Provided for non-commercial research and education use. Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Games and Economic Behavior 69 (2010) 458-474

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Why income comparison is rational

David H. Wolpert

NASA Ames Research Center, Mail Stop 269-1, Moffett Field, CA 94035-1000, United States

ARTICLE INFO

Article history: Received 11 February 2009 Available online 5 December 2009

JEL classification: D03 D81 D82 D83

C44 Keywords:

Social comparison Hedonic utility Easterlin paradox U-shaped lifetime curve Keep up with the Joneses Herding Income comparison Search algorithm Decision theory

ABSTRACT

A major factor affecting a person's happiness is the gap between their income and their neighbors', independent of their own income. This effect is strongest when the neighbor has moderately higher income. In addition a person's lifetime happiness often follows a "U" shape. Previous models have explained subsets of these phenomena, typically assuming the person has limited ability to assess their own (hedonic) utility. Here I present a model that explains all the phenomena, without such assumptions. In this model greater income of your neighbor is statistical data that, if carefully analyzed, would recommend that you explore for a new income-generating strategy. This explains unhappiness that your neighbor has greater income, as an emotional "prod" that induces you to explore, in accord with careful statistical analysis. It explains the "U" shape of happiness similarly. Another benefit of this model is that it makes many falsifiable predictions.

Published by Elsevier Inc.

I do count my blessings, but then I count those of others who have more, and that pisses me off. R. Mankoff, The New Yorker, Nov. 26, 2007

1. Background

There appear to be many aspects to "well-being," "happiness" or "hedonic utility" (Reichhardt, 2006), for example as reflected in self-reports (Schwarz and Strack, 1999) like the general social survey, and as reflected in the U-index (Kahneman and Krueger, 2006), based on the day reconstruction method (Kahneman et al., 2004). Some even suggest that happiness intrinsically has multiple components (Diener et al., 2006; Kahneman et al., 2006). Moreover, even the most basic empirical results concerning hedonic utility, like whether there is a hedonic setpoint, a hedonic treadmill, etc., are not fully resolved (Easterlin, 2003).

Despite this lack of consensus concerning what "happiness" means and what the associated experimental data says, it *is* widely accepted that, everything else being equal, people are less happy if their neighbors have slightly greater income than they do (Duesenberry, 1949; Luttmer, 2005; Easterlin, 1995; Neumark and Postlewaite, 1998; Stevenson and Wolfers, 2008; Frank, 1999; Frey and Stutzer, 2002; Fehr and Schmidt, 1999; Easterlin, 2003; Easterlin and Plangno, 2008; D'Ambrosio and Frick, 2007). Indeed, John Stuart Mill went so far as to state that "Men do not desire to be rich, but richer than other men."





E-mail address: david.h.wolpert@nasa.gov.

^{0899-8256/\$ -} see front matter Published by Elsevier Inc. doi:10.1016/j.geb.2009.12.001

Marx then commented that "Our wants and pleasures ... (are measured) in relation to society; they are of a relative nature" (Frey and Stutzer, 2005).

Things aren't actually this extreme: it is now understood that both absolute income and relative income affect a person's happiness, as do many other factors (Andersson, 2008; Kahneman et al., 2006). Nonetheless, it is peculiar that relative income matters at all. By definition the income of a person's neighbor has no effect on what material goods that person can consume. This would seem to imply that the income of a person's neighbor should not affect how happy a person is. And indeed, it appears that if the income of a neighbor of a person *i* is smaller than that of *i* (or far greater), then the effect of that neighbor's income on *i*'s happiness is not very significant (Duesenberry, 1949; Andersson, 2008; Ferrer-i Cabonell, 2005). But the income comparison effect can be important for intermediate gaps in income. Moreover, this comparison effect is found in very many cultures. It even seems to arise in fMRI studies (Zink et al., 2008; Takahashi et al., 2009).

The income comparison effect can have substantial effects on behavior. For example, people will often choose lower income – and therefore have less access to material goods – if that guarantees that their neighbors will have even lower income than they do (Solnick and Hemenway, 1998; Luttmer, 2005; Tversky and Griffin, 1991). As another example, the "Easterlin paradox" is the fact that above a certain income threshold, the happinesses of a typical inhabitant of a country seems to only be partly determined by average national income, depending as well on how their income compares to that average national income (Easterlin, 2003, 1995; Stevenson and Wolfers, 2008; Gardes and Merrigan, 2008). This paradox can be partly explained as resulting from the income comparison effect: under that effect the typical inhabitant of a country is made unhappy by having income lower than that of their neighbors, and tautologically, the income of those neighbors contributes to the national average income. Another example is that people often strive to "keep up with the Joneses" (Duesenberry, 1949; Pollak, 1976; Abel, 1990; Campbell and Cochrane, 1990), to remove (or at least reduce) the difference between their wealth displays and that of their neighbors. Again, this can be partly explained by the income comparison effect, if one approximates ability to engage in wealth displays with income. There are many more examples; so pervasive is the income comparison effect that some have even argued that it should be used to determine governmental economic policy (Abel, 2005; Frey and Stutzer, 2007; Boskin and Sheshinski, 1978).

These facts suggest that the income comparison effect is somehow rational, despite the fact that relative income of a neighbor has no effect on an individual's current access to material goods. If it is indeed rational for a person to be made unhappy by the relative income of a neighbor, presumably that is due to the effect of that unhappiness on the person's future behavior. This in turn implies that the change in that person's future behavior induced by the unhappiness must improve their expected future utility.

To explore this hypothesis, we must model how the income gap between an individual and their neighbor is statistically related to the expected utility that individual would receive for any of their possible future behaviors. Such a model would complement work in which concern for relative income (known as "interdependent preferences") is introduced into the utility function in an ad hoc manner, with no conjecture about *why* individuals care about relative income (Clark et al., 2008). Ideally, such a model would even provide insight into some seemingly unrelated peculiarities concerning happiness. For example, ideally such a model would provide insight into the fact that the happiness of people follows a "U-shaped" curve as they age, first shrinking, and then ultimately, as they grow older, rising (Yang, 2008; Blanchflower and Oswald, in press; Clark et al., 1996).

2. Contribution of this paper

Here I consider the difference between the income of an individual i and the income of a neighbor of i. I introduce a statistical model relating that difference to the income that i would likely receive if they changed their current incomegenerating strategy. Under the model, if an individual i observes a neighbor with slightly greater income, then i's optimal behavior is to explore for a new, better income-generating strategy. On the other hand, if the neighbor that i observes has less income than i, then that observation does not recommend that i explore more. Similarly if i observes a neighbor with far greater income, then that observation does not recommend that i explore more. (Loosely speaking, this is because the income of such a neighbor with far greater income is probably governed by a different function mapping income-generating strategies to income than the strategy-to-income function that governs i.)

Roughly speaking, the model maps the incomes of someone and their neighbor to a probability distribution of the maximal income that person could readily achieve under *some* income-generating strategy. If that person is rational (!), this information may then change the person's aspirations. In turn, changes in the person's aspirations change whether they are unhappy with their current strategy (Frey and Stutzer, 2002; Easterlin, 2003). Such unhappiness may then prod them to change their behavior, to try to achieve their new aspirations.

In this way unhappiness induced by observing that neighbors have slightly greater income causes the individual to follow the strategy that is optimal given those observations. We can thus view unhappiness as a computational shortcut for the full calculation of their optimal strategy. It is a "hedonic prod" that induces a change in behavior that is likely to lead to greater income whenever an appropriate income gap is observed.

More broadly, the basic premise of this paper has two components:

1. The utility difference between a person and their neighbor provides statistical information about how their expected utility would change if they modify their behavior;

2. The function of the emotion of "unhappiness" is to induce someone to modify their behavior to improve expected utility.

The contribution of this paper is to build a model that combines the two parts of this premise, and then compare the predictions of that model to experimental data. This model makes many falsifiable predictions. For example, it predicts that the more effort either a person and/or their neighbor has exerted to achieve their income(s), the less effect a given income gap between a person and their neighbor has on the person's happiness.

This model is similar to recent social psychological models that show how observing others may rationally lead a person to engage in seemingly irrational behavior like conformity (Denrell, 2008). More generally, it bears similarities to social learning models involving herding and imitation (Banerjee, 1992; Bikhchandani et al., 1992; Schlag, 1999; Apesteguia et al., 2007; Vega-Redondo, 1997; Selten and Ostmann, 2001; Henrich and Boyd, 2001; Boyd et al., 2003; Fudenberg and Levine, 1998). However these other models directly concern what decision a person should make for given social comparison data. In contrast, in my model the relation between a person's decision and social comparison data is indirect: the data is first reduced to a change in the decision maker's happiness, and that changed happiness then affects his decision-making. Another distinction is that these other models analyze situations where the decision maker not only observes the income of his neighbor but also observes his neighbor's actual choice. Such obtrusive observation is often absent in the real world, and is not assumed in my model.

There are other models that, like the one introduced here, concern the optimal relation between social comparison data and happiness (Rayo and Becker, 2007; Robson, 2001; Netzer, in press), rather than consider the optimal relation directly between the data and decision-making. However these other models often make the assumption that "happiness" can only take on the values in some bounded set *S*, while there is no experimental data validating this assumption. (Note in particular that the carriers of information in the human brain are inter-neuron spike-train patterns, and there are an uncountable number of such patterns.) They further assume the decision maker has observational limitations. In particular, these models assume there is stochastic "imprecision" in a decision maker's perception of *his own* happiness, imprecision that results in having *S* be finite (as in Robson, 2001) or results in bounded rationality (as in Rayo and Becker, 2007).

In contrast, in the model presented here, the only stochasticity is due to uncertainty in the environment. There is no assumption of observational limitations in the decision maker, and in particular no assumptions that the decision maker cannot observe his own internal state.¹ Another contrast is that the model presented here explains all characteristics of the income comparison effect mentioned above – including aspects of the U-shape followed by happiness as one ages – not just some of those characteristics.

There is also some work that models social comparison by modeling how other members of society will change their interaction with a person based on that person's relative income (Postlewaite et al., 1995). This work can be viewed as a strategic analysis of social comparison. In contrast, the model here is not strategic in any sense; it shows how social comparison can be rational even in a decision-theoretic context, without concern for how other people will treat you based on such social comparison. (However see Section 7.1 below for a cursory discussion of possible extensions of the model of this paper to strategic scenarios.)

Finally, there is some work that models the relationship between social comparison and "happiness" in terms of complicated multi-generational evolutionary processes (Samuelson, 2004). Typically these models require that the decision maker be bounded rational, and know something about what strategy their neighbor adopted, together with knowing other population-level information.

In contrast, the model presented here explains income comparison by considering only a single decision, by a single decision maker. Under this model the explanation for the hedonic prod does not invoke reproductive fitness; it arises from considerations of the decision maker's own future utility, not the utility of their progeny. Moreover the model presented here does not require that the decision maker have population-level information, or that they be bounded rational.

By presuming full rationality, the model presented here is more in keeping with the standard considerations of *homo economicus* game theory than models like that of Samuelson (2004). In addition, the model presented here is substantially simpler than the typical model based on evolutionary processes. This means that the model presented here is easier to analyze. It also means that the model can be used to make many predictions for experiments and surveys that have not yet been done but can be readily tested, as elaborated below. In contrast, it can be quite hard to make such predictions using models based on evolutionary processes. The simplicity of the model presented here also means that it is easier to extend it to other situations involving hedonic utility than are models based on evolutionary processes. (See Section 7 below for a discussion of some such extensions.)

