Keeping up with the Joneses and Income Inequality

Richard C. Barnett
Villanova University

Joydeep Bhattacharya
Iowa State University

Helle Bunzel
Iowa State University

March 13, 2009

Abstract

We study a variant of the conventional keeping-up-with-the-Joneses setup in which heterogeneous-ability agents care both about consumption and leisure and receive an utility premium if their consumption exceeds that of the Joneses'. Unlike the conventional setup in which all agents are assumed to want to participate in the rat race of staying ahead of the Joneses, our formulation explicitly permits the option to drop out. Mean-preserving changes in the spread of the underlying ability distribution, via its effect on the economy-wide composition of rat-race participants and drop-outs, have important consequences for induced distributions of leisure and income, consequences that are unobtainable using conventional keeping-up preferences.

Keywords: leisure distribution, rat race, keeping up with the Joneses, income inequality
JEL classifications: J 22, E2

> "It must strike the modern man as strange that the Ten Commandments end with a prohibition against coveting[1], something which is neither criminal nor punishable in any society... A man who does not want to raise his standard of living to that enjoyed by the other fellow is considered ambitionless, and as such reprehensible. To be sure, if coveting results only in unconstructive greed, ostentation, or pushiness (e.g., "keeping up with the Joneses"), and therefore entails an unseemly response to the potentially desirable stimulus of coveting, it may be frowned upon. But when coveting impels us to greater effort, so that we may rise constructively toward the level of our more affluent neighbors, we are on the path that universally leads to approbation. Why, then, does the Bible make a great issue of coveting, grouping it with such evil offenses as murder?"
>
> *Cyrus Gordon (1963)*

1 Introduction

Human beings have always had an ambivalent attitude towards the commonly-observed practice of keeping up with the Joneses. The reason is simple. Keeping up with our neighbors is a worthy goal insofar as it acts as a desirable stimulus, impelling us to greater effort, as noted in the passage above. On the flip side though, the very act of keeping up exposes us to the influence of the Joneses on our lives, and this we may not like. Predictably, this dialectic elicits one of two reactions in people; some choose to run on the "hedonic treadmill" of staying ahead of some reference group – the "rat race" – while others find such influence unwelcome, and opt out. In this paper, we study a preference formulation that captures this idea – that some people, acting in their own self interest, *choose* to participate in the rat race, while others choose not to. We use these preferences to study how individual choices about participation produce consequences for aggregate income and leisure distributions in an an economy. Put another way, we ponder, what would the income/leisure distribution look like in an economy in which agents can choose to either join the rat race or forsake it?

Our framework is a variant on keeping up with the Joneses preferences. In such setups,

[1] The last of the Ten Commandments, reads as follows "You shall not covet your neighbor's house; you shall not covet your neighbor's wife, or male or female slave, or ox, or donkey, or anything that belongs to your neighbor."

the utility of any agent is influenced not just by her own consumption but also via external habits — the consumption of the Joneses.[2] Interpreted in a certain way, conventional keeping-up preference formulations share the underlying presumption that everyone welcomes such influence from the Joneses, and wants to keep up with them if they could; the issue of "what if they don't?" has not received much attention.

In our setup, people care both about consumption and leisure and when it comes to participation in the rat race, they have a choice. For example, they can choose to build a 2000 square feet house in a neighborhood with mainly 1200 square feet houses because they receive some 'extra' satisfaction knowing their house is bigger than their neighbors'. In the bargain, they sacrifice some leisure; more importantly, they expose themselves to a kind of 'jealousy' (Dupor and Liu, 2003) they would feel if their neighbors responded by building 3000 square feet houses. Alternatively, they can choose to build a 1000 square feet house in that same neighborhood, give up the aforementioned utility premium, and enjoy more leisure.[3] Effectively, the presence of the latter option offers individuals a way to escape the influence of their neighbors – from *having to* get caught up in the rat race.[4] Depending on one's innate ability to convert time into income, different people make different choices: some choose to work a lot and build big houses while others choose to work less and live in smaller houses.

More formally, assume everyone cares about consumption and leisure, and preferences over consumption are level-dependent; only when their consumption exceeds a certain benchmark or threshold, agents get an extra utility boost from beating the benchmark, otherwise not. The benchmark is assumed to be proportional to the cross-sectional mean

[2] Models with interdependent preferences have become immensely popular in many areas of economics and finance. Classic references in this literature include the seminal pieces by Abel (1990) and Gali (1994), followed by more recent work by Ljungqvist and Uhlig (2000), Chan and Kogan (2002), and Dupor and Liu (2003), among numerous others.

[3] A November 2004 poll conducted by the US News and World Report found that 48 per cent of Americans have done at least one of the following in the past five years: "cut back their hours at work, declined or did not seek a promotion, lowered their expectations for what they need out of life, reduced their work commitments, or moved to a community with a less hectic way of life."

[4] That the negative influence of neighbors may create inefficiency has been studied. Corneo and Jeanne (1997), in a model with preferences for relative wealth, showed that status concerns induce overaccumulation of capital but higher growth.

consumption in the economy; evidently, it is this endogeneity of the benchmark that makes the subsequent analysis especially challenging.[5] We embed these preferences into a model economy inhabited by heterogenous innate-ability agents. The level-dependence in consumption introduces a non-convexity in overall preferences – even though keeping-up is restricted to consumption, its effects are felt on leisure too. While the marginal benefit from an additional unit of leisure is the same under the option of participation or of dropping out, the accompanying marginal cost of lost consumption is *higher* under participation. A rise in the benchmark (average consumption) raises the marginal valuation of 'effective' consumption over leisure. A rat-race participant responds by substituting out of leisure (into work), earning more income, and raising her own consumption. A drop-out is spared all this.

Because of the aforediscussed nonconvexity, inequality in innate ability – what we'll refer to as fundamental inequality – induces inequality in other derived variables, such as income and leisure. We show that the means of these induced distributions depend on the *entire* innate-ability distribution — as such, fundamental inequality (captured, for example, by the second moment) affects the first *and* second moments of the income and leisure distributions. In particular, when ability-inequality rises (modeled as a mean-preserving increase in the spread of the innate-ability distribution), more people may participate in the rat race, causing average labor supply and income to increase. The variance of labor supply and income may also increase.[6]

Our analysis informs the more general discussion, posed for example, by Banerjee and Newman (1991), Galor (1996), and others: "Does a market economy exacerbate the level of inequality in wealth and income, or does it merely reproduce variation in individual

[5] We are not the first to use a version of these preferences. Creedy (1998), following up on pioneering work by Lewis and Ulph (1988), analyze consumption and leisure choices in a model in which agents get a utility premium after their consumption crosses an exogenous threshold. This preference formulation has also received some indirect empirical backing. For example, Dynan and Ravina (2007) analyze time-series General Social Survey data for the U.S. and find that self-reported "happiness is higher the greater one's income is relative to one's neighbors, *with the effect concentrated among those with above-average incomes.*" (italics ours)

[6] Kuhn and Lozano (2008) find that "between 1970 and 1990, the share of employed-for-pay U.S. men who worked more than 48 hours in the U.S. Census reference week rose from 15.4% to 23.3%....the increase in long hours was [...] concentrated among highly educated, high-wage, and older men, and it was largely confined to workers paid on a salaried basis."

attributes?" They go on to argue that "[u]nder the complete markets of Arrow and Debreu, there are essentially two sources of variation in income, namely differences in preferences and differences in endowments...".[7] These features are not present in our model: the model is static, there is no risk and no market aggregation to speak of, and everyone has identical preferences. Agents differ only in their endowments of innate-ability (fundamental inequality). Yet, it is precisely the non-convexity in preferences, brought on by the choice to drop out, that leads people with different endowments to take disproportionately different decisions on labor supply, thereby generating disproportionate differences in income.[8] In fact, in our setup income inequality is greater than the underlying fundamental inequality. In an otherwise-identical model with conventional keeping-up preferences (where everyone is *assumed* to be in the rat race), the first moment of the income distribution would depend solely on the first moment of the innate-ability distribution, and income equality simply mirrors the underlying fundamental inequality in endowments.

We go on to allow for the possibility that all agents have access to fair gambles – consumption lotteries – that work in the following way: agents can place a bet with their entire income and get, either a high prize (which allows them to consume more than the benchmark) or a low prize (which forces them to drop out), and their odds depend on the size of the bet placed. While such lotteries 'convexify' preferences to a large extent, they cannot achieve strict convexity – after all, a lottery replaces the non-convex portion with a linear segment. The implication is that fundamental inequality continues to matter for the first and second moments of the induced distributions. As before, when ability-inequality rises, average labor supply and income may increase and so too their variances. Of course, the pass-through from fundamental to income inequality is lower in the presence of lotteries.

[7] There is a parallel literature – motivated by the so-called Easterlin paradox – that studies the distribution of happiness or utility (given an exogenous income distribution) in models with status concerns or keeping-up preferences. See Hopkins (2008) for a survey.

[8] The literature connecting income distribution and economic development (see Galor (1996) or Benabou (1996)), which follows the dominant capital-market-imperfection paradigm, finds that second moment of the income distribution affects the first moment, provided the production technology of human capital or final output is non-convex. Our approach investigates the role played by non-convexities in preferences thereby sidestepping market frictions or indivisibilities.

In studying voluntary risk-taking behavior in the presence of lotteries, especially in connection with endogenous income inequality, we are following a long line of research associated with Friedman and Savage (1948), with more recent contributions by Freeman (1996), Sadler (2000), and Becker, Murphy, and Werning (2005). In Sadler (2000), agents face two technologies, a high-return modern technology and a low-return traditional technology, and entry into the former requires an up-front payment of a fixed cost. In such a setting, some poor risk-averse agents may save up to participate in a lottery; those who win receive a prize big-enough to pay the fixed cost, and thereby escape poverty. In Becker, Murphy, and Werning (2005), agents care both about consumption and status, the latter being a good sold in 'hedonic status markets' but in lumpy amounts – hence the nonconvexity. In such a setting, both rich and poor agents would voluntarily participate in lotteries; winners get both higher consumption and higher status, and losers get lower consumption and lower status, thereby generating income inequality.

Our work is also in line with the literature investigating how initial inequality affects future growth in the presence of status concerns. In an important recent contribution, García-Peñalosa and Turnovsky (2008), following up on García-Peñalosa and Turnovsky (2006), study a Ramsey model and show that under conventional keeping-up-with-the-Joneses preferences, simple restrictions on homogeneity of utility (in consumption and leisure) or separability of the reference consumption level, are enough to make the initial distribution of capital irrelevant for the aggregate equilibrium. Our focus is more on what these preferences imply for inequality, and to what extent they amplify fundamental inequality. In this sense, our work is closer to Kawamoto (2009) who studies a two-class, two-period overlapping generations model in which all agents have status concerns, and an agent's human capital depends on learning effort in youth and the (warm-glow induced) educational spending by parents. Young agents in richer households receive higher educational expenditure than those in poorer households; the latter spend more time on learning though. Status-seeking affects the learning incentives of young agents, and via this channel, the long-run income distribution.[9]

[9] Long and Shimomura (2004) showed that, in the Ramsey model with initial-wealth heterogeneity, if relative wealth concerns appear in the utility function, poorer agents may catch up with wealthier agents.

The plan for the rest of the paper is as follows. Section 2 outlines the model and characterizes the set of agents that participate in the rat race and those that drop out. Section 3 studies the impact of a change in the spread of the innate-ability distribution on the induced distributions; it contains our main results. Section 4 introduces consumption lotteries into the framework. Concluding remarks are contained in Section 5. Proofs of major results are relegated to the appendix.

