



Moonshots, investment booms, and selection bias in the transmission of cultural traits

David Hirshleifer^{a,1}  and Joshua B. Plotkin^{b,c} 

^aMerage School of Business, University of California, Irvine, CA 92617; ^bDepartment of Biology, University of Pennsylvania, Philadelphia PA 19103; and ^cDepartment of Mathematics, University of Pennsylvania, Philadelphia PA 19103

Edited by Andrew W. Lo, Massachusetts Institute of Technology, Cambridge, MA, and accepted by Editorial Board Member Simon A. Levin February 10, 2021 (received for review August 14, 2020)

Biased information about the payoffs received by others can drive innovation, risk taking, and investment booms. We study this cultural phenomenon using a model based on two premises. The first is a tendency for large successes, and the actions that lead to them, to be more salient to onlookers than small successes or failures. The second premise is selection neglect—the failure of observers to adjust for biased observation. In our model, each firm in sequence chooses to adopt or to reject a project that has two possible payoffs, one positive and one negative. The probability of success is higher in the high state of the world than in the low state. Each firm observes the payoffs received by past adopters before making its decision, but there is a chance that an adopter’s outcome will be censored, especially if the payoff was negative. Failure to account for biased censorship causes firms to become overly optimistic, leading to irrational booms in adoption. Booms may eventually collapse, or may last forever. We describe these effects as a form of cultural evolution, with adoption or rejection viewed as traits transmitted between firms. Evolution here is driven not only by differential copying of successful traits, but also by cognitive reasoning about which traits are more likely to succeed—quantified using the Price Equation to decompose the effects of mutation pressure and evolutionary selection. This account provides an explanation for investment booms, merger and initial public offering waves, and waves of technological innovation.

cultural evolution | evolutionary finance | social finance | Price Equation | mutation pressure

We study how biases in social transmission of information about the actions and payoffs of others induce innovation and risk taking by firms. Our model is based on two premises. The first is the tendency for large successes, and the actions that led to them, to be more visible and salient to others than failures. For example, if 999 out of 1,000 small start-ups fail completely, and one grows to become as large and successful as Google, each of the failures, being small, is seldom noticed and remembered, whereas discussion of the huge success becomes ubiquitous. Successes tend to be more visible because they are associated with extensive continuing economic transactions, which garner attention, whereas failures tend to vanish.

The second premise is the psychological phenomenon of *selection neglect*, the failure of observers to adjust for bias in the process that generates the data they observe.* We apply these two premises to business initiatives, such as creating a start-up firm, undertaking a corporate “moonshot” or “sure bet” investment project, or making a large acquisition. We show that this evolutionary process can result in boom and bust dynamics, consistent with merger and initial public offering (IPO) waves, and sudden waves of innovative technological activity.

In our setting, firms in sequence decide whether or not to undertake (adopt or reject) a risky project, based on observation of the actions and payoffs of previous firms. If the project is adopted, it yields one of two possible payoffs—a positive success payoff or a negative failure payoff. We assume *biased censorship*—a higher probability of observing successful projects and

their outcomes than failed ones. Owing to selection neglect, a manager will have a biased assessment of the prospects of the project. We analyze how these factors shape boom/bust dynamics and long-term survival of different behaviors: adoption versus rejection of projects with different characteristics. We consider these dynamics in the context of cultural evolution of financial traits in a population of firms.

The probability of project success is higher in the high state of the world than in the low state, such that the project is profitable in expectation only in the high state. Each firm observes payoff outcomes of predecessors, but there is some probability that any given adopter’s payoff is censored. If observers were rational, then in the long run, firms would adopt only if the state of the world were high. However, owing to selection neglect, later firms tend to become overly optimistic about the state of the world, so there can be a positive probability that all firms adopt in the low state. In this state, biased censorship can also exacerbate boom/bust dynamics, wherein a long string of adopts occurs before ultimate collapse.

We view adoption or rejection of the risky project as a cultural trait transmitted between firms. We employ the Price Equation to decompose this trait’s evolution into a component due to natural selection and a component due to mutation. Surprisingly, despite the central role of selection bias in the evolution of project choice in the model, the predominant source of cultural change in our context is not natural selection, but, rather, mutation pressure. The importance of mutation during transmission

Significance

What drives innovation, risk taking, and investment booms? We study these phenomena using a model of decision making by firms that make biased observations of prior returns. We assume that firms are more likely to observe large successes than small successes or failures. As a result, firms tend to become overly optimistic, leading to irrational booms in investment. Booms may eventually collapse, or they may last forever. We describe the cultural evolutionary sources of these effects. Evolution of investment behavior is driven not only by copying successful strategies, but also by cognitive reasoning about which behaviors are more likely to succeed. This account provides an explanation for investment booms, merger and IPO waves, and waves of technological innovation.

Author contributions: D.H. and J.B.P. designed research, performed research, contributed new reagents/analytic tools, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. A.W.L. is a Guest Editor invited by the Editorial Board.

Published under the [PNAS license](#).

¹To whom correspondence may be addressed. Email: david.h@uci.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2015571118/-/DCSupplemental>.

Published June 25, 2021.

*See, e.g., refs. 1 and 2 and, in the financial context, ref. 3.

differs sharply from cultural evolutionary models with biased imitation, in which there is only natural selection. This feature of our analysis highlights the role of cognitive reasoning in the cultural evolution of risk-taking behaviors.

Our approach differs from a literature on entrepreneurship that assumes that entrepreneurs are exogenously overoptimistic. Beyond the conclusion of overadoption of projects, we derive implications about boom/bust dynamics and a variety of distinctive comparative statics predictions. More broadly, our analysis suggests that it is important to account for the psychology of attention and social interaction rather than solely of overconfidence.

We are not the first to examine sequential social learning. For example, in the models of refs. 4 and 5, rational agents update their beliefs by observing the actions of predecessors. Owing to so-called information cascades, decisions are often mistaken, and social outcomes are fragile. The current paper differs from prior work by examining social learning from past *payoffs* as well as actions, the possibility of biased censorship of past observations, and imperfect rationality in the form of neglect of this selection bias. These features have distinctive consequences for boom–bust dynamics in investment choice.

We are also not the first to examine selection bias and learning by firms. Denrell (6) also examines a setting in which observers neglect selection bias in the information they observe about firms. As in our model, failure is less likely to be observed than success. Denrell’s focus is on how this biases learning about the traits that are characteristic of the upper tail of successful firms, which he argues will disproportionately consist of variance-increasing strategies. So Denrell concludes that selection bias will cause the spread of risky and unreliable management practices.

Our study differs in several ways. First, we focus on beliefs about the benefits of project adoption rather than about general managerial practices. Second, we allow for sequential choices and selection bias. In other words, we model explicitly how neglect of selection bias affects an arbitrary observer’s behavior, not just beliefs; and how the observer, in turn, becomes the target of observation for the next agent, and so forth. This allows us to study the implications of selection bias for the dynamics of investment booms and collapses. Third, our modeling allows us to analyze the evolutionary process by which behavioral traits are transmitted across agents in terms of selection and mutation pressure. Fourth, Denrell’s focus is on how differences in variance biases choices, whereas we analyze payoff asymmetry and show that greater moonshotness promotes project overadoption.

