

NBER WORKING PAPER SERIES

VISIBILITY BIAS IN THE TRANSMISSION OF CONSUMPTION BELIEFS AND
UNDERSAVING

Bing Han
David Hirshleifer
Johan Walden

Working Paper 25566
<http://www.nber.org/papers/w25566>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019

We thank Pierluigi Balduzzi, Martin Cherkas, Marco Cipriani, Douglas Bernheim, Andreas Fuster, Giacomo de Giorgi, Mariassunta Giannetti, Luigi Guiso, Robert Korajczyk, Theresa Kuchler, Chris Malloy, Timothy McQuade, Peter Reiss, Nick Roussanov, Erling Skancke, Andy Skrzypacz, Tano Santos, Siew Hong Teoh, Bernie Yeung, Youqing Zhou, David Benjamin Zuckerman, and participants at the American Economics Association meetings in Boston, the MIT Conference on Evolution and Financial Markets, the National Bureau of Economic Research Behavioral Finance Working Group meeting at UCSD, the OSU Alumni Conference at Ohio State University, the Miami Behavioral Finance Conference at the University of Miami, the SFS Finance Cavalcade at the University of Toronto; the WISE18 Workshop on Information and Social Economics at Caltech; and seminar participants at Arizona State University, Boston College, Calgary University, Caltech, Cheung Kong Graduate School of Business, Columbia University, Hong Kong Polytechnic University, London Business School, National University of Singapore, New York Federal Reserve Board, Peking University, Renmin University, Shanghai University of Finance and Economics, Southwest University of Finance and Economics, Stanford University, University of Hong Kong, University of Houston, University of International Business and Economics, University of Texas Austin, and York University for very helpful comments; and Lin Sun and Qiguang Wang for very helpful comments and research assistance. David Hirshleifer is a paid advisor to EvoShare, a firm that helps employees save more for retirement. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Bing Han, David Hirshleifer, and Johan Walden. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Visibility Bias in the Transmission of Consumption Beliefs and Undersaving

Bing Han, David Hirshleifer, and Johan Walden

NBER Working Paper No. 25566

February 2019

JEL No. D14,D83,D84,D85,D9,D91,G02,G11,G41

ABSTRACT

We study how bias in the social transmission process affects contagion of consumption beliefs and behavior. In the model, consumption is more salient than non-consumption. This visibility bias causes people to perceive that others are consuming heavily and have favorable information about future wealth prospects. These inferences increase aggregate consumption. In contrast with other approaches, the visibility bias approach suggests that relatively simple disclosure policy interventions can ameliorate undersaving. In contrast with the Veblen wealth-signaling approach, information asymmetry about wealth *reduces* overconsumption. Our approach offers new implications about the effects on saving of social connectedness, observation biases, and demographic structure; and offers a novel explanation for the dramatic drop in the savings rate in the US and several other countries in the last thirty years.

Bing Han
Department of Finance
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, Ontario
Canada
Bing.Han@Rotman.Utoronto.Ca

Johan Walden
Haas School of Business
University of California at Berkeley
545 Student services building, #1900
Berkeley, CA 94720
walden@haas.berkeley.edu

David Hirshleifer
The Paul Merage School of Business
University of California, Irvine
4291 Pereira Drive
Irvine, CA 92697
and NBER
david.h@uci.edu

1 Introduction

In acquiring attitudes about the world, people are heavily influenced by their cultural milieu, and by interactions with others, especially when a clear conclusion is not evident by introspection. Several authors have argued that people have little idea how much they should save for retirement (e.g., Akerlof and Shiller (2009)), owing either to lack of relevant information, or failure to process it effectively. It is hard to know what stream of satisfaction will actually result from a consumption/savings rule chosen today.¹ It is also hard to forecast remaining lifespan or health in old age; most people do not process the relevant public but technical information contained in mortality tables and medical research. Finally, it is hard to predict risky future wealth realizations, and therefore, how much saving is needed today.

There is a great deal of evidence suggesting that people are indeed often ‘grasping at straws’ in their savings decisions, which suggests that they may look to social cues for help.² Consistent with this, there is also considerable evidence that social interactions affect several dimensions of consumption, saving, and investment choices.³

A large analytical literature shows that socially inefficient outcomes can arise from rational or biased social learning (e.g., see the survey of Golub and Sadler (2016)), and theoretical models have shown that social interaction can affect investment and saving behaviors such as market participation, house purchases, and the aggressiveness of trading in individual stocks (Hong, Kubik, and Stein 2004; Burnside, Eichenbaum, and Rebelo 2016; Han, Hirshleifer, and Walden 2018). Surprisingly, however, there has been little formal modeling of how bias in social learning processes affects lifetime consumption/savings choices. There is evidence of contagion of consumption and investment behaviors, but contagion can potentially spread either a decision to consume more or a decision to consume less. There is surprisingly little research on whether social learning induces a general *directional* bias toward over- versus underconsumption, and on the associated questions, what determines the size of such bias, and can policy interventions help remedy it?

We address this topic in a model in which social learning about others’ consumption expenditures is biased by the fact that engaging in a consumption activity is more salient to others than not doing so. This leads to biases in inferring others’ beliefs about future wealth

¹Allen and Carroll (2001) point out that “...the consumer cannot directly perceive the value function associated with a given consumption rule, but instead must evaluate the consumption rule by living with it for long enough to get a good idea of its performance. ... it takes a very large amount of experience ... to get an accurate sense of how good or bad that rule is.”

²People make very basic mistakes, and rely on noisy cues, in deciding how much to save (Samuelson and Zeckhauser (1988), Shefrin and Thaler (1988), Madrian and Shea (2001), Beshears et al. (2008), Benhassine et al. (2015)).

³Duflo and Saez (2002, 2003), Hong, Kubik, and Stein (2004), Brown et al. (2008), Charles, Hurst, and Roussanov (2009), Kaustia and Knüpfer (2012), Georgarakos, Haliassos, and Pasini (2014), Shemesh and Zapatero (2016), Giorgi, Frederiksen, and Pistaferri (2016), and the evidence reviewed in Hirshleifer and Teoh (2009).

prospects. For example, a boat parked in a driveway draws the attention of neighbors more than the absence of a boat. Similarly, it is more noticeable when a friend or acquaintance is encountered eating out or reports taking an expensive trip than when not, or buys an enjoyable product as compared with not doing so. We call the greater availability and salience of potential consumption events that do rather than do not occur *visibility bias*. We further assume that people do not adequately adjust for the selection bias in favor of noticing the consumption rather than nonconsumption events of others. This causes updating toward the belief that others are consuming heavily. So observers conclude that future consumption prospects are good, and therefore that low saving is appropriate. Observers therefore increase their own actual consumption.

Such visibility bias effects are self-reinforcing, as each individual becomes an overconsuming model for others. So biased learning generates a positive feedback that can result in severe undersaving in society as a whole, even when visibility bias is mild. In market equilibrium, the reluctance of individuals to save implies a higher interest rate.

The two premises of our model—that consumption activities are more available and salient to others than nonconsumption; and that people do not adequately adjust for the selection bias in their attention toward consumption—are motivated by the psychology of attention, salience, and social communication (see Section 2). With regard to communication, transmission of information in conversation is greater for positive information and for more arousing information (Berger and Milkman 2012); we expect enjoyable consumption in general to be more positive and arousing than the passive choice not to consume.

Visibility bias in our model need not be viewed as a cognitive failure; it is a source of bias in the *social transmission* of information. There are good reasons to allocate more attention to occurrences (or more generally, to salient events). However, failing to adjust appropriately for this selection bias in attention/observation *is* a clear mistake that produces a directional bias in inferences.

The model offers a new explanation for a well-known puzzle in savings rates. Personal saving rates in the U.S. have declined dramatically since the 1980s, from 10% in the early 1980s to a low of about 3% in 2007, while national debt has increased. This has raised concerns among many observers about whether Americans will be able to sustain their standards of living in retirement.⁴ A similar trend has occurred in many OECD countries, with ratios of household debt to disposable income often reaching well over 100% (OECD 2014). Economists have long argued that this drop is hard to explain (Parker 1999; Guidolin

⁴Many economists (e.g., Laibson, Repetto, and Tobacman (1998), Madrian and Shea (2001), Poterba, Venti, and Wise (2012), Poterba (2014), and Stanford Center on Longevity (2016)) argue that US households undersave, but some authors have different viewpoints about whether there is substantial undersaving (e.g., Scholz, Seshadri, and Khitatrakun (2006)); see also *Wall Street Journal*, June 22, 2018, “A Generation of Americans Is Entering Old Age the Least Prepared in Decades.” Over 50% of U.S. adults say that they could not easily come up with \$400 to cover an emergency expense (Federal Reserve Board 2018). More than half of households with bank cards carry debt from month to month, almost always at high interest rates; a substantial fraction borrow at close to their credit limits (Gross and Souleles 2002).

and Jeunesse 2007).⁵ There are, however, a wide range of possible explanations, and we recognize that different scholars may have different perspectives about this conclusion.

The visibility bias approach offers a novel explanation. The model is driven by observation of the consumption of others; greater observability of consumption intensifies the overconsumption effect. For example, the drop in costs of long-distance communication, the rise of cable television and VCRs (video cassette recorders), and subsequently the rise of the internet, greatly increased people's ability to observe others' consumptions, as people are able to hear, view, or report via social networks about consumption experiences. The rise of an increased diversity of cable television offerings (including channels devoted to shopping, travel, home remodeling, and other costly leisure pursuits, as well as dramas that less directly highlight consumption activities) further increased visibility. People often report by phone or other electronic networks on such activities as traveling, eating out, and recent product purchases. Social media for sharing pictures and videos, such as Instagram, have heavy emphasis on travel, fashion, and celebrities, all of which are associated with high observation of others' consumption. Such biased observation is the driving force behind overconsumption in our model.

There are also notable differences in savings rates across countries and ethnic groups which are not well-explained by traditional economic models (Bosworth 1993). A new possible explanation suggested by our approach is that there cultural differences affect communication about, or observability of others' consumption or wealths.⁶ Our model also implies that degree of urbanization will be negatively related to savings rate, as urbanization is associated with a higher intensity of social interaction and observation of the consumption of others. This is consistent with the evidence of Loayza, Schmidt-Hebbel, and Serven (2000), and is not directly predicted by non-social theories of consumption.

Overconsumption in our approach derives from mistaken beliefs about economic shocks and, in some cases, the behaviors of others. Such beliefs can potentially be corrected. So a distinctive empirical and policy implication of the visibility bias approach is that salient public disclosure about prospects for future consumption (which are, e.g., influenced by the risk of health and unemployment shocks), or about how much others actually consume, can help reduce overconsumption. Specifically, overconsumption in the model comes from complacency—people underestimate the probability of adverse wealth shocks. So vividly publicizing more accurate estimates about, e.g., risks of layoffs, or high health care bills,

⁵Many possible explanations have been offered for the decline in savings rates. However, Parker (1999) concludes that “Each of the major current theories of the decline in the U.S. saving rate fails on its own to match significant aspects of the macroeconomic or household data.” Guidolin and Jeunesse (2007) argue that factors such as greater capital mobility, new financial instruments, and aging populations do not suffice to explain the phenomenon, and conclude: “The recent decline of the U.S. private saving rate remains a puzzle.”

⁶Carroll, Rhee, and Rhee (1994) do not find an effect of culture on savings. They describe this as a tentative conclusion owing to data limitations. In contrast, using a similar methodology, Carroll, Rhee, and Rhee (1999) conclude that there *are* cultural effects.

should reduce overconsumption.

However, in practice, announcements of probability estimates may be hard for people to process and convert into consumption plans. This suggests, as an alternative policy intervention, saliently disclosing information about *how much others actually consume*. Under appropriate conditions, such disclosures reduces overconsumption. Furthermore, accurate disclosures that make saving behavior more salient (i.e., disclosures that are visibility-biased toward saving) can have an especially strong effect.

There is extensive evidence from social psychology that people often have biased perceptions about the attitudes and behaviors of others. For example, studies find that college students overestimate how much other students engage in and approve of uncommitted or unprotected sexual practices (Lambert, Kahn, and Apple 2003) and heavy alcohol use (Prentice and Miller (1993), Schroeder and Prentice (1998), Perkins and Haines (2005)), and overestimate the use of various other drugs (Perkins et al. (1999)). These studies argue that these misperceptions encourages such behavior, and that in some cases disclosure interventions can help remedy the problem.

In our model, agents observe upward-biased samples of the consumption of others. Agents neglect this selection bias, and update toward overoptimistic beliefs that on average favor higher consumption more than their priors would suggest. But surprisingly, in equilibrium actual consumption is so high that it confirms agents' high beliefs about others' consumption. So in contrast with the intuition from the abovementioned social psychology literature, the base version of our model demonstrates that social influence can induce more of a behavior without any overestimation of how much others engage in it. Since there is no overestimation of others' consumption, disclosure of others' average consumption does not correct agents beliefs.

However, some natural generalizations of the base model suggest conditions under which there will also, on average, be overestimation of others' consumption. For example, in reality some people have more accurate prior knowledge than others about the risk of adverse wealth shocks, or are less naive than others about visibility bias. In a simple setting with such "smart agents," agents on average overestimate the consumption of others, and disclosure of actual average consumption reduces overconsumption. Furthermore, in reality some people are more heavily connected than others in social networks. When we allow for this, we show that there can on average be overestimation of others' consumption, so that again disclosure of actual average consumption has a corrective effect.

Misestimation of the average actions of others also affects behavior in the model of Jackson (2018). In his framework, there is a strategic complementarity between agents' actions, a possible example being recreational drug use. This means that a higher action level by one agent rationally leads to more of the action by others. In the version of the model in which agents do not account for the *friendship paradox* in their observations of

others, agents overestimate how much others engage in the complementary action.⁷ This overestimation in equilibrium causes agents to engage more in the complementary behavior.

A key distinction between our model and Jackson's is that there is no strategic complementarity in our setting—utility of consumption does not depend upon others' consumption. Overconsumption in Jackson's setting derives from overestimation of consumption complementarities. Overconsumption in our approach instead derives from agents' naivete about visibility bias. So where Jackson's approach is based on the interplay between naivete and payoff complementarities, our approach is purely belief-focused.

In the network extension of our model, the friendship paradox is present, further amplifying overconsumption. The fact that naivete about the friendship paradox can be an amplifier of certain behaviors is in the spirit of Jackson's approach. However, even here, the mechanism works very differently. In Jackson (2018), highly-connected individuals (those with many social links) engage more heavily in the activity because their utility complementarities are stronger. In our setting, there are no such complementarities. Highly-connected agents are especially heavily influenced by visibility bias because they have more opportunities to observe their friends' consumptions and update beliefs in an upward-biased way.

The conclusion that agents, when young, overconsume extends to an overlapping generations setting in which the young can observe old as well as young agents. Overconsumption by the young is decreasing with the extent to which their observations are tilted toward the old. This tilt depends on the age distribution of the population, and on how visible and salient the consumption of old versus young agents is to young observers. So this finding provides a further set of distinctive empirical implications of our approach. Furthermore, since the old on average consume less than the young, in this setting there is another policy intervention that can reduce overconsumption—salient disclosure of the consumption of the old.

A plausible alternative theory of overconsumption and undersaving is that people are present-biased (i.e., subject to hyperbolic discounting, Laibson (1997)). Present bias is a preference effect, whereas the visibility bias approach is based on belief updating. Also, present bias is an individual-level bias, whereas the visibility bias approach is based upon social observation and influence. The visibility bias approach therefore has the distinctive implications that the intensity of social interactions and shifts in the technology for observing the consumption of others affect how heavily people consume. It also implies that population level characteristics such as wealth dispersion matter, in contrast with approaches based upon pure individual-level biases.

Another appealing approach to overconsumption is based on Veblen effects (Cole,

⁷The friendship paradox in general agent networks refers to the fact that the average number of friends (connections) of agents' friends is higher than the average number of friends agents have. The paradox arises generically, because well-connected agents are friends of more agents, and therefore disproportionately influential.

Mailath, and Postlewaite 1995; Bagwell and Bernheim 1996; Corneo and Jeanne 1997; Charles, Hurst, and Roussanov 2009), wherein people overconsume to signal high wealth to others. In wealth signaling models, beliefs are rational, whereas the visibility bias approach is based upon misperceptions. The visibility bias approach has distinct empirical implications as well. Information asymmetry about other people's wealths is the source of Veblen effects, which are not present when wealths are equal. In contrast, as shown in Section 5.2, in the visibility bias approach, overconsumption is strongest when there is *low* wealth dispersion and information asymmetry about wealth.

A third approach is based on agents deriving utility as a function of the consumptions of other agents (Abel 1990; Galí 1994; Campbell and Cochrane 1999). The concern for relative consumption is often referred to as the 'keeping up with the Joneses' approach. The preference interaction approach in general does not necessarily imply equilibrium overconsumption (Dupor and Liu (1993), Beshears et al. ((2018))), but this does arise in some settings (Harbaugh 1996; Ljungqvist and Uhlig 2000).⁸ In addition to unambiguously predicting a specific direction, overconsumption, the visibility bias approach offers various distinctive implications about the effects of disclosure and of shifts in visibility bias, wealth dispersion, and demographics.

As we have mentioned, a further distinctive empirical and policy implication of the visibility bias approach is that salient public disclosure can help correct people's beliefs, reducing overconsumption. That a relatively simple policy intervention can potentially ameliorate the undersavings problem is specific to the visibility bias approach.

Finally, another approach that can lead to overconsumption is based on speculative disagreement Heyerdahl-Larsen and Walden (2017). When investors with heterogeneous beliefs bet against each other in an asset market, they may all expect to profit, at least some of them mistakenly. Depending on agents' elasticity of intertemporal substitution, this can result in equilibrium overconsumption. Several of the implications discussed above also distinguish our approach from theirs. For example, the the speculative disagreement approach does not share the implications here about network properties and overconsumption.

2 Psychology Background

The two key assumptions of our model are that consumption activities are more available and salient to others than nonconsumption; and that people do not adequately adjust for the selection bias in their attention toward these consumption events. We now discuss evidence from the psychology of attention and salience that motivate and support our assumptions.

⁸Even with conventional preferences, externalities can also induce 'keeping up with the Joneses'-like effects (DeMarzo, Kaniel, and Kremer (2004, 2008)), though models based on this approach do not focus on the issue of over- or under- consumption.

With regard the first assumption, there is extensive evidence that occurrences are more salient and more fully processed than nonoccurrences (e.g., Neisser (1963), Healy (1981), the review of Hearst (1991), and Enke (2017)). Occurrences provide sensory or cognitive cues that trigger attention. In the absence of such triggers, an individual will only react if (as is usually not the case) the individual is actively monitoring for a possible absence. This is what is striking about the famous phrase “The dog that did not bark” in the Sherlock Holmes story; it takes a genius to detect and recognize the importance of an absence. An example of the low salience of non-occurrences is neglect of opportunity costs, i.e., hypothetical benefits that would occur under alternative courses of action. Indeed, the opportunity cost concept is something that students struggle with. Neglect of absences is also reflected in the principle of WYSIATI, “What you see is all there is,” one of the key features of System 1 thinking (Kahneman 2011). Consistent with the application of these ideas to consumption, Frederick (2012) concludes that “purchasing and consumption are more conspicuous than forbearance and thrift,” and gives the example that “Customers in the queue at Starbucks are more visible than those hidden away in their offices unwilling to spend \$4 on coffee.”

With regard to the second key assumption of our model, evidence from both psychology, experimental economics, and field studies of *selection neglect* confirms that observers often fail to adjust appropriately for data selection biases (Nisbett and Ross 1980; Brenner, Koehler, and Tversky 1996).⁹ In general, neglect of selection bias is implied by the representativeness heuristic of Kahneman and Tversky (1972). Owing to limited cognitive resources, adjusting for selection bias requires attention, and effort. Selection bias is especially hard for people to correct for because adjustment requires attending to the *non-occurrences* that shape a sample. A model of how neglect of selection bias affects economic decisions is provided in Hirshleifer and Teoh (2003). The combination of visibility bias and selection neglect in our model can be viewed as endogenizing the *availability heuristic* of Kahneman and Tversky (1973), so the tendency in the model to update toward thinking others are consuming heavily can alternatively be interpreted as coming from the use of this heuristic.¹⁰

Potentially consistent with the idea that the combination of visibility bias and selection neglect can affect perceptions, Frederick (2012) provides experimental evidence that

⁹People often naively accept sample data at face value (Fiedler 2008). Mutual fund families advertise their better-performing funds; in the experimental laboratory both novice investors and financial professionals misinterpret reported fund performance owing to selection neglect (Koehler and Mercer 2009). Auction bidders in economic experiments tend to suffer from the winner’s curse (neglect of the selection bias inherent in winning), and hence tend to lose money on average (Parlour, Prasnika, and Rajan 2007).