2 The Model

2.1 Preliminaries

Consider a simple static model in which people (indexed by i) work for a fixed wage w. Agents differ in their innate ability (equivalently, effective units of time endowment), $1+h^i$. Agent i draws h^i from an exogenously-specified continuous distribution with cdf $\mathcal{F}\left(h^i\right)$ and density $f\left(h^i\right)$ with support $[h_l, h_u]$ where $0 \leq h_l < h_u < \infty$. Let $\bar{h}$ denote the mean value of h, and let σ denote the standard deviation of the distribution of h.

Agents' preferences over consumption (c) and leisure (l) are summarized by the utility function $U(c,l)$ as follows:

$$U(c,l) \equiv \begin{cases} u(c+\theta(c-\kappa)) + \phi\, v(l) & \text{if } c > \kappa \\ u(c) + \phi\, v(l) & \text{if } c \leq \kappa \end{cases}.$$

Here $\theta \geq 0$ is a scalar capturing the importance of staying ahead (see below) and $\kappa > 0$ is the to-be-specified consumption benchmark the agent takes as given and aspires to beat. Also, $\phi > 0$ is a parameter capturing the relative importance of leisure.[10] For now, assume the functions u and v are increasing and concave.

[10] Our preference formulation shares many similarities with those studied in Dupor and Liu (2003). In their paper, an increase in per-capita consumption causes a) marginal utility from own consumption to rise relative to leisure (what they term "Keeping up with the Joneses" or KUJ) and/or b) ceteris paribus, it lowers the agent's own utility ("jealousy"). In our setup, those who consume more than the per-capita consumption behave exactly as people with KUJ and jealousy preferences in Dupor and Liu (2003). The critical difference is that Dupor and Liu (2003) do not allow agents to drop out. In our framework, as will become clear below, those who choose to drop out, do so, because they do not wish to work as hard as those who participate.

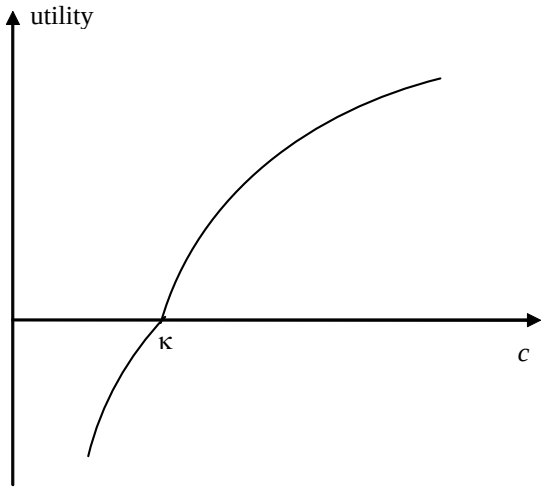

Figure 1a: The utility from consumption

For the present, focus attention on the "consumption portion" of these preferences. If $\theta = 0$, there is no utility premium to beating κ, and as such, κ is irrelevant. In this case, the agent has a single utility function defined over the entire domain of consumption – the standard textbook setting. However, when θ is positive, the agent has a two-piece utility function, the pieces defined over non-overlapping portions of the consumption domain. When her consumption exceeds κ, the agent receives utility from it as well as a bonus from beating κ (success in the rat race). This is the satisfaction from having "made it in life". To carry on with the house example, *only* when her house is bigger than her neighbors' does the agent get extra utility from this knowledge; otherwise, all she gets is utility from the house itself.[11] Two other points deserve mention. First, the utility she gets from "housing services" is not separable from the utility premium she gets because her house is bigger than those of her neighbors'. Second, *for fixed leisure*, the marginal utility of consumption is higher to the right of κ than to the left. Consequently, a kink (illustrated in Figure 1a),

[11] We ignore the possibility that some agents may feel worse off if their own consumption exceeds their neighbors' average – inequity aversion (Hopkins, 2008). This is justified in the current paper, because unlike in status-contest models, there is no notion that some consumption is ostentatious.

appears at κ.

In other respects, this formulation of preferences is fairly standard. Utility from consumption and leisure is assumed to be continuous and concave in each of the two segments. As is common in specifications of keeping-up preferences, when $c > \kappa$, higher κ hurts the agent because it is harder to stay ahead of the neighbors. Also, for the range $c > \kappa$, our preferences exactly match the subtractive form of the external habits formulation used in Abel (1990), and discussed in Hopkins (2008).[12]

The rest of the model is deliberately kept extremely simple. Agent i allocates her effective time endowment between work (x) and leisure (l), and so $x^i + l^i = 1 + h^i$ holds. Then agent i's budget constraint is given by $c^i = wx^i$. For future reference, note that in this highly-simplified setting, consumption, income, and labor supply are essentially one and the same thing.

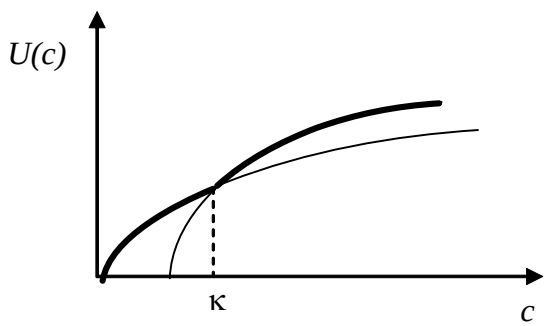

Figure 1b: The nonconvexity

One can use the budget constraint to solve for leisure and rewrite utility in terms of consumption alone. Figure 1b plots $U(c)$ for agent i, assuming a given κ : in the figure, it is drawn as the upper envelope (thick black curve) of the utilities in the two consumption domains, $c \leq \kappa$, and $c > \kappa$. As is clear, the kink at κ creates a non-convexity in utility.[13]

[12] Luttmer (2005) finds convincing evidence that "an increase in neighbors' earnings and a similarly sized decrease in own income each lead to a reduction in happiness of about the same order of magnitude."

[13] Clark and Oswald (1998) get around this nonconvexity by postulating a utility function on a costly action a, where $U(a) = \alpha\, u(a - a^*) + (1 - \alpha)\, u(a) - c(a)$ where a^* is, say, the mean action in society, $c(a)$ is the convex cost function, and $\alpha \in [0, 1]$.

Notice that the non-convexity in the consumption portion of the preferences (see Figure 1a) translates into a non-convexity in overall utility. For agent i, the theoretical possibility arises that she is made better off in an ex-ante sense were she to accept a fair gamble, one which gives her consumption above κ with some probability and below κ with the remaining probability. We appeal to concerns of tractability and ignore the existence of such consumption lotteries for now; we return to this important issue in Section 4 below.

In what follows, people enjoying consumption greater than κ are informally dubbed *participants* – those participating in the rat race of staying ahead of their neighbors; all others are *drop-outs*, a catch-all phrase intended to replace the mouthful "non-participant in the rat race of keeping up with the Joneses".[14] The important thing is that when $c < \kappa$, the agent perceives no *direct* influence of the Joneses on her actions – she is out of the rat race.

As is evident, some agents may not have innate ability high enough to finance a level of consumption higher than κ, *even* when they enjoy no leisure. This implies that agents with h such that

$$h^i < h^F \equiv \max\left(\left(\frac{\kappa}{w} - 1\right), h_l\right)$$

holds are those who are forced (by their low ability) to stay out of the rat race.[15] All others face a choice; they can work a lot to earn enough income so as to beat κ, or they can enjoy more leisure and drop out.

We now consider the possibility that some agents may *choose* to drop out of the rat race. In other words, restrict attention to those i for whom $h^i > h^F$. When such an agent is *contemplating* dropping out, her leisure and consumption are computed by solving

$$\max_{c,l} \quad u(c) + \phi v(l)$$

[14] A word on terminologies is in order. First, "staying out" of the rat race and "dropping out" are exactly equivalent; the latter phrase has no pejorative implication and does not imply pursuit of an "alternative" reclusive Thoreau-esque lifestyle. Second, the term "participant" always refers to participant in the rat race. Finally, our notion of a rat race is not the same as that in the literature on adverse selection; see for example, Landers, Rebitzer, and Taylor (1996). There, agents may try to signal their high type by agreeing to work longer than most, and firms may use this willingness to work long hours as an indicator of high type.

[15] Of course, if $h^F = h_l$, no agent is forced out of the rat race.

subject to $c_d^i = w\left(1 + h^i - l_d^i\right)$, taking κ as given; the subscript d refers to "dropping out". This yields the following first order condition:

$$w\ u'\left(w\left(1 + h^i - l_d^{i*}\right)\right) = \phi\, v'\left(l_d^{i*}\right). \tag{1}$$

A *valid* interior optimum l_d^{i*} additionally satisfies $w\left(1 + h^i - l_d^{i*}\right) \equiv c_d^{i*}$. The consumption choice will satisfy $c_d^{i*} \leq \kappa$ for all i who drop out. When that same agent contemplates joining the rat race, her leisure and consumption are computed by solving

$$\max_{c,l} \quad u\left(c + \theta\left(c - \kappa\right)\right) + \phi v\left(l\right)$$

subject to $c_r^i = w\left(1 + h^i - l_r^i\right)$, taking κ as given; the subscript r refers to "rat race participant". This yields

$$(1 + \theta)\, w\ u'\left((1 + \theta)\, w\left(1 + h^i - l_r^{i*}\right) - \theta\kappa\right) = \phi\, v'\left(l_r^{i*}\right). \tag{2}$$

A *valid* interior optimum l_r^{i*} here satisfies $w\left(1 + h^i - l_r^{i*}\right) \equiv c_r^{i*}$. The consumption choice will satisfy $c_r^{i*} > \kappa$ for all i who participate. It follows her consumption will be higher and leisure lower under participation than not. It is also clear from (1) and (2) that for agent i, while the marginal benefit from an additional unit of leisure is the same under the option of participation or of dropping out, the accompanying marginal cost of lost consumption is higher (because $\theta > 0$) under participation.

Ceteris paribus, a rise in κ reduces the effective consumption of a participant and raises her marginal valuation of effective consumption over leisure. This is the sense in which participation in the rat race entails allowing oneself to be negatively influenced by the consumption of a reference group – a complementarity also noted by Becker, Murphy, and Werning (2005). When the marginal valuation of effective consumption over leisure rises, a participant's effective consumption falls, and she counters by substituting out of leisure (into work), earning more income, and raising her own consumption. Needless to say, a drop-out is spared all this.

Under the usual concavity assumptions of u and v, it is relatively straightforward to verify that for any i, there is a unique l_d^* that solves (1) and a unique l_r^* that solves (2).

Also, l_d^* and l_r^* are each non-decreasing in h suggesting that leisure is a normal good for both participants and drop-outs.[16] Note that it is quite possible that a sufficiently-rich agent under the option of participation will have higher leisure than a sufficiently poor agent under the option of dropping out.