Han et al. (7) examine a setting in which stock-market investors randomly meet to discuss their strategies, and the probability that an investor reports the investor’s return performance is increasing in return. An investor has an exogenous probability of copying another investor’s strategy that is increasing with reported return, if the investor receives a report. As a result, high-variance investing strategies spread through the population. Our focus here is on project choices by firms that update beliefs in a quasi-Bayesian fashion, based on sequential observation of a history of past payoffs by other firms. In contrast, in ref. 7, investors are randomly drawn to meet, and message receivers have an exogenous switching probability based upon the single observation during that meeting. Our study also differs in deriving boom/bust dynamics, comparative statics about moonshotness, and decomposition of outcomes into evolutionary components.

The Basic Setting

Consider a setting in which agents in sequence choose between two actions, adopt or reject an investment project. Each agent learns about the state of the world from the payoff experiences

of predecessors (8–10). These assumptions fit many practical applications in which firms’ investment choices are distributed over time, so that later firms can observe the decisions and performances of predecessors.

Reject always generates a payoff of zero, whereas adopting generates one of two possible payoffs, which are distributed independently across firms, conditional upon state. Let the possible actions for the n^{th} agent in the sequence, denoted I_n , be $a_n = A, R$ (Adopt or Reject). When an agent rejects, the payoff is zero. The binary state of the world is $\theta = H, L$, with prior probability q of state H . In each state, there are two possible net payoffs to adoption, $v = V > 0$ or $v = -1$. Project success has probability $p > 1/2$ in state H and $1 - p$ in state L . Therefore, high payoff is a symmetric binary signal about state, with payoff probabilities shown in Table 1.

We assume that when indifferent, an agent rejects, and we impose parameter constraints such that adopting is, in expectation, strictly profitable in state H and unprofitable in state L . We let the expected profit from adopting conditional on state be Π^θ ,

$$\begin{aligned} \Pi^H &\equiv E[v|\theta = H] = pV - (1 - p) > 0 \\ \Pi^L &\equiv E[v|\theta = L] = (1 - p)V - p < 0, \end{aligned} \quad [1]$$

which together imply that $\frac{1-p}{p} < V < \frac{p}{1-p}$.

Agents maximize expected profits conditional on their information, which derives from observing past actions and payoffs. To avoid the trivial outcome that all agents always reject, we assume that the prior expected value of action A is positive,

$$\begin{aligned} E[v] &= q[pV - (1 - p)] + (1 - q)[(1 - p)V - p] \\ &= q(2p - 1)(V + 1) + (1 - p)V - p > 0. \end{aligned} \quad [2]$$

This implies that the first agent, I_1 , adopts.

Full Observability of Past Payoffs

If at any point an agent I_i ’s belief about the state becomes sufficiently pessimistic, I_i rejects the investment project, and I_i ’s payoff of zero contains no information about state. Since I_{i+1} has no additional information, I_{i+1} also rejects, as do all later agents. So in the stochastic process of actions and payoffs, the action R is an absorbing outcome. With full observability of past payoffs, there is a random walk on beliefs, as measured by the log-likelihood ratio (LLR) between the two states, with a lower absorbing barrier that produces long-run project rejection (see, e.g., ref. 10).

Specifically, conditional on state, so long as agents adopt, the belief follows a random walk with drift. Let d_i be the difference between the number of high-payoff realizations (V) and the number of low realizations (-1) observed by agent I_i (i.e., up to and including the payoff of I_{i-1}), where, by convention, $d_1 = 0$. Since v realizations are conditionally independent given state $\theta = H$ or L , a standard result for Bayesian updating with binary state and signals is that updated beliefs depend only on the difference between the counts of the two types of outcomes,

Table 1. Payoff probabilities

State (θ)	Payoff outcome (v)	
	V	-1
H	p	$1 - p$
L	$1 - p$	p

$$\lambda_i \equiv \log \left(\frac{P(H|d_i)}{P(L|d_i)} \right) = d_i \log \left(\frac{p}{1-p} \right) + \log \left(\frac{q}{1-q} \right). \quad [3]$$

So the relevant payoff history from past adopts is fully summarized by d_i . (Rejects generate a deterministic payoff, and therefore are uninformative.)

The belief state in our model, as measured by the LLR λ_i , therefore behaves according to the standard Gambler's Ruin additive random walk. The LLR changes at each step by the increment $\log \left(\frac{P(v|H)}{P(v|L)} \right)$, which is positive when $v = V$ and negative when $v = -1$, and which has the same absolute value in both cases. In state $\theta = H$, the probabilities of these increments are p for $v = V$ and $1 - p$ for $v = -1$. In state L , the probabilities are $1 - p$ for $v = V$ and p for $v = -1$.

The lower absorbing barrier in this Gambler's Ruin is the cutoff on the LLR above which adoption has a positive expected payoff. Intuitively, under the prior belief, an agent adopts, and so bad news ($d_i < 0$) is required for an agent to reject. We can calculate the continuous cutoff value d^* above which an agent I_i would be just willing to adopt by imposing the positive profit condition

$$E[v|d_i] = P(H|d_i)[pV - (1-p)] + (1 - P(H|d_i)) \times [(1-p)V - p] = 0. \quad [4]$$

By Eqs. 1 and 3, the indifference cutoff on d is

$$d^* \equiv \frac{\log \left(\frac{p-(1-p)V}{pV-(1-p)} \right) - \log \left(\frac{q}{1-q} \right)}{\log \left(\frac{p}{1-p} \right)}. \quad [5]$$

Since an indifferent agent rejects, the first reject occurs when d_i reaches or crosses d^* , $d_i \leq \lfloor d^* \rfloor$. For example, if $d^* = -1.1$, the first reject occurs when $d_i = -2$. The numerator of Eq. 5 is negative (see *SI Appendix, section 2* for details), which implies that $d^* < 0$. In summary, the count of past high minus low payoffs follows a random walk with a lower absorbing barrier $\lfloor d^* \rfloor < 0$.

The model therefore endogenously generates a Gambler's Ruin random walk in beliefs. In general, for a random walk with a lower absorbing barrier at zero and an up move probability $r \leq 1/2$, the probability of ultimate ruin starting from position $d_0 > 0$ is one, and with $r > 1/2$, is

$$P(\text{Ruin}) = \left(\frac{1-r}{r} \right)^{d_0}, \quad [6]$$

which is decreasing in d_0 (11). Intuitively, unless the probability of an up move exceeds that of a down move ($r > 1/2$), the gambler will eventually wander below the barrier and go broke. In contrast, if $r > 1/2$, the upward drift in the walk gives the gambler a chance of exceeding the lower barrier forever.

The standard Gambler's Ruin random walk is an equivalent transformation of our model in which the absorbing barrier $\lfloor d^* \rfloor < 0$ is increased to zero, and the starting position is accordingly increased to $d_0 = -\lfloor d^* \rfloor$. Here, the probability of an up move corresponds to $p > 1/2$ or $1 - p$ in the H or L state, respectively. We therefore have:

Proposition 1. *With observation of all past actions and payoffs:*

1. In state L , with probability one, the project will eventually be abandoned by an agent (and by all subsequent agents).
2. In state H , there is a strictly positive probability that all agents will adopt the project.

The exact formula for the probability that all agents adopt in state H is given in *SI Appendix, section 4*. The proof is also given in *SI Appendix*.