¹⁰According to the availability heuristic, people overestimate the frequency of events that come to mind more easily, such as events that are highly memorable and salient. The availability heuristic is therefore a failure to adjust for the selection bias in information brought to conscious attention—this being the subset of information that was stored into memory and is easy to retrieve from it. In our model, this is disproportionately information about the consumption activities that were engaged in rather than not engaged in.

the salience of consumption results in overestimation by observers of how much other individuals value certain consumer products. Consistent more broadly with the idea that visibility bias affects consumption behavior, there is evidence that people are influenced in car purchase decisions by observation of the purchases of others (Grinblatt, Keloharju, and Ikäheimo (2008), Shemesh and Zapatero (2016)), and such effects are stronger in areas where commuting patterns make the cars driven by others more visible (McShane, Bradlow, and Berger 2012).

Consumption activities of others may also be more cognitively available than non-consumption because someone who is consuming chooses to talk about it more than someone who is not consuming. Generally it is more interesting to hear about an action than inaction. Berger and Milkman (2012) provide evidence that online content is more likely to go viral when it is positive than negative, and more rather than less arousing. For various consumer products, positive word-of-mouth discussion of user experiences tends to predominate over negative discussion (see the review of East, Hammond, and Wright (2007)). A plausible reason is that users would like to persuade others of their expertise at product choice (Wojnicki and Godes 2008). This evidence suggests that people are more prone to sharing news about enjoyable consumption activities than news about stoical restraint from consuming.

3 The Base Model

In the base model, all agents have the same initial wealth and age. We start by focusing on individual optimization with interest rate given.

3.1 Optimal Consumption

Each individual maximizes a quadratic expected utility function with zero subjective rate of discount over two dates,

$$U = c_0 - \left(\frac{\alpha}{2}\right) c_0^2 + E \left[c_1 - \left(\frac{\alpha}{2}\right) c_1^2 \right]. \quad (1)$$

(We have verified in Appendix B that results that are qualitatively similar to the central results of our base model as in Proposition 1 below also hold with power and log utility.) At date 0, each individual chooses how much to consume and how much to borrow or lend at the riskfree interest rate $r = 0$, so the budget constraint is

$$c_1 = W - c_0 - \epsilon, \quad (2)$$

where, c_0 and c_1 are consumptions at dates 0 and 1, W is date 0 wealth, and ϵ is a potential wealth shock at time 1. We assume that

$$\epsilon = \begin{cases} 0 & \text{with probability } p \\ W & \text{with probability } 1 - p, \end{cases} \quad (3)$$

so with probability $0 < p < 1$, the agent's date 1 wealth is high, and with probability $1 - p$ it is low. We permit possible negative consumption, $c_1 < 0$. We assume that $\alpha W < 1$, to ensure that utility is increasing in consumption.

The negative wealth shock, which we view as being rare ($1 - p \ll 1$), can represent a systematic event, such as a major depression, or an underfunded pension system; or an idiosyncratic event that all agents are symmetrically exposed to, such as the possibility of a financially costly illness, disability or job loss. The key is that agents draw inferences about the probability of such events (even if their occurrence is independent across agents) from their observations of the consumption of others.¹¹

The agent's estimated probability that date 1 wealth will be high ($\epsilon = 0$) is \hat{p} . Based on this estimate, the agent chooses date 0 consumption to satisfy the first order condition, which by (1) yields optimal consumption

$$c_0 = \hat{p} \left(\frac{W}{2} \right). \quad (4)$$

In other words, date 0 consumption is proportional to the estimated probability that future consumption will be high. It follows that if people were sure of a high outcome, they would consume half their total wealth; if $\hat{p} < 1$ they consume less than half.

To analyze saving out of date 0 income, it is natural to think of total wealth, W , as being generated through time. Since the interest rate is zero, the agent's opportunity set as given in equations (2) and (3) is consistent with the agent receiving a cash flow of $W/2$ in each of the two periods, where if the wealth disaster occurs, a further cash flow of $\epsilon = -W$ is incrementally obtained at date 1. With this interpretation of the model, it then follows that optimal saving (at time 0) is

$$s_0 = (1 - \hat{p}) \left(\frac{W}{2} \right). \quad (5)$$

Empirically, there is evidence that fear of adverse wealth shocks strongly affects consumption/savings decisions (Malmendier and Shen 2018). Recent survey evidence suggests

¹¹An alternative modeling approach that would yield similar results would have agents learning from others about the probability of dying young, which would also affect the benefits from saving. Yet another approach would be to assume that owing to visibility bias, people overestimate the subjective discount rates (preferences) of others; and that owing to conformism, people update their own subject discount rates accordingly.

that, at least from an ex post perspective, people do think that the prospect of adverse economic shocks should be an important consideration for their savings decisions. Boersch-Supan et al. (2018) find that the experience of adverse economic shocks (such as prolonged unemployment, divorce, or a health shock) is a key determinant of “savings regret” (wishing one had saved more earlier in life), and argue that misperceptions about the probability of shocks is a root cause of undersaving. There is also evidence that people learn from the personal bankruptcies of neighbors about the risk of financial disaster, that learning of a neighbor’s bankruptcy causes people to cut back on their credit card expenditures, and that this results in a strong social multiplier effect (Agarwal, Qian, and Zou 2017).

3.2 Visibility Bias and Learning About Others’ Consumption

There are $N \gg 1$ identical agents, all facing identically distributed ϵ risks. The total date 0 potential consumption of an agent is divided into K different activities which we call “bins” (K large), where each bin represents potential consumption of $W/(2K)$. There are thus in total NK agent-consumption bins. We refer to a bin as *full* if it contains consumption and *empty* otherwise. An agent who chose to consume $W/2$ at date 0 (consistent with belief $\hat{p} = 1$, i.e., no risk of a negative shock) would then have all bins full, whereas an agent who chooses to consume 0 (consistent with belief $\hat{p} = 0$, i.e., certainty of the adverse outcome) would have all bins empty. (We have verified in Appendix C that our basic finding of overconsumption is robust to generalizing to allow for the possibility that when $\hat{p} = 1$ only a fraction of bins is full.)

We refer to the agent’s prior as the agent’s perceived distribution of p based only on private signals, not social observation. The prior reflects private information, which makes it useful for agents to update based upon the observation of others. Specifically, each agent has a Beta-distributed prior for p , which is based on observation of Q private signals about the level of p . So $p \sim \text{Beta}(Q_n, Q - Q_n) = \text{Beta}(Qq_n, Q(1 - q_n))$, where Q is a common natural number for all agents, $Q \in \mathbb{N}$, $0 \leq Q_n \leq Q$, and $q_n \stackrel{\text{def}}{=} Q_n/Q$.¹² Based on his signals, the agent’s prior estimate of the probability of a high outcome is then q_n . We call q_n agent n ’s *prior type*.

We view the prior type distribution as arising when each agent, starting with an improper $\text{Beta}(0, 0)$ “initial prior distribution” over p , takes Q drawings from a Bernoulli distribution with probability p , and uses the information from these outcomes to update his beliefs about p . The agent conditions on these observations to update to a prior $\text{Beta}(Q_n, Q - Q_n)$ distribution, based on Q_n successes (where each success provides information suggesting that disaster is less likely). So Q is a proxy for agents’ prior precision,

¹²In the cases $Q_n = 0$ or $Q_n = Q$, the prior, like the initial prior, is improper, but this does not lead to any technical problems with Bayesian updating of the beta prior. The Beta distribution, whose support is $[0, 1]$, is commonly used to describe the distribution of unknown probabilities of a Bernoulli random variable (in this case, the occurrence of the favorable wealth outcome).

and q_n for agent n 's optimism. The number of successes Q_n observed by agent n is ex ante stochastic and binomially distributed. Since N is arbitrarily large, the fractions of agents of different prior types are deterministic and follow a binomial distribution across agents, as follows by the Glivenko-Cantelli theorem.

Specifically, each agent's prior type is ℓ/Q , where $\ell \sim \text{Binom}(Q; p)$. By the Law of Large Numbers, it follows that the deterministic fraction f_ℓ , of agents associated with prior type ℓ/Q is

$$f_\ell = \binom{Q}{\ell} (1-p)^{Q-\ell} p^\ell, \quad \ell = 0, 1, \dots, Q.$$

By standard properties of binomial distributions over count variables (in this case, ℓ), it follows that the average prior type is

$$\sum_{\ell=0}^Q \left(\frac{\ell}{Q}\right) f_\ell = p. \quad (6)$$

So on average, agents' prior estimates are correct.

Since the average prior estimate in the population is correct, if agents were to choose their consumptions without observation of others, the average consumption level in the population would be $pW/2$. So if just one deviant agent in the large population could observe a sample of other agents' consumption bins, this would be informative, with a higher fraction of full bins indicative of higher p . The same point applies in our actual setting, in which agents simultaneously observe samples of others' consumption bins. Each agent updates his belief about p based upon the fraction of full bins in his observations of others.

Specifically, each agent observes a subset M bins of the other agents' $B = (N-1)K$ bins, and updates his belief about p based upon the fraction of full bins in his observations. He views his observations as an unbiased sample, which may not be the case. Crucially, we assume that observation is tilted toward those activities in which consumption did occur. This derives from what we call *visibility bias*, the tendency to notice and recall occurrences rather than non-occurrences.¹³

One reason that consumption activities are highly visible is that many are social, such as eating at restaurants, wearing stylish clothing to work or parties, and traveling. Furthermore, physical shopping is itself a social activity. Both physical and electronic shopping

¹³The occurrence versus non-occurrence distinction that we focus upon is not the only source of differences in the salience of different consumption behaviors. Extreme outcomes also tend to be more salient. Other things equal, we might expect this to cause observers to notice especially when others have either unusually low or unusually high total consumption, with no clear overall bias toward either over- nor under- estimation of others' consumption. Such effects (which we do not model) would basically be orthogonal to those we focus on. Our focus is on an attentional bias—neglect of nonoccurrences—that has a clearcut directional implication.

and product evaluation are also engaging topics of conversation, either in person or online. In contrast, saving is often a private activity and investing are often undertaken privately through banks, brokers, or retirement account software. Many television dramas display glamorous consumption activities, travel, entertaining, and dining, and some media channels explicitly focus on shopping and other costly leisure activities. There are of course exceptions to these generalizations, such as investment clubs, but overall, consumption tends to be more observable and salient to others than is saving.¹⁴

Specifically, if B^C of the B bins are full and B^N are not, then we assume that the chance that an *observed* bin is full is

$$\frac{k^C B^C}{k^C B^C + k^N B^N} = \frac{\frac{B^C}{B}}{\frac{B^C}{B} + \frac{k^N}{k^C} \left(1 - \frac{B^C}{B}\right)} = \frac{x}{x + \frac{1-x}{\tau}} \stackrel{\text{def}}{=} S_\tau(x), \quad (7)$$

where k^C is the probability that a bin is observed conditional upon it being full, k^N is the probability that a bin is observed conditional upon it being empty, $\tau = k^C/k^N \geq 1$, and $x = B^C/B$ is the consumption fraction. The parameter τ measures the overrepresentation of full bins in the observer's sample, i.e., *visibility bias*. When $\tau = 1$, the random observations match the actual distribution of consumption bins. When $\tau > 1$, there is overrepresentation of draws of consumption bins over non-consumption bins.¹⁵ The failure of the agent to adjust for this overrepresentation is a type of selection neglect—neglect of visibility bias. The number of full bins observed by agent n is Z_n , $0 \leq Z_n \leq M$. We define $z_n = Z_n/M$ as the fraction of bins that agent n observes to be full, and $z = \frac{1}{N} \sum_n z_n$. Visibility bias tends to increase z_n . Agents not understanding that there is visibility bias in consumption observations is the only deviation from rationality in the model. It is straightforward to show that the function $S_\tau(x)$ is strictly increasing in τ and $x \in (0, 1)$, is concave in x , and satisfies $S_\tau(0) = 0$, $S_\tau(1) = 1$, and $S_0(x) \equiv x$.

As we saw in equation (6), the average prior in the population is the true value p . After updating to belief \hat{p}_n , each agent n consumes $\hat{p}_n W/2$. If all agents were to update rationally, the average of their beliefs would still be p , so that average consumption in the

¹⁴Observability of two specific behaviors merit explicit discussion. First, retirement savings have very low visibility to others, so our approach suggests that people will underestimate such saving. Second, buying a house is highly visible. This is another example of the higher visibility of engaging in a consumption activity than not doing so. The purchase of a house is usually a shift to a higher flow of current consumption of housing services financed by a major increase in indebtedness (mortgage down payments are usually much smaller than the size of the loan), i.e., higher current consumption at the expense of future consumption. Indeed, real estate equity is often accessed to finance non-housing consumption expenditures as well (Chen, Michaux, and Roussanov 2013). Lusardi, Mitchell, and Oggero (2018) report that in recent years, older Americans close to retirement hold more debt than earlier generations, primarily owing to the purchase of more expensive homes with smaller down payments.

¹⁵We refer to the observer as observing a biased sample of target activities. However, the algebra of the updating process in the model can equally be interpreted as reflecting a setting in which observers draw unbiased random samples of observations, but where there is a bias in the ability to *retrieve* different observations for cognitive processing and the formation of beliefs.

large population would be the full-information optimal level, $pW/2$.

We assume that each agent uses Bayesian updating to estimate p based on his bin observations, under the belief that the fraction of full bins in the population is p . In other words, agents are unaware of visibility bias, so each agent thinks that average consumption of others in the large population is the full-information optimal level, $pW/2$. Defining $\xi = M/Q$, it follows that agent n 's posterior belief is

$$\hat{p}_n = \frac{Qq_n + Mz_n}{Q + M} = \frac{q_n + \xi z_n}{1 + \xi}. \quad (8)$$

The parameter ξ is the weight the agent puts on the new observations relative to the agent's prior. This captures the intensity of an agent's social interaction or social observation, and will be a source of some of the model's distinctive empirical implications. An agent who observes others more updates more based on social observation.

Owing to visibility bias, agents tend to update heavily from their priors toward a belief that others have high consumption. The variables z_n and q_n for agent n are ex ante stochastic, but in the limit with many agents, the average updated estimate across agents $\bar{p} = \frac{1}{N} \sum_n \hat{p}_n$, is, by the law of large numbers,

$$\bar{p} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_n \frac{q_n + \xi z_n}{1 + \xi} = \frac{p + \xi E[z]}{1 + \xi}. \quad (9)$$

This calculation takes as given the consumptions of all agents and the distribution of priors. Each agent consumes in proportion to his probability estimate (as shown in (4)), and the average agent estimate is \bar{p} , so this determines average consumption. In a static simultaneous equilibrium, this level of consumption determines the observations that agents draw from each other.¹⁶ So the expected fraction of full bins observed by an agent under visibility bias is just $S_\tau(\bar{p})$, i.e.,

$$E[z] = S_\tau(\bar{p}). \quad (10)$$

Specifically, each agent observes M bins, each with probability $S_\tau(\bar{p})$ of being full, leading to an expected observed consumption fraction of $E[z] = E[z_n] = S_\tau(\bar{p})$.

We define the mapping T from observation of consumption fraction $S_\tau(x)$ to the posterior belief as

$$T(x) \stackrel{\text{def}}{=} \frac{p + \xi S_\tau(x)}{1 + \xi}.$$

An equilibrium is defined as a solution \bar{p} to (9,10), i.e., a fixed point $\bar{p} = T(\bar{p})$. It is easy

¹⁶We examine an equilibrium in which all agents determine their consumptions simultaneously. In a variation of the model, agents choose consumption sequentially based on (biased) observations of previous agents' consumption. The large sample equilibrium with sequential observations converges to the equilibrium with simultaneous observations that we study.

to verify that the unique equilibrium \bar{p} when $\tau > 1$ is

$$\bar{p} = \frac{(\tau - 1)(p + \xi) - 1 + \sqrt{V}}{2(1 + \xi)(\tau - 1)}, \quad \text{where} \quad (11)$$

$$V = [(\tau - 1)(p + \xi) - 1]^2 + 4p(1 + \xi)(\tau - 1). \quad (12)$$

When $\tau = 1$, equilibrium is simply $\bar{p} = p$. With no visibility bias, agents are correct in their beliefs that the observed fraction of full bins in the entire population is p . When $\tau > 1$, agents are mistaken in updating as if there were no visibility bias in anyone's observations. We write

$$\bar{p} = B(\tau, p, \xi) \quad (13)$$

for the function defined by (11,12) for $\tau > 1$, where $B(1, p, \xi) = p$.

In equilibrium, different agents have different \hat{p}_n 's because of randomness in the number of Bernoulli successes built into each agent's prior, Q_n and the number of full bins observed socially, Z_n . But the aggregate estimate, \bar{p} , and corresponding population per capita consumption, \bar{c}_0 , are nonrandom by the law of large numbers. We call \bar{p} the *equilibrium probability estimate*. It is proportional to the population per-capita consumption, $\bar{c}_0 = \bar{p} \left(\frac{W}{2}\right)$. The *overconsumption factor*, the ratio of consumption to optimal consumption (which is $p \left(\frac{W}{2}\right)$), is therefore $\bar{p}/p \geq 1$.

It is easy to verify that when $\tau > 1$, there is a positive feedback effect:

$$\bar{p} > \frac{p + \xi S_\tau(p)}{1 + \xi}. \quad (14)$$

In equilibrium an agent has higher consumption owing to visibility bias, thereby inducing higher consumption by other agents. This in turn encourages even higher consumption by the original agent. This feedback effect is reflected in the difference between the left-hand-side and right-hand-side of (14). The feedback effect can be powerful, especially when ξ is high. The equilibrium has the following properties.

Proposition 1 *In equilibrium:*

1. *The equilibrium probability estimate, \bar{p} , and aggregate consumption are increasing in visibility bias, τ , i.e., $\partial\bar{p}/\partial\tau > 0$, with $\bar{p} = p$ when $\tau = 1$;*
2. *As $\tau \rightarrow \infty$, $\bar{p} \rightarrow (p + \xi)/(1 + \xi)$, so that $0 < \bar{p}_\infty < 1$;*
3. *If $\tau > 1$, the average estimated probability of high consumption, \bar{p} , and aggregate consumption are increasing with ξ , the intensity of observation of others' consumption bins, and in the limit approaches 1 as ξ becomes large.*

Intuitively, Part 1 says that owing to visibility bias in consumption observations, and neglect of sample selection bias (or equivalently, use of the availability heuristic) in assessing

frequencies, people update heavily away from their priors toward a belief that others are consuming heavily. In consequence, observers update too favorably about the information others have about the probability of wealth non-disaster. This causes people to overconsume, and the greater the visibility bias, the larger the effect.

Part 2 indicates that when visibility bias becomes maximally strong, beliefs become maximally overoptimistic, but that agents' prior beliefs have a moderating effect, so that the equilibrium beliefs do not spiral upward toward $\bar{p} = 1$. Agents put some weight on their priors, so even if 100% of observed bins are full, observers only update their beliefs to $(p + \xi)/(1 + \xi) < 1$. The prior beliefs thus limit the severity of overconsumption.

Part 3 says that owing to visibility bias, greater observation of others as reflected in ξ implies more optimistic beliefs and greater aggregate consumption. As observation of others becomes very large relative to the prior precision, so that ξ approaches infinity, there is drastic overconsumption (agents consume as if they were sure there were no risk of a bad outcome). New biased observations dominate prior information, so that people become certain of a high outcome, even if visibility bias is small (τ close to one) and the probability for a high outcome, p , is low. This is because when agents place heavy weight on socially derived information, the feedback effect becomes very strong. Since $\bar{p} = p$ when $\tau = 1$, this also means that equilibrium consumption is very sensitive to changes in τ for large ξ . Together, Parts 2 and 3 suggest that the feedback effect inherent in social transmission may be more important in generating severe overconsumption than visibility bias itself, as long as there is some visibility bias.¹⁷

As discussed in the introduction, personal saving rates have plunged in the U.S. and several other OECD countries over the last 30 years, and existing rational theories do not seem to fully explain this phenomenon. Parts 1 and 3 of Proposition 1 provide a possible explanation.

