Social Influence Leads to the Formation of Diverse Local Trends

ZIV EPSTEIN, MATTHEW GROH, ABHIMANYU DUBEY, and ALEX PENTLAND, MIT Media Lab, USA

How does the visual design of digital platforms impact user behavior and the resulting environment? A body of work suggests that introducing social signals to content can increase both the inequality and unpredictability of its success, but has only been shown in the context of music listening. To further examine the effect of social influence on media popularity, we extend this research to the context of algorithmically-generated images by re-adapting Salganik et al's Music Lab experiment. On a digital platform where participants discover and curate AI-generated hybrid animals, we randomly assign both the knowledge of other participants' behavior and the visual presentation of the information. We successfully replicate the Music Lab's findings in the context of images, whereby social influence leads to an unpredictable winner-take-all market. However, we also find that social influence can lead to the emergence of local cultural trends that diverge from the status quo and are ultimately more diverse. We discuss the implications of these results for platform designers and animal conservation efforts.

CCS Concepts: • **Human-centered computing** → **Web-based interaction**; **Empirical studies in collaborative and social computing**.

Additional Key Words and Phrases: platform design, social influence, crowdsourcing

ACM Reference Format:

Ziv Epstein, Matthew Groh, Abhimanyu Dubey, and Alex Pentland. 2021. Social Influence Leads to the Formation of Diverse Local Trends. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 409 (October 2021), 18 pages. <https://doi.org/10.1145/3479553>

1 INTRODUCTION

The explosion of information contained on modern online platforms requires users to use heuristics to both efficiently search through this information, and to make informed decisions. One such heuristic is the social signals of how other users have engaged with the platform. *Social influence* occurs when the decisions of a user are impacted by those of other users [10], and has been shown to be a key design dimension for contexts as varied as health behavior [9], political engagement [5], collective behavior [41], online book purchasing [8], food ordering [30], and digital news engagement [46, 62]. The ubiquity of social influence suggests how crucial a factor it is for platform designers seeking to jointly optimize for the quality and diversity of content online [28].

Perhaps the most influential study on how social influence and information hierarchy impact online platforms is that of Salganik et al. [54], informally dubbed the “Music Lab” experiment. In this study, the authors created an “artificial cultural market,” where participants could listen to and download previously unknown songs. Critically, some participants were provided a layout which displayed information about previous participant's choices, while the others had no such

Authors' address: Ziv Epstein, zive@mit.edu; Matthew Groh; Abhimanyu Dubey; Alex Pentland, MIT Media Lab, 75 Amherst Street, Cambridge, MA, 02143.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).
2573-0142/2021/10-ART409. <https://doi.org/10.1145/3479553>

knowledge. This experimental design allowed for causal identification of the role of social influence on both an individual's propensity to download songs, and on the dynamics of the ecosystem as a whole. In particular, Salganik et al found that introducing social influence increased the inequality of song success, as defined by the number of times they were downloaded. This suggests a cascading "winner take all" phenomenon, whereby social influence increased the availability of songs that were perceived as successful by past participants. They also found that social influence increased the *unpredictability* of success, as defined by the variation in a song's success across the worlds in an experimental condition. From these two findings, the authors infer that the underlying quality of a song only partly determines its final success: social influence causes a snowball effect that results in the emergence of local preferences.

The Music Lab experiment called to attention what is at stake when designing social influence into online platforms, but it remains unclear how to apply its conclusions to the design of modern social media platforms. For one, the original paper did not specify any mechanism or model to explain how social influence operates [39]. Moreover, it is unclear how the findings in the context of music translate to other forms of media such as images.

In the present paper, we show a conceptual replication of the original Music Lab study in a context that is fundamentally different from music — the *in silico* evolution of AI-generated hybrid animals (which we call "ganimals") [16]. In particular, we ask if similar patterns of results are observed in such a different context. This allows us to test the generalizability of the Music Lab to adjacent contexts (for a full characterization of the similarities and differences between the Music Lab and the present study, see Related Work). To do so, we built MEET THE GANIMALS, an online platform where users could generate and curate their own ganimals. Previous work introduced this platform as a "casual creator," and evaluated how its random stimulus approach can efficiently search the possibility space of a GAN generator [16]. Rather than focus on the usability of the platform itself, this work introduces an alternative layout to the "Feed 'Em" page of the system (see Figure 1), and presents experimental results on the impact of such design interventions in the field.

As a collection of images of synthetically generated hybrid animals, ganimals represent a unique context to study social influence. Interpolations of image-based GANs are a new form of media about which, due to their unfamiliarity, most people have no preconceived notions. In studies of social influence, previous knowledge of the content can introduce a key confound. Since ganimals are uniformly novel, both because of both the novelty of GAN-technologies, and also the vast possibility space of potential hybrid animals (see [16] for a characterization of this possibility space), here we can study social influence excluding the potential confounding factor of previous experience (see Section 2.2 for a full discussion). While the images as a whole are unfamiliar, the component parts — bodies, faces, eyes, mouths, colors, backgrounds, positioning, etc. — are well-studied affectively salient features [15, 20, 22]. As such, we hypothesize that the emotional valence that ganimals induce (falling in the uncanny valley of cute/creepiness) is highly *subjective*, and therefore may be subject to social influence.

The use of ganimals also allows us to assess how the findings of the Music Lab might translate to the medium of images. Relative to their text-based counterparts, image-based social media is on the rise (especially during the COVID-19 pandemic) [32, 44, 51]. While images of ganimals of course differ from images on social media along many important axes, the process of rapidly searching through troves of emotionally salient and unfamiliar content and unconsciously deciding which to attend to (and engage with) may mirror some of the cognitive patterns of surfing social media [6, 12, 49, 61].

A final key ingredient of the Ganimals platform is that it allows users to annotate the ganimals with morphological features. These rich ganimal-level covariates allow us to directly quantify the diversity and divergence of this online media ecosystem.

This paper has five main contributions. First, we introduce the HCI community to the methods and results of the Music Lab study, and experiments that build on it. Second, we show a conceptual replication and generalization of the Music Lab to the entirely different context of images, which substantially increases the extent to which the HCI community can base theories and built systems on the original findings. Third, we employ morphological embeddings to provide in-depth insight into both the diversity and divergence of digital ecosystems. Fourth, we introduce a new visual display layout, called cloud view, which allows us to isolate the mechanistic features of ranked lists that drive the effects. Finally, we discuss how our findings and methodologies can be applied by systems designers to both quantify the emergent outcome of platform designs, and evaluate new visual layout designs.