The results are intuitive. If an arbitrarily large number of agents were to adopt, beliefs would converge to virtual certainty in state L , causing rejection. As for state H , either early bad news makes adopt seem so bad that the reject barrier is crossed (in which case, no further information is generated, so all subsequent agents reject), or else there is an infinite number of adopts, so that in the long run, agents become highly confident of the H state.

From a welfare perspective, a key property of the model is underexploration. When an agent chooses A , the payoff outcomes provide information about the state that is useful for later agents. Each agent ignores this external benefit, so that, from the viewpoint of social welfare, agents choose A too rarely. Specifically, even if the expected private gain to adopting A is slightly negative, from a social perspective, A should be adopted. A similar form of underexploration arises in models of social two-armed bandit problems (9).

Biased Censorship and Rational Updating

Suppose now that high-payoff outcomes are more likely to be observed than low-payoff outcomes. Specifically, we assume that a payoff of $v = V$ is observed with probability $\delta \leq 1$, whereas a payoff of $v = -1$ is observed only with probability $\pi \leq \delta$. So the upside censorship probability is $1 - \delta \geq 0$, and the downside censorship probability is $1 - \pi \geq 0$. We assume strict inequalities, where the case of $\delta = \pi = 1$ is a benchmark for comparison.

We call greater censorship of downside outcomes *upside salience*, defined algebraically as the ratio $\beta \equiv \delta/\pi$. Intuitively, projects that succeed are associated with a high scale of continuing economic transactions, which garners attention, whereas projects that fail tend to vanish. Consistent with the example in the introduction of 1,000 start-ups, one of which has a big success, there is extensive evidence from psychology that rare events are salient and tend to be overweighted (12) and, similarly, that people disproportionately recall realizations with extreme payoffs (13). This motivates our assumption that $\beta > 1$, especially for low q and large V .

We can think of censorship of observation of a firm's project outcome as meaning that the outcome is not reported conspicuously in the media. So we assume that censorship of an agent is universal: The agent's payoff is either visible to all successors or to none of them.[†]

A related motivation for the assumption $\beta > 1$ is that managers and firms have an incentive to disclose good news immediately, and to defer the disclosure of bad news. If investors are not fully attentive to these strategic incentives, then the withholding of bad news about project payoffs will cause investors to process high-payoff outcomes more often than low-payoff outcomes (14).

Let $v_j = V$ or -1 be the realized payoff outcome of agent I_j if j adopts. Since I_1 adopts, observation by rational I_2 of I_1 's payoff updates I_2 's belief about state H to beliefs that are independent of π and δ (see *SI Appendix, section 6* for details). This is because, conditional on observation of a payoff, the greater censorship of low payoffs makes no difference for inferences. The only information contained in the sheer fact of observation is that the payoff is less likely to be -1 . But conditional on a payoff that is directly observed, this information is redundant.

In contrast, the absence of an observation of an earlier agent contains useful information for agent I_i , as in the story in which detective Sherlock Holmes draws a key inference from "the dog that did not bark." For example, if I_2 does not observe I_1 , then the rational inference is tilted toward state $\theta = L$, since the

[†]As a result, agents cannot glean any private information from the action choices of predecessors. This point is not relevant for our main analysis, since, in our setting with imperfect rationality, agents do not even realize that there is any such indirect information to be extracted.

occurrence of no observation tends to come from the payoff -1 (since $\pi < \delta$). However, there is evidence that, in reality, financial decision makers and markets sometimes underreact to the information implicit in event nonoccurrence (15).

Neglect of Payoff-Biased Censorship

What if agents neglect the fact that there is biased censorship of the payoffs of past investment projects? We first consider how neglect of censorship affects short-term and long-term dynamics of project adoption. In *Neglect of Payoff-Biased Censorship*, we examine basic comparative statics on long-run adoption. In *Identifying Sources of Cultural Evolution using the Price Equation*, we examine the sources of cultural evolution toward adopt or reject using the Price Equation.

Neglect of Biased Censorship and Adoption Dynamics. When agents are unaware of censorship, if agent I_j observes $j' < j - 1$ past actions and payoffs, agent I_j mistakenly believes that he is the $(j' + 1)^{\text{th}}$ agent in the decision queue. We expect that disproportionate censorship of low payoffs will make agents overly optimistic about project adoption.

We will show that the beliefs of an appropriate sequence of agents follows a random walk with drift, with an absorbing barrier that induces rejection. Consider the *subsequence* of the belief sequence that removes all agents whose payoffs are censored. This *uncensored subsequence* contains the beliefs of all agents whose payoffs matter for later agents. We will see that the beliefs (LLRs) of agents in this subsequence follow a random walk.

Let a B (for biased) superscript denote an imperfectly rational expectation. A biased agent mistakenly fails to condition on observation versus nonobservation of payoffs when forming expectations. Let O_{21} denote the event that I_2 observes I_1 . Then, for I_2 (see *SI Appendix, section 3* for details),

$$\begin{aligned}
 P^B(\theta = H | O_{21}, v_1 = V) &= \frac{P^B(\theta = H, V_1 = V)}{P^B(v_1 = V)} \\
 &= \frac{1}{1 + \frac{1-q}{q} \frac{1-p}{p}} \\
 P^B(\theta = H | O_{21}, v_1 = -1) &= \frac{P^B(\theta = H, v_1 = -1)}{P^B(v_1 = -1)} \\
 &= \frac{1}{1 + \frac{1-q}{q} \frac{p}{1-p}}.
 \end{aligned}$$

We see that the beliefs conditional on seeing V or -1 are the same as for a rational agent, independent of noncensorship probabilities δ and π . So neglecting the information implicit in the fact that a payoff was observed does not affect conditional beliefs.

In contrast, neglect of the information implicit in the absence of an observation crucially affects action dynamics and the evolution of beliefs, because unlike a rational agent, an inattentive agent does not draw inferences from censored observations.

When agents naively neglect nonobservation, they update based on the observed difference between the numbers of high and low payoffs, just as in the case with no censorship. Specifically, let j be the index for the subsequence, i.e., a count in order of the uncensored agents. Let d_j be the difference in the number of V and -1 outcomes through the $(j - 1)^{\text{th}}$ agent in this subsequence. Then, Eq. 3 holds with i replaced with j .

However, neglect of censorship modifies the probabilities of an up or down move in beliefs conditional on state in the LLR walk. In the uncensored subsequence, the probabilities of an up move and down move in the next step are constant over time. So in the random walk for the uncensored subsequence, conditional on state $\theta = H$, the probabilities of an up move, p^{H*} , and of a down move, $1 - p^{H*}$, are (see *SI Appendix, section 3* for details)

$$\begin{aligned}
 p^{H*} &= P(v_1 = V | H, O_{21}) = \frac{p}{p + (1-p)/\beta} > p \\
 1 - p^{H*} &= \frac{1-p}{1 + p(\beta-1)} < 1-p,
 \end{aligned} \tag{7}$$

where the inequalities follow because upside salience $\beta > 1$. The probability of an up move in the subsequence random walk conditional upon state H is increased by upside salience.

