

“Rational Fatalism”: Non-monotonic choices in response to risk

Preliminary and Incomplete – Please do not Circulate

Jason T. Kerwin¹

University of Michigan Department of Economics and Population Studies Center

March 2012

Many decisions involve definite benefits but uncertain, probabilistic costs. Examples of such risky choices range from chemical exposure to unprotected sex. Previous empirical research has assumed monotonic, negative responses to risks, where people always decrease the number of risks they choose to take when the per-act chance of a bad outcome rises. This paper shows that this need not be true if rational agents have a history of risk-taking with a still-unrealized outcome. In that case there exists a tipping point above which agents switch from negative (self-protective) to positive (fatalistic) responses to risks, even if they do not use the true total probability of the bad outcome but instead use some other risk-aggregation function that satisfies a few basic conditions. I also demonstrate that the typical test for violations of monotonic responses – checking whether a quadratic term added to a linear regression is statistically significant – can yield misleading results, and develop an alternative empirical framework that is more effective. I then apply my framework to decisions about risky sex in Malawi, finding suggestive evidence of non-monotonic behavior.

¹I am grateful to Yusufcan Masatlioglu, Raj Arunachalam, Joe Golden, Emily Nicklett, and seminar participants in the IGPI conference and the University of Michigan Informal Development Seminar for their invaluable feedback. Rebecca Thornton, Ajay Shenoy, Shaun McGirr, Isaac Sorkin, and participants at the University of Michigan Extremely Informal Development Seminar provided helpful comments on an earlier version of this paper. Research on this project was supported by a Student Global Health Engagement Grant from Center for Global Health, the Population Studies Center Initiatives Fund, and the Weinberg International Travel Fund. The Situational Analysis of Sexual Behaviors and Alternative Safer Sex Strategies was funded by a grant from the Eva L. Mueller New Directions in Economics and Demography Fund. Mr. Kerwin acknowledges fellowship support from the Population Studies Center and the Rackham Graduate School.

1. Introduction

Life is full of decisions that involve weighing the definite benefits of an act against uncertain, but potentially large consequences. When people consider everything from smoking and drinking to exposure to chemicals to risky sex to merely driving to work, they face a probabilistic chance of a bad outcome, and decide whether it's worth it to take a chance. Typically research has assumed "self-protective" responses to risks – that when the per-act risk goes up, people take fewer chances.² In this paper, I show that rational actors may also respond "fatalistically" – by increasing their risk-taking when per-act risks rise, and describe a general model of "rational fatalism" in which the choice of the number of risky acts is a non-monotonic function of the per-act risk, so that behavior has a both a self-protective and a fatalistic region. This leads to a fundamental non-monotonicity in people's behavior – their responses switch from downward-sloping to upward-sloping depending on the value of the per-act risk. This non-monotonicity is driven by a history of past risks taken for which outcome is still unrealized. Logically, this might apply to HIV infection, contracting cancer from smoking, or lung damage from asbestos exposure. The logic is fairly straightforward: people who are convinced that the bad outcome is sure to happen perceive no benefit from reducing their risk-taking. In contrast with previous research I show that this result holds for any reasonable benefit function for the risky acts, and for a wide range of risk aggregation functions including those that might be employed by people with a limited background in mathematics. I also demonstrate that fatalistic behavior can exist for interior solutions, rather than solely in situations where agents choose the maximal number of risks allowed in the model.

² Economists commonly use "risk" or "riskiness" to refer to the volatility of an outcome or asset that has some upside or downside risk, and focus on the curvature of the utility function as a summary of aversion to this sort of risk. In this paper I use "risk" in the more colloquial sense of "taking a risk", in which doing something has a known (beneficial) outcome and carries some chance of a bad result.

Although theoretical work has acknowledged that fatalism can be rational, an extensive empirical literature on risky choices has almost uniformly imposed a linear functional form assumption, therefore ruling out potentially non-monotonic relationships. The portion of this literature that has studied the relationship between unprotected sex and HIV in Africa has found perplexingly small responses. Many explanations have been advanced for this limited response, but rational fatalism suggests another possibility that has yet to be studied. If a negative, self-protective response by some is balanced by a positive, fatalistic response by others, the average magnitude of the response will be lower and potentially near zero unless the researcher explicitly allows for non-monotonic responses. This is rare – only a few studies have allowed for non-monotonicity, commonly by including a quadratic term in the per-act risk. A conventional approach to looking for non-monotonicity is to examine the statistical significance of this squared term, but this can be misleading. Using data generated by the rational fatalism model, I show that a parametric test of the presence of a non-monotone relationship also performs poorly. I develop a method to identify non-monotonic responses that is effective for this model, and then apply it to preliminary data from Malawi. While the data has is purely observational, and reverse causality between risk-taking and per-act risks is likely, I do find suggestive evidence of non-monotone choices, especially among people in rural areas. This implies that research on sexual behavior and HIV transmission risks should consider the potential of rational fatalism, and look for possible non-monotonicity in responses.

I begin in Section 2 by looking at the existing body of research on choices in the face of risk, with a particular emphasis on HIV as well as on models of fatalistic behavior. Section 3 lays out a general theoretical framework that captures the possibility of rational fatalism, and discusses the conditions under which we expect fatalistic (as opposed to self-protective)

behavior. In Section 4, I develop an empirical strategy for identifying potential fatalistic behavior from data on risk-taking and perceived per-act risks. I emphasize that linear models will find attenuated responses in the face of underlying non-monotonicity, and that standard approaches to testing for a non-monotonic relationship have important limitations when applied to the rational fatalism model. Section 5 applies that strategy to preliminary observational data from Malawi's Zomba District, and shows that there is suggestive evidence of rational fatalism in that region, and Section 6 concludes.

2. Background

There is an extensive empirical literature in economics that considers decisionmaking in a context where actions have a known benefit but carry some risk of a bad outcome (a failure in the probability terminology). This section gives an overview of that literature, with an emphasis on what the present work contributes and how it differs from related research. I begin with a discussion of research on responses to risks in general (Section 2.1), and then turn to a discussion of the possibility of fatalism as a rational choice (Section 2.2). I then discuss the performance of standard empirical approaches when used on data that may be non-monotonic due to rational fatalism (Section 2.3).

2.1. Measuring responses to risks

Among the earliest studies to examine choices in the face of some per-act risk is Viscusi (1990), which shows that people who think the risk of acquiring lung cancer from smoking is higher are much less likely to smoke. Research on responses to other risks has found similar responses. Using a fixed-effects approach that exploits repeated home sales to reduce omitted-variable bias, Gayer et al. (2002) show that the release of information that cancer risks are lower than expected leads to increases in home prices, implying a rise in demand for housing. A positive response of wages to job-related mortality risks is one of the central predictions of the theory of compensating differentials. Also working with home prices, Linden and Rockoff (2008) use data from Megan's Law to show that the arrival of a sex offender in a neighborhood decreases home prices by about 4 percent, which they attribute to a decline in demand due to the increased perceived risk of crime. Viscusi (2004) uses variations in pay and risk across occupations to estimate the statistical value of a human life. He finds a significant and positive

response of wages to occupational mortality risks, consistent with a *ceteris paribus* reduction in labor supply in response to the risk of death.

Unprotected sex in the face of potential HIV infection is arguably the most important risky choice from a public policy perspective. Accordingly, the response of people's sexual behavior to risks has received extensive attention from economists, beginning with Philipson and Posner (1993), who develop a model of "rational epidemics," wherein infectious diseases are spread principally by voluntary behavior. This approach led to later work that attempted to measure the responsiveness of sexual behavior to the risk of HIV. Empirical studies of HIV and sexual behavior in the United States consistently find strong negative responses. Geoffard and Philipson (1996) estimate the parameters of a rational epidemic model of HIV using data on homosexual men from 1980s San Francisco, finding significant differences from a traditional epidemiological predictions. Based on the same dataset, Auld (2006) uses a structural model of sexual behavior to show that the rate of sexual partner change dropped rapidly in response to higher HIV prevalence in the homosexual population of that city, with a 10% rise in the prevalence decreasing the rate of partner change by 5%. Ahituv et al. (1996) use the National Longitudinal Survey of Youth to study behavioral responses across the United States as a whole, and find significant increases in condom use as HIV prevalence rises.

In contrast, there is decidedly less evidence of people responding to the risk of HIV by curtailing sexual risk-taking in sub-Saharan Africa, where the epidemic is at its worst. A study of 14 countries in the region using data on HIV prevalence and sexual risk-taking from the Demographic and Health Surveys (DHS) finds that higher HIV prevalence does decrease risky sex at several margins, after instrumenting for HIV prevalence using the distance to the origin of

the virus (Oster 2012). However, the reductions are small in magnitude: a doubling of the HIV infection rate decreases the probability of unprotected sex by just 2 percentage points.

