similarity scores in this table are calculated using a perceptual hash algorithm. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table 9: Latent-utility estimates for each action, as a function of $\Pr(\text{Win})$

	Tweak	Original	Abandon
Latent utility when $\Pr(\text{Win}) = 100\%$:	4.065 (1.113)	3.210 (0.775)	0.000 n.a.
Latent utility when $\Pr(\text{Win}) = 80\%$:	2.745 (1.011)	2.368 (0.687)	0.000 n.a.
Latent utility when $\Pr(\text{Win}) = 60\%$:	1.425 (0.909)	1.526 (0.599)	0.000 n.a.
Latent utility when $\Pr(\text{Win}) = 40\%$:	0.105 (0.806)	0.684 (0.511)	0.000 n.a.
Latent utility when $\Pr(\text{Win}) = 20\%$:	-1.215 (0.704)	-0.157 (0.422)	0.000 n.a.
Latent utility when $\Pr(\text{Win}) = 0\%$:	-2.535 (0.602)	-0.999 (0.334)	0.000 n.a.

Notes: Table shows latent-utility estimates from a choice model relating players' actions to their contemporaneous win probability, evaluated at six values of $\Pr(\text{Win})$. These estimates represent the latent utility of each action relative to the outside option of abandonment, which has utility normalized to zero. See text for discussion. Boxes identify action with greatest utility. Robust standard errors in parentheses.

Appendix for Online Publication

A Theoretical Background

Before proving the propositions of the paper, it is useful to establish an identity. Proposition 1 requires that $q \in \{\frac{1}{1+\alpha}, \frac{1}{2}\}$. In the text of the paper, it is stated that this condition implies that exploration has higher expected quality than exploitation. To see this, observe that:

$$\begin{aligned} q > \frac{1}{1+\alpha} &\implies q > \frac{\alpha-1}{\alpha^2-1} \implies q(\alpha^2-1) - (\alpha-1) > 0 \implies q(\alpha - \frac{1}{\alpha}) - (1 - \frac{1}{\alpha}) > 0 \\ &\implies q\alpha + (1-q)\frac{1}{\alpha} > 1 \implies \underbrace{q\alpha\beta_1 + (1-q)\frac{1}{\alpha}\beta_1}_{E[\beta_2|\text{Explore}]} > \underbrace{\beta_1}_{E[\beta_2|\text{Explore}]} \end{aligned}$$

Note that here, as in the proofs below, the j subscript is omitted to simplify notation.

Proposition 1: Suppose $q \in \{\frac{1}{1+\alpha}, \frac{1}{2}\}$. Then, there exists a μ^* such that for all $\mu_j < \mu^*$,

$$\underbrace{F(\beta_{j1})}_{E[Pr(\text{Win})|\text{exploit}]} > \underbrace{[qF(\beta_{j2}^H) + (1-q)F(\beta_{j2}^L)]}_{E[Pr(\text{Win})|\text{explore}]}$$

and for all $\mu_j > \mu^*$,

$$\underbrace{[qF(\beta_{j2}^H) + (1-q)F(\beta_{j2}^L)]}_{E[Pr(\text{Win})|\text{explore}]} > \underbrace{F(\beta_{j1})}_{E[Pr(\text{Win})|\text{exploit}]}$$

Proof:

To prove this statement, it is sufficient to show that (i) the difference in returns to exploration over exploitation is zero when $\mu = 0$, (ii) the first derivative of this function is negative when $\mu = 0$, and (iii) the function has exactly one positive, real root. These three conditions imply a function that is negative for low μ and positive for high μ . The proof proceeds in sequence.

Proof of (i): When $\mu = 0$,

$$\begin{aligned} &qF(\beta_2^H) + (1-q)F(\beta_2^L) - F(\beta_1) \\ &= q\left(\frac{(1+\alpha)\beta_1}{(1+\alpha)\beta_1+0}\right) + (1-q)\left(\frac{(1+\frac{1}{\alpha})\beta_1}{(1+\frac{1}{\alpha})\beta_1+0}\right) - \left(\frac{2\beta_1}{2\beta_1+0}\right) = q + (1-q) - 1 = 0 \end{aligned}$$

Proof of (ii): When $\mu = 0$,

$$\begin{aligned}
& \frac{\partial}{\partial \mu} [qF(\beta_2^H) + (1-q)F(\beta_2^L) - F(\beta_1)]_{\mu=0} \\
&= \frac{\partial}{\partial \mu} \left[q \left(\frac{(1+\alpha)\beta_1}{(1+\alpha)\beta_1 + \mu} \right) + (1-q) \left(\frac{(1+\frac{1}{\alpha})\beta_1}{(1+\frac{1}{\alpha})\beta_1 + \mu} \right) - \left(\frac{2\beta_1}{2\beta_1 + \mu} \right) \right]_{\mu=0} \\
&= q \left(\frac{-(1+\alpha)\beta_1}{((1+\alpha)\beta_1 + 0)^2} \right) + (1-q) \left(\frac{-(1+\frac{1}{\alpha})\beta_1}{((1+\frac{1}{\alpha})\beta_1 + 0)^2} \right) + \frac{2\beta_1}{(2\beta_1 + 0)^2} \\
&= q \left(\frac{-1}{(1+\alpha)\beta_1} \right) + (1-q) \left(\frac{-1}{(1+\frac{1}{\alpha})\beta_1} \right) + \frac{1}{2\beta_1} \\
&= \frac{1}{2(1+\alpha)(1+\frac{1}{\alpha})\beta_1} \cdot \left[q \cdot -2 \left(1 + \frac{1}{\alpha} \right) + (1-q) \cdot -2(1+\alpha) + (1+\alpha) \left(1 + \frac{1}{\alpha} \right) \right] \\
&= \frac{1}{2\alpha(1+\alpha)(1+\frac{1}{\alpha})\beta_1} \cdot \left[q \cdot -2(1+\alpha) + (1-q) \cdot -2(1+\alpha)\alpha + (1+\alpha)^2 \right] \\
&= \frac{1}{2\alpha(1+\frac{1}{\alpha})\beta_1} \cdot [-2q - 2(1-q)\alpha + (1+\alpha)] \\
&= \frac{1}{2\alpha(1+\frac{1}{\alpha})\beta_1} \cdot [-2q + 2q\alpha - \alpha + 1] \\
&= \frac{1}{2\alpha(1+\frac{1}{\alpha})\beta_1} \cdot [2q(\alpha - 1) - (\alpha - 1)] \\
&= \frac{\alpha - 1}{2\alpha(1+\frac{1}{\alpha})\beta_1} \cdot [2q - 1] < 0 \quad , \text{ because } \alpha > 1 \text{ and } q < 1/2
\end{aligned}$$

Proof of (iii): To find the roots, we must solve

$$\begin{aligned}
& qF(\beta_2^H) + (1-q)F(\beta_2^L) - F(\beta_1) \\
&= q \left(\frac{(1+\alpha)\beta_1}{(1+\alpha)\beta_1 + \mu} \right) + (1-q) \left(\frac{(1+\frac{1}{\alpha})\beta_1}{(1+\frac{1}{\alpha})\beta_1 + \mu} \right) - \left(\frac{2\beta_1}{2\beta_1 + \mu} \right) = 0
\end{aligned}$$