Similarly, in the subsequence random walk, conditional on state $\theta = L$, the probabilities of an up move and a down move, respectively, are

$$\begin{aligned}
 1 - p^{L*} &= \frac{1-p}{1-p+p/\beta} > 1-p \\
 p^{L*} &= \frac{p}{p + \beta(1-p)} < p,
 \end{aligned} \tag{8}$$

since $\beta > 1$.

There are two key differences in the evolution of beliefs from the case of no censorship. First, we must characterize beliefs for censored agents as well. These are determined trivially from the beliefs in the uncensored subsequence. Any censored agent's belief is identical to the belief of the next uncensored agent, since the next uncensored agent does not see intervening censored payoffs. Second, the probability of an up move in the LLR walk is higher with neglect of biased censorship. This effect can be arbitrarily large. As β becomes large ($\pi \approx 0$), both p^{H*} , $1 - p^{L*} \rightarrow 1$. So with neglect of biased censorship, beliefs can tend to march upward, even in the L state, where complete information would inevitably lead to rational rejection.

We summarize the evolutionary dynamics of adoption under biased censorship in the following proposition.

Proposition 2. *Under biased censorship, if:*

1. *In state L:*
 - a. (Weak upside salience) *If $\beta \leq p/(1-p)$, then with probability one the project is eventually abandoned.*
 - b. (Strong upside salience) *If $\beta > p/(1-p)$, then the probability that all agents adopt is strictly between zero and one, and it is increasing in δ and decreasing in π . Otherwise, the project is eventually abandoned.*
2. *In state H, the probability that all agents adopt is strictly between zero and one, and it is increasing in δ and decreasing in π . Otherwise, the project it is eventually abandoned.*

In sharp contrast with Proposition 1, owing to neglect of censorship of low payoffs, there is a positive probability that all agents adopt, even when the project has negative expected value (state L). Moreover, the chance that all agents adopt is greater in state H than without censorship.

Let $\hat{\beta} \equiv p/(1-p)$ be the cutoff value on upside salience, such that with positive probability, agents adopt in the long run, even in state L . Clearly, $\hat{\beta}$ is increasing with p , the superiority in the success rate in the high versus low state. When project success is highly sensitive to state, greater relative selection bias on the upside versus the downside is needed to induce long-run adoption in state L .

Proposition 2 shows that the model generates boom dynamics, with strings of adoption, even in state L . Owing to upside salience, such booms occur more often when agents neglect censorship. Such neglect causes irrational extra booms. Booms may be followed by collapse, or they may be sustained permanently. Adoption continues until (if ever) the preponderance of uncensored evidence, in the form of low payoffs, favors reject.

In some realizations, there is no boom (a single adopt followed by all reject). In others, there is a boom period followed by a collapse (several or many adopts and then all reject). In other cases,

there is a persistent boom (all adopt). In the L state, the boom component of the boom/collapse pattern is mistaken relative to full knowledge of state.

Since agents do not know the state, booms and busts do not, in themselves, indicate bad decision making. Without censorship, all decisions are rational, but such mistakes still occur with positive probability. However, when there is biased censorship and selection neglect, there are *irrational* booms and busts. This is illustrated in Fig. 1, which depicts realizations belief LLR over time. Without censorship, after two Adopts, the Reject barrier is reached, and all subsequent agents reject. In contrast, under censorship, there is a temporary boom, until the late sequence of negative payoffs brings about collapse. This must eventually happen in the L state if upside salience is sufficiently weak. Fig. 1, *Lower* makes a similar comparison under a set of realizations in which, under censorship, the boom takes longer to collapse. Indeed, the boom may continue forever, even in state L , provided upside salience is sufficiently strong.

Under selection neglect, the probability of an early string of adopts is increased—there are adoption bubbles. Suppose, for example, that there is relatively little censorship. By Proposition 2, with probability one, in the L state, eventually all agents reject. So censorship causes extra booms that would not occur were agents rational. If censorship is not too severe, these booms later crash. If censorship is more severe, matters are even worse—there are permanent mistaken booms.[‡] The model therefore offers an explanation for real investment booms and busts (16), IPO waves and overoptimism (17–19), and value-reducing merger waves (20, 21). This explanation differs from some past explanations that require payoff externalities (22) or shifts in investor sentiment (23). There are, of course, other possible explanations as well.

Since key parameters of a model, such as β , can never be estimated perfectly, it is important to know whether model implications are robust to parameter variations. The qualitative implications of Propositions 1 and 2 are robust to small variations in β , except for the measure-zero set of critical values described in the propositions.

Comparative Statics of Long-Run Adoption. How does the upside payoff, the censorship probability, and the prior likelihood of success affect the chance of eventual adoption? To address these questions, in either state H or L , we analyze the partial derivatives of the chance of persistent adoption with respect to model parameters.

By Eq. 6, the long-run log rejection probability in the two states (where for state L , we impose the condition that $\beta > p/(1-p)$ so that ruin is not assured) is

$$\log(P(\text{Ever Reject}|\theta)) = \begin{cases} -(\lfloor d^* \rfloor) \log\left(\frac{1-p}{\beta p}\right) & \text{if } \theta = H \\ -(\lfloor d^* \rfloor) \log\left(\frac{p}{\beta(1-p)}\right) & \text{if } \theta = L. \end{cases} \quad [9]$$

Since $\log(p/(1-p))$ and $\log((1-p)/p)$ have opposite signs, it is evident that all of the parameters, except perhaps p and β (which reflects δ and π), have opposite directional effects on the long-run reject probabilities.

We use a version of Eq. 9 that ignores the floor function as a continuous approximation for the chance long-run rejection, after substituting Eq. 5 for d^* . To derive comparative statics, let

[‡]The conclusion that booms last forever will not hold for some applications in which payoffs are interdependent, in contrast with our model assumption. For example, if firms undertake projects that are in competition with each other, then adoption by early firms reduces the expected payoffs to later firms of adopting. In such a setting, booms would eventually end, as firms started to experience lower and lower payoffs.

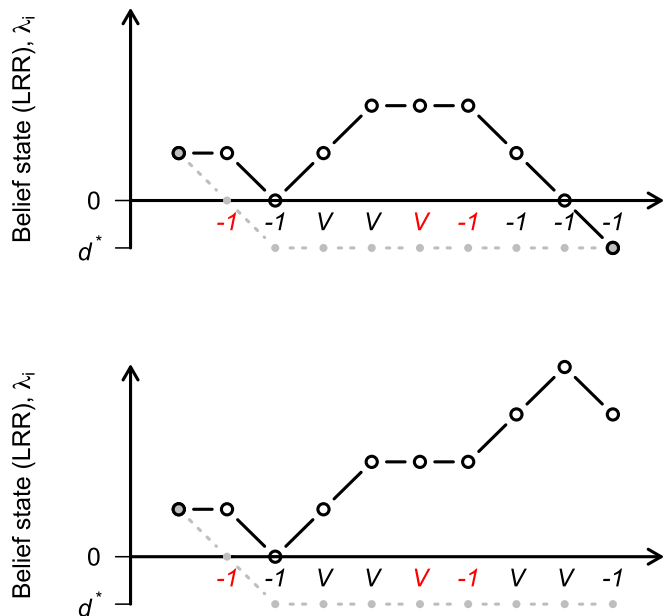


Fig. 1. Booms of project adoption, with and without biased censorship. (*Upper*) The graph compares belief states of successive firms (x axis) in the case of censorship (black) to the case of no censorship (gray), under a set of payoff realizations, in which the boom eventually collapses. (*Lower*) The graph makes a similar comparison under a set of realizations in which, under censorship, the boom takes longer to collapse, or may never collapse. Red payoffs correspond to censored outcomes.