2 RELATED WORK

2.1 Replications in Human Computer Interaction

In response to the growing concern within the HCI community of prioritizing “novelty” over “consolidation,” the contribution of replications have been re-articulated [24, 25, 29, 48, 65]. Recent HCI papers have replicated studies from online labor markets (like MTurk and Lucid) on topics such as misinformation [17], visualization [26, 31], input devices [19], and usable security [52], but replications of large scale, virality based field experiments are more infrequent. By replicating the Music Lab, we show that such replications are well-scoped and useful for designing new systems.

2.2 Music Lab Experiments

The Music Lab inspired a generation of experiments in artificial cultural markets designed to assess the impact of social influence and information design on collective behavior [1, 2, 27, 40, 46, 55]. Some have attempted to decouple social influence and item position, which were confounded in the original experiment [1, 27, 40]. Hogg and Lerman [27] found that the impact of position is twice that of social influence itself. Abeliuk et al. [1] found that ranking positively affects unpredictability more than social influence does, and that combining ranking by quality and social influence allows high quality stories to become “blockbusters.”

In a follow-up paper to the Music Lab, Salganik and Watts [55] seeded songs with false and arbitrary initial signals of popularity, and looked at how those initial signals affected the market equilibrium. They found that while certain songs had a “self-fulfilling prophecy,” the best songs were able to recover their popularity in the long run. They also found that the initial distortion of the market information reduced correlations between appeal and popularity, and led to fewer overall downloads. Building on this work, Shulman et al. [57] found that while it is hard to predict an item’s final popularity, “peaking” at early adopters provides a highly effective framework for predicting future success.

Antenore et al. [2] conducted a Music Lab experiment where they found no evidence of an effect of social influence. Critically, their experiment only contained 10 songs to “avoid as much as possible the interference attributable to choice overload.” The differences they observed are due to the fact that the small number of songs meant every participant could try each and every song, and did not need to use social signals as a heuristic to “avoid the high cognitive cost of exploration.” Yet this heuristic is a critical feature of social influence, since most markets of interest involve too many items to try every one (see Section 2.3, below). In addition, their experiment was not a web-based study and occurred in a computer lab under the direct supervision of the experimenter. This induced a focused mindset divorced from the actual cognitive context where most cultural markets occur (e.g. where people are distracted, overwhelmed and must rely on heuristics).

In contrast, like the Music Lab, our study recruited subjects from the internet via their intrinsic interest in the subject matter (not a financial incentive). We also took precautions to ensure the content participants saw would be unfamiliar to them (since previous knowledge would introduce a confound). The original Music Lab experiment went to great lengths to ensure the songs selected were unknown to subjects, such as restricting to bands that played in less than 10 states, played less than 15 concerts in the past 30 days, had less than 30k hits on their PureVolume page, and had not played at the Warped Tour. The authors themselves admit that their restriction criteria are “ultimately arbitrary.” By focusing on ganimals, we could be sure that the stimuli were inherently and uniformly unfamiliar without relying on ad hoc restrictions like the Music Lab did.

Our conceptual replication of the Music Lab has several key differences from the original study that are important to highlight. First and foremost, we focus on the domain of (AI-generated) images, rather than music. A critical theoretical difference between these mediums is that in the original Music Lab, participants first decided whether they wanted to listen to song based on social and other signals, whereas in the image context the impression of the media is immediate and inherently entangled with other signals (see Section 2.3 for a full characterization of this two-stage process, and how the image context relates). A second difference is we use user-annotated ganimal morphology to directly assess the divergence and diversity of the media environment. Finally, we also manipulate the visual display layout in order to isolate the components of ranked lists that drive the effects.

2.3 A theoretical model for cultural markets

Krumme et al. [39] observe that a market for songs involves a two-stage process: the participant first chooses which song they will listen to, and then after listening, decides whether or not to download that song. In the first stage, the only information the participant has to decide if they will listen to the song is the name of the song and band, and also the social signals if they are in the social influence condition. Krumme et al. [39] found that social influence is only present in the first stage, and that the probability a user downloads a song is conditionally independent to if they clicked on it [39]. This two-stage model explains the findings of Antenore et al. [2]: with ample time and only 10 songs, participants were able to “try” each and every song and thus social influence did not factor into the “buying” stage. Abeliuk et al. [1] use this formulation to derive a metric for quality – the conditional probability of downloading a song given that it was sampled – that they recommend for optimizing.

This “try and buy” model of cultural markets has also been mathematically characterized as a individual-level heuristic that results in collective Bayesian rationality [38]. In particular, individual agents locally utilizing a “try and buy” heuristic corresponds to a regret-minimizing solution to a population-level exploration-exploitation dilemma (e.g. Thompson sampling).

In the context of ganimals, where the content are images, we study *only the first stage* of this two-stage process. It is appropriate to compare this first stage of the process to the Music Lab, since it is in this stage (not the second) where social influence is present.

2.4 Quantifying diversity and divergence

A key aspect of social influence is its capacity to decrease the diversity of an information environment [21, 43, 58, 66]. Sunstein [59] has argued that social media has engendered a polarized culture where people do not seek out new information. The salience of group identity online can impact the perceived and actual diversity of the resulting ecosystem by fostering an in-group/out-group mentality [6, 66].

Recent work has shed light on design features that moderate the relationship between social influence and diversity. In the context of music listening, Holtz et al. [28] found that personalized

recommendation decreased diversity within users, but increased diversity across users. Pescetelli et al. [50] found that increased diversity increases collective intelligence for large groups (≈ 25 people) but decreases it for small groups (≈ 5 people). Lorenz et al. [43] found that social influence undermines the wisdom of the crowd by mitigating the diversity of the crowd's responses without improving upon its collective errors.

However, the impact of social influence on diversity is understudied in the Music Lab context. This is because the content in those artificial cultural markets does not typically include any covariates, so the authors focus their analyses only on engagement metadata, such as popularity. Ganimals, however, are annotated with their morphological characteristics, which in turn allows us to assess how the experimental conditions affect the distribution of these characteristics across worlds.

2.5 Information design for social influence

A growing body of work within HCI has explored how information design can impact how users interact with systems [13, 14, 35, 37, 63]. Toth [60] explored the modality of feedback in a group discussion paradigm, and found that 2D graphics can augment normative and inhibit informational social influence. Hullman et al. [33] showed that social signals affected graphical perception accuracy in a linear association task. They also demonstrate a cascade pattern, such that initial inaccurate guesses can erroneously affect the responses of subsequent participants. Romero found a substantial effect of early respondents in Doodle polls [53], whereby the first few respondents of a poll can dramatically influence the behaviors of subsequent respondents. In the context of online gift giving, Kizilcec et al. [36] found that receiving a gift causes individuals to give more gifts in the future, and that designing observability into a system made gift-giving more socially acceptable [36]. Sharma and Cosley [56] introduced a statistical procedure for distinguishing between personal preferences and social imitation behavior, and find that a large majority of user actions reflect personal preference rather than copy-influence on a music recommendation website. Wijenayake et al. [64] investigated how design features like user representation, interactivity, and response visibility impact conformity. They find not only main effects in differences in group size, task objectivity, and perceived self-confidence, but also interactions between interactivity and response visibility.

