

December 2, 2015

Visibility Bias in the Transmission of Consumption Norms and Undersaving

Bing Han*

David Hirshleifer**

We study how bias in the social transmission process affects contagion of time preference norms. In the model, consumption is more salient than non-consumption. This visibility bias causes people to perceive that others are consuming heavily and to infer that others have a high discount rate. The transmission of norms for high discounting increases consumption and the equilibrium interest rate. Information asymmetry about the wealth of others dilutes the inference from high observed consumption that the discount rate is high. In consequence, in contrast with the Veblen wealth-signaling approach, information asymmetry about wealth *reduces* overconsumption. The visibility bias approach offers a novel explanation for the dramatic drop in the savings rate in the US and several other countries in the last thirty years. In contrast with other approaches, the visibility bias approach implies that relatively simple policy interventions can ameliorate undersaving.

*Rotman School of Management, University of Toronto

**Merage School of Business, UC Irvine

We thank Pierluigi Balduzzi, Martin Cherkas, Marco Cipriani, Andreas Fuster, Giacomo de Giorgi, Mariassunta Giannetti, Luigi Guiso (the AEA discussant), Robert Korajczyk, Theresa Kuchler (the Miami Behavioral Finance Conference discussant) Chris Malloy (the NBER discussant), Timothy McQuade, Peter Reiss, Andy Skrzypacz, Tano Santos, Siew Hong Teoh, Bernie Yeung, Youqing Zhou, participants at the American Economics Association meetings in Boston, 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; and seminar participants at Boston College, Columbia University, Peking University, Southwest University of Finance and Economics, Hong Kong Polytechnic University, Cheung Kong Graduate School of Business, London Business School, National University of Singapore, New York Federal Reserve Board, and Stanford University for very helpful comments; and Lin Sun and Qiguang Wang for very helpful comments and research assistance.

1 Introduction

In acquiring attitudes about the world, people are heavily influenced by their cultural milieu, and by interactions with other individuals. We would expect such influence to be stronger when it is hard to form a clear assessment through introspection. For example, several authors have argued that people find it hard to decide how heavily to discount the future in making savings decisions (e.g., Akerlof and Shiller (2009)). People may lack relevant information, or be unwilling to process relevant publicly available information.

As a result, people can't be sure what stream of satisfaction will actually result from a consumption/savings rule chosen today.¹ Similarly, it is hard to forecast remaining lifespan; most people do not process the relevant public but technical information contained in mortality tables and medical research. The same point applies to forecasting future health in old age, and to how much one's housing and recreation expenditures are likely to change.

This suggests that people may 'grasp at straws' by relying on noisy cues, including the behavior of others. A great deal of evidence suggests that people do make basic mistakes and rely on noisy cues in their savings decisions (Samuelson and Zeckhauser (1988), Shefrin and Thaler (1988), Madrian and Shea (2001), Beshears et al. (2008), Benhassine et al. (2013)). There is also considerable evidence that social interactions affect several dimensions of consumption and saving choices, sometimes in dysfunctional ways.²

The effects of social interactions involve more than just rational information transmission, as they include contagion of behaviors about which there is ample public information (participating in the stock market), and there is transmission of trading in individual stocks, an activity which individual investors are not, on average, good at.³ Surprisingly, however, there has been little formal modeling of how bias in social learning processes affect consumption choices over time.

In our model, social learning about the time preferences of others is biased by the fact that consumption is more salient than non-consumption. For example, a boat parked in a neighbor's driveway draws more attention than the absence of a boat. Similarly,

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."

²Dufo and Saez (2002, 2003)), Hong, Kubik, and Stein (2004), Brown et al. (2008), Kaustia and Knüpfer (2012)), Georgarakos, Haliassos, and Pasini (2014), and the evidence reviewed in Hirshleifer and Teoh (2009).

³See Barber and Odean (2000). In an experimental consumption/savings setting, social interaction caused subjects to deviate more from the optimal consumption path over time (Carbone and Duffy 2013).

it is more noticeable when a friend or acquaintance wears designer apparel or reports taking an expensive trip than when not, or acquires an unusual or interesting product as compared with not doing so. We call the greater salience of potential consumption events that occur rather than do not occur *visibility bias*. In addition, we assume that people do not adequately adjust for the selection bias in favor of noticing the consumption rather than nonconsumption events of others. This results in overestimation of others' consumption expenditures.

These key premises of our model—that consumption activities are more 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 and salience (see Section 2). Visibility bias in our model is not inherently an error; it is a source of bias in the *social transmission* of information. For example, being prone to noticing a neighbor's new car rather than the old one need not signify any cognitive failure. There are often good reasons to allocate attention to occurrences (or more generally, to salient events). However, failing to adjust appropriately for this selection bias in attention/observation *is* a mistake that biases inferences.

Visibility bias makes consumption more available than non-consumption for later retrieval and cognitive processing. In consequence, people infer high consumption and low savings rates by others, and conclude that a high discount rate is appropriate. Observers in our model therefore increase their own subjective discount rates accordingly, which increases actual consumption.

In the model, overestimation of the consumption of others is self-reinforcing, as each individual becomes an overconsuming model for others. So misperceptions of consumption norms can result in severe undersaving in society as a whole. In equilibrium, the reluctance of individuals to save implies a higher interest rate. Each individual ends up viewing the subjective discount rates of others as higher than his own. This comes from overestimation of others' discount rate, and the fact that each individual shifts only partially in the direction of the perceived norm.

Such a mismatch between beliefs and social reality, wherein everyone individually rejects a norm (in this case, a norm of high discount rates), yet believes that others embrace it, is called *pluralistic ignorance* (Katz and Allport 1931). For example, social psychology studies find that college students overestimate how much other students engage in and approve of heavy alcohol use (Prentice and Miller 1993) and uncommitted or unprotected sexual practices (Lambert, Kahn, and Apple 2003). These studies argue that pluralistic ignorance encourages such behaviors. In our model it is overconsumption that is promoted

by pluralistic ignorance.

This feature of our model can also help explain why parents, social observers, and religious leaders tend to preach in favor of thrift, and to criticize the now-focused consumer culture. To the extent that society is subject to pluralistic ignorance, people believe that others have an unduly high discount rate. Moral authorities will therefore believe that they can improve matters by using their influence against current consumption.

Advertising and media biases can further reinforce overconsumption. Advertisers have an incentive to depict consumers using their products heavily. News media serve their clientele by highlighting interesting high-end products or of consumption events (consider, e.g., the “Travel” section of newspapers). These further contribute to the higher visibility of consumption than nonconsumption. Of course, there is advertising of financial products as well. But it is much easier to vividly depict individuals consuming heavily at restaurants or exotic locations than to depict individuals saving heavily.

Our model helps explain puzzling patterns in the 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. 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 struggled to explain this drop (Parker 1999; Guidolin and Jeunesse 2007). 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.”⁴

The great savings decline presents a challenge to behavioral as well as rational models, since it is not obvious why psychological traits would shift over time toward greater impatience. 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. The drop in costs of long-distance communication, the rise of cable television, and then the rise of the internet, greatly increased the ability of

⁴Existing explanations for the decline in savings rates include recent gains in household wealth coming from the stock market and real estate; upward revisions in individuals estimate of their permanent-income due to technological advances and higher productivity, growing income inequality, increasing social safety net, increasing medical expenditure, aging population, easier access and improvements in the credit markets due to financial innovations, and trends in the way companies pay out to shareholders (disappearing dividends and more stock repurchases).

individuals to observe others' consumption, 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 networks devoted to shopping, travel, home remodeling, and other costly leisure pursuits, as well as dramas that less directly highlight consumption activities) further increased visibility. It is very common for people, in communicating on electronic networks, to report on such activities as travelling, eating out, and recent product purchases. Such observation is the driving force behind overconsumption in our model. Of course, the effects we model are far from the only drivers of savings rates, but it is one that offers a possible explanation for the large and anomalous drop in saving.

There are also notable differences in savings rates across countries and ethnic groups which are not well-explained by traditional economic models (Bosworth 1993). An implication of the visibility bias approach to time preference is that relatively modest differences in inherent discount rates can be amplified through social influence. This can help explain the diversity of savings rates across group. It suggests, for example, that cultural differences can have a surprisingly strong effect.⁵ 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 consumption theories that are not based upon social interaction.

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 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/or interaction. 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 the extent of overconsumption. It also implies that population level characteristics such as wealth dispersion matter, in contrast with pure individual-level biases such as present bias.