R^H and R^L denote the log probabilities of ever rejecting in state H or L :

$$\begin{aligned} R^H &\equiv \log(P(\text{Ever Reject}|\theta = H)) \\ R^L &\equiv \log(P(\text{Ever Reject}|\theta = L)). \end{aligned} \quad [10]$$

We study the partial derivatives of R^H and R^L with respect to model parameters, p , q , V , and β (or π and δ separately) in regime with a positive chance of all agents adopting even in state L (i.e., $\beta > p/(1-p)$).

The effect of increasing upside salience follows immediately from differentiating Eq. 9 with respect to β . This gives $\frac{\partial R^H}{\partial \beta}, \frac{\partial R^L}{\partial \beta} < 0$. Since δ and π enter the reject probability expressions only through the definition of β , we also immediately have $\frac{\partial R^H}{\partial \delta}, \frac{\partial R^L}{\partial \delta} < 0$ and $\frac{\partial R^H}{\partial \pi}, \frac{\partial R^L}{\partial \pi} > 0$. So as upside salience becomes stronger (owing to either lower upside censorship or stronger downside censorship), the long-run chance of adoption under L increases. Intuitively, upside salience breeds overoptimism, reducing rejection probability.

We can show several other simple conclusions (see *SI Appendix, section 5* for derivations). High upside payoff, V , and high probability of the H state, q , both promote long-run adoption. These facts have opposite implications about the effects of moonshotness, which is, by definition, associated with both high V and low q . For state H , as the upside reward V of the project increases, or the prior likelihood of state H increases, the probability of long-run adoption increases. Likewise, for state L , so long as biased censorship is strong enough that there is no certainty of eventual rejection, the probability of eventual rejection is increasing in chance of failure p , decreasing in upside payoff V , and decreasing in q .

Higher p promotes rejection because payoffs become more accurate indicators of the actual state, L . Agents also understand that payoffs are more accurate, and therefore update more strongly to payoff outcomes. This increases the sizes of up

and down moves in the LLR random walk, reducing the mean number of steps required to reach the absorbing barrier. This effect also promotes rejection in the long run. Some less obvious comparative statics, including those pertaining to interactive effects of parameters obtained from mixed partial derivatives, are derived and discussed in more depth in *SI Appendix, section 5*. These effects provide a rich set of further potentially testable empirical implications.

Proposition 3. *In both states $\theta = H, L$, where for state L , we require parameter values such that $R^L < 1$, we have the following comparative statics: $\frac{\partial R^\theta}{\partial \beta} > 0$, $\frac{\partial R^\theta}{\partial \delta} < 0$ and $\frac{\partial R^\theta}{\partial \pi} < 0$; $\frac{\partial R^\theta}{\partial V} > 0$, $\frac{\partial R^\theta}{\partial q} < 0$. Furthermore, in state L , for parameter values such that $R^L < 1$, we have: $\frac{\partial R^L}{\partial p} > 0$; $\frac{\partial^2 R^L}{\partial \beta \partial V} > 0$ and $\frac{\partial^2 R^L}{\partial \delta \partial V} < 0$ and $\frac{\partial^2 R^L}{\partial \pi \partial V} > 0$; $\frac{\partial^2 R^L}{\partial \beta \partial q} > 0$.*

To understand the cross-partials, recall that rejection occurs when the random walk d'_k (the net count of up versus down steps observed by agent I_k) hits the lower boundary $d^* < 0$. Increasing the upside payoff V decreases d^* , making eventual rejection less likely ($\partial R^L / \partial V < 0$). Increasing π (reducing censorship) increases the chance of a downward step, making rejection more likely ($\partial R^L / \partial \pi > 0$). However, if V is already very large, the chance of ever reaching the lower boundary $d^* \ll 0$ is almost zero, and so the marginal effect of an increase in π is small. So the cross-partial on R^L is positive. A similar intuition applies with opposite sign for the cross-partials with respect to β or δ .

Moonshots, Sure Bets, and Upside Salience. A moonshot is a project that has low probability of success, as reflected in the probability of the H state, q , and a high upside payoff V . This is the opposite of a sure bet project, which has high q and low V . The key psychological premise of our analysis of moonshotness is that upside salience β is greater for moonshots than for sure bets. Both a low ex ante probability of success and a high conditional payoff make success more surprising and newsworthy.

As stated in Proposition 3, long-run adoption is more likely when upside salience β is larger. So the model offers the empirical implication that moonshots will tend to be adopted more heavily than sure bet projects, even after controlling for net expected value.

Moonshotness can only be varied by varying q and V , each of which affects rejection probability. To provide a sharp assessment of the effects of moonshotness on rejection probability via its effect on upside salience β , we control for the direct effects of q and V by considering excess rejection probabilities in the two states.

As before, we assume that $\beta > 1$ and, for the analysis of state L , focus on the case in which $\beta > p / (1 - p)$, so that eventual rejection is not assured. Specifically, for two possible values of upside salience, β and $\bar{\beta}$ (where both satisfy the relevant inequality above), let the *excess log reject probability* relative to benchmark $\bar{\beta}$ in state H or L be defined as

$$e^H(\beta, \bar{\beta}) \equiv -\frac{R^{H,\beta}}{R^{H,\bar{\beta}}} = -\log\left(\frac{1-p}{\beta p}\right) / \log\left(\frac{1-p}{\bar{\beta} p}\right)$$

$$e^L(\beta, \bar{\beta}) \equiv -\frac{R^{L,\beta}}{R^{L,\bar{\beta}}} = -\log\left(\frac{p}{\beta(1-p)}\right) / \log\left(\frac{p}{\bar{\beta}(1-p)}\right), \quad [11]$$

where $R^{\theta,\beta}$ and $R^{\theta,\bar{\beta}}$ are the values of R^θ for two given values of β , and where the algebraic expressions follow from Eq. 9. The excess log reject probability captures the idea that, for any given values of q , V , and p , greater upside salience biases agents away from rejecting. So these expressions are decreasing with upside

salience, β .⁵ Notably, these are exact expressions, and they are independent of q and V .

To understand the effect of moonshotness in more depth, consider the probability of eventual rejection when varying V and q inversely, while holding constant the prior expected payoff, $\bar{v} \equiv E(v)$. As argued above, an increase in moonshotness, captured here by an increase in V , increases upside salience β . Since the expressions in Eq. 11 have no direct dependence upon V and q , varying moonshotness affects these expressions only through its effect on upside salience β .

Specifically, differentiating Eq. 11 with respect to β yields

$$\frac{\partial e^H(\beta, \bar{\beta})}{\partial \beta} = \frac{1}{\beta \log\left(\frac{1-p}{\bar{\beta} p}\right)} < 0$$

$$\frac{\partial e^L(\beta, \bar{\beta})}{\partial \beta} = \frac{1}{\beta \log\left(\frac{p}{\bar{\beta}(1-p)}\right)} < 0. \quad [12]$$

So greater upside salience β decreases the excess rejection probability. By increasing upside salience, an increase in moonshotness decreases excess rejection, meaning it decreases rejection after controlling for the direct effects of V and q (apart from there indirect effect via beta β). Intuitively, moonshotness biases observation toward past successes rather than failures, which promotes adoption. In *SI Appendix, section 6*, we discuss in greater detail the conceptual experiment underlying the comparative statics on excess log reject probabilities.

