2008-WP-01
February 2008

The Subsistence Constraint and Endogenous Risk Aversion

Joel M. Guttman

Abstract: This paper contributes to the literature on endogenous preferences by showing that, when income is close to the minimum required for subsistence, individuals rationally will behave as if they were risk averse. It is suggested that observed risk aversion in empirical studies can be explained by the fact that, in the distant past, most people lived close to the subsistence constraint. It is shown that this approach can account for a number of empirical phenomena, including increasing income equality over time, decreasing risk aversion as income increases, and the economic role of insurance companies.

About the Author: Joel M. Guttman is a Senior Fellow at Networks Financial Institute. Guttman is currently serving as an associate professor in the department of economics in Bar-Ilan University. He has also served as a visiting professor at the University of Chicago, Brown University and the Helsinki School of Economics. His research areas include evolution of trust and cooperation, economic growth and conflict (internal and external), public economics, applied game theory, and defense economics. He has published numerous articles in scholarly journals including *American Economic Review*, *Journal of Political Economy*, *Journal of Public Economics*, and *Economic Journal*. He received his Ph.D. in economics from the University of Chicago.

Keywords: risk aversion, endogenous preferences, insurance.

JEL Classification: D01, D81, Z10.

The views expressed are those of the individual author and do not necessarily reflect official positions of Networks Financial Institute. Please address questions regarding content to Joel M. Guttman at guttman@mail.biu.ac.il. Any errors or omissions are the responsibility of the author.

NFI working papers and other publications are available on NFI's website (www.networksfinancialinstitute.org). Click "Research" and then "Publications/Papers."

The Subsistence Constraint and Endogenous Risk Aversion

Joel M. Guttman*

1 Introduction

In the theory of the consumer of modern economics, the consumer is assumed to maximize her utility subject to one constraint—that her expenditures cannot exceed her income, or in a multiperiod context, that the present value of her expenditures cannot exceed her wealth. There is another constraint which is routinely ignored: her consumption must be sufficient to ensure her survival. This “subsistence constraint,” while ignored by the modern theory of the consumer, featured prominently in classical nineteenth-century economics. Thomas Malthus (1798) argued that whenever income was higher than the minimum required for subsistence, population—and thus the supply of labor—would grow and per-capita income would then decline (due to an assumed fixed technology and fixed quantities of the other factors of production, land and capital). Eventually, per-capita income would fall below the subsistence level, and then a combination of war (i.e., violent competition for scarce means of sustenance) and famine would reduce population, raising per-capita income back to the subsistence level. If per-capita income were again to rise above the subsistence level, the cycle would repeat itself. This “Iron Law of Wages” had a prominent role in classical economics. Per-capita income was predicted to cycle eternally around the minimum required for subsistence.

Since the days of Malthus, per-capita income, at least in the Western world, has risen high above the subsistence level, making the subsistence constraint non-binding and thus explaining its neglect in modern economic theory. The present paper suggests that the subsistence constraint nevertheless continues to be relevant to economics, for two reasons. First, a large proportion of the world’s population still lives close to the subsistence level. Second, a relatively new field of research in economics—the theory of endogenous preferences—inquires into the roots of our preferences in the distant past of the human race, when indeed people typically lived close to the subsistence level. As we shall see, one of the most common assumptions

*Department of Economics, Bar-Ilan University, 52900 Ramat-Gan, Israel. E-mail: guttman@mail.biu.ac.il. The author thanks Sergiu Hart, Yuval Shilony, John A. Tatom, and Shlomo Yitzhaki for useful comments and discussions, while retaining responsibility for all remaining errors.

made in everyday economic theorizing—the assumption of risk-aversion—can be justified by recognizing the rationality of behaving in a risk-averse manner when one’s wealth is close to the minimum required for subsistence. Thus the assumption of risk-aversion can be shown to be no mere arbitrary assumption, equally compelling *a priori* as risk-neutral or risk-loving preferences. On the contrary, risk-aversion is a *consequence* of assuming that preferences evolved over time to maximize the survivability and growth of the human species.

Section 2 of this paper provides a quick review of expected utility theory. Section 3 argues, on the basis of some simple calculations and recent research by economic theorists, that risk aversion is the most plausible assumption that one can make about risk preferences, once the subsistence constraint is taken into account. Section 4 outlines some economic consequences of the conclusion of Section 3. Section 5 offers concluding observations.