To simplify notation, let

$$\begin{aligned}
X &= (1+\alpha)\beta_1 \\
Y &= \left(1 + \frac{1}{\alpha} \right) \beta_1 \\
Z &= 2\beta_1
\end{aligned}$$

Multiplying out the denominators, we can then rewrite the equation as:

$$qX(Y + \mu)(Z + \mu) + (1-q)Y(X + \mu)(Z + \mu) - Z(X + \mu)(Y + \mu) = 0$$

Expanding and combining terms, the equation transforms to the following quadratic:

$$\mu^2(qX + (1-q)Y - Z) + \mu(qXZ + XY - qYZ - XZ) = 0$$

From this expression, it is immediately apparent that one of the two potential roots is $\mu = 0$ – as found in part (i) of this proof. To find the sign of the nonzero root, we can write:

$$\mu = -\frac{q(X-Y)Z + X(Y-Z)}{q(X-Y) + (Y-Z)}$$

Reverting back to the original notation,

$$\mu = -\frac{q \cdot \left((1+\alpha)\beta_1 - \left(1 + \frac{1}{\alpha}\right)\beta_1\right) \cdot 2\beta_1 + (1+\alpha)\beta_1 \cdot \left(\left(1 + \frac{1}{\alpha}\right)\beta_1 - 2\beta_1\right)}{q \left((1+\alpha)\beta_1 - \left(1 + \frac{1}{\alpha}\right)\beta_1\right) + \left(\left(1 + \frac{1}{\alpha}\right)\beta_1 - 2\beta_1\right)}$$

Which simplifies to:

$$\mu = -\frac{2q\left(\alpha - \frac{1}{\alpha}\right)\beta_1^2 + (1+\alpha)\left(\frac{1}{\alpha} - 1\right)\beta_1^2}{q\left(\alpha - \frac{1}{\alpha}\right)\beta_1 + \left(\frac{1}{\alpha} - 1\right)\beta_1}$$

Multiplying and dividing by α , we can write:

$$\mu = -\frac{2q(\alpha^2 - 1)\beta_1^2 - (\alpha^2 - 1)\beta_1^2}{q(\alpha^2 - 1)\beta_1 - (\alpha - 1)\beta_1}$$

Then, canceling out $(\alpha - 1)\beta_1$, we get:

$$\begin{aligned}\mu &= -\frac{2q(1+\alpha)\beta_1 - (1+\alpha)\beta_1}{q(1+\alpha) - 1} \\ &= \frac{(1-2q)(1+\alpha)\beta_1}{q(1+\alpha) - 1} > 0,\end{aligned}$$

because $q < \frac{1}{2}$ (which implies that the numerator is positive) and $q > \frac{1}{1+\alpha}$ (which implies that the denominator is positive). Thus, the remaining root is positive.

Proposition 2: The returns to a player's second design decline to zero as $\mu_j \rightarrow \infty$.

Proof:

The returns to the second submission can be measured in terms of the increase in win probability that it generates. For any value of β_2 (i.e., for both exploration and exploitation), as $\mu \rightarrow \infty$:

$$F(\beta_2) - F(\beta_1) = \left(\frac{\beta_1 + \beta_2}{\beta_1 + \beta_2 + \mu}\right) - \left(\frac{\beta_1}{\beta_1 + \mu}\right) \rightarrow \left(\frac{\beta_1 + \beta_2}{\mu}\right) - \left(\frac{\beta_1}{\mu}\right) \rightarrow \left(\frac{\beta_2}{\mu}\right) \rightarrow 0$$

B Dataset Construction

Data were collected on all logo design contests with open (i.e., public) bidding that launched the week of September 3 to 9, 2013, and every three weeks thereafter through the week of November 5 to 11, 2013. Conditional on open bidding, this sample is effectively randomly drawn. The sample used in the paper is further restricted to contests with a single, winner-take-all prize and with no mid-contest rule changes such as prize increases, deadline extensions, and early endings. The sample also excludes one contest that went dormant and resumed after several weeks, as well as a handful of contests whose sponsors simply stopped participating and were never heard from again. These restrictions cause 146 contests to be dropped from the sample. The final dataset includes 122 contests, 4,050 contest-players, and 11,758 designs.¹

To collect the data, I developed an automated script to scan these contests once an hour for new submissions, save a copy of each design for analysis, and record their owners' identity and performance history from a player profile. I successfully obtained the image files for 96 percent of designs in the final sample. The remaining designs were entered and withdrawn before they could be observed (recall that players can withdraw designs they have entered into a contest, though this option is rarely exercised and can be reversed at the request of a sponsor). All other data were automatically acquired at the conclusion of each contest, once the prize was awarded or the sponsor exercised its outside option of a refund.

B.1 Variables

The dataset includes information on the characteristics of contests, contest-players, and designs:

- Contest-level variables include: the contest sponsor, features of the project brief (title, description, sponsor industry, materials to be included in logo), start and end dates, the prize amount (and whether committed), and the number of players and designs of each rating.
- Contest-player-level variables include: the player's self-reported country, his/her experience in previous contests on the platform (number of contests and designs entered, contests won), and that player's participation and performance in the given contest.
- Design-level variables include: the design's owner, its submission time and order of entry, the feedback it received, the time at which this feedback was given, and whether it was eventually withdrawn. For designs with images acquired, I calculate originality using the procedures described in the next section. The majority of the analysis occurs at the design level.

Note that designs are occasionally re-rated: five percent of all rated designs are re-rated an average of 1.2 times each. Of these, 14 percent are given their original rating, and 83 percent are re-rated within 1 star of the original rating. I treat the first rating on each design to be the most informative, objective measure of quality, since research suggests first instincts tend to be most reliable and ratings revisions are likely made relative to other designs in the contest rather than an objective benchmark.

¹While 268 contests were originally sampled, only 122 of these survived the filters described above. The number of contests sampled was constrained by high costs of data collection, which was performed in real-time by making round-the-clock hourly scans for new activity without putting a heavy strain on the platform's servers.

B.2 Image Comparison Algorithms

This paper uses two distinct algorithms to calculate pairwise similarity scores. One is a perceptual hash algorithm, which creates a digital signature (hash) for each image from its lowest frequency content. As the name implies, a perceptual hash is designed to imitate human perception. The second algorithm is a difference hash, which creates the hash from pixel intensity gradients.

I implement the perceptual hash algorithm and calculate pairwise similarity scores using a variant of the procedure described by the Hacker Factor blog.² This requires six steps:

1. Resize each image to 32x32 pixels and convert to grayscale.
2. Compute the discrete cosine transform (DCT) of each image. The DCT is a widely-used transform in signal processing that expresses a finite sequence of data points as a linear combination of cosine functions oscillating at different frequencies. By isolating low frequency content, the DCT reduces a signal (in this case, an image) to its underlying structure. The DCT is broadly used in digital media compression, including MP3 and JPEG formats.
3. Retain the upper-left 16x16 DCT coefficients and calculate the average value, excluding first term.
4. Assign 1s to grid cells with above-average DCT coefficients, and 0s elsewhere.
5. Reshape to 256 bit string; this is the image's digital signature (hash).
6. Compute the Hamming distance between the two hashes and divide by 256.