2.2 Who participates, who doesn't

It is useful to learn more about the identities of who chooses to participate in the rat race and who does not. To that end, and for reasons of analytical tractability, we henceforth assume a quasi-linear form[17]:

$$U(c,l) \equiv \begin{cases} (c+\theta(c-\kappa)) + \phi \ln l & \text{if } c > \kappa \\ c + \phi \ln l & \text{if } c \leq \kappa \end{cases}.$$

For this specification, the marginal utility of consumption of anyone who drops out is 1, and for anyone who participates, is constant at $1+\theta$. Let $\phi_u \equiv w(1+h_l)$ and

$$\phi \leq \phi_u. \tag{A.1}$$

It is easy to check that the first order conditions (1) and (2) imply that

$$l_d^i = \frac{\phi}{w}; \quad c_d^i = w\left(1 + h^i - \frac{\phi}{w}\right) \tag{3}$$

$$l_r^i = \frac{\phi}{w(1+\theta)}; \quad c_r^i = w\left(1 + h^i - \frac{\phi}{w(1+\theta)}\right) \tag{4}$$

hold. Then (A.1) implies that $c_d^i, c_r^i \geq 0$ and $l_d^i, l_r^i \leq 1 + h^i$ for each i. For these candidate solutions to be valid optima, we must verify that $c_d^i \leq \kappa$ and $c_r^i > \kappa$ hold. From (3)-(4), it follows that $l_r^i < l_d^j$ for any i, j implying that any agent who chooses to participate would enjoy less leisure than *any other* agent who chooses to drop out. The principal benefit of

[16] From (1), it follows that

$$w^2 u''(\cdot)\left(1 - \frac{\partial l_d^i}{\partial h^i}\right) = \phi v''(\cdot)\frac{\partial l_d^i}{\partial h^i} \Longleftrightarrow \frac{\partial l_d^i}{\partial h^i} = \frac{w^2 u''(\cdot)}{\phi v''(\cdot) + w^2 u''(\cdot)}.$$

Since $u'' \leq 0$ and $v'' \leq 0$, it follows that $\frac{\partial l_d^i}{\partial h^i} \geq 0$. Analogously, it follows that $\frac{\partial l_k^i}{\partial h^i} \geq 0$.

[17] Our results generalize straightforwardly to most standard forms for u and v.

our choice of a quasi-linear form for utility is now apparent. Neither l_d^i nor l_k^i depends on h_i; hence, from (3)-(4), it is clear that the leisure of *anyone* who participates in the rat race is less than that of someone who drops out. Additionally, note that κ does not enter c_d^i or c_r^i, a convenient feature from the standpoint of analytical tractability. This last feature is not true in general – see Appendix A as well as Section 4 below.

An agent will choose to stay out of the rat race if the indirect utility from staying out exceeds that from participating. Define $b\left(h^i;\kappa\right)$ to be the net benefit from participation. Then it is easily verified that

$$b\left(h^i;\kappa\right) = \frac{\theta\phi}{1+\theta} + \theta\left(w\left(1+h^i-\frac{\phi}{w\left(1+\theta\right)}\right)-\kappa\right) - \phi\ln\left(1+\theta\right). \tag{5}$$

The first term represents the additional utility from higher consumption received from participation; the second term is the utility 'kick' the agent gets from the very act of participation, and the third term is the net loss in utility incurred from needing to work harder so as to be able to stay in. Formally, agents for whom the following inequality holds choose to participate: $b\left(h^i;\kappa\right) \geq 0$ for any $h^i \geq \hat{h}$, where $\hat{h}$ is defined by $b\left(\hat{h};\kappa\right) \equiv 0$. Straightforward algebra shows[18]

$$h^i > \hat{h} \equiv \left(\frac{\kappa}{w}-1\right) + \frac{\phi\ln\left(1+\theta\right)}{\theta w} \geq h^F. \tag{6}$$

Notice that (5) implies that $b\left(h^i;\kappa\right) = \theta w\left(1+h^i-\kappa\right) - \phi\ln\left(1+\theta\right)$. It follows that as ϕ rises, the net loss in utility incurred from needing to work harder (so as to be able to participate) rises. This raises the cut-off, $\hat{h}$; only those who get a sufficiently high utility premium from participation can rationally forego the accompanying loss in utility from having to work harder.

We now verify that $c_d^i \leq \kappa$ for anyone who chooses to drop out and $c_r^i > \kappa$ for any agent who chooses to participate, and hence any agent who chooses to participate would enjoy more consumption than *any other* agent who chooses to drop out.

[18] In passing, note that when $\phi = 0$, $\hat{h} = h^F$; i.e., if no one derives utility from leisure, then a person participates if and only if they *can*. When ϕ is positive, people care about leisure and only then does dropping out become a choice for some.

Lemma 1 *If $h^i \leq \hat{h}$, agent i drops out and chooses consumption $c_d^i \leq \kappa$; if $h^i > \hat{h}$, agent i participates in the rat race and chooses consumption $c_r^i > \kappa$.*

The previous analysis suggests that, for a given reference point κ, the entire population gets divided into two categories. Those with effective time endowment $h^i < \hat{h}$ drop out of the rat race (either by choice or by force) and they are the "drop-outs"; those with $h^i \in \left[\hat{h}, h_h\right]$ stay in, and they are the participants. Formally, $\mathcal{D} \equiv \left\{i : h^i < \hat{h}\right\}$ is the fraction of agents in the population who drop out of the rat race and $\mathcal{K} \equiv \left\{i : h^i \geq \hat{h}\right\}$ is the set of participants. Even if κ does not directly affect c_d^i or c_r^i, changes in κ affect $\hat{h}$, and via this channel, the size of the sets $\mathcal{D}$ and $\mathcal{K}$. Also note, if κ is a fixed constant, $\hat{h}$ would be independent of the distribution of h.

Denote by $\bar{c}$, the cross-sectional mean consumption in the economy. Then,

$$\bar{c} = \int_{h_l}^{\hat{h}} c_d^i \, d\mathcal{F}\left(h^i\right) + \int_{\hat{h}}^{h_u} c_r^i \, d\mathcal{F}\left(h^i\right).$$

Using the expressions for c_d^i and c_r^i derived earlier, it is straightforward to check that

$$\bar{c} = w\left(1+\bar{h}\right) - \phi \int_{h_l}^{\hat{h}} d\mathcal{F}\left(h^i\right) - \frac{\phi}{(1+\theta)} \int_{\hat{h}}^{h_u} d\mathcal{F}\left(h^i\right) = w\left(1+\bar{h}\right) - \frac{\phi\left(1+\theta\,\mathcal{F}\left(\hat{h};\sigma\right)\right)}{(1+\theta)}, \quad (7)$$

where $\mathcal{F}\left(\hat{h};\sigma\right)$ is the measure of agents (for given spread, σ) with ability less than or equal to the cut-off $\hat{h}$, or more succinctly, the probability of an agent being in $\mathcal{D}$. Since $\frac{\phi\left(1+\mathcal{F}\left(\hat{h};\sigma\right)\theta\right)}{(1+\theta)} \leq \phi$, it follows from (A.1) that $\bar{c} > 0$. It is clear from (7) that average consumption in the economy falls when $\mathcal{F}\left(\hat{h};\sigma\right)$ rises. This is intuitive because those in the set $\mathcal{D}$ consume less than those in $\mathcal{K}$ and hence an increase in the probability mass of the former group lowers average consumption. Also, notice that $c_i = wx_i = w\left(1+h^i-l_i\right)$ $\forall i$, implying

$$\int_i c_i = \int_i w\left(1+h^i-l_i\right) \Rightarrow \bar{c} = w\left[\left(1+\bar{h}\right)-\bar{l}\right]. \quad (8)$$

From (8), it is clear that average consumption and average leisure are inversely related.

A few additional words about eq.(7) are in order. Notice when $\phi = 0$, people do not value leisure and hence average consumption in the economy is just $w\left(1+\bar{h}\right)$, the income

of the agent with average h. When $\phi > 0$, but $\theta = 0$, people care about leisure but not about the rat race. In that case, preferences become homothetic for the entire range of consumption, and hence the distribution of h does not influence the determination of average consumption, i.e., $w\left(1+\bar{h}\right) - \phi < w\left(1+\bar{h}\right)$ holds.[19] Notice when $\theta > 0$ holds, then $\mathcal{F}\left(\hat{h};\sigma\right) \leq 1 \Rightarrow \frac{1+\theta\mathcal{F}\left(\hat{h};\sigma\right)}{(\theta+1)} \leq 1$, and hence $\phi\left(\frac{1+\theta\mathcal{F}\left(\hat{h};\sigma\right)}{(\theta+1)}\right) \leq \phi$ obtains, implying (see eq. (7)) that average consumption is higher when people care about keeping up. In short, $\bar{c}_{\phi=0} > \bar{c}_{\phi>0,\theta>0} > \bar{c}_{\phi>0,\theta=0}$ obtains.

It remains to specify the societal consumption benchmark, κ. While there is obviously no uniquely acceptable way to choose κ, we posit without loss of much generality, that $\kappa = \lambda\bar{c}$, where λ is a positive scalar. Using $\kappa = \lambda\bar{c}$ in (6) along with (7), we get a single equation in the cut-off, $\hat{h}$:

$$\left(\hat{h}-\bar{h}\right) + \frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right)\mathcal{F}\left(\hat{h};\sigma\right) = \frac{1}{w}\left(\frac{\phi\ln\left(1+\theta\right)}{\theta} - \frac{\lambda\phi}{(1+\theta)}\right) + (\lambda-1)\left(1+\bar{h}\right). \quad (9)$$

The next lemma outlines a necessary and sufficient condition for the existence of a unique cut-off.

Lemma 2 *Denote*

$$\underline{\lambda} \equiv \frac{1+h_l - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta}}{1+\bar{h} - \frac{\phi}{w}\frac{1}{(1+\theta)}}; \quad \overline{\lambda} \equiv \frac{1+h_u - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta}}{1+\bar{h} - \frac{\phi}{w}}.$$

There exists a unique $\hat{h} \in [h_l, h_u]$ *iff*

$$\underline{\lambda} < \lambda < \overline{\lambda}. \quad (A.2)$$

Henceforth, we maintain Assumption (A.2). Notice that $\underline{\lambda} < 1 < \overline{\lambda}$ holds. An implication of Lemma 2 is that the population gets split *exactly* into two groups, $\mathcal{D}$ and $\mathcal{K}$.[20] What is additionally noteworthy is the dependence of $\hat{h}$ on the entire distribution, $\mathcal{F}$, not just $\bar{h}$ – this last feature is a consequence of the assumed dependence of κ on average consumption.

[19] Average consumption is lower when $\phi > 0$. The opportunity cost of leisure in terms of goods is $\frac{\phi}{1+\phi}$ and this is subtracted from average consumption.

[20] This renders impossible a situation, for example, in which an agent with a very high h chooses to drop out and enjoy leisure; the loss in utility from the resultant decline in her consumption would simply hurt her too much.

To foreshadow, a change in the exogenous innate-ability distribution will change $\hat{h}$ via the change in κ (and hence, influence the size of the sets $\mathcal{D}$ and $\mathcal{K}$), which in turn will affect the induced distributions of income/consumption and leisure. Indeed, the effect will pervade beyond the first moments of these induced distributions. As will be evident, a model in which agents have no option to drop out will not generate such outcomes.