2 A review of expected utility theory

Consider an individual whose initial wealth is W_0 and is offered a gamble which yields a wealth of $W_0 + x_1$ with a probability p and a wealth of $W_0 - x_2$ with a probability $1 - p$. Let us denote the wealth of the individual if he “wins” the bet as W_1 and his wealth if he “loses” as W_2 . That is, $W_1 \equiv W_0 + x_1$ and $W_2 \equiv W_0 - x_2$. Should the individual accept this gamble?

Suppose, as a first possible hypothesis, that the individual cares only about what would happen to his wealth if he were to repeat this gamble a great many times. The average *change* in the individual’s wealth, caused by accepting the gamble, is called the *expected value* of the gamble, which we define as

$$EV \equiv px_1 - (1 - p)x_2. \quad (1)$$

Thus, if the individual cares only about increasing his wealth over time, he should accept the gamble as long as EV is positive. If EV is zero, the individual would be just indifferent between accepting and rejecting the gamble. For example, suppose $x_1 = x_2$, i.e., the gain from winning equals the loss from losing. Then $EV = 0$ when $p = 0.5$. That is, the individual would be indifferent between accepting and rejecting an “even bet” of gaining and losing the same sum of money if the probability of winning (and losing) is one-half or “fifty-fifty,” since, on the average, he would neither gain nor lose wealth by accepting such a gamble. A gamble whose expected value is zero

is called “fair” since, on the average, one would neither gain nor lose by accepting such a gamble.¹

We observe, however, that many individuals systematically avoid accepting fair gambles. In fact, such individuals pay money to *avoid* risks of unfortunate events (that is, they buy insurance) even though the price of buying such insurance is high enough to allow the insurance company to make a profit, on the average, by accepting such risks. In other words, insurance companies take on gambles that are unfair *in their favor*: the expected value of such gambles to the insurance company is positive. Economists interpret such behavior by making an assumption about such people’s preferences, and by coupling this assumption to a plausible theory of how individuals rationally behave in the face of uncertainty. In particular, economists infer that people who buy insurance are *risk-averse*.

To understand the meaning of risk-aversion, consider Figure 1. This figure shows a possible utility function of wealth, $U(W)$. Since the function is concave, its slope—which is the marginal utility of wealth—decreases as wealth increases. As before, W_0 is the individual’s wealth before accepting the gamble, W_1 is his wealth if he accepts the gamble and wins, and W_2 is his wealth if he accepts the gamble and loses. Let us assume that the gamble is fair ($EV = 0$), so that, if the individual accepts the gamble, his wealth—on the average, if the gamble were repeated many times—would stay constant at W_0 . In technical terminology, we say that the individual’s *expected wealth* of accepting the gamble, which is calculated similarly to the expected value of the gamble as in (1), is

$$\begin{aligned} EW &\equiv pW_1 + (1 - p)W_2 \\ &= W_0 + px_1 + (1 - p)x_2 \\ &= W_0 + EV = W_0. \end{aligned}$$

Thus, if the individual was one who cared only about increasing his wealth, he would be just indifferent about taking this gamble.

Let us suppose, in contrast, that the individual wants to maximize his *expected utility*, not his expected wealth. The expected utility of the gamble can be defined as the individual’s average utility consequent to accepting the gamble, if the gamble were repeated many times. It is calculated similarly

¹The adjective “fair” can be understood by recalling that there are often two sides to a gamble. For example, John may bet against Jill that horse A will win a race. In the example of the even bet given above, neither John nor Jill will increase his or her wealth, on the average, if the probability of A winning the race is exactly 0.5. In this sense, the bet is fair.

Figure 1: Utility of wealth function, risk-averse individual

to the expected wealth of the gamble:

$$EU \equiv pU_1 + (1 - p)U_2,$$

where $U_1 \equiv U(W_1)$ is the individual's utility if he "wins" and $U_2 \equiv U(W_2)$ is his utility if he "loses," as shown in Figure 1. Graphically, we can calculate EU by drawing a straight line between points A and B , which are the points on the individual's utility function when he loses and wins the gamble, and locating the point C on line segment $\overline{AB}$ exactly above the expected wealth EW . The height of point C above the horizontal axis is EU , as indicated in the figure.