Over the last several decades, improvements in electronic communications by such means as phone (the drop in cost of long-distance telephone service), the rise of cell phones and email in the early 1990s, the rise of internet in the late 1990s, and blogging and social networking (such as Facebook) over the last decade have dramatically reduced the cost of conveying information about personal consumption activities. This is reflected in our model as an increase in both τ and ξ , as in Parts 1 and 3 of Proposition 1. Greater observation and communication in general about the behavior of others is reflected by higher ξ in the model. Greater ξ intensifies the effects of visibility bias by increasing the weight on social observation relative to the prior, and implies a reduction in the savings rate.

Crucially, these technological changes also strongly suggest an increase in bias toward observing consumption over nonconsumption, i.e., visibility bias τ . The activities that

¹⁷Empirically, social learning can indeed induce strong feedback effects in consumption behavior. Moretti (2011) provides evidence that social learning about movie quality induces a large 'social multiplier,' wherein observation of others greatly increases the sensitivity of aggregate demand to quality.

are noteworthy to report on very often involve expensive purchases, as with eating out or traveling. Indeed, numerous television dramas and reality shows have long had a focus, implicit or explicit, on such consumption activities. The explicit side includes travel and shopping channels. For example, the first national shopping network began in 1985 as the Home Shopping Network. The implicit side includes dramas, not limited to those centered upon the antics of the wealthy (“Who shot JR?”). The shift to reality television also induced greater observation of the consumption activities of others.

In more recent years, social media and review sites have been organized around consumption activities, such as Yelp and TripAdvisor. The universe of YouTube video postings includes travel and other consumption activities. On Facebook, a posting about a consumption event triggers a notification to friends; a non-posting about not engaging in a consumption event does not. Participants in special interest online discussion sites (e.g., focused on high tech or classical music) often post about associated product purchases. Such postings are more interesting, and therefore more likely to occur, than a posting to announce the news that the individual did not buy anything today.

In contrast, in-person unmediated observation of physically proximate friends or acquaintances are likely to often include even nonconsumption activities. So the rise in modern communications results in an increase in visibility bias (i.e., larger τ) and lead to higher overconsumption.^{18,19}

Past social research has also used other proxies for the intensity of social interaction and observation (ξ), such as population density (e.g., urban versus rural).²⁰ This leads to the empirical implication that after appropriate controls, greater population density is associated with lower saving. In the time series, this suggests that the secular increase in U.S. population over time may also have contributed to the decline in the savings rate.

The social influence parameter ξ is identical across individuals. With diverse ξ 's, we expect that individuals who engage in greater social observation will overconsume more than those with lower ξ . Such individuals update their beliefs more optimistically. A similar point holds for individuals who are more subject to visibility bias, i.e., greater τ .

¹⁸Increased internet usage—especially through online social networking platforms—is associated with a larger number of ‘weak ties’ (merely casual acquaintances) in ones’ social network. Such weak ties are especially useful for acquiring information and ideas (Donath and Boyd 2004; de Zúñiga and Valenzuela 2011). Also, a social networking platform that relies on advertising for its revenues may have an incentive to disproportionately convey notifications that relate to consumption activities.

¹⁹Hirsh (2015) provides evidence that the drop in savings rates was accompanied by increasing population-level extraversion in many countries. Hirsh’s shifting extraversion explanation is compatible with our approach, since greater sociability causes greater observation of others’ consumption. However, even in the absence of shifts in population-level psychological traits, our model can explain the drop in the savings rates by improvements in communication technologies.

²⁰People who are geographically closer tend to interact more (Borgatti et al. (2009)), even after the rise of the internet and low-cost telephony (Mok et al. (2010)), and even in online social networks (Scellato et al. (2010)). Sociologists have argued that people in urban areas have more voluntaristic social linkages, as contrasted, e.g., with family ties (White and Guest 2003). These findings suggest that greater population density increases the opportunities for people to interact and observe each other.

It is evident that these predictions hold for the case in which ξ or τ is identical for almost everyone.

Proposition 2 *Consider a society with common social observation parameter ξ and visibility bias parameter τ , with the exception of a deviant individual who has a social observation parameter value of ξ' , or alternatively a visibility bias parameter value τ' . Then the expected consumption of the ξ -deviant is increasing with ξ' , and the expected consumption of the τ -deviant is increasing with τ' . A ξ -deviant on average consumes more than the others if and only if $\xi' > \xi$. A τ -deviant on average consumes more than the others if and only if $\tau' > \tau$.*

The result follows from (8), since

$$E[\hat{p}] = \frac{p + \xi' S_{\tau'}(\bar{p})}{1 + \xi'},$$

which is increasing in ξ' and τ' (and a small deviant fraction does not alter the average probability estimate \bar{p}).

Proposition 2 suggests that people who engage in greater social observation or are more subject to visibility bias will overconsume more. These implications are empirically testable. For example, survey data has been used to study reported investment behavior in relation to households' sociability or intensity of social interaction, in the form of self-reports of interactions with neighbors or regular church-going (Hong, Kubik, and Stein 2004; Georgarakos and Pasini 2011)). Some studies have exploited information about actual social media connections (Heimer (2016), Bailey et al. (2016, 2018)).

Since neglect of visibility bias is an error, we expect τ to be higher for individuals and groups that are more subject to psychological bias, such as those with lower education and IQ. Psychometric indices such as scores based upon the Cognitive Reflection Task (see the discussion in Frederick (2005)) provide more direct ways of measuring whether an individual is likely to fail to adjust for selection bias (in this case, visibility bias).

A possible objection to the conclusion of overconsumption is that houses serve as investment as well as consumption vehicles, and are highly visible to others. However, as discussed in footnote 14, the purchase of a house tends to be associated with an increase in the consumption of housing services, financed heavily by debt—i.e., an increase in current consumption at the expense of future consumption.

A related objection is that in a multiperiod setting, the purchase of a house could be an indicator that an individual had saved heavily to accumulate enough for a substantial down-payment. We extend the model to allow for observation of the old by the young in Section 5.1 to allow for inferences about past saving. We show there that in equilibrium, unambiguously, there is still overconsumption.

An interesting feature of the base model is that, despite naivete about visibility bias, agents end up with correct beliefs about others' average consumption and beliefs. To see

why, recall that all agents expect others to on average consume based on the correct value of p . In other words, agents do not recognize that others overconsume. Similarly, each agent thinks that his own consumption is on average based on the correct value p . So each agent believes that the consumption of peers is on average the same as his own. Since all agents are ex ante identical (apart from their unbiased prior signals), they are correct in thinking so. In other words, on average agents correctly assess the average consumption of others.

This may seem counterintuitive, since agents are updating naively about the consumption of others based upon upward-biased samples. However, since agents think that others are not overconsuming, based upon priors, agents on average *underestimate* others' equilibrium consumption. On average these two effects exactly offset, so that people end up on average with correct assessments of others' average consumption.²¹

3.3 Policy Interventions and Smart Agents

Overconsumption in our model derives from overestimation of safety with respect to adverse wealth shocks (overestimation of p). This suggests that a relatively simple policy intervention—saliently publicizing valid information about the risk of wealth shocks—can help alleviate overconsumption. For example, publicity about the frequency of layoffs or of expensive illness could be beneficial. Interpreted more broadly, low p could be the risk of living a long time, resulting in higher-than-expected post-retirement consumption needs. So salient publicizing of life expectancy information could help.

However, in practice it may be hard to make such disclosure salient and easy for people to interpret. Research on heuristics and biases consistently finds that people tend to put little weight on base rate probability information such as a numerical report about risk of layoffs (Kahneman and Tversky 1973; Borgida and Nisbett 1977). Furthermore, people may have trouble interpreting mortality or life-expectancy tables, which require significant cognitive processing to translate into an optimal plan for how much to save.

We therefore consider other possible types of disclosure, and their effectiveness when there are 'smart agents.' In Subsubsection 3.3.1 we consider disclosures of the average consumption or saving rate of peers, where there is no visibility bias in the disclosure. By this we mean that the disclosure does not do anything special to highlight consumption versus nonconsumption activities, nor the actions of those who are consuming heavily versus those who are saving heavily. In Subsubsection 3.3.2, we consider such visibility-biased disclosures.

²¹To put this another way, observers update from their priors toward on-average-high consumption observations of others, where these observations tend to be high even relative to others' actual high consumption. But observers still place positive weight on their priors, so their updates are only partial. As is standard in Bayesian updating, they attribute the high level of the signal in part to randomness (or in the context of observing consumption bins, sampling error).

3.3.1 Public Disclosure of Average Consumption of Peers

A possible policy intervention suggested by the visibility bias approach is to publicize something simple which translates fairly directly into a consumption/savings recommendation: the average consumption or saving rate of peers. Since people in our approach do not take into account that others are biased, under plausible variations of the base model assumptions, as we shall see, people may end up with biased perceptions about what others believe and how much they consume. If so, saliently publicizing accurate information about others can help alleviate overconsumption. Specifically, if people overestimate how optimistic others are, and how much others consume, then accurate information about others would correct that mistake, reducing overconsumption.

Empirically, in tests covering a very wide range of activities, the intervention of providing accurate information about peer beliefs or behavior tends to cause behavior to conform more closely to the disseminated peer norm (e.g., Frey and Meier (2004), Cialdini et al. (2006), Salganik, Dodds, and Watts (2006), Goldstein, Cialdini, and Grieskevicus (2008), Cai, Chen, and Fang (2009), Gerber and Rogers (2009), and Chen et al. (2010)).²² Social norms marketing, or dissemination of information about one's peers, can be an effective policy tool for correcting inaccurate beliefs about peers, and to makes peer actions more salient.

In an example of the effects of peer disclosure, informing college freshmen of survey results about the attitudes of other college students toward heavy drinking is associated with lower levels of self-reported drinking 4-6 months later, as well as lower acceptance of pro-drinking norms (Schroeder and Prentice 1998). The explanation offered by the authors was that students were conforming to a better and more realistic norm. Bursztyn, González, and Yanagizawa-Drott (2018) report that a very large majority of young married men in Saudi Arabia support female labor force participation outside of home, but substantially underestimate how much similar men support this. In an incentivized experiment, randomly correcting beliefs about other men increases men's willingness to let their wives join the labor force, and actual labor force participation by their wives. In the context of our model, these findings raise the question of whether salient reporting of actual consumption of peers would reduce consumption.

We have seen that in the base model as developed so far, agents on average have correct assessments of the average consumption of others. It follows that disclosure of the

²²However, there are exceptions in which people adjust their behavior away from the disclosed actions of others. In an experiment on retirement savings behavior in a large manufacturing firm, Beshears et al. (2015) document that information about the high savings rates of other employees can sometimes lead low-saving individuals to shift away from the disclosed savings rates, which Beshears et al. suggest may derive from a discouragement effect. This result holds only for the subpopulation of employees with low relative incomes who had never participated in the firm's 401(k) plan. Such employees may regard the higher-income employees who were plan participants as not truly peers (in the context of our model, e.g., not having identically distributed wealth shocks). So our theory does not make a prediction about the outcome of this experiment.

actual average consumption of others will not change the average level of overconsumption. Specifically, an agent with prior q_n and bin observations z_n , who observes a publicized signal of \bar{p} (or equivalently of aggregate consumption), arrives at the posterior

$$\hat{p}_n = \frac{q_n + \xi z_n + \alpha \bar{p}}{1 + \xi + \alpha}.$$

Here, the parameter $\alpha \geq 0$ determines how much weight the agent puts on the public signal.²³

The equilibrium condition then becomes

$$\bar{p} = \frac{p + \xi S_\tau(\bar{p}) + \alpha \bar{p}}{1 + \xi + \alpha}, \quad (15)$$

which, as is easily verified, leads to the same equilibrium as in the base model (which corresponds to $\alpha = 0$).

However, in a slight generalization of the base model, there is systematic overestimation of the consumption of others. Suppose that there is a group of more knowledgeable or sophisticated agents (“smart agents”). Specifically, we now allow for a second group of agents consisting of fraction ϕ of the population who end up with unbiased expectations, i.e., their average posterior belief is p . This would occur, for example, if these are smart agents who rationally adjust for visibility bias in their observations. Alternatively, even if these agents are subject to visibility bias, this outcome will occur if they have strong prior information about p (i.e., the number of signals Q reflected in their priors is arbitrarily large). For example, these agents may have studied the statistics on the frequency of expensive illness or of job loss. Then their beliefs will tend to be close to p , and they won’t update much when they learn the population average consumption. The remaining $1 - \phi$ agents, just as before, neglect visibility bias.

Now the average population belief \bar{p} , and therefore also aggregate consumption, reflect a balance of beliefs between the smart agents and the other agents (whom we will refer to as ‘biased’). Let \bar{p}^V be the average belief of the biased agents (V for visibility bias). Intuitively, in the absence of public disclosure, the average belief is dragged down by the smart agents (beliefs near p), and dragged up by the optimism of the biased agents ($\bar{p}^V > p$). This implies that the beliefs of the biased agents are above average ($\bar{p}^V > \bar{p}$). So a salient disclosure that indicates that the average belief is \bar{p} pulls down the beliefs of the biased agents substantially, and only modestly lifts the beliefs of the smart agents. So overall \bar{p} declines, reducing overconsumption.

²³In a large population, a perfectly accurate disclosure of the average behavior would reflect an extremely large sample. This would be so informative that agents would put arbitrarily heavy weight upon the public signal. However, in reality empirical proxies for aggregate consumption are noisy, so that a Bayesian (or quasi-Bayesian) would put only a finite weight on it. For simplicity, instead of modeling noise in detail, we assume that each agent places only a finite weight α on the disclosed average consumption of others.

To formalize these intuitions minimally, consider the case in which the smart agents (whom we have described as either rational or well-informed) know the true p . Then the biased agents observe $S_\tau(\phi p + (1 - \phi)\bar{p}^V)$ and neglect the selection bias in this observation. Also, just as above, they put weight α on the public disclosure of \bar{p} . So equilibrium with the public signal can be defined as the modification of (15),

$$\bar{p}^V = \frac{p + \xi S_\tau(\bar{p}) + \alpha \bar{p}}{1 + \xi + \alpha}, \quad \bar{p} = \phi p + (1 - \phi)\bar{p}^V. \quad (16)$$

It follows that biased agents believe that average consumption is higher than it is, $\bar{p}^V > \bar{p} = \phi p + (1 - \phi)\bar{p}^V$. Comparing the case $\alpha > 0$ with the no-disclosure case of $\alpha = 0$, we see that in the smart agents model, publicizing a (potentially noisy) signal about the average belief \bar{p} or average consumption of others in equilibrium decreases aggregate consumption.

The above discussion implies:

Proposition 3 *Consider the modification of the base model from Section 3.2 in which there is a fraction ϕ of ‘smart’ agents who know the true probability p of no disaster, and where average consumption is publicized. Then*

1. *Agents on average overestimate the average consumption of others.*
2. *Average consumption is decreasing in the weight, α , that agents with visibility bias assign to the public signal.*

We are not aware of evidence about whether people tend to overestimate the aggregate consumption of others (as in this smart agents model), or not (as in the base model equilibrium). Which of these is the case has important policy implications.

The conclusions of Proposition 3 would also hold even if ‘smart’ agents were less smart than we have assumed, as long as their average belief, in the absence of disclosure, is below the population average, and they are sufficiently resistant to updating upward in response to disclosure of the population average. In the proposition, these two features are achieved by having all smart agents know the true $p < \bar{p}$. This certain knowledge makes them resistant to updating upward after disclosure.

Another possibility is that smart agents are individuals who are overconfident in the sense that they mistakenly have more faith in the accuracy of their own signals than in the signals of others.²⁴ As such, they are little influenced by their visibility-upward-biased personal observations of others, causing their average beliefs to be below the population average. Furthermore, when the public signal about average consumption is disclosed, they

²⁴We consider here overconfidence about priors that are, on averaged unbiased. I.e., we are not referring here to systematic overoptimism. Overconfidence in the sense of overestimating the accuracy of one’s own beliefs is extremely well-documented in the psychology and economics literatures (see, e.g., the surveys of Daniel and Hirshleifer (2015), Moore, Tenney, and Haran (2015)). It is also well-documented that people are heterogeneous in their degrees of overconfidence.

place little weight on this information as well, so they do not revise their beliefs upward very much.

Another pathway to the result is if ‘smart’ agents are smart only in the sense that they understand that there is visibility bias (where it is not crucial whether or not smart agents understand that other agents neglect it). If so, then smart agents will not update upward as strongly as other agents do based upon personal observations of others, so again, the beliefs of smart agents will be below the population mean.

When we consider our extension to a setting with overlapping generations, we will see why a different kind of disclosure—of the consumption of only a subset of the population—can also help reduce overconsumption.

A different possible approach to correcting overconsumption, which we explore in the next section, is to fight visibility bias with visibility bias. If there is a way to make saving behavior more salient, observers who neglect visibility bias will tend to update more toward a belief that others save heavily, and hence that saving is desirable. This approach can be effective even when there are no smart agents in the model, and no misperceptions about others’ average consumption.

3.3.2 Visibility-Biased Disclosures

We now examine disclosure that is subject to visibility bias, in the sense of increasing the salience of consumption versus nonconsumption behaviors. We first examine analytically the effect of accurate public disclosures that make the saving behavior of others more salient. We verify that in equilibrium such visibility-biased disclosures encourage saving. We then discuss in practical terms how to make a disclosure that highlights saving behavior, and how effective such disclosures are likely to be. Since the effects we describe here do not require the smart agents considered in the preceding subsection, we return to the base assumption that all agents are *ex ante* identical.

As we have seen, when a plain vanilla disclosure is added to the base model with all agents *ex ante* identical, as in equation (15), equilibrium consumption is unaffected. Suppose instead that the signal calls attention more to nonconsumption than to consumption behavior.²⁵ We will refer to this as a signal about “saving,” but clearly in substantive terms, a disclosure about saving is equivalent to a disclosure about consumption, since one can be inferred from the other. Consumption is proportional to \bar{p} , and saving to $1 - \bar{p}$ (see equations (4) and (5)). Observing a signal about average saving is thus equivalent to observing a signal about $1 - \bar{p}$. So we are now referring to a disclosure that makes *salient* the saving rather than consumption activity of others.

What does it mean to make an accurate disclosure which is biased toward visibility of

²⁵Importantly, we are only examining accurate disclosures. They direct attention to saving behavior, but are not misrepresenting the amount of it that is occurring.

saving rather than consumption? People could be given stickers that they are free to post on their cars or in personal spaces, saying “Proud Retirement Saver.” The policymaker could have an advertising campaign explaining that these signs or stickers are given to anyone who is saving more than some prespecified absolute amount or fraction of income. There is visibility bias, since the presence of the sign or sticker is more noticeable than the absence of one.

In the model, if there were no visibility bias to the saving disclosure, in (15) above, agents would infer that a signal about saving (basically $1 - \bar{p}$) is equivalent to a signal about consumption (basically $\bar{p} = 1 - (1 - \bar{p})$). As we have seen earlier, in the absence of visibility bias, a disclosure about consumption has no effect on average saving. However, if there is visibility bias in the disclosure toward the occurrence of saving behavior, and observers neglect this visibility bias, then observed saving is overestimated. Specifically, the visibility bias function in the disclosure has the same value, τ , as the visibility bias in the direct observation of others, so that the visibility bias function S_τ is applied to the amount saved. It follows that equilibrium beliefs satisfy

$$\bar{p} = \frac{p + \xi S_\tau(\bar{p}) + \alpha(1 - S_\tau(1 - \bar{p}))}{1 + \xi + \alpha}. \quad (17)$$

The $1 - S_\tau(1 - \bar{p})$ -term in the expression reflect the agent’s neglect of visibility bias about saving. The agent believes that he is observing a signal about $1 - \bar{p}$ (proportional to saving), when he is actually observing a signal about the visibility-biased quantity $S_\tau(1 - \bar{p})$. This leads to the inferred belief (proportional to consumption) of those observed via the disclosure of $1 - S_\tau(1 - \bar{p})$, rather than $\bar{p} = 1 - (1 - \bar{p})$.

Equation (17) has a unique closed form solution, and it is not hard to derive the following:

Proposition 4 *Equilibrium consumption when there is visibility bias toward a disclosure of higher saving, as given in (17), is lower than when there is no the disclosure, as given in (13).*

Publicizing the visibility-biased signal about savings thus unambiguously results in lower consumption. In other words, the disclosure fights visibility bias with visibility bias. Neglect of visibility bias about the disclosure encourages saving by partly offsetting the effects of visibility bias in agents’ direct observations of others’ consumption activities.