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

⁵Carroll, Rhee, and Rhee (1994) do not find an effect of culture on savings. The authors 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, but that these effects do not explain patterns in cross-country differences in savings rates.

others. In wealth signaling models, beliefs are rational, whereas the visibility bias approach is based upon biased inferences. The visibility approach has distinct empirical implications as well. For example, if all wealths were equal, Veblen effects would be eliminated, but the effects in our approach still apply. So in a visibility bias approach we expect to see overconsumption even within peer groups with low wealth inequality. More generally, as discussed in Section 6, wealth dispersion and asymmetric information about others' wealths is the source of Veblen effects, whereas the effects 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 investors deriving utility as a function of the consumptions of other investors (Abel 1990; Galí 1994; Campbell and Cochrane 1999), potentially resulting in overconsumption (Dupor and Liu 1993).⁶ This is a preference approach, as contrasted with our belief-based approach, which offers several distinctive implications.⁷

A further distinctive empirical and policy implication of the visibility bias approach is that exposing pluralistic ignorance by disclosing and saliently emphasizing information about the actual consumption or consumption attitudes of others will reduce overconsumption. That a relatively cheap and simple policy intervention can ameliorate the undersavings problem is specific to the visibility bias approach. Overconsumption in this approach derives from mistaken beliefs about others that can potentially be corrected. In contrast, the present bias approach is not based on social observation, and the other two approaches are based upon rational beliefs.

2 Psychology Background

The two key assumptions of our model are that consumption activities are more salient to others than nonconsumption; and that people do not adequately adjust for the selection bias in their attention toward consumption. This section briefly summarizes evidence from the psychology of attention and salience that motivate and support our assumptions.

There is extensive evidence that occurrences are more salient and easier to process than nonoccurrences (e.g., Neisser (1963), Healy (1981), and the review of Hearst (1991)). Occurrences provide sensory or cognitive cues that trigger attention. In the absence of

⁶Dupor and Liu (1993) show that jealousy (where an increase in aggregate consumption lowers an individual's utility level) can generate equilibrium overconsumption. In DeMarzo, Kaniel, and Kremer (2004, 2008) 'keeping up with the Joneses' effects can arise in settings with conventional preferences.

⁷Our model also provides a possible motivation for the assumption of external habit formation (utility interactions) in asset pricing models. This is usually viewed as being for tractability, but something akin to this arises naturally in our setting because of social influence in consumption.

such triggers, an individual will only react if (as is usually not the case) the individual was actively monitoring for a possible absence. This is why the famous phrase in the Sherlock Holmes story about “The dog that did not bark” is memorable; it is only the brilliant detective who recognizes the importance of the absence of an event.

An example of the low salience of non-occurrences is neglect of opportunity costs, i.e., non-occurrences of benefits that would be received under alternative courses of action. Economics instructors are well aware that the opportunity cost concept is something that students only master with effort, if at all. Another example of underweighting non-occurrences relative to occurrences is omission bias, the tendency of people to dislike and disapprove of actions that can result in adverse consequences, as compared with refraining from taking actions, when refraining can also result in adverse consequences. A classic example is an irrational reluctance to vaccinate (Ritov and Baron 1990).

In general, neglect of selection bias is implied, for example, by the representativeness heuristic of Kahneman and Tversky (1972). There is evidence from both psychology, experimental economics, and field studies of *selection neglect*, a failure of observers to adjust appropriately for data selection biases (Nisbett and Ross 1980; Brenner, Koehler, and Tversky 1996).⁸ Owing to limited cognitive resources, doing so requires time, 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. The evidence mentioned above that people underweight non-occurrences follows naturally from the low salience of non-occurrences together with selection neglect.

The combination of visibility bias and selection neglect in our model can explain the *availability heuristic* of Kahneman and Tversky (1973), so overestimation of consumption in the model can be viewed as a consequence of this heuristic. According 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. For example, people overestimate the frequency of shark attacks because such attacks are vivid and heavily reported in the media relative to other causes of death such as auto accident. The availability heuristic can be viewed as a failure to adjust for the selection bias in information brought to conscious attention—information that is stored into memory or easy to retrieve from it. In our model, this is disproportionately information about the consumption activities that were engaged

⁸People 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 fund performance owing to selection neglect (Koehler and Mercer 2009). Auction bidders in economic experiments tend to suffer from a winner’s curse (neglect of the selection bias inherent in winning), and hence tend to lose money on average Parlour, Prasnikar, and Rajan (2007).

in rather than not engaged in.

Experiments in the field confirm that consumption activity is more salient than non-consumption. As the study of Frederick (2012) concludes, “purchasing and consumption are more conspicuous than forbearance and thrift.” Consistent with the model’s implication that this results in overestimation of others’ consumption, in Frederick’s study this salience results in overestimation by observers of how much other individuals value certain consumer products. He explains the difference in salience between consumption and non-consumption vividly: “Customers in the queue at Starbucks are more visible than those hidden away in their offices unwilling to spend \$4 on coffee. We are repeatedly exposed to commercials of people enthusiastically gulping soda and gyrating to their iPods, but the large segment of nonusers is not so memorably depicted.” Also consistent with visibility bias, people are influenced in car purchases decisions by the purchases of others most strongly in areas where commuting patterns make the automobile purchase choices of others more visible (McShane, Bradlow, and Berger 2012).

3 The Basic Model

We first consider the effects of visibility bias in learning about the consumption of others with no uncertainty or production.

3.1 Optimal Consumption of Individuals

There are two consumption dates. At date 0, the individual chooses how much to consume and how much to borrow or lend at the riskfree interest rate r . Each individual i solves

$$\max_{c_{i0}, c_{i1}} U(c_{i0}) + \delta_i U(c_{i1})$$

subject to the intertemporal budget constraint

$$c_{i0} + \frac{c_{i1}}{1+r} = y_{i0} + \frac{y_{i1}}{1+r}, \tag{1}$$

where the y_i ’s are endowed levels of the consumption good at the two dates.⁹ The first order condition is

$$u'(c_{i0}) = \delta_i(1+r)u'(c_{i1}). \tag{2}$$

⁹The subjective discount factor δ_i should be distinguished from the subjective discount rate, which can be defined as $(1/\delta_i) - 1$. It is common to refer to discount rates in intuitive discussions; we do so at several points in the paper.

For most of the paper we assume logarithmic utility, $U(c) = \log(c)$. Then optimal consumptions c_{i0} and c_{i1} satisfy

$$\frac{c_{i1}}{c_{i0}} = \delta_i(1+r). \quad (3)$$

Define wealth as

$$W_i = y_{i0} + \frac{y_{i1}}{1+r}.$$

Combining (3) with the budget constraint (1) gives the individual's optimal consumption

$$c_{i0} = \frac{W_i}{1+\delta_i}, \quad (4)$$

so heavier subjective discounting (lower δ) causes current consumption to increase.

In the rest of this section, our focus is on the determination of date 0 consumption c_{i0} . For ease of notation, we omit the time subscript 0.

3.2 Visibility Bias and Learning About Others' Consumption

Suppose that there are N potential publicly observable consumption activities. Let the consumption intensity c be the propensity to engage in each of the N available consumption activities, where the probability that individual i undertakes any given activity is increasing in c . For any given activity, the probability that it is engaged in is $p(c) = c/\kappa$, where $0 \leq c \leq \kappa$. Each activity costs $K = \kappa/N$, $\kappa > 0$. (Having a different multiplied constant here would not qualitatively affect the results.) So the total expected consumption expenditure is

$$Np(c)K = c, \quad (5)$$

which is independent of N . In other words, we allow the number of activities to grow and the cost per activity to shrink correspondingly so that the expected cost remains unchanged. As N become large, the expenditure on consumption is close to its expectation c almost surely. We therefore refer to c henceforth as 'consumption' rather than 'consumption intensity.'

An individual randomly selects a target individual to observe, draws a sample of these potential consumption activities, and observes whether the target did or did not undertake each. (In the basic model, we will assume identical individuals, so that it would make no difference if an individual were to observe a sample from several targets.) Crucially, 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. We view the event of engaging in a consumption activity as generally more salient to others than the event of not doing so.¹⁰

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 travelling. Furthermore, physical shopping for consumption goods is itself a social activity. Both physical and electronic shopping and product evaluation are also engaging topics of conversation. In contrast, saving is often a private activity between an individual and his bank, brokerage, or retirement account software. Many television dramas display glamorous consumption activities, travel, entertaining, and dining, and some networks explicitly focus on shopping and other costly leisure activities. There are exceptions of course, such as investment clubs, but overall, consumption tends to be more observable and salient to others than is saving.¹¹

In particular, the probability that the observer samples any given potential activity is q^C if the individual did undertake the consumption activity and q^N if the individual did not, where $q^C > q^N$.¹² The intensity of visibility bias is captured by $\tau \stackrel{def}{=} q^C/q^N > 1$.

Letting $f(c)$ be the fraction of the activities sampled by the observer in which consumption occurs, the expected fraction is therefore

$$E[f(c)] = \frac{p(c)q^C}{p(c)q^C + [1 - p(c)]q^N} \quad (6)$$

$$= \frac{p(c)\tau}{p(c)\tau + [1 - p(c)]} > p(c). \quad (7)$$

The numerator in the first line is the probability that in any given activity consumption occurs and is detected, and the denominator is the probability that anything is detected—

¹⁰The occurrence versus non-occurrence distinction that we focus upon is not the only source of differences in the salience of different consumption behaviors. For example, information that is more action-relevant tends to be more salient. Extreme outcomes also tend to be more salient. Other things equal, this will cause observers to notice especially when others have either unusually low or unusually high 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 model does not require such effects; our focus is on an attentional bias—neglect of nonoccurrences—that has a clearcut directional implication.