The similarity score is obtained by subtracting this fraction from one. In a series of sensitivity tests, the perceptual hash algorithm was found to be strongly invariant to transformations in scale, aspect ratio, brightness, and contrast, albeit not rotation. As described, the algorithm will perceive two images that have inverted colors but are otherwise identical to be perfectly dissimilar. I make the algorithm robust to color inversion by comparing each image against the regular and inverted hash of its counterpart in the pair, taking the maximum similarity score, and rescaling so that the scores remain in $[0,1]$. The resulting score is approximately the absolute value correlation of two images' content.

I follow a similar procedure outlined by the same blog³ to implement the difference hash algorithm and calculate an alternative set of similarity scores for robustness checks:

1. Resize each image to 17x16 pixels and convert to grayscale.
2. Calculate horizontal gradient as the change in pixel intensity from left to right, returning a 16x16 grid (note: top to bottom is an equally valid alternative)
3. Assign 1s to grid cells with positive gradient, 0s to cells with negative gradient.
4. Reshape to 256 bit string; this is the image's digital signature (hash).
5. Compute the Hamming distance between the two hashes and divide by 256.

²See <http://www.hackerfactor.com/blog/archives/432-Looks-Like-It.html>.

³See <http://www.hackerfactor.com/blog/archives/529-Kind-of-Like-That.html>.

The similarity score is obtained by subtracting this fraction from one. In sensitivity tests, the difference hash algorithm was found to be highly invariant to transformations in scale and aspect ratio, potentially sensitive to changes in brightness and contrast, and very sensitive to rotation. I make the algorithm robust to color inversion using a procedure identical to that described for the perceptual hash.

Though the perceptual and difference hash algorithms are both conceptually and mathematically distinct, and the resulting similarity scores are only modestly correlated ($\rho = 0.38$), the empirical results of Section 3 are qualitatively and quantitatively similar under either algorithm. This consistency is reassurance that the patterns found are not simply an artifact of an arcane image processing algorithm; rather, they appear to be generated by the visual content of the images themselves.

B.3 Why use algorithms?

There are two advantages to using algorithms over human judges. The first is that the algorithms can be directed to evaluate specific features of an image and thus provide a consistent, objective measure of similarity, whereas individuals may be attuned to different features and can have different perceptions of similarity in practice (Tirilly et al. 2012). This argument is supported by a pilot study I attempted using Amazon Mechanical Turk, in which I asked participants to rate the similarity of pairs of images they were shown; the results (not provided here) were generally very noisy, except in cases of nearly identical images, in which case the respondents tended to agree that they were similar. The second advantage of algorithms is more obvious: they are cheap, taking only seconds to execute a comparison.

The obvious concern is that computer vision may not be comparable to human perception, in which case the measures are not relevant to human behavior. This concern is mitigated by the fact that they are used to approximate human perception in commercial software, as well as two pieces of evidence from the above studies: (i) when two images are similar, humans and algorithms tend to agree, and (ii) when two images are dissimilar, neither humans nor algorithms find them similar, but they may also disagree on the degree of dissimilarity. Thus, human and computer judgment tend to align within coarse categories, especially at extremes – which is the margin of variation that matters most for this paper.

The evidence of disagreement in subjects’ assessments of similarity nevertheless raises a deeper question: is it sensible to apply a uniform similarity measure at all? Squire and Pun (1997) find that *expert* subjects’ assessments of similarity tend to agree at all levels, and the designers in this paper could reasonably be classified as visual experts. Even so, it is reassuring that the results throughout this paper are similar in sign, significance, and magnitude under two fundamentally different algorithms, and thus emerge no matter which features we choose to focus on when measuring similarity.

B.4 How do the algorithms perform?

In my own experience browsing the designs in the dataset, images that look similar to the naked eye tend to have a high similarity score, particularly under the perceptual hash algorithm. But as Tirilly et al. (2012) show, similarity is in the eye of the beholder – particularly at intermediate levels and when it is being assessed by laypersons. Figure B.1 illustrates the performance of the algorithms for three logos entered in the order shown by the same player in one contest (not necessarily from the sampled platform):

Figure B.1: Performance of image comparison algorithms

Notes: Figure shows three logos entered in order by a single player in a single contest. The perceptual hash algorithm calculates a similarity score of 0.313 for logos (1) and (2) and a score of 0.711 for (2) and (3). The difference hash algorithm calculates similarity scores of 0.508 for (1) and (2) and 0.891 for (2) and (3).

The first two images have several features in common but also have some notable differences. Each is centered, defined by a circular frame with text underneath, and presented against a similar backdrop. However the content of the circular frame and the font of the text below are considerably different, and the first logo is in black and white while the second one is in color. The perceptual hash algorithm assigns these two logos a similarity score of 31 percent, while the difference hash gives them 51 percent.

In contrast, the second two images appear much more similar. They again have similar layouts, but now they share the same color assignments and the same content in the frame. Lesser differences remain, primarily with respect to the font style, but the logos appear broadly similar. The perceptual hash algorithm assigns these two logos a similarity score of 71 percent; the difference hash, 89 percent.

The algorithms thus pass the gut check in this example, which is not particularly unique: further examples using better-known brands are provided below. In light of this evidence, and the consistency of the paper's results, I believe that these algorithms provide empirically valid measures of originality.

Figure B.2: Volkswagen logo in 1937, 1967, 1995, 1999

Notes: Figure shows the evolution of Volkswagen logos since 1937. The perceptual hash algorithm calculates similarity scores of 0.055 for the 1937 and 1967 logos, 0.430 for the 1967 and 1995 logos, and 0.844 for the 1995 and 1999 logos. The difference hash algorithm calculates similarity scores of 0.195, 0.539, and 0.953, respectively.

Figure B.3: Microsoft Windows 95, XP, 7, and 8 logos

Notes: Figure shows a sequence of Windows logos. The perceptual hash algorithm calculates similarity scores of 0.195 for the Windows 95 and XP logos, 0.531 for the Windows XP and 7 logos, and 0.148 for the Windows 7 and 8 logos. The difference hash algorithm calculates similarity scores of 0.055, 0.563, and 0.117, respectively. The reason why the similarity of the Windows XP and 7 logos is not evaluated to be even higher is because the contrast generated by the latter's spotlight and shadow changes the structure of the image (for example, it changes the intensity gradient calculated by the difference hash algorithm).

Appendix References:

- [1] Tirilly, Pierre, Chunsheng Huang, Wooseob Jeong, Xiangming Mu, Iris Xie, and Jin Zhang. 2012. "Image Similarity as Assessed by Users: A Quantitative Study." *Proceedings of the American Society for Information Science and Technology*, 49(1), pp. 1-10.
- [2] Squire, David and Thierry Pun. 1997. "A Comparison of Human and Machine Assessments of Image Similarity for the Organization of Image Databases." *Proceedings of the Scandinavian Conference on Image Analysis*, Lappeenranta, Finland.

C Additional Contest Characteristics

To highlight some of the basic relationships in the contests on this platform, I reproduce a subset of results from Gross (2016), which studies a larger sample from the same setting. Table C.1 estimates the relationship between contest characteristics such as the prize or frequency of feedback and key outcomes, and Table C.2 estimates the relationship between a design’s rating and its probability of being selected.