Other research shows similarly limited responses. Also using data from the DHS and instrumenting for the distance to the origin of the virus, Juhn et al. (2009) estimate separately both the direct biological effect of HIV infection on fecundity, and the indirect behavioral effect of higher HIV prevalence through reduced unprotected sex. They find evidence of significant biological reductions in fecundity but no meaningful change in the fertility rate. Stoneburner and Low-Beer (2004) argue that with the exception of Uganda, no African country has exhibited substantial behavioral changes in response to the HIV epidemic. Consistent with this pattern of limited behavioral change in response to the HIV epidemic, Padian et al. (2010) conduct a systematic review of RCTs that attempt to reduce HIV transmission, finding that only one in seven show any impact, either positive or negative. “In fact, the overwhelming majority of completed RCTs are ‘flat’ – unable to demonstrate either a positive or adverse effect.”

Some recent research in Africa has found more encouraging results. Delavande and Kohler (2011) use panel data on probabilistic expectations in rural Malawi to study the impact of beliefs about the HIV transmission rate. Relying on optional HIV testing as a shock to people's beliefs, they find a significant negative relationship between the perceived risk of HIV transmission and the decision to have multiple sex partners. Godlonton et al. (2012) run an experiment in rural Malawi that demonstrates that when people are told that circumcised men have a relatively lower risk of HIV transmission, circumcised men do not change their sexual risk-taking but uncircumcised men have significantly less risky sex. A study of Kenyan teenagers also explored responses to relative risks: Dupas (2011) finds that a program that provided information on the relative risks of HIV infection by the age of one's partner prompted large

changes in sexual behavior in that group. Despite these encouraging results, however, the estimated response of risky sex to per-act risks in Africa is still strikingly smaller than in the United States. A variety of explanations have been advanced for this difference. . Oster (2012) argues that the limited response she finds is due in part to lower life expectancy in Africa; if people expect to live short lives anyway (due to say, malaria), HIV infection may not be as salient a threat. An alternate view is that condoms and other forms of safer sex are culturally unacceptable in an African context, thus impeding uptake; for example, in Malawi condom use in marriage is seen as an accusation or admission of infidelity and is commonly compared to eating candy with the wrapper still on (Chimbiri 2007). In this paper I argue that a third factor may also contribute to a limited average response: fatalism. If some people are responding fatalistically to risks – by increasing their risk-taking when risks rise – then this will cancel out some of the self-protective responses in the data.

2.2. Fatalistic responses to risks

Empirical research on responses to risks, whether in Africa or elsewhere, has almost uniformly assumed that these responses conform to the “self-protection” model in which choices are a declining linear function of riskiness. In contrast with this is the possibility of the opposite pattern, in which response are actually positive: as the per-act risk goes up, people take more risks. This kind of behavior is commonly thought of as an irrational, and referred to as “fatalism”. The connotation is that it is unreasonable and lies soundly in the realm of behavioral theories that break with rationality. In the context of HIV, for example, “fatalism” is usually used to refer to the seemingly irrational pattern of just giving up on avoiding infection due to hopelessness or a lack of regard for personal safety. This behavior was observed as long ago as the late 1980s in Uganda: Barnett and Blaikie (1992) discuss men who were aware of the risk of

contracting HIV and simply appeared not to care, asking “Who is never going to die?”

Economists have also focused on non-rational explanations for fatalistic behavior. Leon and Miguel (2011), for example, demonstrate that travelers in Sierra Leone reveal a lower willingness to pay for reductions in mortality risk than Americans, despite the fact that they have comparable incomes and remaining life expectancies; they argue that this may be explained by the perceived role of fate in determining life outcomes in West African societies.

However, it is not necessarily the case that fatalism must arise from people behaving irrationally. Fatalistic behavior can actually be perfectly rational: in their early treatise on the economics of HIV, Philipson and Posner (1993) point out that selfishly rational actors will tend to demand more risky sex as their probability of already having HIV goes up. In the limit, the logic is simple: if I already have HIV, I gain no benefit from using a condom and doing so carries some cost. More generally, O’Donoghue and Rabin (2001) develop a simple model of the expected cost of risk-taking, showing that for some parameter values the marginal cost of another risky act actually decreases when the per-act risk rises. While their analysis focuses on people being unwilling to go below some minimum level of risky choices, in Section 3 I show that the same logic will apply when people have already engaged in risk-taking in the past and do not yet know the outcome of that risk-taking. Using a special stepwise functional form for the benefit of the risky behavior, they show that fatalistic responses can be rational, and note that along with other potential applications such as drug use, their model may have relevance for the risk of HIV infection as well. More recently, Sterck (2011) develops a theoretical framework that uses the same cost function as O’Donoghue and Rabin (2001), but in a dynamic setting. Using parameter values derived from data on Burundi, he argues that believing the risk of HIV transmission is high can lead to rises in risky sex.

The idea of rational fatalism can be seen as a specific form of the Mickey Mantle effect, in which people invest less in their health when their life expectancy is lower for reasons unrelated to the investment in question (Fang et al. 2007). Fang et al. find that this decline in health investments is heterogeneous across health behaviors, with smoking being more responsive than heavy drinking. In a rationally fatalistic model, the specific shifter of life expectancy is past choices of the very same risk under consideration.

Is rationally fatalistic behavior relevant to the study of responses to risks? It may well be, at least in the context of the HIV epidemic in Southern Africa. Qualitative evidence from high-prevalence areas in the region suggests that indeed rationally fatalistic behavior is potentially common there. In research on how rural Malawian men discuss HIV and risky sex, Kaler (2003) documents many cases in which single “freelancer” men use fatalistic reasoning when thinking about the disease. One informal conversation on the topic recorded in the Kaler (2003) study proceeded as follows:

Friend: I don't fear AIDS because I know that I have it already.

Diston: How do you know that you have got AIDS?

Friend: I have malaria and some coughs so I know that I have it.

Diston: Do you use condoms when [sleeping with] these bargirls?

Friend: What for, since I know that I am already infected? (Kaler, 2003)

Kaler describes many cases in which high perceived risks are leading sexually experienced individuals to fatalism: even though they have never tested positive for HIV, they give up on safer sex, having decided that they cannot avoid developing AIDS.

2.3. Measuring the response of risk-taking to per-act risks

The possibility of fatalism as a rational response to increased risks has commonly been discounted in the literature on estimating risk responses empirically. For example, Viscusi (1990) uses the one-tailed version of the t-test for his statistical significance calculations, thereby assuming that responses can only be either negative or zero. In the case of HIV, researchers typically follow Philipson and Posner (1993) in assuming that the overall prevalence of the virus is sufficiently low that few if any people believe they are HIV-positive. This justifies the assumption of a linear risk-response relationship, since for low values of both the number of risks taken and the per-act risk, the probability of infection is approximately linear in the number of risks. Given the low prevalence in most US populations, ignoring potential fatalism may be justifiable in this context. However the prevalence of HIV in Africa is much higher, making this assumption harder to justify, and as noted above at least some people in Malawi use rationally fatalistic reasoning to explain their own behavior.

This implies that any empirical study of the response of sexual behavior in Africa to the risk of HIV prevalence must allow for responses to be potentially non-monotonic. Almost no existing research on the effect of risks on behavior allows for non-monotone effects. Some studies have considered heterogeneity in responses, for example by gender and marital status (Oster 2012) and by level of risk-taking (Auld 2006).³ In an experiment studying the effect of providing HIV test results to people in Malawi, Thornton (2008) explicitly considers combating fatalism as one of the mechanisms through which HIV testing could potentially affect sexual behavior. Selfishly rational people who believe they are HIV positive, and find out they are not due to a test result, are likely to reduce the amount of risky sex they have. She therefore

³ The Auld (2006) approach comes close to allowing us to directly examine the rational fatalism model laid out in Section 3, but he explores heterogeneous responses by *current* level of risk-taking, whereas rational fatalism is driven by variation in *past* risks taken.

explicitly looks for heterogeneous effects by HIV status, but finds that only HIV-positive individuals respond significantly to the test results, and that the effect is small in magnitude. One study related to HIV and sexual behavior that does not impose linear responses is de Paula et al. (2009). However, they are examining the impact of beliefs about one's own current HIV status on sexual behavior, and so are not directly comparable with the other literature discussed above.

Since almost none of literature on the response of sexual behavior to risks allows for that response to be non-monotonic, this suggests an additional explanation for the lower estimated responses to risks in Africa than in the US. If people in Africa are more likely to be fatalistic, then their positive responses to risks will cancel out some of the negative response on average. Since a linear regression of risk-taking on per-act risks measures the average slope across the population, this would tend to attenuate the estimated response.