Notice that point C is lower than the point on the utility function corresponding to the individual's initial wealth, namely point D . The fact that C is lower than D is a direct consequence of the assumed concavity of the utility function, which is expressed by the fact that its slope, the marginal utility of wealth, decreases as wealth increases. This fact implies that the expected utility of the gamble, EU , is less than the individual's utility of his initial wealth, W_0 , namely $U(W_0)$. The concavity of the utility function in Figure 1 reflects the assumption that the individual is risk-averse. A risk-averse individual who maximizes his expected utility would clearly reject the gamble, since accepting this gamble would reduce his expected utility to a level below his (certain) utility before accepting the gamble. Unlike a risk-neutral individual, a risk-averse agent would *not* be indifferent between accepting and rejecting a fair gamble, but rather unambiguously would reject such a gamble.

The shape of the utility function shown in Figure 1 is only one theoretical possibility. If the agent's utility function were linear, as in Figure 2, the agent would be said to be *risk-neutral*. Such an individual would be indifferent between accepting and rejecting a fair gamble. If the individual's utility function were convex, as in Figure 3, we would say that this agent is *risk-loving*. Such an individual would accept even unfair gambles (but would reject gambles that are sufficiently unfair).

Expected utility theory says nothing about which type of agent (or utility function) is more likely to be observed in the real world. This silence on the type of utility function likely to be observed reflects the fact that economic theory traditionally takes individuals' preferences as exogenously given. The preferences themselves are not explained by the theory. In recent years, however, theorists have started to develop models that make predictions

Figure 2: Utility of wealth function, risk-neutral individual

Figure 3: Utility of wealth function, risk-loving individual

about preferences.² In the following section, it will be argued that risk-averse preferences are more likely to be observed than either risk-neutral or risk-loving preferences.

3 Risk aversion as a consequence of the subsistence constraint

In most work that endogenizes preferences, the fundamental assumption is that preferences will evolve from one generation to the next so as to maximize expected income, or, as the biologists call it, “fitness.” Such models usually do not specify the precise mechanism by which a given preference ordering is supposed to spread in the population if it induces higher fitness. Two possible mechanisms have been proposed: *biological* and *cultural* evolution. In models of biological evolution, agents with higher lifetime incomes are assumed to have more offspring, and preferences are inherited genetically from the previous generation. In models of cultural evolution (e.g., Boyd and Richerson, 1985) the younger generation is assumed to adopt the preferences of those members of the older generation who have relatively high lifetime incomes. Alternatively, one can assume that parents attempt to instill preferences in their offspring that lead to relatively high incomes, either because the parents are altruistically motivated to maximize their children’s incomes, or because they expect a pecuniary return on their investment in their offspring in the form of old-age support (Guttman, 2001a and 2001b). In all of these versions of cultural evolution, there is an underlying assumption that agents can somehow influence their own preferences or those of their offspring, for example by educational activities, moral exhortation, etc.

Over the course of their lives, agents must make many decisions under conditions of uncertainty. As long as incomes are well above the subsistence level, maximizing lifetime income or fitness requires maximizing the expected income or wealth of these decisions under uncertainty. Intuitively, maximizing lifetime income means maximizing the *average* income of individual decisions made under uncertainty, which is mathematically equivalent to maximizing *expected* income or value. As we saw in Section 2, an individual who maximizes expected *income* behaves exactly like a risk-neutral agent maximizing expected *utility*. Since risk-neutral agents will make deci-

²For surveys of this literature, see Bowles (1998) and Robson (2001). The first paper to endogenize risk preferences seems to be Rubin and Paul (1979). For another perspective on accounting for preferences, see Becker (1996).

sions under uncertainty that maximize expected income, while risk-averse or risk-loving agents, by maximizing their expected utility, will make decisions which do not maximize expected income, it follows that risk-neutral agents will have higher fitness than risk-averse or risk-loving individuals, and thus will tend to displace them in an evolutionary context.

Now consider a subsistence economy, in which individuals routinely make decisions whose consequences are subject to uncertainty (for example, which crops to plant, where and when to hunt, etc.). If the agent's income prior to deciding where to hunt, etc., is close to the subsistence level, then the wrong decision can mean starvation. To explore the consequences of this fact with some simple mathematics, suppose that an individual in a subsistence economy must make n identical decisions under uncertainty, one per time period, over the course of his lifetime. In each of these decisions, he receives an income y if the decision leads to "success," and an income of zero if the decision leads to "failure." Assume that y is greater than the minimum income required for subsistence, while zero is less than the subsistence income. Thus, if the individual even once makes the "wrong" decision, he starves and dies. The probability of success is $p \in (0, 1)$. Then the individual's expected income, consequent to accepting the gamble, is

