

PLEBEIAN BIAS:
SELECTING CROWDSOURCED CREATIVE DESIGNS FOR
COMMERCIALIZATION

Anirban Mukherjee

Ping Xiao

Hannah H. Chang

Li Wang

Noshir Contractor¹

This version: September 20, 2017.

¹ Anirban Mukherjee is Assistant Professor of Marketing at the Lee Kong Chian School of Business, Singapore Management University. E-mail: anirbanm@smu.edu.sg. Ping Xiao is Senior Lecturer at University of Technology Sydney. E-mail: ping.xiao@uts.edu.au. Hannah H. Chang is Associate Professor of Marketing at the Lee Kong Chian School of Business, Singapore Management University. E-mail: hannahchang@smu.edu.sg. Li Wang is Assistant Professor of Marketing at Shanghai University of Finance and Economics. E-mail: wang.li@mail.shufe.edu.cn. Noshir Contractor is Jane S. & William J. White Professor of Behavioral Sciences, at the McCormick School of Engineering & Applied Science, the School of Communication and the Kellogg School of Management at Northwestern University. E-mail: nosh@northwestern.edu. This material is based upon work supported by the National Science Foundation under grant no. IIS-1514427. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

Abstract

We identify a new phenomenon – “Plebeian bias” – in the crowdsourcing of creative designs. Stardom, an emphasis on established individuals, has long been observed in many offline contexts. Does this phenomenon carry over to online communities? We investigate a large-scale dataset tracking all submissions, community votes on submissions, and revenues from commercialized submissions on a popular crowdsourcing portal, Threadless.com. In contrast to stardom, we find that the portal selects designs from “Plebeians” (i.e. users without an established fan base and track record) over “Stars” (i.e. users with an established fan base and track record). The tendency is revenue and profit sub-optimal. The evidence is consistent with incentives for the portal to demonstrate procedural fairness to the online community.

Keywords: crowdsourcing, stardom, procedural fairness, managerial conservatism.

Almost a century of scholarship—spanning a diverse set of academic literatures (theoretical economics, industrial organizational economics, amongst others) and popular press—finds that in many contexts, firms prefer to employ a small set of established individuals over a vast pool of unknown individuals. This phenomenon seems particularly pronounced in the creative industries (e.g., movies, books, and music). In the spirit of Rosen (1981), we refer to this phenomenon as “stardom.” Various theoretical models propose that stardom arise from low marginal costs (Rosen 1981), learning (Adler 1985), and managerial conservatism (Scharfstein & Stein 1990; Zweibel 1995; Holmström 1999). Empirical work has documented stardom in fields as diverse as finance (Gabaix & Landier 2008), entertainment (Einav 2010), software development (Volmer & Sonnentag 2011), sports (Hausman & Leonard 1997), and law (Sunstein, Murphy, Frank, & Rosen 2000). Taken together, the extant literature strikes a pessimistic note: it predicts the increased prevalence of stardom and consequently a more economically stratified society (Frank & Cook 2010).

Recently, the emergence of crowdsourcing has brought new hope of a more equitable future. “Crowdsourcing,” a portmanteau neologism, is the sourcing of organizational functions from the “crowd”: a large, undefined community of the firm’s consumers, partners, and collaborators (p. 226, Bayus 2013). The online crowdsourcing of new venture funding, crowdfunding, for example, has enabled the flow of capital to entrepreneurs in previously underserved regions (Sorenson, Assenova, Li, Boada, & Fleming 2016). In the context of new product development, crowdsourcing portals issue an open call for ideas from an online community and develop the ideas that seem commercially viable (Ogawa & Piller 2006). Given the egalitarian nature of internet communities, scholars have expressed optimism that online

crowdsourcing may allow novice entrepreneurs and emerging artists, an alternative path to success (Howe 2006).

Little, however, is formally known of stardom in crowdsourcing. There are two theoretical paradigms of stardom, which lead to conflicting predictions. On the one hand, extant research finds that stardom in the creative industries has significant informational value to the firm, its customers and its partners (Adler 1985; Liu, Mazumdar, & Li 2014). These papers suggest that stars increase the economic value of creative products beyond the expense of additional expense of hiring stars; thus, the observed emphasis on stardom is profit optimal (Hofmann, Clement, Völckner, & Hennig-Thurau 2017). Based on this stream of research, one would expect the informational role of stardom to be even greater in crowdsourcing than in traditional enterprise. This is because crowdsourcing often attracts many submissions that vary considerably in quality. Furthermore, established members of the community are typically well-known and respected within the community. These factors should amplify the informational value of stardom in crowdsourcing (p. 2752, Liu 2017).

On the other hand, crowdsourcing portals focused on new product development have different incentives than traditional firms. In traditional firms, new product development is centralized and involves a few, relatively homogenous individuals. Prior research suggests that stardom arises from to a principal-agent problem in the traditional firms. Managers follow the industry norm and make conservative hiring choices – they favor stars – to further their own career interests (Mukherjee & Kadiyali 2017). In particular, new product development is plagued by outcome uncertainty—it is difficult to predict the commercial prospects of a new product (Eliashberg, Hui, & Zhang 2007). By focusing on stars, managers mitigate the effect of outcome uncertainty on their career. In crowdsourcing, however, new product development is

decentralized and determined by many heterogeneous, self-selected individuals: the crowd. Hence, the crowdsourcing portal has to consider the incentives of individuals in the crowd.