How can one tell if such a linear regression is likely to be misleading? One common approach to testing for non-monotonic relationships is to run some regression specifications that include a quadratic term, which allows for the common technique of examining the statistical significance of the second-order coefficient to determine whether a relationship is non-monotonic. However, even this approach may fail to reveal the non-monotone responses typified by rational fatalism.⁴ Recent work by Lind and Mehlum (2010) shows that a significant quadratic term may arise even if the relationship is monotonically negative, and propose a formal parametric test for a U-shaped relationship. Their technique also has its drawbacks, because it requires an approximately quadratic functional form for the data. As I discuss in detail in Section 4, if rational fatalism is potentially at work in generating data on responses to risks, careful analysis is necessary to determine whether the relationship is monotonic.

⁴ De Paula et al. (2009) also look at heterogeneous effects by quantile of perceived risk, which would perform far better than the use of a quadratic term in a regression.

3. Theoretical Framework: The Rational Fatalism Model

How can existing approaches to studying behavioral responses to risks be modified in order to accommodate the possibility of rational fatalism? In this section I will develop a model of rational fatalism that extends the logic employed in the previous literature in a simple but powerful way. This model replaces the linear probability function used in most research on risk-taking with a risk aggregation function that represents an agent's belief about the probability of a failure given the number of risky acts chosen and the perceived riskiness of each act. Using this model, I will show that a tipping point into fatalistic behavior emerges naturally from a wide range of possible risk aggregation functions, including the true total probability derived from the binomial distribution.

I begin this section by laying out the basic form of the model, and the conditions that must be satisfied for an agent's chosen number of risky acts to be optimal (Section 3.1). Using a minimal set of assumptions about the optimization problem, I show the conditions under which an interior optimum will and will not exist. I demonstrate that a non-zero fixed price for each risky act (which could also be thought of as a time cost or an emotional cost) will exclude the possibility that agents just choose as many risky acts as possible. (Section 3.2) I then explore the model's general comparative statics, focusing in particular on how the optimal number of risks taken changes in response to variation in the perceived per-act risk. To explore these comparative statics, I prove that any risk-aggregation function that satisfies basic conditions will have a tipping point, above which additional increases in riskiness decrease the marginal cost of risk-taking. I use this finding to show that there are tipping points in individuals' behavior as well: above a certain point, additional risks lead to more risk-taking rather than less (Section 3.3). After showing these tipping points exist, I then discuss two fairly simple heuristics that

people might use to think about how risks add up, and show that despite not being differentiable they also exhibit the tipping point that is central to my result, implying that the behavior I model could hold even for unsophisticated agents (Section 3.4). In Section 3.5, I show that the rational fatalism model implies that risk responses are qualitatively different depending on the domain being considered: the standard negative response occurs in when individuals face low per-act risks or have a small amount of previous risk-taking, while individuals facing a combination of both high perceived risks as well as substantial past risk-taking will tend toward fatalistic behavior. Finally, I discuss the ways in which this model differs from previous theoretical work, in particular the fact that it holds for any valid risk aggregation function and that it shows that fatalistic risk responses can occur for interior solutions, and not just in situations where agents take the maximum possible number of risks (Section 3.6).

3.1. Model Basics

In this model, I assume that people weigh the benefits of choosing a level of risk-taking, n , against both any monetary costs as well as the expected cost of a stochastic bad outcome occurring due to their choice.⁵ Each risky act carries some perceived per-act risk r (its “riskiness”) that it will cause the bad outcome to occur.⁶ The benefit of n risky acts is described by a continuously differentiable benefit function, $B(n)$, with $B'(n) > 0$ and $B''(n) < 0$ so the marginal benefit of risk-taking is positive with diminishing returns. For simplicity I normalize $B(0) = 0$. There is some non-probabilistic cost $p > 0$ for each act that might be thought of as a monetary cost, the cost of the time devoted to the act, or even an emotional cost or guilt, so n acts cost pn . The expected cost of the bad outcome is the *perceived* probability of it occurring,

⁵ For the sake of simplicity I will assume that n is continuous, rather than a discrete number of acts. This follows O’Donoghue and Rabin (2001). The simulations in Section 4 will demonstrate that the qualitative results derived in this section for continuous n will also hold for discrete-valued n .

⁶ This perceived risk does not have to equal the true risk; most people overestimate per-act risks across a wide range of activities from smoking (cf. Viscusi 1990) to unprotected sex (see Section 5 of this paper).

P , times its perceived cost, c . Given a number of acts and a per-act risk, the *true* probability can be computed using the binomial distribution. I reserve this for later discussion, however, and allow people's perceived probability of the bad outcome (a "failure", to use terminology from probability theory) to be some general (continuously differentiable) risk-aggregation function $P = P(r, n + m)$. Here r , the perceived per-act risk, is assumed to lie between 0 and 1. m is the number of previous risks taken for which the outcome has not yet been realized, and is weakly positive. In other words, this model explicitly considers agents for whom some of their risk-taking history is still unresolved; for example, people in areas with no accessible HIV testing who are still in the window period between risky sex they might have had and the point at which they would develop AIDS symptoms. In order to ensure that $P(r, n + m)$ corresponds to well-formed probabilities, I impose the following basic assumption:

Assumption 1

- A. $P(r, n + m) \in [0, 1] \forall r, n, \text{ and } m$.
- B. $P(r, n + m) = 0$ if $r = 0$, or both n and $m = 0$.
- C. $P(1, n + m) = 1$ iff $n \neq 0$ or $m \neq 0$. $\lim_{n \rightarrow +\infty} P = 1$ and $\lim_{m \rightarrow +\infty} P = 1$ iff $r \neq 0$.

Probabilities must never be negative or greater than one. The probability of the bad outcome is zero if either the activity is risk-free or if the individual does not engage in the activity at all. Taking any risks will lead to the bad outcome occurring with certainty if the per-act risk is 100%. Likewise, choosing an unbounded number of risky acts will eventually lead to the bad outcome happening for sure, as long as the act has some non-zero risk associated with it.

This simply imposes that the probabilities produced by the risk-aggregation function start from zero and rise to one as we increase either the riskiness of an individual act or total risk-taking. In

addition, I assume that P is never decreasing in any of its arguments, and strictly increasing to begin with. I also impose that additional acts do not raise the total probability of a failure if the acts are riskless, and that raising the per-act risk does not increase the total probability if no risks are taken.

Assumption 2

$P_1 \geq 0$, with $P_1(0, n + m) > 0$ if $n + m > 0$ and $P_1(r, 0) = 0$; $P_2 \geq 0$, with $P_2(r, 0) > 0$ if $r > 0$ and $P_2(0, n + m) = 0$.

Increasing the per-act risk will never decrease the overall probability of a failure; likewise raising the number of risky acts chosen or the previous stock of risky acts always (weakly) increases the total probability of a failure. An initial increase in riskiness or risk-taking strictly raises the total probability of a failure (as long as total risk taking or riskiness, respectively, are non-zero). Conversely, increasing the number of riskless acts chosen, or the riskiness of a risky act that is not chosen, does not affect the total probability of the bad outcome occurring.

Taken together, these assumptions simply state that people must have a general understanding of risks, so that they understand that probabilities never fall outside of a 0-100% range and that additional risk taking is bad, up to the limit imposed by maximal probability of 1. They immediately lead to an initial lemma:

Lemma 1

A. $\lim_{r \rightarrow 1} P_1 = 0$

B. $\lim_{n \rightarrow +\infty} P_2 = 0$

The effect of increasing the per-act risk on the total probability of the bad outcome is zero if the per-act risk is one. The effect of the number of risks chosen or the existing stock of risks approaches zero as the sum of those variables approaches infinity.

Lemma 1A holds trivially if $n+m = 0$, and likewise for Lemma 1B if $r = 0$. To see why they must hold in the non-trivial case, assume they do not hold. Then P is unbounded. But by assumption P is bounded above at 1, so we have a contradiction. Therefore Lemma 1 must hold in general.

Note that because P is continuously differentiable, Lemma 1A also implies that

$P_1(1, n + m) = 0$. Conceptually, Lemma 1 says that increasing the riskiness of the act high enough, or taking a sufficient number of risks, pushes the likelihood of the bad outcome to 100%. Once it has reached that point, additional risk-taking does not increase the probability any further.

The agent's optimization problem is therefore the following:

$$\max_{n>0} \{U(n; r, m, p, c)\} = \max_{n>0} \{B(n) - pn - P(r, n + m)c\} \quad (1)$$

By the assumption that n is continuous, the maximand $U(n; m, p, c, r)$ is the sum of continuously differentiable functions and therefore continuously differentiable itself.

3.2. Conditions for a Non-trivial Optimum

Since the object of interest in this analysis is the response of n to changes in r , one concern is whether the solutions to the problem are purely trivial, with fatalism representing jumps to some maximal level of risk taking. In this section I show that a non-stochastic price for each risky act guarantees that we will find interior solutions unless risk-taking is not beneficial at all at the agent chooses to take zero risky acts.

The above optimization problem admits many conceivable forms for the benefit function $B(n)$, including some that make little intuitive sense. To restrict the discussion to reasonable benefit functions, I assume that at some point taking additional risks yields no utility gains.