As another point of contrast, the model presented here does not implicitly hypothesize that there are genes on the human chromosome that can promote income comparison – genes whose existence has never been confirmed in the laboratory. Nor does the model involve suppositions considering Pleistocene societies, societies which we cannot observe, and which may well differ in many important (and not currently known) respects from all societies that we can observe. In addition, the model can be tested directly without needing to wait for many successive generations of humans to breed. So the model

¹ The only assumption in the model presented here concerning cognitive limitations of the decision maker is the implicit assumption that there is some cost to computation, which results in the use of a hedonic computational shortcut. However this assumption concerns a person's computation rather than their observation. Moreover, the resultant computational shortcut still results in the correct decision. In this sense, it is not a "limitation."

is experimentally refutable in all respects, in contrast to models based on evolutionary processes that take place over many generations of societies we cannot observe. A final important distinction is that the model presented here does not require that the decision maker be able to observe their neighbor's strategy, in which sense its starting point more closely matches many of the income comparison scenarios about which we have experimental data.

Despite the foregoing, reproductive fitness issues clearly are important in human behavior. They undoubtedly play a role – perhaps a major one – in income comparison. The point of this paper is that they do not appear to be necessary to explain income comparison; much can be understood without their complexity and experimental difficulties.

In Section 3 I present an overview of my model. In Section 4 I provide the fully detailed specification of the model. In Section 5 I consider the case where the decision maker compares themselves to the "Vanderbilts" (those with far higher income) and to the "Bundys" (those with less income). Section 6 considers the case where the decision maker compares themselves to the "Joneses" (those with comparable income). After this, in Section 7 I discuss the connection between the model and the U-shape of lifetime happiness curves, and between the model and ordinal (rather than cardinal) social comparison. Section 8 comprises a general discussion of the model. A list of some of the model's testable predictions are presented in this section. This section also analyzes some of the model's simplifications and how they might be removed in a more sophisticated analysis. This section also presents a hypothesis inspired by the model for why utility and hedonic utility differ and how they are related. This section ends by discussing how the model might fit into a broader theory of the rational basis of human emotions.

3. Model synopsis

I consider a Decision maker (**D**) who must decide what strategy to adopt to maximize income, e.g., what job to adopt to maximize income. To do this **D** will use whatever information he has about the "environment" function mapping his strategies to associated incomes.² For example, **D** will use his knowledge of what incomes he received for previously chosen strategies. However such information that **D** has about the environment function is limited. Accordingly, he faces a search problem involving an exploration–exploitation tradeoff, similar to the tradeoffs faced by a species undergoing natural selection, by a decision maker in a multi-armed bandits scenario (Macready and Wolpert, 1998) and by blackbox computational search algorithms (Back et al., 1997; Kirkpatrick et al., 1983).

Say **D** has a neighbor who is searching the same environment function as **D** is. The income of that neighbor provides information to **D** concerning whether he should explore for a new strategy or instead exploit his current strategy. In particular, as elaborated below, if **D**'s neighbor is searching the same environment function and has moderately higher income, then a detailed statistical calculation shows that **D** should explore rather than exploit. On the other hand, if **D**'s neighbor is searching an environment function that differs from **D**'s, then the income of that neighbor provides no information to **D**.

Accordingly, in his detailed statistical calculation of how to use information from his neighbor's search, **D** has to assess the probability that his neighbor's environment function differs from **D**'s. In particular, if **D**'s income is far less than that of his neighbor, then it is likely that they are searching different functions. So if the income gap is very large, that gap does not provide information that **D** should explore.

Now consider real-world behavior of people rather than their idealized behavior. Due to the income-comparison phenomenon, sometimes in the real world a decision maker **D** is unhappy to see that his neighbor has somewhat higher income. Such unhappiness will in turn prod **D** to search for a new strategy. But such a search is exactly what the detailed statistical calculation tells **D** to do, if it is the case both that the neighbor has moderately higher income and that the data indicates that the neighbor's environment function is the same as **D**'s. The premise of this paper is that when it arises in the real-world, unhappiness by **D** due to income comparison is just a computational shortcut, one that causes **D** to follow the "change-your-strategy" advice of that detailed statistical calculation without consciously carrying out that calculation. Such a shortcut is a heuristic, in essence (Hart, 2005). Under this premise, insofar as the income comparison phenomenon makes **D** unhappy to have a neighbor with somewhat higher income, but does not make him unhappy to have a neighbor with a statistical calculation.³

Note that this model considers *un*happiness only. In this, its explanation of experimental data concerning happiness is indirect. In an informal form, such an indirect explanation of the experimental data goes back at least to Frank (2000).

The possibility that income comparison may be partly explained as a hedonic prod suggests that many other kinds of social comparison might also be partly explained this way. In particular, in some scenarios "strategies" x concern how you go about improving some quantity, so that f(x) concerns a rate of improvement. As an example, the hedonic prod model can be applied to an athlete, telling him how to use the rate of performance improvement of near-peers to help decide whether he should try a new coach, a new training regimen, or some such. In this example, x is the coach or playing regimen, and f(x) is how fast **D**'s performance is improving.⁴ Accordingly, from now on I will consider the "payoff" that a

 $^{^2}$ To avoid awkward grammar, I refer to **D** as if they are male.

³ No claim is made here as to whether a hedonic prod shortcut is learned by an individual during their lifetime, or is instead encoded in their genes. The point rather is that a hedonic prod shortcut can be advantageous for an individual, without regard to its effects on their mating ability and/or progeny.

⁴ Note that in such situations, *improving performance by* D *is built in*, and the prod does not induce improving performance, but rather an increase in the rate of that improving performance.

decision maker receives for making a particular choice; the income that a decision maker receives for a particular incomegenerating strategy is just a special case. Indeed, the "payoff" associated with a hedonic prod doesn't even have to be readily quantifiable based on field data, as income is. For example, we would expect a hedonic prod if someone is exposed to a neighbor who is more self-assured than they are.

How strong hedonic prods should be to result in optimal behavior will vary depending on the precise scenario in which the prod arises. Loosely speaking, the more an individual's payoff depends on their choice of strategy in a given scenario, the stronger the hedonic prod effect should be in that scenario. In general, that dependence will vary both with the type of payoff (pecuniary, non-pecuniary, etc.) and with the space of possible strategies. This variability agrees with the fact that the importance of social comparison to determining happiness depends on what is being compared ("positional" vs. "non-positional" goods).

4. Model details

I formalize the foregoing with a model that is as simple as possible while still capturing these effects. Let X be the set of all possible strategies that **D** can adopt, and $f_{\mathbf{D}}(x)$ the environment function mapping **D**'s choice of strategy $x \in X$ to an associated payoff. Due to uncertainty about his environment, **D** must treat $f_{\mathbf{D}}(x)$ as a random variable. At some present time, **D** has $N_{\mathbf{D}}$ samples of $f_{\mathbf{D}}$, giving a data set $L_{\mathbf{D}}$ of pairs of X values and associated values of $f_{\mathbf{D}}$: $\{L_{\mathbf{D}}^{X}(i), L_{\mathbf{D}}^{Y}(i) = f_{\mathbf{D}}[L_{\mathbf{D}}^{X}(i)]: i = 1, \ldots, N_{\mathbf{D}}\}$. These represent the set of payoff-generating strategies that **D** has explored during his lifetime to date. Define the best such value as $L_{\mathbf{D}}^{\max} \equiv \max_{i \in 1, \ldots, N_{\mathbf{D}}}[L_{\mathbf{D}}^{Y}(i)]$, and define $L_{\mathbf{D}}^{\min}$ similarly, as the worst value sampled so far. So $L_{\mathbf{D}}^{\max}$ is the best payoff that $L_{\mathbf{D}}^{\max}$ knows how to get.

The data set $L_{\mathbf{D}}$ provides some statistical information to **D** about $f_{\mathbf{D}}$. To model this information, for simplicity **D** assumes that the search algorithm he used to generate L^X , combined with the stochastic nature of how $f_{\mathbf{D}}$ was generated, means that the formation of $L_{\mathbf{D}}$ can be modeled as a set of $N_{\mathbf{D}}$ Independent and Identically Distributed (IID) samples of a single "window" density function having width σ and centered at some value $\mu_{\mathbf{D}}$:

$$P(L_{\mathbf{D}}|\mu_{\mathbf{D}}) = \prod_{i=1}^{N_{\mathbf{D}}} \frac{W_{\sigma}(\mu_{\mathbf{D}}, L_{\mathbf{D}}^{Y}(i))}{\sigma}$$
(1)

where

$$W_{\sigma}(\mu_{\mathbf{D}}, y) = 1 \quad \text{if } |y - \mu_{\mathbf{D}}| \leq \sigma/2$$

= 0 otherwise.

(2)

As an example, this model would apply if **D**'s search algorithm were purely random search over *X* and $f_{\mathbf{D}}$ were a fixed function whose range is $[\mu_{\mathbf{D}} - \sigma/2, \mu_{\mathbf{D}} + \sigma/2]$, with each value in that range occurring just as frequently as every other value.⁵

D knows σ but is uncertain about $\mu_{\mathbf{D}}$. **D** models this uncertainty by assuming that $\mu_{\mathbf{D}}$ is itself a random variable, generated before the search began by sampling some associated probability density function *h*. (**D** can view $\mu_{\mathbf{D}}$ as a stochastic State of Nature, so that *h* runs over possible States of Nature.)

Further assume that for any time at least ΔT after a making a sample, **D** can return to the best of his previously seen samples, but he cannot return before then. (This is a crude model of "the time needed to back out of a sampled move.") Also assume we're considering the decision of **D** at a time at least ΔT since his most recent sample, i.e., $L_{\mathbf{D}}$ consists of samples all of which are at least ΔT ago. So assuming he is rational, the current payoff of **D** is $L_{\mathbf{D}}^{\text{max}}$.

On the other hand, if **D** decides to make a new sample, *x*, he will be forced to endure the associated value $f_{\mathbf{D}}(x)$ for ΔT before he can return to $L_{\mathbf{D}}^{\max}$. If we assume that **D** has a discounting rate that is large on the scale of ΔT , this means that in deciding whether to form a new sample, **D** should act as though he is only concerned with how the payoff of that new sample would likely compare to his current payoff, $L_{\mathbf{D}}^{\max}$.

The Neighbor (**N**) is another decision maker facing the same kind of search problem as **D**. **N**'s space of possible moves is also *X*, and his payoff is given by a function $f_{\mathbf{N}}(x)$, which may or may not be the same as $f_{\mathbf{D}}$ (see below). Due to environmental uncertainty, **D** views $f_{\mathbf{N}}$ as a random variable, just like $f_{\mathbf{D}}$. The set of samples of $f_{\mathbf{N}}$ that **N** has at present are written as $L_{\mathbf{N}} \equiv \{L_{\mathbf{N}}^{X}(i), L_{\mathbf{N}}^{Y}(i) = f_{\mathbf{N}}[L_{\mathbf{N}}^{X}(i)]$: $i = 1, ..., N_{\mathbf{N}}\}$, and $L_{\mathbf{N}}^{\max} \equiv \max_{i \in 1,...,N_{\mathbf{N}}}[L_{\mathbf{N}}(i)]$. So $L_{\mathbf{N}}^{\max}$ is the current payoff achieved by **N**, and intuitively, $N_{\mathbf{N}}$ is how hard **N** worked to achieve that payoff.

Just like with L_D , **D** assumes that the formation of L_N can be modeled as N_N IID samples of a window distribution with parameters (μ_N , σ). Just as **D** does not know μ_D , nor does he know μ_N . And so, just like with μ_D , **D** models μ_N as set randomly, before **N**'s search began, by sampling some associated density function.