Crowdsourcing participants express the need to perceive the selection process as fair to be willing to contribute to a portal (Franke, Keinz, & Klausberger 2013). Prior research has established that the perceived fairness of the allocation process—known as procedural fairness—is an important factor in determining recipients' responses (Barrett-Howard & Tyler 1986; Leventhal 1980; Thibaut & Walker 1975; Gilliland 1993). Procedural fairness has been studied extensively, spanning the contexts of legal procedures (Lind, Kurtz, Musante, Walker, & Thibaut 1980), dispute resolutions (Tyler & Folger 1980), job applicant selections (Gilliland 1993), pay raise decisions (Folger & Konovsky 1989), student evaluations of teachers (Tyler & Caine 1981), and supplier-reseller relationships (Kumar, Scheer, & Steenkamp 1995). Various theories on procedural fairness (e.g., Folger & Cropanzano 2001; Lind 2001; Van den Bos, Lind, & Wilke 2001) propose that information about procedural rules signals to members whether the group and, more generally, the broader environment are fair and whether members should expect fair treatment.

Procedural fairness is a particularly important consideration in crowdsourcing because crowdsourcing participants are drawn by intrinsic motivations (Boudreau & Lakhani 2013; Jeppesen & Frederiksen 2006) rather than the (nominal) monetary compensation (Bullinger et al. 2010; Füller 2006; Nambisan & Baron 2010). Therefore, to encourage participation, a crowdsourcing portal needs to ensure that it is perceived as creating a neutral arena (i.e., a level playing field) for all participants, ensuring equality of opportunity and not favoring some users over others based on preconceptions. This is in line with the notions of neutrality (Tyler 1989) and procedural rules of consistent allocation and bias-suppression (Leventhal 1980).

To examine the role of stardom in crowdsourcing, we study a large-scale dataset from Threadless.com (henceforth Threadless). Threadless has an open call for new designs, which draws a large number of submissions from its community. The community votes on each design. Threadless then selects a small number of submitted designs for commercialization. Our data describes all designs, votes, and revenues on Threadless between January 2004 and July 2010.²

Our results indicate that Threadless is biased in its selections—the portal chooses designs from its unestablished users (“Plebeians”) over its established users (“Stars”). This tendency is revenue-suboptimal: Threadless selects lower commercial potential designs from Plebeians while rejecting higher commercial potential designs from Stars.

What makes Threadless different? There are three important factors to consider. First, crowdsourcing contributors differ in their track record (i.e., whether their previous design submissions were selected by the portal). While some contributors have a limited track record, other contributors—the stars of the crowdsourcing community—have an established track record. Second, stars (established users) on Threadless submit designs with higher commercial potential. This suggests that if Threadless were to choose submissions with the highest commercial potential, it would likely end up selecting more designs from the stars. Third, the status of the contributors (i.e., their stardom) is visible to the public, but neither the disaggregate votes submitted by the community nor the revenues from different designs, are visible to the public.

Therefore, selecting designs on a revenue basis alone might make the selection system appear to favor the stars. To be perceived as fair, Threadless’ selection decision needs to be seen

² The data period is constrained by our arrangements with the firm. Competitive concerns hinder our ability to get more recent data.

as egalitarian and not elitist, akin to fairness perception in subgroup hiring rate in personnel selection (Hartigan & Wigdor 1989; Hunter & Schmidt 1996; Quillian, Pager, Hexel, & Midtbøen 2017). In sum, unlike traditional firms, in crowdsourcing, the portal has both a profit motive and a communal motive—it needs to ensure that its decisions inspire the community to remain engaged in problem-solving activities for the firm. The latter incentive is a likely reason for the observed bias.

Institutional Context

Details on the submission and voting process on Threadless during our sample period are as follows: All registered users (registration is free and open to the public) can submit designs. The submission process involves uploading a digital image of the design and a title for the design. Submitted designs are put up for voting for seven days. Any registered user (excluding the user who submitted the design) may vote once on a submitted design. To ensure designs receive a fair vote, Threadless randomizes the order in which users encounter designs open for voting, and does not provide an option to sort designs open for voting. This ensures that all designs get a similar chance of being voted on. Users vote on a 6-point scale from 0 (“I don’t like this design”) to 5 (“I love this design”). Voting consists only of a numerical score and users do not provide any other formal feedback to the Threadless portal.

To reduce gaming, the disaggregate votes (scores) are private and never revealed to the public. Threadless reveals the mean vote and the number of votes cast for a submitted design at the end of the voting process. At the end of voting, Threadless selects the designs that it wishes to retail. Threadless has discretion on how many (if any) designs it chooses, without being bound to a specific decision criterion.

Users whose design were selected for retail are given a modest monetary reward (US\$2,000 in 2010). Users whose designs are not selected for retail are not compensated monetarily. Importantly, regardless of star status, users cannot privately negotiate a contract with Threadless. Therefore, there is no difference in marginal cost to Threadless of selecting a design from a star over a plebeian. This is different than stardom in traditional contexts.

Data and Empirical Strategy

We rely on a carefully collected large-scale dataset of all votes, all submissions, and all revenues on Threadless from January 1, 2004 to July 31, 2010. From these, we drop less than 0.05% of votes where the numerical value of the vote is missing in our data³. From the 150,093 designs submitted to Threadless, we drop 62 designs (less than 0.05%) where the identity of the user who submitted the design is missing, and 1 design (less than 0.01%) where the date of the submission is missing in our data. Our final dataset tracks 150,030 designs submitted by 48,556 unique users.

Our data provides an excellent test bed to study stardom in the crowd. Two factors are crucial to the analysis. First, we observe all candidate designs at a relatively complete stage of the design process. In contrast, in most extant empirical applications, researchers only observe candidate designs at an early stage of the development process (for example, the script of a movie, as in Eliashberg, Hui, & Zhang 2007, or a raw product idea, prior to iteration and change, as in Kornish & Ulrich 2014).