Assumption 3

$\lim_{n \rightarrow +\infty} B'(n^*) = 0$. As the number of risky acts chosen approaches infinity, the marginal benefit from an additional risky act approaches zero.⁷

Under Assumptions 1-4 the problem still admits trivial corner solutions where $n^*=0$. In order to discuss interior solutions, I impose one additional assumption.

Assumption 4

$B'(0) > p + P_2(r, 0 + m)c$. Risk-taking is desirable: given the stochastic and non-stochastic costs of risky acts, agents will choose a non-zero level of risk-taking.

Empirically, Assumption 4 seems reasonable in many applications: for example, a large proportion of people have had unprotected sex at some point in their lives. If the converse of Assumption 4 holds, agents will (weakly) prefer to set $n=0$, and the problem becomes trivial.

Given Assumption 4, however, the model allows a fairly powerful statement to be made:

Proposition 1

$\exists n^* \in (0, \infty) : n^* = \arg \max_{n \geq 0} \{U(n; r, m, p, c)\}$ if $p > 0$.

An interior solution to the optimization problem described in (1) is guaranteed whenever the non-stochastic cost (e.g. the price) of a risky act is not zero.

Proposition 1 follows because $\lim_{n \rightarrow +\infty} [B'(n^*) - p - P_2(r, n^* + m)c] = -p < 0$ by

Assumption 3 and Assumption 4, and because $B'(0) - p - P_2(r, 0 + m)c > 0$. This, along with

⁷ This assumption is substantively identical to the sixth Inada condition used to guarantee the stability of neoclassical growth models.

the continuity of U , allows me to use the extreme value theorem to state that U has at least one optimum where $n^* \in (0, \infty)$, as long as $p > 0$.

This eliminates the possibility of trivial corner solutions, in which the optimal response to an increase in risk is always to either choose $n^* = 0$ or $n^* = n_{max}$ (where n_{max} is some upper bound on n that prevents it from reaching infinity). Conversely, if $p=0$, then given the other conditions the optimal n^* will be arbitrarily large: U is initially upward-sloping and its slope never becomes negative, so additional risk-taking is always weakly beneficial. The other analyses of optimal risk-taking that admit fatalistic responses (Sterck 2011; O'Donoghue and Rabin 2001) have shown fatalism only as a corner case, in which the individual pursues the maximum feasible level of risk-taking. While corner solutions are a fairly intuitive response – they align with the reasoning that once one is doomed, one might as well indulge as much as possible – they are not empirically relevant: there is little evidence that individuals ever truly seek out the *maximal* level of available risk-taking. Moreover, the reason for this is exactly that given above – taking additional risky acts, whether that means smoking more or seeking out sex partners, carries pecuniary costs so that there are tradeoffs with out goods an individual might desire.

Proposition 1 guarantees that the optimum will be non-trivial if the price of risk-taking is positive. It does not rule out interior optima in other cases; O'Donoghue and Rabin do have an interior optimum in their model's non-fatalistic case, for example. However, it is a fairly intuitive economic result: people are constrained by resources from pursuing the high extreme in risk-taking. The results that follow will hold for the commonly-seen case in which people pursue some intermediate level of risk-taking irrespective of their perception of the per-act risk r . In the following section I will show that fatalism can occur even for these interior solutions.

3.3. Comparative Statics

Given that an interior solution exists, the optimal choice of n , n^* , must satisfy the first-order condition:

$$B'(n^*) - p - P_2(r, n^* + m)c = 0$$

where $B'(n)$ and $P_2(r, n + m)$ are the derivatives with respect to n of $B(n)$ and $P(r, n + m)$. n^* also needs to satisfy the second-order condition that the utility function be concave at that point.

$$B''(n^*) - P_{22}(r, n^* + m)c \leq 0$$

It is not possible to solve for n^* without additional assumptions about the functional form of $B(n)$ and $P(r, n + m)$, and there will be no closed-form solution for n^* for most possible functional forms of the benefit and risk aggregation functions.⁸

Despite the intractability, in general, of the precise optimum n^* , it is possible to explore the response of n^* to changes in other variables without solving for the optimum analytically by employing the implicit function theorem (IFT). In particular, there is a function $G = B'(n^*) - p - P_2(r, n^* + m) * c = 0$, and therefore the IFT allows us to compute the comparative static for changes in n^* in response to changes in r , as well as the way that response changes when c varies.⁹

$$\begin{aligned} \text{I. } \frac{\partial n^*}{\partial r} &= -\frac{\frac{\partial G}{\partial r}}{\frac{\partial G}{\partial n^*}} = \frac{P_{21}(r, n^* + m)c}{B''(n^*) - P_{22}(r, n^* + m)c} \\ \text{II. } \frac{\partial^2 n^*}{\partial r \partial c} &= \frac{\partial}{\partial c} \left(-\frac{\frac{\partial G}{\partial r}}{\frac{\partial G}{\partial n^*}} \right) = \frac{P_{21}(r, n^* + m) * B''(n^*)}{(B''(n^*) - P_{22}(r, n^* + m)c)^2} \end{aligned}$$

For comparative static I the denominator is $B''(n) - P_{22}(r, n^* + m)c$. This is precisely the left-hand side of the second-order condition (1.4), and is therefore weakly negative. It is not possible

⁸ In particular, if I impose that P be the true risk-aggregation function derived from the binomial distribution and that B have a logarithmic form, then no closed-form solutions for n^* are possible.

⁹ In Appendix A I derive additional comparative statics with respect to p and c .

to rule out the possibility that the second-order condition is exactly zero in general, so the optimum occurs at a flat region of the utility function. This would mean that neither comparative static would exist, and the model would predict neither self-protective nor fatalistic responses. Since such flat regions of the utility function seem unlikely, I will assume the second-order condition holds strictly.¹⁰ The denominator of comparative static II is just the square of the same expression and therefore strictly positive. Its numerator is the product of $B''(n)$ and $P_{21}(r, n^* + m)$; $B''(n)$ is negative as long as $n < +\infty$. As a result, comparative statics I and II will have the same sign, which is opposite to the sign of $P_{21}(r, n^* + m)$.

Before deriving the sign of $P_{21}(r, n^* + m)$ in general, I first consider the specific case where agents use the true risk-aggregation function $\Phi(r, n + m)$. This function comes from the binomial distribution, and can be constructed as follows. Given a stock of past risk-taking m , a choice of the number of additional risks to take n , and the per-act risk r , the probability of a failure after one act is r ; after two acts, $r + r(1 - r)$; and after three, $r + r(1 - r) + r(1 - r)^2$. The probability of avoiding a failure after three acts is $1 - [r + r(1 - r) + r(1 - r)^2] = (1 - r)^3$. Likewise the probability of avoiding a failure after $n + m$ acts is $(1 - r)^{n+m}$. Therefore the chance of a failure having occurred by after $n + m$ acts is:

$$\Phi(r, n + m) = 1 - (1 - r)^{n+m}$$

The relevant second derivative from Comparative Statics I and II for this functional form is then

$$\Phi_{21}(r, n^* + m) = (1 - r)^{n^*+m-1} [1 + (n^* + m)\ln(1 - r)]$$

Interestingly, the expression for $\Phi_{21}(r, n^* + m)$ may be positive or negative: the first factor, $(1 - r)^{n^*+m-1}$, is weakly positive, but $[1 + (n^* + m)\ln(1 - r)]$ is either greater or less than zero

¹⁰ Note that even in the limit as the denominator approaches zero all the results in this section will go through.

depending on the values of n , m , and r . It is possible to solve the value of r that determines the sign of the second term.

$$\begin{aligned}
 1 + (n^* + m)\ln(1 - r) &< 0 \\
 \ln(1 - r) &< -\frac{1}{n^* + m} \\
 1 - r &< e^{-\frac{1}{n^* + m}} \\
 r &> 1 - e^{-\frac{1}{n^* + m}}
 \end{aligned}$$

This result holds even if we assume the agent is at $n = 0$. Then the threshold value of r becomes $1 - e^{-\frac{1}{m}}$. This threshold depends on the level of past and current risk-taking. Using the true risk-aggregation function $\Phi(r, n + m)$, as $(n + m)$ approaches zero, the value of the per-act risk which leads to fatalism approaches 1. As $(n + m)$ becomes sufficiently large, the critical value of r can become arbitrarily close to zero. Imagine a decisionmaker who has a long history of taking chances. If he believes the per-act risk is low, and he finds out it is quite high, then he will be nearly-certain he is already doomed. Thus the cost of taking one more chance goes down substantially.