Now assume that the random variables $\mu_{\mathbf{N}}$ and $\mu_{\mathbf{D}}$ are coupled *a priori*, i.e., their formations before the searches begin are statistically dependent. This reflects the fact that the environments of the two decision makers are related in some manner. Write the probability (density) that the sampling produced the values ($\mu_{\mathbf{D}}, \mu_{\mathbf{N}}$) as $P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}})$. For simplicity, assume

⁵ More realistic models would assume **D** spends most of his search at X values that are relatively close to optimal. Such models would assign higher probability to $L_{\mathbf{D}}$ if its elements are "bunched up" at high values, with a long "tail" toward lower values.

that the joint probability of $\mu_{\mathbf{D}}$ and $\mu_{\mathbf{N}}$ is a mixture: With probability ρ , $\mu_{\mathbf{D}} = \mu_{\mathbf{N}}$ and they are generated by sampling (once) the density function *h*, while with probability $1 - \rho$, $\mu_{\mathbf{D}}$ and $\mu_{\mathbf{N}}$ are independent, being generated by two independent samples of *h*.⁶

In the first of these two cases, where $\mu_{\mathbf{D}} = \mu_{\mathbf{N}}$, information concerning $L_{\mathbf{N}}$ may tell **D** something concerning the value of $\mu_{\mathbf{D}}$. In the second case, where $\mu_{\mathbf{D}}$ and $\mu_{\mathbf{N}}$ are statistically independent, information about $L_{\mathbf{N}}$ provides no information concerning $\mu_{\mathbf{D}}$. As two examples, the first case, where $f_{\mathbf{N}} = f_{\mathbf{D}}$, can model the situation where (**D** knows that) **D** and **N** have the same age and have been reared in the same socio-economic environment, and therefore have the same payoff function. The second case can model a situation where **D** knows that they do not have this commonality, and so are searching different environment functions.

Say that **D** knows $L_{\mathbf{D}}$ in full, $N_{\mathbf{N}}$ and $L_{\mathbf{N}}^{\max}$. **D** must use this information to choose between two options. Under the first, he gets payoff $L_{\mathbf{D}}^{\max}$, i.e., he "exploits" the best strategy he has found so far. Under the second option he "explores" by forming $\hat{N}_{\mathbf{D}}$ more IID sample of $W_{\sigma}(\mu_{\mathbf{D}}, .)$ to generate a new set, $\hat{L}_{\mathbf{D}} \equiv \{\hat{L}_{\mathbf{D}}^{X}(i), \hat{L}_{\mathbf{D}}^{Y}(i) = f_{\mathbf{D}}[\hat{L}_{\mathbf{D}}^{X}(i)]$: $i = 1, ..., \hat{N}_{\mathbf{D}}\}$. He then gets payoff $\hat{L}_{\mathbf{D}}^{\max}$.

For simplicity, take $\hat{N}_{\mathbf{D}} = 1$ (similar results hold for other choices). So $\hat{L}_{\mathbf{D}}^{Y}$ is a single number. Accordingly, given what he knows, **D** should explore if and only if the associated payoff exceeds the payoff for exploiting, i.e., if $\hat{L}_{\mathbf{D}}^{\max} - L_{\mathbf{D}}^{\max} > 0$. Define the posterior expected value of this difference, conditioned on what **D** knows (namely $N_{\mathbf{N}}$, $L_{\mathbf{N}}^{\max}$, and $L_{\mathbf{D}}$), as the *hedonic prod* for what **D** knows. **D** should search whenever his hedonic prod is positive.

Intuitively, the information that $L_{\mathbf{D}}$ provides concerning $\mu_{\mathbf{D}}$ cannot tell **D** that exploring is likely to get a better payoff than the best one in $L_{\mathbf{D}}$. However the pair $(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}})$ may provide some additional information concerning $\mu_{\mathbf{D}}$. Such additional information may tell **D** that $\hat{N}_{\mathbf{D}}$ more samples of $W_{\sigma}(\mu_{\mathbf{D}}, .)$ are likely to produce a payoff better than $L_{\mathbf{D}}^{\max}$. When this is the case the hedonic prod will be positive, and therefore **D** should explore.⁷

As shown below, the relation between the sign of the hedonic prod and the various parameters specifying the search problem agree with much of the experimental data concerning the income comparison phenomenon. This is the basis for suggesting that that phenomenon is just a computational shortcut.

5. The Vanderbilts and the Bundys

The analysis in Appendices A–D shows that the behavior of the hedonic prod varies depending on the relative sizes of the payoffs in $L_{\mathbf{D}}$ and $L_{\mathbf{N}}$. First, if either $L_{\mathbf{D}}^{\max} - L_{\mathbf{N}}^{\max} > \sigma$ or $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min} > \sigma$, then the hedonic prod is given by

$$\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = \frac{\int_{L_{\mathbf{D}}^{\min} - \sigma/2}^{L_{\mathbf{D}}^{\min} + \sigma/2} dt \, th(t)}{\int_{L_{\mathbf{D}}^{\min} - \sigma/2}^{L_{\mathbf{D}}^{\min} + \sigma/2} dt \, h(t)} - L_{\mathbf{D}}^{\max}.$$
(3)

Note that this expression is independent of L_N^{max} and N_N , which are all that **D** knows concerning **N**'s search. Intuitively, the reason for this independence is that if either $L_D^{max} - L_N^{max} - \sigma$ or $L_N^{max} - L_D^{min} > \sigma$, then **D** concludes that f_N is independent of f_D , and so the results of **N**'s search have nothing to tell **D** about how to conduct his search. As dictated by Eq. (3), maybe he should explore, maybe not, but examining L_N does not help him decide.

Now $L_{\mathbf{D}}^{\max} - L_{\mathbf{N}}^{\max} > \sigma$ means that **N** has substantially less payoff than **D**. So our results means that **D**'s hedonic prod should not respond to someone with substantially less payoff. Colloquially, **D** should not get prodded by the payoff of any "Bundy" (and so in particular how happy **D** is should not be affected by his knowing about a Bundy).

"Bundy" (and so in particular how happy **D** is should not be affected by his knowing about a Bundy). Conversely, the condition $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min} > \sigma$ means that $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\max}$ is greater than 0, and perhaps much more so. Indeed, without considering the value of $L_{\mathbf{D}}^{\min}$, we are guaranteed that $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min} > \sigma$ whenever $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\max} > \sigma$. This inequality is the case where **N**'s payoff is at least σ greater than **D**'s. So our result means that **D** should not get prodded by any "Vanderbilt" (and so in particular how happy **D** is should not be affected by his knowing about a Vanderbilt).

As an example of these two results, a third world farmer should not get a hedonic prod by knowing that there is a first world physician with a far greater payoff. Nor should that physician get a hedonic prod by knowing about the farmer with low income.

To investigate Eq. (3) further, fix L_{N}^{max} and L_{D}^{max} , and in particular fix the payoff gap $L_{N}^{max} - L_{D}^{max}$ to a small, positive value. However don't fix any other aspects of L_{D} , in particular allowing both N_{D} and L_{D}^{min} to vary. Now the larger N_{D} is, the more likely it is *a priori* that $L_{D}^{max} - L_{D}^{min}$ is almost as large as σ . In turn, for our fixed $L_{N}^{max} - L_{D}^{max}$, the closer that $L_{D}^{max} - L_{D}^{min}$ is to σ , the larger $L_{N}^{max} - L_{D}^{min}$ is. So for a fixed payoff gap, the more searching **D** has already done, the more likely it is that $L_{N}^{max} - L_{D}^{min} > \sigma$ — which is exactly one of the conditions for Eq. (3) to apply. So the more searching **D** has

⁶ Formally, $P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}} | \rho) \triangleq \rho \delta(\mu_{\mathbf{D}} - \mu_{\mathbf{N}}) h(\mu_{\mathbf{N}}) + (1 - \rho) h(\mu_{\mathbf{D}}) h(\mu_{\mathbf{N}}).$

⁷ In many real world situations there are also costs associated with searching, and other values in $L_{\mathbf{D}}^{Y}$ and $\hat{L}_{\mathbf{D}}^{Y}$ besides their two maxima are important. Also, potentially at some cost **D** might be able to search, but then revert to the old strategy that gave payoff $L_{\mathbf{D}}^{\text{max}}$ before the time ΔT has passed. Here for pedagogical clarity I do not consider such situations.

already done, the more likely that the results of N's search are irrelevant to whether D should explore or exploit, i.e., the more likely D views N as a Vanderbilt.

Intuitively, this result reflects the fact that if **D** has worked hard to achieve his current payoff, he is likely to have seen it rise substantially, and therefore is unlikely to get a hedonic prod from any remaining gap between his current payoff and that of **N**. This suggests that in the real world, people should be less unhappy for a given gap between their payoff and that of their (richer neighbor) if they have worked hard to raise their payoff to its current value than if they haven't.

On the other hand, fix L_{N}^{max} , L_{D}^{max} , N_{N} , and N_{D} . Then as L_{D}^{min} shrinks, **D** eventually comes to view **N** as a Vanderbilt. Intuitively, this reflects the fact that if **D** has encountered a value of $f_{D}(x)$ in his past that is small compared to L_{D}^{max} , so that $L_{D}^{max} - L_{D}^{min}$ is close to σ , then he knows that if he were to explore, he would likely get a new payoff worse then his current one. So he should be conservative, and not explore. This is in broad agreement with empirical data that a person's happiness depends on previous earnings (Frederick and Loewenstein, 1999). See Appendices A–D for further discussion of Eq. (3).

6. The Joneses

I now consider the remaining scenario, where the conditions required for Eq. (3) are not met. So in particular, **N** neither has substantially less payoff than **D** nor has substantially greater payoff. In this remaining scenario, we cannot invoke Eq. (3) to rule out the possibility that the value of the hedonic prod depends on $L_{\rm N}^{\rm max}$ and/or $N_{\rm N}$.

I analyze this scenario under an approximation for *h*. To motivate this approximation, note that in many societies, the range of incomes encountered by a given person as they search over their strategy space is far smaller than the range of incomes of all people in that society. If we assume that *h* does not vary too much on scales much smaller than its range, this implies that *h* is approximately flat on the scale σ . Accordingly, here I make that approximation that *h* is perfectly flat over all regions of interest. More precisely, I take $h(.) \simeq W_T(\mu, .)/T$ for some μ and large *T*, where $[L_{\mathbf{p}}^{\max} - \sigma/2, L_{\mathbf{p}}^{\min} + \sigma/2] \subset [\mu - T/2, \mu + T/2]$.

The results described below are all derived in Appendices A–D. I divide the analysis into three special cases. The first special case is where **D**'s best payoff so far is better than **N**'s, but by an amount less than σ . Neither of the conditions for Eq. (3) are met in this case. However due to our assumption for *h*, it turns out that the payoff gap cannot provide a hedonic prod telling **D** to explore, just as it cannot provide such a prod when either of the conditions for Eq. (3) hold. Combining with the discussion of Eq. (3), we see that under our assumed *h*(.) *everyone* with less income than **D** is a "Bundy," as far as **D** is concerned; the precise details concerning **N**'s search are of no interest to **D** if the best payoff **N** found in that search is less than the best that **D** has found.

I now consider the remaining situation where $L_{\mathbf{D}}^{\max} < L_{\mathbf{N}}^{\max} \leq L_{\mathbf{D}}^{\min} + \sigma$. So the best payoff of **N** is larger than that of **D**, but not too much so. First, if $\rho \rightarrow 0$ while all other variables don't change, then the hedonic prod disappears. Intuitively, the greater **D**'s *a priori* belief that his environment function differs from that of his neighbor, the less likely he is to conclude that the results of his neighbor's search are relevant. (As an example, **D** could form such a belief by observing characteristics of **N**.)