$$EI \equiv py + (1 - p) \cdot 0 = py.$$

We can thus treat py as the agent's expected income in the first time period of his lifetime. What would be his expected income in the second time period, *as calculated at the beginning of the first time period*? It would be the same, py , conditional on his surviving into the second time period—which has a probability of p , since, as we have assumed, with a probability $1 - p$ the individual starves in the first time period. Thus his expected income in the second time period is $p(py)$, or p^2y . Similarly, his expected income in the third time period is p^3y , and in general his expected income in time period t is p^ty . Therefore the individual's expected lifetime income, or expected wealth (ignoring time discounting for simplicity)³ is

$$EW = \sum_{t=1}^n p^ty. \quad (2)$$

Suppose that, in each time period, the individual can choose an alternative course of action, or strategy, that is less risky in the sense that the probability of success is higher, but also less lucrative when he succeeds (y is smaller). Let the expected income be the same in the two strategies. That

³Incorporating time discounting would not alter our results.

is, let $\bar{p}$ be the probability of success in the alternative strategy and $\bar{y}$ be the income yielded by the alternative strategy when the individual succeeds, while the individual's income in the case of failure remains equal to zero (implying starvation). Then we assume that $py = \bar{p}\bar{y}$. It can easily be verified that the individual will prefer the safer strategy, even though the expected income in each time period is the same in the two strategies. To see this, note first that the individual's expected wealth resulting from consistently choosing the more risky strategy is

$$EW = \sum_{t=1}^n p^t y = py \sum_{t=0}^{n-1} p^t, \quad (3)$$

while his expected wealth in the alternative strategy is

$$\overline{EW} = \sum_{t=1}^n \bar{p}^t \bar{y} = \bar{p}\bar{y} \sum_{t=0}^{n-1} \bar{p}^t. \quad (4)$$

Since $py = \bar{p}\bar{y}$ by assumption, while $p < \bar{p}$, it follows that the right-hand side of (3) must be less than the right-hand side of (4). Thus an expected wealth-maximizing agent will strictly prefer the safer strategy, even though the expected income in each time period is the same in the two strategies.⁴

If we look only at the per-period choices of the individual, then the individual's consistent preference for the safer strategy, despite the equivalence in expected income yielded by the two strategies, will make the individual *appear* to be risk-averse. Note, however, that this one-period or short-run "apparent" risk-aversion is a consequence of long-run *risk-neutrality*, since we have been asking which strategy will maximize lifetime expected wealth, which is precisely what a (nonmyopic) risk-neutral individual would maximize.

Nevertheless, we can now suggest a possible explanation of the prevalence of risk-aversion observed in experimental studies, based on the above discussion. Consider an extreme case that will simplify the exposition. Suppose that individuals are myopic: they care only about their current-period payoffs. The fact that short-run risk-aversion maximizes lifetime expected wealth means that such risk-aversion induces greater *fitness* than short-run risk-neutrality or risk-loving preferences. Therefore (myopic) risk-averse individuals drive out (myopic) risk-neutral and risk-loving agents in an evolutionary context.

⁴The same point can be made regarding a one-time deviation from one strategy to the other: a one-time deviation from the more risky strategy to the less risky strategy will increase the agent's expected wealth, and conversely for a one-time deviation in the opposite direction.

The above explanation, while simple and useful as a “first approximation,” has two drawbacks. First, it is based on the assumption that when the agent’s strategy “fails,” his income falls below the subsistence level and he therefore dies. This is an extreme assumption. Even if the subsistence constraint is relevant to individual decisions made under uncertainty, it need not be the case that failure in each individual decision would entail starvation. The second drawback is the assumption that individuals are myopic. While this assumption can be weakened without invalidating the explanation proposed above, it still needs to be assumed that individuals place excessive weight on current consumption relative to future consumption, where “excessive” is relative to the rate of time preference that will maximize fitness. (The continued dominance of risk-averse agents, even when the subsistence constraint is no longer binding, may be a result of genetic transmission of such preferences. In the very long run, risk-averse preferences will no longer dominate if the subsistence constraint continues to be non-binding, but it may be that not enough time has elapsed in order for this to happen.)