Salience and Firm Size. The noncensorship parameter π in the model is a proxy for the salience of downside outcomes and δ for the salience of upside outcomes, where we have defined $\beta = \delta / \pi$ as the relative attention to the upside payoff (upside salience). Greater overall attention by media and observers will tend to increase both π and δ , with potentially offsetting effects on β .

However, at least in the limit, higher attention tends to reduce β . For example, other things equal, large firms tend to receive much greater attention than small firms (24). In the limit, if a well-known firm such as Ford receives high attention to both its successes and failures, $\pi = \delta \approx 1$, so $\beta \approx 1$, i.e., there is no upside salience. In other words, for a large firm, failure of a major project can be notable enough to be reported upon in the media. In contrast, small start-ups often fail unnoticed. What is mainly reported in the media are extraordinary successes that start in garages and become tech giants. So, other things equal, β will be higher for small firms than for large firms. The model therefore implies that the bias in favor of adopting risky projects will be especially strong for small start-ups. A further implication is that moonshots that are initiated with great fanfare and heavy investment may not generate as much mythology and overestimation as stories about firms that started in a garage.

If the effects of the model are stronger for small start-ups, then naive observers will strongly overestimate the probability of such start-ups succeeding, resulting in active and impetuous entrepreneurial activity. There is survey evidence that entrepreneurs are highly overoptimistic about their likely success (25). Our model suggests that there will tend to be much less overoptimism about the innovative projects of large firms, and therefore less frequent undertaking of such projects.

Identifying Sources of Cultural Evolution using the Price Equation

Our model is an evolutionary system, in which later agents stochastically inherit the trait of adopt versus reject from earlier

⁵Since the benchmark denominator expression $R^{\theta,\bar{\beta}} < 0$, a negative sign is needed to ensure that higher $e^\theta(\beta, \bar{\beta})$ is indicative of higher reject probability.

agents. If we designate some set of earlier agents as ancestral, and some set of later agents as descendant, then we can glean insight into the drivers of cultural evolution using the Price Equation (26). The Price Equation decomposes evolutionary change into selection and nonselection effects. The nonselection component is often called *mutation pressure*—the degree to which traits shift through the inheritance process instead of fitness-biased replication. The standard Price Equation applies to realizations; we take expectations to employ it in ex ante form.

Most studies on cultural evolution focus on trait change driven by the cultural analogue of natural selection: namely, biased imitation of traits that produce high payoff. Under biased imitation, some traits are more likely to be imitated than other traits. In the large field of evolutionary game theory, for example, biased imitation of high-payoff traits is often the only force considered (27).

However, the Price Equation highlights that evolutionary change in general also derives from mutation pressure—the tendency for traits to be systematically modified in the process of being transmitted from an ancestor to a descendant, irrespective of ancestor fitness. We therefore use the Price Equation to explore the roles of selection versus mutation pressure on the evolution of the financial trait of adopting versus rejecting an investment project. As we will see, the conclusions are somewhat surprising.

The Ex Ante Price Equation. The insights provided by the Price Equation depend on how key categories are defined: ancestors, descendants, inheritance, and traits. We view *inheritance* in our setting as a potential causal dependence of the adopt/reject trait among descendants on the traits of ancestors. Under this interpretation, there is no inheritance from a censored agent, because later agents derive no information from such an agent.

With our definition of inheritance in mind, we now apply a stochastic version of the Price Equation. Let the ancestral generation be denoted a and the descendant generation d . Let q^i be the frequency of type i in the a population, where $i=0$ indicates reject and $i=1$ indicates adopt. Let q'^i be the frequency of type i in the d population. We use overlines to denote expected frequencies, so let $\overline{q^i} \equiv E[q^i]$. In a realization where type $i=0$ has zero frequency in a , fitness is undefined. We therefore use a form of the Price Equation emphasized by ref. 28 that contains ancestral and descendant frequencies instead of fitnesses.

Letting z^i denote the trait value of ancestral agent i , the average trait values in the a and d populations are $\hat{z} \equiv \sum q^i z^i$, $\hat{z}' \equiv \sum q'^i z^i$. In our context, the sum is over the two alleles, adopt and reject. We code these types with index $i=0, 1$, which have ancestral alleles $z^0=0$ and $z^1=1$, where zero indicates reject and one indicates adopt. We study the change in trait value, $\Delta \hat{z} \equiv \hat{z}' - \hat{z}$ and its conditional expectation.

Let $\Delta q^i \equiv q'^i - q^i$ be the frequency change due to natural selection. Let the trait value change be $\Delta z^i \equiv z'^i - z^i$. The Price Equation decomposes average trait change into two terms

$$\widehat{\Delta z} = \underbrace{\sum q'^i (\Delta z^i)}_{\text{mutation pressure}} + \underbrace{\sum (\Delta q^i) z^i}_{\text{selection}}, \quad [13]$$

as given by ref. 28.

Consider, in state θ , a sequence of A or R realizations through agent I_n . We consider this population of n agents to be the ancestral generation, and we analyze the change or expected change in the average trait value to the descendant generation, I_{n+1} , which is a population consisting of a single agent.

Under our definition of inheritance as a potential causal dependence of traits, an ancestral agent who is censored has no descendants. Furthermore, an agent with trait $i=0$ rejects,

resulting in a deterministic payoff of zero, which is uninformative to later agents. Since reject agents have no influence on later beliefs or actions, reject agents have no descendants. Finally, an *uncensored* agent with trait $i=1$ does generate information that influences the information, and potentially the behavior, of the descendant I_{j+1} . So all uncensored agents collectively share the same, single descendant. That is, the descendant I_{j+1} inherits from all ancestral agents.[†] Since i refers to type in the above equations, the terms reflect the aggregated inheritance derived from all of the agents of each type.

Selection and Mutation Pressure in Project Adoption. We study the expected evolution of traits conditional on the behaviors of early agents. We examine two conditionings: 1) All ancestral agents through I_n have adopted; or 2) not all ancestral agents adopted. Then, we examine unconditional trait evolution in an illustrative special case.

Case 1: All Ancestral Agents Adopted. Consider first the case, in which, in state θ , all agents adopt through agent I_n . There is only type $i=1$ in the a generation, so $q^0=0$ and $q^1=1$. Trivially, there can be no selection in the evolutionary sense, because the a population has only one allele.

Furthermore, there is stochastic mutation pressure, wherein the $z^i=1$ ancestors can map into a $z'^i=0$ descendant. Observe that $q'^1=1$; all agents in d are descendants of type $i=1$ in generation a .

Since it is not meaningful to study evolution when there are no ancestors, let \mathcal{B} denote the event that there is at least one uncensored ancestor. Substituting into the expectation of Eq. 12 conditional upon \mathcal{B} and on all past adopt, we obtain

$$\begin{aligned} \overline{\Delta z} &= \underbrace{E[z'^1 - 1 | z_1 = \dots = z_n = 1, \mathcal{B}]}_{\text{mutation pressure}} \\ &= \underbrace{P(z'^1 = 1 | z_1 = \dots = z_n = 1, \mathcal{B}) - 1}_{\text{mutation pressure}}, \quad [14] \end{aligned}$$

where on the right-hand side, the selection term of the Price Equation is zero. (The overline notation leaves implicit the conditioning of the expectation.)