Correlation of contest characteristics and outcomes

The estimates in Table C.1 suggest that an extra \$100 in prize value on average attracts an additional 13.3 players, 47.7 designs, and 0.1 designs per player and increases the odds that a retractable prize will be awarded by 1.6 percent at the mean of all covariates. There is only a modest incremental effect of committed prize dollars, likely because the vast majority of uncommitted prizes are awarded anyway. The effects of feedback are also powerful: a sponsor who rates a high fraction of the designs in the contest will typically see fewer players enter but receive more designs from the participating players and have a much higher probability of finding a design it likes enough to award the prize. The effect of full feedback (relative to no feedback) on the probability the prize is awarded is greater than that of a \$1000 increase in the prize – a more than quadrupling of the average and median prize in the sample.

Table C.1: Correlations of contest outcomes with their characteristics

	(1) Players	(2) Designs	(3) Designs/Player	(4) Awarded
Total Prize Value (\$100s)	13.314*** (0.713)	47.695*** (2.930)	0.050*** (0.016)	0.101*** (0.027)
Committed Value (\$100s)	2.469** (1.205)	8.674* (4.932)	0.038 (0.025)	
Fraction Rated	-7.321*** (0.813)	15.195*** (3.098)	1.026*** (0.041)	1.121*** (0.102)
Contest Length	0.537*** (0.073)	2.109*** (0.283)	0.013*** (0.004)	0.021** (0.010)
Words in Desc. (100s)	0.130 (0.092)	3.228*** (0.449)	0.063*** (0.006)	-0.143*** (0.013)
Attached Materials	-0.943*** (0.173)	-1.884*** (0.692)	0.048*** (0.013)	-0.021 (0.015)
Prize Committed	1.398 (3.559)	4.539 (14.552)	-0.007 (0.087)	
Constant	-2.445 (2.309)	-63.730*** (9.045)	1.916*** (0.072)	1.085*** (0.155)
N	4294	4294	4294	3298
R ²	0.57	0.54	0.22	

Notes: Table shows the estimated effect of contest attributes on overall participation and the probability that the prize is awarded, using the larger sample of Gross (2016). The final specification is estimated as a probit on contests without a committed prize. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Monthly fixed effects included but not shown. Robust SEs in parentheses.

Estimating the success function: details

Recall the following model from the text: let R_{ijk} denote the rating on design i by player j in contest k , and (in a slight abuse of notation) let $R_{ijk} = \emptyset$ when design ijk is unrated. The value of each design, ν_{ijk} , can then be written as follows:

$$\nu_{ijk} = \gamma_0 \mathbb{1}(R_{ijk} = \emptyset) + \gamma_1 \mathbb{1}(R_{ijk} = 1) + \dots + \gamma_5 \mathbb{1}(R_{ijk} = 5) + \varepsilon_{ijk} \equiv \psi_{ijk} + \varepsilon_{ijk} \quad (3)$$

As in the theoretical model, the sponsor is assumed to select as winner the design with the highest value. In estimating the γ parameters, each sponsor's choice set of designs is assumed to satisfy I.I.A.; in principle, the submission of a design of any rating in a given contest will reduce competing designs' chances of winning proportionally. For contests with an uncommitted prize, the choice set also includes an outside option of not awarding the prize, with value normalized to zero. Letting I_{jk} be the set of designs by player j in contest k , and I_k be the set of all designs in contest k , player j 's probability of winning is:

$$Pr(j \text{ wins } k) = \frac{\sum_{i \in I_{jk}} e^{\psi_{ijk}}}{\sum_{i \in I_k} e^{\psi_{ik}} + \mathbb{1}(\text{Uncommitted prize})}$$

I use the sample of 496,401 designs in 4,294 contests from Gross (2015) to estimate this model by maximum likelihood. The results are reproduced in Table C.2.

Table C.2: Conditional logit of win-lose outcomes on ratings

Model: Latent design value $\nu_{ijk} = \gamma_5 + \gamma_4 + \gamma_3 + \gamma_2 + \gamma_1 + \gamma_0 + \varepsilon_{ijk}$					
Fixed effect	Est.	S.E.	t-stat	Implied β (Appendix A)	
Rating==5	1.53	0.07	22.17	4.618	
Rating==4	-0.96	0.06	-15.35	0.383	
Rating==3	-3.39	0.08	-40.01	0.034	
Rating==2	-5.20	0.17	-30.16	0.006	
Rating==1	-6.02	0.28	-21.82	0.002	
No rating	-3.43	0.06	-55.35	0.032	

Notes: Table provides results from a conditional logit estimation of the win-lose outcome of each design as a function of its rating, using the larger sample of Gross (2016). Outside option is not awarding the prize, with utility normalized to zero. The results can be used to approximate a player's probability of winning a contest as a function of her ratings. As a measure of fit, the design predicted by this model as the odds-on favorite wins roughly 50 percent of contests. See Gross (2016) for further discussion.

Table C.3 below sheds more light on the source of the conditional logit estimates, which are difficult to interpret directly. The table shows a cross-tabulation of contests, by the highest rating granted (columns) and the rating of the winning design (rows). The table shows that sponsors typically select the highest-rated design as winner, especially when the highest rating is 4- or 5-stars, but sponsors also often select unrated designs or the outside option. Rarely are 1- or 2-star designs ever awarded.

Table C.3: Frequency of contests, by highest rating and winning rating

<i>Rating of winner</i>	<i>Highest rating in contest</i>						Total
	Unrated	1-star	2-star	3-star	4-star	5-star	
Not awarded	66	4	12	92	202	85	461
Unrated	142	5	10	59	347	276	839
1-star	.	.	.	3	6	5	14
2-star	.	.	3	11	16	8	38
3-star	.	.	.	43	146	53	242
4-star	.	.	.	.	836	379	1,215
5-star	.	.	.	.	.	1,485	1,485
Total	208	9	25	208	1,553	2,291	4,294

Notes: Table shows the frequency of contests in the Gross (2016) sample by the highest rating granted and the rating of the winning design.

Evidence that the samples are comparable

The dataset in Gross (2015) consists of nearly all logo design contests with open bidding completed on the platform between July 1, 2010 and June 30, 2012, excluding those with zero prizes, multiple prizes, mid-contest rule changes, or otherwise unusual behavior, and it includes nearly all of the same information as the sample in this paper – except for the designs themselves. Although this sample comes from a slightly earlier time period than the one in the present paper (which was collected in the fall of 2013), both cover periods well after the platform was created and its growth had begun to stabilize.

Table C.4 compares characteristics of contests in the two samples. The contests in the Gross (2015) sample period are on average slightly longer, offer larger prizes, and attract a bit more participation relative to the sample of the present paper, but otherwise, the two samples are similar on observables. These differences are mostly due to the presence of a handful of outlying large contests in the Gross (2015) data. Interestingly, although the total number of designs is on average higher in the Gross (2015) sample, the number of designs of each rating is on average the same; the difference in total designs is fully accounted for by an increase in unrated entries. The most notable difference between the two samples is in the fraction of contests with a committed prize (23 percent vs. 56 percent). This discrepancy is explained by the fact that prize commitment only became an option on the platform halfway through the Gross (2015) sample period. Interestingly, the fraction of contests awarded is nevertheless nearly the same in these two samples.