Second, the data allow us to map submitted designs to their commercial potential. As described prior, Threadless crowdsources votes on submitted designs. Threadless randomizes the order in which users see designs. This ensures that there are no order effects in voting. Users do

³ To the best of our knowledge, the missing data is at random and due to data corruption during warehousing.

not observe other users' votes and do not observe the voting history of the submitting user. This ensure that there are no herding effects in voting. In sum, the votes of the crowd likely reflect its preferences.

The conjunction of the voting data and the revenue data allows us to predict the commercial potential of all submissions, including those that were not selected for manufacture and retail by Threadless. Note that we use information available to Threadless at the time of selecting designs. Therefore, we are able to infer and evaluate its selection strategy. These features are unique to our data and context. In extant applications, however, it is challenging to both obtain commercial data on new products and to evaluate the commercial potential of product ideas that were not selected for commercialization, due to the lack of a comprehensive evaluation (voting) mechanism.

We divide submissions into three categories based on submitting user's track record: (1) submissions where the submitting user has not had a design selected by Threadless, (2) submissions where the submitting user has had 1 to 3 prior submissions selected by Threadless, and (3) submissions where the submitting user has had 4 or more prior submissions selected by Threadless. Users with more than one prior submission were separated into two categories (i.e., the second and third categories) to better illustrate the findings. In the remainder of the text, for ease of exposition, we refer to the first group as "Plebeians", and the joint of the second and third groups as "Stars". We focus on three groups of variables: (a) the votes submitted by the crowd on the design, (b) the track record of the submitting user, and (c) a year-specific fixed effect.

Table 1 summarizes the descriptive statistics of our variables. Table 1 shows that designs submitted by Stars get consistently higher votes than designs submitted by Plebeians. For

example, the median design from a Plebeian receives 40 votes of 5 on a 0 to 5 scale, while the median design from a Star receives between 145 and 155 votes of 5.

--- TABLE 1 ABOUT HERE ---

Figure 1 is a boxplot of the number of negative votes (i.e. sum of the number of 0, 1, or 2 votes on a 0 to 5 scale) and the number of positive votes (i.e. sum of the number of 3, 4, or 5 votes on a 0 to 5 scale) received by each submission, by each category of submission. Figure 1 shows that submissions from Stars receive considerably more positive votes, but about the same number of negative votes, as submissions from Plebeians. Importantly, about half of the unselected submissions from the Stars received a comparable number of positive votes as submissions that were selected from the Plebeians. There is no comparable trend in the negative votes. This suggests that there is a large pool of unselected submissions from the Stars that garner positive attention from the crowd, but which are not selected for commercialization. Instead, submissions that received less positive attention, and comparable negative attention, were picked by Threadless.

--- FIGURE 1 ABOUT HERE ---

We use machine learning methods to examine the data. The crowd's votes are of high dimensionality (they are on a 6-point scale, from 0 to 5). Furthermore, the relationship between the crowd's votes and success may be non-linear and may vary over time. Therefore, it is difficult to a priori identify the appropriate statistical model structure relating the crowd's votes to revenues. Machine learning models search over both model structure and data features to determine the most appropriate statistical model for a predictive model. Therefore, they are ideally suited to developing the empirical model. Specifically, we rely on a class of (supervised) machine learning models called Support Vector Regression (henceforth SVR) to predict

revenues (Drucker, Burges, Kaufman, Smola & Vapnik 1996). SVRs are able to efficiently perform non-linear regressions due to the “kernel trick,” which allows a mapping of the inputs into high-dimensionality space (Rasmussens & William 2006). We use a radial basis and conduct three-fold cross validation to select the model.

Results

Figure 2 compares the predicted revenues for selected and unselected submissions across the three groups of users (no prior selected submissions, 1 to 3 prior selected submissions, and 4 or more prior selected submissions). Figure 3 shows that Threadless chooses submissions from Plebeians that are substantially lower in forecasted revenues than Stars. Figure 3 is a quantile-to-quantile plot (Q-Q plot) of predicted revenues for submissions by Stars and Plebeians, which were selected (or not selected) by Threadless. Specifically, Figure 3 plots each percentile by predicted revenues for submissions by Stars against the corresponding percentile by predicted revenues for submissions by Plebeians. It overlays a similar plot for the unselected submissions from Stars and Plebeians. Last, it includes a line passing through the origin with slope equal to 1, which represents equal opportunity across stardom.

--- FIGURE 2 ABOUT HERE ---

--- FIGURE 3 ABOUT HERE ---

If stardom did not play a role on Threadless, we would expect the quantile-to-quantile points to (on average) center on the line with slope equal to 1. Instead, Figure 3 shows that across all quantiles, the predicted revenue for selected designs from Stars is higher than designs from Plebeians. In particular, across all percentiles, only higher commercial potential designs from Stars are selected, relative to the designs selected from Plebeians. To formalize this intuition, we compute the Kolmogorov-Smirnov test statistic. This statistical test compares the cumulative

distribution functions of two variables. In our case, the test corresponds to a test of fairness. The test rejects the null of similarity ($D = 0.21, p < 2.2e-16$) for the distribution of predicted revenues from selected designs from Stars and from Plebeians.