¹¹Two other comparisons are of interest. First, retirement savings has 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 rate of current consumption of housing services financed by a major increase in indebtedness (and with the mortgage down payment usually much smaller than the size of the loan).

¹²We 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.

either consumption or non-consumption.

In the model, the observer fails to discount fully for the selection bias in observation, so the observer uses $f(c)$ to evaluate the frequency of the target's consumption activities. So visibility bias causes the observer to overestimate the fraction of the time that the target engages in consumption activity. As N , the number of activities observed, becomes large, $\lim_{N \rightarrow \infty} E[f(c)] = f(c)$, so that the fraction of activities sampled in which consumption occurs, as given by (6), is nonstochastic.

The observer inverts from the consumption activity fraction given in (6) to infer the target's total consumption. When the true consumption of others is c , the inferred consumption is

$$\hat{c} = h(c) \stackrel{def}{=} p^{-1}(f(c)) = \frac{cq^C \kappa}{cq^C + (\kappa - c)q^N} = \frac{c\kappa}{c + (\kappa - c)/\tau} > c, \quad (8)$$

where the last inequality follows since $\kappa - c > 0$ and $\tau > 1$. So owing to visibility bias, observers overestimate others' consumption levels, and the amount of overestimation increases with visibility bias, $\partial h(c)/\partial \tau > 0$.

3.3 Updating Discount Rates Based Upon Observation of Others

Let \hat{c}_i denote individual i 's inference about the level of others' consumptions, and suppose that all individuals have the same wealth, $W_i = W$ for all i . Then by (4), the time preference parameter $\hat{\delta}_i$ inferred by i satisfies

$$\hat{c}_i = \frac{W}{1 + \hat{\delta}_i},$$

so

$$\hat{\delta}_i(\hat{c}_i) = \frac{W}{\hat{c}_i} - 1. \quad (9)$$

Let δ^* be the common *inherent* time preference parameter for all individuals. (In a review of literature in consumer psychology, Simonson (2008) argues that preferences do have an 'inherent' core as well as constructed aspects.) After learning about others' consumption and time preference, the individual updates his own time preference parameter to δ_i by taking a weighted average of his inherent time preference, δ^* , and the inferred time preference of others, $\hat{\delta}_i$:

$$\delta_i = g(\hat{\delta}_i) \equiv (1 - \gamma)\delta^* + \gamma\hat{\delta}_i, \quad (10)$$

where $0 \leq \gamma < 1$, and where γ depends on the degree of social interactiveness/observability. This updating rule is based on the idea that an individual who thinks that others are consuming heavily infers that consuming heavily is a good idea.

This can be viewed as a reduced form for a setting in which individuals acquire information from others about the benefit from saving. For example, an individual may be learning from others (not necessarily rationally) about the probability of dying at a young age or about how consumption needs will change in retirement. This emphasizes that the driving force of our model is a bias in information processing rather than a direct preference for early consumption.

Alternatively, the inference could be moralistic, i.e., learning about whether being a good person demands providing for the future—recall Aesop’s fable of the ant and the grasshopper. Akerlof (2007) emphasizes that moral attitudes affect consumption decisions; under the moralistic interpretation of our model, such attitudes are influenced by the perceived consumption of others.

We refrain from imposing a rational-expectations-like condition that an individual understands that others share the individual’s inherent time preferences. We do not think of individuals as even recognizing the distinction between others’ inherent and actual time preferences. Each individual simply updates based upon inferences about the actual time preferences of others. Such simplified reasoning, wherein an observer fails to take into account fully the thought processes of other individuals, is broadly consistent with models and experimental studies of limited cognition about the thought processes of others (e.g., Hirshleifer and Teoh (2003), Camerer, Ho, and Chong (2004), Eyster and Rabin (2005) and Jehiel (2015)).

Also, our model captures conformism by assuming that observation of others causes people to update their discount factors. An alternative approach would be based on a direct preference for consuming to match the consumption timing of others. Such a keeping-up-with-the-Joneses’ approach, however, potentially generates multiple equilibria with early or late consumption (Stracca and al Nowaihi 2005), whereas our approach predicts a specific direction, overconsumption.

The updating rule (10), together with (4), implies that individual i selects a current consumption level of

$$c_i = \frac{W}{1 + g(\hat{\delta}_i)}. \quad (11)$$

3.4 The Symmetric Equilibrium

Assume that all individuals have identical utility functions, inherent discount factors, and consumption endowments. We seek a symmetric equilibrium, i.e., a fixed point in consumption and discount factor, for given wealth W and model parameters: $0 < \delta^* < 1$ (inherent

discount factor), γ (weight on the inferred discount factor of others) and τ , intensity of visibility bias.

We define (symmetric) equilibrium as follows.

Definition of Equilibrium *In an equilibrium, for all i , $c_i = c$, $\hat{\delta}_i = \hat{\delta}$ satisfy:*

$$(1 + g(\hat{\delta}))c = W \quad (12)$$

$$(1 + \hat{\delta})h(c) = W, \quad (13)$$

with the functions g and h defined as

$$g(\hat{\delta}) \stackrel{def}{=} (1 - \gamma)\delta^* + \gamma\hat{\delta},$$

and

$$h(c) \stackrel{def}{=} p^{-1}(f(c)).$$

In this definition of equilibrium, (12) is the requirement that the individual optimize based upon the individual's updated discount rate as in (11), and (13) is the requirement that the individual draws inferences about others' discount rates recognizing that others are optimizing in their consumption choices, i.e., follow the solution as given in (4) with δ_i replaced with $\hat{\delta}$.

This is an equilibrium in the sense that individuals are satisfied given their observations of others and the exogenous interest rate. But we have not yet imposed market clearing to determine prices.

Consistent with intuition, in this setting visibility bias causes overconsumption.

Proposition 1 *In a setting with exogenous interest rate and with visibility bias, under log utility and the other assumptions of the model:*

1. *There exists a unique symmetric equilibrium.*
2. *The consumption level c is higher than that without visibility bias.*
3. *The amount of overconsumption increases with visibility bias (τ) and with social interactiveness/observability (γ).*

Proof: From (13), $\hat{\delta} = W/h(c) - 1$. Substituting this into (12), equilibrium consumption c satisfies

$$[1 + (1 - \gamma)\delta^*]c + \gamma \left(\frac{W}{h(c)} - 1 \right) c = W. \quad (14)$$

Let c_{δ^*} denotes the optimal consumption level corresponding to the inherent time preference δ^* ,

$$c_{\delta^*} = \frac{W}{1 + \delta^*}. \quad (15)$$

Let

$$F(x) \stackrel{def}{=} [1 + (1 - \gamma)\delta^*]x + \gamma \left(\frac{W}{h(x)} - 1 \right) x - W \quad (16)$$

be the difference between the LHS and RHS of (14) as a function of possible consumption levels x . F is a continuous function with the property that

$$F(c_{\delta^*}) = \frac{\gamma W}{1 + \delta^*} \left(-\delta^* + \frac{W}{h(c_{\delta^*})} - 1 \right) = \frac{\gamma W^2}{1 + \delta^*} \left(\frac{1}{h(c_{\delta^*})} - \frac{1}{c_{\delta^*}} \right) < 0,$$

because $h(c_{\delta^*}) > c_{\delta^*}$; and $F(W) = g(\hat{\delta})W > 0$. Therefore, there exists at least one $c \in (c_{\delta^*}, W)$ satisfying $F(c) = 0$. Hence, the equilibrium exists, i.e., there exists a fixed point in consumption c and discount factor $\hat{\delta}$ satisfying (12) and (13).

Uniqueness of the equilibrium follows from the fact that F is monotonically increasing, so that there can be at most one solution to $F(c) = 0$. To verify monotonicity, differentiate F to get

$$F'(c) = (1 - \gamma)(1 + \delta^*) + \gamma W \frac{h(c) - ch'(c)}{h^2(c)}.$$

By (8),

$$h(c) - ch'(c) = \frac{c^2 q^C (q^C - q^N) \kappa}{[cq^C + (\kappa - c)q^N]^2} > 0,$$

so $F' > 0$.

To verify Part 2, recall by (8) that $h(c) > c$. Substituting for $h(c)$ and c from (12) and (13), it follows that

$$\hat{\delta} < g(\hat{\delta}) = (1 - \gamma)\delta^* + \gamma\hat{\delta}.$$

So $\hat{\delta} < \delta^*$, and hence

$$g(\hat{\delta}) < \delta^*. \quad (17)$$

By (12) and (15), we conclude that $c > c_{\delta^*}$.