For several reasons, interventions by policymakers to make saving behavior more salient, as reflected in the α term in the numerator of equation (17), is unlikely to fully offset the spontaneous visibility bias toward observing the consumption activities of others, as reflected in the ξ term. People are heavily exposed to the consumption of others many times each day as people interact with others in-person or electronically. In contrast, policy campaigns are likely to be episodic and to generate a relatively limited number of

observations. Nevertheless, if the campaign succeeds in making observations about others' saving behaviors sufficiently salient, there can be a beneficial effect.

A drawback of the sticker intervention that we have described is that people have little incentive to post the stickers. Indeed, some might fear that posting such a sticker would look like bragging. However, a more nuanced scheme could potentially be designed to minimize these problems.²⁶

Disclosure that is visibility-biased toward saving is, in a sense, a more powerful way of encouraging saving than nonselective disclosure, because it enlists visibility bias to shift beliefs in favor of saving. Suppose, for example, that a nonselective disclosure about consumption levels reports that among some set of 100 bins, 75 are full, whereas a saving-visibility-biased disclosure reports that within some set of 100 bins, 40 are full. Since observers neglect visibility bias, clearly the second intervention will push beliefs farther downward.

4 Observation of Others in a Social Network

We can describe an agent's linkages in a social network in terms of whose consumption an agent can observe. So it is interesting to study how network location affects consumption and perceptions of the consumption of others. We therefore extend the model to allow for an arbitrary social network. This allows us to derive empirical implications for how individual centrality and overall network connectedness affect beliefs and consumption.

Agents are connected in an undirected social network represented by the graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where \mathcal{N} is the set of agents and \mathcal{E} is the set of edges connecting them. The set of agents $\mathcal{N} = \{1, \dots, N\}$, and $(m, n) \in \mathcal{E} \subset \mathcal{N} \times \mathcal{N}$ if agents m and n are connected in the sense that they potentially can observe each other's consumption. By convention, the network is undirected, i.e., $(m, n) \in \mathcal{E} \Leftrightarrow (n, m) \in \mathcal{E}$, and agents are not connected to themselves ($(n, n) \notin \mathcal{E}$). The set of agents that n is socially linked to is $\mathcal{D}_n = \{m : (n, m) \in \mathcal{E}\} \subset \mathcal{N} \setminus \{n\}$, and n 's connectedness is $d_n = |\mathcal{D}_n|$. The maximal degree in the network is $D = \max_n d_n$.²⁷

²⁶For example, the scheme could be localized, with members of a neighborhood receiving the "Proud Retirement Saver" stickers. Signs would be set up in the neighborhood as locations for people to post their stickers anonymously, but people would also be encouraged to post them on their own cars, or at their homes or workplaces. The policymaker counts the number of stickers posted on the public sign locations. For posts in personal locations, people self-report to the policy-maker. Savers as a group are given monetary bonus when there are more posted stickers in the area. Again, the presence of the stickers leads to a clear visibility bias toward saving over consumption.

²⁷To ensure that there is a large enough number of agents so that the law of large numbers can be used (as in the previous section), we make the technical assumption that there are a many agents at each node position in the network, where each agent's prior is still drawn independently of all other agents. Each agent randomly observes the consumption bins of agents in its neighboring node positions. The approach is similar to the replica network approach in Walden (2018). Each node in the network thus represents a whole equivalence class of agents who are ex ante identical, and there is a sufficient large number of agents

Associated with the network is the symmetric *adjacency matrix* $\mathbf{E} \in \mathbb{R}^{N \times N}$, with $\mathbf{E}_{mn} = 1$ if $(m, n) \in \mathcal{E}$, and $\mathbf{E}_{mn} = 0$ otherwise. We focus on a connected network (meaning that there is a path between any two agents). Each agent therefore has at least one neighbor.

Agent n , with prior type q_n , randomly observes $d_n M$ consumption bins of his neighbors' $d_n K$ bins. Here, we assume that the number of bins per agent, K , is sufficiently large that all agents treat these observations as effectively independent, i.e., as if the agent were sampling with replacement. An agent with more neighbors has more observations, and therefore updates his consumption behavior more aggressively than an agent with few neighbors. This is captured by the variation of d_n , as contrasted with the base model in which all agents have the same number of observations, M . Given observations of their neighbors, each agent forms posterior beliefs \hat{p}_n , that govern their own consumption. Specifically, z_n is the fraction of agent n 's observations of n 's neighbors's bins that are full. It follows that when there is visibility bias, $\tau > 1$,

$$E[z_n] = S_\tau \left(\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \hat{p}_m \right). \quad (18)$$

Definition 1 A network consumption equilibrium is a vector, $\bar{\mathbf{p}} = (\bar{p}_1, \bar{p}_2, \dots, \bar{p}_N)' \in [0, 1]^N$, such that

$$\bar{p}_n = \frac{p + d_n \xi S_\tau \left(\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m \right)}{1 + d_n \xi}, \quad n = 1, \dots, N. \quad (19)$$

This is the natural generalization of the equilibrium concept used in the preceding section.

The population equilibrium probability estimate is

$$\bar{p} = \frac{1}{N} \sum_n \bar{p}_n,$$

and the per capita consumption is $\bar{c}_0 = \bar{p}(W/2)$. The network economy is characterized by the tuple $\mathcal{T} = (\mathbf{E}, \xi, \tau, p)$, where \mathbf{E} is the adjacency matrix of the connected network, $\xi > 0$, $\tau > 1$, and $0 < p < 1$. Owing to neglect of visibility bias, each agent behaves as if the economy is actually $\mathcal{T}' = (\mathbf{E}, \xi, 1, p)$.

The following proposition characterizes the equilibrium:

Proposition 5 Consider an economy represented by $\mathcal{T} = (\mathbf{E}, \xi, \tau, p)$.

1. There exists a network equilibrium vector, $\bar{\mathbf{p}} \in [p, 1]^N$, with correct consumption, i.e., $\bar{p} = (p, p, \dots, p)'$ if and only if $\tau = 1$.

in the economy so that expectations rather than realizations can be used in the subsequent equilibrium fixed point definition, as in the base model of Section 3.2.

2. The equilibrium vector is unique if

$$\left(1 + \frac{1}{\xi D}\right) [(\tau - 1)p + 1]^2 > \tau. \quad (20)$$

Alternative sufficient conditions for uniqueness are that:

- (i) The prior probability for high consumption is sufficiently high, $p > \frac{1}{2}$, or
- (ii) Visibility bias is sufficiently weak, i.e., τ is sufficiently close to one, so that $\tau < 1 + \frac{1}{\xi D}$, or
- (iii) Visibility bias is sufficiently strong, i.e., τ is sufficiently large, so that $1 < p(\tau - 1)$.

3. As $\tau \rightarrow \infty$, the equilibrium vector converges to

$$\bar{\mathbf{p}} = \left(\frac{p + d_1 \xi}{1 + d_1 \xi}, \frac{p + d_2 \xi}{1 + d_2 \xi}, \dots, \frac{p + d_N \xi}{1 + d_N \xi} \right)'.$$

From here on, we focus on the case when $p > 1/2$, justified by our assumption that the negative wealth shock (which occurs with probability $1 - p$) is relatively rare. This implies that the equilibrium vector is unique.

For reasons of tractability, network models often focus on symmetric networks. In this context we define a network as symmetric if all agents having the same connectivity, d . (This is a milder assumption than most definitions of network symmetry in the literature.)

Recalling the definition of the probability assessment function B from equations (11)-(13), for symmetric networks, we also have:

Proposition 6 *When the social network is symmetric, equilibrium satisfies $\bar{p}_n = B(\tau, p, d\xi)$, $n = 1, \dots, N$, so that all agents have the same beliefs and consumption.*

Specifically, in a symmetric network, all agents consume $(W/2)B(\tau, p, d\xi)$. It follows that all the results in Proposition 1 generalize to symmetric networks. It is intuitively clear why equilibrium is symmetric in a symmetric network. In the model, each agent's belief updating as a function of bin observations depends directly only on how many neighbors the agent has, and in a symmetric network, all agents have the same number of neighbors. So if we propose a symmetric equilibrium, agents update symmetrically across nodes, consistent with the symmetry of the proposed equilibrium.

Moreover, the following results hold with respect to connectivity, d :

Corollary 1 *When the social network is symmetric:*

- Equilibrium consumption \bar{c}_0 is increasing in connectivity, d .

- As connectivity, $d \rightarrow \infty$, equilibrium consumption approaches $W/2$, corresponding to $\bar{p}_n = 1$ for all n .

So, overconsumption is more pronounced in more well-connected societies. Intuitively, greater connectivity creates greater observation of others, which promotes heavy upwards updating; and this effect is further intensified by the social feedback effect.

4.1 Individual consumption and centrality

The network equilibrium relation (19) suggests that an agent's equilibrium consumption increases in his connectedness, d_n , because the more connections, the more weight placed on observations of others and the less weight on the (lower) prior. This is the only mechanism through which an agent's number of connections is important in our model. This distinguishes our model from that of Jackson (2018), which assumes a direct payoff complementarity wherein the incremental payoff from engaging in the behavior is increasing in the product of the number connections an agent has and the average behavior of those connections.

In our network setting, an agent's consumption is also increasing with the consumption of network neighbors. The consumption of these neighbors, in turn, tends to be increasing in their connectedness. So an agent's consumption depends upon a potentially unlimited iteration of dependencies, where each stage tends to be increasing with the relevant agents' connectedness.

Measures of *centrality* from network theory are sometimes designed to take into account such iterated dependencies. This suggests that agents that are more central (well-connected) will overconsume more. This can only be evaluated in a network where agents differ in their connectivity. To examine such effects, we therefore now consider asymmetric networks.

We study a class of networks in which there is a core of highly connected agents surrounded by peripheral, less connected, agents who are mainly connected to the core. In a social context, we can think of the network's core as consisting of highly social agents who have many connections among themselves and to others. Such networks are thus asymmetric.²⁸

For tractability, we study core-periphery networks in which all core agents have the same number of connections to other core agents, namely $d^C > 0$, and the same number, $d^P > 0$ to peripheral agents. Each peripheral agent is connected to only one core agent. An example of a network with $d^C = 3$, $d^P = 3$ is shown in Figure 1. Note that d^P also denotes the number of peripheral agents per core agent in the economy.

²⁸Core-periphery networks arise in many different real world contexts, such as over-the-counter dealer markets.




Figure 1: Core-periphery network, with 4 (black) core agents in center and 12 (red) peripheral agents. Each core agent has $d^C = 3$ connections to other core agents, and $d^P = 3$ connections to peripheral agents.

By Definition 1, the equilibrium probability estimates of the core and peripheral agents, \bar{p}^C and \bar{p}^P , satisfy

$$\bar{p}^C = \frac{p + (d^C + d^P)\xi S_\tau \left(\left(\frac{d^C}{d^C + d^P} \right) \bar{p}^C + \left(\frac{d^P}{d^C + d^P} \right) \bar{p}^P \right)}{1 + (d^C + d^P)\xi}, \quad (21)$$

$$\bar{p}^P = \frac{p + \xi S_\tau(\bar{p}^C)}{1 + \xi}. \quad (22)$$

Also, the per capita consumption in the economy is $\bar{c}_0 = \bar{p} \left(\frac{W}{2} \right)$, where, since there are d^P peripheral agents for each central agent,

$$\bar{p} = \left(\frac{1}{d^P + 1} \right) \bar{p}^C + \left(\frac{d^P}{d^P + 1} \right) \bar{p}^P$$

is a weighted average of the agents' probability estimates. Under our assumption that $p > 1/2$, by Proposition 5, the equilibrium is unique, and network structure determines equilibrium overconsumption.

Proposition 7 *Core agents consume more than peripheral agents, $\bar{p}^C > \bar{p}^P$, and aggregate consumption is increasing in the connectivity of core agents, d^C .*

Proposition 7 indicates that when the core agents are more heavily connected, there is greater overconsumption. So if social trends, such as the rise of Facebook or Twitter,

result in core agents becoming more heavily connected, we expect overconsumption to increase.

4.2 Social Networks, The Majority Illusion, and Policy Interventions

We have seen in Subsection 3.3 that when there are smart agents, on average agents overestimate the average consumption of others. It followed that salient public disclosure of the consumption of others helps reduce overconsumption. We will now see that even without smart agents, when we make the realistic assumption that the social network is asymmetric, again there are misperceptions of the average beliefs and behavior of others. In consequence, the disclosure of public information about what others think and do again can reduce overconsumption.

We illustrate this in a simple example in which agents end up overestimating the consumption of others. We again will evaluate the policy intervention of introducing a public signal. So we now consider an extension of the network model in which all agents receive an additional unbiased public signal about per capita consumption or the average probability estimate, \bar{p} , which they incorporate into their Bayesian posteriors. This leads to the following definition of network equilibrium:

Definition 2 *A network consumption equilibrium with a public signal about aggregate consumption is a vector, $\bar{\mathbf{p}} = (\bar{p}_1, \bar{p}_2, \dots, \bar{p}_N)' \in [0, 1]^N$, such that*

$$\bar{p}_n = \frac{p + d_n \xi S_\tau \left(\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m \right) + d_n \alpha \bar{p}}{1 + d_n \xi + d_n \alpha}, \quad n = 1, \dots, N, \quad (23)$$

where $\bar{p} = \frac{1}{N} \sum_n \bar{p}_n$.

In this extension, agents' posteriors are based on three components: their priors, their observations, and the public signal. As before, the parameter $\alpha \geq 0$ determines how much weight the agents put on the public signal. When $\alpha = 0$, the extension reduces to the original network model.

As in the base model without smart agents, in the network model, the presence of a public signal about aggregate consumption does not correct overconsumption *if the network is symmetric*.

Proposition 8 *In any symmetric network, aggregate consumption is unaffected by the weight agents assign to the public signal, $d\bar{p}/d\alpha = 0$.*

This comes from the fact that agents consume based on posterior beliefs, which on average are equal to \bar{p} . It follows that on average the public signal just reconfirms these average

posterior beliefs, and therefore has no effect on aggregate consumption. This intuition is essentially the same as the intuition for the effect of disclosure in the base model.

Proposition 8 provides conditions under which publicizing aggregate consumption data does not help address overconsumption. However, this relies on the crucial assumption that networks are symmetric. A key property of symmetric networks in our model is that the average posterior belief is consistent with actual aggregate consumption. This is in general not the case in asymmetric networks. As we have seen, in our model central agents in asymmetric networks consume more, and also disproportionately influence the consumption of other agents. Therefore, the consumption of the average *observed* agent tends to be higher than the consumption of the average agent. This point is a reflection of the *majority illusion* in social networks (Lerman, Yan, and Wu 2016), wherein observers disproportionately see the characteristics of better-connected agents. (This effect is a generalization of the friendship paradox, which was discussed in the introduction.) The higher average consumption of observed agents in an asymmetric network is greater than the public signal of actual per capita consumption. So the public signal acts as a corrective to overconsumption.

We verify this conclusion in the core-periphery network in the previously studied case, where it is easy to verify that \bar{p} , and thereby also average consumption, is decreasing in the weight assigned to the public signal.

Result 1 *In a core-periphery network with $d^C = 3$, $d^P = 3$, $\xi = 1$, $\tau = 2$, and $p = 3/5$, \bar{p} , and thereby also average consumption, is decreasing in the weight assigned to the public signal, α .*

To sum up, just as in the ‘smart agents’ setting of Subsection 3.3, in an asymmetric networks setting, a public signal helps expose mistaken beliefs that agents have about the consumption of others. This makes public perceptions more accurate and reduces overconsumption. The conclusion that public disclosure can help remedy misperceptions about the beliefs and behaviors of others is in the spirit of Jackson (2018), who finds that beliefs and aggregate actions can be corrected by public disclosure when there are positive strategic complementarities in the actions of different agents. Our findings differ in two key ways.

First, our finding in the setting with smart agents is not based on network asymmetry. Even if all agents are equally well-connected, disclosure can help correct misperceptions about others. In contrast, network asymmetry (and its consequence, the friendship paradox) is the reason why disclosure has an effect in Jackson’s model.

Second, our model is not based upon strategic complementarities, which drive the effects in Jackson’s model. The driving force behind our results is naivete about visibility bias.²⁹ So

²⁹This can be relevant for empirical applications, since strategic complementarities for some consumption activities are likely to be modest or even negative. For example, congestion effects can induce negative strategic complementarity. If others consume heavily by buying cars, the resulting traffic makes it less desirable to own a car.

whereas Jackson's approach is about the interaction between naivete about the friendship paradox and preference complementarities, our approach is focuses solely on belief biases. Even in our network model, in which the friendship paradox is involved, the mechanism by which it operates to amplify overconsumption (which is present even without it) is quite different.

5 Extensions

Both to address the generality of our conclusions and to address interesting additional issues, we now consider extensions of the base model which, for tractability, make some stronger assumptions.

First, in reality people differ in age. Might observation of the old undermine the conclusion that people overconsume when young, perhaps because observation of the old people provides a reality check? Also, does the degree of tilt toward observing the young versus the old have empirical implications for how heavily people consume? We address these topics in Subsection 5.1.

Second, in reality people do not perfectly know each other's wealth, which changes the learning problem because observed consumption in many bins is an indication that the targets of observation have high wealth, not just favorable signals about the probability of no wealth disaster. This leads to empirical implications about wealth dispersion and overconsumption in Subsection 5.2.

Third, we have so far assumed a fixed riskfree interest rate. We allow for increasing supply of debt as a function of the interest rate in Subsection 5.3. We show that overconsumption is obtained in this setting too, and that interest rates are higher when visibility bias is present than when it is not.

As a matter of robustness, we also show in the appendix that results similar to those in the base model arise under some technical variations. We consider other utility functions in Appendix B. We then depart from the assumption that the maximal fraction of full consumption bins is 100%, in Appendix C. Specifically, our base model made the assumption that when the agent is maximally optimistic, and therefore consumes $W/2$ at date 1, that this is achieved when all the bins are full. Our extension allows for consumption of $W/2$ to leave some bins empty. This variation is also useful for the analysis in Sections 5.2 and 5.3.

For the variations we study in this section, we make the additional assumption that the number of prior and consumption observations, Q and M , are very large, so that agents' priors are very close to p and the number of observed full bins, z_n , is very close to $E[z]$. The fraction $\xi = \frac{M}{Q}$ is still an arbitrary positive number. We may think of this as studying the limit of a sequence of economies as $Q \rightarrow \infty$, with $M = \xi Q$ in each economy.

5.1 Age Differences: An Overlapping Generations Setting

In the base model all agents observe each other and make their savings decision at the same time, when young. In reality, young people sometimes observe the consumption of old people who are consuming from their savings. If the young are overconsuming, then as in the base model, observations of the young promote an inference of low disaster risk, which encourages the young to consume heavily. However, if the young overconsume, the old will, on average, underconsume. If a young observer sees *low* consumption of the old (this will not necessarily be the case, owing to visibility bias), the observer may infer that disasters were realized heavily, reducing consumption. This leads to an inference of high disaster risk, which discourages current consumption by young agents. This raises the question of whether overconsumption by the young can be self-confirming.

An alternative intuitive perspective also raises a doubt about whether in equilibrium there must be overconsumption. Suppose that owing to visibility bias, young observers think they see *high* consumption by old agents. Then young agents may infer that old agents had saved a lot when they were young, which would occur if they had viewed the risk of a wealth disaster as high. This inference discourages young observers from consuming heavily. Alternatively, a young observer might conclude that in the current period old agents have generally had favorable wealth realizations (not disasters). That suggests an inference of low disaster risk, which encourages young agents to consume heavily. The overall outcome is not immediately obvious.

When modelled in a straightforward way, we will see that just as in the base model, unambiguously, the young overconsume. To allow for observation of the old, we extend the base model to include an overlapping generations (OLG) structure in which there are both young and old agents at any given point in time. Specifically, the fraction $\lambda \in [0, 1]$ of the bin observations are of the young, and the remaining fraction $1 - \lambda$ of observations are of the old.

The case $\lambda = 1$ corresponds to the base model. The young might, for example, disproportionately observe each other rather than the old, owing to homophily (the tendency for people to interact with others who are similar), leading to higher λ . On the other hand, the old may act as role models for the young, leading to lower λ . In addition to bias toward observing young or old, λ reflects the fractions of the population that are in these two groups. If, for example, there is no bias in observation of young versus old, then since all agents live exactly two dates, we can think of λ as a summary statistic for population growth. The population pyramid will be such that λ is high in rapidly growing populations.