Tables C.5 and C.6 compare the distribution of ratings and batches in the two samples. The tables demonstrate that individual behavior is consistent across samples: sponsors assign each rating, and players enter designs, at roughly the same frequency. The main differences between the two samples are thus isolated to a handful of the overall contest characteristics highlighted in Table C.4.

Table C.4: Comparing Samples: Contest characteristics

	Gross (2016)	This paper
<i>Sample size</i>	<i>4,294</i>	<i>122</i>
Contest length (days)	9.15	8.52
Prize value (US\$)	295.22	247.57
No. of players	37.28	33.20
No. of designs	115.52	96.38
5-star designs	3.41	2.59
4-star designs	13.84	12.28
3-star designs	22.16	22.16
2-star designs	16.04	17.61
1-star designs	10.94	12.11
Unrated designs	49.14	29.62
Number rated	66.38	66.75
Fraction rated	0.56	0.64
Prize committed	0.23	0.56
Prize awarded	0.89	0.85

Table C.5: Comparing Samples: Distribution of ratings

	Gross (2016)	This paper
<i>Sample size</i>	<i>285,052</i>	<i>8,144</i>
1 star (in percent)	16.48	18.15
2 stars	24.16	26.39
3 stars	33.38	33.19
4 stars	20.84	18.39
5 stars	5.13	3.88
	100.00	100.00

Table C.6: Comparing Samples: Design batches by size of batch

	Gross (2016)	This paper
<i>Sample size</i>	<i>335,016</i>	<i>8,072</i>
1 design (in percent)	72.46	71.84
2 designs	17.04	18.62
3 designs	5.75	5.57
4 designs	2.50	2.19
5+ designs	2.25	1.77
	100.00	100.00

D Additional Support for Identification

D.1 Timing of feedback unrelated to rating granted

This subsection provides further support for the empirical strategy of estimating the effects of information about relative performance on individual choices. The paper shows that players do not behave in any ways consistent with their being able to forecast forthcoming feedback: behavior is uncorrelated with forthcoming (future) ratings on their previous entries. But we can also show that the timing of feedback is itself difficult to predict, and that it offers no information on performance.

To do so, in Table D.1 I take the subsample of all rated designs and regress the lag between the time of submission and time of feedback on indicators for the rating granted. The dependent variable in Column (1) is this lag, in hours; Column (2), the lag as a percent of contest duration; and Column (3), an indicator for whether the design was rated before the contest ended. All specifications account for contest and player fixed effects, include the standard controls used in the paper, and cluster standard errors by contest. Across all specifications, I find no relationship between feedback lags and ratings.

Table D.1: Correlation of feedback lags with rating granted

	(1) Lag (hours)	(2) Lag (pct. of contest)	(3) Rated before end?
Rating==5	1.473 (3.461)	0.007 (0.016)	-0.022 (0.031)
Rating==4	-2.453 (1.792)	-0.013* (0.008)	0.010 (0.020)
Rating==3	-0.759 (2.131)	-0.004 (0.009)	0.007 (0.015)
Rating==2	1.130 (1.922)	0.005 (0.008)	0.006 (0.011)
Constant	22.779** (11.009)	0.242*** (0.066)	1.353*** (0.122)
N	7388	7388	8144
R^2	0.45	0.48	0.63
Contest FEs	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes

Notes: Table illustrates tendency for designs of different ratings to be rated more or less quickly. The results suggest that sponsors are not quicker to rate their favorite designs. Dependent variable in Column (1) is the lag between submission and feedback, in hours; Column (2), this lag as a fraction of contest length; and Column (3), an indicator for whether a design receives feedback before the contest ends. All columns control for the time of entry, the number of previous designs entered by the given player and competitors, and contest and player fixed effects. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. SEs clustered by contest in parentheses.

D.2 Written comments are rare

A distinct threat to identification arises if ratings are accompanied by unobserved, written feedback, and these comments provide explicit instruction that generates the patterns found in this paper.

To evaluate this possibility, I draw a new sample of contests from this platform in which written feedback was visible to the public, seemingly by glitch/error. This sample consists of contests from early in the platform’s history (significantly preceding the sample period for this paper), but it can nevertheless shed light on the frequency and potential confounding effects of written comments.

Within this sample, sponsors provided written comments to fewer than 8 percent of submissions, though as Table D.2 shows, the frequency is substantially higher for designs rated 4- or 5-stars (20 percent) than for those with poor ratings. Comments take a range of flavors, with many echoing the rating given (e.g., “This is on the right track” or “Not what I’m looking for”), but some make a more explicit request or suggestion of content changes. Because the latter present the risk of a confound, I had individuals read every comment in this sample and determine whether the sponsor suggested specific changes.

Table D.2: Typical distribution of ratings on designs receiving comments

Rating	All designs in sample (1)	w/ Comments		w/ Instructive comments		
		Number (2)	% of all (2)/(1)	Number (3)	% of all (3)/(1)	% of comm. (3)/(2)
5 stars	719	141	19.61	90	12.52	63.83
4 stars	3,045	551	18.10	401	13.17	72.78
3 stars	5,334	762	14.29	553	10.37	72.57
2 stars	5,205	519	9.97	369	7.09	71.10
1 star	5,114	498	9.74	262	5.12	52.61
Unrated	25,753	1,066	4.14	665	2.58	62.38
Total	45,170	3,537	7.83	2,340	5.18	66.16

Notes: Table tabulates rated designs by rating, within the sample of designs for which written feedback was accessible. The sample was compiled under a separate data collection effort and comprises contests from early in the platform’s history (significantly preceding the sample period for this paper), but it is nonetheless informative. The table provides the frequency of each rating within this sample as a whole, then among: (i) designs receiving written comments, and (ii) those receiving written comments deemed instructive. Though higher-rated designs are more likely to receive written comments, the risk of a confounding effect arises only if these comments instruct players to make specific changes. The table shows that instructive comments are disproportionately given to designs with middling ratings – where this commentary may prove most constructive. The net effect is that 5-star designs receive instructive comments at roughly the same rate as 4- and 3-star designs.

On average, roughly two-thirds of comments are instructive in this way. Among designs receiving any comment, the ones most likely to receive an instructive comment are those with middling ratings: ideas that are incomplete but can be improved. As a result, on net we find that 5-, 4-, and 3-star designs all receive instructive comments at comparable rates (10-13 percent, see Table D.2).

Table D.3 formalizes this result with a regression, which provides standard errors. Column (1) regresses an indicator for whether a design in this sample received written feedback on indicators for its rating. Column (2) replaces the dependent variable with an indicator for a design receiving an instructive comment; Column (3) repeats this model for the subsample of designs that received any comment. In all cases, standard errors are clustered by contest. With standard errors in hand, it can be seen that 5-, 4-, and 3-star designs receive instructive comments at statistically similar rates (Table D.3, Column 2).