The above proof of the existence of a tipping point for the true risk-aggregation function $\Phi(r, n + m)$ is due to O'Donoghue and Rabin (2000). However, whereas they consider only the true risk-aggregation function $\Phi(r, n + m)$, it is possible to show that the same qualitative result holds for any credible risk-aggregation function $P(r, n + m)$. To prove this I first show that the cross-partial is initially positive:

Lemma 2

$P_{21}(0, n + m) > 0$ and $P_{21}(r, 0) > 0$. The cross-partial derivative of the total probability of a failure with respect to riskiness and number of risky acts chosen is positive when the number of risky acts

This follows straightforwardly from Assumption 2: P_1 is zero if n and m are both zero and positive if at least one of n or m is positive, so the initial cross-partial is positive; a symmetric analysis holds for P_2 .

Given Lemma 2, we can therefore prove that this cross-partial changes sign in general, for all functions P that meet the conditions laid out above.

Proposition 2

$\exists \tilde{r} = r(n+m)$ with $n+m < +\infty$ s.t. $P_{21}(r, n+m) > 0$ if $r < \tilde{r}$ and $P_{21}(r, n+m) < 0$ if $r > \tilde{r}$. For sufficiently high values of the per-act risk, increasing the per-act risk actually diminishes the marginal impact of additional risk-taking.

To prove this, I consider two functions $P_{r_L}(n+m) = P(r_L, n+m)$ and

$P_{r_H}(n+m) = P(r_H, n+m)$ with $r_L < r_H$. By Lemma 2,

$$P'_{r_L}(0) < P'_{r_H}(0)$$

Assumption 1 also gives us

$$P_{r_L}(0) = P_{r_H}(0) = 0$$

and

$$\lim_{n+m \rightarrow +\infty} P_{r_L}(n+m) = \lim_{n+m \rightarrow +\infty} P_{r_H}(n+m) = 1$$

Then these two continuous functions begin at the same value and converge to the same value, but the slope of P_{r_H} is initially higher than that of P_{r_L} . This implies that there must be some point at which the slope of P_{r_L} exceeds that of P_{r_H} . If not then the value of P_{r_L} can never catch up with that of P_{r_H} .

Formally, consider a point n_1 sufficiently close to zero that $P_{r_L}(n_1) < P_{r_H}(n_1)$, which must be possible because the second function's slope is initially higher. Then the average slopes of the two functions between n_1 and some higher point n_2 are $\frac{P_{r_L}(n_2) - P_{r_L}(n_1)}{n_2 - n_1}$ and $\frac{P_{r_H}(n_2) - P_{r_H}(n_1)}{n_2 - n_1}$,

so the ratio of the two slopes is $\frac{P_{r_L}(n_2) - P_{r_L}(n_1)}{P_{r_H}(n_2) - P_{r_H}(n_1)}$. Taking the limit as n_2 approaches infinity, this ratio approaches $\frac{1 - P_{r_L}(n_1)}{1 - P_{r_H}(n_1)}$, which is greater than one. This implies that there is a point above which the average slope of P_{r_L} exceeds that of P_{r_H} ; therefore, by the extreme value theorem there must be a point where the instantaneous slope $P_{r_L}' > P_{r_H}'$. Figure 1 illustrates why this must be the case. The solid blue line gives the known initial shape of P_{r_L} and likewise the solid red line for P_{r_H} . Above the breakpoint at infinity, the two-colored line shows their common value of 1. The dashed lines show the implied average slopes in the intermediate region; because P_{r_L} is initially shallower, it must be steeper on average over this range.

This ensures that a tipping point must exist in any valid risk-aggregation function $P(r, n + m)$. It does not rule out multiple tipping points, which could conceivably arise from sophisticated curvature of the risk-aggregation function, but the number of such tipping points must be odd. I will ignore the possibility of multiple tipping points, motivated by the fact that for the true risk-aggregation function Φ the cross-partial derivative changes sign only once.

Because the impact of riskiness on the marginal cost of risk-taking has a tipping point, responses to riskiness will have tipping points as well. Formally we have the following two comparative statics:

Proposition 3

$$\frac{\partial n^*}{\partial r} < 0 \text{ if } r < \hat{r}, \text{ and } \frac{\partial n^*}{\partial r} > 0 \text{ if } r > \hat{r} \text{ (Comparative Static I)}$$

There exists a threshold value for the per-act risk, \hat{r} , at which rational behavior switches from self-protection (negative responses to risks) to fatalism (positive responses).

Proposition 4

$$\frac{\partial^2 n^*}{\partial r \partial c} < 0 \text{ if } r < \hat{r}, \text{ and } \frac{\partial^2 n^*}{\partial r \partial c} > 0 \text{ if } r > \hat{r} \text{ (Comparative Static II)}$$

Increasing the cost of a failure will increase the magnitude of the responsiveness of risk-taking to per-act risks, making it more negative when agents are self-protective and more positive when they are fatalistic.

Comparative Static I is my central result, which is that there is a tipping point not just in the marginal cost of risk-taking (as O'Donoghue and Rabin (2001) show for the true risk aggregation function $\Phi(r, n + m)$) but also in the optimal choice of n . Comparative Static II is an extension of the finding of Oster (2012), who shows develops a model of rational responses to HIV risks to show that agents will respond more to the per-act risk of HIV infection when the costs are higher, e.g. when non-HIV mortality in their area is lower. In this model, a similar result holds, but only for self-protective agents - those who face risks below \tilde{r} . Above \tilde{r} , higher costs will tend to encourage more fatalism.¹¹

Thus the model implies that under fairly broad and plausible assumptions, rationally fatalistic responses will occur for sufficiently high combinations of the per-act risk r and the past stock of risk-taking m . While somewhat surprising, this result is consistent with the intuitive notion, expressed by the men quoted in Kaler (2003), that having made enough mistakes in the past can doom you to HIV infection, no matter what you do to protect yourself now. If HIV is unavoidable, attempting to mitigate your own risk of contracting it is useless.

This applies to any situation where there is a risk from each act chosen, and the outcome goes unrealized for an extended period of time. Consider an individual who knows he has engaged in some risk-taking in the past in the past. If his accumulated stock of risky acts, and his perceived per-act risk, are sufficiently high, then learning that the per-act risk is lower than he had thought can actually lead him to take fewer risks.

¹¹ Another result from Oster (2012) is that all else equal higher costs will lead to less risk-taking. This also holds for my theoretical framework. See Appendix A for a derivation.

3.4. Other Risk Aggregation Functions

The results above hold for a broad range of possible risk-aggregation functions that satisfy a minimal set of conditions, including the true function $\Phi(r, n + m)$. However, the central point – that behavior will swing from self-protection to fatalism for sufficiently high values of r – is driven by a tipping point in impact of riskiness on the marginal cost of riskiness. This kind of tipping point may exist even for far simpler heuristic risk aggregation functions that agents might employ, in particular ones that are not differentiable and therefore not amenable to the techniques employed in Sections 3.1-3.3. I therefore cannot prove that an interior optimum exists for such functions, or that optimal risk-taking will switch from self-protective to fatalistic. Instead, I demonstrate that two very simple heuristic risk aggregation functions exhibit this tipping point phenomenon.

It might seem that this sort of tipping point is an esoteric mathematical feature of how probabilities add up that people cannot be expected to understand, but in fact such tipping points arise naturally and in a comprehensible way from some fairly basic heuristic risk aggregation functions. Consider the simple linear function used in much of the literature, where the assumption is made that levels of risk-taking and per-act risks are sufficiently low that the probability never approaches 1. Agents might use a similar rule, but also assume that if the probability does reach 1 then it stays there forever:

$$P(r, n + m) = \begin{cases} \gamma r(n + m) & : \gamma r(n + m) < 1 \\ 1 & : \gamma r(n + m) \geq 1 \end{cases}$$

This function might appear to lack a tipping point as defined in Proposition 2, but the same basic behavior actually obtains. Consider two agents, one who believes $r = 0$ and one who believes $r = 1/\gamma(n + m) - \varepsilon$. If both agents increase their risk belief by 2ε , the marginal cost of

increasing n rises for the first agent and falls for the second. Any shift in r that increases its value to at least $1/\gamma(n + m)$ will induce fatalism, with further increases having no additional effect on behavior.

An even simpler alternative is the "exposed enough" heuristic discussed in MacGregor et al. (1999), wherein people think they are totally safe as long as they stay below some level of activity, and then doomed with certainty if they take too many risks:

$$P(r, n + m) = \begin{cases} 0 & : \gamma r(n + m) < 1 \\ 1 & : \gamma r(n + m) \geq 1 \end{cases}$$

In this case only the act that shifts an agent over the threshold, $n = 1/(\gamma r) - m$, has a direct marginal cost – all other acts carry no cost at all. Increasing r will in general push agents closer to the margin of being “sufficiently exposed” to suffer harm, thus carrying an indirect marginal cost. But if r reaches or crosses $1/\gamma(n + m)$, the agent believes he or she is already sure to suffer the bad outcome and hence this decreases the marginal cost of an additional act to zero.

Despite not being amenable to analysis through standard optimization techniques, these functions both exhibit the crucial tipping-point phenomenon, implying that the results of Section 3.3 could hold even if agents handle the addition of risks in a very simple and heuristic way.

