Social Psychology Quarterly
2016, Vol. 79(2) 168–180

DOI: 10.1177/0190272516645603
http://spq.sagepub.com

Is Popular More Likeable? Choice Status by Intrinsic Appeal in an Experimental Music Market

Freda B. Lynn[1], Mark H. Walker[2], and Colin Peterson[3]

Abstract

There is widespread agreement from many areas of status research that evaluators' judgments of performances can be distorted by the status of the performer. The question arises as to whether status distorts perceptions differently at different levels of performance quality. Using data from the Columbia Musiclab study, we conduct a large-scale test of whether the effect of popularity on private perceptions of likeability is contingent on songs' intrinsic appeal. We discover that choice status (i.e., popularity) can boost perceptions of a song's likeability but only for songs of lower quality. In effect, the likeability halo created by popularity is one mechanism for why it is that "bad" songs can sometimes become more successful than songs that are intrinsically more appealing. But this same mechanism does not explain why "good" songs sometimes turn into superstars. This study suggests that status theories be refined to consider heterogeneous effects.

Keywords

popularity, social influence, status and quality, cultural markets, choice status

To study status advantage is really to study the interplay between *status* and *quality*. Often discussed as the Matthew Effect (DiPrete and Eirich 2006; Merton 1968), a core theme is that high status actors benefit because they "get more for the same." Experimental research, for example, clearly documents that evaluations of an actor's quality changes in the presence of certain status markers, such as motherhood (Correll, Benard, and Paik 2007), gender (Moss-Racusin et al. 2012), and sociometric status (Lynn et al. 2016). In broad strokes, this behavior is posited to occur because evaluators' perceptions of a candidate are affected not only by actual performance quality but also by the performer's status position (Berger et al. 1977).

A central question that arises then is the form of this weighting process: how much does status matter relative to

[1]University of Iowa, Iowa City, USA
[2]Louisiana State University, Baton Rouge, LA, USA
[3]Stanford University, Stanford, CA, USA

Corresponding Author:
Freda B. Lynn, University of Iowa, W143 Seashore Hall, Iowa City, Iowa 52241, USA.
Email: freda-lynn@uiowa.edu

quality? In the present study, we focus on the fundamental and yet relatively ignored issue of whether the effect of status is contingent on quality (see Foschi, Sigerson, and Lembesis 1995). For example, in the creation of the expectation states paradigm, Berger et al. (1977) laid out a formal procedure for how different task-relevant and status characteristics would combine to form performance expectations, where the form of aggregation was assumed to be consistent across all levels of task proficiency (see Foschi et al. 1995:336–37). More recently, Gould (2002) developed a generalized model of the Matthew Effect in which *i*'s perception of *j*'s quality—a precursor to whether *i* decides to "attach" to *j* in the future (e.g., adopt, hire, purchase)—is calculated as a weighted combination of *j*'s "true" quality and *j*'s status, which he operationalized with respect to the extent to which *j* is already sought after or deferred to by others in the market (for similar theories, see Becker 1991; Cialdini 1984; Hedström 1998; Podolny 1993). But neither Gould (2002) nor those building from his model of hierarchy formation (Lynn, Podolny, and Tao 2009; Manzo and Baldassarri 2015) explored whether the weight given to status might differ at high and low levels of quality. If there is heterogeneity in how status can distort perceptions of quality (by level of actual quality) and yet we assume that there is not, this oversight diminishes our predictive capability, and we ignore an opportunity to refine our theoretical understanding of why status information matters to decision makers.

The extant literature suggests that status markers are more likely to shape how evaluators perceive low or average quality performers compared to top performers. For example, with respect to the effect of gender on applicant evaluation, Foschi et al. (1995) report a reverse double standard that favors the female applicant when the applicant is "excellent"; however, female applicants are typically disadvantaged when candidates are average performers (Foschi and Valenzuela 2012; Heilman, Martell, and Simon 1988; but see Benard and Correll 2010). Similarly, with respect to choice status, Lynn et al. (2016) suggest that lower-skilled candidates can benefit more from being presented as highly sought after relative to higher-skilled candidates.[1]

In the present study, we take a step toward unifying these insights by conducting a large-scale study of whether the effect of status on perceptions of quality is moderated by the underlying quality of the performance itself. The key limitation of current studies (see Foschi and Valenzuela 2012; Lynn et al. 2016) is that the number of performance levels under investigation has generally been limited to just two (e.g., high and low performers). The data we employ here come from a large, experimental market (Salganik 2008) and allow us to test the interaction over a broad spectrum of quality levels. Specifically, we test whether the effect of *choice status* on *perceived likeability* varies by *intrinsic appeal*.