Table D.3: Raw frequencies of written feedback, by rating

	Commented on (1)	Instructive (2)	Instructive (3)
Rated 5-stars	0.196*** (0.026)	0.125*** (0.018)	0.638*** (0.049)
Rated 4-stars	0.181*** (0.014)	0.132*** (0.011)	0.728*** (0.025)
Rated 3-stars	0.143*** (0.013)	0.104*** (0.010)	0.726*** (0.023)
Rated 2-stars	0.100*** (0.013)	0.071*** (0.010)	0.711*** (0.024)
Rated 1-star	0.097*** (0.019)	0.051*** (0.010)	0.526*** (0.029)
Unrated	0.041*** (0.005)	0.026*** (0.003)	0.624*** (0.040)
N	45170	45170	3537
R^2	0.11	0.08	0.67
Sample	All designs	All designs	Commented

Notes: Table provides results from regressing an indicator for whether a design received any comment (Column 1) or received a constructive comment (Columns 2 and 3) on indicators for its rating. Column (3) restricts the sample to only designs that received any comment. The table reproduces the tabulations in Table D.2, but provides standard errors. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. SEs clustered by contest in parentheses.

Given that only around 10 percent of 5-star designs receive instructive comments, and that 5-star designs receive instruction at similar rates to 4- and 3-star designs, these comments are unlikely to be generating the results of this paper, which are large and increase dramatically from 3- to 5-stars. As a final check, in Table D.4 I expand the specifications in the previous table to match the format of the main results in the paper, interacting with competition and adding fixed effects and controls. The patterns in the previous table persist, and – importantly – we see no evidence that 5-star designs receive comments of any kind at different rates in the presence versus absence of top-rated competition.

Table D.4: Adding measures of competition, interactions, and controls

	Commented on (1)	Instructive (2)	Instructive (3)
Rated 5-stars	0.148*** (0.022)	0.100*** (0.018)	0.088 (0.067)
* 1+ competing 5-stars	-0.013 (0.027)	-0.004 (0.020)	-0.203 (0.273)
* prize value (\$100s)	0.002 (0.008)	-0.004 (0.006)	-0.013 (0.069)
Rated 4-stars	0.136*** (0.015)	0.108*** (0.011)	0.134*** (0.047)
Rated 3-stars	0.092*** (0.015)	0.077*** (0.012)	0.133*** (0.041)
Rated 2-stars	0.055*** (0.016)	0.047*** (0.012)	0.114** (0.046)
Rated 1-star	0.045** (0.020)	0.020** (0.010)	-0.066 (0.059)
One or more competing 5-stars	0.014 (0.016)	0.008 (0.011)	-0.039 (0.067)
Pct. of contest elapsed	-0.179*** (0.014)	-0.130*** (0.011)	-0.148** (0.061)
N	45170	45170	3537
R^2	0.26	0.20	0.45
Controls	Yes	Yes	Yes
Contest FEs	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes
Sample	All designs	All designs	Commented

Notes: Table provides results from regressing an indicator for whether a design received any comment (Column 1) or received a constructive comment (Columns 2 and 3) on indicators for its rating as well as an interaction for the presence of top-rated competition, as in the paper. Column (3) restricts the sample to only designs that received any comment. All columns include contest and player fixed effects and the standard set of controls. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. SEs clustered by contest in parentheses.

E Robustness Checks (1)

The following tables provide robustness checks on the focal results in Section 3 estimating the effects of competition on originality, using the difference hash algorithm. These estimates demonstrate that the results are not sensitive to the procedure used to calculate similarity scores. Table E.1 is a robustness check on Table 4; Table E.2, on Table 5; Table E.3, on Table 6; Table E.4, on Table 7; and Table E.5, on Table 8. The results in these tables are statistically and quantitatively similar to those in the body of the paper, despite the use of the computationally distinct similarity measure.

Table E.1: Similarity to player's best previously-rated designs (diff. hash)

	(1)	(2)	(3)	(4)	(5)
Player's prior best rating==5	0.329*** (0.117)	0.373*** (0.116)	0.201 (0.124)	0.246* (0.133)	0.242* (0.141)
* 1+ competing 5-stars	-0.118 (0.084)	-0.152** (0.074)	-0.126 (0.079)	-0.169** (0.086)	-0.177** (0.087)
* prize value (\$100s)	-0.021 (0.027)	-0.019 (0.028)	-0.006 (0.036)	-0.018 (0.038)	-0.024 (0.042)
Player's prior best rating==4	0.168*** (0.036)	0.171*** (0.035)	0.095** (0.039)	0.067* (0.039)	0.049 (0.041)
Player's prior best rating==3	0.124*** (0.033)	0.128*** (0.033)	0.069* (0.037)	0.044 (0.038)	0.033 (0.039)
Player's prior best rating==2	0.098*** (0.034)	0.104*** (0.034)	0.034 (0.039)	0.014 (0.040)	0.007 (0.040)
One or more competing 5-stars	-0.018 (0.020)	-0.028 (0.028)	-0.003 (0.024)	-0.010 (0.031)	-0.019 (0.032)
Prize value (\$100s)	-0.023** (0.009)		-0.016 (0.013)		
Pct. of contest elapsed	-0.055 (0.036)	-0.059* (0.034)	-0.012 (0.036)	-0.021 (0.038)	-0.057 (0.104)
Constant	0.496*** (0.052)	0.426*** (0.035)	0.506*** (0.055)	0.480*** (0.088)	0.491*** (0.119)
N	3871	3871	3871	3871	3871
R^2	0.04	0.15	0.48	0.53	0.53
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of a design's similarity to the highest-rated preceding entry by the same player, taking values in [0,1], where a value of 1 indicates the design is identical to another. The mean value of this variable in the sample is 0.52 (s.d. 0.30). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a difference hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table E.2: Change in similarity to player's best previously-rated designs (diff. hash)

	(1)	(2)	(3)	(4)	(5)
$\Delta(\text{Player's best rating}=5)$	0.635*** (0.203)	0.658*** (0.218)	0.682*** (0.257)	0.688** (0.268)	0.693*** (0.267)
* 1+ competing 5-stars	-0.352** (0.159)	-0.349** (0.175)	-0.374* (0.207)	-0.364* (0.218)	-0.368* (0.218)
* prize value (\$100s)	-0.046 (0.044)	-0.048 (0.046)	-0.060 (0.055)	-0.062 (0.057)	-0.064 (0.057)
$\Delta(\text{Player's best rating}=4)$	0.244*** (0.066)	0.262*** (0.070)	0.235*** (0.081)	0.232*** (0.086)	0.232*** (0.086)
$\Delta(\text{Player's best rating}=3)$	0.175*** (0.058)	0.192*** (0.062)	0.168** (0.073)	0.162** (0.077)	0.162** (0.077)
$\Delta(\text{Player's best rating}=2)$	0.121** (0.053)	0.133** (0.058)	0.109 (0.067)	0.104 (0.071)	0.104 (0.071)
One or more competing 5-stars	-0.002 (0.008)	-0.002 (0.017)	0.004 (0.017)	-0.001 (0.028)	-0.001 (0.029)
Prize value (\$100s)	0.002 (0.003)		0.008 (0.009)		
Pct. of contest elapsed	-0.001 (0.014)	-0.007 (0.020)	-0.010 (0.029)	-0.014 (0.036)	-0.104 (0.110)
Constant	-0.019* (0.011)	-0.009 (0.010)	-0.033 (0.032)	0.031 (0.089)	0.116 (0.123)
N	2694	2694	2694	2694	2694
R^2	0.01	0.04	0.10	0.13	0.13
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of the *change* in designs' similarity to the highest-rated preceding entry by the same player, taking values in $[-1,1]$, where a value of 0 indicates that the player's current design is as similar to her best preceding design as was her previous design, and a value of 1 indicates that the player transitioned fully from innovating to recycling (and a value of -1, the converse). The mean value of this variable in the sample is -0.01 (s.d. 0.25). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a difference hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table E.3: Similarity to other designs in the same submission batch (diff. hash)