We assume that the ϵ -shock is independent across agents (though still identical in distribution), to avoid systematic variations in aggregate consumption across time. Moreover, we study stationary equilibrium in which the average estimated probability of no wealth disaster, \bar{p} , is constant over time.

Young agents observe a random sample of consumption from each cohort, i.e., λM observations are from the young generation, and $(1 - \lambda)M$ from the old. Introducing observation of the old requires a slight extension of the base model, because an old agent who is unlucky and hit with disaster potentially has a negative level of consumption, but it does not make sense to talk about a negative fraction of full consumption bins. On the opposite side, an old agent who is lucky and not hit with a disaster potentially consumes more than $W/2$, the expenditure corresponding to all bins being full. Since the main reasoning of the model is based on average levels of consumption within the population or subpopulations, these problems can be addressed by assuming a date-1 transfer of consumption from the lucky old to the unlucky old that brings the consumption level of the unlucky up to zero, and ensures that the consumption levels of the lucky old are never greater than $W/2$.³⁰

We assume that visibility bias, as previously specified in (7), is the same for observations of either the young or the old. As in the base model, observers think there is no visibility bias, and believe that the average consumption of the young is optimal, $pW/2$, where as before, p is the true probability of non-disaster. It follows that observers believe that the average consumption of the old is also $pW/2$, corresponding to the fraction p of full consumption bins by the old.³¹ The young therefore view consumption bin observations of the old as having identical information content as observations of the young. In consequence, it does not matter for the analysis whether observers can see whether any given consumption bin observation is drawn from the old or from the young; even if an observer sees the identity of the agent corresponding to an observed bin, the observer ignores this information.

Given a cohort's estimated \bar{p} when young and their associated consumption of $\bar{p} \left(\frac{W}{2}\right)$, by the law of large numbers their average consumption in the next period, when old, is $(2p - \bar{p}) \left(\frac{W}{2}\right)$, where p is the true probability. When $\bar{p} > p$, there is underconsumption by the old generation compared with the social optimum, $pW/2$, since $2p - \bar{p} < p$. Without visibility bias, by reasoning similar to that leading to (8), equilibrium average beliefs satisfy

$$\bar{p} = \frac{p + \xi[\lambda\bar{p} + (1 - \lambda)(2p - \bar{p})]}{1 + \xi}, \quad (24)$$

which has the unique solution $\bar{p} = p$. In this equilibrium, agents in both cohorts consume on average $pW/2$, observations of young and old consumption are consequently equally informative, and young agents update accordingly. Thus, without visibility bias agent

³⁰Specifically, we assume that at date 1, unlucky old agents work for lucky old agents (e.g., shopping for them or mowing their lawns), the payment thereby increasing the consumption level of the unlucky to zero, and reducing the average consumption of the lucky old accordingly. We further assume that the gains from this exchange are very small, so that the disutility of work of the unlucky old offsets their consumption benefit, and similarly the benefit to the lucky old of hiring the unlucky old is close to zero. Since these transactions leave everyone virtually indifferent, the ex ante optimization problem at date 0 is unchanged.

³¹The average consumption of old who consumed $pW/2$ when young is by the law of large numbers $p(W - pW/2) + (1 - p)(W - pW/2 - W) = pW/2$.

behavior is rational, as in the base model.

When there is visibility bias, observation of the bins of the young and the old are biased toward full bins, as reflected in the S_τ function, so the equilibrium condition (24) is replaced by

$$\bar{p} = \frac{p + \xi S_\tau[\lambda \bar{p} + (1 - \lambda)(2p - \bar{p})]}{1 + \xi}. \quad (25)$$

This is the OLG extension of the model with visibility bias. The expression reduces to (24) when $\tau = 1$.

Proposition 9 *In the OLG extension of the model with visibility bias, there is a stationary equilibrium satisfying the following properties:*

1. *The equilibrium probability estimate of the young generation satisfies $\bar{p} > p$, so the young generation overconsumes.*
2. *The equilibrium probability estimate, \bar{p} is increasing in the fraction of young agents, λ .*
3. *When $\lambda = 0$,*

$$\bar{p} = \frac{1 + \xi(\tau + 1) + p(3 + 2\xi)(\tau - 1) - \sqrt{V}}{2(1 + \xi)(\tau - 1)}, \quad \text{where} \quad (26)$$

$$V = (1 + \xi(\tau + 1) + p(3 + 2\xi)(\tau - 1))^2 - 4p(1 + \xi)(\tau - 1)(1 + 2p(\tau - 1) + 2\xi\tau). \quad (27)$$

An implication of Proposition 9 is that we expect overconsumption to be more severe in economies with rapid population growth or in which observation is more heavily tilted toward the young.

Proposition 9 shows that the young unambiguously overconsume in equilibrium, even when the young predominantly observe the old (λ is close to zero). This may seem surprising given the intuitions at the start of this subsection, which suggested that the outcome might be ambiguous. The first intuition was that if the young are on average overconsuming, then the old are on average underconsuming. This suggests that when observation is tilted toward the old, there is an inference of low p (high disaster risk), which favors underconsumption. So if there is heavy observation of the old, an equilibrium that proposes heavy overconsumption by the young is not self-confirming.

However, no matter how heavy the tilt toward observing the old, this effect only limits equilibrium overconsumption by the young—it cannot reverse it. If, out of equilibrium, the amount of overconsumption by the young were arbitrarily small, then the average consumption of the old would be almost as high as average consumption of the young. Visibility

bias in observation of others—even of the old—would then result in a high inference about the consumption of others, which non-negligibly favors overconsumption.

The alternative argument given at the start of this subsection for why there might be underconsumption was that owing to visibility bias, observers think they see the old consuming heavily. Such apparent high consumption by the old might be taken to mean that the old, when young, had adverse information about the wealth shock. Why doesn't this lead young observers to conclude that they need to save heavily rather than consume heavily?

The flaw in this argument is that observers believe that the old, when young, were on average consuming optimally, i.e., consuming $pW/2$, which implies that the average consumption of the older generation when old is the same, $pW/2$. In other words, the young think that the probability that any observed bin is full (regardless of whether it is drawn from a young or old agent) is p . So a full bin is always indicative of a *high* probability of non-disaster. So visibility bias in observations of others, old as well as young, encourages *high* consumption.

This intuition makes clear why in equilibrium, overconsumption is greater when observation is more heavily tilted toward the young (λ high). Relative to observation of the young, observation of the old acts as a partial reality-check on belief bias. Observers mistakenly think that on average consumption is equally divided between an agent's youth and old age, but owing to overconsumption, in equilibrium, *actual* average consumption is lower for old agents. So observations of the consumption bins of the young are more often full than observations of the bins of the old. So higher λ (sampling from the young) leads to more favorable inferences about p , and therefore greater overconsumption.

Our findings suggest that salient public disclosure of the consumption of the old can help address the problem of overconsumption. This intervention differs from that considered earlier, as it involves disclosing the average consumption of a subset of the population, not of the entire population. The effect of such disclosure is effectively to push the model in the direction of low λ (in which there is more observation of consumption of the old). As we have shown, lower λ decreases aggregate consumption. It is interesting that even when disclosing aggregate consumption does not help, disclosing the consumption of the right subset of the population does help—but does not fully remedy the problem. At best it only reduces overconsumption to that of the $\lambda = 0$ case.

5.2 Uncertainty about Wealth

So far, we have assumed that all agents have the same non-disaster wealth level W . We now generalize to allow for ex ante wealth dispersion in the population (even apart from ex post disaster realizations), and ignorance of the wealths of others. Intuitively, the inference an observer draws about others' signals based upon observing others' consumption is di-

luted by ignorance of the wealth of the observation target. High apparent consumption of a target could come either from the target possessing a favorable signal (indicating low risk of disaster), or from having higher ex ante wealth. Observers will therefore not, on average, revise their estimate of p upward as aggressively as they do when there is no information asymmetry about wealth. Wealth dispersion therefore reduces equilibrium overconsumption. This contrasts sharply with the Veblen wealth-signaling approach, in which it is precisely the fact that there is uncertainty about wealth that causes overconsumption to serve as a signal.

To allow for wealth dispersion, we now assume that a fraction λ of the population has non-disaster wealth level $(1 + \Delta)W$ (the wealthy group), where $\Delta > 0$, a fraction λ has non-disaster wealth $(1 - \Delta)W$ (the poor group), and the remaining fraction $1 - 2\lambda$ has non-disaster wealth W (the medium group). Henceforth we refer to “non-disaster wealth” more briefly as “wealth.” The average wealth is then still W , but the higher Δ is, the higher the economy’s wealth dispersion. We continue to assume that there is a probability $1 - p$ that any given individual (rich, poor, or medium) experiences a disaster (negative value of ϵ) that entirely wipes out an individual’s wealth at date 1.

To explore this setting, we now need to allow for the possibility that someone with maximally optimistic beliefs, $\hat{p} = 1$, still has some empty consumption bins. So far, we have assumed that when the agent consumes half of potential wealth at date 0, $c_0 = W/2$, that all consumption bins are full. We now instead assume that in this circumstance, only a fraction $0 < f \leq 1$ of the bins are full. To prevent the wealthy group from consuming so heavily that more than 100% of the consumption bins are full, we assume that $(1 + \Delta)f < 1$. In addition, we impose the technical condition that

$$\Delta \leq \frac{1}{1 + \frac{2}{\tau-1}}. \quad (28)$$

For large τ , this implies $\Delta < 1$. Since poor agents have wealth $(1 - \Delta)W$, for large τ this imposes the very weak restriction that even poor agents do not have wealth very close to zero.

Agents know the economy’s wealth distribution (λ and Δ), and for simplicity we assume that each agent’s consumption bin observations come from one or more agents of the same wealth type. Based on these observations, an agent forms posterior beliefs about p , taking into account that an observation of high consumption could reflect high wealth, not just optimistic beliefs, on the part of the target of observation.

The following proposition confirms the intuition that wealth dispersion reduces overconsumption:

Proposition 10 *Under the above assumptions, the equilibrium probability estimate \bar{p} is decreasing in wealth dispersion, $\partial\bar{p}/\partial\Delta < 0$, as is the overconsumption factor.*

Proposition 10 predicts that savings rates increase with wealth dispersion. This is the opposite of what is expected based upon Veblen wealth-signaling considerations.

Specifically, in the Veblen approach to overconsumption, people consume more in order to signal the level of wealth to others (Bagwell and Bernheim 1996; Corneo and Jeanne 1997), so there is no motive to overconsume when wealth dispersion is zero. A comparison of the cases of zero versus positive wealth dispersion indicates an average tendency for greater wealth dispersion to induce greater overconsumption, though not necessarily monotonically.³² This intuition is reflected, for example, in a comparative statics of Charles, Hurst, and Roussanov (2009) in which a parameter shift that increases wealth dispersion by reducing the lower support of wealth results in greater consumption signaling.³³ Furthermore, and also in contrast with our approach, in at least one version of the keeping-up-with-the-Joneses approach, income dispersion encourages the non-wealthy to consume more in emulation of the wealthy (Bertrand and Morse 2016).³⁴

Using survey evidence from Chinese urban households, Jin, Li, and Wu (2011) find that greater income inequality is associated with lower consumption and with greater investment in education, where income inequality is measured within age groups by province. Similarly, using high geographical resolution 2001-12 data, Coibion et al. (2014) provide strong evidence that low-income households in high-inequality U.S. locations accumulated less debt (relative to income) than their counterparts in lower-inequality locations. These findings are consistent with Proposition 10, in contrast with the idea that greater information asymmetry about wealth increases wealth-signaling via consumption, or with the intuitive idea that low income individuals borrow and consume more in order to try to keep up with high income households.

Several studies report that wealth dispersion has increased in the United States since the 1980s (e.g. Card and DiNardo (2002), Piketty and Saez (2003), Lemieux (2006)). Given an increase in wealth dispersion, all else equal, Proposition 10 counterfactually implies a rising savings rate. However, the time series shift in *information asymmetry* about others' wealth may have been downward rather than upward, potentially reversing the implication. Our result that wealth dispersion reduces overconsumption derives from asymmetric information—the unobservability of others' wealths—rather than dispersion per se. The

³²Consistent with this idea, Charles, Hurst, and Roussanov (2009) find empirically that greater dispersion of reference group income is associated with higher visible spending for minorities, consistent with our theory. On the other hand, they report that higher dispersion in reference group income significantly lowers White visible spending.

³³A key parameter in wealth signaling models is the lower support of wealth, which acts as a starting point for the signaling schedule. For any class of bounded wealth distributions that are symmetric with given mean, and such that higher dispersion is associated with more extreme wealth outcomes, it is equivalent to express a result in terms of wealth dispersion or lower support.

³⁴Concern for relative consumption, as in 'keeping up with the Joneses' preferences can also induce a fear of falling behind which raises precautionary savings (Harbaugh 1996). Other than Bertrand and Morse (2016), we are not aware of any results in the keeping-up-with-Joneses approach relating overconsumption to wealth dispersion (holding constant the average level of wealth).

rise of the internet has likely made observation of others' wealths or incomes easier than in the past in some countries through search of government or other archives. To the extent that this is true, this effect reinforces the other time series shifts we've described, implying a shift over time toward greater overconsumption.

Furthermore all else was not equal in this time series shift. From the standpoint of our model, a more fundamental effect (which holds even in our base model) comes from the dramatic transformation of electronic communications and social networks. As discussed in Section 3.2, this has increased the visibility of the consumption activities of others (both absolutely, and relative to non-consumption), which implies greater overconsumption in our model. Also as discussed earlier, the trend toward rising population density, and associated increase in social observation intensity, further reinforces the implication of stronger overconsumption in the base model.

As discussed in the introduction, agents who disagree may overconsume owing to the expectation of profiting by trading against others Heyerdahl-Larsen and Walden (2017). Although our model allows for difference in beliefs, our overconsumption finding is not driven by speculative trading. The speculative trading theory also shares with our model the prediction that greater wealth dispersion reduces overconsumption. However, as discussed in the introduction, the models are empirically distinguishable in various other ways.

5.3 The Equilibrium Interest Rate

In the base model, the riskfree rate is exogenously set to zero. This corresponds to having storable consumption or, equivalently, to having riskfree bonds in perfectly elastic supply offered at a zero interest rate. We now modify the model to allow for endogenous determination of the interest rate.

When the interest rate can vary, potentially a high interest rate could imply negative date 0 consumption. That does not correspond well with the idea that at worst all consumption bins are empty. We therefore adjust the model to preclude this possibility.

As in the base model, agent utility is defined by (1), where we focus on the case $\alpha = 1/2$. Given a one-period interest rate of r , an agent's budget constraint is now

$$c_1 = (1 + r)(W - c_0) - \epsilon, \tag{29}$$

For tractability, in this section we assume a less severe possible wealth disaster than in the base model to preclude negative date-0 consumption over a range of potential interest rates. We therefore assume that

$$\epsilon = \begin{cases} 0 & \text{with probability } p \\ \frac{W}{2} & \text{with probability } 1 - p, \end{cases} \tag{30}$$

and without loss of generality we focus on the case $W = 1$.

Solving for the optimum of an agent whose probability estimate for a high outcome is \hat{p} yields consumption

$$\frac{1 - r + 2r^2 + (1 + r)\hat{p}}{4 + 4r + 2r^2} \stackrel{\text{def}}{=} g + f\hat{p}, \quad (31)$$

with g and f defined in the obvious way. With average agent probability estimate of \bar{p} , aggregate per capita consumption is then

$$\bar{c}_0 = g + f\bar{p}. \quad (32)$$

We assume that there is an equally large set of agents who are not subject to ϵ risk, i.e., for whom $c_1 = (1+r)(1-c_0)$. Since these outsiders face no disaster risk, their consumption does not depend on their inferences about p . We can think of them as outsiders such as institutional investors or foreign lenders that supply capital to the individual investors that our analysis focuses upon. Their role in the model is to increase trading opportunities in our exchange economy, so that optimistic beliefs of the individual investors can increase equilibrium per capita consumption, rather than just the interest rate.³⁵

Institutional investors are willing to lend to (or borrow from) the agents that are the main focus of our analysis. Since they have no disaster risk, the optimal time-0 consumption of institutions, given r , is

$$\bar{c}_0^I = \frac{1 + r^2}{2 + 2r + r^2}. \quad (33)$$

We assume free disposal of the consumption good, so the equilibrium interest rate satisfies $r \geq -1$. We focus on the region of interest rates in which the institutional investors' lending increases in the interest rate, and therefore require that $r \leq 1/2$.³⁶

The aggregate endowed quantity of the time-0 consumption good per individual investor, \bar{c}^e is fixed. Specifically, we assume that the total endowment is such that in an equilibrium with unbiased beliefs (i.e., $\bar{p} = p$), the market clears at interest rate $r = 0$. The market clearing condition is

$$\bar{c}^e = \bar{c}_0^I + \bar{c}_0, \quad (34)$$

where the terms on the right are functions of r . By (32,33), and by our assumption that

³⁵The consumption of outsiders are excluded from our measure of aggregate consumption; their sole role is to supply capital as an increasing function of the interest rate. Including outsiders in the model is a simplified way of reflecting the idea that in general when the current consumption good is scarce, more can be generated via an aggregate production function for transformation between current and future consumption.

³⁶For $r > 1/2$, it is easy to verify that lending decreases in the interest rate. This comes from the standard result in intertemporal choice that an increase in the interest rate has both a substitution effect (which encourages lending) and a wealth effect (which can discourage lending). To illustrate basic insights simply, we focus on the case in which the substitution effect, which is highly intuitive, dominates.

when $r = 0$ a market with unbiased agents would clear, we have

$$\bar{c}^e = \frac{1}{2} + \frac{1+p}{4}, \quad (35)$$

so (34) becomes

$$\bar{c}_0^I + \bar{c}_0 = \frac{1}{2} + \frac{1+p}{4}. \quad (36)$$

The left hand side is the population per capita endowed consumption.

Regardless of the level of visibility bias, market clearing implies that r adjust so that aggregate demand is unchanged, i.e., is equal to the endowment, so in equilibrium

$$\bar{p} = \frac{(8 - 5r)r + p(2 + 2r + r^2)}{2(1 + r)}. \quad (37)$$

It is easy to verify that \bar{p} is strictly increasing in r in the relevant region of r .

Along the lines of the arguments in the base model, given r and \bar{p} , the fraction of bins that are full is given by (32). Since agents suffer from visibility bias, the fraction of bins *observed* to be full is $S_\tau(g + f\bar{p})$. By Bayes' rule, the agents then arrive at the posterior probability estimate

$$\hat{p} = R(p, S_\tau(g + f\bar{p}), \xi, f, g), \quad (38)$$

given that the fraction of bins that agents observe are full is z , where the function R is defined in equation (52) in the appendix (see the proof of Proposition 11). An equilibrium is then an outcome in which markets clear, so that (32) holds, and agents' posterior beliefs are in line with their biased observations, $\bar{p} = R(p, S_\tau(g + f\bar{p}), \xi, f, g)$. For tractability, we focus on the case when $\xi = 1$ (so that agents put comparable weight on prior information and on social observations).

We show the existence of equilibrium with the following properties:

Proposition 11 *Under the above assumptions, the equilibrium probability estimate, over-consumption, and the interest rate, are all increasing in visibility bias, τ .*

6 Concluding Remarks

We examine how bias in social learning endogenously shapes how people trade off current versus future consumption. In our model, people observe the consumption activities of others and use this to update beliefs about whether there is a high or low need to save for the future. Consumption is more salient than non-consumption, resulting in greater observation and cognitive encoding of others' consumption activities. This visibility bias makes episodes of high consumption by others more salient and easier to retrieve from memory than episodes of low consumption. So owing to neglect of selection bias (or a

well-known manifestation of it, the availability heuristic), people infer that low saving is a good idea. This effect is self-reinforcing at the social level, resulting in overconsumption and high interest rates. With many opportunities to observe others, this feedback effect can be intense. The effects in the model can also bring about biased assessments of the savings rates of others, wherein people think that others are consuming even more heavily than they really are. In consequence, a distinctive implication of our approach is that accurate disclosure of the beliefs or average consumption of others can, in some cases, help remedy overconsumption. Furthermore, accurate disclosures that increase the salience of saving behavior can have an especially strong effect.