To verify Part 3, observe that the overconsumption ratio is

$$c/c_{\delta^*} = \frac{1 + \delta^*}{1 + g(\hat{\delta})},$$

so it suffices to show that $\partial g(\hat{\delta})/\partial \tau < 0$ and $\partial g(\hat{\delta})/\partial \gamma < 0$. This follows from

$$\begin{aligned}\frac{\partial g(\hat{\delta})}{\partial \tau} &= \gamma \frac{\partial \hat{\delta}}{\partial \tau} = -\gamma \frac{W}{h(c)} \frac{\partial h(c)}{\partial \tau} < 0 \quad (\text{by (8)}) \\ \frac{\partial g(\hat{\delta})}{\partial \gamma} &= -\delta^* < 0.\end{aligned}\tag{18}$$

||

Intuitively, owing to visibility bias in what is observed, and neglect of sample size/use of the availability heuristic in assessing frequencies, people overestimate others' consumption, and therefore overestimate others' discount rates. Based on a misperception that the social norm is less thrifty than it really is, people update their own time preferences toward current consumption.

As discussed in the introduction, personal saving rates have plunged in the U.S. and several other OECD countries over the last 30 years. Existing rational and behavioral theories do not seem to fully explain this phenomenon. Part 3 of Proposition 1 provides 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), email, blogging, and social networking (such as Facebook) 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 γ , as in Part 3. Greater observation and communication about consumption is reflected by higher γ in the model. Greater observability of the consumption of others, as reflected in γ , intensifies the effects of visibility bias and implies a reduction in the savings rate.

These changes also suggest an increase in τ , bias in favor of observing consumption over nonconsumption. The activities that are noteworthy to report on very often involve expensive purchases, as with eating out or travelling. Indeed, many social networking and review sites are organized around consumption activities, such as Yelp. In-person observation of nearby individuals or close friends may often include even nonconsumption activities, whereas electronic reports tend to select especially strongly for consumption activities. So the rise in modern communications results in an increase in visibility bias (i.e., larger τ) and lead to higher overconsumption.¹³

¹³Hirsh (2015) provides evidence suggesting that the drop in savings rates was caused by increasing population-level extraversion in many countries. This 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 electronic communication technologies.

The analysis suggests that a relatively cheap and simple policy intervention—saliently publicizing the consumption rates of peers—can help alleviate overconsumption. In the model people conform to what they perceive as the high discount rate of peers. Learning that peers actually save more, such pluralistic ignorance and mistaken updating is reduced.

Empirically, the intervention of providing accurate information about peers tends to cause behavior to conform more closely to the disseminated peer norm in tests covering a very wide range of activities. In the studies discussed in the introduction of pluralistic ignorance about college student drinking and sexual attitudes/behaviors, dissemination therefore resulted in a shift toward the reported norms (which were less favorable to hazardous activities). If people overestimate the consumption of peers, similarly suggests that salient reporting of actual consumption of peers would reduce consumption.¹⁴

The assumption that individuals are identical, yet misperceive the attitudes of others, is stark. However, similar findings would apply in settings with heterogeneous discount rates. Furthermore, as mentioned earlier, such mismatches between own beliefs and social reality capture the empirical phenomenon of pluralistic ignorance, wherein everyone individually rejects a norm, yet believe that others favor it.

The social influence parameter γ is identical across individuals. With diverse γ 's, we conjecture that individuals with greater susceptibility to social influence will overconsume more than those with lower γ . Such individuals update their preferences more in the direction of their overestimated perception of the consumption norm. It is evident that this is true for the special case in which γ is identical for almost everyone.

Result 1 *For given social influence parameter γ of other investors, an individual with zero mass with parameter γ' consumes more than others if and only if $\gamma' > \gamma$, and less than others if and only if $\gamma' < \gamma$. This individual's consumption is increasing with γ' .*

This implication that individuals with greater susceptibility to social influence will overconsume more is empirically testable.

¹⁴Studies in which peer information promotes conformity toward the disclosed norm include Frey and Meier (2004), Cialdini et al. (2006), Salganik, Dodds, and Watts (2006), Goldstein, Cialdini, and Griskevicius (2008), Cai, Chen, and Fang (2009), Gerber and Rogers (2009), and Chen et al. (2010). However, sometimes people adjust their behavior away from the disclosed actions of others. In an experiment on retirement savings behavior in a large manufacturing firm, Beshears, Choi, Laibson, Madrian, and Milkman (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 firms 401(k) plan. Such employees may regard the higher-income employees who were plan participants as not truly peers, and hence instead of updating toward the high-income norm, low-income employees might be discouraged and shift in the opposite direction.

The comparative statics on varying endowments are intuitive. Details are provided in the appendix. Under homothetic preference (such as log utility), visibility bias increases the marginal propensity to consume from current income. Visibility bias and resulting misperceptions of social norms cause the individual's current consumption c_0 to be more sensitive to changes in income than without visibility bias, because higher wealth tends to increase both current and future consumption. With homothetic preferences, the ratio between the two is constant. Social influence and visibility bias in our setting increase the ratio of current to future consumption, and therefore also the marginal propensity to consume out of income (or wealth).

This finding (as summarized in Proposition 6 in the appendix) is consistent with the finding in Heffetz (2011) that income elasticity of a consumption good increases with its sociocultural visibility, because visibility bias is stronger for more visible/noticeable consumption.

3.5 The Equilibrium Interest Rate

The reasoning so far focuses on individual decisions with the riskfree rate r taken as given. We next solve for r by clearing the bond market. Intuitively, since visibility bias in social observation results in too low a time discount factor (δ), the equilibrium interest rate (r) is too high.

In a pure exchange economy, people can't all consume more today. The pro-consumption effects of visibility bias on δ are offset by a corresponding rise in r , so that the representative individual consumes the (exogenous) per capita income (although each individual thinks that his neighbors consume more than himself). In Section 5, allowing for intertemporal production reinstates the result that visibility bias increases current consumption, as individuals are able to satisfy their amplified preferences for current consumption by investing less.

From the optimization condition (3) and the budget constraint (1), the optimal consumption of the representative individual at date 0 satisfies

$$c_0(1 + \delta) = y_0 + \frac{y_1}{1 + r}.$$

In the symmetric equilibrium, each individual perceives that others use a time discount parameter $\hat{\delta}$, and updates his own discount factor to $g(\hat{\delta}) = (1 - \gamma)\delta^* + \gamma\hat{\delta}$. By (12), equilibrium satisfies

$$c_0(1 + g(\hat{\delta})) = y_0 + \frac{y_1}{1 + r}. \tag{19}$$

Since individuals are identical, in equilibrium the interest rate r zeroes out borrowing and lending. This implies that $c_0^* = y_0$, so by (19),

$$1 + r = \frac{y_1}{y_0 g(\hat{\delta})}. \quad (20)$$

Comparing to the case without visibility bias, where equilibrium interest rate satisfies (20) with $g(\hat{\delta})$ replaced by δ^* , it is evident that visibility bias raises the equilibrium interest rate because $g(\hat{\delta}) < \delta^*$ (equation (17)).

Proposition 2 *Under pure exchange and log utility, the equilibrium riskfree interest rate is higher when individuals are subject to visibility bias in social observation than when they are not.*

We argue at the end of Section 3.4 that with the rise in modern communications there has been greater visibility bias and social learning, leading to a higher overconsumption and decline in the U.S. savings rate over the last two decades. By Proposition 2, the same argument would also imply an increasing trend in the real interest rate, whereas the data shows a general decline in the real interest rate since the early 1980's (e.g., International Monetary Fund, April 2014). However, in reality equilibrium interest rates are determined in a global market, and there are many forces at work determining the U.S. interest rate that are outside our model (e.g., FED monetary policy and increased supply of saving from Germany and from emerging market economies such as China, owing to institutional changes in these countries). Our model suggests that if individuals are not subject to visibility bias, or if its effects are eliminated owing to an absence of social interactions or social observations, then the interest rate would have been even lower.

Also, it could be argued that the declining trend in the real interest rate offers an alternative explanation to ours for the saving rate trend. However, in general the effect of a decline in real interest rate on saving is theoretically ambiguous since the substitution and income effects work in opposite directions, and there is mixed empirical evidence on the sign of interest rate elasticity of saving. For example, under the permanent income hypothesis the elasticities of consumption and saving to interest rate depends on the model parameters such as the intertemporal elasticity of substitution. Weber (1970) finds that aggregate U.S. consumption increases in response to an increase in interest rate (so that the income effect dominates the substitution effect). Hall (1988) finds that intertemporal substitution is weak in postwar data for the United States, so that a higher real interest rate does not induce significant postponement of consumption. In addition, Parker (1999)

points out that the real interest rate was high during the consumption boom in the 1980's and 1990's relative to the previous two decades.

4 General Utility Function and Income Uncertainty

In the model of Section 3, there is no uncertainty, and we assumed log utility. We now show that visibility bias causes overconsumption in a setting with uncertainty about future income and with an arbitrary increasing concave utility function $U(c)$. We examine this issue in a partial equilibrium with exogenous interest rate r .