Recently, Foster and Hart (2007), as a byproduct of their development of a measure of the riskiness of choices made under uncertainty, have proposed a much more sophisticated explanation for risk-aversion. In common with the simpler approach outlined above, Foster and Hart (FH for short), in effect, introduce a subsistence constraint, and analyze a situation in which an individual must make repeated risky choices. In their setup, an agent starts with an initial wealth W_1 . At each time t , the agent can accept or reject a gamble which will either increase or decrease his wealth. For example, at $t = 1$, the agent can accept a gamble that will *increase* his wealth by x_1 with probability p or *decrease* his wealth by x_2 with probability $1 - p$. If the individual rejects the gamble, his wealth at time $t = 2$, W_2 , remains equal to W_1 . If he accepts the gamble and “wins,” his wealth at $t = 2$ is $W_1 + x_1$, while if he loses, his wealth at $t = 2$ is $W_1 - x_2$. At time $t = 2$, the agent is again offered a gamble, with possibly different payoffs and their associated probabilities, and he can again either accept or reject the gamble. The process then repeats itself, ad infinitum. The only restriction that FH place on the series of gambles is that the payoffs in each gamble, x_1 and $-x_2$ in our example, are related to the corresponding payoffs at some other time by a nonnegative multiple λ . That is, if we denote the payoffs at time t by $x_1(t)$ and $-x_2(t)$, then, for some pair of times t_1 and t_2 , $x_1(t_1) = \lambda x_1(t_2)$ and $x_2(t_1) = \lambda x_2(t_2)$.

Note that a risk-neutral agent would accept any gamble whose expected value is non-negative, i.e., $px_1 - (1 - p)x_2 \geq 0$. A major implication of FH is that if the agent wants to ensure that his wealth will never reach zero, even in

the limit as time passes, he should *not* accept every gamble whose expected value is non-negative. (FH call a zero level of wealth “bankruptcy,” while in the context of our discussion, a zero level of wealth is the subsistence level.) To ensure that the agent’s wealth will never reach zero, even in the limit, it is not even sufficient to accept all gambles with non-negative expected value *provided* that $x_2(t) < W_t$, meaning that the potential loss entailed by the gamble is less than his current wealth. While this proviso would indeed ensure that, in the *current stage*, the individual’s wealth will not sink to the subsistence level, in the *course of time* such a policy would almost surely ensure that his wealth will approach zero, in the limit, if the expected value of the gamble at each time is zero. To see why, observe that the wealth of an individual adopting this policy would follow a stochastic process similar to a random walk. Like any variable in a random walk, the agent’s wealth eventually will reach a level (if only once) infinitesimally close to zero. At this point, the agent could not accept (almost) any gamble, since any such gamble would have a potential loss x_2 that is greater than the individual’s current wealth. The agent would then get “stuck” at this infinitesimally low level of wealth.

If, then, our agent restricts himself only to gambles whose expected value is strictly positive, he behaves as if he were risk-averse. (I use the term “as if” because the theoretical framework of FH does not require the introduction of utility functions or expected utility theory.) Indeed, FH prove—subject to a mild “homogeneity” restriction on the strategies the agent is allowed to adopt—that in order to ensure that his wealth will not approach zero over the course of time, the individual should not accept a gamble g unless his current wealth is at least a specified minimum value $R(g)$, where $R(g)$ denotes the “riskiness” of the gamble (g denotes the vector of payoffs x_1 and $-x_2$ and their associated probabilities p and $1 - p$). This specified minimum level of wealth solves the equation

$$E \left[\log \left(1 + \frac{1}{R(g)} g \right) \right] = 0, \quad (5)$$

where E is the expectation operator. If the agent rejects any gamble when his current wealth is less than $R(g)$, FH prove that his wealth will almost surely *not* approach zero, even in the limit. A characteristic of $R(g)$ is that, as the expected value of the gamble g approaches zero, $R(g)$ approaches infinity. Since the agent’s wealth is finite, this means that he should reject any gamble whose expected value is zero. Therefore, a necessary, but not sufficient, condition to accept the gamble, in FH’s framework, is that the

gamble should have a positive expected value. This, of course, implies risk-averse behavior.

From an economist's standpoint, perhaps the most striking feature of $R(g)$ is that it is the minimum level of wealth required to induce an expected-utility-maximizing agent to accept the gamble g if the agent has a logarithmic utility function, $U = \log(W)$. That is, if the agent happens to have a logarithmic utility function, and faces the choice at each time t of rejecting the gamble and receiving utility $\log(W_t)$ with certainty in the next time period, or accepting the gamble g and receiving expected utility $p \log(W_t + x_1) + (1 - p) \log(W_t - x_2)$, then the agent (if he maximizes expected utility) will be just indifferent between accepting and rejecting the gamble, if his current wealth W_t equals $R(g)$.⁵ If his current wealth is less than $R(g)$, then he will reject the gamble.⁶ Since the logarithmic utility function is concave, an individual with such a utility function is risk-averse. Thus an individual whose utility function of wealth is logarithmic, if he behaves according to expected utility theory, can rest assured that he will never be driven to extinction (or even infinitesimally close to extinction) if he rationally accepts and rejects gambles in FH's series of gambles. This provides us with an evolutionary clue to the explanation of risk aversion.