2.3 Moments of the induced distributions

For future use, define average leisure $\left(\bar{l}\right)$ in the economy as follows:

$$\bar{l} = \int_{h_l}^{\hat{h}} l_d^i \, d\mathcal{F}\left(h^i\right) + \int_{\hat{h}}^{h_u} l_r^i \, d\mathcal{F}\left(h^i\right). \tag{10}$$

This is the weighted average of the leisure enjoyed by those who drop out and those who stay in; in our case, these weights are endogenous since $\hat{h}$ is endogenous. Using the expressions for l_d^i and l_r^i derived earlier, it follows that

$$\bar{l} = \frac{\phi}{(1+\theta)\, w}\left[1 + \theta\mathcal{F}\left(\hat{h};\sigma\right)\right]. \tag{11}$$

Similarly, the variance of leisure may be calculated from

$$V(l) = \int_{h_l}^{\hat{h}} \left(\frac{\phi}{w} - \frac{\phi}{(1+\theta)\, w}\left[1 + \theta\mathcal{F}\left(\hat{h};\sigma\right)\right]\right)^2 d\mathcal{F}\left(h^i\right)$$
$$+ \int_{\hat{h}}^{h_u} \left(\frac{\phi}{w\,(1+\theta)} - \frac{\phi}{(1+\theta)\, w}\left[1 + \theta\mathcal{F}\left(\hat{h};\sigma\right)\right]\right)^2 d\mathcal{F}\left(h^i\right),$$

which, after routine simplification yields,

$$V(l) = \left(\frac{\phi}{w}\right)^2 \left(\frac{\theta}{(1+\theta)}\right)^2 \left(1 - \mathcal{F}\left(\hat{h};\sigma\right)\right) \mathcal{F}\left(\hat{h};\sigma\right). \tag{12}$$

Finally, denote by $\bar{y}$, the cross-sectional mean income in the economy. Then, since income and consumption are identical, $\bar{y} = \bar{c}$. The variance of income (the same as the variance of consumption) is given by

$$V(y) = \int_{h_l}^{\hat{h}} \left(w\left(1 + h^i - \frac{\phi}{w}\right) - w\left(1+\bar{h}\right) + \phi\int_{h_l}^{\hat{h}} d\mathcal{F}\left(h^i\right) + \frac{\phi}{(1+\theta)}\int_{\hat{h}}^{h_u} d\mathcal{F}\left(h^i\right)\right)^2 d\mathcal{F}\left(h^i\right)$$
$$+ \int_{\hat{h}}^{h_u} \left(w\left(1 + h^i - \frac{\phi}{(1+\theta)\, w}\right) - w\left(1+\bar{h}\right) + \phi\int_{h_l}^{\hat{h}} d\mathcal{F}\left(h^i\right) + \frac{\phi}{(1+\theta)}\int_{\hat{h}}^{h_u} d\mathcal{F}\left(h^i\right)\right)^2 d\mathcal{F}\left(h^i\right)$$

which, after routine simplification yields,

$$V(y) = w^2\sigma^2 + \left(\frac{\theta}{(1+\theta)}\right)^2 \phi^2 \mathcal{F}\left(\hat{h};\sigma\right)\left(1 - \mathcal{F}\left(\hat{h};\sigma\right)\right). \tag{13}$$

Among the first things to notice is that the mean and variance of income and leisure depend on the entire distribution of h – as noted earlier, this is a consequence of the dependence of κ on average consumption. Next, notice from (13) that variance of income is clearly more-than-proportionately higher than the variance of the underlying innate-ability distribution, σ^2. And this is true when $\theta > 0$, i.e., when people care about staying ahead of the Joneses. Put another way, the option to drop out of the rat race *exacerbates* fundamental inequality because it increases income inequality disproportionately – the pass-through (if you will) from fundamental to income inequality is more than one-to-one. A model with assumed/forced keeping-up cannot generate this. In such a setup, average leisure (assuming the quasi-linear preference structure of Section 2.2) would not even depend on the first moment of the innate-ability distribution, and the second moment of the income distribution (see (13)) would exactly mirror the same for the innate-ability distribution – merely 'reproducing variation in individual attributes'.

3 Changes in the innate-ability distribution

We are now in a position to ask the question, how do changes in the exogenous innate-ability distribution affect the induced distributions of income and leisure in the economy? Our focus is restricted to a mean-preserving increase in the spread (MPIS) of the ability distribution. Such an increase redistributes mass away from the center and towards the tails. For a fixed cut-off, the mass redistribution would have a first-order effect on the number of drop-outs. In our setup, since the cutoff *itself* depends on average consumption, there would be an additional (second-order) effect; that is, the first-order change in average consumption would change the cutoff and hence the mass of drop-outs, and ultimately the means and variances of the induced distributions of income and leisure.

3.1 Position of the cut-off ability relative to the mean

It is well known that if a new unimodal distribution, $\mathcal{F}(h, \sigma_2)$, is obtained from an old one, $\mathcal{F}(h, \sigma_1)$, via a mean-preserving increase in spread from σ_1 to σ_2, it exhibits the single crossing property. That is, there exists an h^* such that $\mathcal{F}(h, \sigma_2) \gtreqless \mathcal{F}(h, \sigma_1)$ whenever $h \lesseqgtr h^*$. Also, when the mean is preserved, and $\mathcal{F}(.)$ is a symmetric and unimodal distribution, the single crossing takes place at the mean, i.e., $h^* = \bar{h}$.[21] See Figures 2a and 2b for an illustration. Intuitively, for a symmetric and unimodal distribution, a mean-preserving increase in spread redistributes the mass evenly away from the mean towards the tails, implying draws near the mean under the new distribution are less likely and draws near the tails are correspondingly more likely.

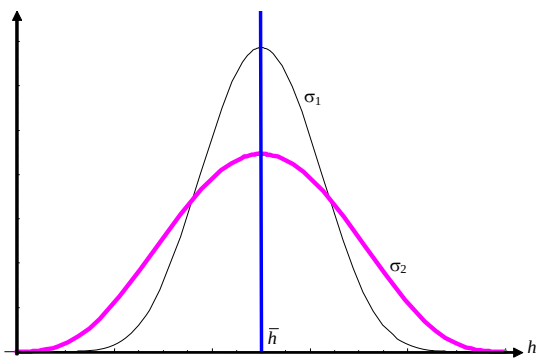

Fig 2a: PDF

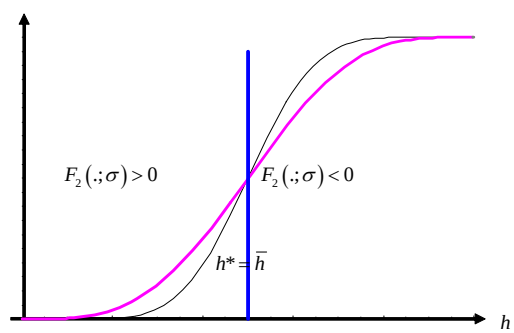

Fig 2b: CDF

The analytical results below are derived under the assumption that $\mathcal{F}(.)$ is a symmetric and unimodal distribution.[22] In that case, the position of the cut-off relative to the mean will be critical in what follows.[23] This is because for a symmetric distribution, when σ increases in a mean preserving manner, $\mathcal{F}\left(\hat{h}; \sigma\right)$ rises (falls) for all $\hat{h}$ below (above) the single-crossing point (which, in this case, is the mean $\bar{h}$). In short, for a symmetric, unimodal distribution, $\mathcal{F}_2\left(\hat{h}, \sigma\right) \gtreqless 0$ as $\hat{h} \lesseqgtr \bar{h}$.

[21] Recall, that for a density function f symmetric about its mean $\bar{h}$, the skewness given by $\mu_3 \equiv E\left(h^i - \bar{h}\right)^3$ is zero, and $\mathcal{F}\left(\bar{h}\right) = 1/2$. If $\mu_3 > 0$, the density is often called skewed to the right, or having long tails to the right.

[22] The robustness of our results extends to skewed distributions; see Section 4 below.

[23] More generally, what is crucial is the position of the cut-off $\hat{h}$ relative to the single crossing point h^* (which, for asymmetric distributions, may not coincide with the mean).

Lemma 3 *Let $\mathcal{F}(;\sigma)$ be a symmetric and unimodal distribution. Define*

$$\lambda^{\dagger} \equiv \frac{1+\bar{h}-\frac{\phi}{w}\frac{\ln(1+\theta)}{\theta}}{1+\bar{h}-\frac{\phi}{w}\frac{1}{(1+\theta)}\left(1+\frac{\theta}{2}\right)}.$$

Then, a) $\hat{h} \gtreqless \bar{h}$ *iff* $\lambda \gtreqless \lambda^{\dagger}$, *and b)* $\underline{\lambda} < \lambda^{\dagger} < \overline{\lambda}$.

This means for any λ in the range $\left(\lambda^{\dagger}, \overline{\lambda}\right)$, $\hat{h} > \bar{h} \Leftrightarrow \mathcal{F}_2\left(\hat{h}, \sigma\right) < 0$ will obtain; similarly, for any λ in the range $\left(\underline{\lambda}, \lambda^{\dagger}\right)$, $\hat{h} < \bar{h} \Leftrightarrow \mathcal{F}_2\left(\hat{h}, \sigma\right) > 0$ will obtain.[24] We are now ready to state a very important result.

Proposition 1 *a) The sign of* $\frac{d\hat{h}}{d\sigma}$ *is opposite that of* $\mathcal{F}_2\left(\hat{h}, \sigma\right)$. *b) Suppose* $\lambda > \lambda^{\dagger}$. *Then, for a symmetric and unimodal distribution,* $\frac{d\hat{h}}{d\sigma} \geq 0$ *holds, implying a mean-preserving increase in the spread raises the cut-off,* $\hat{h}$.

To summarize: from Lemma 3, we know that for a symmetric and unimodal distribution of innate ability, the cut-off, $\hat{h}$, will exceed the mean, $\bar{h}$, for high-enough λ. Proposition 1 takes it further by claiming that for symmetric distributions, and for high-enough λ, a mean-preserving increase in the spread *raises* the cut-off $\hat{h}$.

As discussed above, for sufficiently high λ, $\hat{h} > \bar{h}$ holds and this implies $\mathcal{F}\left(\hat{h};\sigma\right)$ falls when σ increases. In other words, a mean-preserving increase in spread reduces the mass on those not in the rat race. Since the consumption of the latter is lower than that of those who are, average consumption rises, and hence the cut-off rises.

3.2 Effect on the induced distributions

We start by exploring the effect on average leisure in the economy. In a sense, this is our flagship result; hence we devote a fair amount of words to explain it and provide intuition. Many of the results that follow, work off similar intuition, and hence are not subsequently repeated.

[24] By inspection, $\lambda^{\dagger} > 1$. This is an artifact of our simplifying assumptions of symmetric distributions and quasi-linear utility. We have verified that for a more general CRRA form utility, $\lambda^{\dagger} \leq 1$ is possible. In any case, it is not unreasonable to imagine that people who *are* in the rat race aspire to reach consumption levels that are "above average".

Proposition 2 *a)* $\frac{\partial \bar{l}}{\partial \sigma}$ *has the same sign as* $\mathcal{F}_2\left(\hat{h}, \sigma\right)$. *b) Suppose* $\lambda > \lambda^{\dagger}$. *Then, for a symmetric and unimodal distribution,* $\frac{\partial \bar{l}}{\partial \sigma} \leq 0$.

In other words, for sufficiently high λ, a mean preserving increase in the spread of the innate-ability distribution reduces average leisure. Loosely speaking, Proposition 2 says that when people desire to stay ahead of a benchmark that exceeds average consumption, an increase in the underlying inequality pushes them to work harder on average.