Next, fix ρ , and divide up the situation where $L_{\mathbf{D}}^{\max} < L_{\mathbf{N}}^{\max} \leq L_{\mathbf{D}}^{\min} + \sigma$ into two special cases. The first of these special cases is where $L_{\mathbf{N}}^{\max}$ is just below $L_{\mathbf{D}}^{\min} + \sigma$ (on the scale of σ). In this case, say *T* and ρ are fixed and ρ/T is negligible. Then if $L_{\mathbf{N}}^{\max}$ rises closer to $L_{\mathbf{D}}^{\min} + \sigma$, the hedonic prod shrinks to zero. Intuitively, **D** views **N** more and more as a Vanderbilt as **N**'s payoff rises.

In the last case $L_{\mathbf{N}}^{\text{max}}$ is larger than $L_{\mathbf{D}}^{\text{max}}$, but not too close to $L_{\mathbf{D}}^{\text{min}} + \sigma$ (on the scale of σ). For this case, in the limit where $(1 - \rho)/T^2$ becomes negligible, we can write

$$\frac{\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\sigma} \simeq \gamma + 1/2 - \frac{N_{\mathbf{N}}}{N_{\mathbf{N}} + 1} \left[\frac{1 - \beta^{N_{\mathbf{N}} + 1}}{1 - \beta^{N_{\mathbf{N}}}} \right]$$
(4)

where $\beta \equiv (L_{N}^{\max} - L_{D}^{\min})/\sigma$ is non-negative and by hypothesis substantially less than 1, and where $\gamma \equiv (L_{N}^{\max} - L_{D}^{\max})/\sigma$ is non-negative and bounded above by β .

This hedonic prod can be positive, for example if $\gamma = 3/8$, $\beta = 1/2$, and $N_{\rm N} = 1$. In this situation, due to the size of *T*, it is very unlikely that $L_{\rm N}^{\rm max}$ and $L_{\rm D}^{\rm max}$ would be as close as they are unless $f_{\rm D} = f_{\rm N}$. Presuming that in fact $f_{\rm D}$ does equal $f_{\rm N}$, the greater payoff of N means that if D were to explore for a new strategy, he would likely get larger payoff. Accordingly, D becomes unhappy with his current choice of strategy, and as a result is prodded to search for a new strategy.

It may be that if **D** had done the calculation in Eq. (3) (in which he uses no information concerning **N**), he would also come to the conclusion that he should search for a new strategy. In this case, in essence the hedonic prod serves as a confirmation that such search is called for. That confirmation is communicated to **D** via his emotions, and is based on a social comparison heuristic.

The hedonic prod of Eq. (4) can also be negative. This happens for example if β and γ are close to 0 and N_N is very large. In this situation, the size of T and relative smallness of the gap in payoff again makes it likely that $f_D = f_N$. However now, the fact that **N** has searched so hard to find his strategy (N_N being so large) means that L_N^{max} is likely to be near the top of the window of values of f_N . Given that L_D^{max} is close to L_N^{max} , this means that **D** would likely end up with smaller

465

payoff if he were to replace his current strategy with the result of a new search. Accordingly, he should not be unhappy with his current strategy, and as a result is not prodded to search for a new strategy.

Note that if the hedonic prod is less than zero, its actual value is irrelevant. This implies that its effect on unhappiness has a lower bound of zero, loosely speaking. On the other hand, the hedonic prod can have different positive values. This implies that while one should not get prodded by the precise relative payoff of the Bundys (or Vanderbilts), the precise relative payoff of the Joneses is important. It's not just that observing the Joneses should prod **D** to search. Rather by ascertaining precisely how much more payoff the Joneses make, **D** can determine just how crucial it is that he search.⁸ This is in agreement with the fact that some evidence suggests that people are more concerned with upward social comparison than with downward social comparison (Duesenberry, 1949).

As a final, semi-formal comment on this analysis of the Joneses case, note that in the simple model of this paper, ρ reflects all prior information that **D** has concerning whether he and his neighbor are searching the same function f, and σ similarly reflects his prior information concerning the width of f(x) values. In different scenarios, reflecting different prior information, those values will be different, and therefore so will **D**'s perception of whether a given income gap Δ means **N** is a Vanderbilt or a Joneses.

As an example, say that in one scenario, the prior information available to "Joe Public" makes him believe that σ is small, as is the ρ relating him and a Wall Street banker. In this case, for a given (huge) gap in incomes, Δ , he will conclude that the banker is a Vanderbilt, and so will not be upset by the banker's income. (See discussion just before Eq. (4).) Now consider a different scenario, where Joe Public concludes that the banker is no smarter than he is, but rather is likely to make his income by bilking the US taxpayer, or via some similar process that Joe Public could have engaged in. In this case, Joe Public might believe that σ is actually quite large, and ρ is close to 1. These values mean that for the same gap in incomes of Δ , Eq. (4) would now apply, and **D** might conclude that the banker is a Jones. In other words, say that public perception makes **D** conclude that the banker's income is based on ill-gotten gains (which **D** himself could have gotten if he had chosen the appropriate strategy). Then *independent of "moral" reasoning* by **D**, **D** might be subject to a hedonic prod, making him be unhappy at the same income gap of Δ that would not bother him if there weren't that public perception.

7. The breadth of phenomena involving hedonic prods

The model analyzed above makes many simplifying assumptions. Two of the more important ones are:

1. The decision maker has only one neighbor.

2. $\hat{N}_{\mathbf{D}} = 1$.

In this section I present two stylized arguments concerning what happens if we relax each of these assumptions. I first argue that by relaxing assumption (1), one gets a model that predicts hedonic prods based on ordinal rather than cardinal social comparison. I then argue that by also relaxing assumption (2), one gets a model that predicts U-shaped lifetime happiness curves.

These stylized arguments are not intended as fully formal models that can be carefully compared, in a full statistical analysis, with experimental data. (That is future work.) Rather their purpose is to illustrate the broad range of phenomena that potentially involve hedonic prods.

7.1. Ordinal social comparison

Say that **D** has a total of *m* neighbors, where now *m* can exceed 1. For simplicity assume all of them generate their data sets { L_i : i = 1, ..., m} in the following way. First, **D** samples *h* to generate $\mu_{\mathbf{D}}$. Next each neighbor *i* independently flips a coin with probability ρ of heads. If it comes up heads, then $\mu_i = \mu_{\mathbf{D}}$. Otherwise, μ_i is formed by IID sampling h.⁹ Next, once $\mu_{\mathbf{D}}$ and all the μ_i are set, the associated data sets are generated by IID sampling the window functions $W_{\sigma}(\mu, .)$ parameterized by those μ 's. (This is just a simple-minded generalization of the single neighbor case analyzed above.)

Say that the information **D** uses to decide whether to search – the argument of his hedonic prod – is the vector $(L_{\mathbf{D}}, L_1^{\max}, N_1, \dots, L_m^{\max}, N_m)$. The question facing **D** is whether the expected value of the difference $\hat{L}_{\mathbf{D}}^{\max} - L_{\mathbf{D}}^{\max}$, conditioned on a particular value of this vector, is greater than zero.

As an example, consider the case where ρ is close to 1. So with high probability, all the μ 's are the same, i.e., $P(\mu_{\mathbf{D}}, \mu_1, \dots, \mu_m) \simeq h(\mu_{\mathbf{D}}) \prod_i \delta(\mu_i - \mu_{\mathbf{D}})$. Say that $\forall i \in \{1, \dots, m\}, N_i = N_{\mathbf{D}} = 1$. So loosely speaking, given $\mu_{\mathbf{D}}, L_{\mathbf{D}}^{\max}$ is a single IID sample of $W_{\sigma}(\mu_{\mathbf{D}}, .)$, as is each L_i^{\max} .

Write the cumulative distribution function (CDF) of $W_{\sigma}(\mu_{\mathbf{D}}, .)$ as

⁸ For example, if **D** were to have any uncertainty in his observation of $L_{\rm D}$, $L_{\rm N}^{\rm max}$ and $N_{\rm N}$, or if were to have other sources of information concerning the likely results of his searching in addition to these quantities, a large value of the prod would more assuredly mean that he should search than a small value.

⁹ Formally, $P(\mu_{\mathbf{D}}, \mu_1, ..., \mu_m) = h(\mu_{\mathbf{D}}) \prod_{i=1}^m [\rho \delta(\mu_i - \mu_{\mathbf{D}}) + (1 - \rho)h(\mu_i)].$

$$C(\kappa) = \int_{-\infty}^{\kappa} dz \, W_{\sigma}(\mu_{\mathbf{D}}, z) = \frac{\kappa - (\mu_{\mathbf{D}} - \sigma/2)}{\sigma}$$
(5)

for $\kappa \in [\mu_{\mathbf{D}} - \sigma/2)$, $\mu_{\mathbf{D}} + \sigma/2)$], 0 for $\kappa < \mu_{\mathbf{D}} - \sigma/2$, and 1 otherwise. Let *m* be quite large. Then given $\mu_{\mathbf{D}}$, on average the empirical CDF, $\hat{C}(\kappa) \equiv \{\text{Fraction of neighbors } i \text{ such that } L_i^{\max} > \kappa \}$, is a good approximation to $C(\kappa)$.

Due to this accuracy of the approximation, on average **D** gets a good estimate of where his current performance $L_{\mathbf{D}}^{\text{max}}$ lies on $C(\kappa)$ by examining where $L_{\mathbf{D}}^{\max}$ lies on $\hat{C}(\kappa)$. Therefore where $L_{\mathbf{D}}^{\max}$ lies on $\hat{C}(\kappa)$ gives **D** a good estimate of the probability that a subsequent sample of the distribution $W_{\sigma}(\mu_{\mathbf{D}}, .)$ would give a value larger than $L_{\mathbf{D}}^{\max}$. (Indeed, since I am using a distribution $W_{\sigma}(\mu_{\mathbf{D}}, .)$ that is uniform across its support, the expected value of a further sample of $W_{\sigma}(\mu_{\mathbf{D}}, .)$ is C(.5), which is well-approximated by $\hat{C}(.5)$, the 50th percentile value of the empirical CDF.) So where $L_{\mathbf{D}}^{\text{max}}$ lies on the empirical CDF gives a good estimate of the hedonic prod, i.e., of the posterior expected difference in utilities if **D** were to re-sample or not.

Similar results would hold even if the distribution parameterized by μ were not a window function. The crucial point is that the value of the hedonic prod is a single-valued function of where $L_{\mathbf{D}}^{\max}$ lies on $\hat{C}(\kappa)$. Broadly speaking, the same kind of reasoning holds even if the data sets had more than a single element: By examining where $L_{\mathbf{D}}^{\max}$ lies on $\hat{C}(\kappa)$, **D** would get a good estimate of whether he should search.

All of this means that there should be only a small expected loss of utility associated with making a choice of whether to search by examining $\hat{C}(L_{\mathbf{D}}^{\max})$ rather than by examining all of $(L_{\mathbf{D}}, L_1^{\max}, N_1, \dots, L_m^{\max}, N_m)$. On the other hand, one would expect that the computational savings for the human brain to evaluate a prod based on just $\hat{C}(L_{\mathbf{D}}^{\max})$ rather than all of $(L_{\mathbf{D}}, L_{1}^{\max}, N_{1}, \dots, L_{m}^{\max}, N_{m})$ may be large. This is a stylized argument that when **D** has multiple neighbors, under some circumstances (e.g., ρ close to 1, *m* large, etc.), **D** should use a prod based on $\hat{C}(L_{\mathbf{D}}^{\max})$. In other words, he should make his choice of whether to search based on an ordinal comparison of his performance with that of his neighbors. In essence, using percentiles to compare your performance with that of your multiple neighbors gives you a strong indication of whether you should search or not. Therefore, by the hypothesis of this paper, it is a major factor determining whether social comparison makes you unhappy.