The visibility bias approach offers a simple explanation for one of the most important puzzles in household finance: the dramatic drop in personal saving rates in the U.S. and many other OECD countries over the last 30 years. In the model, greater observability of the consumption of others intensifies the effects of visibility bias, and therefore increases overconsumption. We argue that over the last thirty years the decline in costs of long-distance telephony, the rise of cell phones, cable television and urbanization, and subsequently the rise of the internet, dramatically increased the extent to which people observe possible personal consumption activities of others by television enactment, phone, email, blogging, and social networking. Specifically, this communication is biased toward making the decision to engage rather than not engage in such activities more salient to others, because travel, dining out, or buying a car tend to be relatively noteworthy to report upon.

In contrast with the present bias (hyperbolic discounting) theory of overconsumption, the effects here are induced by social observation and interaction. Our approach can therefore be distinguished from present bias using proxies for sociability, individual network position, and network connectedness (such as urban versus rural). Our approach also differs in offering predictions about how population-level characteristics such as wealth variance and age distribution affect consumption.

Also, our visibility bias approach is not based on a link between an agent's participation in speculative markets and his overconsumption, in contrast to overconsumption associated with heterogeneous beliefs as in Heyerdahl-Larsen and Walden (2017).

The effect of wealth dispersion in our model contrasts with the implications of the Veblen and keeping-up-with-the-Joneses approaches. In at least some versions of the latter approach, wealth dispersion encourages the poor to consume more in emulation of the wealthy. The Veblen wealth-signaling approach broadly implies that a comparative statics shift from certainty to information asymmetry about others' wealths (i.e., a rise in wealth dispersion) implies greater overconsumption. In our setting greater information asymmetry dilutes the inference from high observed consumption that others have favorable information about the risk of a wealth disaster. In consequence, equilibrium consumption is *lower*, the opposite prediction. The visibility bias approach also helps explain high variation in savings

rates across countries and ethnic groups, because even modest differences in beliefs can be amplified through the social multiplier effect.

In contrast with the signaling, preference-based, and speculative disagreement approaches, our social learning bias approach implies that a relatively simple policy intervention can potentially increase saving. The relevant intervention is to provide—in highly salient form—accurate information about how much peers save, or their attitudes toward consumption. Indeed, there is evidence discussed earlier supporting this implication in specialized settings, such as the decisions of college students of how much to drink.

The model in this paper is static. An interesting extension would be to consider a dynamic setting in which agents consume over time and update in response to common shocks. The feedback/multiplier effect from social learning may then potentially lead to substantial cyclical shifts in overconsumption. Such fluctuations may help explain consumption booms and business cycles. This might potentially provide an interesting contrast to Keynesian ideas about business cycles deriving from resource underutilization and underconsumption.

Appendices

A Proofs

Proof of Proposition 1: Part 1 follows from noting that (11,12) implies that $\bar{p} = p$ if and only if $-4(1-p)p(\tau-1)^2\xi(1+\xi) = 0$, which holds if and only if $\tau = 1$. Now that $\frac{\partial \bar{p}}{\partial \tau} > 0$ can be seen by substituting $x = \frac{1}{\tau-1}$, noting that x is decreasing in τ , and taking the derivative w.r.t. x , leading to $\frac{\partial \bar{p}}{\partial x} = \left(x + p - \xi + 2p\xi - \sqrt{(p-x+\xi)^2 + 4px(1+\xi)} \right) r(x)$, where $r(x) > 0$. It then follows from the fact that $(x+p-\xi+2p\xi)^2 - ((p-x+\xi)^2 + 4px(1+\xi)) = -4(1-p)p\xi(1+\xi) < 0$, that $\frac{\partial \bar{p}}{\partial x} < 0$, and thus $\frac{\partial \bar{p}}{\partial \tau} > 0$. Part 2 and the claim in 3 that \bar{p} approaches 1 as ξ becomes large follow immediately by taking the limit of (11,12) as τ and ξ become large. To prove the other claims in Part 3, note that \bar{p} can be written as

$$\bar{p} = \frac{m + \sqrt{m^2 + 4(k+m)p}}{2(k+m)} \stackrel{\text{def}}{=} V(m),$$

where $m = (p+\xi)(\tau-1)-1 > -1$, and $k = (1-p)(\tau-1)+1 > 1$. Since $\frac{\partial m}{\partial \xi} > 0$, it is therefore sufficient to show that $V'(m) > 0$ when $m > -1$. By calculating $V'(m)$, it follows that $-2mp+k(m-2p+\sqrt{m^2+4p(k+m)}) > 0$ is necessary and sufficient for $V'(m) > 0$ to hold. For $m = 0$, the expression evaluates to $V'(0) = k(-2p+2\sqrt{kp}) > 0$. Moreover, the solution to $V'(m) = -2mp+k(m-2p+\sqrt{m^2+4kp+4mp}) = 0$ is $m_{+/-} = -k < -1$. Thus, since V' is a continuous function of m , $V'(m) > 0$ for all $m \geq -1$. We also verify that the equilibrium probability estimate and aggregate consumption are increasing in the true probability of the high state, p , by calculating $\frac{\partial \bar{p}}{\partial p} = \frac{1}{2(1+\xi)} + \frac{1}{2\sqrt{V}} \left(2 + \frac{1}{1+\xi}((\tau-1)(p+\xi)-1) \right)$, which is obviously positive for $p \in [0, 1]$, since $2 + \frac{1}{1+\xi}((\tau-1)(p+\xi)-1) \geq 2 + \frac{1}{1+\xi}((\tau-1)\xi-1) = 2 + \frac{\tau\xi}{1+\xi} - 1 > 0$ for such p . ■

Proof of Proposition 3: Equilibrium with the public noisy signal is defined via the relation:

$$\bar{p}^V = \frac{p + \xi S_\tau(\bar{p}) + \alpha \bar{p}}{1 + \xi + \alpha}, \quad \bar{p} = \phi p + (1 - \phi) \bar{p}^V,$$

which is equivalent to

$$\bar{p}^V = \frac{p + \hat{\xi} S_\tau(\bar{p})}{1 + \hat{\xi}}, \quad \hat{\xi} = 1 + \alpha \phi,$$

which we in turn rewrite as

$$\hat{\xi} = \frac{\bar{p}^V - p}{S_\tau(\bar{p}) - \bar{p}^V}.$$

It follows that

$$\frac{d\hat{\xi}}{d\bar{p}^V} = \frac{1}{S_\tau(\bar{p}) - \bar{p}^V} + \frac{1 - (1 - \phi) S'_\tau(\bar{p})}{(S_\tau(\bar{p}) - \bar{p}^V)^2} > 0, \quad (39)$$

since $S_\tau(\bar{p}) > \bar{p}^V$ and $S'_\tau(\bar{p}) < 1$, and thus that

$$\frac{d\bar{p}^V}{d\hat{\xi}} > 0.$$

Since $\frac{d\hat{\xi}}{d\alpha} = -\frac{\phi}{(1+\alpha\phi)^2} < 0$, the result follows. ■

Proof of Proposition 4: Denote by \bar{p} the equilibrium probability estimate in (13) and by \bar{p}' the equilibrium in (17). It follows from these equations that

$$\bar{p}' - \bar{p} = \frac{\xi}{1+\xi}(S_\tau(\bar{p}') - S_\tau(\bar{p})) + \frac{\alpha}{1+\xi}(1 - \bar{p}' - S_\tau(1 - \bar{p}')). \quad (40)$$

Define the function

$$R(x) = x - \bar{p} - \left(\frac{\xi}{1+\xi}(S_\tau(x) - S_\tau(\bar{p})) + \frac{\alpha}{1+\xi}(1 - x - S_\tau(1 - x)) \right), \quad (41)$$

and note that $R(\bar{p}') = 0$, and that R is a continuous function of $x \in [0, 1]$. From (14), it follows that $\bar{p} > \frac{\xi}{1+\xi}S_\tau(\bar{p})$, and therefore, since $S_\tau(0) = 0$, $S_\tau(1) = 1$, that $R(0) < 0$. Moreover, since $S_\tau(y) > y$, for $y \in (0, 1)$, it follows that $R(\bar{p}) = \frac{\alpha}{1+\xi}(S_\tau(1 - \bar{p}) - (1 - \bar{p})) > 0$. Thus, by the intermediate value theorem, it follows that the point, \bar{p}' where $R(\bar{p}') = 0$, satisfies $0 < \bar{p}' < \bar{p}$. ■

Proof of Proposition 5: 1. Define the mapping $F : \mathbb{R}_+^N \rightarrow \mathbb{R}_+^N$ by

$$(F(w))_n = \frac{p + d_n \xi S\left(\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} w_m\right)}{1 + d_n \xi}, \quad n = 1, \dots, N. \quad (42)$$

An equilibrium is then a fixed point to this mapping, $\bar{\mathbf{p}} = F(\bar{\mathbf{p}})$. It is easy to see that F is nondecreasing in each of its arguments: $w^2 \geq w^1 \rightarrow F(w^2) \geq F(w^1)$, that $F(p, p, \dots, p) = \left(\frac{p+d_1\xi S(p)}{1+d_1\xi}, \dots, \frac{p+d_N\xi S(p)}{1+d_N\xi}\right)' \stackrel{\text{def}}{=} z \geq (p, p, \dots, p)'$, and that $F(1, 1, \dots, 1) = \left(\frac{p+d_1\xi}{1+d_1\xi}, \frac{p+d_2\xi}{1+d_2\xi}, \dots, \frac{p+d_N\xi}{1+d_N\xi}\right)' \stackrel{\text{def}}{=} y \leq (1, 1, \dots, 1)'$. It follows that F maps the convex compact set $S \stackrel{\text{def}}{=} [z, y]^N$ into itself. Since F is a continuous mapping, Brouwer's theorem implies that F has a fixed point, i.e., that there exists an equilibrium in S , and since $S \subset [p, 1]^N$, the existence result follows.

For $\tau = 1$, (19) reduces to the linear algebraic equation

$$\bar{p} = (I + \xi \text{diag}(d_1, \dots, d_N))^{-1}(p\mathbf{1} + \xi E\bar{p}),$$

or equivalently,

$$(I + \xi \text{diag}(d_1, \dots, d_N) - \xi E)\bar{p} = p\mathbf{1}.$$

Here, the $\mathbf{1}$ is a vector of ones, $\mathbf{1} = (1, \dots, 1)' \in \mathbb{R}^N$. It is easy to verify that $\bar{p} = p\mathbf{1}$ is a solution. Since the matrix $A \stackrel{\text{def}}{=} (I + \xi \text{diag}(d_1, \dots, d_N) - \xi E)$ is diagonally dominant ($A_{nn} - \sum_{m \neq n} |A_{mn}| = 1 > 0$, $n = 1, \dots, N$), it is invertible, so this solution is unique when $\tau = 1$. When $\tau > 1$, it also immediately follows that $p\mathbf{1}$ is not a solution, since $z \gg p\mathbf{1}$ in

this case, and the solution must lie in $S = [z, y]^N$.

2. The case $\tau = 1$ is already covered in part 1 of the proof, so w.l.o.g. assume that $\tau > 1$. Consider the function $G : (0, 1)^N \rightarrow \mathbb{R}$, defined by

$$G(x) = \sum_{n=1}^N d_n R_n \log(x_n) + \frac{1}{2} \sum_{n,m=1}^N x_n E_{nm} x_m - \sum_{n=1}^N g_n x_n, \quad (43)$$

$$R_n = \frac{d_n \xi \tau}{(1 + d_n \xi)(\tau - 1)^2}, \quad (44)$$

$$g_n = \frac{d_n}{\tau - 1} + \sum_{m \in \mathcal{D}_n} f_m, \quad (45)$$

$$f_n = \frac{p}{1 + d_n \xi} + \frac{d_n \xi \tau}{(1 + d_n \xi)(\tau - 1)}. \quad (46)$$

The gradient of G is $\nabla G(x) \in \mathbb{R}^N = a + Ex - g$, where $a = (a_1, \dots, a_N)'$, $a_n = \frac{d_n R_n}{x_n}$, $(Ex)_n = \sum_{m \in \mathcal{D}_n} x_m$, and $g = (g_1, \dots, g_N)'$. A stationary point of G satisfies $\nabla G(x) = 0$. Defining the bijection $\bar{p} \leftrightarrow x$, via $\bar{p}_n \stackrel{\text{def}}{=} f_n - x_n$, it follows that at such a stationary point $\frac{d_n R_n}{f_n - \bar{p}_n} = g_n - \sum_{m \in \mathcal{D}_n} f_m + \sum_{m \in \mathcal{D}_n} \bar{p}_m = \frac{d_n}{\tau - 1} + \sum_{m \in \mathcal{D}_n} \bar{p}_m$, or equivalently

$$\begin{aligned} \bar{p}_n &= f_n - \frac{d_n R_n}{\frac{d_n}{\tau - 1} + \sum_{m \in \mathcal{D}_n} \bar{p}_m} \\ &= \frac{p}{1 + d_n \xi} + \frac{d_n \xi \tau}{(1 + d_n \xi)(\tau - 1)} - \frac{R_n}{\frac{1}{\tau - 1} + \frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m} \\ &= \frac{p}{1 + d_n \xi} + \frac{d_n \xi}{1 + d_n \xi} \left(\frac{1}{1 - \frac{1}{\tau}} - \frac{1}{(1 - \frac{1}{\tau})(\tau - 1) \left(\frac{1}{\tau - 1} + \frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m \right)} \right) \\ &= \frac{p}{1 + d_n \xi} + \frac{d_n \xi}{1 + d_n \xi} \left(\frac{(\tau - 1) \left(\frac{1}{\tau - 1} + \frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m \right) - 1}{(1 - \frac{1}{\tau})(\tau - 1) \left(\frac{1}{\tau - 1} + \frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m \right)} \right) \\ &= \frac{p}{1 + d_n \xi} + \frac{d_n \xi}{1 + d_n \xi} \left(\frac{\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m}{(1 - \frac{1}{\tau}) \frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m + \frac{1}{\tau}} \right) \\ &= (F(\bar{p}))_n. \end{aligned}$$

Thus, every equilibrium point, \bar{p} , is equivalent to a stationary point of G , x , under the mapping $\bar{p} \leftrightarrow x$, with

$$x_n = \frac{p}{1 + d_n \xi} + \frac{d_n \xi \tau}{(1 + d_n \xi)(\tau - 1)} - \bar{p}_n.$$

It also follows that the set S under the $\bar{p} \leftrightarrow x$ mapping corresponds to the set

$$\{x\} \in U \stackrel{\text{def}}{=} \left[0, \frac{d_1 \xi}{1 + d_1 \xi} \left(\frac{\tau}{\tau - 1} - \frac{\tau p}{(\tau - 1)p + 1} \right) \right] \times \left[0, \frac{d_2 \xi}{1 + d_2 \xi} \left(\frac{\tau}{\tau - 1} - \frac{\tau p}{(\tau - 1)p + 1} \right) \right] \\ \times \cdots \times \left[0, \frac{d_N \xi}{1 + d_N \xi} \left(\frac{\tau}{\tau - 1} - \frac{\tau p}{(\tau - 1)p + 1} \right) \right].$$

Thus, if there is a unique stationary point of G in U , then the corresponding equilibrium vector \bar{p} is unique in $[p, 1]^N$.

It is easy to check that the Hessian of G , $H_G(x) \in \mathbb{R}^{N \times N}$, has elements

$$[H_G(x)]_{nm} = \begin{cases} -\frac{d_n R_n}{x_n^2}, & n = m, \\ E_{nm} & n \neq m. \end{cases}$$

It follows that for $x_n \in \left[0, \frac{d_n \xi}{1 + d_n \xi} \left(\frac{\tau}{\tau - 1} - \frac{\tau p}{(\tau - 1)p + 1} \right) \right]$,

$$[H_G(x)]_{nn} \leq -\frac{d_n R_n}{\left(\frac{d_n \xi}{1 + d_n \xi} \left(\frac{\tau}{\tau - 1} - \frac{\tau p}{(\tau - 1)p + 1} \right) \right)^2} \\ = -\frac{d_n \frac{d_n \xi \tau}{(1 + d_n \xi)(\tau - 1)^2}}{\left(\frac{d_n \xi}{1 + d_n \xi} \left(\frac{\tau}{\tau - 1} \right) \left(\frac{1}{(\tau - 1)p + 1} \right) \right)^2} \\ = -d_n \left(\frac{1 + d_n \xi}{d_n \xi} \right) \left(\frac{1}{\tau} \right) ((\tau - 1)p + 1)^2.$$

Since $\sum_m E_{nm} = d_n$ for all n , it follows that if

$$\left(\frac{1 + D\xi}{D\xi} \right) \left(\frac{1}{\tau} \right) ((\tau - 1)p + 1)^2 > 1, \quad (47)$$

then H_G is diagonally dominant with negative diagonal elements, in the whole of U and thus the Hessian is negative definite in this region. Standard theory of optimization then in turn implies that a stationary point of G is unique in U , and thus that \bar{p} is unique in S , and therefore also in $[p, 1]^N$. The condition (47) is obviously equivalent to (20).

3. The result follows immediately from the fact that when $\tau \rightarrow \infty$, both z_n and y_n , as defined in part 1 of the proposition, converge to $\frac{p + d_n \xi}{1 + d_n \xi}$ for all n . ■

Proof of Proposition 6: Conjecture an equilibrium in which $\bar{p}_1 = \bar{p}_2 = \cdots = \bar{p}_N = w$. Since $\frac{1}{d_n} \sum_{m \in \mathcal{D}_n} \bar{p}_m = w$ for all n , (19) reduces to the condition $\bar{p}_n = \frac{p + d_n \xi S_\tau(\bar{p}_n)}{1 + d_n \xi}$, which has solution $\bar{p}_n = B(\tau, p, d\xi)$, $n = 1, \dots, N$. ■

Proof of Proposition 7: As evident from Proposition 5, equilibrium consumption satisfies $\bar{p}^P > p$, $\bar{p}^C > p$. Assume that $\bar{p}^P \geq \bar{p}^C$. It follows that $\zeta = \frac{d^C}{d^C + d^P} \bar{p}^C + \frac{d^P}{d^C + d^P} \bar{p}^P$ satisfies

$\bar{p}^C \leq \zeta \leq \bar{p}^P$, and consequently that $S_\tau(\zeta) \geq S_\tau(\bar{p}^C)$. Consequently (since $d^C + d^P > 1$),

$$\frac{p + (d^C + d^P)\xi S_\tau(\zeta)}{1 + (d^C + d^P)\xi} > \frac{p + \xi S_\tau(\bar{p}^C)}{1 + \xi},$$

and thus $\bar{p}^C > \bar{p}^P$, leading to a contradiction. It follows that $\bar{p}^C > \bar{p}^P$.

For the second part of the theorem, note that it follows from (21, 22) that in equilibrium $F(\bar{p}^C, d^C) = 0$, where

$$\begin{aligned} F(\bar{p}^C, d^C) &= p + (d^C + d^P)\xi S_\tau \left(\frac{d^C}{d^C + d^P} \bar{p}^C + \frac{d^P}{d^C + d^P} \left(\frac{p + \xi S_\tau(\bar{p}^C)}{1 + \xi} \right) \right) - \bar{p}^C(1 + (d^C + d^P)\xi) \\ &= p + d\xi S_\tau(\zeta) - \bar{p}^C(1 + d\xi), \end{aligned}$$

where $d = d^C + d^P$. Now,

$$\begin{aligned} \frac{\partial F}{\partial \bar{p}^C} &= d\xi S'_\tau(\zeta) \left(\frac{d^C}{d^C + d^P} + \frac{d^P}{d^C + d^P} \frac{\xi}{1 + \xi} S'_\tau(\bar{p}^C) \right) - (1 + d\xi) \\ &\leq d\xi - (1 + d\xi) \\ &< 0, \end{aligned}$$

where the inequality follows from the fact that $\bar{p}^C > \frac{1}{2}$, $\zeta > \frac{1}{2}$, and $S'_\tau(x) < 1$ when $x > \frac{1}{2}$.