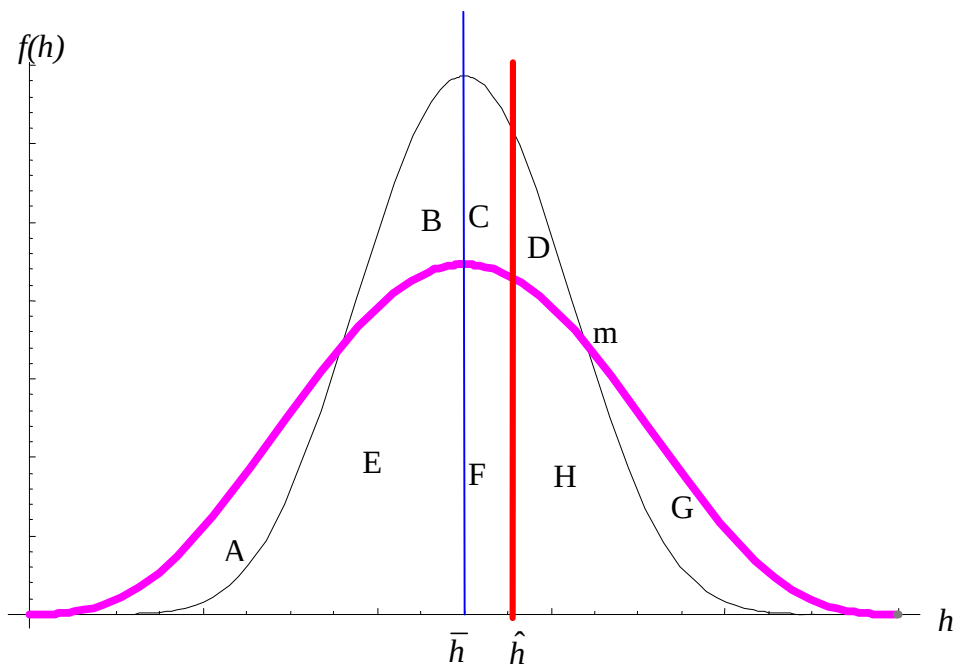

Figure 3: A mean-preserving increase in spread

To get a sense of the intuition for Proposition 2, focus attention on Figure 3. In this figure, the heavier purple curve represents a mean-preserving increase in spread of the original distribution (shown as the thin black line). For ease of presentation, we start by considering a fixed $\hat{h}$ (the bold red vertical line drawn here to the left of the point m; arguments similar to the ones presented below hold when $\hat{h}$ is to the right of m). Under the original distribution, the mass of people dropping out is given by the area $B + C + E + F$;

under the new distribution, the mass becomes $A + E + F$. Since both distributions are symmetric, $B + E = A + E = 0.5$ implying that the mass of agents below $\bar{h}$ – all drop-outs – is the same. Above the mean, the mass of drop-outs in the original distribution is $C + F$ while it is F under the new distribution; this implies that the total mass of drop-outs is now *reduced* by the area C.[25] Since those that participate in the rat race consume less leisure than those who drop out, the net effect is that average leisure goes down.[26]

More generally though, we know from Lemma 3 that a mean-preserving increase in the spread raises the cutoff $\hat{h}$ when $\hat{h}$ lies above $\bar{h}$ (i.e., it moves the red line to the right). Why? Above, we established that the net effect of an increase in spread is that average leisure goes down for a given $\hat{h}$; this means average consumption rises and this causes the cutoff $\hat{h}$ to rise as well. The increase in $\hat{h}$ further serves to increase average leisure; evidently, this last effect is of second-order magnitude.

Next, we study the effect of a MPIS of the innate-ability distribution on the second moment of the leisure distribution.

Proposition 3 *Suppose $\lambda > \lambda^{\dagger}$. Then, for a symmetric and unimodal distribution,*

$$\frac{\partial V(l)}{\partial \sigma} > 0.$$

In other words, for sufficiently high λ, a mean-preserving increase in the spread of the innate-ability distribution raises the variance of the leisure distribution.

How does a change in the innate-ability distribution affect the first and second moments of the income distribution?

[25] Under the old distribution, the mass of participants is $D + H$ and under the new, it is $H + G$. Since $C + D + F + H = F + H + G = 0.5$, it follows that $C + D = G$, i.e., the net increase in the mass of participants is $C = G - D$.

[26] Using panel data from 10 OECD countries, Bowles and Park (2005) find that income inequality is a significant predictor of average hours worked, and this relationship is robust across a variety of model specifications. Bowles and Park (2005) also study a simple model of the choice of work hours. Theirs is a formalization of Veblen's persuasive argument that everyone in society tries to emulate the richest – the leisure class – who in Veblen's time did not do much work. Their model predicts that work hours are increasing in the degree of income inequality.

While there are several superficial similarities, the differences between our papers stand out prominently. In Bowles and Park (2005), the need to emulate the leisure class is modeled as a consumption externality in the utility function of the poor; the rich are in no way impacted by the actions of the poor. Any stratification of the population along the lines of consumption and leisure is purely exogenous to their model, since no one *chooses* whether or not to participate in the rat race.

Proposition 4 *a)* $\frac{\partial \bar{y}}{\partial \sigma}$ *has the opposite sign of* $\mathcal{F}_2\left(\hat{h}, \sigma\right)$. *b) Suppose* $\lambda > \lambda^{\dagger}$. *Then, for a symmetric and unimodal distribution,* $\frac{\partial \bar{y}}{\partial \sigma} > 0$ *and* $\frac{\partial V(y)}{\partial \sigma} > 0$.

In other words, for sufficiently high λ, a mean-preserving increase in the spread of the innate-ability distribution raises the average income in the economy. It also increases the variance of income.

4 Consumption lotteries

The analysis thus far has ignored the possibility that agents may have access to consumption lotteries – a central theme in Sadler (2000) and Becker, Murphy, and Werning (2005). We start with some general discussion regarding consumption lotteries and their potential to make some people, with our kind of keeping-up preferences, better off.

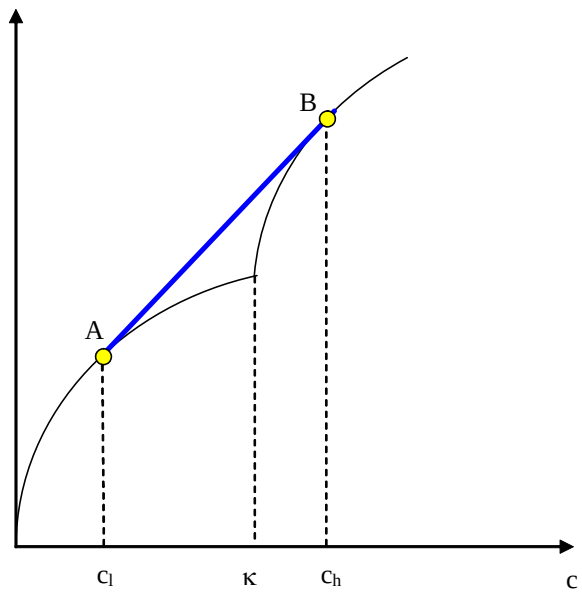

Figure 4: Consumption lotteries

To that end, focus attention on Figure 4 which plots the utility from consumption. A person with income y_i faces a kinked utility function (the kink is at κ) in the consumption

space. The kink introduces a non-concave segment in the utility function. Suppose a lottery was available that offered a top prize of $c_h > \kappa$ and a low prize of $c_l < \kappa$; also suppose that if c_l (c_h) was put into the lottery, the chance of winning the top prize was exactly 0 (1). In an ex-ante sense, an agent who can invest the income-equivalent of an amount in the range (c_l, c_h) would be better off participating in the lottery. People with income-equivalents in the range (c_l, κ) would benefit from having a shot at joining the rat race, while those with income-equivalents in the range (κ, c_h) would benefit from being able to reach c_h.

We now derive the terms of such lotteries.[27] Consider a setting in which such lotteries are freely accessible by all agents, though, in equilibrium, not all will choose to participate. We assume the lotteries are fair in the sense that the expected payout to a lottery participant equals the amount invested in it.

Consider agent i who invests her income y_i in such a lottery.[28] The lottery pays c_h if agent i wins the lottery; otherwise it pays $c_l < y_i$. Let α_i be the probability of agent i winning the lottery. Then, since the lottery is fair to agent i,

$$\alpha_i c_h + (1 - \alpha_i) c_l = y_i \Rightarrow \alpha_i = \frac{y_i - c_l}{c_h - c_l}.$$

Since $\alpha_i \in [0, 1]$, it follows that agent i participates in such a lottery if $c_h \geq y_i \geq c_l$. In other words, the subset of agents that potentially participate in such a lottery are those with income-equivalents $y_i \in [c_l, c_h]$. In passing, notice that the odds of winning increase with the amount invested. Those whose incomes are below c_l don't participate because they prefer the sure thing (their income) to any fair lottery offered to them. The same is true of people with incomes above c_h.

How are c_l and c_h selected? We posit a perfectly competitive lottery industry where individual lottery firms are Nash competitors. In other words, c_l and c_h are chosen by maximizing the expected utility of any lottery participant. The lottery firm is aware that its choices of c_l and c_h would affect the labor supply decision of the participants.

[27] See Hartley and Farrell (2002) and Sadler (2000) for detailed discussion.

[28] Of course, y_i is a function of the agent's choice of leisure, i.e., $y_i(l_i) = w(1 + h^i - l_i)$; for now, notationally we suppress the dependence of y_i on l_i.

Proposition 5 *The lottery prizes, c_l and c_h, satisfy*

$$u'(c_l) = (1+\theta)\, u'((1+\theta)\, c_h - \theta\kappa)\,. \tag{14}$$

An agent with income $y_i \in [c_l, c_h]$ faces the odds

$$\alpha_i = \frac{y_i - c_l}{c_h - c_l}$$

of winning the top prize, c_h.

Notice that the prizes, c_h and c_l, are independent of i. In passing, notice that (14) represents the line segment AB in Figure 4, a tangency condition if you will, with the tangent touching points A and B.

In our framework, as is readily apparent, an agent can influence her odds of winning the lottery, α, by working harder. The question then arises, conditional on an agent participating in the lottery, what is the optimal amount of leisure she enjoys? If she participates, her expected utility as a function of her choice of leisure is given by

$$U_i(l) = \alpha_i(l)\, u(c_h) + (1 - \alpha_i(l))\, u(c_l) + \phi v(l)$$

where $\alpha_i(l) = \frac{y_i(l) - c_l}{c_h - c_l}$. Then, the first order condition for an interior solution is given by

$$\alpha_i'(l)\,[u(c_h) - u(c_l)] + \phi v'(l) = 0. \tag{15}$$

Since $y_i(l) = w\left(1 + h^i - l\right)$, it follows that $\alpha_i'(l)$ is independent of l and i. This implies the solution of (15) is independent of i. In short, all lottery participants enjoy the same leisure; those with high h get higher incomes (and hence, better odds) than those with lower h.

From here on, we focus attention on a special functional form for preferences:

$$U(c,l) \equiv \begin{cases} \ln(c + \theta(c - \kappa)) + \phi \ln l & \text{if } c > \kappa \\ \ln c + \phi \ln l & \text{if } c \leq \kappa \end{cases}. \tag{16}$$

For the preferences described by (16), it is easy to check that (14) reduces to

$$c_h - c_l = \frac{\theta}{1+\theta}\kappa. \tag{17}$$

Using (26)-(27) in the appendix, and (17), it can be shown that

$$c_l = \frac{\theta\kappa}{1+\theta}\frac{1}{\ln(1+\theta)};\ c_h = \frac{\theta}{1+\theta}\kappa\left[1+\frac{1}{\ln(1+\theta)}\right].$$

Using (15), it also follows that the leisure of all lottery participants is given by

$$l_{lot} = \frac{\phi\frac{\theta}{1+\theta}\frac{\kappa}{w}}{\ln(1+\theta)}.$$

Then, using (18)-(19) in Appendix A, it can be shown that

drop-outs	$l_d^i = \frac{\phi}{(\phi+1)}(1+h_i)$	$\left(1+h^i\right) < \left(\frac{\theta}{1+\theta}\frac{\kappa}{w}\right)\frac{1+\phi}{\ln(1+\theta)} \equiv \left(1+h^\dagger\right)$
lottery players	$l_{lot} = \frac{\theta}{1+\theta}\frac{\phi}{\ln(1+\theta)}\frac{\kappa}{w}$	$\left(\frac{\theta}{1+\theta}\frac{\kappa}{w}\right)\frac{1+\phi}{\ln(1+\theta)} < \left(1+h^i\right) < \left(\frac{\theta}{1+\theta}\frac{\kappa}{w}\right)\left(1+\frac{(1+\phi)}{\ln(1+\theta)}\right) \equiv (1+h^*)$
participants	$l_r^i = l_d^i - \frac{\theta}{(\theta+1)}\frac{\phi}{(\phi+1)}\frac{\kappa}{w}$	$\left(1+h^i\right) > \frac{\theta}{1+\theta}\frac{\kappa}{w}\left[1+\frac{(1+\phi)}{\ln(1+\theta)}\right]$

The log-log formulation is a departure from the preferences used earlier in the paper. It is discussed in greater detail in Appendix A. The reason for the change is simple: for the preferences described in Section 2.2, *everyone* participates in the lottery, and hence, no one *can* make the ex-ante choice to participate in the rat race or drop out; their ex post status is determined *entirely* by the lottery outcome. The preferences outlined in (16) are less limiting – there emerges an ex-ante endogenous separation of the population into three groups, the drop-outs, the lottery participants, and the rat race participants. Of course, some of the lottery participants would be labeled drop-outs in a world sans lotteries.