The Music Lab and its follow-ups have primarily focused on linear lists and grids of music [2, 54, 55] and scientific articles [1]. This standard display layout, also employed by social media platforms as a “newsfeed,” involves scrolling through large lists of content. We maintain the use of the list, but also introduce a new visual display, inspired by tag clouds and the designers outpost [37], called “cloud view.” In contrast to the Music Lab's two experimental conditions (independent and social influence), we cross those conditions with showing the participant either the standard ranked list view, or the alternative cloud view (see Section 3.1 for more details on the experimental design). By experimentally varying social influence and the type of layout, we can decouple the relative effects of each, which serves two critical purposes. First, it allows us to separate screen location from popularity information, which were confounded in the original Music Lab experiment (that is, for the social influence condition, the more popular items were both higher in the list, and designated by their popularity — here, the ecosystem view allows us to disentangle these two factors). Second, it allows us to see what results are dependent on having to scroll through all the ganimals individually, versus all of them being presented together.

3 METHODS

MEET THE GANIMALS is an online platform where individuals can generate and curate “ganimals” - AI-generated hybrid animals [16]. A schematic for the system is shown in Figure 1. Ganimals

are generated by blending animal categories in BigGAN [7] in a way that balances exploring new hybrids, and exploiting existing signals for ganimal quality (see [16] for a characterization of this algorithm). In the Discover 'Em page, participants could interact with the generated ganimals and breed their own. Once they found a ganimal they like, they could “discover” it by naming their ganimal and rating how cute/creepy/realistic/memorable it is on the Name 'Em page.¹ Discovered ganimals appeared in the Feed 'Em page, where users could “feed” (i.e. cast a vote for) the ganimals they liked the best. Separately, in the Catalogue 'Em page, participants could rate the morphological traits of the ganimals (see Section 3.3 for more details). Screenshots of and more information about all of the pages can be found in the Supplementary Materials.

Fig. 1. Schematic of the MEET THE GANIMALS architecture. Ganimals are generated via BigGAN and a process that balances exploring new and existing ganimals. Participants can discover and breed ganimals in the Discover 'Em page, and name and rate their favorites in the Name 'Em page. The experiment took place in the Feed 'Em page, which also includes 47 seed ganimals. Participants can also characterize ganimals in the Morphology Quiz. Adapted from [16] with permission.

3.1 Experimental Design

The experiment itself took place in the Feed 'Em page. As participants arrived to the platform, they were randomly assigned to one of four conditions (independent list, independent cloud, social influence list, social influence cloud) distinguished solely by the availability of information about the prior decisions of others, and the visual display of the ganimals in the Feed 'Em page.

In the independent conditions, each ganimal was displayed at the same size, whereas in the social influence conditions, the ganimal's size was proportional to the number of votes it had from previous participants, and displayed this number as well as its name (see Figure 2). Therefore, participants in the social influence conditions were provided a signal on the preferences of past participants, which they could use to make their own decisions. All users could “feed” (i.e. cast a vote for) the ganimals they liked the best, and the interface included the instructions: “Ganimals need food to survive. To feed a ganimal, click on its image. To learn more about that ganimal, click on its name.” For both the independent and social influence conditions, whenever a participant voted on a ganimal, it grew a bit larger. In the list conditions, ganimals are rank ordered and

¹these subjective signals are then recycled for future ganimal generation

displayed by number of votes in a grid which has two columns in desktop view and one column in mobile view. In the cloud view, animals are displayed in a spatial circle pack, with larger animals often (but not always) in the center.

Within each of the four experimental conditions, participants were randomly assigned to one of four “worlds” (for a total of 16 — see Figure 2), each of which evolved independently of the other fifteen. In particular, participants only saw animals discovered and votes cast by others in their world, and the ranking (for list view) and visualization (for cloud view) of animals was based only on votes in that world. A randomly chosen “seed set” of 47 animals was used to initialize each of the sixteen worlds (such that all worlds started with the same set of initial animals, see Section 1 of the Supplementary Materials for more information).

Fig. 2. Overview of our experimental design, with (stylized) screenshots of the Feed ‘Em page in the four conditions designated by the letters: A) independent list, B) independent cloud, C) social influence cloud, D) social influence list. There are four worlds for each the experimental conditions.

Irrespective of condition, participants could choose between 60 different animals to feed. At any given time, this set of 60 included 30 of the top voted animals in that world, and 30 of the most recently discovered animals in that world. The top rated animals were initialized with the seed set, but at most 30 of the them were shown to the user, and that number was much less once participants in a given world started voting (i.e. no animals with zero votes were displayed after thirty animals were voted on in that world).

3.2 Recruitment, reliability and robustness

From April 26th to June 26th 2020, 44,791 animals were generated, 8,547 animals were bred, and 743 animals were named by a total of 10,657 users. In the Feed ‘Em page 2,370 votes were placed on 434 animals by 549 users. Of these 549 users, 18% were on mobile, while 81% where on desktop, and they predominately hailed from Russia (27%), USA (23%), Ukraine (16%), and Japan (3%). We did not collect other demographic information.

Participants were recruited through word of mouth and social media, bolstered by a climate fiction (cli-fi) world-building campaign (learn more here: https://www.youtube.com/watch?v=IFc4nQK_5Q).

A critical part of our experiment was making sure there was no information contamination between worlds and conditions. We used cookies to ensure that each user would be placed in the same world if they returned to the website at a later time. To prevent contamination from exposure to other ganimals after they first experience the website (e.g. from elsewhere online), we only counted votes that occurred within 2 hours of the participants first visiting the site. To mitigate the impact of a few power users, we only counted the first 10 votes of each participant, as recorded through cookies.

We preregistered our primary hypotheses, primary analyses and sample size, which are available at <https://aspredicted.org/65nv7.pdf>.² Because the experiment received less media attention than anticipated due to Covid-19, we deviated from our preregistration in several ways. First, we stopped data collection early, but did not look at any results before doing so (due to the number of monthly active users, we would never reach 1K people — the number specified by our pre-registration). Second, this smaller N meant we were unable to block on cohort, since there was only a single cohort with its own set of worlds. As such, for all our analyses, we report results only for the first cohort, and also do not look at the interaction between treatment arms, as we are underpowered to do so with only 16 observations. Finally, the limited number of participants also meant that computing a joint distribution between popularity and a high-dimensional feature embedding for the diversity measure was under-specified. Therefore, we instead directly compute the entropy of the morphological traits, which is ultimately a more interpretable metric.