7.2. U-shaped lifetime happiness curves

Say we generate a value $\mu_{\mathbf{D}}$ by randomly sampling h. Next generate an $L_{\mathbf{D}}$ by IID sampling the associated distribution $W_{\sigma}(\mu_{\mathbf{D}},.)$ a total of $N_{\mathbf{D}}$ times. Now, since $\hat{N}_{\mathbf{D}} = 1$, form a single next, "search" value of $f_{\mathbf{D}}$ by IID sampling $W_{\sigma}(\mu_{\mathbf{D}},.)$ one more time. Note that both the new search value and all the values in $L_{\mathbf{D}}$ are formed by IID sampling the same distribution. So for $N_{\mathbf{D}} > 1$, averaging both over $L_{\mathbf{D}}$ and that search value, we will find that that search value is less than $L_{\mathbf{D}}^{\max}$. Since **D** is not allowed to return to the *x* that produced $L_{\mathbf{D}}^{\max}$, this means that on average, the change in utility arising in the search would be negative.¹⁰

This reasoning has nothing to do with any data set $L_{\rm N}$. So it is unchanged if we also average over data sets $L_{\rm N}$ formed by the same sampling process. In other words, even if we also have access to such an $L_{\rm N}$, we would still conclude that, on average, the new search value of **D** is less than $L_{\mathbf{D}}^{\max}$.

This suggests (without formally proving) that since $\hat{N}_{\mathbf{D}} = 1$, even if $N_{\mathbf{D}}$ is small, usually the hedonic prod will not be positive. The crucial thing is the relative sizes of $\hat{N}_{\mathbf{D}}$ and $N_{\mathbf{D}}$. If $\hat{N}_{\mathbf{D}}$ is large enough (e.g., far bigger than $N_{\mathbf{D}}$), then the same kind of reasoning tells us that (on average) the hedonic prod will be positive.

This means that allowing $\hat{N}_{\rm D} > 1$ might partially explain U-shaped lifetime happiness curves. To see this, assume that if a prod is stimulated by a triple ($L_{\rm D}$, $L_{\rm N}^{\rm max}$, $N_{\rm N}$), then the prod persists for the entire time it takes to make the $\hat{N}_{\rm D}$ further samples of $f_{\mathbf{D}}$. Assume as well that **D** can have multiple neighbors, with new ones arising through a stochastic IID process. These two assumptions mean that there is a possibility that **D** would get a new prod while an earlier one still applies. For simplicity, assume that in such a situation, the prods (which all are either positive or zero) simply add. In other words, the unhappiness of **D** accumulates linearly.¹¹

Now early in **D**'s search $L_{\mathbf{D}}^{\text{max}}$ will not be very large, and as remarked above, for $\hat{N}_{\mathbf{D}}$ large enough, the hedonic prod stimulated by any single randomly chosen neighbor should be positive, on average. Accordingly, early in life, as they run into more and more neighbors, **D** keeps getting hedonic prods, which accumulate. The cumulative effect of those prods would mean that **D**'s unhappiness increases early in life, i.e., his happiness decreases. As the search continues, $L_{\mathbf{D}}^{\max}$ will eventually stop increasing much (presumably coming close to its maximum, given by

 $\mu_{\mathbf{D}} + \sigma/2$). However $N_{\mathbf{D}}$ will continue to increase (perhaps getting far bigger than $\hat{N}_{\mathbf{D}}$). As recounted above in the discussion

466

¹⁰ Note though that this is not the same as proving that the average *sign* of the change in utility value arising in the search would be negative.

¹¹ A more sophisticated hedonic prod would combine the data provided by all neighbors that **D** is observing at a given time. However once we allow for such sophistication, the notion that $\hat{N}_{\mathbf{D}}$ is fixed to some value (greater than 1) looks a bit preposterous. More generally, a proper analysis would also include a cost for searching, and allow **D** to return to a previously searched value before waiting the full time ΔT , at some cost. Then the hedonic prod would combine the observations of all of **D**'s neighbors with those costs to determine the optimal associated \hat{N}_{D} . Such a detailed calculation is beyond the scope of the stylized argument here. See the discussion in Section 8.3 below.

of Eq. (3), this increase in the number of points already searched would mean that **D**'s hedonic prod would shrink on average. The cumulative effect is that as $N_{\mathbf{D}}$ grows, **D** should receive hedonic prods less frequently. As a consequence, he will be made unhappy less frequently, i.e., his happiness will increase.

This general character of hedonic prods, in which they take a while to kick in, but after they do kick in they eventually shrink, should simultaneously hold for many of the different kinds of payoff that a human being tries to maximize in life. So there should be a cumulative effect of increasing unhappiness early in life, followed by decreasing unhappiness later in life. In this sense, the model predicts the U-shaped curve of happiness versus age found in recent social surveys.

Note that the hedonic prods in this explanation of U-shaped curves are actually just pushing **D** to do what is rational. Early in a search, on average searching further will be beneficial (assuming that $\hat{N}_{\mathbf{D}}$ is sufficiently greater than 1). So **D** should be prodded into such search. However later on, it will not be rational to search anymore; **D** has already done a very broad survey of their possible strategies. So **D** should no longer be prodded to search at that later time.

8. Discussion

8.1. The difference between utility and hedonic utility

As elaborated above, the central hypothesis of this paper has two parts. The first is that a person's environment provides statistical information about which strategies are best for him, for many implicit decision problems. The second is that it often makes sense to synopsize that information as computational shortcuts, aka "emotions," a particular type of which is the emotion of unhappiness. Under this hypothesis, many emotional responses by a person to his environment that appear pointless may actually be utility-maximizing, if we view them as computational shortcuts, and if we also consider *all* the implicit decision problems facing the person, and *all* the information the environment provides about how to respond to those problems.

As a particular instance of this hypothesis, first note that, tautologically, social comparison provides many kinds of statistical information to a decision maker. Next, note that experiments have found that such information sometimes induce "unhappiness." In turn, such unhappiness typically cause changes in future behavior. So it causes changes in *future* utility. Even if unhappiness based on social comparison has no effect on *current* utility, it may be rational because it induces them to change their behavior in a way that improves their *future* utility.

Note that this hypothesis may explain why utility and happiness are so often uncorrelated in the real world. (Un)happiness has to do with *changes* in utility values, not those utility values themselves. This makes unhappiness inherently a relative concept. As an example, this hypothesis concerning unhappiness suggests why people have hedonic set points — a set point is simply a "calibrated baseline," that arises as a comparison point when estimating the future expected utilities associated with possible changes in behavior. If your utility goes up and stays up (or goes down and stays down), that provides no reason for your happiness to change; it is only in considering actions that can affect the dynamics of your utility that your happiness level should vary from the baseline.

8.2. The hedonic model introduced here

To investigate this hypothesis, I have analyzed an extremely simple model concerning social comparison and the emotion of unhappiness. Despite all of its simplifications, the model agrees with many of the stylized facts reported in the experimental and survey literatures. (Because there are so many studies in that literature, using very many different protocols, it is hard to make a detailed numerical comparison of how well those studies agree with the model presented here.)

In addition, this model makes predictions that can be tested in future studies. Some of these are:

- 1. On average, a person **D** is less unhappy for a given gap between their payoff and that of a neighbor **N** with higher payoff if **D** searched over many strategies to raise their payoff to its current value. So for example, for a given gap in performance at a sport between an athlete **D** and their neighbor, the longer **D** has been playing the sport, the less unhappy the gap should make **D**. (See discussion at the end of Section 5.)
- 2. Similarly, on average, a person **D** is less unhappy for a given gap between their payoff and that of a neighbor **N** with higher payoff if **N** searched over many strategies to raise their payoff to its current value. So for example, for a given gap in performance at a sport between an athlete **D** and their neighbor **N**, the longer **N** has been playing the sport, the less unhappy the gap should make **D**.
- 3. Another set of predictions are presented after Eq. (4). In particular, one of those predictions is that the more similar the life circumstances of a decision maker are with that of a "neighbor," the higher ρ will be, and therefore the stronger social comparison hedonic effects should be. For example, you should be made more unhappy by a sibling who always performed as well as you when you were growing up but now has more income than by a next-door neighbor who has more income, who in turn should make you more unhappy than the knowledge that an abstract "average citizen" of your country has more income.
- 4. People are less unhappy for a given gap between their payoff and that of their (richer neighbor) if they have encountered particularly bad payoff values earlier in their life.

There are also several implications of the informal analysis of U-shaped lifetime happiness curves that suggest testable predictions. For example, that analysis suggests that the more searching over professional strategies a person has done by a given age, and the more professional comparisons of their payoff to that of others, the more quickly they should go through the U-shaped lifetime happiness curve. As an example, a professional musician or athlete, engaging in frequent competition, might be expected to go through that curve more quickly than a policeman or fireman. (Such predictions should be treated with care however; the analysis presented above of U-shaped curves was informal and incomplete.)

If any of these predictions are refuted by experiments, that would mean the precise model presented here is wrong in some of its details. This is almost undoubtedly true to some degree; the only question is how wrong the details are. However such experimental refutations would not by themselves mean that the central hypothesis of this paper is wrong; that could only be determined by carefully considering all implicit aspects of the precise problem faced by the decision maker in the experiment refuting the particular formal model presented above. Instead discrepancies between the precise model presented here and experimental tests may simply mean that the model presented here needs to be refined.

8.3. Simplifications of the model

To better understand how to use experiments to refine the hedonic prods model, it is necessary to better understand all the simplifications of that model. Two of those simplifications were discussed above in Section 7, but there are many others.

Two of the more obvious of these other simplifications are that the model assumes that there is zero cost to search, and that **D** cannot return to the strategy that produced $L_{\mathbf{D}}^{\max}$ once he decides to search without waiting a full time ΔT . Modifying the model to fix the first simplification would be straightforward; one simply subtracts some appropriate search cost constant from the definition of hedonic prod presented at the end of Section 4. The effect of this fix would be to reduce the number of instances in which **D** should search, and therefore reduce the number (and severity) of hedonic prods.

Fixing the second simplification would lead to less trivial changes in the analysis. In particular, the ability of **D** to return to $L_{\mathbf{D}}^{\max}$ quickly (on the scale of his discounting rate) would make him not be interested in the expected payoff for any new x, $\mathbb{E}(f_{\mathbf{D}}|x)$, but rather in something like the expected improvement, $\mathbb{E}(f_{\mathbf{D}}|x) - L_{\mathbf{D}}^{\max}$. In general, this modification should make **D** more willing to search. (This contrasts with the modification that fixes the first simplification, which has the effect of making **D** less willing to search.)

There are many other simplifications in the model that have significant implications. One is the fact that the model ignores all statistical correlations in the samples forming $L_{\mathbf{D}}^{X}$. This means that the model assumes, in essence, that **D** cannot improve his sampling as he searches more, and it is not at all clear that this is a property of real-world search. The use of a window density function, rather than an everywhere differentiable density function having infinite support, is also problematic. So is the fact that the density functions governing **N**'s search (e.g., the prior over $\mu_{\mathbf{N}}$, the width of the window distribution, etc.) are the same as those governing **D**'s search. Yet another dubious simplification is that all that **D** knows is *L*, $N_{\mathbf{N}}$, and $L_{\mathbf{N}}^{\max}$. (A particularly dubious instance of this simplification is that in the model presented here, **D** knows nothing about the *x* that corresponds to $L_{\mathbf{N}}^{\max}$.)

At least as important as these simplifications of the search process are what aspects of search are ignored entirely. For example, in the current model there is no noise in either $f_{\mathbf{D}}$ or $f_{\mathbf{N}}$. However often in the real world there will be a large amount of such noise. The result of such noise would presumably be that the more times that **D** samples his current *x*, the more accurately he can estimate $\mathbb{E}(f_{\mathbf{D}}(x))$. One might expect that this would result in situations where a single sample of *x* does not result in a hedonic prod, but multiple samples do. So the longer **D** sticks with an *x*, the more likely he is to get an unhappiness prod, and therefore have his happiness reduced. This effect might explain some of the empirical data concerning habituation (Clark, 1999; Di Tella et al., 2003).