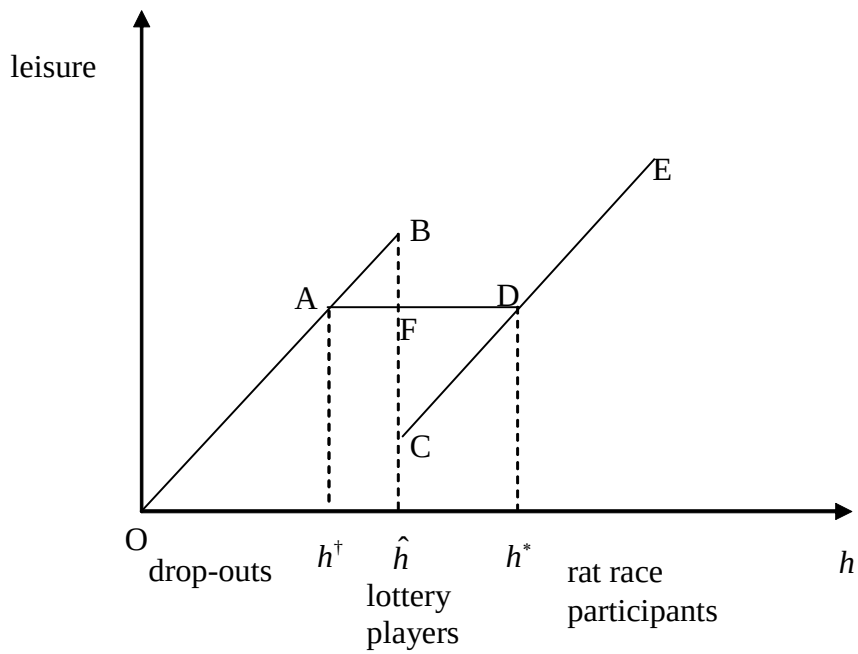

Figure 5: Leisure in the presence of lotteries

We focus on explaining how our flagship result regarding average leisure is affected with the introduction of lotteries. Consider Figure 5 which is drawn under *the assumption of a fixed* κ. (Note: $h^\dagger$ and h^* are defined in the box above, and $\hat{h}$, the ability corresponding to the agent who is indifferent between dropping out of the rat race and participating in a world with no lotteries, is defined in (20) below). In the absence of lotteries, agents with ability $h < \hat{h}$ are drop-outs and their leisure is described by the line segment, OB [a line with slope $\phi/(1+\phi)$]; similarly, those with ability $h > \hat{h}$ are participants and their leisure is described by the line segment CE.

When lotteries are available, the lottery players' leisure is described by the line segment AD. Agents with ability in the range $\left(h^\dagger, h^*\right)$, the erstwhile drop-outs, now participate in the lottery; their leisure is described by AF and no longer by AB. Similarly, agents with ability in the range $\left(\hat{h}, h^*\right)$, the erstwhile participants, now play the lottery; their leisure is described by the segment FD and no longer by CD. It follows that if the distribution of h puts sufficient weight on agents with low h (those who enjoy reduced leisure), the overall

effect of introducing lotteries is a likely reduction in average leisure.

It is not possible to make much analytical progress from here on. Below, we present an example illustrating the effects of MPIS in h on the first and second moments of the induced distributions. It is assumed that h follows a skewed Beta distribution; the figures below support our claim that the theoretical results derived earlier for symmetric distributions generalize to skewed distributions.

Example 1 *Suppose* $w = 1, \phi = 1, \theta = 0.5,\ \lambda = 1.5,$ *and let* h^i *be drawn from a Beta distribution with fixed mean 0.4 and variance* σ*. Figure 6a and 6b plots average leisure and income against* $\sigma^2 \in [0.05, 0.5]$*. Figure 6c and 6d plot the variance of leisure and income against fundamental inequality –* $\sigma^2 \in [0.05, 0.5]$*. The black (red) line represents variables from the model economy in which lotteries are absent (present).*

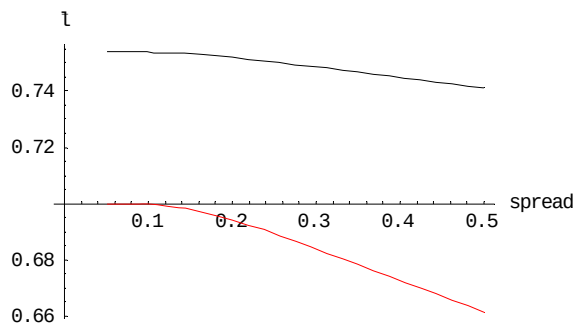

Figure 6a: Average leisure against σ^2

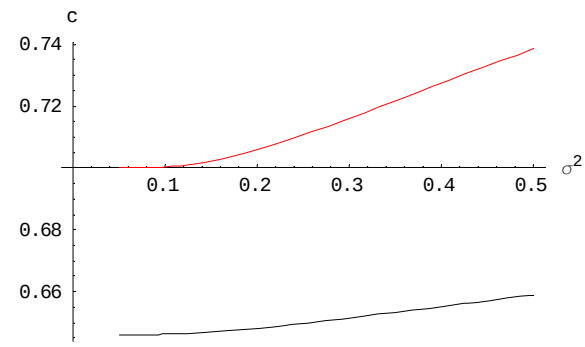

Figure 6b: Average income against σ^2

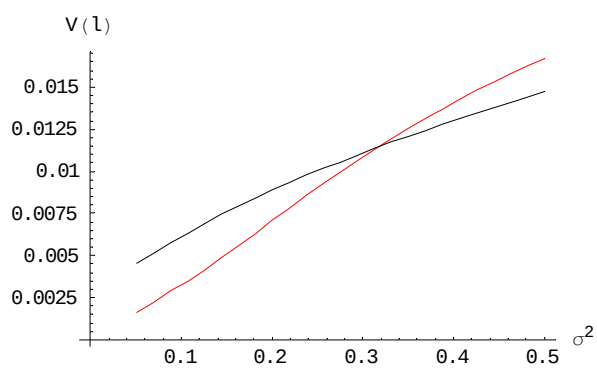

Figure 6c: Variance of leisure against σ^2

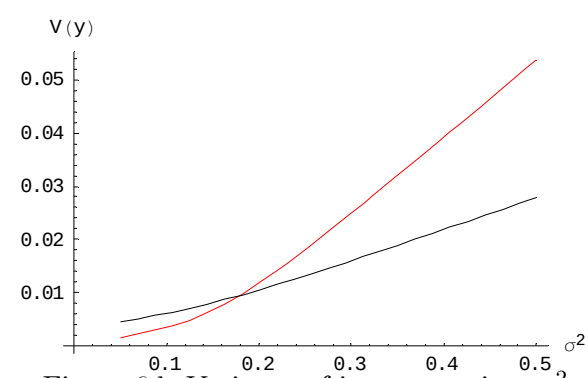

Figure 6d: Variance of income against σ^2

A few words about Figure 6 are in order. First, for the parametric specification in

Example 1, average leisure falls while average income rises with a mean-preserving increase in the spread. Second, the variance of both income and leisure rise with a MPIS of h. These pictures indicate that the main thrust of our previously-derived analytical results (from the model sans lotteries) continue to hold. It deserves mention that income inequality in the presence of lotteries (see Fig 6d) is initially lower but as fundamental inequality rises, it catches up and then exceeds income inequality from the model sans lotteries. The same pattern emerges in leisure inequality. Somewhat predictably, the presence of lotteries reduces the pass-through from fundamental to income/leisure inequality considerably.

5 Concluding remarks

From biblical times, human beings have struggled with an innate desire to covet their neighbors' possessions. For many, this translates into a seemingly endless pursuit of keeping up to the consumption standards set by their neighbors. Others simply drop out of this rat race. In this paper, we take the stance that people who successfully keep up with their neighbors receive an extra boost in utility relative to those who drop out. We embed this preference formulation in a simple framework in which heterogenous-ability agents get utility not just from consumption but also from the very act of keeping up with a benchmark proportional to the average level of consumption. We ask, what do such preferences imply for the distribution of income (and leisure) in an economy? We find that only the high-ability agents choose to keep up and they enjoy higher consumption but lower leisure than those who stay out of the rat race. The rat race amplifies the fundamental inequality in the economy by making incomes far more diverse than underlying innate abilities. When ability-inequality rises in a mean-preserving manner, average leisure and income in the economy may fall and leisure and income may get more dispersed. Such connections between inequality, work, and consumption do not emerge from a standard keeping-up-with-the-Joneses framework.

Embedded in a growth model, our framework has the potential to address themes connecting current second moments with future first moments of the income distribution.

These themes recur in the works of Benabou (1996) and Galor and Zeira (1993) among numerous others. It is important to note that the source of inequality in these models is typically a non-convex investment technology, preventing convergence of incomes. Our formulation suggests that non-convexity in preferences may also be important.

The analysis also sheds some light on the transatlantic leisure-differential issue.[29] Prescott (2004), for example, has argued that higher marginal income tax rates in Europe relative to the U.S. are closely correlated with their shorter workweeks and expanding vacations. While no doubt tax rates have strong effects on the first moment of the leisure distribution, our analysis suggests that fundamental inequality may have additional effects on average leisure above and beyond those implied by average tax rates. Our analysis also suggests that the entire distribution of tax rates (aka the tax structure), not just the average tax rate, that affects the various moments of the leisure distribution.

Our preference formulation has other potential applications, such as in the area of asset pricing. Similarly, the model could be extended to allow different income classes to keep up with consumption levels closer to their own. A dynamic version of this economy that includes a consumption-saving choice may shed light on the connection between income inequality and aggregate savings. Our formulation may also prove to be important in interpreting the recent evidence on leisure inequality in the U.S. (see Kuhn and Lozano, 2008). We leave these, and other issues, for future research.

[29] Prescott (2004) and others ask: in recent times, why are Americans working so much harder than the Europeans? Suppose one loosely proxies the income distribution by the innate-ability distribution. In that case, one could argue that all else being similar across the two shores of the Atlantic, the higher income inequality in the United States may be inducing more Americans to join the rat race and work harder on average than their European counterparts.