	Unweighted		Weighted	
	(1)	(2)	(3)	(4)
Player's prior best rating==5	0.225 (0.299)	0.246 (0.293)	0.235 (0.286)	0.260 (0.281)
* 1+ competing 5-stars	-0.328** (0.146)	-0.324** (0.144)	-0.313** (0.147)	-0.308** (0.145)
* prize value (\$100s)	-0.022 (0.093)	-0.023 (0.092)	-0.025 (0.087)	-0.026 (0.085)
Player's prior best rating==4	-0.015 (0.031)	-0.003 (0.032)	-0.013 (0.030)	0.004 (0.031)
Player's prior best rating==3	0.011 (0.034)	0.019 (0.035)	0.010 (0.031)	0.020 (0.032)
Player's prior best rating==2	-0.019 (0.047)	-0.012 (0.049)	-0.022 (0.045)	-0.014 (0.045)
One or more competing 5-stars	-0.017 (0.033)	-0.018 (0.034)	-0.015 (0.032)	-0.017 (0.033)
Pct. of contest elapsed	-0.001 (0.039)	-0.024 (0.085)	0.003 (0.038)	0.007 (0.079)
Constant	0.643*** (0.121)	0.673*** (0.156)	0.670*** (0.099)	0.661*** (0.128)
N	1987	1987	1987	1987
R^2	0.59	0.59	0.59	0.59
Contest FEs	Yes	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes	Yes
Other Controls	No	Yes	No	Yes

Notes: Observations are design batches, which are defined to be a set of designs by a single player entered into a contest in close proximity (15 minutes). Dependent variable is a continuous measure of intra-batch similarity, taking values in $[0,1]$, where a value of 1 indicates that two designs in the batch are identical. The mean value of this variable in the sample is 0.69 (s.d. 0.28). Columns (3) and (4) weight the regressions by batch size. Columns (2) and (4) control for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a difference hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table E.4: Similarity to player's best not-yet-rated designs (placebo test; diff. hash)

	Similarity to forthcoming			Residual
	(1)	(2)	(3)	(4)
Player's best forthcoming rating==5	0.203 (0.241)	0.069 (0.116)	0.022 (0.127)	0.060 (0.119)
* 1+ competing 5-stars	-0.040 (0.145)	-0.024 (0.073)	0.000 (0.080)	-0.023 (0.079)
* prize value (\$100s)	-0.064 (0.042)	-0.009 (0.028)	-0.001 (0.032)	-0.006 (0.029)
Player's best forthcoming rating==4	0.023 (0.074)	0.051 (0.066)	0.059 (0.064)	0.045 (0.067)
Player's best forthcoming rating==3	0.043 (0.050)	0.069 (0.055)	0.069 (0.055)	0.055 (0.053)
Player's best forthcoming rating==2	0.031 (0.048)	0.025 (0.051)	0.026 (0.051)	0.025 (0.049)
One or more competing 5-stars	-0.077 (0.076)	-0.089 (0.115)	-0.087 (0.119)	-0.085 (0.124)
Pct. of contest elapsed	-0.210 (0.261)	0.033 (0.442)	0.093 (0.422)	0.023 (0.469)
Constant	0.735*** (0.243)	0.448 (0.465)	0.423 (0.473)	-0.025 (0.521)
N	1147	577	577	577
R^2	0.69	0.87	0.88	0.69
Contest FEs	Yes	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes

Table provides a placebo test of the effects of future feedback on originality. Observations are designs. Dependent variable in Columns (1) to (3) is a continuous measure of a design's similarity to the best design that the player has previously entered that has yet to *but will eventually be* rated, taking values in $[0,1]$, where a value of 1 indicates that the two designs are identical. The mean value of this variable is 0.50 (s.d. 0.29). Under the identifying assumption that future feedback is unpredictable, current choices should be unrelated to forthcoming ratings. Note that a given design's similarity to an earlier, unrated design can be incidental if they are both tweaks on a rated third design. To account for this possibility, Column (2) controls for the given and unrated designs' similarity to the best previously-rated design. Column (3) allows these controls to vary with the highest rating previously received. Dependent variable in Column (4) is the residual from a regression of the dependent variable in the previous columns on these controls. These residuals will be the subset of a given design's similarity to the unrated design that is not explained by jointly-occurring similarity to a third design. All columns control for days remaining and number of previous designs by the player and her competitors. Similarity scores in this table are calculated using a difference hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table E.5: Similarity to any of player's previous designs: 4-vs-4 (diff. hash)

	(1)	(2)	(3)
Player's prior best rating==4	0.087 (0.068)	0.221* (0.122)	0.745*** (0.253)
* 1+ competing 4- or 5-stars	-0.105** (0.047)	-0.190** (0.084)	-0.304 (0.201)
* prize value (\$100s)	0.006 (0.020)	-0.007 (0.037)	-0.162* (0.084)
Player's prior best rating==3	-0.014 (0.023)	-0.009 (0.039)	-0.048 (0.070)
Player's prior best rating==2	-0.016 (0.026)	0.049 (0.053)	0.058 (0.097)
One or more competing 4- or 5-stars	0.045* (0.025)	0.109 (0.081)	0.111 (0.177)
Pct. of contest elapsed	-0.006 (0.101)	-0.394 (0.267)	1.045 (0.864)
Constant	0.517*** (0.168)	1.033*** (0.294)	-0.621 (0.717)
N	2926	1557	879
R^2	0.55	0.61	0.70
Contest FEs	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Restriction	All	2nd half	4th qtr

Notes: Table shows the effects of 4-star feedback and competition on originality when no player has a 5-star rating. Observations are designs. Dependent variable is a continuous measure of a design's maximal similarity to previous entries in the same contest by the same player, taking values in [0,1], where a value of 1 indicates the design is identical to another. All columns include contest and player fixed effects and control for the number of days remaining and number of previous designs entered by the player and her competitors. Columns (2) and (3) restrict the sample to submissions in the second half or fourth quarter of a contest, when the absence of 5-star ratings may be more meaningful and is increasingly likely to be final. Similarity scores in this table are calculated using a difference hash algorithm. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

F Robustness Checks (2)

The following tables show that the full effect of competition on high performers' originality is realized with one high-quality competitor, and does not vary as competition intensifies. Tables F.1 to F.3 demonstrate this result with the perceptual hash similarity measures. In all cases, I estimate differential patterns in the presence of one vs. two or more top-rated, competing designs and find no such difference.