In addition, the findings show that (1) the predicted revenues for unselected designs (at all percentiles) are higher for Stars than Plebeians, and (2) the predicted revenues for a significant number of designs by Stars are higher than those for designs by Plebeians. This implies that, as suggested by Figure 1 where for a number of unselected submissions by the Stars, the number of positive votes is substantially higher than the number of positive votes for selected submissions by Plebeians, Threadless is under-selecting (high commercial potential) submissions from Stars in favor of (low commercial potential) designs from Plebeians.

Table 2 describes the deciles of these groups over the years of the dataset. Across all years (rows) and all deciles (columns), we see the same trend as depicted in Figure 3. Therefore, the bias identified in Figure 3 is both pervasive and persistent across the 90 months covered in our data.

--- TABLE 2 ABOUT HERE ---

Discussion

Inequality due to stardom is a distressing, yet ubiquitous, phenomenon. Today, corporate America's stars—its top CEOs, ace investment bankers, and hotshot lawyers—receive a greater share of total remuneration than any time prior in modern history. The rising inequality in wages and opportunity has led to increasing calls for governmental action, in part due to a perception that without intervention, inequality may beget more inequality (Sands 2017). Broadly, scholars are pessimistic about the future (Piketty 2017).

An important exception is the role of the internet. Scholars have expressed hope that the internet may help make available a wide-variety of resources to entrepreneurs in disadvantaged neighborhoods, reducing inequality (Boudreau & Lakhani 2013). Of these tools, perhaps the most prevalent and discussed phenomenon is the crowdsourcing of new venture funding, known as crowdfunding (Sorenson et al. 2016). More generally, crowdsourcing is a form of open innovation, where firms and customers collaborate in the development of new products and services (Bayus 2013). These new business practices may help democratize access to success for unestablished entrepreneurs, artists, and professionals.

Our findings are very encouraging for a more equitable future. We observe that Threadless favors Plebeians to Stars: it favors unestablished users over established users. This strategy is revenue and profit suboptimal for Threadless. However, it is likely undertaken to encourage contribution and participation by the online community. Specifically, the prior evidence suggests that if a crowdsourcing system is perceived as unfair, potential contributors are unlikely to join the system in the first place (Franke, Keinz, & Klausberger 2013), and current contributors are likely to exit the system (Felstiner 2011). Thus, the reduction in discrimination between Plebeians and Stars is likely because the online community values fairness.

In sum, our findings suggest that open innovation may help reduce inequity. Stardom is rooted in information asymmetry (Adler 1985) and managerial conservatism (Zwiebel 1995). Our findings suggest that open innovation may both help mitigate information uncertainty by asking the crowd for feedback on alternatives, and overcome managerial conservatism by injecting procedural fairness into the decision calculus of managers.

Crucially, this is uplifting news because it implies that the open innovation may be more important than theorized previously. That is, not only does Threadless allow anyone to submit a design from anywhere, but the community oversight also leads to its emphasis on unestablished users over established users. Therefore, there is reason to hope that open innovation may act as a foil for the star-centered business model of many modern industries. In particular, open innovation may lead to fairer outcomes in the creative industries, where the effects of managerial conservatism are particularly pernicious.

REFERENCES

1. Rosen S (1981) The economics of superstars. *The American economic review* 71(5):845-858.
2. Adler M (1985) Stardom and talent. *The American economic review* 75(1):208-212.
3. Scharfstein DS & Stein JC (1990) Herd behavior and investment. *The American Economic Review* 80(3):465-479.
4. Zwiebel J (1995) Corporate conservatism and relative compensation. *Journal of Political economy* 103(1):1-25.
5. Holmström B (1999) Managerial incentive problems: A dynamic perspective. *The review of Economic studies* 66(1):169-182.
6. Gabaix X & Landier A (2008) Why has CEO pay increased so much? *The Quarterly Journal of Economics* 123(1):49-100.
7. Einav L (2010) Not all rivals look alike: Estimating an equilibrium model of the release date timing game. *Economic Inquiry* 48(2):369-390.
8. Volmer J & Sonnentag S (2011) The role of star performers in software design teams. *Journal of Managerial Psychology* 26.3: 219-234.
9. Hausman JA & Gregory KL Superstars in the National Basketball Association: Economic value and policy. *Journal of Labor Economics* 15.4: 586-624.
10. Sunstein C, Murphy K, Frank R, & Rosen S (2000) The Wages of Stardom: Law and the Winner-Take-All Society: A Debate, Roundtable Discussion at the University of Chicago Law School. *U. Chi. L. Sch. Roundtable* 6:1.
11. Frank RH & Cook PJ (2010) *The winner-take-all society: Why the few at the top get so much more than the rest of us* (Random House).
12. Bayus BL (2013) Crowdsourcing new product ideas over time: An analysis of the Dell IdeaStorm community. *Management science* 59(1):226-244.
13. Sorenson O, Assenova V, Li G-C, Boada J, & Fleming L (2016) Expand innovation finance via crowdfunding. *Science* 354(6319):1526-1528.
14. Ogawa S & Piller FT (2006) Reducing the risks of new product development. *MIT Sloan management review* 47(2):65.
15. Howe J (2006) The rise of crowdsourcing. *Wired magazine* 14(6):1-4.
16. Liu A, Mazumdar T, & Li B (2014) Counterfactual decomposition of movie star effects with star selection. *Management Science* 61(7):1704-1721.
17. Hofmann J, Clement M, Völckner F, & Hennig-Thurau T (2017) Empirical generalizations on the impact of stars on the economic success of movies. *International Journal of Research in Marketing* 34(2):442-461.
18. Liu HK (2017) Crowdsourcing Design: A Synthesis of Literatures. *Proceedings of the 50th Hawaii International Conference on System Sciences*.
19. Mukherjee A & Kadiyali V (2017) The Competitive Dynamics of New DVD Releases. *Management Science*.