PREVIOUS RESEARCH

With regard to choice status, prior research suggests that the effect of status will be more pronounced for lower rather than higher performers (see Foschi et al. 1995). The literature offers two rationales for this hypothesis. First, to the extent that positive herding is more likely than

[1]In addition, although not a comparative study of different quality levels, Kovács and Sharkey (2014) find that for the very best books in the market (i.e., nominees for a prestigious award), books that actually win the award do not experience an increase in ratings; in fact, ratings tend to decline faster in the post-award period compared to books that were nominated but did not win.

negative herding (Muchnik, Aral, and Taylor 2013), popularity information is more likely to boost rather than erode favorability (for a given performance level). Because of this asymmetrical effect, and assuming that there is a theoretical ceiling on what constitutes performance excellence, this means that low quality performers have more to gain from status distortion than excellent performers. For example, if a candidate's true skill or performance level is at the 85th percentile of the distribution, the opportunity for status to upwardly distort is much smaller than if the candidate is at the 45th percentile.

A second line of thinking posits that evaluators are effectively more "motivated" to incorporate status information at lower levels of quality. The key idea is that in the presence of "indisputable evidence" establishing "merit," evaluators no longer need to call upon the information embedded in status characteristics to help them form judgments (e.g., Benard and Correll 2010; Foschi and Valenzuela 2012; Heilman et al. 1988). That is, above a certain threshold of performance ability, status bias will not be activated, which is to say that no weight will be given to status information in the formation of quality perceptions. Conversely, when performance information is not available or is demonstrated to be average or of lower quality, doubt arises as to the candidate's worthiness, and status information is called in to help adjudicate this doubt. In summary, both because (1) evaluators may be more motivated to rely on status cues to judge lower quality performances and (2) there is "more room" for choice status to elevate low versus high quality performances, choice status cues should ultimately play a greater role in boosting perceptions of lower quality performances compared to higher quality performances.

THE MUSICLAB EXPERIMENT

A decade ago, Salganik, Dodds, and Watts (2006; see also Salganik and Watts 2008) published what has become a landmark experimental study on social influence. The Musiclab project consisted of four experiments that drew over 27,000 participants in 2004 and 2005. On the website, participants could listen to—then rate and download—any of 48 songs, sampled in any order. The key manipulation was whether participants had access to popularity information: those in the control conditions were given only the names of the songs and bands, whereas those in the social influence conditions could see the number of times each song had been downloaded by previous participants. A key finding from Salganik et al. (2006) is that popularity information generated greater market-level inequality in download distribution—a result that clearly implies that participants' adoption choices were, at least on average, positively influenced by what other participants before them had chosen.

The data have never been used, however, to examine the extent to which popularity actually affected private likeability ratings. We do know from aggregate measures (song-level or market-level statistics), that (1) popularity strongly predicted whether participants decided to listen to a song (Krumme et al. 2012; Salganik et al. 2006; Salganik and Watts 2008) and (2) a participant's private rating of a song's likeability was strongly associated with whether that song was ultimately downloaded (Salganik and Watts 2008). But the effect of popularity on private ratings of likeability has not been the subject of its own investigation. Does a song's choice status boost its likeability ratings, as Gould (2002) predicts? Moreover, is this effect the same for songs that are intrinsically "good" versus "bad"? In our empirical analyses, we analyze the data from Experiment 1 (weak popularity signal) and Experiments 2 and 3 (strong popularity signal) and find that popularity can boost perceptions of likeability much more for "bad" songs than "good" songs and that this interaction is not an artifact of a mere statistical ceiling.

Experimental Design

Participants were given the opportunity to download any number of 48 previously unknown "indie rock" songs for free (Salganik 2008). Experiment 1 began in October 2004 and ran for 69 days, and Experiment 2 ran for the 83 days immediately following Experiment 1; participants for Experiments 1 and 2 were recruited mainly from a teen interest website. Experiment 3, which recruited participants from a different population, began in March 2005 and ran for 23 days (for details, see Salganik 2008).[2]

Upon arrival, participants were randomly assigned to one of two conditions: independent versus social influence. In Experiment 1, participants in the independent (control) condition were given a menu of songs formatted in a 16 × 3 rectangular grid (Salganik 2008: Figure 7). Songs were randomly assigned to 1 of the 48 positions in that grid and identified only by band name and song title. Those assigned to the social influence condition were further assigned to one of eight separate "worlds" (Salganik 2008: Figure 1). In each of these eight worlds, respondents were shown the names of the songs and bands as well as the number of times each song had been downloaded by others arriving before them in their specific world; all songs began with zero downloads and then accumulated downloads in real time. Thus, because of this multiple world design, each song has eight "lives" (i.e., we observe the counterfactual eight times), which means that we can isolate the effect of a song's popularity from its intrinsic characteristics.

The key difference between Experiments 2 and 1 is the format in which the 48 songs were presented to the participants. In contrast to the 16 × 3 grid used in Experiment 1, songs in Experiments 2 and 3 were presented in a single column (48 × 1) and rank ordered by the number of times they were previously downloaded, with the most popular song listed at the top (Salganik 2008: Figure 8). Thus, because popularity is coupled with a visibility advantage in Experiment 2, the popularity signal is stronger in Experiment 2 than in 1. Experiment 3 is identical to Experiment 2 except that Experiment 3 draws from a different population and contains only two social influence worlds instead of eight, which drastically reduces the amount of variability on popularity per song.