With y_1 uncertain at date 0, each individual solves

$$\max_{c_0} U(c_0) + \delta E[U(c_1)]$$

subject to the intertemporal budget constraint

$$c_0 + \frac{c_1}{1+r} = y_0 + \frac{y_1}{1+r}. \quad (21)$$

Optimal consumption satisfies

$$u'(c_0) = \delta(1+r)E[u'(c_1)]. \quad (22)$$

Differentiating both sides of (22) with respect to δ gives

$$u''(c_0) \frac{\partial c_0}{\partial \delta} = (1+r)E[u'(c_1)] + \delta(1+r)E[u''(c_1)] \frac{\partial c_1}{\partial \delta}. \quad (23)$$

The budget constraint (1) implies that

$$\frac{\partial c_1}{\partial \delta} = -(1+r) \frac{\partial c_0}{\partial \delta}.$$

Substituting this into (23) yields

$$\{u''(c_0) + \delta(1+r)^2 E[u''(c_1)]\} \frac{\partial c_0}{\partial \delta} = (1+r)E[u'(c_1)]. \quad (24)$$

The RHS of (24) is positive because $u' > 0$; the coefficient on $\partial c_0 / \partial \delta$ on the LHS is negative because $u'' < 0$. Hence, even under general risk averse preference and uncertain future income, it remains the case that as an individual becomes more impatient (δ becomes smaller), he consumes more today:

$$\frac{dc_0}{d\delta} < 0. \quad (25)$$

It is then straightforward to show that the optimal consumption c_0 with biased social transmission is higher than the case without visibility bias, because in the former case, the individual uses an updated discount factor $g(\hat{\delta})$ that is lower than the discount factor δ^* that obtains without visibility bias. So visibility bias and social influence increases equilibrium consumption.

Proposition 3 *In a setting with general von-Neumann Morgenstern risk averse utility, uncertain future income y_1 , and an exogenous interest rate, visibility bias increases consumption relative to a setting with no visibility bias.*

5 The Model with Production

We have shown in a setting with interest rate exogenously fixed that visibility bias increases consumption; and, in a pure exchange equilibrium setting, that visibility bias increases the interest rate, where equilibrium consumption is fixed by the market clearing condition at endowed consumption.

We now extend the basic model to allow for productive transformation between current and future consumption. In addition to the riskfree asset, individuals can invest some of their savings in a production technology which produces Y_1 units of future consumption goods using I units of investment. At date 0, each individual chooses the level of consumption (c_0), allocates a positive or negative amount to riskfree bonds (b) and to real investment (I). Each individual is endowed with exogenous incomes y_0 and y_1 at the two dates.

Individual i solves the optimization problem

$$\max_{c_0, c_1} U(c_0) + \delta U(c_1)$$

subject to the technological constraint

$$Y_1 = H(I) \stackrel{def}{=} AI^\alpha, \quad 0 < \alpha < 1, \quad (26)$$

and the budget constraints

$$c_0 = y_0 - I - b \quad (27)$$

$$c_1 = Y_1 + y_1 + b(1 + r). \quad (28)$$

The optimization can be done over two controls c_0 and b . Once these two are chosen, by (27), the initial investment in the production technology is

$$I = y_0 - c_0 - b.$$

Together with (26) and (28), this implies that the date 1 consumption is

$$c_1 = H(y_0 - c_0 - b) + y_1 + b(1 + r). \quad (29)$$

Substituting (29) into the objective function $L = U(c_0) + \delta U(c_1)$ of the maximization problem, we obtain the first order conditions with respect to c_0 and b :

$$\frac{\partial L}{\partial c_0} = U'(c_0) - \delta U'(c_1)H(I) = 0 \quad (30)$$

$$\frac{\partial L}{\partial b} = \delta U'(c_1)[-H'(I) + 1 + r] = 0. \quad (31)$$

Equation (31) and the definition of $H(I) = A(I)^\alpha$ imply that

$$\alpha AI^{\alpha-1} = 1 + r, \quad \text{or} \quad (32)$$

$$I = \left(\frac{1+r}{\alpha A} \right)^{\frac{1}{\alpha-1}}. \quad (33)$$

The log utility function $U(c) = \log(c)$ and (30) imply that

$$c_1 = \delta H'(I)c_0 = \delta(1+r)c_0. \quad (34)$$

By (27), it follows that

$$c_0 + b = y_0 - \left(\frac{1+r}{\alpha A} \right)^{\frac{1}{\alpha-1}} = y_0 - I. \quad (35)$$

By (28) and (34),

$$\delta c_0 = \frac{y_1 + H(I)}{1+r} + b. \quad (36)$$

We can solve for date 0 optimal consumption c_0 and allocation to riskfree asset b from (35) and (36). Let W be the individual's wealth, defined as the discounted value of the sum of present and future endowment y 's and of production of future consumption $Y = H(I)$, net of expenditure on investment:

$$W = y_0 + \frac{y_1 + H(I)}{1+r} - I.$$

Adding (35) and (36), and cancelling b from both sides gives

$$c_0 = \frac{W}{1+\delta}. \quad (37)$$

By (35) and (37), the optimal bond investment is

$$b = \frac{\delta}{1 + \delta}(y_0 - I) - \frac{1}{1 + \delta} \frac{y_1 + H(I)}{1 + r}. \quad (38)$$

In equilibrium, the bond market clears, $b = 0$, so the equilibrium interest rate satisfies

$$1 + r = \frac{y_1 + H(I)}{\delta(y_0 - I)}. \quad (39)$$

Corresponding to this interest rate, the optimal consumption is $c_0 = y_0 - I$. Since a log utility investor will never consume a negative amount, the equilibrium condition that $b = 0$ implies that $I < y_0$.

Combining (33) and (39), the equilibrium level of investment I satisfies

$$G(I; \delta) = 0, \quad (40)$$

where the function G is defined as

$$G(I; \delta) \stackrel{def}{=} AI^\alpha(1 + \alpha\delta) - \alpha\delta Ay_0 I^{\alpha-1} + y_1. \quad (41)$$

Differentiating shows that

$$\frac{\partial G}{\partial \delta} = \alpha AI^{\alpha-1}(I - y_0) < 0 \quad (42)$$

$$\frac{\partial G}{\partial I} = \alpha AI^{\alpha-2}[(1 + \alpha\delta)I - (\alpha - 1)\delta y_0] > 0, \quad (43)$$

because $I < y_0$ and $0 < \alpha < 1$. To see that this implies a solution for I within its support $(0, y_0)$, observe that $\lim_{I \rightarrow 0} G(I; \delta) = -\infty$, that $G(y_0; \delta) > 0$, and that $\partial G / \partial I > 0$.

By (42), (43), and the implicit function theorem, $\partial I / \partial \delta > 0$. Since, by (17), visibility bias reduces δ , the equilibrium investment I is lower, and hence the consumption level $c_0 = y_0 - I$ is higher.

Let equilibrium be defined by each individual's date 0 consumption c_0 and his inference of others' time preference $\hat{\delta}$ satisfying

$$(1 + g(\hat{\delta}))c_0 = W \quad (44)$$

$$(1 + \hat{\delta})h(c_0) = W, \quad (45)$$

with the function g given by

$$g(\hat{\delta}) \stackrel{def}{=} (1 - \gamma)\delta^* + \gamma\hat{\delta},$$

and the inferred consumption of others $h(c)$ by

$$h(c) \stackrel{\text{def}}{=} p^{-1}(f(c)).$$

Here owing to visibility bias $h(c) > c$.

The following proposition summarizes these results.

Proposition 4 *In an equilibrium setting with log utility, visibility bias and intertemporal production:*

(1) *There exists a unique symmetric equilibrium. At date 0 all individuals invest*

$$I = \left(\frac{1+r}{\alpha A} \right)^{\frac{1}{\alpha-1}}.$$

Each individual invests b in the riskfree asset at date 0, where

$$b = \frac{g(\hat{\delta})}{1+g(\hat{\delta})}(y_0 - I) - \frac{1}{1+g(\hat{\delta})} \frac{y_1 + AI^\alpha}{1+r},$$

and where the riskfree rate r is

$$1+r = \frac{y_1 + A(I)^\alpha}{g(\hat{\delta})(y_0 - I)}.$$

(2) *With visibility bias, the equilibrium consumption level c_0 is higher, and investment I is lower, than in the absence of visibility bias.*

So as in the basic setting with exogenous interest rate of Subsection 3.4, visibility bias increases consumption, and as in the equilibrium analysis of Subsection 3.5 (in which the pure exchange setting precluded an effect on equilibrium consumption), visibility bias increases the interest rate. Here visibility bias also decreases saving and real investment.

6 Information Asymmetry

We now generalize to allow for wealth dispersion in the population, and for individuals not knowing the wealths of others. Intuitively, the inference an observer draws about the discount rate of others based on observation of another's consumption is weaker if the observer does not know the target's wealth, owing to conflation between the possibilities that the discount rate is high or that wealth is high. 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 model information asymmetry, we assume that each individual, as an observer, has a nondegenerate prior distribution over the common actual discount rate of others. As in the earlier sections, the setup reflects pluralistic ignorance; an observer does not understand that everyone is essentially similar to the observer with respect to discount rate. This creates an opening for visibility bias and the availability heuristic to influence perceptions.