References

[1] Abel, A."Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review*, 80(2), 38-42 (1990)

[2] Banerjee, Abhijit V & Newman, Andrew F, "Risk-Bearing and the Theory of Income Distribution," *Review of Economic Studies*, 58(2), 211-35 (1991)

[3] Becker, Gary S. , Kevin M. Murphy, and Iván Werning. "The Equilibrium Distribution of Income and the Market for Status", *Journal of Political Economy,* 113(2), 282-310, (2005)

[4] Benabou, R "Inequality and Growth", NBER Macroeconomics Annual, MIT Press, Cambridge, MA , 11–74 (1996)

[5] Bowles, S. and Park, Y., "Emulation, Inequality and Work Hours: Was Thorstein Veblen Right?" *Economic Journal*, 115, Issue 507, 397-412 (2005)

[6] Clark, Andrew E., Andrew J. Oswald "Comparison-concave utility and following behaviour in social and economic settings", *Journal of Public Economics*, 70(1), 133-15. (1998)

[7] Corneo, G., Jeanne, O.: "Pecuniary emulation, inequality and growth", *European Economic Review*, 43, 1665–1678 (1999)

[8] Creedy, John *The Dynamics of Inequality and Poverty*, Edward Elgar, MA (1998)

[9] Dupor, Bill, and Wen-Fang Liu, "Jealousy and Equilibrium Overconsumption," *American Economic Review*, XCIII, 423-428 (2003)

[10] Dynan, Karen E., and Enrichetta Ravina. "Increasing Income Inequality, External Habits, and Self-Reported Happiness." *American Economic Review*, 97(2), 226–231.(2007)

[11] Freeman, S. "Equilibrium Income Inequality among Identical Agents," *Journal of Political Economy*, 104, 1047-1064 (1996)

[12] Friedman, M., and Savage, L. "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy* 56, 279-304 (1948)

[13] Galí, J. "Keeping up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices", *Journal of Money, Credit and Banking*, 26 (1), 1-8 (1994)

[14] Galor Oded, "Convergence? Inferences from Theoretical Models," *Economic Journal*, 106, 1056-1069. (1996)

[15] – and Zeira, J. "Income Distribution and Macroeconomics," *Review of Economic Studies* 60, 35-52 (1993)

[16] García-Peñalosa C. and Turnovsky S.J. "Growth and income inequality: a canonical model", *Economic Theory* 28, 25–49 (2006)

[17] – "Consumption externalities: a representative consumer model when agents are heterogeneous", *Economic Theory*, 37, 439-467 (2008)

[18] Gordon, Cyrus H. "A Note on the Tenth Commandment", *Journal of Bible and Religion*, 31(3), 208-209 (1963)

[19] Hartley, Roger, and Farrell, Lisa. "Can Expected Utility Theory Explain Gambling?," *American Economic Review*, 92 613-624 (2002)

[20] Hopkins, E. "Inequality, happiness and relative concerns: What actually is their relationship?" *Journal of Economic Inequality*, 6, 351–372 (2008)

[21] Kawamoto, Koichi "Status-seeking behavior, the evolution of income inequality, and growth", *Economic Theory*, 39, 269–289 (2009)

[22] Kuhn, P., and Lozano, F. "The expanding workweek? Understanding trends in long work hours among U.S. men, 1979-2006", *Journal of Labor Economics*, 26(2), 311-343 (2008)

[23] Landers, R., Rebitzer, J., Taylor. L., "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms." *American Economic Review*, 86(3), 329-48 (1996)

[24] Lewis, G.W.. and Ulph, D.T. "Poverty, inequality, and welfare", *Economic Journal*, 98, 117-131 (1988)

[25] Long, N.V., Shimomura, K. "Relative wealth, status-seeking, and catching-up", *Journal of Economic Behavior and Organization*, 53, 529–542 (2004)

[26] Ljungqvist, Lars and Harald Uhlig, "Tax Policy and Aggregate Demand Management Under Catching Up with the Joneses," *American Economic Review* 90, 356–66 (2000)

[27] Luttmer, E., "Neighbors as Negatives: Relative Earnings and Well-Being", *Quarterly Journal of Economics*, 20(3), 963-1002 (2005)

[28] Prescott, E.C. "Why Do Americans Work So Much More Than Europeans?" *Quarterly Review*, Federal Reserve Bank of Minneapolis, 28(1), 2-13 (2004).

[29] Sadler, Michael A. "Escaping Poverty: Risk-Taking and Endogenous Inequality in a Model of Equilibrium Growth", *Review of Economic Dynamics* 3, 704-725 (2000).

Appendix

A Additive log preferences

We briefly sketch the model in the case where preferences take the log-log form (used in Section 4) above:

$$U(c,l) \equiv \begin{cases} \ln(c+\theta(c-\kappa)) + \phi \ln l & \text{if } c > \kappa \\ \ln c + \phi \ln l & \text{if } c \leq \kappa \end{cases}.$$

In this case, it is easy to verify that (3)-(4) is replaced by

$$l_d^i = \frac{\phi\left(1+h^i\right)}{1+\phi}; \quad l_r^i = l_d^i - \frac{\phi\theta\kappa}{w(\phi+1)(\theta+1)} \tag{18}$$

$$c_d^i = w\frac{\left(1+h^i\right)}{1+\phi}; \quad c_r^i = c_d^i + \frac{\phi\theta\kappa}{(\phi+1)(\theta+1)}. \tag{19}$$

Notice that both c_r^i and l_r^i now depend on κ. Also, the gap between l_r^i and l_d^i (parenthetically, the gap between c_r^i and c_d^i) depends on κ : as κ rises, these gaps rise. Also note that it is no longer the case that $l_r^i < l_d^j$ for any i, j but $c_r^i > c_d^j$ obviously continues to hold. An agent will choose to stay out of the rat race if the utility from participating exceeds that of staying out. Formally, those h^i for whom the following inequality hold will participate:

$$\ln\left(c_r^i + \theta\left(c_r^i - \kappa\right)\right) + \phi \ln l_r^i > \ln c_d^i + \phi \ln l_d^i.$$

Straightforward algebra shows that when $\kappa = \lambda\bar{c}$, the previous inequality reduces to

$$h^i > \hat{h} \equiv \frac{\theta}{(\theta+1) - (\theta+1)^{\frac{\phi}{1+\phi}}}\left(\frac{\bar{c}}{\lambda w}\right) - 1, \tag{20}$$

where $(\theta+1) - (\theta+1)^{\frac{\phi}{1+\phi}} > 0$ and $\theta/[(\theta+1) - (\theta+1)^{\frac{\phi}{1+\phi}}] > 1$.[30] Using the expressions for c_d^i and c_r^i derived earlier, it is straightforward to check that

$$\bar{c} = \frac{w(\theta+1)\left(\bar{h}+1\right)}{1+\phi+\theta+\theta\lambda\phi\mathcal{F}\left(\hat{h};\sigma\right)}. \tag{21}$$

[30] Note that

$$l_r^i > 0 \Leftrightarrow 1 + h_i - \frac{\theta\lambda\bar{c}}{w(\theta+1)} > 0.$$

It is straightforward to verify that the previous inequality holds for all $h_i \geq \hat{h}$. This means every rat race participant enjoys a strictly interior level of leisure.

Using (20) and (21), we derive the analog of eq. (9):

$$\frac{w\left(\theta+1\right)\left(\bar{h}+1\right)}{1+\phi+\theta+\theta\lambda\phi\mathcal{F}\left(\hat{h};\sigma\right)}=w\left(\frac{\left(\theta+1\right)-\left(\theta+1\right)^{\frac{\phi}{1+\phi}}}{\theta\lambda}\right)\left(1+\hat{h}\right).$$

B Proof of Lemma 1

Since

$$1+\hat{h}=\frac{\kappa}{w}+\frac{\phi\ln\left(1+\theta\right)}{\theta w}=\frac{\kappa}{w}+\frac{\phi}{w}\frac{\ln\left(1+\theta\right)}{\theta}$$

and drop-outs are those with $h^i<\hat{h}$, we have

$$1+h^i<1+\hat{h}=\frac{\kappa}{w}+\frac{\phi}{w}\frac{\ln\left(1+\theta\right)}{\theta}<\frac{\kappa+\phi}{w}.$$

Since

$$c_d^i=w\left(1+h^i-\frac{\phi}{w}\right)<\kappa\Longleftrightarrow 1+h^i<\frac{\kappa+\phi}{w}$$

we get

$$h^i<\hat{h}\Rightarrow c_d^i<\kappa$$

Similarly,

$$c_r^i=w\left(1+h^i-\frac{\phi}{w\left(1+\theta\right)}\right)>\kappa\Longleftrightarrow 1+h^i>\frac{\kappa}{w}+\frac{\phi}{w\left(1+\theta\right)}$$

and since, rat race participants are those with $h^i>\hat{h}$, we have

$$1+h^i>1+\hat{h}=\frac{\kappa}{w}+\frac{\phi}{w}\frac{\ln\left(1+\theta\right)}{\theta}=\frac{\kappa}{w}+\frac{\phi}{w\left(1+\theta\right)}\frac{\left(1+\theta\right)\ln\left(1+\theta\right)}{\theta}>\frac{\kappa}{w}+\frac{\phi}{w\left(1+\theta\right)}$$

because $\frac{(1+\theta)\ln(1+\theta)}{\theta}>1$.

C Proof of Lemma 2

First note that $\underline{\lambda}<1<\overline{\lambda}$ such that the relevant interval for λ is non-empty. Let the left hand side of (9) be denoted by

$$l\left(h\right)\equiv\left(h-\bar{h}\right)+\frac{1}{w}\left(\frac{\phi\lambda\theta}{1+\theta}\right)\mathcal{F}\left(h;\sigma\right) \tag{22}$$

and the right hand side by

$$g \equiv \frac{1}{w}\left(\frac{\phi \ln(1+\theta)}{\theta} - \frac{\lambda\phi}{(1+\theta)}\right) + (\lambda - 1)\left(1+\bar{h}\right). \tag{23}$$

Further, it follows from (22) that $l'(h) > 0$, implying there is a unique equilibrium if $l(h_l) < g$ and $l(h_u) > g$. First consider $l(h_l) < g$. Since $\mathcal{F}(h_l) = 0$,

$$l(h_l) < g \Leftrightarrow h_l + 1 - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta} < \left(1 + \bar{h} - \frac{\phi}{w}\frac{1}{(1+\theta)}\right)\lambda \tag{24}$$

By (A.1), $\phi < \phi_u \equiv w(1+h_l)$. Then, it follows that

$$h_l + 1 - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta} \geq h_l + 1 - \frac{\phi_u}{w}\frac{\ln(1+\theta)}{\theta} = (h_l+1)\left(1 - \frac{\ln(1+\theta)}{\theta}\right) > 0$$

and

$$1 + \bar{h} - \frac{\phi}{w}\frac{1}{(1+\theta)} \geq 1 + \bar{h} - \frac{\phi_u}{w}\frac{1}{1+\theta} = (1+h_l)\left(\frac{1+\bar{h}}{1+h_l} - \frac{1}{1+\theta}\right) > 0$$

implying both sides of the above inequality in (24) are positive. Hence, it follows that

$$l(h_l) < g \Leftrightarrow \underline{\lambda} \leq \lambda.$$

Now consider $l(h_u) > g$. Since $\mathcal{F}(h_u) = 1$,

$$g < l(h_u) \Leftrightarrow \lambda\left(1 + \bar{h} - \frac{\phi}{w}\right) < 1 + h_u - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta} \Leftrightarrow \lambda < \overline{\lambda}.$$

■

D Proof of Lemma 3

Under (A.2), the condition provided in Lemma 2, there exists a unique $\hat{h}$. Since $l(h)$ is increasing and initially smaller than g, we know that $\hat{h} > \bar{h}$ is equivalent to the knowledge that the crossing of l and g has not yet occurred at $\bar{h}$. If this is the case, $l\left(\bar{h}\right)$ must be smaller than g. The argument for $\hat{h} < \bar{h}$ is exactly equivalent. Therefore, it follows that $\hat{h} \gtreqless \bar{h}$ iff $l\left(\bar{h}\right) \gtreqless g$. Noting that for a symmetric distribution, $\mathcal{F}\left(\bar{h}\right) = 1/2$, we get:

$$l\left(\bar{h}\right) < g \Leftrightarrow \frac{1}{w}\left(\frac{\phi\lambda\theta}{1+\theta}\right)\mathcal{F}\left(\bar{h}\right) < \frac{1}{w}\left(\frac{\phi\ln(1+\theta)}{\theta} - \frac{\lambda\phi}{(1+\theta)}\right) + (\lambda-1)\left(1+\bar{h}\right)$$