In each experiment, decision making is structured as a (one-way) three-step sequence: listen⟶rate⟶download. Participants were able to listen to any song simply by clicking on it. While the song was playing, participants were asked to rate the song's likeability on a five-star scale. After submitting a rating, participants were given the opportunity to download the song (yes or no). For our purposes, our analysis examines all of the song ratings that are available in each experiment; Table 1 provides descriptive statistics both at the person and ratings level.

Measures

Song popularity. A song's popularity is defined as the number of times it was previously downloaded in the participant's

[2]Experiment 4 is not analyzed here. Compared to 1 through 3, Experiment 4 contains a layer of complexity that complicates the interpretation of findings regarding the interaction of interest.

Table 1. Descriptive Statistics

	Experiment 1		Experiment 2		Experiment 3	
	M	SD	M	SD	M	SD
Ratings						
Private likeability rating (5-star scale)	3.0	1.2	3.3	1.2	2.7	1.0
Mean number of downloads at time of rating	7.9	5.3	10.3	6.2	11.0	6.3
N	19,954		18,229		10,131	
Raters						
Number of listens	7.2	9.7	6.8	9.4	10.5	10.9
Number of ratings	6.8	9.2	6.2	8.7	9.8	10.2
Number of downloads	2.2	4.5	2.7	5.8	2.0	4.9
Percentage male (1 = yes, 0 = no)	65		26		63	
Age	22.4	9.8	19.3	7.4	30.8	11.6
Percentage ever downloaded music online (1 = yes, 0 = no)	63		64		75	
N	2,951		2,953		1,035	

world. Because downloads accumulate in real time, the download distribution seen by the participant is contingent upon time of arrival to the site. For example, the 5th person to join a world, by definition, will see counts that range (at most) from 0 to 4, whereas the 85th person to participate could potentially see downloads that range from 0 to 84. Thus, for robustness purposes, we model popularity two different ways: as raw download counts as well as counts converted into *z*-scores. This transformation preserves the relative standing of songs but accounts for differences in the average number of downloads and the spread in downloads. For example, for the 85th person, the difference between 4 downloads might seem trivial whereas for the 5th person, a lead of 4 might seem substantial.

Perceived likeability. When listening to a song, participants were asked to rate its likeability using a five-star rating scheme: I love it = 5, I like it = 4, It's OK = 3, I don't like it = 2, I hate it = 1. The distribution of stars was approximately normal in each experiment. See Table 1 for a descriptive summary of the distribution by experiment.

Intrinsic appeal. To measure a song's "true" or intrinsic appeal, we calculate the average rating for each song from the corresponding control condition (i.e., the one condition in each experiment without popularity information).[3] The distribution of intrinsic appeal for the 48 songs varied slightly across experiments: Experiment 1 (M = 2.9, SD = .40), Experiment 2 (M = 3.2, SD = .38), Experiment 3 (M = 2.5, SD = .37). Experiments 1 and 2 are drawn from the same population, and not surprisingly, "natural" tastes are highly correlated (r_{12} = .86). Experiment 3 is drawn from a different population, and its tastes are not as strongly correlated with the first two (r_{13} = .78; r_{23} = .53).

Decision order. Respondents may change their standards for likeability over the course of rating many songs. With the Musiclab data, we know the order in which songs were selected for listening and thus the order that songs were rated. See Table 1 for a descriptive summary of the number of songs rated in each experiment. A large proportion of

[3]See Salganik, Dodds, and Watts 2006 as well as Salganik and Watts 2008 for an alternative formulation of intrinsic appeal.

raters in each experiment rated ten or more songs (58 percent, 57 percent, and 68 percent in Experiments 1, 2, and 3, respectively).

RESULTS

The Musiclab data are well suited to test for an interaction between popularity and intrinsic appeal given the extent of variation observed for both intrinsic appeal as well as popularity. The dotplots in Figure 1 illustrate this for Experiment 1. Each vertical cluster corresponds to a song, anchored at its level of intrinsic appeal (x-axis). Each dot represents a respondent, and the y-axis is the level of popularity attached to the song when the respondent listened to it. Figure 1 clearly shows that (1) intrinsically more appealing songs are likelier to become popular (Salganik et al. 2006; Salganik and Watts 2008) but that (2) regardless of intrinsic appeal, each song is seen by respondents at a variety of popularity levels. These data are thus akin to an experiment with 48 levels of true quality "crossed" with a substantial range of popularity levels.