In order to understand (5) intuitively, recall that if the gambles that the agent accepts always have zero expected value, the agent's wealth will behave over time like a random walk. Thus, if it were not for the constraint that the agent's wealth never can reach zero or become negative, one would

⁵This can be seen as follows. Start with the condition just stated:

$$p \log(W_t + x_1) + (1 - p) \log(W_t - x_2) = \log W_t.$$

Subtracting $\log W_t$ from both sides,

$$[p \log(W_t + x_1) + (1 - p) \log(W_t - x_2)] - \log W_t = 0,$$

which is equivalent to

$$p \log \left(1 + \frac{x_1}{W_t} \right) + (1 - p) \log \left(1 - \frac{x_2}{W_t} \right) = 0.$$

This is precisely (5) when g consists of two only possible payoffs, x_1 and $-x_2$, where $R(g) = W_t$.

Note that the definition of $R(g)$ precludes the possibility that the agent's loss when he loses the gamble, x_2 , is equal to or greater than his current wealth (thus reducing his wealth to zero), since in this case the ratio x_2/W_t in the right-hand term in the above equation would be at least equal to unity, making $\log(1 - x_2/W_t)$ undefined.

⁶This is due to the fact that the logarithmic utility function displays decreasing absolute risk aversion.

expect that, in the very long run, his wealth would stay constant, on the average. In statistical terminology, the agent's wealth would have neither a positive nor a negative time trend. The difficulty, as we noted above, is that if—due to the possibility of losing in any gamble—the agent's wealth even once becomes extremely close to zero, the agent gets “stuck” at this low level, since he will not accept a gamble whose “downside” loss, x_2 , is greater or equal to his current wealth, in order to ensure that his wealth never actually reaches zero. (Alternatively, if the agent does not impose this restriction on his acceptance of gambles, then his wealth eventually will reach exactly zero after losing in some gamble, since wealth behaves like a random walk. At that point in time, the individual will starve.) In order to avoid this “technical” difficulty imposed by the “boundary constraint” that wealth W_t always must stay positive, one can transform W_t by taking its logarithm. As $W_t \rightarrow 0$, $\log W_t \rightarrow -\infty$. Thus by adopting the requirement that the mathematical expectation of the change in $\log W_t$, consequent to accepting the gamble, is zero, one can ensure that in the very long run, $\log W_t$ —which now evolves over time in a manner similar to a random walk—will stay constant. The “trick” of the logarithmic transformation solves the problem introduced by the boundary constraint, since there is no lower or upper bound to $\log W_t$. The requirement that the expectation of the change in $\log W_t$ be zero is precisely what is specified in equation (5), where $R(g)$ is the critical W_t required to make this expectation equal to zero.

Of course, if the same expectation is positive, then over time $\log W_t$ will grow without bound, in the long run. The larger is the initial wealth W_t , the greater will be the expectation of the change in $\log W_t$ for a given gamble, since the downside loss x_2 becomes smaller in percentage terms (the change $\log W_t$ is approximately the percentage change in W_t , for small changes), relative to the upside gain x_1 as a percentage of W_t . Therefore, for any W_t larger than the critical value $R(g)$ required to solve (5), the expectation of the change in $\log W_t$ will be positive.

Let us compare the explanation of risk aversion implied by FH's contribution to the simpler explanation presented earlier. As noted above, the two explanations both involve a series of gambles, and both incorporate a subsistence constraint. Unlike the simpler explanation, however, FH do not assume that all gambles (or any gambles) will reduce the agent's current wealth to zero if the agent loses the gamble. This is an obvious gain in realism. The subsistence constraint is present, but does not necessarily “bind” the agent in any individual gamble. Also, the agent does not choose a strategy in each period in order to maximize expected lifetime wealth in FH's model, unlike in the simpler explanation. Thus there is no difficulty, in terms

of explaining preferences, of reconciling “apparent” short-run risk aversion and long-run risk neutrality. On the other hand, there is a certain lack of realism in FH’s assumption that current gains or losses are fully “reinvested”, i.e., if the agent accepts the gamble, his wealth in the next time period is simply augmented or reduced by his gain or loss in the current period. In subsistence economies, it is often difficult or impossible to store gains and losses from previous time periods, if the time period is as long as a week or a month. If risky decisions are made on a daily basis, however, then the assumption that gains and losses carry over from one period to the next becomes much more plausible.