Moreover,

$$\begin{aligned} \frac{\partial F}{\partial d^C} &= \xi S_\tau(\zeta) + d\xi S'_\tau(\zeta) \frac{d\zeta}{dd^C} - \bar{p}^C \xi \\ &> 0 \end{aligned}$$

which follows from the fact that $S_\tau(\zeta) \geq \bar{p}^C$ (which in turn follows from (21)), and that $\frac{\partial \zeta}{\partial d^C} = \frac{1}{d^2}(\bar{p}^C - \bar{p}^D) > 0$. Therefore, by the inverse function theorem, $\frac{\partial \bar{p}^C}{\partial d^C} = -\frac{\frac{\partial F}{\partial d^C}}{\frac{\partial F}{\partial \bar{p}^C}} > 0$. Finally, from (22) it follows that \bar{p}^D is strictly increasing in \bar{p}^C . Both \bar{p}^C and \bar{p}^D are therefore increasing in d^C , as is then aggregate consumption. ■

Proof of Proposition 8: For a symmetric network, equation (23) reduces to

$$\bar{p} = \frac{p + d\xi S_\tau(\bar{p}) + d\alpha \bar{p}}{1 + d\xi + d\alpha}. \quad (48)$$

since $\bar{p}_n = \bar{p}$ and $d_n = d$ for all agents. It follows immediately that any solution to (48) equivalently satisfies

$$\bar{p} = \frac{p + d\xi S_\tau(\bar{p})}{1 + d\xi}, \quad (49)$$

i.e., is also an equilibrium in the economy without public signal. ■

Proof of Proposition 9: It is easy to verify that the solution to the equilibrium condition (25) is

$$\bar{p} = \frac{-1}{2(2\lambda - 1)(1 + \xi)(\tau - 1)} \left(1 - 3p + 4\lambda p + \xi - 2p\xi + 2\lambda p\xi + 3p\tau - 4\lambda p\tau + \xi\tau + 2p\xi\tau - 2\lambda p\xi\tau - \sqrt{V} \right), \quad (50)$$

where

$$V = -4(2\lambda - 1)p(1 + \xi)(\tau - 1) \left(-1 + 2(\lambda - 1)p(\tau - 1) + 2(\lambda - 1)\xi\tau \right) + \left(1 - p(-3 - 2\xi + 2\lambda(2 + \xi)(\tau - 1) + \xi(1 + \tau - 2\lambda\tau)) \right)^2.$$

It is also easy to verify that the solution reduces to (26,27) when $\lambda = 0$, and to (11,12) when $\lambda = 1$. Moreover, it is easy to verify that $\bar{p}^{\lambda=0} > p$, and that for any $\lambda \in [0, 1]$, $\bar{p}^\lambda = p \Rightarrow p \in \{0, 1\}$. It also follows immediately that \bar{p} is a continuous function of λ , except possibly at $\lambda = 1/2$.

We next define $x = 2\lambda - 1 \in [-1, 1]$, and rewrite (50) as

$$\begin{aligned} \bar{p} &= a + \frac{b(\sqrt{1 - cx + dx^2} - 1)}{2x(\tau - 1)}, & \text{where} \\ a &= \frac{p(2 + \xi)(\tau - 1) + \xi\tau}{2(1 + \xi)(\tau - 1)}, \\ b &= 1 + p(\tau - 1) \\ c &= \frac{2\xi(-p(\tau - 1)^2 + p^2(\tau - 1)^2 + \tau)}{(1 + \xi)(1 + p(\tau - 1))^2}, \\ d &= \frac{\xi^2(p(1 - \tau) + \tau)^2}{(1 + \xi)^2(1 + p(\tau - 1))^2}. \end{aligned}$$

A Taylor expansion of $\sqrt{1 - cx + dx^2} - 1$ around $x = 0$ i.e., $\lambda = 1/2$, yields $\sqrt{1 - cx + dx^2} - 1 = \frac{1}{2}cx + O(x^2)$, and thus \bar{p} is a continuous function of λ at $\lambda = 1/2$ too. Thus, since $\bar{p}^{\lambda=1} > p$, \bar{p}^λ depends continuously on λ , and $\bar{p}^\lambda \neq p$ for $\lambda \in [0, 1]$, it follows that $\bar{p}^\lambda > p$ for all $\lambda \in [0, 1]$. We have shown $\bar{p} > p$, i.e., (1), and (3).

To show (2), we note that

$$\frac{d\bar{p}}{dx} = \frac{b}{4(\tau - 1)} \times \frac{\sqrt{1 - cx + dx^2} + \frac{c}{2}x - 1}{\sqrt{1 - cx + dx^2}},$$

so $\sqrt{1 - cx + dx^2} > 1 - \frac{c}{2}x$ is necessary and sufficient for $\frac{d\bar{p}}{dx} > 0$. This implies the following sufficient condition:

$$1 - cx + dx^2 > \left(1 - \frac{c}{2}x \right)^2 = 1 - cx + \frac{c^2}{4}x^2,$$

or equivalently,

$$4d^2 - c^2 > 0.$$

It is easy to verify that

$$4d^2 - c^2 = \frac{16(1-p)p\xi^2(\tau-1)^2\tau}{(1+\xi)^2(1+p(\tau-1))^2} > 0,$$

so the condition is indeed satisfied. ■

Proof of Proposition 10:

We first state and prove the following lemma, which characterizes the equilibrium probability estimate:

Lemma A.1 *The equilibrium probability estimate is the solution to the equation*

$$\begin{aligned} \bar{p} &= qR(p, S_\tau(f\bar{p}(1-\Delta)), \xi, f(1+\Delta)) \\ &+ (1-2q)R(p, S_\tau(f\bar{p}), \xi, f(1+\Delta)) \\ &+ qR(p, S_\tau(f\bar{p}(1+\Delta)), \xi, f(1+\Delta)), \end{aligned} \quad (51)$$

where the function R is defined by

$$R(p, z, \xi, f) = \frac{1}{2f(1+\xi)} \left(1 + fp + f\xi + z\xi - \sqrt{(1 + fp + f\xi + z\xi)^2 - 4f(1+\xi)(p + z\xi)} \right).$$

The lemma states that an agent who observes higher-than-expected consumption updates beliefs as if the agent were observing only wealthy agents (who consume the fraction $\bar{p}f(1+\Delta)$ of bins rather than the average, $\bar{p}f$). The reason why the agent so strongly concludes that the wealthy were observed is that the number of observations Q and M are large. The observer finds the strength of the evidence of high consumption very surprising; the likelihood is low under the hypothesis that observations are of either high or low wealth agents. But the likelihood is especially low when the observation targets have low wealth, so the posterior belief puts all the weight on observing wealthy agents.

Proof of Lemma A.1: For $\alpha, \beta, f_1, f_2 \in (0, 1]$, $\xi > 0$, define

$$L(\alpha, \beta, \xi, f_1, f_2) = \lim_{Q \rightarrow \infty} \frac{J(\alpha, \beta, \xi, f_1, Q, 0)}{J(\alpha, \beta, \xi, f_2, Q, 0)},$$

where J was previously defined. It follows from standard properties of Beta distributions, that

$$L(\alpha, \beta, \xi, f_1, f_2) = \begin{cases} \infty, & |f_1 - \frac{\beta}{\alpha}| < |f_2 - \frac{\beta}{\alpha}|, \\ 0, & |f_1 - \frac{\beta}{\alpha}| > |f_2 - \frac{\beta}{\alpha}|. \end{cases} \quad (52)$$

An agent with prior p , who observes a fraction β of bins with consumption, believing that the observations provide an unbiased estimate of the consumption of others, and who

believes that the distribution of wealth groups (poor, medium, rich) among the population is $(q, 1 - 2q, q)$ who consume the fraction $(f(1 - \Delta), f, f(1 + \Delta))$ of the bins, respectively, will update—using Bayes rule—to the posterior:

$$\begin{aligned}\hat{p} &= \frac{qJ(\alpha, \beta, \xi, (1 - \Delta)f, Q, 1) + (1 - 2q)J(\alpha, \beta, \xi, f, Q, 1) + qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 1)}{qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0) + (1 - 2q)J(\alpha, \beta, \xi, f, Q, 0) + qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)} \\ &= \frac{J(\alpha, \beta, \xi, (1 - \Delta)f, Q, 1)}{J(\alpha, \beta, \xi, (1 - \Delta)f, Q, 0)}g_1 + \frac{J(\alpha, \beta, \xi, f, Q, 1)}{J(\alpha, \beta, \xi, f, Q, 0)}g_2 + \frac{J(\alpha, \beta, \xi, (1 + \Delta)f, Q, 1)}{J(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)}g_3,\end{aligned}$$

where

$$\begin{aligned}g_1 &= \frac{qJ(\alpha, \beta, \xi, (1 - \Delta)f, Q, 0)}{qJ(\alpha, \beta, \xi, (1 - \Delta)f, Q, 0) + (1 - 2q)J(\alpha, \beta, \xi, f, Q, 0) + qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)}, \\ g_2 &= \frac{(1 - 2q)J(\alpha, \beta, \xi, f, Q, 0)}{qJ(\alpha, \beta, \xi, (1 - \Delta)f, Q, 0) + (1 - 2q)J(\alpha, \beta, \xi, f, Q, 0) + qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)}, \\ g_3 &= \frac{qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)}{qJ(\alpha, \beta, \xi, (1 - \Delta)f, Q, 0) + (1 - 2q)J(\alpha, \beta, \xi, f, Q, 0) + qJ(\alpha, \beta, \xi, (1 + \Delta)f, Q, 0)}.\end{aligned}$$

It follows from (52) and assumption (28), that as $Q \rightarrow \infty$, $g_1, g_2 \rightarrow 0$ and $g_3 \rightarrow 1$. In words, the observing agent puts all the weight on having observed a wealthy agent's consumption, regardless of which agent he actually observes. Moreover, from (60) and it follows that an agent observing poor, average, and wealthy agents consuming based on posterior beliefs \bar{p} will have posterior belief

$$\hat{p} = R(\alpha, \beta, \xi, (1 + \Delta)f),$$

where β equals $S_\tau((1 - \Delta)f\bar{p})$, $S_\tau(f\bar{p})$, and $S_\tau((1 + \Delta)f\bar{p})$ with probabilities q , $1 - 2q$, and q , respectively. The fixed point problem that matches aggregate posterior beliefs with agents' updating is therefore (51).

Existence of a solution to (51) follows from the easily verifiable fact that $R(p, S_\tau(g \times 0), \xi, f) > 0$ and $R(p, S_\tau(g \times 1), \xi, f) < 1$, regardless of $g, p, f \in (0, 1)$, and $\xi > 0$. Therefore, the r.h.s of (51), which is a continuous function, is strictly greater than \bar{p} when \bar{p} close to zero, and strictly less than \bar{p} when \bar{p} is close to one. Existence therefore follows from the intermediate value theorem. This completes the proof of Lemma A.1. \blacksquare

It is straightforward to verify that the function R satisfies $\frac{\partial R}{\partial z} > 0$, since

$$\frac{\partial R}{\partial z} = \frac{\xi(1 + c)}{2f(1 + \xi)},$$

where

$$c^2 = \frac{(-1 - z\xi + f(2 - p + \xi))^2}{(f^2(p + \xi)^2 + (1 + \beta\xi)^2 + 2f(\xi - \beta\xi(2 + \xi) + \alpha(-1 + (\beta - 2)\xi))} < 1,$$

implying positivity of the derivative. It also follows that $\frac{\partial^2 R}{\partial z^2} = -\frac{2(1-f)(1-p)\xi^2}{((1+fp+f\xi+z\xi)^2-4f(1+\xi)(p+z\xi))^{3/2}} < 0$

0, so R is concave in z .

Moreover, one can show that $\frac{\partial R}{\partial f} < 0$. Specifically, it is easy to verify that $\frac{\partial R}{\partial f} = \frac{\kappa_1}{\kappa_2 f^2}$, where the function $\kappa_2 > 0$, and $\kappa_1 = 0 \Leftrightarrow f = 0$ on $f \in [0, 1]$, for the smooth function κ_1 . Thus, R is a monotone function for positive $0 \leq f \leq 1$. A Taylor expansion of κ_1 in f around $f = 0$ implies that κ_1 is of the form $R'(f) = -c_1 f^2 + O(f^3)$, where the constant $c_1 > 0$, altogether implying that $\frac{\partial R}{\partial f} < 0$ for small positive f , and thereby for all $0 \leq f \leq 1$ (since R is monotone).

Now, we use these properties of R to show that the total derivative of the r.h.s. of (51) w.r.t. Δ is negative. Specifically, using the notation R_i for the partial derivative of the function R w.r.t. its i th argument, from the calculus of total derivatives it follows that this r.h.s. derivative is of the form

$$q\bar{p}f(-S'_\tau(f\bar{p}(1-\Delta))R_2(\cdot) + S'_\tau(f\bar{p}(1+\Delta))R_2(\cdot)) + f(qR_4(\cdot) + (1-2q)R_4(\cdot) + qR_4(\cdot)).$$

Since $R_4(\cdot) < 0$, the second part of this expression is negative. Moreover, S_τ is concave and R is concave in its second argument, so the first part of the expression is also negative. Thus, the r.h.s. of (51) is decreasing in Δ .

Because R is increasing and concave in its second argument, it follows that the r.h.s. of (51) is concave and increasing in \bar{p} , and since $R(p, 0, \xi, f) > 0$, it follows that at the equilibrium point $0 < \frac{\partial R}{\partial \bar{p}} < 1$. Altogether, the inverse function theorem then implies that $\frac{\partial \bar{p}}{\partial \Delta} < 0$ for the fixed point \bar{p} defined by (51). ■

Proof of Proposition 11: The proof of the Bayesian updating follows similar lines as in Proposition C.2. Define

$$J(\alpha, \beta, \xi, f, g, Q, x) = \int_0^1 t^{\alpha Q - 1 + x} (1-t)^{(1-\alpha)Q-1} (g+ft)^{\beta \xi Q} (1-g-ft)^{(1-\beta)\xi Q} dt. \quad (53)$$

Standard properties of Beta distributions, implies that an agent's posterior estimate is

$$\hat{p} = R(\alpha, \beta, \xi, f, g) \stackrel{\text{def}}{=} \lim_{Q \rightarrow \infty} \frac{J(\alpha, \beta, \xi, f, g, Q, 1)}{J(\alpha, \beta, \xi, f, g, Q, 0)}, \quad (54)$$

Moreover, taking the derivative of the term inside the integral of (53) with respect to t , and using the factor that for large Q , J converges to a scaled Dirac distribution, it follows that \hat{p} satisfies:

$$\frac{\alpha}{\hat{p}} - \frac{1-\alpha}{1-\hat{p}} + f \left(\frac{\beta \xi}{g+f\hat{p}} - \frac{(1-\beta)\xi}{1-g-f\hat{p}} \right) = 0. \quad (55)$$

for large Q . Substituting in the equilibrium condition $\hat{p} = \bar{p}$, \bar{p} as a function of r defined in (36), setting $\alpha = p$, $\beta = S_\tau(g+f\bar{p})$, $\xi = 1$, f and g as defined in (31), and solving for τ in (55) leads to the functional relation:

$$\begin{aligned}
\tau(p, r) &= \left(2p + 2p^2 - 32r + 20pr + 4p^2r + 36r^2 + 3pr^2 + \right. & (56) \\
&+ \left. 5p^2r^2 - 50r^3 - 5pr^3 + 3p^2r^3 + 25r^4 - 10pr^4 + p^2r^4 \right) \\
&/ \left(-2p + 2p^2 + 20pr + 4p^2r + 48r^2 + 7pr^2 + \right. \\
&+ \left. 5p^2r^2 - 54r^3 - pr^3 + 3p^2r^3 + 15r^4 - 8pr^4 + p^2r^4 \right).
\end{aligned}$$

This relation thus represents the level of visibility bias that is consistent with equilibrium, given p and r .

It is easy to verify that $\tau(p, 0) = 1$, and thus that the unbiased equilibrium with $r = 0$ is obtained in this case. Moreover, one verifies that τ is strictly increasing in r in a neighborhood of 0, regardless of p , and that τ approaches infinity for some $r < 1/2$, so equilibrium is defined for all parameter values p and τ , and r increases in τ , as does then \bar{p} . Finally, since \bar{c}_0^I is decreasing in r , see (33), and $\bar{c}_0 = \bar{c}^e - \bar{c}_0^I$, it follows that \bar{c}_0 is increasing in r , and then also in τ , since r is increasing in τ . ■

B Other Utility Functions

The combination in the base model of the utility specification in (1), which leads to a linear consumption function in wealth (4), and the assumption that when $\frac{W}{2}$ is consumed at time 0 all bins contain consumption, makes the relationship between \bar{p} and $E[z]$ in (9) especially tractable, which allows for a strong characterization of equilibrium.

We now verify that qualitatively similar results as in Proposition 1 also hold under more common utility specifications. For example, consider the case in which agents have power utility, $U = \frac{c_0^{1-\gamma}}{1-\gamma} + \frac{c_1^{1-\gamma}}{1-\gamma}$, with risk aversion coefficient $\gamma \geq 1$ (where in the case $\gamma = 1$, log-utility is used). The consumption shock, ϵ is assumed to take on value $\frac{W}{2}$ with probability $1 - p$ (to avoid negatively infinite utility), and 0 with probability p . As before, the agent's estimated probability for a high outcome is \hat{p}

The agent's first order condition is in this case is

$$c_0^{-\gamma} = \hat{p}(W - c_0)^{-\gamma} + (1 - \hat{p}) \left(\frac{W}{2} - c_0 \right)^{-\gamma},$$

leading to the mapping $c_0 = G(\hat{p})\frac{W}{2}$. In the base model case with quadratic utility, $G(\hat{p}) = \hat{p}$. In the case of power utility, G is a nonlinear function for which a closed form solution

is not available, bare a few special values of γ .³⁷ However, the following behavior of G is easy to show:

Lemma B.2 *The function G satisfies $G(0) = \frac{1}{2}$, $G(1) = 1$, and is strictly increasing and convex. Its inverse is*

$$G^{-1}(c) = \frac{(1 - \frac{c}{2})^\gamma (\frac{c}{2})^{-\gamma} ((\frac{c}{2})^\gamma - (\frac{1}{2} - (\frac{c}{2}))^\gamma)}{(1 - \frac{c}{2})^\gamma - (\frac{1}{2} - \frac{c}{2})^\gamma}.$$

Proof: The form of G^{-1} follows immediately from the f.o.c. Differentiation of G^{-1} implies that G^{-1} is strictly increasing and concave on $c \in (1/2, 1)$. Moreover, $G^{-1}(1/2) = 0$, and $G^{-1}(1) = 1$. It follows that $G(0) = 1/2$, $G(1) = 1$, and from the inverse function theorem that G is invertible on $p \in (0, 1)$, being increasing and convex. ■

If an agent observes a fraction x of consumption bins, his posterior expected value of p is

$$\hat{p} = \frac{p + \xi G^{-1}(x)}{1 + \xi}.$$

Owing to visibility bias, if other agents' consumptions are based on the posterior expected probability \bar{p} , then $x = S_\tau(G(\bar{p}))$. Finally, in equilibrium, $\hat{p} = \bar{p}$, leading to the following fixed point equilibrium condition:

$$\bar{p} = \frac{p + \xi G^{-1}(S_\tau(G(\bar{p})))}{1 + \xi}. \quad (57)$$

The following proposition shows the existence of an equilibrium with over consumption in this setting:

Proposition B.1 *For $\tau > 1$, there exists an equilibrium probability estimate, $\bar{p} > p$, with associated consumption $G(\bar{p}) \frac{W}{2} > G(p) \frac{W}{2}$.*

Proof: Note that $y = G(\bar{p}) \in (\frac{1}{2}, 1)$ satisfies

$$G^{-1}(y) = \frac{p + \xi G^{-1}(S_\tau(y))}{1 + \xi}. \quad (58)$$

To show the existence of a $y \in (G(p), 1)$ solving (58), we note that

$$p = G^{-1}(G(p)) < \frac{p + \xi G^{-1}(G(p))}{1 + \xi} < \frac{p + \xi G^{-1}(S_\tau(G(p)))}{1 + \xi},$$

and that

$$1 = G^{-1}(G(1)) > \frac{p + \xi G^{-1}(S_\tau(G(1)))}{1 + \xi} = \frac{p + \xi}{1 + \xi}.$$

³⁷For the special case when $\gamma = 1$, the closed form solution is $G(\bar{p}) = \frac{5}{4} (5 - \bar{p} - \sqrt{9 - 10\bar{p} + \bar{p}^2})$.