We grant individuals enough rationality to understand that when they see high consumption, this could come from either a high discount rate or high wealth. As a simple benchmark to highlight the effects of visibility bias, we assume that people have correct prior beliefs about the joint distribution of others' wealths and their actual discount rate. People update from this correct prior to an incorrect posterior about these variables.

As in the earlier sections, observers draw inferences about others' actual discount rate, not their inherent discount rate. Observers view others' actual discount rate as normative, and update their own discount rates in the direction of their inferences. However, here even after observing others an individual does not feel sure about others' discount rate, so the individual updates in the direction of the conditional expected discount rate of others.

We return to the pure exchange setting of Section 3, except that we allow for wealth dispersion. Let $f_\delta(\delta)$ be the prior probability density that each individual has about the common actual discount factor of others. All observers share this prior about possible targets of observation, and this density matches the true underlying density. Similarly, let $g_W(W)$ be the prior density for individuals' wealths (independently distributed across targets of observation). Wealth and actual discount factors are independently distributed, and everyone correctly perceives this to be the case. An observer makes observations of a single target; since the wealth distribution is identical across individuals, we omit subscripts identifying which target.

For notational simplicity we now omit hats for variables indicating individual perceptions. By (4) and independence of δ and W , an observer who infers that a target's consumption is c for sure updates his belief about δ to

$$\begin{aligned} f_\delta(\delta|c) &= \frac{f_{\delta,c}(\delta, c)}{f_c(c)} \\ &= \frac{f_\delta(\delta)g_W((1+\delta)c)}{\int_\delta f_\delta(\delta)g_W((1+\delta)c)d\delta}. \end{aligned} \tag{46}$$

Assume that $\delta \sim U[0, 1]$ and $W \sim U[1 - K, 2 + K]$, where $0 \leq K \leq 1$. K measures the amount of wealth dispersion. In the numerator of the RHS of (46), $f_\delta(\delta) = 1$ on the

support of δ , and otherwise is zero. $g_W((1 + \delta)c) = 1$ iff $(1 + \delta)c \in [1 - K, 2 + K]$, and otherwise is zero.

The condition that $W \geq 1 - K$ implies that $(1 + \delta)c \geq 1 - K$, so

$$\delta \geq \max\left(0, \frac{1 - K}{c} - 1\right) \stackrel{def}{=} \underline{\delta}. \quad (47)$$

The condition that $W \leq 2 + K$ implies that $(1 + \delta)c \leq 2 + K$, so

$$\delta \leq \min\left(\frac{2 + K}{c} - 1, 1\right) \stackrel{def}{=} \bar{\delta}. \quad (48)$$

We can therefore calculate the expected discount factor as perceived by an observer who believes he has observed another individual with consumption c , as

$$E[\delta|c] = \frac{\int_{\underline{\delta} \leq \delta \leq \bar{\delta}} \delta(1 + \delta)d\delta}{\int_{\underline{\delta} \leq \delta \leq \bar{\delta}} (1 + \delta)d\delta}. \quad (49)$$

In this calculation, the $1 + \delta$ terms come from the fact that the g density is a function of $(1 + \delta)c$.

We consider three cases.

Case 1: $0 < c < 1 - K$.

Then the range of the integrals becomes $\delta \in [\frac{1-K}{c} - 1, 1]$, so

$$\begin{aligned} E[\delta|c] &= \frac{\int_{\frac{1-K}{c}-1}^1 \delta(1 + \delta)d\delta}{\int_{\frac{1-K}{c}-1}^1 (1 + \delta)d\delta} \\ &= \frac{1}{3} \left(1 + \frac{2(1 - K)^2}{2c^2 + (1 - K)c}\right). \end{aligned} \quad (50)$$

It is evident from this expression that inferred time preference parameter decreases with observed consumption of others. We will compare the amount of updating about others' time preference based on observed consumption c as measured by the sensitivity of inferred time preference parameter to c , $dE[\delta|c]/dc$, when there is wealth dispersion to the benchmark case when there is no wealth dispersion. It is straightforward to show that the absolute value of this sensitivity $dE[\delta|c]/dc$ decreases with the amount of wealth dispersion K , and

$$\left.\frac{dE[\delta|c]}{dc}\right|_{K=0} = -\frac{2(4c + 1)}{3c^2(2c + 1)^2}. \quad (51)$$

Case 2: $1 - K < c < 1 + \frac{K}{2}$, which implies that $\underline{\delta} = 0$ and $\bar{\delta} = 1$. In this case, $E[\delta|c] = \frac{5}{9}$, so the inferred discount rate of the others does not vary with the observed consumption activities.

Case 3: $1 + \frac{K}{2} < c < 2 + K$, which implies $\bar{\delta} = \frac{2+K}{c} - 1$, $\underline{\delta} = 0$, so

$$\begin{aligned} E[\delta|c] &= \frac{\int_0^{\frac{2+K}{c}-1} \delta(1+\delta)d\delta}{\int_0^{\frac{2+K}{c}-1} (1+\delta)d\delta} \\ &= \frac{8 + 8K + 2K^2}{6c + 3c^2} - \frac{1}{3}. \end{aligned} \quad (52)$$

Again, the inferred time preference parameter $E[\delta|c]$ decreases with observed consumption of others, the absolute value of the sensitivity $dE[\delta|c]/dc$ decreases with the amount of wealth dispersion K , and

$$\left. \frac{dE[\delta|c]}{dc} \right|_{K=0} = -\frac{16(1+c)}{3c^2(c+2)^2}. \quad (53)$$

As a benchmark for comparison, suppose that there is no wealth dispersion, and that the known level of wealth \bar{W} is equal to the expected value of wealth in the model with wealth dispersion, $\bar{W} = 1.5$. In the model without wealth dispersion, by (4), the inferred value of δ is

$$\hat{\delta}(c) = \frac{\bar{W}}{c} - 1 = \frac{3}{2c} - 1, \quad (54)$$

so

$$\hat{\delta}'(c) = -\frac{3}{2}c^{-2}. \quad (55)$$

It is straightforward to show by direct comparison that in all three cases,

$$\hat{\delta}'(c) < \frac{dE[\delta|c]}{dc} \leq 0. \quad (56)$$

In other words, with wealth dispersion, the discount factor that the observer perceives about the target of observation does not decrease as rapidly with perceived target consumption as in the model without wealth dispersion.

Proposition 5 *In the numerical example of Section 6, comparing the setting with wealth dispersion with a setting with constant wealth equal to the expected wealth in the uncertainty setting, an observer's perception of the target's discount factor is less sensitive to perceived target consumption than when there is no wealth dispersion.*

Empirically, Proposition 5 predicts that savings rates increase with wealth dispersion. This is the opposite of what is expected based upon Veblen wealth-signaling considerations. 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). Greater information asymmetry about wealth intensifies the effect, by increasing the potential improvement in wealth perceptions that can be achieved by signaling. This is reflected, for example, in the finding of Charles, Hurst, and Roussanov (2009) that an increase in wealth dispersion that takes the form of a reduction in the lower support of wealth results in greater signaling.¹⁵

For example, in the limiting case of no information asymmetry, the Veblen effect would disappear and people would consume only for their direct utility benefits. More generally, in a simple setting in which the upper bound of the support of the wealth distribution becomes higher, then the range of possible equilibrium consumption signal levels is higher, so there will be more overconsumption on average.

Empirically, Jin, Li, and Wu (2011) find that greater income inequality is associated with lower consumption and with greater investment in education, in survey evidence from Chinese urban households, 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 the visibility bias approach, in contrast with the implication of wealth-signaling via consumption, or with the idea that low income individuals borrow and consume more in order to try to keep up with high income households.¹⁶

¹⁵A more detailed intuition is that when dispersion increases, the lower support of the wealth distribution decreases, so the signalling schedule starts increasing from an earlier beginning (a lower wealth level), raising the signaling schedule for any given wealth level. In their Veblen-style model, Charles, Hurst, and Roussanov (2009) show that an increase in the dispersion of the wealth distribution that derives from a reduction in the lower support (making the poorest poorer) causes greater conspicuous consumption. Their explanation is essentially the same: “The intuition is that as poorer people are added to a population, persons of every level of income must now signal more to distinguish themselves from those immediately poorer, because those people are themselves now compelled to spend more to distinguish themselves from persons who are even poorer still.” However, Charles, Hurst, and Roussanov (2009) show that the effect of a more general increase in wealth dispersion is ambiguous. What we expect to apply quite generally in Veblen-style models is that wealth-signaling through consumption vanishes when wealth dispersion is zero. So such models reflect a general tendency for greater wealth dispersion to induce greater overconsumption, though not necessarily monotonically.