This study was approved by the MIT COUHES committee, and the code/data used can be found at https://github.com/zivepstein/ganimals_cscw.

3.3 Crowd annotation of ganimal morphology

Participants could also provide information about the morphology of ganimals. We worked with a professional zoologist to assemble 10 traits that would characterize the variability in possible ganimals (for the full list of traits, see Table 1). Within the “Morphology Quiz” section, the user was provided with 16 ganimals for each morphological trait, and was asked to select the ones that exhibited that trait. For each ganimal, we computed the average response for each morphological

Table 1. Morphological traits used.

Trait name	Question Asked	Trait name	Question Asked
Head	Does this ganimal have a head?	Size	Is this ganimal bigger than a house cat?
Eyes	Does this ganimal have eyes?	Underwater	Does this ganimal live underwater?
Mouth	Does this ganimal have a mouth?	Feathers	Does this ganimal have feathers?
Nose	Does this ganimal have a nose?	Scales	Does this ganimal have scales?
Legs	Does this ganimal have legs?	Hair	Does this ganimal have hair?

trait across responses. With a 1 coded as exhibiting the trait and 0 coded as not, the average response represents the likelihood or extent to which a given ganimal exhibits a given trait. 14,348 ratings were provided for 1,250 ganimals by 177 users of the MEET THE GANIMALS platform (48 of these raters were also participants in the experiment). Many of these 1,250 ganimals did not appear in the actual experiment, so we restrict our attention to the 449 ganimals that received at least one vote.

²Pre-registration is a framework that allows researchers to specify which analyses are *a priori* and which are *post hoc*. This strengthens the validity of statistical analyses, and can mitigate publication bias [47].

Table 2. World-level linear regression predicting inequality.

	Estimate	Standard Error	<i>t</i> -value	<i>p</i> -value
Social Influence	0.077	0.019	3.981	0.002
Cloud	0.066	0.019	3.379	0.002
Intercept	0.623	0.016	36.9	<0.001
Adjusted $r^2 = 0.628$, $N=16$				

21% of the 449 ganimals were missing at least one morphological trait (due to limits in data labeling), and among this subset, there was an average of 6.13/10 non-missing traits. Furthermore, multiple users often rated the same trait of a given ganimal: among this subset, an average of 2.4 users rated each non-missing trait. We find no statistical difference in the number of missing traits across conditions ($p = 0.351$ for social influence, $p = 0.175$ for information design).

4 RESULTS

We investigate the role of social influence and information design on five outcomes: inequality, unpredictability, diversity, divergence and engagement. As a replication of Salganik et al. [54], we find evidence that social influence increases inequality and unpredictability. We also find evidence that social influence increases the morphological divergence and diversity of worlds. For each set of analyses, we focus on the list view results since those are directly comparable to the literature, but also show the cloud view results for contrast.

4.1 Inequality

To make our results comparable to Salganik et al. [54], we also use the Gini Coefficient [4] to assess the concentration of the votes of the ganimals in each world. Figure 3 and Table 2 show the effects of social influence and layout on inequality, as measured by the Gini coefficient. For both the list and cloud displays, we find that worlds with social influence exhibit more inequality than the worlds where participants made independent decisions ($p = 0.002$, preregistered), which is a direct replication of the Music Lab experiment.

Fig. 3. Inequality of success for independent list (orange left), social influence list (blue left), independent cloud (orange right) and social influence cloud (blue right) worlds, with number corresponding to world. Dashed line corresponds to the average Gini coefficient of all worlds within a given condition.

As a preregistered robustness check, we use Fisherian randomization inference (FRI) to compute an exact p -value [34]. Fisherian randomization inference is a non-parametric approach to computing

p-values that does not require modeling assumptions about potential outcomes. To perform FRI, we create 10,000 permutations of the assigned world treatments, and recompute the t-statistic for each. We then compute a p-value by assessing the fraction of permutations that yielded t-statistics larger than the t-statistic observed in the actual data. We find that $p_{\text{fisher}} < 0.001$.

As a posthoc robustness check, we perform bootstrapping at the world level and count the fraction of bootstrap samples with mean Gini greater for social influence than independent. We find that in 100% of the 10,000 bootstrap samples, the mean Gini is greater for social influence worlds than independent worlds.

We also find a main effect of design, whereby worlds with the cloud display exhibited more inequality than the worlds with the list display ($p = 0.005$ for regression and $p_{\text{fisher}} = 0.0015$, preregistered, and the mean Gini is greater for cloud view worlds than list view worlds for 100% of the 10,000 bootstrap samples).

When restricting only to list view worlds, we find that social influence is significantly associated with inequality ($p = 0.004$, posthoc). This suggests a stronger association between social influence and inequality in the list view, where participants rely more heavily on social signals instead of scrolling through all 60 ganimals (versus cloud view where all are visible and thus social influence is a less important cue).

As additional robustness checks, we reran the main analyses restricting only to ganimals with one or more vote, and using additional measures of concentration (following the robustness checks of Salganik et al. [54]), which are reported in Section 3.1 of the Supplementary Materials. The results are similar across model specifications, except for layout design, where the effect of Cloud View is not significant for the Herfindahl index model (this may be because, unlike Gini or the coefficient of variation, the Herfindahl index is correlated with the number of ganimals in a given world.)

4.2 Unpredictability

To measure the unpredictability of each condition, we follow the Music Lab and compute the average difference in market share for that ganimal between pairs of realizations. The one critical difference in our setup is, unlike the Music Lab, not all ganimals necessarily appear in all worlds. Thus, we only consider pairs of worlds *for which that ganimal actually appeared*. In particular, we first compute the market share of votes $m_i^{(j)} = v_i^{(j)} / \sum_{\ell} v_{\ell}^{(j)}$ for each ganimal i in each world j . Then, we compute the average difference in market share for all pairs of worlds j and k within that condition S , and then average across ganimals to get a unpredictability score u for each pair of worlds:

$$u_{j,k} = \frac{1}{N} \sum_i \sum_{j \in S} \sum_{k \in S: j \neq k} \left| m_i^{(j)} - m_i^{(k)} \right|.$$

This gives us $\binom{4}{2} = 6$ unpredictability scores per condition. To reduce noise in our unpredictability estimates, we deviated from our preregistration plan in two ways. First, because the vast majority of engagement was with ganimals not in the seed set (71%), we considered all ganimals, not just seed ganimals. Second, since many ganimals were not seen and thus could not even be voted on, we restricted only to ganimals with at least one vote (see SI Section 3.2 for full justification for these changes, as well as additional analyses for our original measure).