20. Eliashberg J, Hui SK, & Zhang ZJ (2007) From story line to box office: A new approach for green-lighting movie scripts. *Management Science* 53(6):881-893.
21. Franke N, Keinz P, & Klausberger K (2013) “Does this sound like a fair deal?”: Antecedents and consequences of fairness expectations in the individual’s decision to participate in firm innovation. *Organization Science* 24(5):1495-1516.
22. Barrett-Howard E & Tyler TR (1986) Procedural justice as a criterion in allocation decisions. *Journal of Personality and Social Psychology* 50(2):296.
23. Leventhal GS (1980) What should be done with equity theory? *Social exchange*, (Springer), pp 27-55.
24. Thibaut JW & Walker L (1975) *Procedural justice: A psychological analysis* (L. Erlbaum Associates).
25. Gilliland SW (1993) The perceived fairness of selection systems: An organizational justice perspective. *Academy of management review* 18(4):694-734.
26. Lind EA, Walker L, Kurtz S, Musante L, & Thibaut JW (1980) Procedure and outcome effects on reactions to adjudicated resolution of conflicts of interest. *Journal of Personality and Social Psychology* 39(4):643.
27. Tyler TR & Folger R (1980) Distributional and procedural aspects of satisfaction with citizen-police encounters. *Basic and Applied Social Psychology* 1(4):281-292.
28. Folger R & Konovsky MA (1989) Effects of procedural and distributive justice on reactions to pay raise decisions. *Academy of Management journal* 32(1):115-130.
29. Tyler TR & Caine A (1981) The influence of outcomes and procedures on satisfaction with formal leaders. *Journal of Personality and Social Psychology* 41(4):642.
30. Kumar N, Scheer LK, & Steenkamp J-BE (1995) The effects of perceived interdependence on dealer attitudes. *Journal of marketing research* 32(3):348-356.
31. Folger R & Cropanzano R (2001) Fairness theory: Justice as accountability. *Advances in organizational justice*, eds Greenberg J & Gropanzano R (Stanford University Press, Stanford, California), Vol 1, pp 1-55.
32. Lind EA (2001) Fairness heuristic theory: Justice judgments as pivotal cognitions in organizational relations. *Advances in organizational justice*, eds Greenberg J & Gropanzano R (Stanford University Press, Stanford, California), Vol 56, p 88.
33. Van den Bos K, Lind E, & Wilke H (2001) The psychology of procedural and distributive justice viewed from the perspective of fairness heuristic theory. in press In R. Cropanzano (Ed.), Justice in the workplace: Vol. 2. From theory to practice. *Justice in the workplace*, ed Gropanzano R (Lawrence Erlbaum Associates, Publishers, New Jersey), Vol 2.
34. Boudreau KJ & Lakhani KR (2013) Using the crowd as an innovation partner. *Harvard business review* 91(4):60-69, 140.
35. Jeppesen LB & Frederiksen L (2006) Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments. *Organization science* 17(1):45-63.

36. Bullinger AC, Neyer AK, Rass M, & Moeslein KM (2010) Community - based innovation contests: Where competition meets cooperation. *Creativity and innovation management* 19(3):290-303.
37. Füller J (2006) Why consumers engage in virtual new product developments initiated by producers. *Connie Pechmann and Linda Price*, ed 33 N-AiCRV (Association for Consumer Research, Duluth, MN), pp 639-646.
38. Nambisan S & Baron RA (2010) Different roles, different strokes: Organizing virtual customer environments to promote two types of customer contributions. *Organization Science* 21(2):554-572.
39. Tyler TR (1989) The psychology of procedural justice: A test of the group-value model. *Journal of personality and social psychology* 57(5):830.
40. Hartigan J & Wigdor A (1989) Fairness in employment testing. *Science* 245(4913):14-14.
41. Hunter JE & Schmidt FL (1996) Intelligence and job performance: Economic and social implications. *Psychology, Public Policy, and Law* 2(3-4):447.
42. Quillian L, Pager D, Hexel O, & Midtbøen AH (2017) Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences*.
43. Kornish LJ & Ulrich KT (2014) The importance of the raw idea in innovation: Testing the sow's ear hypothesis. *Journal of Marketing Research* 51(1):14-26.
44. Drucker H, Burges CJ, Kaufman L, Smola AJ, & Vapnik V (1997) Support vector regression machines. *Advances in neural information processing systems 6 (NIPS 1996)*, eds Mozer MC, Jordan MI, & Petsche. T (MIT Press), pp 155-161.
45. Rasmussen CE & Williams CK (2006) *Gaussian processes for machine learning* (MIT Press, Cambridge, Massachusetts).
46. Felstiner A (2011) Working the crowd: employment and labor law in the crowdsourcing industry. *Berkeley Journal of Employment and Labor Law* 32(1).
47. Boudreau K & Lakhani K (2011) Field Experimental Evidence on Sorting, Incentives and Creative Worker Performance. *Harvard Business School*:11-107.
48. Sands, M L (2017) Exposure to inequality affects support for redistribution. *Proceedings of the National Academy of Sciences*.
49. Piketty, T (2017) *Capital in the twenty-first century*. Harvard University Press.