¹⁶The standard keeping-up-with-Jones utility where people care about their consumption relative to the average consumption leads to overconsumption, but the amount of overconsumption does not change with an increase in the dispersion of wealth among people as long as the average level of wealth in the economy is held constant. A variant of the keeping-up-with-Jones model where a given household’s consumption is

Our approach also implies contagion in consumption across individuals with different wealths. Using state-level survey data, Bertrand and Morse (2013) document that increased income of the top quintile or decile in the income distribution is associated with more consumption by less wealthy individuals. In our approach, when the wealths of others are imperfectly observable, such contagion in consumption occurs because individuals who see high consumption of others attribute this in part to a high discount rate rather than high wealth. The authors further find that non-rich households' consumption of more visible goods and services is especially responsive to increases in top-end wealth, consistent with visibility bias effects of the type that we model. However, contagion in consumption is also implied by some versions of the Veblen and keeping-up-with-the-Joneses approaches to overconsumption.

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 5 counterfactually implies a rising savings rate. However, all else was not equal. We believe that a much more important shift (from the standpoint of the effects in our model) has been the dramatic transformation of electronic communications and social networks. As discussed in Section 3.4, 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.

The effects of wealth dispersion here derive from the unobservability of others' wealths rather than dispersion per se. The model therefore predicts that when the wealth of neighbors is harder to observe directly, there is less overconsumption.

7 Concluding Remarks

We examine how social influence endogenously shapes time preferences. In our model, 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 (and a well-known manifestation of it, the availability heuristic), people infer that low saving is normative and increase their own discount rates accordingly. This effect is self-reinforcing at the social level, resulting

directly positively affected by the consumption of the households whose income is just above theirs can generate an expenditure cascade, resulting in a negative relationship between income inequality and the savings rate of non-rich households (Bertrand and Morse 2013).

in overconsumption and high interest rates. The model therefore results in pluralistic ignorance about the time preferences of others; each person thinks that the others have a higher discount rate. This ignorance in equilibrium has adverse consequences for resource allocation.

The visibility bias approach offers a simple explanation for one of the most puzzling and important stylized facts about 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 the decline in the last thirty years in costs of long-distance telephone service, and subsequently the rise of the internet, dramatically increased the extent to which people convey information about personal consumption activities to others by phone, email, blogging, and social networking. This is because consumption activities such as 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 and observability, such as urban versus rural, and survey questions about sociability (see, e.g., Hong, Kubik, and Stein (2004), Brown et al. (2008), Christelis, Georgarakos, and Haliassos (2011), and Georgarakos and Pasini (2011)). Our approach also offers predictions about how population-level characteristics such as wealth variance affect consumption.

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 greater overconsumption when there is greater information asymmetry about the wealth of others (as occurs with high wealth dispersion). In our setting greater information asymmetry dilutes the inference from high observed consumption that the discount rate of others is high. 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 inherent discount rates can be amplified through social influence.

In contrast with the signaling and preference-based approaches, our biased-belief-based approach implies that a cheap and easy policy intervention will substantially increase saving. This is to provide—in highly salient form—accurate information about how much peers save, or their attitudes toward consumption. This implication about the beneficial

effects of salient disclosure could be tested in the field. Indeed, there is evidence discussed earlier supporting this implication in a specialized setting, the decision of college students of how much to drink.

In addition to providing insights about undersaving in general, the social interaction approach to consumption/saving norms potentially can contribute to our understanding of the dynamics of consumption as well. For example, suppose that there are lags between people observing others and updating their own consumption plans, or lags between their consumption plans and actual consumption. If there are shocks to the system that encourage high or low spending, then the response lags can create momentum in shifts in consumption and consumption norms. This can potentially cause patterns of overshooting and correction, which would provide a possible basis for an overconsumption theory of local consumption booms and business cycles.

More generally, our approach to understanding the evolution of consumption and saving attitudes is based upon *misperception of norms*—pluralistic ignorance. This approach is potentially applicable to various other settings, which suggests a rich direction for future research.

Appendix

Comparative Statics on Varying Endowments

Omitting i subscripts, let y_0 and y_1 be dates 0 and 1 endowed income. We now consider the comparative statics on varying y_0 , the individual's date 0 income endowment, in settings with and without visibility bias.

Proposition 6 *Under homothetic preferences, visibility bias increases the marginal propensity to consume from current income.*

This comparative statics is based upon an exogenous interest rate r , as marginal propensity to consume is defined based on variation of a single individual's income. To prove this comparative statics, we use individual's budget constraint (1), the optimality condition (2) and the property of homothetic preferences. Differentiating the first order condition (2) with respect to y_0 ,

$$u''(c_0) \frac{dc_0}{dy_0} = \delta(1+r)u''(c_1) \frac{dc_1}{dy_0}. \quad (57)$$

Differentiating the budget constraint (1) with respect to y_0 gives

$$\frac{dc_0}{dy_0} + \frac{dc_1}{dy_0} \frac{1}{1+r} = 1. \quad (58)$$

so

$$\frac{dc_1}{dy_0} = (1+r) \left(1 - \frac{dc_0}{dy_0} \right). \quad (59)$$

Substituting for dc_1/dy_0 into (57), gives

$$u''(c_0) \frac{dc_0}{dy_0} = \delta(1+r)^2 \left[u''(c_1) - u''(c_1) \frac{dc_0}{dy_0} \right]. \quad (60)$$

So

$$\frac{dc_0}{dy_0} = \frac{1}{\frac{u''(c_0)}{u''(c_1)\delta(1+r)^2} + 1}. \quad (61)$$

Under the homotheticity assumption, the marginal rate of substitution is constant on a ray through the origin, so

$$\frac{u'(c_1)}{u'(c_0)} = K \left(\frac{c_0}{c_1} \right), K > 0. \quad (62)$$

So

$$u'(c_1)c_1 = Ku'(c_0)c_0 \quad (63)$$

Partially differentiating with respect to c_1, c_0 gives

$$u''(c_1)c_1 + u'(c_1) = 0 \quad (64)$$

$$u''(c_0)c_0 + u'(c_0) = 0. \quad (65)$$

It follows from (62), (64) and (65) that

$$\frac{u''(c_1)}{u''(c_0)} = K \left(\frac{c_0}{c_1} \right)^2. \quad (66)$$

So

$$\frac{dc_0}{dy_0} = \frac{1}{\left(\frac{c_1}{c_0} \right)^2 \left[\frac{1}{K\delta(1+r)^2} + 1 \right]} > 0. \quad (67)$$

As for the social effect on the marginal propensity to consume, differentiating with respect to γ ,

$$\text{sign} \left(\frac{\frac{dc_0}{dy_0}}{d\gamma} \right) = \text{sign} \left(\frac{\left(\frac{c_0}{c_1} \right)^2}{d\gamma} \right) = \text{sign} \left(\frac{dc_0}{d\gamma} \right) > 0. \quad (68)$$

References

- Abel, A. B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* **80**, 38–42.
- Akerlof, G. A., 2007, The missing motivation in macroeconomics, *American Economic Review* **97**, 3–36.
- Akerlof, G. A. and R. J. Shiller, 2009, *Animal Spirits: How Human Psychology Drives the Economy and Why it Matters for Global Capitalism* (Princeton University Press, Princeton, NJ).
- Allen, T. W. and C. D. Carroll, 2001, Individual learning about consumption, *Macroeconomic Dynamics* **5**, 255–271.
- Bagwell, L. S. and B. D. Bernheim, 1996, Veblen effects in a theory of conspicuous consumption, *American Economic Review* **86**, 349–373.
- Barber, B. and T. Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* **55**, 773–806.
- Benhassine, N., F. Devoto, E. Duflo, and P. D. and Victor Pouliquen, 2013, Turning a shove into a nudge? A labeled cash transfer for education, NBER working paper No.19227, World Bank.
- Bertrand, M. and A. Morse, 2013, Trickle-down consumption, Working paper, University of Chicago.
- 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 University Press, Oxford).
- Beshears, J. L., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman, 2015, The effect of providing peer information on retirement savings decisions, *Journal of Finance* **70**, 1161–1201.
- Bosworth, B. P., 1993, *Saving and Investment in a Global Economy* (Brookings Institution Press, Washington, D.C.).
- Brenner, L. A., D. J. Koehler, and A. Tversky, 1996, On the evaluation of one-sided evidence, *Journal of Behavioral Decision Making* **9**, 59–70.
- Brown, J. R., Z. Ivkovich, P. A. Smith, and S. Weisbenner, 2008, Neighbors matter: Causal community effects and stock market participation, *Journal of Finance* **63**, 1509–1531.
- Cai, H., Y. Chen, and H. Fang, 2009, Observational learning: Evidence from a randomized natural field experiment, *American Economic Review* **99**, 864–882.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong, 2004, A cognitive hierarchy model of games, *Quarterly Journal of Economics* **119**, 861–898.