The left panel of Figure 4 and Table 3 show the effects of social influence and layout on unpredictability, as measured by the average unpredictability across world pairs. We find that worlds with social influence exhibit significantly more unpredictability than the worlds where participants made independent decisions ($p = 0.005$, posthoc), but found no effect of the cloud layout ($p = 0.786$).

Fig. 4. Unpredictability of success for the four conditions, computing using all ganimals with more than one vote.

Table 3. Linear regression predicting unpredictability.

	Estimate	Standard Error	<i>t</i> -value	<i>p</i> -value
Social Influence	0.010	0.003	3.112	0.005
Cloud	0.001	0.003	0.275	0.786
Intercept	0.016	0.002	5.831	<0.001
Adjusted $r^2 = 0.252$, N=24				

When restricting only to list view worlds, we find a significant association between social influence and unpredictability for all ganimals with more than one vote ($p = 0.003$). This suggests a stronger association between social influence and unpredictability in the list view, where instead of browsing all ganimals, participants relied more on social signals.

4.3 Divergence and Diversity

The 48 songs from the original Music Lab experiment did not include a rich set of covariates, so the authors focused their analyses only on the songs' popularity. The MEET THE GANIMALS platform, in contrast, allows users to annotate ganimals across 10 morphological traits (see Figure 6 for the morphology of four exemplary ganimals, and Table 1 for more details). These ganimal-level covariates allows us to characterize and compare the kind of ganimals that evolved across worlds. In particular, we focus on the divergence of these features between worlds and the diversity of these features within worlds.

For each world, we fit a multivariate gaussian with mean μ_j and covariance Σ_j to the feature vectors of each ganimal in that world with one or more vote. Thus μ_j represents the average morphology of world j (e.g. its local trend). We use principal component analysis (PCA) to collapse these world-level average feature embeddings μ_j into a 2D space for visualization (offset such that the average feature embeddings of the seed ganimals corresponds to the origin). In this 2D space, the PCA parameters define orthonormal lines for each of the 10 morphological features. The map of the 16 worlds in this coordinate space, as well as exemplary ganimals, are shown in Figure 5. As shown, the majority of worlds centered around ganimals with eyes, a nose, a head, and that do not live underwater. Several worlds (all social influence) diverged and are centered around underwater ganimals without eyes, nose or a head.

Fig. 5. Morphological embeddings of the 16 worlds, relative to the morphology of the seed set of animals. Lines correspond to orthonormal projections for each of the 10 morphological features. The morphological embeddings of four exemplary animals is also shown.

Table 4. Linear regressions predicting diversity and divergence. $\cdot$ refers to $p \leq 0.1$, $*$ refers to $p \leq .05$, $**$ refers to $p \leq .01$, $***$ refers to $p \leq .001$. The value in parentheses under each coefficient is the standard error.

Outcome variable (DV)	Divergence (2D Euclidean)	Diversity
Social Influence	0.032 $*$ (0.014)	3.75 $**$ (1.50)
Cloud	0.008 (0.014)	-1.20 (1.50)
Intercept	0.041 $***$ (0.012)	4.55 $***$ (1.32)
N	24	16
Adjusted R^2	0.014	0.245

To quantify these visual intuitions, we measure the Euclidean distance between the 2-D points in Figure 5, and then compute the average distance between pairs of worlds within a given condition. These results are shown in Table 4. In the additive regression model, we find a significant effect of social influence on divergence ($p = 0.030$). We find no effect for display type ($p = 0.568$).

As additional robustness checks, we quantify divergence using two additional measures, 10-D euclidean distance and Fréchet distance, which are reported in Section 3.3 of the Supplemental Materials. We find a marginal ($p_{10D} = 0.053$) and no effect ($p_{\text{fréchet}} = 0.607$) of social influence, respectively, and a marginal effect of display type in both cases ($p_{10D} = 0.095$ and $p_{\text{fréchet}} = 0.053$).

However, when restricting only to list worlds, we find a significant effect of social influence on divergence for all three measures ($p_{2D} = 0.013$, $p_{10D} = 0.021$, $p_{f_{rechet}} = 0.021$). This suggests that in the list view, where scrolling through all the ganimals is cumbersome, social influence is an important cue. However in the cloud view, all the ganimals are easily available via a quick scan, so those worlds have high divergence regardless of social influence.

These results suggest that there is dramatic variation in morphology across worlds. But how does that compare to the variation in morphology within worlds? To answer that question, we calculate and compare the morphological diversities of each world. To compute the diversity of a given world j , we again use the multivariate gaussians with mean μ_j and covariance Σ_j fitted to the feature vectors of each ganimal in that world with one or more vote. Then, we calculate the entropy of that distribution:

$$h_j = \frac{1}{2} \log \det(\pi e \Sigma_j). \quad (\text{diversity})$$

Fig. 6. Left: Morphological features of four ganimals: an alligator/basenji hybrid (green), a kite/spider monkey hybrid (light blue), a hammerhead shark/macaw hybrid (dark blue) and a lynx/meerkat hybrid (brown). Center: Divergence across conditions, measured using 2-D Euclidean distance. Right: Diversity across conditions, measured by entropy of the multivariate gaussian fit to each world, with the dashed line corresponding to the imputed entropy of the seed set.

The average morphological diversity of the four conditions is shown in Figure 6. We find that worlds with social influence are more diverse than worlds where participants made independent decisions ($p = 0.02208$, $p_{\text{fisher}} = 0.0099$). We also bootstrap the diversity at the world level, and find that the mean diversity is greater for social influence worlds than independent worlds for 99.991% of the 10,000 bootstrap samples. We find no difference in morphological diversity between the list and cloud views for the main regression ($p = 0.3842$), FRI ($p_{\text{fisher}} = 0.8161$), or through bootstrapping (14.51% of bootstrap samples have more diversity in cloud view conditions than list view).

4.4 Engagement

We start by assessing the engagement across conditions, which we measured using the total number of votes each participant cast. We find that participants in worlds with the cloud design engaged with more ganimals than those in the worlds with the list design ($p < 0.001$, posthoc, $p_{\text{fisher}} = 0$ — the left side of Figure 7). We also find a main effect of social influence, whereby participants in worlds with social influence engaged with more ganimals than those in the independent worlds ($p = 0.0182$, posthoc, $p_{\text{fisher}} = 0.0081$). We also look at the position of the votes for each of the conditions, as shown on the right side of Figure 7. For both the independent and social influence list

Fig. 7. Left: Number of votes participants casts across conditions. Right: Distribution of the position of votes across the four conditions: independent list (A), independent cloud (B), social influence list (C) and social influence cloud (D). A and C show the distribution of votes over order (0 is the at the top of the page, 60 is at the bottom), and B and D show heatmaps of user votes. All are colored such that white corresponds to the most engagement, and purple the least.