This application of the Price Equation (Eq. 14) reveals that, even though there is *selection bias* on project payoffs that affects the evolution of the adopt/reject trait, there is no *selection* in the evolutionary sense. Selection bias induces cultural evolution by inducing mutation pressure in the trait value between ancestors and their descendants, rather than through selection on survival of ancestors.

This may seem counterintuitive, since a key driver of evolution here is that agents with low payoffs are “selected out.” However, all ancestral agents have the same trait, $i=1$. There is no variance for evolutionary selection to act upon. In contrast, censorship is based on whether the agents who adopted experienced high versus low payoffs. That censorship decreases the tendency for I_{n+1} to switch to reject.

The mutation pressure that drives expected trait change in Eq. 14 is negative, meaning that there is always a chance that the descendant of an A agent will reject. Mutation pressure here is caused by cognitive reasoning, as later agents assess state based on the payoffs received by earlier agents. Notably, stronger selection bias weakens the pressure toward R . Moreover, the form of mutation pressure here differs from mutation pressure in genetic settings, in that its strength depends on the number

[†]This is consistent with versions of the Price Equation in which a descendant can have multiple ancestors. In the biological context, the number of ancestors is usually one or two for asexual or sexual reproduction, but in cultural contexts, an agent can have many influencers.

of (uncensored) type-*A* agents in the ancestral generation. In genetic inheritance, a mutation bias has a fixed strength, and it applies to a single parent–offspring replication event. In our context, all of the uncensored type-*A* ancestors are collectively the causal parents of a single offspring. Furthermore, the greater their number, the greater the chance the offspring will convert to type *R*.

Case 2: A Past Reject Has Occurred. We now consider trait evolution conditional on at least one past reject. There is also at least one past adopt, since I_1 always adopts. Since the a generation has a mixture of adopt and reject alleles, there may be selection in the evolutionary sense.

Since there has been a past reject, all subsequent agents reject, so in the d generation, agent I_{n+1} always rejects. In fact, with no additional complexity, we now generalize trivially so that the d generation includes any number of agents I_{n+1} , I_{n+2} , and so forth.

Since a reject does not generate informative payoffs, and therefore has no causal effect on later agents, the reject type ($i = 0$) has no cultural descendants. In contrast, when an a agent adopts and is not censored, payoff information is transmitted to descendants and can influence behavior. Collectively, a sufficiently high number of low payoffs derived from uncensored past adopts causes the d agent(s) to reject. So, as before, each agent in the d generation descends from all of the agents of type $i = 1$, and only type $i = 1$ has descendants. So by definition, the descendant frequencies are $q^0 = 0$, $q^1 = 1$.

Since, under Case 2 conditioning, the descendants of type $z^1 = 1$ always reject, we have $z^{1'} = 0$. In other words, trait $i = 1$ has *perfect negative heritability*. This is reflected in the mutation pressure component of the Price Equation. Substituting these values into the expectation of the Price Equation, Eq. 12, conditional on at least one past reject and at least one uncensored adopt, gives

$$\overline{\Delta z} = \underbrace{-1}_{\text{mutation pressure}} + \underbrace{1 - q^1}_{\text{selection}}. \quad [15]$$

(The overline notation leaves implicit the conditioning of the expectation.)

The initial mean value of z^i in a is $q^0(0) + (1 - q^0)1 = 1 - q^0$. The initial a mean and the change in mean are negatives of each other, which reflects the fact that all descendants reject (mean trait value of zero).

The decomposition shows that evolution toward the reject allele derives from the opposition of two strong effects: evolutionary selection and mutation pressure. There is strong selection for A , since only A types leave descendants. However, this is overwhelmed by even stronger mutation pressure toward R . Specifically, evolutionary selection results in no descendants of the $i = 0$ type, yielding a positive selection term. But mutation pressure in the descendants of the $i = 1$ type is overwhelming; all of them shift from $z^1 = 1$ to $z^{1'} = 0$. This generates the -1 term. So the R allele becomes fixed in the d generation.

The mutation pressure in this setting is more extreme than in the previous case (Eq. 14), because this setting is deterministic: The descendent generation definitively rejects, and the rejection assuredly arose by mutation from an ancestor who had adopted.

Case 3: Without Conditioning on Whether Past Rejects. In *SI Appendix, section 7, part D*, we also perform a Price Equation decomposition in a case with minimal conditioning—only on the presence of at least one uncensored ancestor. This allows for the cases of All Past Adopt (Case 1) or Some Past Reject (Case 2) as possible realizations. The insights from Cases 1 and 2 carry through to this minimal-conditioning case. There is still opposition between the effects of selection and mutation pressure, with selection favoring adopt and mutation pressure favoring reject.

Taken together, the three applications of the Price Equation above provide a notable contrast between cultural and genetic evolution. Moreover, evolution in our context also differs from trait dynamics in other behavioral settings, such as in evolutionary game theory, where inheritance is determined by direct copying (27). Accurate copying leads only to selection, i.e., differential reproduction of traits. In contrast, agents here process information thoughtfully, which enriches the causality of trait transmission. Information transmission from parent to offspring results in mutation pressure, which can even overwhelm selection.

This application of the Price Equation offers two lessons. The first is that the Price Equation decomposition can apply to many social economic models, so long as some agents are influenced by other agents and a causal linkage between actions is used to specify inheritance.

The second lesson is that mutation pressure is just as important as selection in the cultural transmission process. In evolutionary economics and other applications, cultural transmission is often presented as simply a matter of differential survival of traits that compete to be copied. This is indeed part of the evolutionary process, even in settings with complex strategic interactions. However, as we have seen, there is also often systematic, endogenous mutation pressure arising from cognition. Such pressure can operate even when there is no selection, and it can overwhelm selection when the two forces are opposed.

Discussion

Biased information about others can profoundly influence the dynamics of investment risk taking, which we have studied as a cultural trait transmitted among firms. We have shown that when low-payoff outcomes are censored with higher probability than high-payoff outcomes, firms that do not account for this censorship bias become overly optimistic and undertake projects too often. This causes booms of overadoption, followed either by an eventual bust or by permanent long-run adoption, even in a state where complete information would assuredly lead to (rational) rejection of a risky project.

These dynamics are a form of cultural evolution in which parentage reflects causality. This is a richer form of cultural evolution than simple copying of successful traits, because it allows for cognitive reasoning about which traits are more likely to yield high payoff based on (biased) observations of prior outcomes.

Even for given size of project, some types of firms (such as large firms) tend to receive more attention than others (e.g., greater coverage by financial and media analysts). Since attention probabilities are bounded above by one, the upside salience of high-profile firms is necessarily limited; even their failures are noticed.

Specifically, when attention is high, censorship of low outcomes is less likely, which reduces overadoption. For example, if a huge moonshot initiative by a large, established firm (e.g., self-driving cars) were to fail, this could be as conspicuous as a success. So the model implies that overadoption will be more pronounced among start-ups and small firms than among large established enterprises. The extreme difference in visibility of success versus failure among start-ups provides an explanation for the overoptimistic expectations of entrepreneurs, and for the empirical anomaly of low returns to private equity (29).