3.5. Domains of Self-Protection and Fatalism

Comparative Statics I and II identify a threshold value of r (that depends on $n + m$) above which agents will become fatalistic, and increased per-act risks will lead to more risk-taking rather than less. That is, the rational fatalism model implies that risk responses will not be monotonic, but will shift from negative to positive when r reaches \tilde{r} . This differs sharply from much of the empirical literature on risks, in which the expected cost of infection, is linear in the per-act risk: $P(r, n + m) = \gamma r(n + m)$. Under the linear model, risk responses will be

monotonic and negative; there is no tipping point. In other words, the rational fatalism model and the linear model give qualitatively identical predictions for any behavior that occurs below the tipping point \tilde{r} .

As a result, the rational fatalism model implies that we should expect to see three distinct patterns depending on the domain in which it is applied. In settings that are either low-risk or low-activity, it predicts negative responses consistent with the conventional linear model. In settings that are both high-risk and high-activity, it predicts positive, fatalistic responses. Finally, in heterogeneous settings, we expect to find a roughly U-shaped response, with some amount of self-protection as well as some fatalism.

The first domain includes the Kenyan teenagers studied by Dupas (2011) - they are sexually inexperienced, and therefore even fairly high perceived per-act risks still cause them to behave in a self-protective fashion. The second domain is exactly that discussed in Kaler (2003) - men with extensive sexual experience who perceive high per-act HIV risks end up rationally fatalistic; they are doomed no matter what, so why bother using condoms. The third domain is the most interesting, and most comprehensive: most populations include both people who think they have some past exposure and those who think they have never taken a risk. In the majority of cases I expect observed patterns of risk responses to include both self-protective and fatalistic behavior.

3.6. Comparison with Existing Theoretical Work

As discussed in Section 2.2 above, previous theoretical work has studied fatalism as a potentially rational response to risks, with both O'Donoghue and Rabin (2001) and Sterck (2011) developing formal mathematical models that predict fatalistic responses at some margins. The

rational fatalism model outlined in the above sections improves on those models in several valuable ways.

First, previous research has only been able to demonstrate fatalistic behavior as a corner solution, where agents either pursue a low, self-protective level of risk-taking or jump to a point where they fatalistically take as many risks as possible. Sterck (2011) shows that without random mistake or a condom failure, people will choose either no risk-taking or maximal risk-taking and stay there forever, while O'Donoghue and Rabin use a contrived utility function to show that agents will sometimes react to increased per-act risks by taking as many risks as they possibly can. While O'Donoghue and Rabin do discuss the implications of their results for interior solutions, they do not demonstrate that interior solutions actually exist when responses are fatalistic, as I did in Section 3.2. In contrast the results of Section 3.3 are explicitly conditioned on interior solutions; the predictions hold specifically for values of n^* that are *not* at a maximum or minimum allowed value. This is guaranteed to be possible because individual behavior is constrained by an implicit budget constraint; every risky act carries a cost p . In Section 4 I will show that a wide range of interior solutions exists for plausible parameter values.

Second, both Sterck (2011) and O'Donoghue and Rabin (2001) combine the number of risky acts being chosen and the unavoidable minimum number of acts into a single variable, and neither separates out potential past mistakes from current decisionmaking (although Sterck portends this line of reasoning when discussing mistakes). The rational fatalism model draws an explicit separation between the unavoidable number of risky acts, m , and the current number of acts being chosen, n . It also focuses on the case where the m unavoidable acts probably occurred in the past, but that the outcome (malaria infection, lung cancer, HIV transmission) has not yet been revealed.

Third, both previous papers rely on specific, simple benefit functions for the risky acts that potentially raise questions about the robustness of the results to other functional forms, whereas the model laid out in this section is robust to any concave, increasing benefit function $B(n)$ with sufficiently strongly diminishing returns.

Finally, the previous theoretical work has used the true risk aggregation function to show interesting effects of varying per-act risks on choices of risky behavior. This is a valuable line of inquiry, but extensive research has demonstrated that in addition to overestimating per-act risks, individuals often do not understand how those risks compound into the total probability of a failure. O'Donoghue and Rabin (2001) note that evidence on HIV transmission beliefs indicates that the risk aggregation functions individuals employ tend to be far more concave than the true function. The rational fatalism model shows that similar non-monotonic behavior can be found for *any* valid risk aggregation function that satisfies very simple conditions. In addition, I illustrate similar tipping point behavior even for very simple, non-differentiable risk aggregation functions that people might realistically use even if they know very little math.

4. Empirical Approach

Based Comparative Static I above, any analysis of data on risk-taking must consider the possibility of a non-monotonic response to the per-act risk. In this section I lay out the regression specifications commonly employed in the literature on risk-taking (Section 4.1). I then use the model from Section 3 generate simulated data (Section 4.2) in order to examine the effectiveness of the basic approach. I show that under plausible parameter values, a squared term in per-act risks can be statistically insignificant even when the underlying model has a tipping point. I also find that in addition to missing the fact that the relationship has a U shape, running simple linear regressions will tend to attenuate the magnitude of the measured risk response (Section 4.3). I then describe two methods that may be more successful – the U-shape test of Lind and Mehlum (2010) and a semi-parametric technique that lets the user look for non-monotonicity directly (Section 4.4). I find that the Lind and Mehlum method is not particularly successful unless the functional form of the relationship is close to quadratic, but that the semi-parametric approach has promise. I propose a combined method, which uses the Lind and Mehlum test but checks that the relationship is close to quadratic semi-parametrically. This substantially out-performs the conventional approaches from the literature: a first-order relationship between risk-taking and per-act risks can show a significant relationship even when the underlying behavior is non-monotone, and a second-order term in risks can be insignificant even when non-monotonicity exists. Finally, in Section 4.5, I move to a discussion of several threats to any attempt to use observational data to identify the relationship perceived per-act risks and risk-taking, and solutions to those problems.

4.1. Regression specification

The standard approach used in the literature is to run regressions of the form

$$n_i = \beta_0 + \beta_1 r_i + Z_i' \delta + \varepsilon_i$$

where Z_i is a vector of controls. This is fine so long as the relationship between r and n is monotonic, but as discussed in Section 3 above there is reason to believe it may not be. While relatively little research has considered the possibility of a non-linear relationship between r and n , a typical strategy for doing so is to add a squared term to the above regression (cf. the De Paula et al. 2009 study of the relationship between people's own perceived HIV status and risky sex). This yields the following specification.

$$n_i = \beta_0 + \beta_1 r_i + \beta_2 r_i^2 + Z_i' \delta + \varepsilon_i$$

One method for determining whether an estimated relationship is non-monotone is to examine the statistical significance of the quadratic term, r^2 .

4.2. Simulated data

To explore the effectiveness of these regression specifications, I construct simulated data using the model described in Section. Specifically, I draw a pseudorandom sample of 3000 individual-level observations, comparable to the size of many field surveys on sexual behavior in Africa, with the following parameter value distributions. I set the past level of risk-taking m to 2 acts for all individuals. I then set the mean of t to the threshold value for declining marginal costs, $1 - e^{-\frac{1}{2}} \approx 0.3935$, and $SD(r)$ arbitrarily to 0.2.¹² I use a logarithmic utility function $\alpha \ln(n)$, setting $\alpha = 1$ for all observations, and impose the true risk-aggregation function, $P(r, n + m) = \Phi(r, n + m) = 1 - (1 - r)^{n+m}$. I set the cost of a failure to 1 for all individuals. I choose the mean of p to be 0.15 and its standard deviation to be 0.03 arbitrarily. Using these values, I engage in a grid search over possible integer values of n ranging from 0 to some upper

¹² Note that this is only the tipping point if agents choose zero risky acts ($n=0$). This is tautologically never observed since if $n=0$ no relationship is present. Since agents will in general choose some level of risk-taking greater than zero, the observed tipping point in the data will differ for each individual, and will be smaller than 0.3935.

limit (which was selected randomly to lie from 15 to 21 for each individual) in order to find the highest total utility value. The values picked above generate almost entirely interior solutions: none of the simulated data points had $n^* = 0$, and only three were at their individual-specific limit for n .

4.3. Parametric regression results

The first two columns of Table 1 show the simple linear model and a model with both a linear and a quadratic term respectively, both without any regression controls. Both specifications find a significant positive first-order relationship between risk-taking and per-act risks. The estimated quadratic term in specification 2 is positive but statistically insignificant. Together with the first-order term this would imply a convex parabola with a minimum below zero. The only relevant potential control is the cost per risky act, p . Controlling for the per-act cost does not substantially affect our coefficient estimates, but it does improve the precision with which they are estimated. While the estimated sign of the first-order term is uniformly positive, implying a positive average relationship between risks and costs, this is simply the result of the specific parameter values chosen. Other simulations (not shown) give a negative first-order term but results that are qualitatively similar to what follows.