Table 5. Linear regression predicting average number of votes for the four conditions, with robust standard errors clustered on world.

	Estimate	Standard Error	t value	p value
Social Influence	0.738	0.271	2.72	0.0182
Cloud	1.322	0.2608	5.084	<0.001
Intercept	3.208	0.202	15.860	<0.001
$r^2 = 0.0427$, $N=549$				

views, we see a power law which closely matches the availability plot (Fig 2) from Krumme et al. [39] and the position bias plot (Fig 7) from Abeliuk et al. [1]. For the independent cloud, we see a constellation of white and yellow dots scattered uniformly across the cloud view. This stands in contrast to the social influence cloud, where we observe clustering near the center, with satellites around the periphery. The circle packing algorithm placed most top voted animals near the center, (e.g. see the social influence cloud in Figure 2). This explains why worlds in the social influence cloud condition had the highest Gini coefficient on average: social influence induced users to click on top voted animals near the center, further increasing inequality.

5 DISCUSSION

For both inequality and unpredictability, we observed the same pattern observed in the Music Lab, that social influence can increase the inequality and unpredictability of the success of content. The pattern was the same across the list view conditions which mirrored the original layout, and the cloud view conditions, which had a different layout. This suggests that the role of social influence in increasing inequality and unpredictability is separable from item position, and generalizes well to different layout patterns, and contexts with image media. Indeed, the original two-stage process of the Music Lab may be less relevant to today's digital platforms, since now images are the primary content and the primary outcome of interest is engagement with those images. Thus, replicating the Music Lab results for engagement with images shows that they generalize to modern contexts.

Our results suggest that without social influence, many worlds converged towards a singular set of animal features. This status quo contained features that conform with morphological

conventions of quality (e.g. ganimals with eyes, a head, and dog-like features). Social influence, however, led to the rapid evolution of local cultures that dramatically diverged from this status quo, and that were ultimately more diverse. Many studies interpret the unpredictability that social influence induces to be a negative externality. Yet we found that this unpredictability corresponds to exploring novel areas of the possibility space, and led to more diverse and divergent local cultures.

These findings stand in contrast to recent results which suggest that social influence decreases diversity, and thus undermines the wisdom of the crowds [43]. It's important to note that the context in Lorenz et al. [43] involved objective performance tasks with ground truth, such as geographical facts and crime statistics. In such objective contexts, divergence from the status quo will "lead the herd astray" by skewing responses away from the crowd average. But in contexts like ours, where the popularity of novel cultural objects is inherently subjective, such divergence can actually be a boon.

The inclusion of the cloud view allowed us to isolate one possible mechanism of how social influence impacts divergence and diversity: the number of ganimals visible at a given time. Since the list view shows one or two ganimals at a time and requires scrolling to see more, social influence is an effective cue for identify ganimals to engage with. In contrast, cloud view shows a large number of ganimals at a given time. This may explain why cloud view worlds were less diverse than list view worlds, regardless of social influence, since ganimals with morphological features traditionally associated with popularity could easily grab user's attention. Indeed, since the size of the ganimals is so different between the two displays (quite large for list view and quite small for cloud view), it is hard to compare the two directly. This is one reason we focused primarily on the list view only, and future work should compare layouts that hold both the image size and number of images displayed constant across conditions.

Our work has implications for the design of social media platforms. While social influence can lead to a "winner take all" market where smaller actors suffer, it can also lead to the rapid evolution of unexplored, diverse trends. Designers of social media platforms should use social influence responsibly to foster more heterogeneous notions of quality. In particular, designers must recognize when divergence is desirable (e.g. fashion or other cultural artifacts) and when it is not (e.g. for objective facts) and design interfaces that are appropriate to that situation. For instance, in the case of subjective culture and trends, presenting feedback on the proportion of similar people may help promote divergence, whereas in situations with objective facts, feedback such as histogram of other people's opinion might be more appropriate. In addition, the cloud view represents a new paradigm for newsfeed design that may be appropriate in contexts where it is desirable to mitigate social influence or the availability bias (such as objective contexts, as in Lorenz et al. [43]).

Our work also has implications for how animal morphology relates to engagement. It has been shown that people have learned preferences for animal morphology on an evolutionary timescale and during development [11, 23, 45] and indeed many of the most popular ganimals have morphological traits of common pets like cats and dogs (see panels C and D of Figure 2). However, our world-level experimental design allows us to decouple the role of visual familiarity/innate preference and social influence in ganimal popularity. Our results suggest that people do indeed use social influence to inform their preferences for animal morphology (at least in the context of AI-generated hybrid animals) and that social influence can lead to the formation of divergent local preferences. In a world where "charismatic megafauna" — animals that play into those more conventional evolutionary and developmental preferences (like the Giant Panda) — absorb much of the public attention and funds for conservation [3, 11, 18, 45], this suggests that social influence may be a powerful mechanism to invigorate attention towards the conservation of animals that do not have such morphological features, like the Chinese Sturgeon [67] or the blobfish [42].

Our work has several important limitations. First, the small number of worlds (4 per condition = 16) makes conducting statistical inference at the world level challenging. Future work could provide more precise effect sizes with a larger number of worlds (e.g. 50 worlds per condition, as recommended by Abeliuk et al. [1]). One particular way of achieving this for web-based experiments that rely on virality for recruitment (as this and the original Music Lab were) is to start with a small, fixed number of worlds per condition, and dynamically branch new worlds (with the same fixed seed set) once a max cap of participants in existing worlds has been reached. This would also naturally allow for cohort blocking to account for variation in outcomes over the course of the experiment's run. Another methodological limitation was the fact that we did not collect demographics to preserve both the privacy and fun of the experience. This lack of user demographics and covariates prevented us from assessing the representativeness of our users, as well as any heterogeneous effects.

We believe that the approach introduced by the Music Lab — randomization at the world level, each with its own independently evolving local ecology — offers platform designers a rigorous way to assess how design interventions not only affect individual behavior as in standard A/B testing, but also complex, collective behavior. We hope this work demonstrates that this experimental paradigm generalizes to outcomes beyond inequality and unpredictability, and in turn can promote interventions to increase the diversity of information ecosystems.

6 ACKNOWLEDGEMENTS

We thank Océane Boulais, Aurélien Miralles, David Rand, Dean Eckles, Matt Salganik, Adam Bear, Anna Chung, Neil Gaikwad, Morgan Frank, Esteban Moro, Skylar Gordon, Josh Hirschfeld-Kroen, Micah Epstein, Jack Muller, Dima Smirnov and Adam Bear for helpful feedback and comments.