By the intermediate value theorem, there is therefore a $y \in (G(p), 1)$ that solves (58), with associated equilibrium probability estimate $\bar{p} = G^{-1}(y) > G^{-1}(G(p)) = p$. ■

C The model with different fraction of consumption bins

The assumption that an agent consumes in all bins when $c_0 = W/2$ makes the analysis tractable, since the calculus of Bayesian updates with Beta distributed priors and observations is straightforward. A generalization is to assume that when $c_0 = W/2$, a fraction $0 < f \leq 1$ of the bins are full. This allows us to analyze situations where there is heterogeneity in consumption behavior, for example, because of wealth dispersion. Specifically, if a rich agent with probability estimate $\hat{p} = 1$ consumes in 100% of the consumption bins, then a poor agent with the same probability estimate will consume strictly less. The base model assumes $f = 1$, leading to the posterior estimate (8).

The following proposition covers the case when $f < 1$:

Proposition C.2 *The posterior expected probability of high consumption of an agent with prior p , who observes fraction z of bins being full, where each bin is full with probability pf , is*

$$\begin{aligned} \hat{p} &= R(p, z, \xi, f) \\ &= \frac{1}{2f(1+\xi)} \left(1 + fp + f\xi + z\xi \right. \\ &\quad \left. - \sqrt{(1 + fp + f\xi + z\xi)^2 - 4f(1+\xi)(p + z\xi)} \right). \end{aligned} \quad (59)$$

Proof: Define

$$J(\alpha, \beta, \xi, f, Q, x) = \int_0^1 t^{\alpha Q - 1 + x} (1-t)^{(1-\alpha)Q-1} (ft)^{\beta \xi Q} (1-ft)^{(1-\beta)\xi Q} dt.$$

Using standard properties of Beta distributions, it follows that

$$R(\alpha, \beta, \xi, f) = \lim_{Q \rightarrow \infty} \frac{J(\alpha, \beta, \xi, f, Q, 1)}{J(\alpha, \beta, \xi, f, Q, 0)}, \quad (60)$$

and that an agent who updates according to Bayes rule will arrive at the posterior estimate $\hat{p} = R(p, z, \xi, f)$ when Q is very large. ■

It is easy to verify that when $f = 1$, (59) reduces to the base model formula, $\hat{p} = \frac{p+\xi z}{1+\xi}$. Also, when $z = fp$, the formula reduces to $\hat{p} = p$, since the fraction of full bin observations is consistent with the prior in this case. Moreover, R is increasing in p and z , and is decreasing in f , since the lower f is, the lower the expected value of z is for a given prior

p , which makes any given number z of observed full bins a more favorable indication about p .

Using similar arguments as before, an equilibrium probability estimate when visibility bias is present is then defined as a solution to the fixed point equation:

$$\bar{p} = R(p, S_\tau(\bar{p}f), \xi, f). \tag{61}$$

We now have

Proposition C.3 *There exists a unique equilibrium. In equilibrium there is overconsumption, and the equilibrium probability estimate is*

$$\bar{p} = B(1 + f(\tau - 1), p, \xi), \tag{62}$$

where the function B is defined in (11,12).

Proof: Substituting in the definition of R into the fixed point problem (61) yields a cubic equation in \bar{p} , two roots of which are outside of the unit interval $(0, 1)$. The remaining root has the prescribed form. ■

The comparative statics from the base model therefore also hold in this variation. Moreover, increasing f has the same effect as increasing τ . Both lead to more overconsumption in equilibrium.

Corollary C.1 *The equilibrium probability estimate, \bar{p} is increasing in the consumption fraction, $\partial\bar{p}/\partial f > 0$, as is the overconsumption factor, \bar{p}/p .*

References

- Abel, A. B. (1990). Asset prices under habit formation and catching up with the Joneses. *American Economic Review* 80, 38–42.
- Agarwal, S., W. Qian, and X. Zou (2017, November). Thy neighbor's misfortune: Peer effect on consumption. Working paper, Georgetown University.
- Akerlof, G. A. and R. J. Shiller (2009). *Animal Spirits: How Human Psychology Drives the Economy and Why it Matters for Global Capitalism*. Princeton, NJ: Princeton University Press.
- Allen, T. W. and C. D. Carroll (2001, April). Individual learning about consumption. *Macroeconomic Dynamics* 5(2), 255–271.
- Bagwell, L. S. and B. D. Bernheim (1996, June). Veblen effects in a theory of conspicuous consumption. *American Economic Review* 86(3), 349–373.
- Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2016, May). Social networks and housing markets. Working paper, Facebook.
- Bailey, M., E. Dávila, T. Kuchler, and J. Stroebel (2018). House price beliefs and mortgage leverage choice. *Review of Economic Studies*. Forthcoming.
- Benhassine, N., F. Devoto, E. Duflo, P. Dupas, and V. Pouliquen (2015, August). Turning a shove into a nudge? A "labeled cash transfer" for education. *American Economic Journal: Economic Policy* 7(3), 86–125.
- Berger, J. and K. L. Milkman (2012, April). What makes online content viral? *Journal of Marketing Research* 49(2), 192–205.
- Bertrand, M. and A. Morse (2016, December). Trickle-down consumption. *Review of Economics and Statistics* 98(5), 863–879.
- Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2018). Behavioral household finance. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics: Foundations and Applications* 1, pp. 177–276. Amsterdam: Elsevier.
- Beshears, J. L., J. J. Choi, D. Laibson, and B. C. Madrian (2008). The importance of default options for retirement saving outcomes: Evidence from the USA. In S. J. Kay and T. Sinha (Eds.), *Lessons from Pension Reform in the Americas*, Chapter 3, pp. 59–87. Oxford: Oxford University Press.
- Beshears, J. L., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman (2015, June). The effect of providing peer information on retirement savings decisions. *Journal of Finance* 70(3), 1161–1201.
- Borgatti, S. P., A. Mehra, D. J. Brass, and G. Labianca (2009, February). Network analysis in the social sciences. *Science* 323(5916), 892–895.
- Borgida, E. and R. E. Nisbett (1977, September). The differential impact of abstract vs. concrete information on decisions. *Journal of Applied Social Psychology* 7(3), 258–271.
- Börsch-Supan, A. H., T. Bucher-Koenen, M. D. Hurd, and S. Rohwedder (2018, November). Saving regret. NBER Working Paper No. 25238.
- Bosworth, B. P. (1993, March). *Saving and Investment in a Global Economy*. Washington, D.C.: Brookings Institution Press.

- Brenner, L. A., D. J. Koehler, and A. Tversky (1996, March). On the evaluation of one-sided evidence. *Journal of Behavioral Decision Making* 9(1), 59–70.
- Brown, J. R., Z. Ivković, P. A. Smith, and S. Weisbenner (2008, June). Neighbors matter: Causal community effects and stock market participation. *Journal of Finance* 63(3), 1509–1531.
- Burnside, C., M. Eichenbaum, and S. T. Rebelo (2016, August). Understanding booms and busts in housing markets. *Journal of Political Economy* 124(4), 1088–1147.
- Bursztyn, L., A. L. González, and D. Yanagizawa-Drott (2018, June). Misperceived social norms: Female labor force participation in Saudi Arabia. NBER Working Paper No. 24736.
- Cai, H., Y. Chen, and H. Fang (2009, June). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99(3), 864–882.
- Campbell, J. Y. and J. H. Cochrane (1999, April). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107(2), 205–251.
- Card, D. and J. E. DiNardo (2002, October). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of Labor Economics* 20(4), 733–783.
- Carroll, C. D., B.-K. Rhee, and C. Rhee (1994, August). Are there cultural effects on saving? Some cross-sectional evidence. *Quarterly Journal of Economics* 109(3), 685–699.
- Carroll, C. D., B.-K. Rhee, and C. Rhee (1999, October). Does cultural origin affect saving behavior? Evidence from immigrants. *Economic Development and Cultural Change* 48(1), 33–50.
- Charles, K. K., E. Hurst, and N. Roussanov (2009, May). Conspicuous consumption and race. *Quarterly Journal of Economics* 124(2), 425–467.
- Chen, H., M. Michaux, and N. Roussanov (2013, September). Houses as ATMs? Mortgage refinancing and macroeconomic uncertainty. NBER working paper No. 19421, MIT Sloan School of Management.
- Chen, Y., F. M. Harper, J. Konstan, and S. X. Li (2010, September). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review* 100(4), 1358–1398.
- Cialdini, R., L. Demaine, B. Sagarin, D. Barrett, K. Rhoads, and P. Winter (2006, February). Managing social norms for persuasive impact. *Social Influence* 1(1), 3–15.
- Coibion, O., Y. Gorodnichenko, M. Kudlyak, and J. Mondragon (2014, January). Does greater inequality lead to more household borrowing? New evidence from household data. Working paper, University of Texas Austin.
- Cole, H. L., G. J. Mailath, and A. Postlewaite (1995, Summer). Incorporating concern for relative wealth into economic models. *Quarterly Review, Federal Reserve Bank of Minneapolis* 19(3), 12–21.
- Corneo, G. and O. Jeanne (1997, October). Conspicuous consumption, snobbism and conformism. *Journal of Public Economics* 66(1), 55–71.

- Daniel, K. D. and D. Hirshleifer (2015, Fall). Overconfident investors, predictable returns, and excessive trading. *Journal of Economic Perspectives* 29(4), 61–88.
- de Zúñiga, H. G. and S. Valenzuela (2011, December). The mediating path to a stronger citizenship: Online and offline networks, weak ties, and civic engagement. *Communication Research* 38(3), 397–421.
- DeMarzo, P., R. Kaniel, and I. Kremer (2008, August). Relative wealth concerns and financial bubbles. *Review of Financial Studies* 21(1), 19–50.
- Demarzo, P. M., R. Kaniel, and I. Kremer (2004, August). Diversification as a public good: Community effects in portfolio choice. *Journal of Finance* 59(4), 1677–1716.
- Donath, J. and D. Boyd (2004, October). Public displays of connection. *BT Technology Journal* 22(4), 71–82.
- Duflo, E. and E. Saez (2002, July). Participation and investment decisions in a retirement plan: The influence of colleagues' choices. *Journal of Public Economics* 85(1), 121–148.
- Duflo, E. and E. Saez (2003, August). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics* 118(3), 815–842.
- Dupor, B. and W.-F. Liu (1993, March). Jealousy and equilibrium overconsumption. *American Economic Review* 93(1), 423–428.
- East, R., K. Hammond, and M. Wright (2007, June). The relative incidence of positive and negative word of mouth: A multi-category study. *International Journal of Research in Marketing* 24(2), 175–184.
- Enke, B. (2017, September). What you see is all there is. Working paper, Harvard University.
- Federal Reserve Board (2018, May). Report on the economic well-being of U.S. households in 2017.
- Fiedler, K. (2008, January). The ultimate sampling dilemma in experience-based decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34(1), 186–203.
- Frederick, S. (2005, Autumn). Cognitive reflection and decision making. *Journal of Economic Perspectives* 19(4), 25–42.
- Frederick, S. (2012, June). Overestimating others' willingness to pay. *Journal of Consumer Research* 39(1), 1–21.
- Frey, B. and S. Meier (2004, December). Social comparisons and pro-social behavior: Testing “conditional cooperation” in a field experiment. *American Economic Review* 94(5), 1717–1722.
- Galí, J. (1994, February). Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit and Banking* 26(1), 1–8.
- Georgarakos, D., M. Haliassos, and G. Pasini (2014, May). Household debt and social interactions. *Review of Financial Studies* 27(5), 1404–1433.
- Georgarakos, D. and G. Pasini (2011, October). Trust, sociability, and stock market participation. *Review of Finance* 15(4), 693–725.

- Gerber, A. S. and T. Rogers (2009, January). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *The Journal of Politics* 71(1), 178–191.
- Giorgi, G. D., A. Frederiksen, and L. Pistaferri (2016, June). Consumption network effects. Working paper, NBER Working Paper No. 22357.
- Goldstein, N. J., R. B. Cialdini, and V. Griskevicius (2008, October). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research* 35(3), 472–482.
- Golub, B. and E. Sadler (2016). Learning in social networks. In Y. Bramoullé, A. Galeotti, and B. W. Rogers (Eds.), *The Oxford Handbook of the Economics of Networks*, Chapter 19, pp. 504–542. New York, NY: Oxford University Press.
- Grinblatt, M., M. Keloharju, and S. Ikäheimo (2008, November). Social influence and consumption: Evidence from the automobile purchases of neighbors. *Review of Economics and Statistics* 90(4), 735–753.
- Gross, D. B. and N. S. Souleles (2002, February). Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. *Quarterly Journal of Economics* 117(1), 149–185.
- Guidolin, M. and E. A. L. Jeunesse (2007, November/December). The decline in the U.S. personal saving rate: Is it real and is it a puzzle? *Federal Reserve Bank of St. Louis Review* 89(6), 491–514.
- Han, B., D. Hirshleifer, and J. Walden (2018, February). Social transmission bias and investor behavior. Working paper, University of Toronto.
- Harbaugh, R. (1996, December). Falling behind the Joneses: relative consumption and the growth-savings paradox. *Economics Letters* 53(3), 297–304.
- Healy, A. F. (1981, September). The effects of visual similarity on proofreading for misspellings. *Memory and Cognition* 9(5), 453–460.
- Hearst, E. (1991). Psychology and nothing. *American Scientist* 79(5), 432–443.
- Heimer, R. (2016, November). Peer pressure: Social interaction and the disposition effect. *Review of Financial Studies* 29(11), 3177–3209.
- Heyerdahl-Larsen, C. and J. Walden (2017, April). Welfare in economies with production and heterogeneous beliefs. Working paper, London Business School.
- Hirsh, J. B. (2015, July). Extraverted populations have lower savings rates. *Personality and Individual Differences* 81, 162–168.
- Hirshleifer, D. and S. H. Teoh (2003, December). Limited attention, information disclosure, and financial reporting. *Journal of Accounting & Economics* 36(1-3), 337–386.
- Hirshleifer, D. and S. H. Teoh (2009). Thought and behavior contagion in capital markets. In T. Hens and K. Schenk-Hoppe (Eds.), *Handbook Of Financial Markets: Dynamics And Evolution*, Handbooks in Finance, Chapter 1, pp. 1–46. Amsterdam, The Netherlands: North-Holland.
- Hong, H. G., J. D. Kubik, and J. C. Stein (2004, February). Social interaction and stock market participation. *Journal of Finance* 59(1), 137–163.
- Jackson, M. O. (2018). The friendship paradox and systematic biases in perceptions and social norms. *Journal of Political Economy*. Forthcoming.

- Jin, Y., H. Li, and B. Wu (2011, June). Income inequality, consumption, and social-status seeking. *Journal of Comparative Economics* 39(2), 191–204.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- Kahneman, D. and A. Tversky (1972, July). Subjective probability: A judgment of representativeness. *Cognitive Psychology* 3(3), 430–454.
- Kahneman, D. and A. Tversky (1973, July). On the psychology of prediction. *Psychological Review* 80(4), 237–251.
- Kaustia, M. and S. Knüpfer (2012, May). Peer performance and stock market entry. *Journal of Financial Economics* 104(2), 321–338.
- Koehler, J. J. and M. Mercer (2009, July). Selection neglect in mutual fund advertisements. *Management Science* 55(7), 1107–1121.
- Laibson, D. (1997, May). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2), 443–477.
- Laibson, D. I., A. Repetto, and J. Tobacman (1998). Self-control and saving for retirement. *Brookings Papers on Economic Activity* 1998(1), 91–196.
- Lambert, T. A., A. S. Kahn, and K. J. Apple (2003, May). Pluralistic ignorance and hooking up. *Journal of Sex Research* 40(2), 129–133.
- Lemieux, T. (2006, May). Postsecondary education and increasing wage inequality. *American Economic Review* 96(2), 195–199.
- Lerman, K., X. Yan, and X.-Z. Wu (2016, November). The “Majority Illusion” in social networks. *PLoS ONE* 11(2), 1–13.
- Ljungqvist, L. and H. Uhlig (2000, June). Tax policy and aggregate demand management under catching up with the Joneses. *American Economic Review* 90(3), 356–366.
- Loayza, N., K. Schmidt-Hebbel, and L. Servén (2000, September). Saving in developing countries: An overview. *World Bank Economic Review* 14(3), 393–414.
- Lusardi, A., O. S. Mitchell, and N. Oggero (2018, May). The changing face of debt and financial fragility at older ages. *AEA Papers and Proceedings* 108, 407–411.
- Madrian, B. and D. Shea (2001, November). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116(4), 1149–1187.
- Malmendier, U. and L. S. Shen (2018, June). Scarred consumption. Working paper, NBER Working Paper No. 24696.
- McShane, B. B., E. T. Bradlow, and J. Berger (2012, December). Visual influence and social groups. *Journal of Marketing Research* 49(6), 854–871.
- Mok, D., B. Wellman, and J. Carrasco (2010, November). Does distance matter in the age of the Internet? *Urban Studies* 47(13), 2747–2783.
- Moore, D. A., E. R. Tenney, and U. Haran (2015, December). Overprecision in judgment. In G. Keren and G. Wu (Eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making* (1st ed.), Chapter 6, pp. 182–209. John Wiley & Sons, Ltd.
- Moretti, E. (2011, January). Social learning and peer effects in consumption: Evidence from movie sales. *Review of Economic Studies* 78(1), 356–393.

- Neisser, U. (1963, September). Decision-time without reaction-time: Experiments in visual scanning. *American Journal of Psychology* 76(3), 376–385.
- Nisbett, R. and L. Ross (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- OECD (2014, August). OECD factbook 2014: Economic, environmental and social statistics. *OECD Publishing* 89(6).
- Parker, J. A. (1999, January). Spendthrift in America? On two decades of decline in the U.S. saving rate. *NBER Macroeconomics Annual* 14, 317–370.
- Parlour, C. A., V. Prasnikar, and U. Rajan (2007, August). Compensating for the winner’s curse: Experimental evidence. *Games and Economic Behavior* 60(2), 339–356.
- Perkins, H. W., M. P. Haines, and R. Rice (2005). Misperceiving the college drinking norm and related problems: A nationwide study of exposure to prevention information, perceived norms and student alcohol misuse. *Journal of Studies on Alcohol* 66(4), 470–478.
- Perkins, H. W., P. W. Meilman, J. S. Leichter, J. R. Cashin, and C. A. Presley (1999). Misperceptions of the norms for the frequency of alcohol and other drug use on college campuses. *Journal of American College Health* 47(6), 253–258.
- Piketty, T. and E. Saez (2003, February). Income inequality in the United States, 1913–1998. *Quarterly Journal of Economics* 118(1), 1–39.
- Poterba, J. M. (2014, May). Retirement security in an aging population. *American Economic Review* 104(5), 1–30.
- Poterba, J. M., S. F. Venti, and D. A. Wise (2012, May). Were they prepared for retirement? Financial status at advanced ages in the HRS and AHEAD cohorts. In D. A. Wise (Ed.), *Investigations in the Economics of Aging*, pp. 21–69. University of Chicago Press.
- Prentice, D. A. and D. T. Miller (1993, February). Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology* 64(2), 243–256.
- Salganik, M. J., P. S. Dodds, and D. J. Watts (2006, February). Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762), 854–856.
- Samuelson, W. and R. Zeckhauser (1988, March). Status-quo bias in decision making. *Journal of Risk and Uncertainty* 1(1), 7–59.
- Scellato, S., C. Mascolo, M. Musolesi, and V. Latora (2010, June). Distance matters: Geo-social metrics for online social networks. In *Proceedings of the 3rd Conference on Online Social Networks*, WOSN’10, pp. 8–8.
- Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006, August). Are Americans saving optimally for retirement? *Journal of Political Economy* 114(4), 607–643.
- Schroeder, C. M. and D. A. Prentice (1998, December). Exposing pluralistic ignorance to reduce alcohol use among college students. *Journal of Applied Social Psychology* 28(23), 2150–2180.
- Shefrin, H. M. and R. H. Thaler (1988, October). The behavioral life-cycle hypothesis. *Economic Inquiry* 26(4), 609–643.

- Shemesh, J. and F. Zapatero (2016, September). The intensity of keeping up with the Joneses behavior: Evidence from neighbor effects in car purchases. Working paper, Monash University.
- Stanford Center on Longevity (2016, February). Seeing our way to living long, living well in 21st century America. *The Sightlines Project*.
- Walden, J. (2018). Trading, profits, and volatility in a dynamic information network model. *Review of Economic Studies*. Forthcoming.
- White, K. J. C. and A. M. Guest (2003, September). Community lost or transformed? Urbanization and social ties. *City & Community* 2(3), 239–259.
- Wojnicki, A. C. and D. Godes (2008, April). Word-of-mouth as self-enhancement. Working paper, University of Toronto.