The results of this analysis are revealing – although this data was simulated using a model with a built-in tipping point, an approach that looks at the significance of the second-order risk term would conclude that the relationship is probably monotonic.

Another notable aspect of these regressions can be seen by comparing the coefficient on the first-order term in Columns 1-4 to that in Column 5; as I shall discuss in more detail later, Column 5 does a better job of matching the shape of the curved part of the function. Irrespective of whether a squared term is included, these regressions estimate first-order coefficients that are

far less than 20% of the magnitude of the true coefficient, and of the opposite sign. This implies that standard regression specifications may yield misleading results when applied to data where rational fatalism is possible.

4.4. Identifying U-shaped relationships

To address this shortcoming, I first explore the formal parametric test for U-shaped relationships proposed by Lind and Mehlum (2010). They argue that a significant quadratic term is necessary, but not sufficient, for a non-monotone relationship. The logic is simple: if a relationship is monotonic, but concave, the best fit to the data will in general involve a significant second-order coefficient. But the turning point in the predicted parabola might be far outside the data range. Thus it is necessary to formally test whether the fitted values involve both a statistically significant upward-sloping and downward-sloping component. Table 1 includes the p-values for this test in Columns 2 and 4, which in both cases are extremely close to 1.¹³ For this dataset, the Lind-Mehlum test concurs with the simpler method of looking at the coefficient on the squared term, and indicates that there is almost surely no non-monotonicity in the relationship between n and r . This is unsurprising – the squared coefficient in one case does not meet the necessary condition of being statistically significant, and in the other just barely crosses that threshold.

Since the Lind-Mehlum method still relies on the estimated coefficients from fitting a second-order polynomial to the data, one potential concern is that the model fit is sufficiently poor that we are drawing false inferences. In many cases the solution would be to derive a superior regression model based on what is known about the data generating process. However, as discussed in Section 3, the utility maximization that generated this data has no closed-form

¹³ Tests conducted using the `utest` function in Stata, written by Lind and Mehlum. Technically both rejections are trivial, since the estimated minimum point of the parabola would have been outside the data range regardless.

solution so we would be left with a fairly complex challenge of numerical optimization, that might yield a wide variety of solutions depending on the parameter values chosen. A natural approach to this issue is to estimate the relationship non-parametrically, in order to avoid any issues of model selection. There are two general methods for attempting this.

The first is to simply do a scatterplot of the values of the two variables, and possibly fit a non-parametric curve estimate to that data. Figure 2 presents a simple scatterplot (with no curve fit to it) of the number of risky acts against the per-act risk. While a possible U-shape is evident, the bivariate relationship is fairly noisy.

An alternative is to use a semi-parametric approach that relies on parametric regression to reduce the noise and omitted variable bias induced by other variables. I do this by using the partially linear model estimator as developed by Yatchew (1997) and implemented by Lokshin (2006) to plot the conditional relationship of n and r , holding p constant. The results are shown in Figure 3. The plot also includes a plot of the LOWESS-estimated non-parametric curve, but the non-monotone relationship is clear without it.¹⁴

These results imply that the Lind and Mehlum approach was indeed hampered by poor model fit; the relationship, while non-monotonic, is also clearly not a parabola. As an additional check, I repeat the same regression as in Column 4 of Table 1, but apply it to just the portion of the data between $r = 0$ and $r = 0.4$. The results are shown in Column 5 of the same table. Now the linear and quadratic terms are large and statistically significant, and the sign of the linear term is negative. The Lind and Mehlum test rejects monotonicity at beyond the 0.001 level.

Given the limitations of the Lind and Mehlum test for analyzing this model, it is sensible to look for a semi-parametric equivalent. For bivariate non-parametric analysis, options do exist.

¹⁴ The lowest observed value of n is 3, implying a tipping point of 0.1813 or less. This is roughly consistent with the observed values.

One example is Bowman et al. (1998), who develop a formal test for the monotonicity of a non-parametric locally linear regression function. They rely on the use of a critical bandwidth that flips the estimated relationship from monotonic to non-monotonic, and use it to construct This has substantial promise for the analysis of risk-response relationships, but unfortunately there is currently no equivalent for partially linear regressions. That limits the direct applicability of their method, since confounding omitted variables are likely to substantially bias the observed relationship between n and r . Adapting the Bowman et al. approach to partially linear regressions is beyond the scope of this paper, but will be a focus of future work.

Lacking a formal semi-parametric test for non-monotonicity, I will rely on a strategy that combines the Lind and Mehlum test and partially linear regressions. My formal test for non-monotonicity will use the Lind and Mehlum approach, but I will confirm that the relationship is approximately quadratic using semi-parametric regressions. If it is not, but possible non-monotonicity is evident, I will use the data truncation method described above, applying the Lind and Mehlum test to the portion of the data that is potentially non-monotone.

The preceding analysis demonstrates that the econometric methods commonly used in the literature perform poorly when data is generated in a manner consistent with the rational fatalism model of Section 3. The standard regression specifications for measuring risk-taking as a function of per-act risks will tend to identify a slope that is too small in magnitude and potentially of the wrong sign. Moreover, if a regression includes both linear and quadratic terms in risks, a statistically significant squared term is neither necessary nor sufficient for non-monotonicity. Correspondingly, the formal parametric test for non-monotonicity from Lind and Mehlum (2010) can fail to identify the fact that the data has a tipping point if the model is sufficiently mis-specified. In contrast, the approach outlined above – using the Lind-Mehlum test

along with a semi-parametric approach to check for the model fit – appears to perform much better. A semi-parametric approach is also advisable for confirming that the estimated regression coefficients are sensible.

4.5. Threats to identification

Even if an appropriate approach to measuring potential non-monotonicity is adopted, using observational data on the relationship between per-act risks and risk-taking can lead to false inferences for three reasons: the data is affected by two different kinds of reverse causality, which may bias the results in opposite directions, and also subject to substantial measurement error, which will attenuate measured effects toward zero.

The first reverse causality issue has to do with cases where the per-act risk is actually the result of previous risk-taking, which may be correlated with current levels of risky choices. This is most notably the case in the literature on HIV and sexual risk-taking: the per-act risk depends on the prevalence of the virus, which in turn depends on how much risky sex people have. Other infectious diseases share this property - malaria, for example, is likely to be more common in places where mosquito net use is lower. For this sort of risk, simple cross-sectional comparisons would find high risk-taking in high-prevalence places, and draw a false inference. This issue will tend to bias estimates toward positive infinity.

Another sort of reverse causality occurs when individuals form their beliefs about the risks of various acts through experience with risk-taking. To remain with the example of HIV, suppose a decisionmaker believes that she will quickly show symptoms after contracting the virus. After engaging in a sufficient quantity of risky sex over time, and not developing symptoms, her perceived per-act risk would decline, generating a mechanical, negative relationship between risk-taking and per-act risks. This reverse causality problem will have the

opposite effect from the first one, tending to bias estimates toward negative infinity. It is possible that some belief-formation processes might generate positive biases as well.

Separate from reverse causality is the problem of poorly-measured data on sexual risk-taking and risk perceptions. Data on sexual behavior may be even more susceptible to recall errors and biases than other behaviors studied using survey data, and also carries the risk of “social acceptability bias”, where respondents say what they think the enumerator wants to hear. Active research in survey methodology has focused on the development of life-history and diary-based methods, in order to reduce these sorts of measurement errors (Luke et al. 2009). Data on risk perceptions present even more severe measurement problems. People with limited math backgrounds may have trouble understanding questions about probabilistic expectations, and this issue is likely to be exacerbated in developing-country settings (Attanasio 2009). Encouragingly, recent work conducted in Malawi by Delavande and Kohler (2009) has shown that, through carefully designed questions, it is possible to elicit meaningful beliefs about probabilities even when individuals have very limited formal education.

Errors in both reported sexual behaviors and risk beliefs are likely to take many forms, and will often be effectively random. Randomly mismeasured sexual behavior will tend to decrease the precision of the estimated relationship between n and X , while random measurement errors in beliefs will attenuate any estimated relationship, biasing estimates toward zero. There is one particular kind of measurement error in risks that may not be random and is worth observing, however. A certain fraction of individuals in all contexts will answer any probabilistic question with “50%”, indicating not that they think it is a 50-50 chance but that they simply do not know (Lillard and Willis 2001). This is true even if questions are posed deterministically, as in “what share of your friends owns an iPod”. One crude way of estimating the measurement error in a

probabilistic belief variable is to look at the extent to which there is an anomalously high “50%” response rate.

One factor that is likely to increase measurement error in probabilistic beliefs is the use of true probabilities as proxies for beliefs. This has been employed in research on HIV and risk-taking in Africa (Juhn et al. 2009, Oster 2012) as well as in the United States (Ahituv et al. 1996, Auld 2006). Since many individuals will be misinformed, this issue will tend to exacerbate the attenuation bias problem described above. This kind of measurement error also raises the question of what structural relationship is being estimated. Individuals are less likely to be aware of the actual probabilities than policymakers, who may intervene to promote reductions in risk-taking. Thus a significant estimated response may not reflect actual behavioral responses by individual people.