REFERENCES

- [1] Andrés Abeliuk, Gerardo Berbeglia, Pascal Van Hentenryck, Tad Hogg, and Kristina Lerman. 2017. Taming the unpredictability of cultural markets with social influence. In *Proceedings of the 26th International Conf. on World Wide Web*. 745–754.
- [2] Marzia Antenore, Alessandro Panconesi, and Erisa Terolli. 2018. Songs of a Future Past—An Experimental Study of Online Persuaders. In *Twelfth International AAAI Conf. on Web and Social Media*.
- [3] Arnold, Carrie. 2014. Which Endangered Species Would You Save? <http://nautil.us/issue/19/illusions/which-endangered-species-would-you-save>. Online; accessed 29 December 2020.
- [4] RB Bendel, SS Higgins, JE Teberg, and DA Pyke. 1989. Comparison of skewness coefficient, coefficient of variation, and Gini coefficient as inequality measures within populations. *Oecologia* 78, 3 (1989), 394–400.
- [5] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. 2012. A 61-million-person experiment in social influence and political mobilization. *Nature* 489, 7415 (2012), 295–298.
- [6] William J Brady, MJ Crockett, and Jay J Van Bavel. 2020. The MAD model of moral contagion: The role of motivation, attention, and design in the spread of moralized content online. *Perspectives on Psychological Science* 15, 4 (2020), 978–1010.
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2018. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096* (2018).
- [8] Yi-Fen Chen. 2008. Herd behavior in purchasing books online. *Computers in Human Behavior* 24, 5 (2008), 1977–1992.
- [9] Nicholas A Christakis and James H Fowler. 2007. The spread of obesity in a large social network over 32 years. *New England journal of medicine* 357, 4 (2007), 370–379.
- [10] Robert B Cialdini and Melanie R Trost. 1998. Social influence: Social norms, conformity and compliance. (1998).
- [11] Agathe Colléony, Susan Clayton, Denis Couvet, Michel Saint Jalme, and Anne-Caroline Prévot. 2017. Human preferences for species conservation: Animal charisma trumps endangered status. *Biological Conservation* 206 (2017), 263–269.
- [12] Rebecca J Compton. 2003. The interface between emotion and attention: A review of evidence from psychology and neuroscience. *Behavioral and cognitive neuroscience reviews* 2, 2 (2003), 115–129.
- [13] Tao Dong, Mark S Ackerman, and Mark W Newman. 2012. Social overlays: augmenting existing UIs with social cues. In *Proceedings of the ACM 2012 Conf. on Computer Supported Cooperative Work Companion*. 79–82.

- [14] Bardia Doosti, David J Crandall, and Norman Makoto Su. 2017. A deep study into the history of web design. In *Proceedings of the 2017 ACM on Web Science Conf.* 329–338.
- [15] Jason Dydyński and Nelly Mäekivi. 2018. Multisensory perception of cuteness in mascots and zoo animals. *International Journal of Marketing Semiotics* 6, 1 (2018), 2–25.
- [16] Ziv Epstein, Océane Boulais, Skylar Gordon, and Matt Groh. 2020. Interpolating gans to scaffold autotelic creativity. *Joint Proceedings of the ICCV 2020 Workshops* (2020).
- [17] Ziv Epstein, Gordon Pennycook, and David Rand. 2020. Will the crowd game the algorithm? Using layperson judgments to combat misinformation on social media by downranking distrusted sources. In *Proceedings of the 2020 CHI Conf. on Human Factors in Computing Systems*. 1–11.
- [18] Mark J Estren. 2018. The Ethics of Preservation: Where Psychology and Conservation Collide. In *The Palgrave Handbook of Practical Animal Ethics*. Springer, 493–509.
- [19] Leah Findlater, Joan Zhang, Jon E Froehlich, and Karyn Moffatt. 2017. Differences in crowdsourced vs. lab-based mobile and desktop input performance data. In *Proceedings of the 2017 CHI Conf. on Human Factors in Computing Systems*. 6813–6824.
- [20] Gary Genosko. 2005. Natures and cultures of cuteness. (2005).
- [21] Nabeel Gillani, Ann Yuan, Martin Saveski, Soroush Vosoughi, and Deb Roy. 2018. Me, my echo chamber, and I: introspection on social media polarization. In *Proceedings of the 2018 World Wide Web Conf.* 823–831.
- [22] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. 2019. Ganalyze: Toward visual definitions of cognitive image properties. In *Proceedings of the IEEE/CVF International Conf. on Computer Vision*. 5744–5753.
- [23] Sofia Gol, Ramona N Pena, Max F Rothschild, Marc Tor, and Joan Estany. 2018. A polymorphism in the fatty acid desaturase-2 gene is associated with the arachidonic acid metabolism in pigs. *Scientific reports* 8, 1 (2018), 1–9.
- [24] Saul Greenberg and Harold Thimbleby. 1991. The weak science of human-computer interaction. (1991).
- [25] C Greiffenhagen and S Reeves. 2013. Is replication important for HCI? IN: Wilson, ML... et al (eds.). ReplicHI 2013. In *Proceedings of the CHI2013 Workshop on the Replication of HCI Research*, Vol. 976. CEUR WS, 8–13.
- [26] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI Conf. on human factors in computing systems*. 203–212.
- [27] Tad Hogg and Kristina Lerman. 2014. Disentangling the effects of social signals. *arXiv preprint arXiv:1410.6744* (2014).
- [28] David Holtz, Benjamin Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. 2020. The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify. *Available at SSRN* (2020).
- [29] Kasper Hornbæk, Søren S Sander, Javier Andrés Bargas-Avila, and Jakob Grue Simonsen. 2014. Is once enough? On the extent and content of replications in human-computer interaction. In *Proceedings of the SIGCHI Conf. on Human Factors in Computing Systems*. 3523–3532.
- [30] Jinghui Hou. 2017. Can interface cues nudge modeling of food consumption? Experiments on a food-ordering website. *Journal of Computer-Mediated Communication* 22, 4 (2017), 196–214.
- [31] Kevin Hu, Snehal Kumar Neil's Gaikwad, Madelon Hulsebos, Michiel A Bakker, Emanuel Zraggen, César Hidalgo, Tim Kraska, Guoliang Li, Arvind Satyanarayan, and Çağatay Demiralp. 2019. Viznet: Towards a large-scale visualization learning and benchmarking repository. In *Proceedings of the 2019 CHI Conf. on Human Factors in Computing Systems*. 1–12.
- [32] Yuheng Hu, Lydia Manikonda, and Subbarao Kambhampati. 2014. What we instagram: A first analysis of instagram photo content and user types. In *Proceedings of the International AAAI Conf. on Web and Social Media*, Vol. 8.
- [33] Jessica Hullman, Eytan Adar, and Priti Shah. 2011. The impact of social information on visual judgments. In *Proceedings of the SIGCHI Conf. on human factors in computing systems*. 1461–1470.
- [34] Guido W Imbens and Donald B Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- [35] Joshua Introne, Karen Levy, Sean Munson, Sean Goggins, Rick Wash, and Cecilia Aragon. 2012. Design, influence, and social technologies: techniques, impacts, and ethics. In *Proceedings of the ACM 2012 Conf. on Computer Supported Cooperative Work Companion*. 9–10.
- [36] René F Kizilcec, Eytan Bakshy, Dean Eckles, and Moira Burke. 2018. Social influence and reciprocity in online gift giving. In *Proceedings of the 2018 CHI Conf. on Human Factors in Computing Systems*. 1–11.
- [37] Scott Klemmer, Mark W Newman, and Raecine Sapien. 2000. The designer's outpost: a task-centered tangible interface for web site information design. In *CHI'00 extended abstracts on Human factors in computing systems*. 333–334.
- [38] PM Krafft, Erez Shmueli, Thomas L Griffiths, Joshua B Tenenbaum, et al. 2021. Bayesian collective learning emerges from heuristic social learning. *Cognition* 212 (2021), 104469.
- [39] Coco Krumme, Manuel Cebrian, Galen Pickard, and Sandy Pentland. 2012. Quantifying social influence in an online cultural market. *PloS one* 7, 5 (2012), e33785.
- [40] Kristina Lerman and Tad Hogg. 2014. Leveraging position bias to improve peer recommendation. *PloS one* 9, 6 (2014), e98914.