Table 1: Descriptive Statistics

		No prior selected designs	1 to 3 prior selected designs	4 or more prior selected designs
Number of Submitted Designs	Count	134,825	11,580	3,625
Number of Zero Votes	Minimum	0	0	2
	Mean	204.21	230.13	158.41
	Median	157	189	98
	Maximum	1,330	1,592	870
Number of One Votes	Minimum	0	0	0
	Mean	189.33	224.42	161.95
	Median	151	180	116
	Maximum	745	745	641
Number of Two Votes	Minimum	0	0	2
	Mean	185.53	252.11	198.87
	Median	143	196	149
	Maximum	736	685	667
Number of Three Votes	Minimum	0	0	0
	Mean	139.50	232.29	206.05
	Median	100	191	164
	Maximum	724	684	599
Number of Four Votes	Minimum	0	0	1
	Mean	85.12	168.45	168.39
	Median	54	145	145
	Maximum	617	673	643
Number of Five Votes	Minimum	0	0	1
	Mean	74.88	169.52	188.41
	Median	40	135	155
	Maximum	3,183	1,435	1,271
Number of Prior Submissions	Minimum	0	1	6
	Mean	3.63	28.63	70.75
	Median	1	21	61
	Maximum	113	196	212
Number of Prior Selected Submissions	Minimum	0	1	4
	Mean	0	1.51	7.78
	Median	0	1	6
	Maximum	0	3	29
Natural Logarithm of Prior Revenue, if Selected	Minimum	0	7.406	8.58
	Mean	0	9.66	9.79
	Median	0	9.68	9.81
	Maximum	0	12.39	10.88

Table notes:

1. No prior selected designs = Submissions from users whose prior design submissions were not selected by Threadless.
2. 1 to 3 prior selected designs = Submissions from users who have 1 to 3 prior design submissions selected by Threadless.
3. 4 or more prior selected designs = Submissions from users who have 4 or more prior design submissions selected by Threadless.
4. Number of Prior Submissions = Number of prior submissions by the submitting user.
5. Number of Prior Selected Submissions = Number of prior submissions by the submitting user that were selected by Threadless.

Table 2: Predicted Revenue of Submissions

Year	Selected	Status	Decile								
			1st	2nd	3rd	4th	5th	6th	7th	8th	9th
2004	NO	Plebeian	5,460	5,686	5,880	6,076	6,268	6,480	6,717	7,075	7,741
		Star	6,834	7,203	7,399	7,652	8,082	8,334	8,761	9,401	10,356
	YES	Plebeian	6,755	7,427	8,085	8,428	8,843	9,161	9,746	10,330	11,702
		Star	9,322	10,280	10,495	10,648	10,732	10,891	11,266	13,036	14,114
2005	NO	Plebeian	5,837	6,396	6,766	6,947	7,131	7,822	8,813	9,938	11,390
		Star	7,709	8,284	9,528	10,452	11,366	12,230	13,349	14,268	15,592
	YES	Plebeian	9,740	11,385	11,976	12,463	13,394	14,444	15,236	15,972	17,233
		Star	10,360	12,275	13,741	14,874	15,419	15,593	16,542	17,052	18,194
2006	NO	Plebeian	5,627	5,981	6,408	9,488	10,678	11,340	11,983	12,800	13,972
		Star	9,430	10,491	11,214	11,893	12,632	13,380	14,297	15,221	16,344
	YES	Plebeian	12,580	13,235	13,970	14,658	15,136	15,648	16,328	16,836	17,919
		Star	12,995	14,111	15,085	15,582	15,944	16,360	16,944	18,007	19,494
2007	NO	Plebeian	5,681	6,126	6,610	7,426	9,474	10,699	11,805	12,986	14,802
		Star	9,660	11,033	12,141	13,314	14,350	15,213	16,152	17,137	18,618
	YES	Plebeian	13,726	15,907	17,335	18,034	18,816	19,486	20,204	21,032	22,503
		Star	15,142	16,724	17,414	18,225	18,794	19,567	20,581	21,736	23,374
2008	NO	Plebeian	6,588	7,138	7,585	8,000	8,367	8,745	9,167	9,923	11,646
		Star	9,177	10,194	11,053	11,796	12,568	13,374	14,249	15,682	17,765
	YES	Plebeian	12,282	14,315	15,392	16,204	17,668	18,105	20,075	23,056	24,996
		Star	13,920	14,675	15,668	16,706	17,716	19,065	20,683	22,292	24,261
2009	NO	Plebeian	7,144	7,861	8,471	9,059	9,644	10,201	10,739	11,210	11,896
		Star	9,917	11,026	11,727	12,366	12,913	13,511	14,046	14,770	15,713
	YES	Plebeian	11,914	13,464	14,053	14,712	15,260	15,866	16,558	17,158	19,297
		Star	13,434	14,325	14,919	15,443	15,934	16,402	17,151	18,290	19,712
2010	NO	Plebeian	9,413	10,119	10,672	11,208	11,715	12,150	12,503	12,854	13,388
		Star	11,185	12,104	12,848	13,429	13,954	14,499	15,087	15,755	17,065
	YES	Plebeian	13,717	14,705	15,558	16,188	16,492	17,017	18,083	18,796	19,972
		Star	14,676	15,189	15,894	16,514	17,049	17,972	18,863	20,556	23,646

Table notes:

1. Each value is the corresponding decile of the predicted revenue of submissions by a Plebeian / Star, which was selected / not selected by Threadless.
2. Selected = Submissions which were selected by Threadless.
3. Plebeian = Submitting users who have not had a prior design submission selected by Threadless.
4. Star = Submitting users who have had at least one prior design submission selected by Threadless.
5. No = Submissions that are not selected by Threadless; Yes = Submissions that are selected by Threadless.