Further simplification yields

$$1 + \bar{h} - \frac{1}{w}\frac{\phi\ln(1+\theta)}{\theta} < \lambda\left(1 + \bar{h} - \frac{\phi}{2w}\frac{2+\theta}{1+\theta}\right) \Leftrightarrow \lambda^{\dagger} \equiv \frac{1 + \bar{h} - \frac{\phi}{w}\frac{\ln(1+\theta)}{\theta}}{1 + \bar{h} - \frac{\phi}{w}\frac{1}{(1+\theta)}\left(1 + \frac{\theta}{2}\right)} < \lambda.$$

The remaining part of the lemma is trivial to verify.■

E Proof of Proposition 1

Differentiating (9) yields

$$\frac{\partial \hat{h}}{\partial \sigma} + \frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right)\left(f\left(\hat{h},\sigma\right)\frac{\partial \hat{h}}{\partial \sigma} + \mathcal{F}_2\left(\hat{h},\sigma\right)\right) = 0$$

which upon rearrangement produces

$$\frac{\partial \hat{h}}{\partial \sigma} = -\frac{\frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right)\mathcal{F}_2\left(\hat{h},\sigma\right)}{1+\frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right)f\left(\hat{h},\sigma\right)}. \tag{25}$$

Hence, the sign of $\frac{\partial \hat{h}}{\partial \sigma}$ is the opposite of the sign of $\mathcal{F}_2\left(\hat{h},\sigma\right)$. For symmetric distributions, when $\lambda > \lambda^{\dagger}$, $\hat{h}$ lies to the right of the single crossing point $\left(\bar{h}\right)$ and so $\mathcal{F}_2\left(\hat{h},\sigma\right) < 0$; hence $\frac{\partial \hat{h}}{\partial \sigma} < 0$.■

F Proof of Proposition 2

Straightforward differentiation of (11) yields

$$\frac{\partial \bar{l}}{\partial \sigma} = \frac{\phi\theta}{(1+\theta)\,w}\left(f\left(\hat{h},\sigma\right)\frac{\partial \hat{h}}{\partial \sigma} + \mathcal{F}_2\left(\hat{h},\sigma\right)\right)$$

Substituting the expression for $\frac{\partial \hat{h}}{\partial \sigma}$ from (25) and rearranging, we get

$$\frac{\partial \bar{l}}{\partial \sigma} = \frac{\phi\theta}{(1+\theta)\,w}\left(\frac{\mathcal{F}_2\left(\hat{h},\sigma\right)}{1+\frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right)f\left(\hat{h},\sigma\right)}\right)$$

which has the same sign as $\mathcal{F}_2\left(\hat{h},\sigma\right)$.

G Proof of Proposition 3

Straightforward differentiation of $V(l)$ yields

$$\frac{\partial V(l)}{\partial \sigma} = \left(\frac{\phi}{w}\right)^2\left(\frac{\theta}{(1+\theta)}\right)^2\left(\left(\mathcal{F}_2\left(\hat{h}\right) + \mathcal{F}\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma}\right) - 2\mathcal{F}\left(\hat{h}\right)\left(\mathcal{F}_2\left(\hat{h}\right) + \mathcal{F}\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma}\right)\right)$$

$$= \left(\frac{\phi}{w}\right)^2\left(\frac{\theta}{(1+\theta)}\right)^2\left(1 - 2\mathcal{F}\left(\hat{h}\right)\right)\left(\mathcal{F}_2\left(\hat{h}\right) + \mathcal{F}\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma}\right)$$

Recall that the sign of $\frac{\partial \hat{h}}{\partial \sigma}$ is the opposite of the sign of $\mathcal{F}_2\left(\hat{h},\sigma\right)$. For symmetric, unimodal distributions, when $\lambda > \lambda^{\dagger}$, $\hat{h}$ lies to the right of the single crossing point $\left(\bar{h}\right)$ and so $\mathcal{F}_2\left(\hat{h},\sigma\right) < 0$; hence $\frac{\partial \hat{h}}{\partial \sigma} < 0$. In that case $1 - 2\mathcal{F}\left(\hat{h}\right) < 0$ as well, and $\frac{\partial V(l)}{\partial \sigma} > 0$.

H Proof of Proposition 4

Straightforward differentiation of $\bar{y}$ yields

$$\begin{aligned}
\frac{\partial \bar{y}}{\partial \sigma} &= -\phi f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} - \phi \mathcal{F}_2\left(\hat{h}\right) + \frac{\phi}{(1+\theta)} f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} + \frac{\phi}{(1+\theta)}\mathcal{F}_2\left(\hat{h}\right) \\
&= -\frac{\theta\phi}{(1+\theta)}\left[f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} + \mathcal{F}_2\left(\hat{h}\right)\right]
\end{aligned}$$

using the expression for $\frac{\partial \hat{h}}{\partial \sigma}$ derived earlier, this reduces to

$$\frac{\partial \bar{y}}{\partial \sigma} = -\frac{\theta\phi}{(1+\theta)}\mathcal{F}_2\left(\hat{h},\sigma\right)\frac{1}{1+\frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right) f\left(\hat{h},\sigma\right)}$$

implying $\frac{\partial \bar{y}}{\partial \sigma}$ has the opposite sign of $\mathcal{F}_2$. The rest is as before.

Straightforward differentiation of $V(y)$ yields

$$\begin{aligned}
&\frac{\partial V(y)}{\partial \sigma} = 2\sigma w^2 + \phi^2\left(\frac{\theta}{(1+\theta)}\right)^2 \times \\
&\left(f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma}\left(1-\mathcal{F}\left(\hat{h}\right)\right) - \mathcal{F}\left(\hat{h}\right) f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} + \mathcal{F}_2\left(\hat{h},\sigma\right)\left(1-\mathcal{F}\left(\hat{h}\right)\right) - \mathcal{F}\left(\hat{h}\right)\mathcal{F}_2\left(\hat{h},\sigma\right)\right) \\
&= 2\sigma w^2 + \phi^2\left(\frac{\theta}{(1+\theta)}\right)^2\left[\left(1-2\mathcal{F}\left(\hat{h}\right)\right) f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} + \left(1-2\mathcal{F}\left(\hat{h}\right)\right)\mathcal{F}_2\left(\hat{h},\sigma\right)\right] \\
&= 2\sigma w^2 + \phi^2\left(\frac{\theta}{(1+\theta)}\right)^2\left(1-2\mathcal{F}\left(\hat{h}\right)\right)\left[f\left(\hat{h}\right)\frac{\partial \hat{h}}{\partial \sigma} + \mathcal{F}_2\left(\hat{h},\sigma\right)\right] \\
&= 2\sigma w^2 + \phi^2\left(\frac{\theta}{(1+\theta)}\right)^2\left(1-2\mathcal{F}\left(\hat{h}\right)\right)\mathcal{F}_2\left(\hat{h},\sigma\right)\frac{1}{1+\frac{1}{w}\left(\frac{\lambda\phi\theta}{1+\theta}\right) f\left(\hat{h},\sigma\right)}
\end{aligned}$$

For symmetric, unimodal distributions, when $\lambda > \lambda^{\dagger}$, $\hat{h}$ lies to the right of the single crossing point $\left(\bar{h}\right)$ and so $\mathcal{F}_2\left(\hat{h},\sigma\right) < 0$ and $\left(1-2\mathcal{F}\left(\hat{h}\right)\right) < 0$; hence $\frac{\partial V(y)}{\partial \sigma} > 0$.

I Proof of Proposition 5

We start by writing down the indirect expected utility to agent i from participation in the lottery:

$$U_i = \alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\ \left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right)\right)$$
$$+\left(1-\alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\right)u\left(c_l\right) + \phi v\left(l_i^*\left(c_l, c_h\right)\right)$$

where $l_i^*\left(c_l, c_h\right)$ is the optimal leisure of agent i. Then, assuming interior solutions,

$$\frac{\partial U_i}{\partial c_l} = \left(\alpha_1 \frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_l} + \alpha_2\right)\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right)\right) - \left(\alpha_1 \frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_l} + \alpha_2\right)u\left(c_l\right)$$
$$+\left(1-\alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\right)u'\left(c_l\right) + \phi v'\left(l_i^*\left(c_l, c_h\right)\right)\frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_l} = 0,$$

and

$$\frac{\partial U_i}{\partial c_h} = \left(\alpha_1 \frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_h} + \alpha_3\right)\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right)\right) + \alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\left(1+\theta\right)u'\left(\left(1+\theta\right)c_h - \theta\kappa\right)$$
$$-\left(\alpha_1 \frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_h} + \alpha_3\right)u\left(c_l\right) + \phi v'\left(l_i^*\left(c_l, c_h\right)\right)\frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_h} = 0,$$

where α_j is the partial derivative of α with respect to its jth argument. The above first order conditions may be rewritten as

$$\frac{\partial U_i}{\partial c_l} = \left[\alpha_1\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right) - u\left(c_l\right)\right) + \phi v'\left(l_i^*\left(c_l, c_h\right)\right)\right]\frac{\partial l^*\left(c_l, c_h\right)}{\partial c_l}$$
$$-\alpha_2 u\left(c_l\right) + \left(1-\alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\right)u'\left(c_l\right) + \alpha_2\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right)\right) = 0,$$

and

$$\frac{\partial U_i}{\partial c_h} = \left[\alpha_1\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right) - u\left(c_l\right)\right) + \phi v'\left(l_i^*\left(c_l, c_h\right)\right)\right]\frac{\partial l_i^*\left(c_l, c_h\right)}{\partial c_h}$$
$$+\alpha_3\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right) - u\left(c_l\right)\right) + \alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\left(1+\theta\right)u'\left(\left(1+\theta\right)c_h - \theta\kappa\right) = 0.$$

Using the envelope theorem, we get

$$\frac{\partial U_i}{\partial c_l} = \left(1-\alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\right)u'\left(c_l\right) + \alpha_2\left(\cdot\right)\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right) - u\left(c_l\right)\right) = 0,$$

and

$$\frac{\partial U_i}{\partial c_h} = \alpha_3\left(\cdot\right)\left(u\left(\left(1+\theta\right)c_h - \theta\kappa\right) - u\left(c_l\right)\right) + \alpha\left(l_i^*\left(c_l, c_h\right), c_l, c_h\right)\left(1+\theta\right)u'\left(\left(1+\theta\right)c_h - \theta\kappa\right) = 0.$$

Now, since

$$\alpha_i = \frac{y(l_i) - c_l}{c_h - c_l},\ \alpha_2 = -\frac{c_h - y(l_i)}{(c_h - c_l)^2},\ \alpha_3 = -\frac{y(l_i) - c_l}{(c_h - c_l)^2},$$

we get

$$\frac{\partial U_i}{\partial c_l} = 0 \Leftrightarrow (c_h - c_l)\, u'(c_l) = (u((1+\theta)\, c_h - \theta\kappa) - u(c_l)), \tag{26}$$

and

$$\frac{\partial U_i}{\partial c_h} = 0 \Leftrightarrow (c_h - c_l)(1+\theta)\, u'((1+\theta)\, c_h - \theta\kappa) = (u((1+\theta)\, c_h - \theta\kappa) - u(c_l)). \tag{27}$$

Since the r.h.s of (26) and (27) are the same, it follows that $\frac{u'(c_l)}{(1+\theta)u'((1+\theta)c_h - \theta\kappa)} = 1$ implying

$$u'(c_l) = (1+\theta)\, u'((1+\theta)\, c_h - \theta\kappa).$$