More generally, in the real world, in many situations **D** will have information not considered in the model presented here. The model would need to be modified to properly apply to those situations. For example, suppose that **D** knows that **N** was "born with a silver spoon in his mouth" and inherited his father's business, but then ran that business poorly, and now has an income only slightly larger than **D**'s. In this case relative income would provide no information to **D** about whether **D** should search more. It would not (or should not) result in a hedonic prod.

8.4. Hedonic prods in societies

Another broad category of simplifications made by the model is that it ignores all effects of **D**'s being embedded in a full society, in situations involving more decisions than the immediate one confronting him. For example, no concern for reputation effects or repeated games are contained in the model. (So for example the model has nothing to say about why someone might not want others to know that their income is less than maximal.) Nor are any concerns related to reproductive fitness reflected in the model.

However the hedonic prod model can be extended to apply to at least some such social situations. To illustrate this, say that some other person i has a similar best payoff to **D**, and that **D** is quite sure that i is searching the surface $f_{\mathbf{D}}$ (e.g., i may be a close friend of **D**, so that **D** knows i well enough to be confident in this belief about i's surface). Say that the happiness of i is determined by a hedonic prod, and that i is searching over multiple neighbors who are not observable to **D**. Then the happiness of i provides information to **D** about whether the payoffs of those unobservable neighbors should (not)

469

lead **D** to change his *x*. (Intuitively, if *i* has no hedonic prod, this suggests that *i*'s neighbors are not Joneses relative to *i*, which in turn means they are likely not to be Joneses relative to **D**.) Generalizing, the happiness of those people *i* who are close social contacts of **D** should have a major impact on the (un)happiness of **D**. In this way, a simple extension to the model in this paper might shed light on aspects of hedonic contagion (Fowler and Christakis, 2008).

As another example, note that the types of hedonic prod investigated in this paper are non-strategic, involving optimal choice problems. However the idea of hedonic prods can easily be extended to model the hedonic effects of social comparison based on strategic considerations. For example, while couched in terms of mating competition and markets, the basic insight in Cole et al. (1995) can easily be used to build to an extended version of the hedonic prods model that incorporates strategic considerations.

This model involves a player/decision maker, **D**, together with a set of multiple other players/neighbors, $B = \{B_i\}$, and a third person, *C*. *C* must choose how to allocated a set 1 + |B| of different "favors" among player **D** and the players in *B*. *C* must give each player exactly one of those favors. Moreover, each favor has a different worth, which is the same to all players. (As an example, each favor might be a different binding contract between *C* and the recipient of the favor, where all players can satisfy their commitments under the contracts with identical costs, and where utilities are transferable.)

For simplicity, assume that each player has a scalar-valued "appeal" to *C*, which is a function of a strategy choice *x* by the player. Assume further that *C* will assign the favor with *k*'th worth to the player that has *k*'th appeal to *C*. So **D** wants to appear as appealing to *C* as possible, compared to the players $B_i \in B$. If before *C* makes their choice **D** observes that some B_i is *slightly* more appealing than **D** is, this would serve as a hedonic prod pushing **D** to search for a new strategy *x*, in the hope of finding one that will make **D** more appealing to *C* than B_i . In this model $f_{\mathbf{D}}(x)$ is the appeal of **D** to *C* if **D** chooses *x*. Note that unlike in the model discussed above, here the value of $f_{\mathbf{D}}(x)$ does not have *direct* worth to **D**. Rather it has direct "appeal" to *C*, and through the resultant actions of *C* has indirect worth to **D**.

It is important to emphasize that even if extended in the ways discussed in this and the preceding subsections, the basic concept of hedonic prods would not be expected to explain all "non-rational" social comparison. For example, it is hard to see how it could explain envy of those with greater pecuniary endowments – endowments are not something that can be changed by modifying one's income-generating strategy. The model does not analyze emotional responses to inter-personal variations in wealth.¹²

Finally, it is worth emphasizing that there is nothing in the model of hedonic prods that predicts what the cost of any given computation is. (Presumably those costs ultimately are set from neurobiological considerations.) Accordingly, there is nothing in the model that says what the argument of a hedonic prod function should be; in the model, the domain of the hedonic prod function is set *ex ante*, to roughly agree with what it seems to be in experiments.

8.5. Hedonic prods and a theory of emotions

In the broadest sense, a "hedonic prod" is any computational shortcut with the following attributes:

- 1. It replaces a full statistical calculation by **D**, based on all the information provided to him, of whether **D** should retain his current strategy.
- 2. It replaces that calculation with a calculation based on a subset of the information provided to **D**.
- 3. It saves **D** some computational cost.
- 4. It manifests itself in an emotional cost for **D** if he decides to retain his current strategy.

Under this definition, a hedonic prod can exist for information provided to **D** that involves no social comparison. For example, it may exist if the information provided to **D** is changes in time of his payoff.

Being a computational shortcut, a hedonic prod "emotion" would get stimulated in any situation sufficiently similar to the one for which it is appropriate. This raises the possibility that prods are stimulated in rare situations that are similar to the (more common) ones where there are rational, even though they are not rational in those rare situations.¹³ This kind of application of a hedonic prod where it is not appropriate may explain some types of human behavior that truly are non-rational. In particular, game theoretic experiments often present subjects with an artificially contrived situation, almost never occurring in the field, that is similar to a far more common situation for which a hedonic prod would be rational. The premise is that precisely because those prods are shortcuts, they can "be stimulated" by these artificial situations, even if they are actually inappropriate in them.

An example might be the famous laboratory experiments in which a subject elects to be poorer rather than richer, if their doing so ensures that their neighbors are poorer than they are. As described above, an income-based hedonic prod can be a rational response to many field scenarios where you observe a neighbor's income. Since that prod is only based on the values L_{N}^{max} , N_{N} , L_{D} , it can get stimulated by rare scenarios that are similar to those field ones, in that those rare

¹² To be precise, what the model analyzes is emotional responses to inter-personal variations in *the rates of change* of people's wealth (i.e., inter-personal variations in income), since it is those that are dependent on one's income-generating strategy.

¹³ Presumably the computational cost avoided by using a hedonic prod rather than a full statistical calculation, combined with the fact that only rarely does using the prod result in worse payoff than using the full calculation, means that on average, the rule to always use the prod is more rational than the rule to always use the full statistical calculation.

scenarios provide values L_{N}^{max} , N_{N} , L_{D} that stimulate the prod, even though the prod is not rational in those rare situations. The aforementioned laboratory experiments might be examples of such a rare situation; in the field, one is almost never given an explicit choice of whether to be poor if doing so ensures your neighbors are poorer still.

Note though that in all such situations in which the prod is inappropriate, one would expect other factors that *are* appropriate to the situation to come into play, in addition to the hedonic prod. So the hedonic prod would not explain all aspects of behavior in these situations.

This illustrates that the simple model hedonic prod considered here is not intended to explain all phenomena involving the emotion of (un)happiness. More generally, many (the majority?) of emotions may have nothing to do with the general phenomenon of unhappiness, in the sense that that phenomenon is considered by hedonic prods. As an example, none of the models that explain emotions in social interactions in terms of honest signaling, like "persona games" (Wolpert et al., submitted for publication) (more recently called "rational emotions"; Winter et al., 2009, or "gaming emotions"; Andrade and Ho, 2009), seem to be closely related to hedonic prods. Nor do hedonic prods directly related to aspects of behavior involving reproductive fitness.

Clearly a lot of work remains to be done to integrate hedonic prods into a full model of human emotions. But at least it provides us a start at understanding the emotion of unhappiness.

Acknowledgments

I would like to thank Mark Wilber and Julian Jamison for helpful discussion.

Appendix A. Preliminary calculations

Since $\hat{N}_{\mathbf{D}} = 1$ is a single number, given any particular value $\mu_{\mathbf{D}}$, the associated expected value of $\hat{L}_{\mathbf{D}}^{Y}$ is $\mu_{\mathbf{D}}$. As a result, for any given triple $(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ that **D** observes, **D**'s associated hedonic prod is given by

$$\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) \triangleq \mathbb{E}(\mu_{\mathbf{D}} | L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) - L_{\mathbf{D}}^{\max}.$$
(6)

For **D** to evaluate this requires calculating the conditional distribution

$$P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}} | L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) \propto P(L_{\mathbf{N}}^{\max}, L_{\mathbf{D}} | \mu_{\mathbf{D}}, N_{\mathbf{N}}, \mu_{\mathbf{N}}) P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}})$$

$$\triangleq \mathcal{L}(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}, L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$$
(7)

where the argument of each P(.) indicates which density function it is, the proportionality constant, $1/P(L_N^{max}, N_N, L_D)$, is given by integrating $\mathcal{L}(\mu_D, \mu_N, L_N^{max}, N_N, L_D)$ over all μ_D and μ_N , and the prior $P(\mu_D, \mu_N)$ is given by the *a priori* coupling between μ_D and μ_N . (Note the implicit simplifying assumption that both N_N and N_D are parameters, not random variables.) Plugging in,

$$\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = \frac{\int d\mu_{\mathbf{D}} d\mu_{\mathbf{N}} \mu_{\mathbf{D}} \mathcal{L}(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}, L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\int d\mu_{\mathbf{D}} d\mu_{\mathbf{N}} \mathcal{L}(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}, L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})} - L_{\mathbf{D}}^{\max}$$
$$\triangleq \frac{\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})} - L_{\mathbf{D}}^{\max}.$$
(8)

To evaluate the numerator and denominator in Eq. (8), we must evaluate

$$\mathcal{L}(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}, L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = P(L_{\mathbf{D}}|\mu_{\mathbf{D}})P(L_{\mathbf{N}}^{\max}|N_{\mathbf{N}}, \mu_{\mathbf{N}})P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}).$$
(9)

 $P(L_{\mathbf{D}}|\mu_{\mathbf{D}})$ is given by the usual IID sampling formula, and $P(\mu_{\mathbf{D}}, \mu_{\mathbf{N}})$ is given by the ρ -based formula described in the text. The remaining term in Eq. (9) is given by

$$P(L_{\mathbf{N}}^{\max}|N_{\mathbf{N}},\mu_{\mathbf{N}}) = N_{\mathbf{N}} \frac{W_{\sigma}(\mu_{\mathbf{N}},L_{\mathbf{N}}^{\max})}{\sigma} \left[\frac{L_{\mathbf{N}}^{\max} - (\mu_{\mathbf{N}} - \sigma/2)}{\sigma} \right]^{N_{\mathbf{N}}-1}.$$
(10)

Using the resultant formula for $\mathcal{L}(\mu_{\mathbf{D}}, \mu_{\mathbf{N}}, L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$,

$$\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = N_{\mathbf{N}} \sigma^{-N_{\mathbf{D}}-N_{\mathbf{N}}} \left\{ \rho \left[\int dt \, th(t) \left(\prod_{i=1}^{N_{\mathbf{D}}} W_{\sigma}\left(t, L_{\mathbf{D}}^{Y}(i)\right) \right) W_{\sigma}\left(t, L_{\mathbf{N}}^{\max}\right) \left[L_{\mathbf{N}}^{\max} - t + \sigma/2 \right]^{N_{\mathbf{N}}-1} \right] + (1 - \rho) \left[\int dt \, th(t) \left(\prod_{i=1}^{N_{\mathbf{D}}} W_{\sigma}\left(t, L_{\mathbf{D}}^{Y}(i)\right) \right) \int ds \, h(s) W_{\sigma}\left(s, L_{\mathbf{N}}^{\max}\right) \left[L_{\mathbf{N}}^{\max} - s + \sigma/2 \right]^{N_{\mathbf{N}}-1} \right] \right\}.$$

$$(11)$$

 $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same weighted sum of integrals, except that in both integrals over *t*, *th*(*t*) is replaced by *h*(*t*).