4 Implications of the Foster-Hart model

The FH model has a number of intriguing economic implications, some of which will be discussed in this section.

4.1 Increasing income inequality

As pointed out above, the FH model suggests that the logarithmic utility-of-income function will be “selected” in an evolutionary context, since an individual who (in the FH setup) maximizes expected utility and who has a logarithmic utility function will almost assuredly never starve or go bankrupt, i.e., will never experience his or her income reaching or even approaching the subsistence level. The model derives a minimum level of wealth that makes it advisable (in order to avoid starvation) to accept a given gamble. The higher the individual’s wealth, the larger the gambles that the individual should be willing to accept. Or, to use the language of expected utility theory, the individual’s absolute level of risk aversion decreases with increasing income or wealth.

There is a great deal of evidence that absolute risk aversion indeed decreases with increases in wealth. This evidence includes early studies of portfolio composition of investors (e.g., Cohn et al., 1975, Friend and Blume, 1975, and Morin and Suarez 1983), as well as more recent, experimental and survey evidence (e.g., Donkers et al., 2001, and Dwyer et al., 2002).

If we combine this prediction of the model with the conventional assumption of a risk-return tradeoff, we obtain the implication that relatively wealthy individuals will obtain higher returns, on the average, than less wealthy investors, since the more wealthy investors will accept higher risks. Indeed, this implication of the model is supported by Yitzhaki (1987), who found that corporate stock owned by high-income investors appreciates sub-

stantially faster than the stock owned by investors with lower income. This, in turn, implies that income inequality will increase over time, a much-studied phenomenon in the recent economics literature.

4.2 Gender differences in risk aversion

A further implication of the FH model is a consequence of the fact that women historically have less personal wealth than men. The hypothesis that preferences are determined in an evolutionary context, as discussed above in Section 3, implies that since women did not control family assets since the beginning of human history, they will have developed more risk-averse preferences with regard to those decisions that they make autonomously, e.g., before and after marriage. Thus, even after controlling for wealth and income, women are predicted to be more risk-averse than men. This implication, as well, is supported by recent studies of the determinants of risk aversion (e.g., Donkers et al., 2001, and Dwyer et al., 2002).

4.3 The economic function of insurance companies

Samuelson (1963) presented a critique of the conventional notion that “there is safety in numbers.” It is commonly thought that as an individual takes on additional risks, the riskiness of his portfolio decreases. Whether or not this common belief is true depends on how “riskiness” is measured, as well as on the covariance between the individual risks. One measure of riskiness is the variance-to-mean ratio (VMR). If the individual risks are identical and statistically independent, the VMR stays *constant* as additional gambles are added to one’s portfolio. This can be shown quite easily. Suppose that x_i is the random return of individual gamble i . Let Ex_i denote the expected value of x_i , which is the mean of the probability distribution of x_i , and let $\sigma^2(x_i)$ be the variance of x_i . The variance of the portfolio of n identical, independent risks x_i is the sum of the individual variances, $n\sigma^2(x_i)$, and the expected value of the portfolio is nEx_i , so that the VMR of the portfolio is invariant to n .

This point raises a question regarding the economic function of insurance companies. Such companies are commonly thought to make it profitable to take on risks—a “bad” in a population of risk-adverse individuals—simply by virtue of the large number of risks accepted by the insurance firm. Arrow and Lind (1970), building on earlier work by Borch (1960), provided an explanation of the role of insurance firms based on the dividing of individual risks between the stockholders of the insurance firm. As the risk associated

with each insurance policy issued by the insurance firm is divided into a very large number of per-stockholder risks, the difference in the stockholder's income between the "bad" state of the world (in which the insured individual becomes ill, has a car accident, etc.) and the "good" state (in which these events do not occur) becomes extremely small. Like any other continuous function, the stockholder's utility-of-income function, even if it is not linear, can be approximated by a linear function for very small changes in income. As we saw in Section 2, a risk-neutral agent has a linear utility-of-income function. Thus the stockholder behaves as if he were risk-neutral for such small changes in income. Thus the disutility associated with risk bearing is diluted to the degree that it virtually disappears with a very large number of stockholders.