The key empirical questions are whether the popularity of a song affects ratings of its likeability and whether this effect varies by the level of intrinsic quality of that song. We begin with a series of fixed effects ordinary least squares (OLS) regression models to predict the number of stars a song was rated (i.e., private perceptions of likeability) as a function of a song popularity, the song's intrinsic appeal, the interaction between popularity and appeal, person fixed effects, and decision order.[4] Clustered standard errors are used to account for the fact that ratings are nested with 48 songs. A person fixed effects approach means that we are controlling for all time-constant sources of person-level variation and thus estimate the effect of popularity on song ratings by leveraging within-person variation in song ratings.

Table 2 summarizes the results for Experiment 1 with popularity measured as raw download counts (divided by ten) in M1 and z-scores in M2. The results indicate that (1) intrinsic appeal has a major effect on likeability perceptions, which is not surprising, and (2) the extent to which popularity can boost likeability is substantially greater for lower quality songs. For example, with respect to M1, the effect of popularity on likeability for a "true" four-star song is estimated to be $.423 - (4 \times .096) = .039$ (i.e., adding ten more downloads is expected to generate roughly a zero-star increase in likeability). However, for a song with two-star intrinsic appeal, the effect of popularity is meaningfully larger: $.349 - (1 \times .078) = .23$, which corresponds to nearly a quarter star bonus in perceived likeability for ten additional downloads. Note that the same interaction is detected when popularity is measured with respect to z-scores of download counts, which hones in on relative popularity (M2). This pattern of results is virtually the same for Experiments 2 and 3 (M3 and M4, respectively). In sum, even when the popularity cue is more visible, the extent to which popularity boosts likeability is still constrained by the level of intrinsic appeal.

A question that arises is whether this interaction effect is an artifact of a statistical ceiling. That is, in contrast to the theoretical ceiling discussed earlier, perhaps there is no effect of popularity on ratings for "good" songs simply because the ordinal rating scale is capped at five stars ("I love it.") and thus there is literally no room for likeability to increase.

[4]The ordinary least squares results are substantively the same as those from an ordinal logistic regression as well as a logistic regression predicting the probability of receiving five stars versus fewer than five stars.