- Campbell, J. Y. and J. H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* **107**, 205–251.
- Carbone, E. and J. Duffy, 2013, Lifecycle consumption plans, social learning and external habits: Experimental evidence, Working paper, University of Pittsburgh.
- Card, D. and J. E. DiNardo, 2002, Skill-biased technological change and rising wage inequality: Some problems and puzzles, *Journal of Labor Economics* **20**, 733–783.
- Carroll, C. D., B.-K. Rhee, and C. Rhee, 1994, Are there cultural effects on saving? Some cross-sectional evidence, *Quarterly Journal of Economics* **109**, 685–699.
- Carroll, C. D., B.-K. Rhee, and C. Rhee, 1999, Does cultural origin affect saving behavior? Evidence from immigrants, *Economic Development and Cultural Change* **48**, 33–50.
- Charles, K. K., E. Hurst, and N. Roussanov, 2009, Conspicuous consumption and race, *Quarterly Journal of Economics* **124**, 425–467.
- Chen, Y., F. M. Harper, J. Konstan, and S. X. Li, 2010, Social comparisons and contributions to online communities: A field experiment on movielens, *The American Economic Review* **100**, 1358–1398.
- Christelis, D., D. Georgarakos, and M. Haliassos, 2011, Stockholding: Participation, location, and spillovers, *Journal of Banking & Finance* **35**, 1918–1930.
- Cialdini, R., L. Demaine, B. Sagarin, D. Barrett, K. Rhoads, and P. Winter, 2006, Managing social norms for persuasive impact, *Social Influence* **1**, 3–15.
- Coibion, O., Y. Gorodnichenko, M. Kudlyak, and J. Mondragon, 2014, 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, Incorporating concern for relative wealth into economic models, *Quarterly Review, Federal Reserve Bank of Minneapolis* **19**, 12–21.
- Corneo, G. and O. Jeanne, 1997, Conspicuous consumption, snobbism and conformism, *Journal of Public Economics* **66**, 55–71.
- DeMarzo, P., R. Kaniel, and I. Kremer, 2008, Relative wealth concerns and financial bubbles, *Review of Financial Studies* **21**, 19–50.
- Demarzo, P. M., R. Kaniel, and I. Kremer, 2004, Diversification as a public good: Community effects in portfolio choice, *Journal of Finance* **59**, 1677–1716.
- Duflo, E. and E. Saez, 2002, Participation and investment decisions in a retirement plan: The influence of colleagues' choices, *Journal of Public Economics* **85**, 121–148.
- Duflo, E. and E. Saez, 2003, The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment, *Quarterly Journal of Economics* **118**, 815–842.
- Dupor, B. and W.-F. Liu, 1993, Jealousy and equilibrium overconsumption, *The American Economic Review* **93**, 423–428.

- Eyster, E. and M. Rabin, 2005, Cursed equilibrium, *Econometrica* **73**, 1623–1672.
- Fiedler, K., 2008, The ultimate sampling dilemma in experience-based decision making, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **34**, 186–203.
- Frederick, S., 2012, Overestimating others’ willingness to pay, *Journal of Consumer Research* **39**, 1–21.
- Frey, B. and S. Meier, 2004, Social comparisons and pro-social behavior: Testing “conditional cooperation” in a field experiment, *American Economic Review* **94**, 1717–1722.
- Galí, J., 1994, Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices, *Journal of Money, Credit and Banking* **26**, 1–8.
- Georgarakos, D., M. Haliassos, and G. Pasini, 2014, Household debt and social interactions, *Review of Financial Studies* **27**, 1404–1433.
- Georgarakos, D. and G. Pasini, 2011, Trust, sociability, and stock market participation, *Review of Finance* **15**, 693–725.
- Gerber, A. S. and T. Rogers, 2009, Descriptive social norms and motivation to vote: Everybody’s voting and so should you, *The Journal of Politics* **71**, 178–191.
- Goldstein, N. J., R. B. Cialdini, and V. Griskevicius, 2008, A room with a viewpoint: Using social norms to motivate environmental conservation in hotels, *Journal of Consumer Research* **35**, 472–482.
- Guidolin, M. and E. A. L. Jeunesse, 2007, 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**, 491–514.
- Hall, R. E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* **96**, 339–357.
- Healy, A. F., 1981, The effects of visual similarity on proofreading for misspellings, *Memory and Cognition* **9**, 453–460.
- Hearst, E., 1991, Psychology and nothing, *American Scientist* **79**, 432–443.
- Heffetz, O., 2011, A test of conspicuous consumption: Visibility and income elasticities, *Review of Economics and Statistics* **93**, 1101–1117.
- Hirsh, J. B., 2015, Extraverted populations have lower savings rates, *Personality and Individual Differences* **81**, 162–168.
- Hirshleifer, D. and S. H. Teoh, 2003, Limited attention, information disclosure, and financial reporting, *Journal of Accounting & Economics* **36**, 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. (North-Holland, Amsterdam, The Netherlands).
- Hong, H. G., J. D. Kubik, and J. C. Stein, 2004, Social interaction and stock market participation, *Journal of Finance* **59**, 137–163.
- Jehiel, P., 2015, Investment strategy and selection bias: An equilibrium perspective on overconfidence, Working paper, Ecole des Ponts and University College London.

- Jin, Y., H. Li, and B. Wu, 2011, Income inequality, consumption, and social-status seeking, *Journal of Comparative Economics* **39**, 191–204.
- Kahneman, D. and A. Tversky, 1972, Subjective probability: A judgment of representativeness, *Cognitive Psychology* **3**, 430–454.
- Kahneman, D. and A. Tversky, 1973, On the psychology of prediction, *Psychological Review* **80**, 237–251.
- Katz, D. and F. H. Allport, 1931, *Students' Attitudes: A Report of the Syracuse University Reaction Study* (Craftsman Press, Syracuse, NY).
- Kaustia, M. and S. Knüpfer, 2012, Peer performance and stock market entry, *Journal of Financial Economics* **104**, 321–338.
- Koehler, J. J. and M. Mercer, 2009, Selection neglect in mutual fund advertisements, *Management Science* **55**, 1107–1121.
- Laibson, D., 1997, Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics* **112**, 443–477.
- Lambert, T. A., A. S. Kahn, and K. J. Apple, 2003, Pluralistic ignorance and hooking up, *Journal of Sex Research* **40**, 129–133.
- Lemieux, T., 2006, Postsecondary education and increasing wage inequality, *American Economic Review* **96**, 195–199.
- Loayza, N., K. Schmidt-Hebbel, and L. Serven, 2000, Saving in developing countries: An overview, *World Bank Economic Review* **14**, 393–414.
- Madrian, B. and D. Shea, 2001, The power of suggestion: Inertia in 401(k) participation and savings behavior, *Quarterly Journal of Economics* **116**, 1149–1187.
- McShane, B. B., E. T. Bradlow, and J. Berger, 2012, Visual influence and social groups, *Journal of Marketing Research* **49**, 854–871.
- Neisser, U., 1963, Decision-time without reaction-time: Experiments in visual scanning, *American Journal of Psychology* **76**, 376–385.
- Nisbett, R. and L. Ross, 1980, *Human Inference: Strategies and Shortcomings of Social Judgment* (Prentice-Hall, Englewood Cliffs, NJ).
- OECD, 2014, *Oecd factbook 2014: Economic, environmental and social statistics*, *OECD Publishing* **89**.
- Parker, J. A., 1999, Spendthrift in America? On two decades of decline in the U.S. saving rate, *NBER Macroeconomics Annual* **14**, 317–370.
- Parlour, C. A., V. Prasnikař, and U. Rajan, 2007, Compensating for the winner's curse: Experimental evidence, *Games and Economic Behavior* **60**, 339–356.
- Piketty, T. and E. Saez, 2003, Income inequality in the United States, 1913–1998, *Quarterly Journal of Economics* **118**, 1–39.
- Prentice, D. A. and D. T. Miller, 1993, Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm, *Journal of Personality and Social Psychology* **64**, 243–256.

- Ritov, I. and J. Baron, 1990, Reluctance to vaccinate: Omission bias and ambiguity, *Journal of Behavioral Decision Making* **3**, 263–277.
- Salganik, M. J., P. S. Dodds, and D. J. Watts, 2006, Experimental study of inequality and unpredictability in an artificial cultural market, *Science* **311**, 854–856.
- Samuelson, W. and R. Zeckhauser, 1988, Status-quo bias in decision making, *Journal of Risk and Uncertainty* **1**, 7–59.
- Shefrin, H. M. and R. H. Thaler, 1988, The behavioral life-cycle hypothesis, *Economic Inquiry* **26**, 609–643.
- Simonson, I., 2008, Will I like a medium pillow? Another look at constructed and inherent preferences, *Journal of Consumer Psychology* **18**, 155–169.
- Stracca, L. and A. al Nowaihi, 2005,. Working paper series, . Technical Report 444, European Central Bank, United Kingdom.
- Weber, W. E., 1970, The effect of interest rates on aggregate consumption, *American Economic Review* **60**, 591–600.