Table F.1: Similarity to player's best previously-rated designs (p. hash)

	(1)	(2)	(3)	(4)	(5)
Player's prior best rating==5	0.442*** (0.103)	0.461*** (0.093)	0.259*** (0.098)	0.358*** (0.098)	0.362*** (0.103)
* 1+ competing 5-stars	-0.252*** (0.093)	-0.290*** (0.078)	-0.176* (0.090)	-0.225** (0.093)	-0.226** (0.093)
* 2+ competing 5-stars	0.066 (0.075)	0.056 (0.078)	0.024 (0.097)	0.025 (0.094)	0.023 (0.092)
* prize value (\$100s)	-0.025 (0.025)	-0.016 (0.024)	0.005 (0.032)	-0.014 (0.031)	-0.018 (0.033)
Player's prior best rating==4	0.165*** (0.024)	0.160*** (0.022)	0.127*** (0.030)	0.122*** (0.031)	0.116*** (0.032)
Player's prior best rating==3	0.080*** (0.018)	0.077*** (0.019)	0.068** (0.028)	0.061** (0.028)	0.056** (0.028)
Player's prior best rating==2	0.045** (0.021)	0.044** (0.022)	0.022 (0.029)	0.026 (0.030)	0.024 (0.030)
One or more competing 5-stars	-0.025 (0.030)	0.017 (0.033)	0.005 (0.032)	0.002 (0.037)	0.002 (0.037)
Two or more competing 5-stars	0.008 (0.030)	-0.012 (0.036)	-0.011 (0.034)	0.004 (0.042)	0.000 (0.044)
Prize value (\$100s)	-0.014* (0.007)		-0.010 (0.010)		
Pct. of contest elapsed	-0.031 (0.034)	-0.059* (0.033)	-0.010 (0.030)	-0.019 (0.034)	-0.103 (0.084)
Constant	0.238*** (0.039)	0.207*** (0.024)	0.234*** (0.044)	0.231*** (0.062)	0.301*** (0.094)
N	3871	3871	3871	3871	3871
R ²	0.07	0.20	0.48	0.53	0.53
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of a design's similarity to the highest-rated preceding entry by the same player, taking values in [0,1], where a value of 1 indicates the design is identical to another. The mean value of this variable in the sample is 0.28 (s.d. 0.27). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table F.2: Change in similarity to player's best previously-rated designs (p. hash)

	(1)	(2)	(3)	(4)	(5)
$\Delta(\text{Player's best rating}=5)$	0.860*** (0.157)	0.879*** (0.166)	0.936*** (0.194)	0.921*** (0.200)	0.928*** (0.200)
* 1+ competing 5-stars	-0.497*** (0.119)	-0.487*** (0.133)	-0.535*** (0.150)	-0.504*** (0.160)	-0.505*** (0.161)
* 2+ competing 5-stars	0.106 (0.106)	0.099 (0.111)	0.163 (0.123)	0.108 (0.127)	0.107 (0.127)
* prize value (\$100s)	-0.090** (0.036)	-0.093** (0.037)	-0.115*** (0.044)	-0.108** (0.044)	-0.110** (0.045)
$\Delta(\text{Player's best rating}=4)$	0.279*** (0.061)	0.285*** (0.065)	0.274*** (0.071)	0.281*** (0.077)	0.284*** (0.077)
$\Delta(\text{Player's best rating}=3)$	0.146*** (0.054)	0.154*** (0.057)	0.140** (0.063)	0.141** (0.067)	0.142** (0.067)
$\Delta(\text{Player's best rating}=2)$	0.082* (0.042)	0.085* (0.046)	0.069 (0.052)	0.064 (0.056)	0.063 (0.056)
One or more competing 5-stars	-0.004 (0.016)	0.002 (0.026)	-0.017 (0.031)	-0.014 (0.044)	-0.013 (0.043)
Two or more competing 5-stars	0.001 (0.018)	-0.006 (0.025)	0.020 (0.033)	0.031 (0.041)	0.029 (0.043)
Prize value (\$100s)	0.003 (0.003)		0.003 (0.008)		
Pct. of contest elapsed	0.015 (0.012)	0.009 (0.018)	0.014 (0.024)	-0.003 (0.029)	-0.050 (0.074)
Constant	-0.029*** (0.011)	-0.017* (0.010)	-0.030 (0.029)	0.062 (0.092)	0.104 (0.108)
N	2694	2694	2694	2694	2694
R^2	0.03	0.05	0.11	0.14	0.14
Contest FEs	No	Yes	No	Yes	Yes
Player FEs	No	No	Yes	Yes	Yes
Other Controls	No	No	No	No	Yes

Notes: Observations are designs. Dependent variable is a continuous measure of the *change* in designs' similarity to the highest-rated preceding entry by the same player, taking values in $[-1,1]$, where a value of 0 indicates that the player's current design is as similar to her best preceding design as was her previous design, and a value of 1 indicates that the player transitioned fully from innovating to recycling (and a value of -1, the converse). The mean value of this variable in the sample is 0.00 (s.d. 0.23). Column (5) controls for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.

Table F.3: Similarity to other designs in the same submission batch (p. hash)

	Unweighted		Weighted	
	(1)	(2)	(3)	(4)
Player's prior best rating==5	0.213 (0.312)	0.240 (0.307)	0.245 (0.302)	0.285 (0.294)
* 1+ competing 5-stars	-0.437** (0.217)	-0.433** (0.214)	-0.475** (0.203)	-0.469** (0.200)
* 2+ competing 5-stars	0.180 (0.236)	0.177 (0.232)	0.237 (0.223)	0.238 (0.218)
* prize value (\$100s)	0.020 (0.099)	0.015 (0.099)	0.013 (0.095)	0.009 (0.093)
Player's prior best rating==4	0.055* (0.032)	0.066* (0.037)	0.065** (0.032)	0.086** (0.038)
Player's prior best rating==3	0.056 (0.035)	0.062* (0.037)	0.052 (0.035)	0.065* (0.037)
Player's prior best rating==2	0.023 (0.049)	0.029 (0.050)	0.009 (0.047)	0.020 (0.047)
One or more competing 5-stars	0.083 (0.063)	0.086 (0.063)	0.076 (0.063)	0.078 (0.063)
Two or more competing 5-stars	-0.094 (0.072)	-0.100 (0.076)	-0.080 (0.074)	-0.086 (0.076)
Pct. of contest elapsed	-0.016 (0.049)	-0.102 (0.115)	-0.005 (0.050)	-0.064 (0.112)
Constant	0.392*** (0.066)	0.504*** (0.149)	0.385*** (0.061)	0.457*** (0.148)
N	1987	1987	1987	1987
R^2	0.58	0.58	0.58	0.58
Contest FEs	Yes	Yes	Yes	Yes
Player FEs	Yes	Yes	Yes	Yes
Other Controls	No	Yes	No	Yes

Notes: Observations are design batches, which are defined to be a set of designs by a single player entered into a contest in close proximity (15 minutes). Dependent variable is a continuous measure of intra-batch similarity, taking values in $[0,1]$, where a value of 1 indicates that two designs in the batch are identical. The mean value of this variable in the sample is 0.45 (s.d. 0.32). Columns (3) and (4) weight the regressions by batch size. Columns (2) and (4) control for the number of days remaining and number of previous designs entered by the player and her competitors. Similarity scores in this table are calculated using a perceptual hash algorithm. Preceding designs/ratings are defined to be those entered/provided at least 60 minutes prior to the given design. *, **, *** represent significance at the 0.1, 0.05, and 0.01 levels, respectively. Standard errors clustered by player in parentheses.