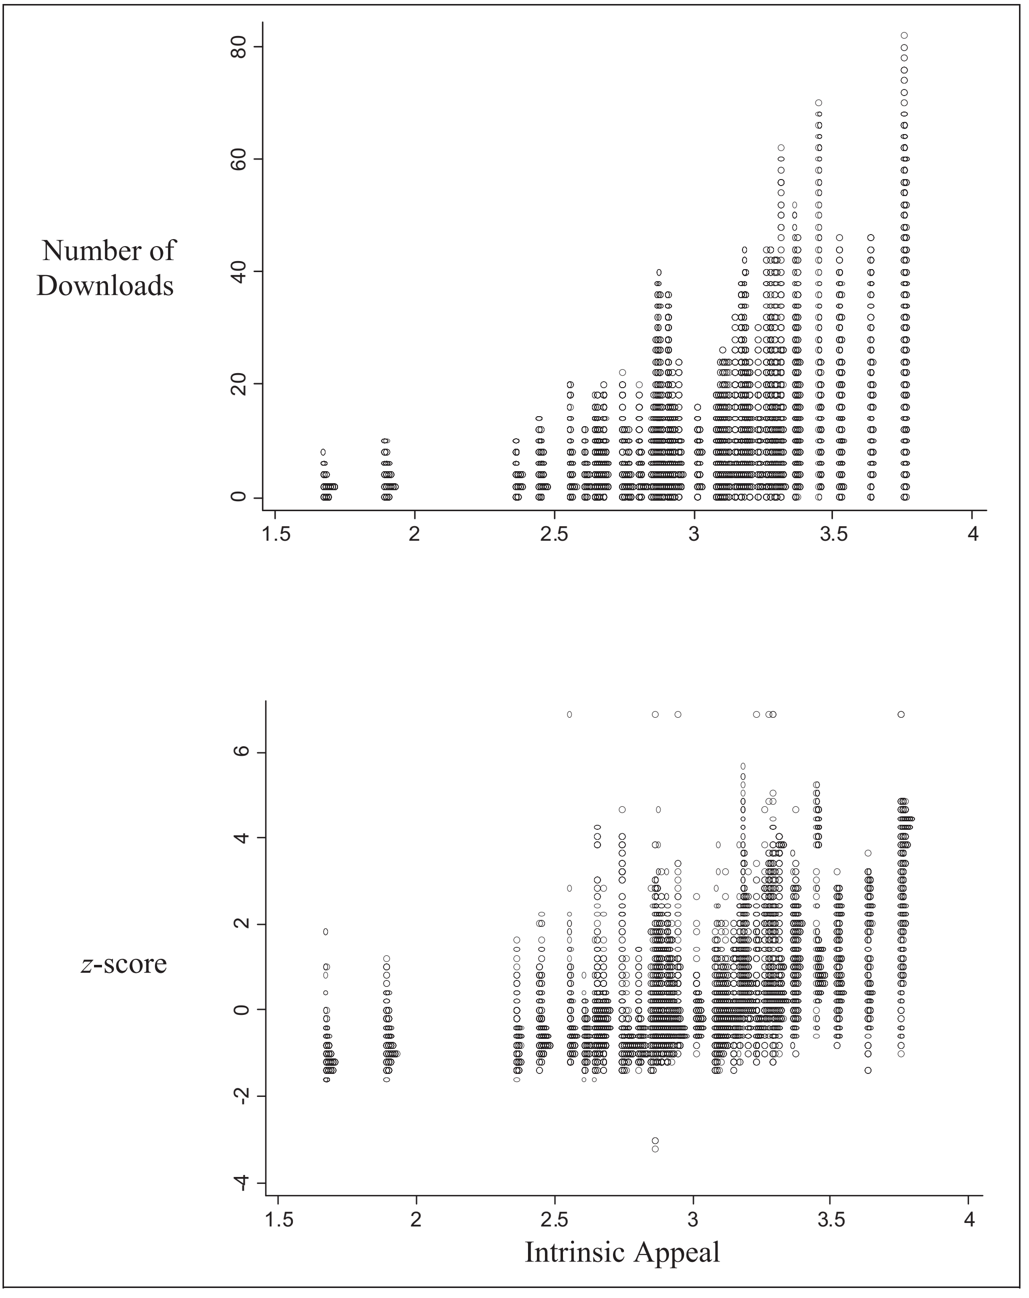

Figure 1. Distribution Plot of Popularity Levels Over Intrinsic Appeal, Experiment 1

We argue that this is unlikely the case here, given that there is still plenty of room for ratings to increase even for the very best songs. For example, for the song with the highest intrinsic appeal in Experiment 1 (song 25, Parker Theory, "She Said"), the proportion of five-star ratings received was still only 23 percent,

Table 2. Ordinary Least Squares Estimates Predicting Likeability Ratings

	M1	M2	M3	M4
Experiment	1	1	2	3
Measure of popularity	Downloads[a]	z-score	z-score	z-score
Popularity	.42***	.31***	.35***	.29***
	(.11)	(.06)	(.09)	(.08)
Intrinsic appeal	.82***	.72***	.74***	.85***
	(.04)	(.04)	(.04)	(.04)
Popularity × intrinsic appeal	–.10**	–.07***	–.08**	–.09**
	(.03)	(.02)	(.02)	(.03)
Decision order	–.00*	–.00	.01***	.00
	(.00)	(.00)	(.00)	(.00)
Person fixed effects	Yes	Yes	Yes	Yes
Constant	.49***	.88***	.69***	.37***
	(.10)	(.10)	(.14)	(.09)
Number of raters	2,951	2,951	2,953	1,035
Observations	19,954	19,954	18,229	10,131
Adjusted R^2 (percent)	40.2	40.5	42.1	35.8

[a]Number of downloads (divided by 10).
*$p < .05$. **$p < .01$. ***$p < .001$.

with a full 40 percent of raters giving it three or fewer stars.

To further investigate the nature of the interaction, we estimated, per experiment, 48 separate regressions (one for each song) to isolate the effect of popularity "within" each song. Table 3 illustrates this approach for three songs in Experiment 1. For example, in the control condition of Experiment 1, the song "Fear" by the band Forthfading was rated a 3.64 on average (the second most "naturally" appealing song in this market). In the eight social influence worlds, 501 respondents rated "Fear" and gave the song a wide range of ratings (see Table 3 part A). We then regressed these 501 ratings on rater characteristics (sex, age, and whether the respondent had ever downloaded songs online) as well as the contextual factors surrounding the decision, including the popularity level (z-score) at which "Fear" was seen, where "Fear" fell in the decision order of ratings, and the mean number of downloads seen by the respondent at the time that "Fear" was rated.[5] For the song "Fear," its popularity level had no positive or negative effect on likeability perceptions on average (b = .006, SE = .034). In contrast, for the other two songs in Table 3, both of which had lower intrinsic appeal compared to "Fear," popularity had a statistically and practically significant effect on likeability perceptions.

To see how this pattern generalizes across all songs and within each experiment, Figure 2 summarizes the point estimate for popularity from each of the 48 × 3 song-specific regressions. Circles and error bars in Figure 2 correspond,

[5]To illustrate, person 1132, a middle aged male in world 4 (Experiment 1), gave "Fear" three stars. It was the fourth song he listened to, and "Fear" had zero downloads at the time (z-score = –1.4). The mean number of downloads when 1132 came to the market was only 1.8. In contrast, person 6705, a teenage female in world 1, saw "Fear" as being downloaded 42 times previously (3.0 standard deviations above the mean number of downloads, which was roughly 13). "Fear" was the fourteenth song that she listened to and she gave it four stars.