Appendix B. Ignore the Bundys and Vanderbilts

The integrand of the first integral in Eq. (11) contains the product $W_{\sigma}(t, L_{\mathbf{D}}^{\max})W_{\sigma}(t, L_{\mathbf{N}}^{\max})W_{\sigma}(t, L_{\mathbf{D}}^{\min})$. Therefore that integral equals zero if either $L_{\mathbf{D}}^{\max} - L_{\mathbf{N}}^{\max} > \sigma$ or $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min} > \sigma$. Therefore under either of those conditions,

$$\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = \frac{\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})} - L_{\mathbf{D}}^{\max}$$

$$= \frac{\int dt th(t)(\prod_{i=1}^{N_{\mathbf{D}}} W_{\sigma}(t, L_{\mathbf{D}}^{Y}(i)))}{\int dt h(t)(\prod_{i=1}^{N_{\mathbf{D}}} W_{\sigma}(t, L_{\mathbf{D}}^{Y}(i)))} - L_{\mathbf{D}}^{\max}$$

$$= \frac{\int_{\mathbf{D}}^{L_{\mathbf{D}}^{\min} + \sigma/2} dt th(t)}{\int_{\mathbf{D}}^{L_{\mathbf{D}}^{\min} + \sigma/2} dt h(t)} - L_{\mathbf{D}}^{\max}.$$
(12)

(Note that there is zero probability of a data set *L* such that $L_{\mathbf{D}}^{\max} - \sigma/2 > L_{\mathbf{D}}^{\min} + \sigma/2$.) This is just Eq. (3) in the text. To illustrate this result, say that its conditions apply, so that either $L_{\mathbf{D}}^{\max} - L_{\mathbf{N}}^{\max} > \sigma$ or $L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min} > \sigma$, and that h(t) is approximately flat across the integration ranges in Eq. (12). Then we can approximate

$$\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) \simeq \frac{L_{\mathbf{D}}^{\min} - L_{\mathbf{D}}^{\max}}{2} \leqslant 0.$$
(13)

So in this case, **D** should decide not to explore. This conclusion is crucially dependent on the modeling choice that $\ddot{N}_{D} = 1$ and that **D** cannot return to any of the previously searched strategies in $L_{\mathbf{D}}^{X}$. If either of those choices are changed, then so might this conclusion.

Appendix C. Keep up with the Joneses

Now consider the situation where N has neither substantially greater payoff than D nor substantially lower, so that $W_{\sigma}(t, L_{\mathbf{D}}^{\max})W_{\sigma}(t, L_{\mathbf{N}}^{\max})W_{\sigma}(t, L_{\mathbf{D}}^{\min}) \neq 0$. In this situation both terms in Eq. (11) are non-zero. First we consider the case where $L_{\mathbf{D}}^{\max} \leq L_{\mathbf{N}}^{\max} \leq L_{\mathbf{D}}^{\min} + \sigma$. For this case,

$$\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = N_{\mathbf{N}}\sigma^{-N_{\mathbf{D}}-N_{\mathbf{N}}} \left\{ \rho \int_{L_{\mathbf{N}}^{\max}-\sigma/2}^{L_{\mathbf{D}}^{\min}+\sigma/2} dt \, th(t) [L_{\mathbf{N}}^{\max}-t+\sigma/2]^{N_{\mathbf{N}}-1} + (1-\rho) \int_{L_{\mathbf{D}}^{\min}-\sigma/2}^{L_{\mathbf{D}}^{\min}+\sigma/2} dt \, th(t) \int_{L_{\mathbf{N}}^{\max}-\sigma/2}^{L_{\mathbf{N}}^{\min}+\sigma/2} ds \, h(s) [L_{\mathbf{N}}^{\max}-s+\sigma/2]^{N_{\mathbf{N}}-1} \right\}.$$
(14)

As before, $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression as $\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$, after replacing $th(t) \rightarrow h(t)$. Now consider the situation where we can approximate $h(.) \simeq W_T(\mu, .)/T$ for some μ and large T, where $[L_{\mathbf{D}}^{\max} - L_{\mathbf{D}}]$ $\sigma/2, L_{\mathbf{D}}^{\min} + \sigma/2] \subset [\mu - T/2, \mu + T/2]$. Under this approximation,

$$\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) = N_{\mathbf{N}}\sigma^{-N_{\mathbf{D}}-N_{\mathbf{N}}} \left\{ \frac{\rho}{T} \int_{L_{\mathbf{N}}^{\max}-\sigma/2}^{L_{\mathbf{D}}^{\min}+\sigma/2} dt t [L_{\mathbf{N}}^{\max}-t+\sigma/2]^{N_{\mathbf{N}}-1} + \frac{1-\rho}{T^{2}} \int_{L_{\mathbf{D}}^{\min}-\sigma/2}^{L_{\mathbf{D}}^{\min}+\sigma/2} dt t \int_{L_{\mathbf{N}}^{\max}-\sigma/2}^{L_{\mathbf{N}}^{\max}+\sigma/2} ds [L_{\mathbf{N}}^{\max}-s+\sigma/2]^{N_{\mathbf{N}}-1} \right\}.$$
(15)

As usual, $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression, only with the monomials *t* removed in the integrands. Note that for $\rho \to 0$, the first integral in Eq. (15) becomes negligibly small. Similarly, for fixed $L_{\mathbf{D}}$, *T* and ρ , as $L_{\mathbf{N}}^{\max} \to L_{\mathbf{D}}^{\min} + \sigma$, that first integral disappears. In these limits, $\mathcal{N}/\mathcal{D} \to [L_{\mathbf{D}}^{\min} + L_{\mathbf{D}}^{\max}]/2$, and the hedonic prod goes negative as **D** starts to view N more and more as a Vanderbilt.

On the other hand, for $L_{\mathbf{N}}^{\max}$ not too close to $L_{\mathbf{D}}^{\min} + \sigma$ (on the scale of σ), the integration limits in the first integral won't be too close to each other. In this case, for large enough T and ρ not too small, the prefactor $(1 - \rho)/T^2$ on the product

of the last two integrals means that product is negligible compared to the first integral. The analogous property holds for $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$. So under these conditions,

$$\frac{\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})} \simeq -\frac{N_{\mathbf{N}}}{N_{\mathbf{N}} + 1} \left[\frac{(L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min})^{N_{\mathbf{N}} + 1} - \sigma^{N_{\mathbf{N}} + 1}}{(L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min})^{N_{\mathbf{N}}} - \sigma^{N_{\mathbf{N}}}} \right] + L_{\mathbf{N}}^{\max} + \sigma/2$$
$$= -\frac{N_{\mathbf{N}}\sigma}{N_{\mathbf{N}} + 1} \left[\frac{1 - \beta^{N_{\mathbf{N}} + 1}}{1 - \beta^{N_{\mathbf{N}}}} \right] + L_{\mathbf{N}}^{\max} + \sigma/2, \tag{16}$$

where $\beta \equiv (L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\min})/\sigma$ is non-negative and by hypothesis substantially less than 1. Under this approximation,

$$\frac{\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\sigma} = \gamma + 1/2 - \frac{N_{\mathbf{N}}}{N_{\mathbf{N}} + 1} \left[\frac{1 - \beta^{N_{\mathbf{N}} + 1}}{1 - \beta^{N_{\mathbf{N}}}} \right]$$
(17)

where $\gamma \equiv (L_{\mathbf{N}}^{\max} - L_{\mathbf{D}}^{\max})/\sigma$ is non-negative and bounded above by β .

Appendix D. Everyone with less payoff than you is a Bundy

We now consider the remaining case, where **D** has greater payoff than **N** but by an amount less than σ . For simplicity, we will use the same approximation for h used in Appendix C.

To begin, consider the special instance of this case where $L_{\mathbf{D}}^{\min} \leq L_{\mathbf{N}}^{\max} \leq L_{\mathbf{D}}^{\min} + \sigma$. For this special case, $\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression as in Eq. (14), except that in the first and third integrals, the lower bound is $L_{\mathbf{D}}^{\max} - \sigma/2$ rather than $L_{\mathbf{N}}^{\max} - \sigma/2$. In turn, just like for the previously considered cases, for this special case $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression as $\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$, after replacing $th(t) \to h(t)$. Plugging our approximation for h into these expressions, we get

$$\frac{\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})}{\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})} = \frac{\frac{\rho}{T} \left[L_{\mathbf{N}}^{\max} + \sigma/2 - \frac{N_{\mathbf{N}}\sigma}{N_{\mathbf{N}+1}} \left(\frac{\beta^{N_{\mathbf{N}}+1} - (1+\gamma)^{N_{\mathbf{N}}+1}}{\beta^{N_{\mathbf{N}}} - (1+\gamma)^{N_{\mathbf{N}}}} \right) \right] + \frac{1-\rho}{T^{2}} \left[\frac{(L_{\mathbf{D}}^{\min} + \sigma/2)^{2} - (L_{\mathbf{D}}^{\max} - \sigma/2)^{2}}{2} \right]}{\frac{\rho}{T} [1] + \frac{1-\rho}{T^{2}} [L_{\mathbf{D}}^{\min} + \sigma - L_{\mathbf{D}}^{\max}]} \\ \triangleq \frac{\frac{\rho}{T}a + \frac{1-\rho}{T^{2}}c}{\frac{\rho}{T}b + \frac{1-\rho}{T^{2}}d}$$
(18)

where β and γ are defined in Section 8.5, and a, b, c and d are defined to equal the square bracket terms in the obvious way. Note that here, $-1 \le \gamma \le 0 \le \beta \le 1$. Now $c/d = [L_{\mathbf{D}}^{\max} + L_{\mathbf{D}}^{\min}]/2 \le L_{\mathbf{D}}^{\max}$. We next consider the other ratio in Eq. (18), a/b:

$$\frac{a}{b} = L_{\mathbf{D}}^{\max} + \sigma \left[\gamma + 1/2 - \frac{N_{\mathbf{N}}}{N_{\mathbf{N}} + 1} \left(\frac{\beta^{N_{\mathbf{N}} + 1} - (1 + \gamma)^{N_{\mathbf{N}} + 1}}{\beta^{N_{\mathbf{N}}} - (1 + \gamma)^{N_{\mathbf{N}}}} \right) \right]$$

$$\triangleq L_{\mathbf{D}}^{\max} + \sigma \left[\gamma + 1/2 - \frac{n}{n+1} \left(\frac{\beta^{N_{\mathbf{N}} + 1} - (1 + \gamma)^{N_{\mathbf{N}} + 1}}{\beta^{N_{\mathbf{N}}} - (1 + \gamma)^{N_{\mathbf{N}}}} \right) \right]$$
(19)

where $n \equiv N_{\rm N}$. To bound the right-hand side first note that the term in circular parentheses is never negative. Next allow $n \neq N_{\mathbf{N}}$ and maximize the right-hand side over all $n \in \mathbb{N}^+$. That maximum occurs for n = 1. So

$$\frac{a}{b} \leq L_{\mathbf{D}}^{\max} + \sigma \left[\gamma + 1/2 - \frac{1}{2} \left(\frac{\beta^{N_{\mathbf{N}}+1} - (1+\gamma)^{N_{\mathbf{N}}+1}}{\beta^{N_{\mathbf{N}}} - (1+\gamma)^{N_{\mathbf{N}}}} \right) \right].$$
(20)