The FH model suggests an alternative explanation of the economic role of insurance firms that does not rely on the assumption that they have large numbers of stockholders. Even an insurance firm owned by a single owner has an economic role, according to the FH model. The fact that large numbers of individuals are insured by the firm, each with a small (and, to a first approximation, statistically independent) probability of enduring the "bad" state of the world, implies that the firm has a large amount of equity at any given time.⁷ The FH model implies that a firm that does not want to go bankrupt should undertake larger risks, the greater its "wealth"—i.e., its equity. Thus an insurance firm, by virtue of its "wealth," can profitably accept larger risks than individual agents, even if the firm is owned by a single owner and even though the riskiness of its portfolio (as measured by the VMR) may be same as that of individual agents.

5 Concluding remarks

We have seen that the assumption regularly made by economists—and supported by a wealth of empirical evidence—that individuals are risk-averse, can be *derived* rather than simply assumed by hypothesizing that preferences evolved over the generations in order to maximize expected wealth or "fitness". Risk aversion leads to greater fitness than risk-loving or risk-neutral preferences specifically when individuals are faced by a subsistence constraint. As per-capita income increases above the minimum required for subsistence, this constraint becomes less relevant, and, as a result, a smaller

⁷That is, the larger the number of insured agents, for a given average size of the insurance policy, the larger will be the equity of the firm, provided that the present value of premia to be paid to the firm exceeds the expected present value of payments to the insured agents and other costs—which must hold if the firm is making a profit.

degree of risk aversion is optimal in terms of fitness. This theory is supported by observations of decreasing absolute risk aversion as income increases, implying increasing income equality over time, as well as smaller degrees of risk aversion among males than among females. The economic function of insurance companies can be explained as a consequence of smaller risk aversion among insurance companies with large numbers of insured individuals, providing insurance firms with large amounts of equity capital, than among individual agents.

References

- Arrow, Kenneth J., and Lind, Robert C., 1970. Uncertainty and the evaluation of public investment decisions. *American Economic Review*, 60, 364-78.
- Becker, Gary S., 1996. *Accounting for Tastes*. Cambridge: Mass: Harvard Univ. Press.
- Borch, Karl, 1960. The safety loading of reinsurance. *Skandinavisk Aktuarietid-skrift*, 163-84.
- Boyd, R. and Richerson, P.J., 1985. *Culture and the Evolutionary Process*. Chicago: Univ. of Chicago Press.
- Bowles, Samuel, 1998. Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, 36, 75-111.
- Cohn, R.A.; Lewellen, W.G.; Lease, R.C.; and Schlarbaum, G.G., 1975. Individual investor risk aversion and investment portfolio composition. *Journal of Finance*, 30, 605-20.
- Donkers, Bas; Melenberg, Bertrand; and van Soest, Arthur, 2001. Estimating risk attitudes using lotteries: A large sample approach. *Journal of Risk and Uncertainty*, 22, 165-95.
- Dwyer, Peggy D.; Gilkeson, James H.; and List, John A., 2002. Gender differences in revealed risk taking: Evidence from mutual fund investors. *Economics Letters*, 76, 151-58.
- Foster, Dean P. and Hart, Sergiu, 2007. An operational measure of riskiness. Center for the Study of Rationality, Hebrew University of Jerusalem, Discussion Paper DP-454.

- Friend, Irwin, and Blume, Marshall E., 1975. The demand for risky assets. *American Economic Review*, 65, 900-22.
- Guttman, Joel M., 2001a. Self-enforcing reciprocity norms and inter-generational transfers: Theory and evidence. *Journal of Public Economics*, 81, 117-51.
- Guttman, Joel M., 2001b. Families, markets and self-enforcing reciprocity norms. *Annales d'Economie et de Statistique*, 63-64, 89-110.
- Malthus, Thomas R., 1798. *An Essay on the Principle of Population*.
- Morin, Roger A., and Suarez, Antonio Fernandez, 1983. Risk aversion revisited. *Journal of Finance*, 38, 1201-16.
- Robson, Arthur J., 2001. The biological basis of economic behavior. *Journal of Economic Literature*, 39, 11-33.
- Rubin, P.H. and Paul, C.W., 1979. An evolutionary model of taste for risk. *Economic Inquiry*, 17, 585-96.
- Samuelson, Paul M., 1963. Risk and uncertainty: A fallacy of large numbers. *Scientia*, April-May, 1-6.
- Yitzhaki, Shlomo, 1987. The relation between return and income. *Quarterly Journal of Economics*, 102, 77-96.