Table 3. Illustration with Three Songs, Experiment 1

A. Descriptives			
Song title and band	"Fear" by Forthfading	"Bolshevik Boogie" by Sum Rana	"Miseries & Miracles" by A Blinding Silence
Intrinsic appeal[a]	3.64 stars (above average)	3.01 stars (average)	2.64 stars (below average)
Likeability ratings[b] (percentages)			
One star	5	5	11
Two stars	14	28	32
Three stars	26	35	27
Four stars	31	26	19
Five stars	24	5	11
B. Ordinary least squares estimates predicting likeability ratings	*b* (SE)	*b* (SE)	*b* (SE)
Popularity	.01	.34**	.49**
	(.04)	(.13)	(.15)
Decision order	.01	.01**	–.01
	(.01)	(.01)	(.01)
Mean number of downloads	.02	.01	.01
	(.01)	(.01)	(.01)
Age	–.02***	.00	.00
	(.01)	(.01)	(.01)
Male	.12	.14	–.30*
	(.11)	(.11)	(.12)
Ever downloaded music online	–.19	.12	.12
	(.10)	(.11)	(.12)
Constant	3.82***	2.64***	3.02***
	(.20)	(.20)	(.23)
Number of raters	501	365	349
R^2 (percent)	4.3	4.8	6.6

[a]Mean likeability rating in the control condition of experiment 1.
[b]Tabulated for all eight social influence worlds.
*$p < .05$. **$p < .01$. ***$p < .001$.

respectively, to the point estimate and 95 percent confidence interval for the effect of popularity "within" each song (e.g., if y = 2, this means that for a 1 standard deviation increase in popularity, there is a corresponding 2-star increase on average in likeability for that particular song). We see that in each experiment, the effect of popularity hovers around zero for the top third of the most intrinsically appealing songs. Only for songs of lower and average intrinsic appeal do we see cases in which popularity produced a likeability halo.[6]

DISCUSSION

The intuitive and scholarly expectation is that an object or actor's choice status will positively affect perception of its quality (e.g., Gould 2002; Lynn et al. 2016).

[6]See online appendix available at spq.sagepub.com/supplemental for additional analyses.