Now

$$\frac{\beta^{N_{\mathbf{N}}+1} - (1+\gamma)^{N_{\mathbf{N}}+1}}{\beta^{N_{\mathbf{N}}} - (1+\gamma)^{N_{\mathbf{N}}}} = (1+\gamma) \left[\frac{(\frac{\beta}{1+\gamma})^{N_{\mathbf{N}}+1} - 1}{(\frac{\beta}{1+\gamma})^{N_{\mathbf{N}}} - 1} \right]$$
$$\geqslant 1+\gamma \tag{21}$$

where the inequality follows from $\beta/(1 + \gamma) \ge 0$. As a result,

$$\left[\gamma + 1/2 - \frac{1}{2} \left(\frac{\beta^{N_{N}+1} - (1+\gamma)^{N_{N}+1}}{\beta^{N_{N}} - (1+\gamma)^{N_{N}}} \right) \right] \leqslant 0.$$
(22)

This means that in addition to $c/d \leq L_{\mathbf{D}}^{\max}$, $a/b \leq L_{\mathbf{D}}^{\max}$. So

$$\frac{\frac{\rho}{T}a + \frac{1-\rho}{T^2}c}{\frac{\rho}{T}b + \frac{1-\rho}{T^2}d} \leqslant L_{\mathbf{D}}^{\max},\tag{23}$$

i.e., $\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) \leq 0.$

The remaining special case to consider is where $L_{\mathbf{D}}^{\max} - \sigma \leq L_{\mathbf{N}}^{\max} \leq L_{\mathbf{D}}^{\min} \leq L_{\mathbf{D}}^{\max} \leq L_{\mathbf{D}}^{\min} + \sigma$. For this special case, $\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression as in Eq. (14), except that in the first and third integrals, not only is the lower bound replaced by $L_{\mathbf{D}}^{\max} - \sigma/2$, but in addition the upper bound is replaced by $L_{\mathbf{N}}^{\max} + \sigma/2$. In turn, as usual, $\mathcal{D}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$ is given by the same expression as $\mathcal{N}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}})$, after replacing $th(t) \to h(t)$. Continuing as before, we again get Eq. (18), just with different *a*, *c*, and *d*. Plugging in for those new values, we now have $c/d = [L_{\mathbf{D}}^{\max} + L_{\mathbf{N}}^{\max}]/2$, which as before is bounded above by $L_{\mathbf{D}}^{\max}$. In addition, plugging in for the new value of *a* gives

$$\frac{a}{b} \leqslant L_{\mathbf{D}}^{\max} + \sigma \left[\gamma + 1/2 - \frac{1}{2} \left(\frac{0 + (1 + \gamma)^{N_{\mathbf{N}} + 1}}{0 + (1 + \gamma)^{N_{\mathbf{N}}}} \right) \right] = L_{\mathbf{D}}^{\max} + \gamma/2.$$
(24)

So as before, $a/b \leq L_{\mathbf{D}}^{\max}$, and therefore we again get $\mathcal{H}(L_{\mathbf{N}}^{\max}, N_{\mathbf{N}}, L_{\mathbf{D}}) \leq 0$ for this remaining special case.

References

Abel, A., 1990. Asset prices under habit formation and catching up with the Joneses. Amer. Econ. Rev. 80, 38-42.

Abel, A., 2005. Optimal taxation when consumers have endogenous benchmark levels of consumption. Rev. Econ. Stud. LXXXVIII, 2009–2042.

Andersson, F.W., 2008. Is concern for relative consumption a function of relative consumption? J. Socio-Econ., 353-364.

Andrade, E., Ho, T., 2009. Gaming emotions in social interactions. J. Cons. Res. (December), 539-550.

Apesteguia, J., Huck, S., Oechssler, J., 2007. Imitation - theory and experimental evidence. J. Econ. Theory 136, 217-235.

Back, T., Fogel, D.B., Michalewicz, Z. (Eds.), 1997. Handbook of Evolutionary Computation. Oxford University Press.

Banerjee, A.V., 1992. A simple model of herd behavior. Quart. J. Econ. 107, 797-817.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Economy 100, 992-1026.

Blanchflower, D., Oswald, A., in press. Is well-being U-shaped over the life cycle? Social Sci. Med.

Boskin, M., Sheshinski, E., 1978. Optimal redistributive taxation when individual welfare depends upon relative income. Quart. J. Econ. 92, 561-589.

Boyd, R., Gintis, H., Bowles, S., Richerson, P., 2003. The evolution of altruistic punishment. Proc. Natl. Acad. Sci. 100, 3531-3535.

Campbell, J., Cochrane, J., 1990. By force of habit: A consumption-based explanation of aggregate stock market behavior. J. Polit. Economy 107, 205-251.

Clark, A.E., 1999. Are wages habit-forming? Evidence from micro data. J. Econ. Behav. Organ. 39, 527-539.

Clark, A.E., Oswald, A., Warr, P., 1996. Is job satisfaction u-shaped in age? J. Occup. Organ. Psychol. 69, 57-81.

Clark, A.E., Frijters, P., Shields, M., 2008. Relative income, happiness, and utility: an explanation for the Easterlin paradox and other puzzles. J. Econ. Lit. 46, 95-144.

Cole, H., Mailath, G., Postlewaite, A., 1995. Incorporating concern for relative wealth into economic models. Fed. Reserve Bank Minneapolis Quart. Rev. 19, 12-21.

D'Ambrosio, C., Frick, J., 2007. Income satisfaction and relative deprivation: an empirical link. Social Indicators Res. 81, 497-519.

Denrell, J., 2008. Indirect social influence. Science 321, 47-48.

Diener, E., Lucas, R., Scollon, C., 2006. C. N. Amer. Psychol. 71, 305-314.

Di Tella, R., Macculoch, R.M., Oswald, A., 2003. The macroeconomics of happiness. Rev. Econ. Statist. 95, 809-828.

Duesenberry, J., 1949. Income, Saving and the Theory of Consumer Behavior. Harvard University Press.

Easterlin, R., 1995. Will raising the incomes of all increase the happiness of all? J. Econ. Behav. Organ. 27, 35-48.

Easterlin, R., 2003. Explaining happiness. Proc. Natl. Acad. Sci. 100, 11176-11183.

Easterlin, R., Plangno, A., 2008. Life satisfaction and economic conditions in East and West Germany pre- and post-unification. J. Econ. Behav. Organ. 68, 433-444.

Fehr, E., Schmidt, K., 1999. A theory of fairness, competition and cooperation. Quart. J. Econ. 114, 817-868.

Ferrer-i Cabonell, A., 2005. Income and well-being: an empirical analysis of the comparison income effect. J. Public Econ. 89, 641-659.

Fowler, J., Christakis, N., 2008. Dynamic spread of happiness in a large social network. Br. Med. J. 337, a2338.

Frank, R., 1999. Luxury Fever. Free Press.

Frank, R., 2000. Why living in a rich society makes us feel poor. New York Times. October 15 magazine section.

Frederick, S., Loewenstein, G., 1999. Hedonic adaptive. In: Well-Being: The Foundations of Hedonic Psychology. Russell Sage Foundation, New York, pp. 302-329.

Frey, B., Stutzer, A., 2002. What can economists learn from happiness research? J. Econ. Lit. 40 (2), 402-435.

Frey, B., Stutzer, A., 2005. Testing theories of happiness. In: Bruni, L., Porta, P. (Eds.), Economics and Happiness, Framing the Analysis, pp. 116-146.

Frey, B., Stutzer, A., 2007. Should economic happiness be maximized? ssrn.com/abstract=936289.

Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games. MIT Press, Cambridge, MA.

Gardes, F., Merrigan, P., 2008. Individual needs and social pressure: evidence on the Easterlin hypothesis. J. Econ. Behav. Organ. 66.

Hart, S., 2005, Adaptive heuristics, Econometrica 73, 1401–1430,

Henrich, J., Boyd, R., 2001. Why people punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. J. Theoret. Biol. 208, 79-89.

Kahneman, D., Krueger, A., 2006. Amer. J. Econ. Perspect. 20, 3-24.

Kahneman, D., Krueger, A., Schkade, D., Schwarz, N., Stone, A., 2004. A survey method for characterizing daily life experience: the day reconstruction method. Science 306, 1776-1780.

Kahneman, D., Krueger, A., Schkade, D., Schwarz, N., Stone, A., 2006. Would you be happier if you were richer? A focusing illusion. Science 312, 1908–1910. Kirkpatrick, S., Gelatt, C.D.J., Vecchi, M.P., 1983. Optimization by simulated annealing. Science 220, 671-680.

Luttmer, E., 2005. Neighbors as negatives: relative earnings and well-being. Quart. J. Econ. 120, 963-1002.

Macready, W.G., Wolpert, D.H., 1998. Bandit problems and the exploration/exploitation tradeoff. IEEE Trans. Evol. Comput. 2, 2-22.

Netzer, N., in press. Evolution of time preferences and attitudes towards risk. Amer. Econ. Rev.

Neumark, D., Postlewaite, A., 1998. Relative income concerns and the rise in married women's employment. J. Public Econ. 70, 157-183.

Pollak, R., 1976. Interdependent preferences. Amer. Econ. Rev. 66, 309-320.

Postlewaite, A., Cole, H., Mailath, G., 1995. Incorporating status concerns in economic models. Quart. Rev. Minneapolis Fed. Reserve Bank. Available at Minneapolis Federal Reserve Bank.

Rayo, L., Becker, G., 2007. Evolutionary efficiency and happiness. J. Polit. Economy 115, 302-337.

Reichhardt, T., 2006. Well-being research: a measure of happiness. Nature 444, 418-419.

Robson, A., 2001. The biological basis of economic behavior. J. Econ. Lit. 39, 11-33.

Author's personal copy

D.H. Wolpert / Games and Economic Behavior 69 (2010) 458-474

Samuelson, L., 2004. Information-based relative consumption effects. Econometrica 72, 93-118.

Schlag, K., 1999. Which one should I imitate? J. Math. Econ. 31, 493-522.

Schwarz, N., Strack, F., 1999. Reports of subjective well-being: Judgmental processes and their methodological implications. In: Well-Being: The Foundations of Hedonic Psychology. Russell–Sage, pp. 61–84.

Selten, R., Ostmann, A., 2001. Imitation equilibrium. Homo Oecon. 43, 111–149.

Solnick, S.J., Hemenway, D., 1998. Is more always better?: A survey on positional concerns. J. Econ. Behav. Organ. 37, 373-383.

Stevenson, B., Wolfers, J., 2008. Economic growth and subjective well-being: reassessing the Easterlin paradox. Working paper.

Takahashi, H., Kato, M., Matsuura, M., Mobbs, D., Suhara, T., Okubo, Y., 2009. When your gain is my pain and your pain is my gain: neural correlates of envy and schadenfreude. Science 323, 937–939.

Tversky, A., Griffin, D., 1991. Endowment and contrast in judgments of well-being. In: Zeckhauser, R.J. (Ed.), Strategy and Choice. MIT Press, pp. 297–318. Vega-Redondo, F.A., 1997. The evolution of walrasian behavior. Econometrica 65, 375–384.

Winter, E., Garcia-Jurado, O., Mendez-Naya, J., Mendez-Nay, L., 2009. Mental equilibrium and rational emotions. Discussion Paper Series dp521. Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem.

Wolpert, D., Jamison, J., Newth, D., Harre, M., submitted for publication. Schelling formalized: Strategic choices of non-rational persons. SSRN:1172602.

Yang, Y., 2008. Social inequalities in happiness in the U.S. 1972-2004: an age-period-cohort analysis. Amer. Sociol. Rev. 73, 204-226.

Zink, C., et al., 2008. Know your place: Neural processing of social hierarchy in humans. Neuron 58, 273-283.