Figure 1: Number of Votes

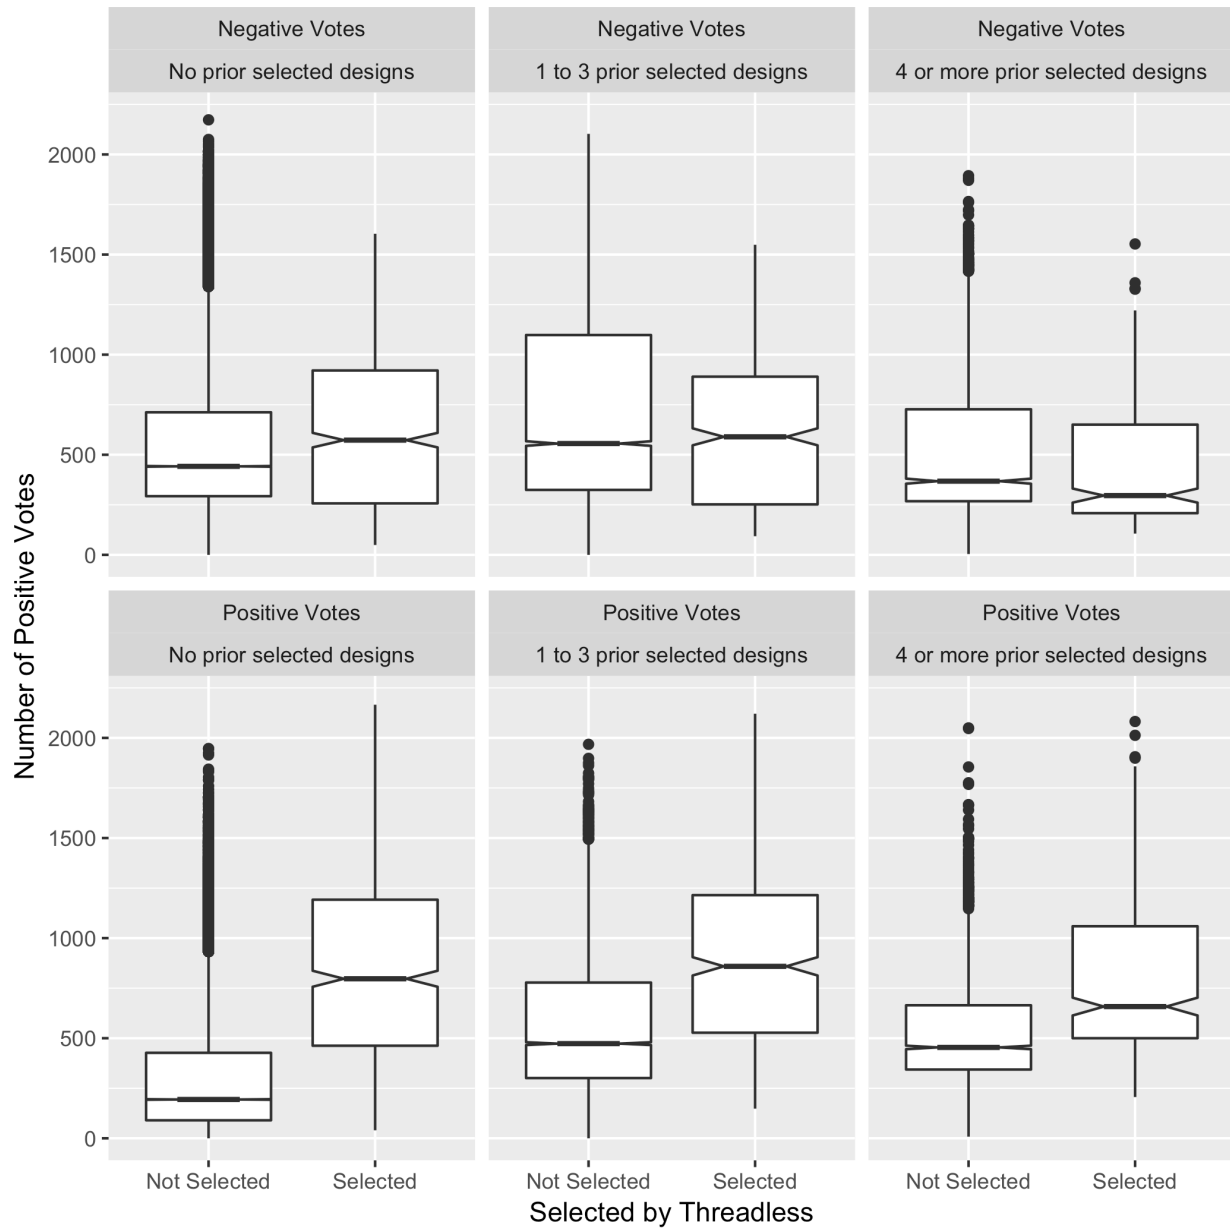


Figure notes:

1. Not selected = Submissions that are not selected by Threadless.
2. Selected = Submissions that are selected by Threadless.
3. Number of negative votes = Number of votes equal to 0, 1, and 2.
4. Number of positive votes = Number of votes equal to 3, 4, and 5.
5. No prior selected designs = Users who have not had a design selected by Threadless.
6. 1-3 prior selected designs = Users who have had between 1 and 3 designs selected by Threadless.
7. 4 or more prior selected designs = Users who have had 4 or more designs selected by Threadless.