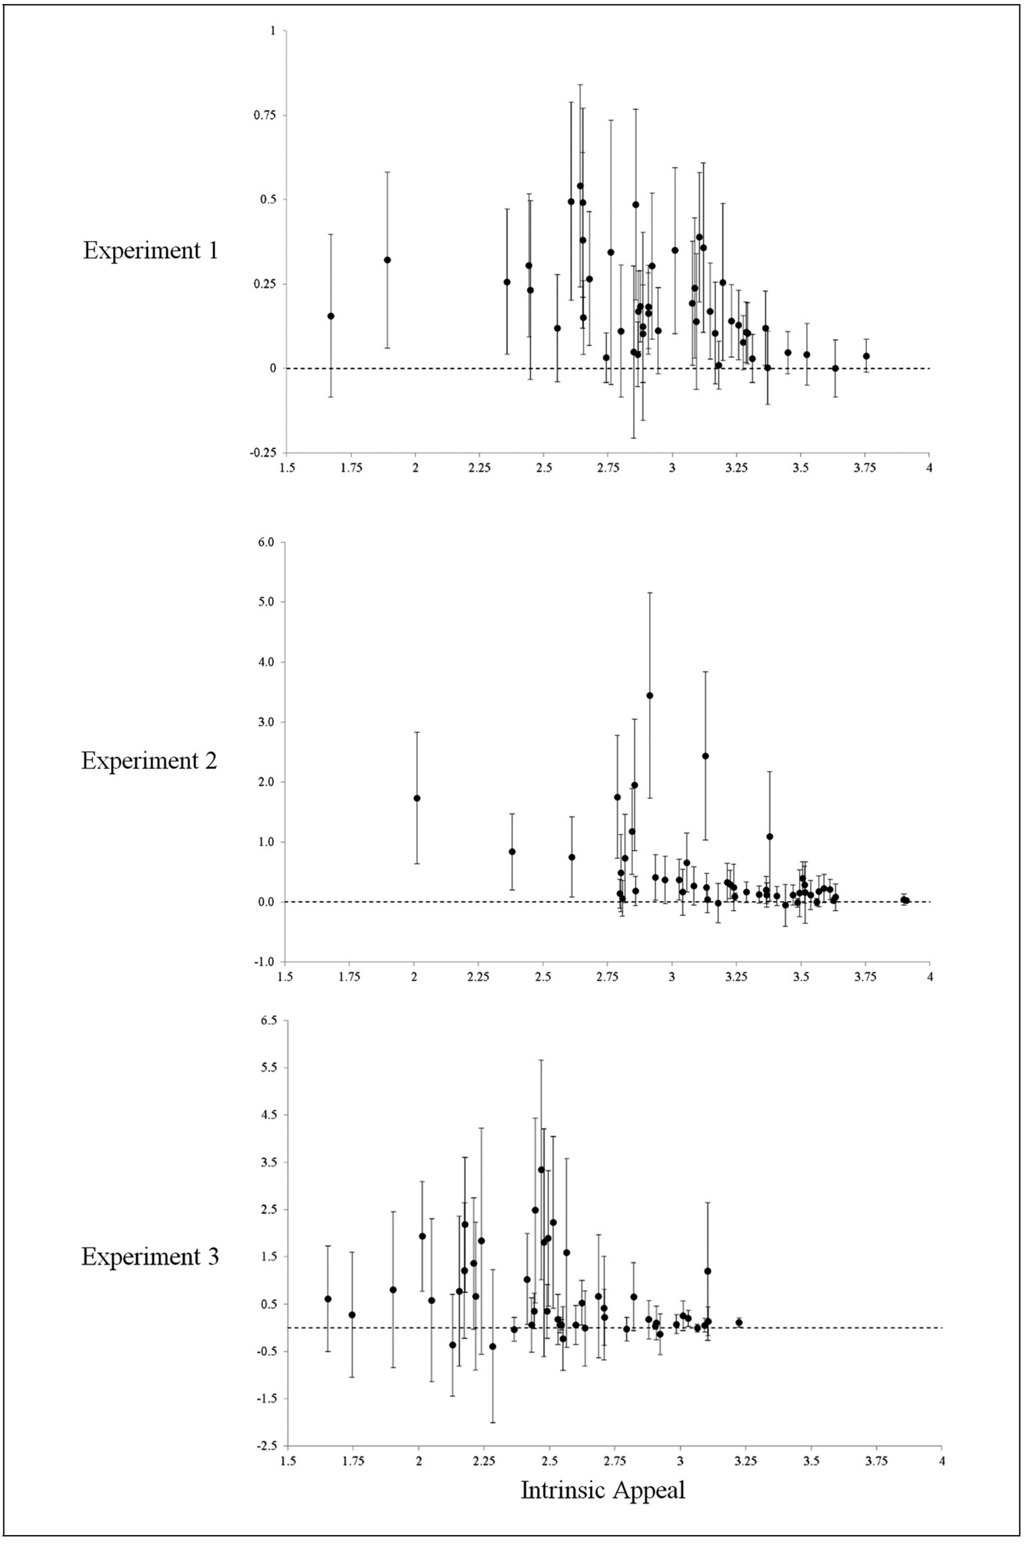

Figure 2. Coefficient for Popularity with 95 Percent Confidence Interval for Individual Songs

Here, we document an important contingency to this relationship, which is that the effect of popularity appears to be limited to objects of lower and average quality. Indeed, in this context, there appears to be no meaningful status advantage with respect to private perceptions of likeability at higher levels of quality: for songs that are intrinsically appealing, being popular (unpopular) does not boost (erode) perceptions of likeability. In contrast, popularity has the potential to upwardly distort the likeability of a song in the case of average and below average songs.

There are two major implications of our finding. First, researchers studying the Matthew Effect should take note that the mechanisms through which success breeds success may not be one size fits all. The basic assumption is that quality perceptions are formed as a weighted combination of actual quality and the extent to which the object is sought after by others (Gould 2002). Here, however, we find no evidence of this for objects above a certain quality threshold. The success of "good" songs was *not* fueled by the fact that popularity boosted private perceptions of likeability. Instead, "good" songs were more likely to become superstars in the market (see Figure 1) because (1) popularity cues heavily affected sampling (Krumme et al. 2012; see also Denrell and Le Mens 2007) and (2) likeable songs were more likely to be downloaded (Salganik and Watts 2008). In other words, popularity information was crucial to the creation of superstars *but not through upward pressure on perceived likeability*.

In contrast, it is the case that social influence on perceptions of likeability was able to elevate the success of songs of lower intrinsic appeal in the market. In other words, taking stock of what the crowd is doing can allow poorer quality performers to thrive in a way that is not possible when the crowd plays no role in decision making. This suggests a specific type of compensatory process with respect to choice status and perceptions of likeability, where a song can "survive" a low quality performance by somehow attracting the attention of many others. An empirical issue that needs further clarification, however, is whether choice status matters differently for poor versus average performances; Figure 2 and the online appendix tentatively suggest that the effect of popularity is greater for average compared to low performances, but the evidence on this point is far from conclusive.

Above all, future research is clearly needed to test whether this interaction exists in other empirical settings and with other kinds of status markers.[7] Although the interaction effect is robust in the Musiclab data and the data source itself is compelling, the market is specific to cultural goods and quality itself is not difficult to ascertain.[8] But deciding whether a song is likeable and worthy of downloading (especially in a cost-free market) is very different than deciding whether, say, a potential employee is skillful and worthy of hire. To this point, we suggest, for example, a simple extension of previous resume studies (Foschi and Valenzuela 2012; Moss-Racusin

[7]In this vein, another question to consider is the extent to which this moderation effect can or cannot be explained by the "inconsistency effect" discussed in the expectation states tradition (see Berger et al. 1977).