Some firms receive higher attention for reasons other than sheer size. For example, as a creator of innovative consumer products, Apple has long been a magnet for attention, starting even before it became a giant. Firms in consumer-product businesses tend to attract greater attention than infrastructure firms. A failure of a project by a firm in a high-attention sector may be more salient than failure in a sector that receives little attention.

When a firm's decision making is influenced by observations of other firms, adopting or rejecting a project is a cultural trait

transmitted with bias between firms. We use the Price Equation to decompose such transmission into a selection component and a mutation component. Notably, although misperceptions of managers are driven by selection bias in what they observe, the Price Equation reveals that behavioral evolution is driven by the opposing effects of mutation pressure and evolutionary selection, and that in some conditions, there is only mutation pressure, without evolutionary selection. This contrasts sharply with a large domain of cultural evolutionary models with accurate copying, in which there is only selection.

The result of excessive adoption in our model is relative to a rational firm-level optimum. However, there are, in general, positive externalities to research and innovation. So excessive adoption that is unprofitable at the firm level may be welfare increasing at the social level. Innovative and moonshot projects may generate especially high externalities. On the other hand, there are also undesirable innovations, such as patent trolling and the use of hijacked airplanes as weapons.

We assume upside salience: that the probability that a high payoff is observed by others be greater than the probability that a low payoff is observed, but less than one. We expect moonshots to have high upside salience, because a rare, very high payoff is especially noticeable. In consequence, the model implies that overadoption will be more severe for moonshot projects. We also expect “sexy” projects (project that are innovative, fun, and exciting, such as self-driving cars) to have high upside salience. For given upside cash flow, we expect upside outcomes to be reported and remembered especially often when people feel the innovation has changed how people live in exciting ways. Therefore, there may be a tendency to overinvest in such sexy projects.

We have shown that upside salience causes managers to overvalue moonshot projects, resulting in overadoption and boom/bust patterns. For similar reasons, security-market investors may overvalue “lottery stocks” (stocks with positive skewness), consistent with evidence from stock returns, as summarized in ref. 7.

Our main focus has been investment at the firm level, but our approach can also be applied at larger scales to explain industry-level or aggregate-level investment boom/bust patterns. Industries may differ in upside salience of payoff outcomes, in part owing to differences in average firm size. So our model suggests that boom/bust patterns can be much more pronounced in some industries than in others. This is an interesting direction for empirical testing.

Stepping beyond the model somewhat, our approach also has organizational implications, including an advantage to managers of recruiting team members who are less heavily censored in observing others. Examples would include directors and venture capitalists who have broader direct experience, rather than observation through the media, than the manager or entrepreneur in past projects, especially for small start-ups.

Data Availability. All study data are included in the article and/or *SI Appendix*.

ACKNOWLEDGMENTS. We thank participants at the Conference on Evolutionary Models of Financial Markets, MIT Laboratory for Financial Engineering (virtual), June 2020; seminar participants at the University of Texas, Dallas, and the Evolutionary Models Conference discussant Terry Burnham; the American Economic Association 2021 Annual Meetings discussant Ilya Strebulaev; and two anonymous referees for insightful comments.

- R. Nisbett, L. Ross, *Human Inference: Strategies and Shortcomings of Social Judgment* (Prentice-Hall, Englewood Cliffs, NJ, 1980).
- L. A. Brenner, D. J. Koehler, A. Tversky, On the evaluation of one-sided evidence. *J. Behav. Decis. Making* **9**, 59–70 (1996).
- J. J. Koehler, M. Mercer, Selection neglect in mutual fund advertisements. *Manag. Sci.* **55**, 1107–1121 (2009).
- A. Banerjee, A simple model of herd behavior. *Q. J. Econ.* **107**, 797–817 (1992).
- S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026 (1992).
- J. Denrell, Vicarious learning, undersampling of failure, and the myths of management. *Organ. Sci.* **14**, 227–243 (2003).
- B. Han, D. Hirshleifer, J. Walden, Social transmission bias and investor behavior. *J. Financ. Quant. Anal.*, in press.
- K. H. Schlag, Why imitate and if so how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theor.* **78**, 130–156 (1998).
- P. Bolton, C. Harris, Strategic experimentation. *Econometrica* **67**, 349–374 (1999).
- H. H. Cao, B. Han, D. Hirshleifer, Taking the road less traveled by: Does conversation eradicate pernicious cascades? *J. Econ. Theor.* **146**, 1418–1436 (2011).
- W. Feller, *An Introduction to Probability Theory and its Applications* (Wiley, New York, NY, 1950).
- A. Tversky, D. Kahneman, Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131 (1974).
- R. Madan, C. L. Spetch, M. Is the enhancement of memory due to reward driven by value or salience? *Acta Psychol.* **139**, 343–349 (2012).
- P. Milgrom, J. Roberts, Relying on the information of interested parties. *Rand J. Econ.* **17**, 18–32 (1986).
- S. Giglio, K. Shue, No news is news: Do markets underreact to nothing? *Rev. Financ. Stud.* **27**, 3389–3440 (2014).
- R. S. Chirinko, H. Schaller, Business fixed investment and “bubbles”: The Japanese case. *Am. Econ. Rev.* **91**, 663–680 (2001).
- J. R. Ritter, The long-run performance of initial public offerings. *J. Finance* **46**, 3–27 (1991).
- R. G. Rajan, H. Servaes, Analyst following of initial public offerings. *J. Finance* **52**, 507–530 (1997).
- M. Lowry, G. W. Schwert, IPO market cycles: Bubbles or sequential learning? *J. Finance* **57**, 1171–1200 (2002).
- S. B. Moeller, F. P. Schlingemann, R. M. Stulz, Wealth destruction on a massive scale? A study of acquiring-firm returns in the recent merger wave. *J. Finance* **60**, 757–782 (2005).
- C. Bouwman, K. Fuller, A. Nain, Market valuations and acquisition quality: Empirical evidence. *Rev. Financ. Stud.* **22**, 633–679 (2009).
- P. DeMarzo, R. Kaniel, I. Kremer, Technological innovation and real investment booms and busts. *J. Financ. Econ.* **85**, 735–754 (2007).
- S. Gilchrist, C. P. Himmelberg, G. Huberman, Do stock price bubbles influence corporate investment? *J. Monetary Econ.* **52**, 805–827 (2005).
- P. O’Brien, R. Bushan, Analyst following and institutional ownership. *J. Account. Res.* **28**, 55–76 (1990).
- A. C. Cooper, C. W. Woo, W. C. Dunkelberg, Entrepreneurs’ perceived chances for success. *J. Bus. Ventur.* **3**, 97–108 (1988).
- G. R. Price, Selection and covariance. *Nature* **227**, 520–521 (1970).
- M. A. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life* (Belknap Press of Harvard University Press, Cambridge, MA, 2006).
- S. A. Frank, The Price equation, Fisher’s fundamental theorem, kin selection, and causal analysis. *Evolution* **51**, 1712–1729 (1997).
- T. J. Moskowitz, A. Vissing-Jørgensen, The returns to entrepreneurial investment: A private equity premium puzzle? *Am. Econ. Rev.* **92**, 745–778 (2002).