Figure 2: Predicted Revenue by Number of Prior Selections, and Selection by Threadless

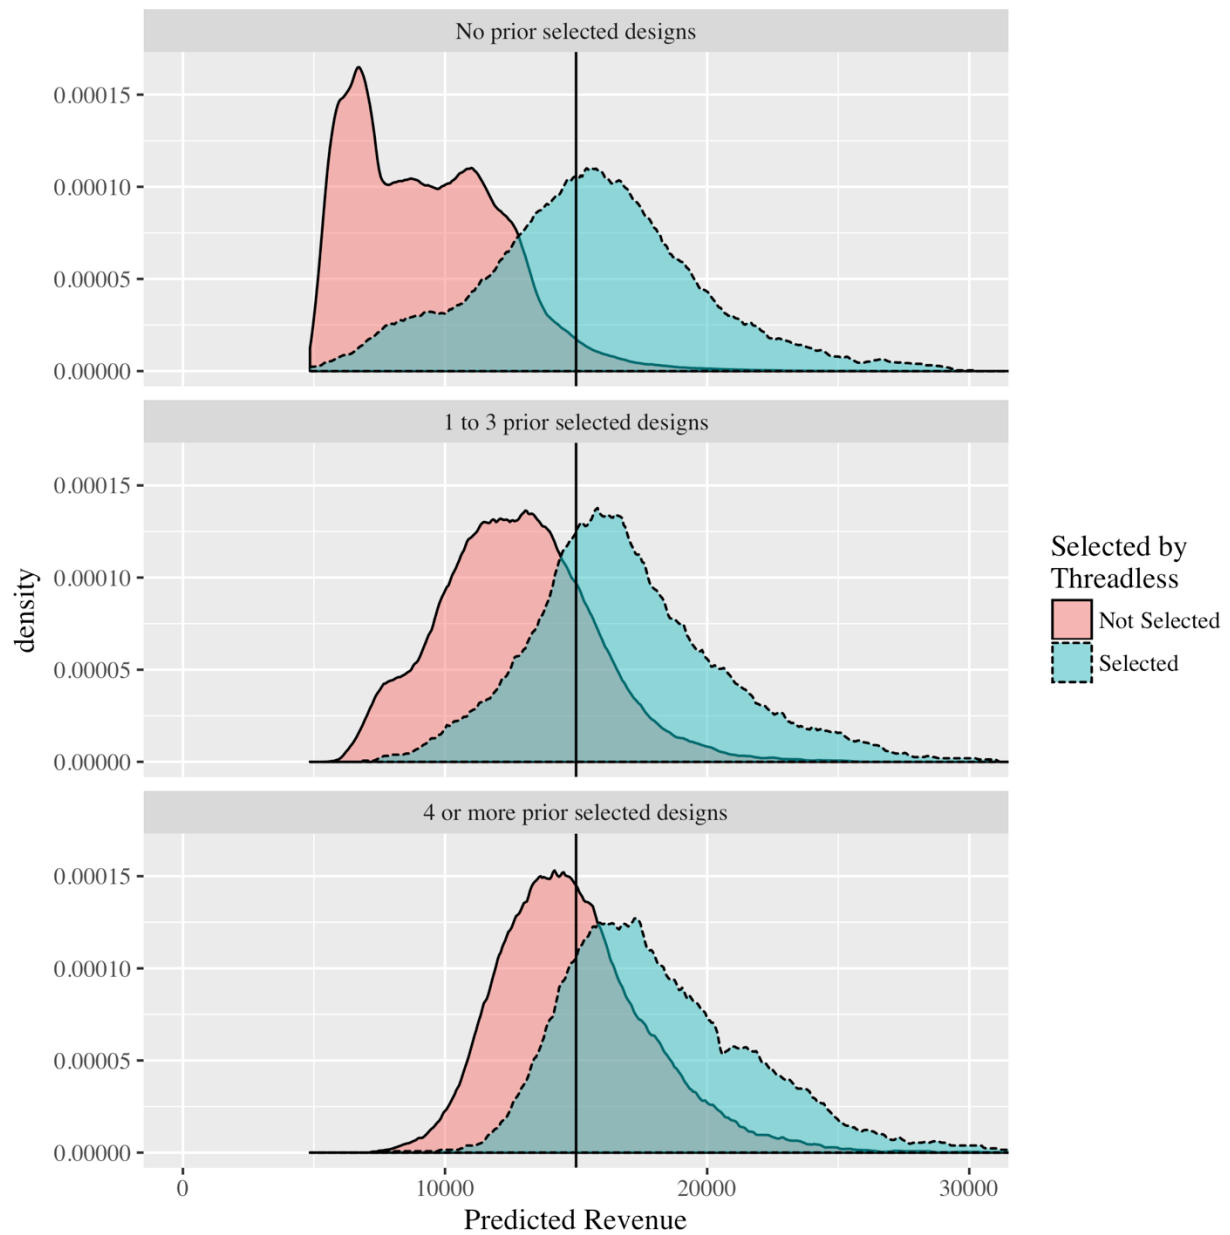


Figure notes:

1. Not Selected = Submissions that are not selected by Threadless.
2. Selected = Submissions that are selected by Threadless.
3. No prior selected designs = Users who have not had a design selected by Threadless.
4. 1-3 prior selected designs = Users who have had between 1 and 3 designs selected by Threadless.
5. 4 or more prior selected designs = Users who have had 4 or more designs selected by Threadless.

Figure 3: Quantile-Quantile Plot of the Predicted Revenue of Designs by Stars and Designs by Plebeians




Figure notes:

1. Not Selected = Submissions that were not selected by Threadless.
2. Selected = Submissions that were selected by Threadless.
3. Predicted Revenue for Designs by Stars = Predicted revenue of designs from Stars.
4. Predicted Revenue for Designs by Plebeians = Predicted revenue of designs from Plebeians.