[8]In addition, the Musiclab is based on a nested decision structure, which means that the songs that participants rate have not been randomly assigned but rather are songs that the participant first decided to listen to (i.e., an evaluation conditional on selection). In contrast, participants in traditional status experiments typically do not have a choice as to which candidate they would like to rate. Future studies should explore how this difference in research design affects evaluation patterns with respect to different status markers.

et al. 2012) where resumes are distributed over a broad range of performance levels. We suspect that that the moderation effect investigated here may be diminished in market settings characterized by greater uncertainty around candidate quality (see Benard and Correll 2010).

SUPPLEMENTAL MATERIAL

Additional supporting information may be found at spq.sagepub.com/supplemental.

ACKNOWLEDGMENTS

We are indebted to Rebecca Durkee, Sarah Harkness, Matt Salganik, Michael Sauder, Brent Simpson, Ezra Zuckerman, and the participants of Iowa's Theory Workshop for their feedback at various stages of this project. A portion of this paper was presented at the annual Group Processes conference in 2014.

REFERENCES

Becker, Gary S. 1991. "A Note on Restaurant Pricing and Other Examples of Social Influences on Price." *Journal of Political Economy* 99:1109–16.

Benard, Stephen, and Shelley Correll. 2010. "Normative Discrimination and the Motherhood Penalty." *Gender and Society* 24:616–46.

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

Cialdini, Robert B. 1984. *Influence*. New York: William Morrow and Company.

Correll, Shelley J., Stephen Benard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112:1297–338.

Denrell, Jerker, and Gaël Le Mens. 2007. "Interdependent Sampling and Social Influence." *Psychological Review* 114:398–422.

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

Foschi, Martha, Kirsten Sigerson, and Marie Lembesis. 1995. "Assessing Job Applicants: The Relative Effects of Gender, Academic Record, and Decision Type." *Small Group Research* 26:328–52.

Foschi, Martha, and Jerilee Valenzuela. 2012. "Who Is the Better Applicant? Effects from Gender, Academic Record, and Type of Decision." *Social Science Research* 41:949–64.

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

Hedström, Peter. 1998. "Rational Imitation." Pp. 306–27 in *Social Mechanisms: An Analytical Approach to Social Theory*, edited by P. Hedström and R. Swedberg. Cambridge: Cambridge University Press.

Heilman, Madeline E., Richard F. Martell, and Michael C. Simon. 1988. "The Vagaries of Sex Bias: Conditions Regulating the Undervaluation, Equivaluation, and Overvaluation of Female Job Applicants." *Organizational Behavior and Human Decision Processes* 41:98–110.

Kovács, Balázs, and Amanda J. Sharkey. 2014. "The Paradox of Publicity How Awards Can Negatively Affect the Evaluation of Quality." *Administrative Science Quarterly* 59:1–33.

Krumme, Coco, Manuel Cebrian, Galen Pickard, and Sandy Pentland. 2012. "Quantifying Social Influence in an Online Cultural Market." *PLoS ONE* 7:e33785.

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

Lynn, Freda B., Brent Simpson, Mark H. Walker, and Colin Peterson. 2016. "Why is the Pack Persuasive? The Effect of Choice Status on Perceptions of Quality." *Sociological Science* 3: 239–263.

Manzo, Gianluca, and Delia Baldassarri. 2015. "Heuristics, Interactions, and Status Hierarchies An Agent-Based Model of Deference Exchange." *Sociological Methods & Research* 44(2):329–87.

Merton, Robert K. 1968. "Matthew Effect in Science." *Science* 159:56–63.

Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. 2012. "Science Faculty's Subtle Gender Biases Favor Male Students." *Proceedings of the National Academy of Sciences* 109:16474–79.

Muchnik, Lev, Sinan Aral, and Sean J. Taylor. 2013. "Social Influence Bias: A Randomized Experiment." *Science* 341:647–51.

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

Salganik, Matthew J. 2008. *Data from Success and Failure in Cultural Markets*. Retrieved April 11, 2016 (http://opr.princeton.edu/archive/cm/data_release.pdf).

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science* 311:854–55.

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

BIOS

Freda B. Lynn is an associate professor of sociology at the University of Iowa. Her research interests include inequality, status dynamics, and networks. In her current work, she explores individual-level perceptions as the building blocks of hierarchical structures. Some of her previous research on status has been funded by the National Science Foundation and appears in the *American Journal of Sociology* and the *Annual Review of Sociology*.

Mark H. Walker is an assistant professor of sociology at Louisiana State University. His research focuses primarily on identity processes, social networks, and mental health outcomes. His research has appeared in *Social Psychology Quarterly, Society and Mental Health*, and *Sociological Science*.

Colin Peterson is a graduate student in the Department of Sociology at Stanford University. His research interests include inequality, social psychology, networks, and computational social science.