- [41] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Governance in social media: A case study of the Wikipedia promotion process. *arXiv preprint arXiv:1004.3547* (2010).
- [42] Lidz, Franz. 2014. Behold the Blobfish. <https://www.smithsonianmag.com/science-nature/behold-the-blobfish-180956967/>. Online; accessed 29 December 2020.
- [43] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the national academy of sciences* 108, 22 (2011), 9020–9025.
- [44] Alexandra Masciantonio, David Bourguignon, Pierre Bouchat, Manon Balty, and Bernard Rimé. 2021. Don't put all social network sites in one basket: Facebook, Instagram, Twitter, TikTok, and their relations with well-being during the COVID-19 pandemic. *PloS one* 16, 3 (2021), e0248384.
- [45] Aurélien Miralles, Michel Raymond, and Guillaume Lecointre. 2019. Empathy and compassion toward other species decrease with evolutionary divergence time. *Scientific reports* 9, 1 (2019), 1–8.
- [46] Lev Muchnik, Sinan Aral, and Sean J Taylor. 2013. Social influence bias: A randomized experiment. *Science* 341, 6146 (2013), 647–651.
- [47] Brian A Nosek, Emorie D Beck, Lorne Campbell, Jessica K Flake, Tom E Hardwicke, David T Mellor, Anna E van't Veer, and Simine Vazire. 2019. Preregistration is hard, and worthwhile. *Trends in cognitive sciences* 23, 10 (2019), 815–818.
- [48] Roger D Peng. 2011. Reproducible research in computational science. *Science* 334, 6060 (2011), 1226–1227.
- [49] Gordon Pennycook, Ziv Epstein, Mohsen Mosleh, Antonio A Arechar, Dean Eckles, and David G Rand. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature* (2021), 1–6.
- [50] Niccolo Pescetelli, Alexis Rutherford, and Iyad Rahwan. 2020. Diversity promotes collective intelligence in large groups but harms small ones. (2020).
- [51] Matthew Pittman and Brandon Reich. 2016. Social media and loneliness: Why an Instagram picture may be worth more than a thousand Twitter words. *Computers in Human Behavior* 62 (2016), 155–167.
- [52] Elissa M Redmiles, Ziyun Zhu, Sean Kross, Dhruv Kuchhal, Tudor Dumitras, and Michelle L Mazurek. 2018. Asking for a friend: evaluating response biases in security user studies. In *Proceedings of the 2018 ACM SIGSAC Conf. on Computer and Communications Security*. 1238–1255.
- [53] Daniel M Romero, Katharina Reinecke, and Lionel P Robert Jr. 2017. The influence of early respondents: information cascade effects in online event scheduling. In *Proceedings of the Tenth ACM International Conf. on Web Search and Data Mining*. 101–110.
- [54] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *science* 311, 5762 (2006), 854–856.
- [55] Matthew J Salganik 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, 4 (2008), 338–355.
- [56] Amit Sharma and Dan Cosley. 2016. Distinguishing between personal preferences and social influence in online activity feeds. In *Proceedings of the 19th ACM Conf. on Computer-Supported Cooperative Work*. 1091–1103.
- [57] Benjamin Shulman, Amit Sharma, and Dan Cosley. 2016. Predictability of popularity: Gaps between prediction and understanding. In *Proceedings of the International AAAI Conf. on Web and Social Media*, Vol. 10.
- [58] Cass R Sunstein. 1999. The law of group polarization. *University of Chicago Law School, John M. Olin Law & Economics Working Paper* 91 (1999).
- [59] Cass R Sunstein. 2001. *Republic.com*. Princeton university press.
- [60] Jozsef A Toth. 1994. The effects of interactive graphics and text on social influence in computer-mediated small groups. In *Proceedings of the 1994 ACM Conf. on Computer supported cooperative work*. 299–310.
- [61] Patrik Vuilleumier. 2005. How brains beware: neural mechanisms of emotional attention. *Trends in cognitive sciences* 9, 12 (2005), 585–594.
- [62] Tim Weninger, Thomas James Johnston, and Maria Glenski. 2015. Random voting effects in social-digital spaces: A case study of reddit post submissions. In *Proceedings of the 26th ACM Conf. on hypertext & social media*. 293–297.
- [63] Alan Wexelblat and Pattie Maes. 1999. Footprints: history-rich tools for information foraging. In *Proceedings of the SIGCHI Conf. on Human Factors in Computing Systems*. 270–277.
- [64] Senuri Wijenayake, Niels Van Berkel, Vassilis Kostakos, and Jorge Goncalves. 2020. Quantifying the effect of social presence on online social conformity. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–22.
- [65] Max L Wilson, Ed H Chi, Stuart Reeves, and David Coyle. 2014. RepliCHI: the workshop II. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. 33–36.
- [66] Sarita Yardi and Danah Boyd. 2010. Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of science, technology & society* 30, 5 (2010), 316–327.
- [67] Xinfa Zhou, Lu Chen, Jing Yang, and Haiqing Wu. 2020. Chinese sturgeon needs urgent rescue. *Science* 370, 6521 (2020), 1175–1175.

Received January 2021; revised April 2021; accepted July 2021