The first two identification threats can best be resolved by instrumental variables approaches. An exogenous shock to beliefs will allow the causal effect of changes in those beliefs on risk-taking to be measured. For example, Oster (2012) and Juhn et al. (2009) use the distance from the origin of the human immunodeficiency virus as an instrument for its prevalence. To study individual risk beliefs, the optimal instrument would actually be some kind of quasi-random information campaign, or an outright experiment that provides information about risks.

Resolving the measurement error issue is somewhat harder. A simple first step is to rely on individual probabilistic beliefs rather than the true per-act risks, since the two may differ substantially. Much recent work in HIV and risk beliefs in Africa has begun to do this. (Godlonton and Thornton 2011, Delavande and Kohler 2011, de Paula et al. 2008). Beyond that, careful questionnaire design, and variables that can serve as cross-checks within a survey, are

useful tools. Question design should focus on minimizing the “50%” response rate (within reasonable bounds), for example by asking a followup question about whether respondents are just unsure (Lillard and Willis 2001). Another approach is to test the sensitivity of the results to excluding respondents who answer “50%” to the relevant question.

5. Preliminary Results

In this section I apply the empirical strategy laid out in Section 4 to a preliminary observational dataset on sexual behavior and perceived HIV risks in Malawi. Section 5.1 describes the data used, detailing the construction of variables to capture the number of risks taken, n , and the perceived per-act risk, r . In Section 5.2 I conduct a basic regression analysis of the form described in Section 4.1, and show that the standard approach of examining the sign of a squared term in risks implies at best fairly fragile evidence for a non-monotonic relationship. Section 5.3 applies the method for testing for a U-shaped relationship described in Section 4.4, finding that the evidence for a U shape is actually fairly robust. Based on this I argue that the evidence is broadly consistent with possible fatalism in rural areas on Malawi's Southern District, which is in line with the findings of previous qualitative research. In Section 5.4 I examine the plausibility of fatalism in this population in two different ways: first, I compute the number of people who would be expected to tip into fatalism based on the true risk-aggregation function Φ ; second, I explore just the subset of individuals who think their sex partners are HIV-positive. Both explorations suggest that fatalism is plausible for this population. Section 5.5 discusses some important limitations of this analysis and argues that my results are likely to be smaller in magnitude than the true relationship, and may actually understate the extent of fatalism in this population.

5.1. Data

I use data collected in the Zomba District of Malawi's Southern Region. Malawi is in the midst of a severe HIV epidemic, with an infection rate of nearly 12%. Its experience has been typical of countries in Southern Africa in that its high infection rate has been sustained for a long period, and has declined only slightly from its peak (National AIDS Commission, 2003). HIV is

the leading cause of non-infant mortality in the country, and is responsible for one in every three adult deaths (PEPFAR 2008).

The data form the peri-urban and rural survey components of the Situational Analysis of Sexual Behaviors and Alternative Safer Sex Strategies which I helped run as part of a collaboration with PIs Professor Rebecca Thornton of the University of Michigan and Dr. Jobiba Chinkhumba of the University of Malawi College of Medicine, and co-investigators Sallie Foley, LMSW, of the University of Michigan and Alinafe Chibwana, MA, of Catholic Relief Services (Kerwin et al. 2011). We conducted the surveys in July of 2011 in areas adjoining a major trading post in Zomba District. The survey enumerators were hired near the survey sites, and all surveys were conducted in Chichewa, the national language, with female enumerators conducting all interviews with female respondents and likewise for males.

We drew a geographic representative sample of 447 sexually active adults (age 18-49) in the area based on the locations of households, with selection rules for picking participants within each household. Our sampling strategy oversampled urban areas and unmarried individuals, so we constructed sample weights to match the population proportions by locale and marital status to the 1998 Malawi census. All reported statistics use these sample weights unless otherwise noted.

Summary statistics for the sample's demographics are given in Table 2, below. The sample is 5% peri-urban, meaning people that live around the major trading post, and about 60% literate. The average age is 29 and people have completed five and half years of school on average. Nearly 90% of the weighted sample is married. The sample is also almost 60% female; anecdotally, people in the region reported that men were more likely to migrate away for work. Respondents have 3 children on average and want just over one more. The sample is 80%

Christian and 20% Muslim, and is dominated by the Lomwe ethnic group at 55%. Other large ethnic groups include the Yao (22%) and Chewa (16%).

Table 3 summarizes the data on the perceived risks and costs of HIV for our respondents, as well as their self-reported sexual behavior. The data contains both a measure of perceived past exposure to HIV (the number of a respondent's past partners that, looking back, they now think were HIV-positive at the time) and of their potential future exposure (the share of attractive people in the region who the respondent thinks have HIV). The former measure appears to be more or less uncorrelated with current sexual behavior, so I rely on the latter. Respondents report that they think that 49% of attractive people, on average, are HIV-positive. The survey also asked respondents how many out of 100 people who had sex with an HIV-positive person last night would contract the virus. Over 50% of respondents answered that 100% would become infected, and the mean response was 85%. I multiply these two variables to construct a measure of the perceived per-act risk from unprotected sex with a possible sex partner, or the variable X from the framework described in Section 3. The average value of r in this sample is 43%.

The survey also asks about people's beliefs about their life expectancy, including without HIV, if they contracted HIV today but did not have ARVs, and if they contracted HIV today and did not have ARVs. In addition, it asks them how many out of 100 people from their area they would expect to have access to ARVs. I use these to construct a measure of the expected cost of an HIV infection in terms of years of life lost. On average, respondents believe HIV infection will cost them 18.9 years of life – they expect to live almost 30 years if HIV-free and 11.4 years if they contract the virus.¹⁵ While probably lower than the expected cost of HIV in a developed country, this is hard to square with the idea that limited responses to the threat of HIV infection

¹⁵ These numbers do not add up properly because some data is missing for each variable.

are the result of low expected costs in terms of foregone years of life. I will use this variable as the equivalent of C from the Section 3 theoretical framework.

Finally, the dataset contains a detailed sex diary that captures detailed information on all sex acts in the past week. I use this to construct a measure of the total number of sex acts and the number in which no condom was used. The average respondent had unprotected sex 1.9 times in the preceding week. In the analyses that follow, I will rely on this variable as the measure of sexual risk-taking from n from the theoretical framework in Section 3.

5.2. Regression Analysis

The focus of this analysis is to explore the relationship between n (Unprotected sex acts in past week) and r (Average Prob. of HIV Xmission per act, attractive people) for my representative sample of individuals from Malawi's Zomba District. As a first pass, I construct a simple, unweighted scatterplot of the two variables (Figure 4). No obvious bivariate relationship is evident. A simple bivariate regression does not find a significant coefficient on n either (not shown).

However, the oversampling of unmarried individuals and people near the trading post may mask an underlying relationship between n and r . I therefore run a set of weighted regressions, using the specifications described in Section 4.1. Table 4 presents the results. Column 1 is a bivariate regression of n on r , but using the sample weights. I estimate a statistically significant *positive* first-order relationship between perceived per-act risks and the number of risks taken. Taken at face value, this would imply that the population is fatalistic on average. However, this result is not robust to the inclusion of a second-order risk term (Column 2), which yields a negative estimated slope. The coefficient on the quadratic term in risks is positive but imprecisely measured.

Columns 3 and 4 repeat the specifications from Columns 1 and 2 but include a broad set of regression controls (suppressed from the tables for space). Again the bivariate relationship (Column 3) is positive and significant. Including a quadratic term in risks (Column 4) reverses the sign of the first-order term and decreases the accuracy with which it is estimated. The quadratic term is statistically significant at the 5% level.

In Columns 5 and 6 I extend the model to capture Comparative Static IV (from Appendix A), which states that the number of risk acts is declining in the cost of a failure (captured by the expected years of life lost if the respondent contracts HIV) and Comparative Static II, which states that the higher the cost of a failure, the stronger the relationship between risk-taking and per-act risks is likely to be. My results contrast with Oster (2012) in that I do not see a particularly strong effect of the cost of HIV on the risk-response relationship. Increasing the number of life years lost to the disease by 10 will decrease the slope of the n - r relationship by just 0.1 in the quadratic specification of Column 6, just 2% of its overall magnitude. However it is evident that cost of an HIV infection was an important omitted variable in the other regressions. The first-order term becomes more negative in both specifications and in the second-order specification the squared term becomes more positive, and in both specifications the coefficients are more precisely estimated.

Comparative Static V also states that there should be a tipping point in the effect of the cost of a failure on the relationship between n and r , and that the two tipping points should be equal. I therefore include an interaction between the cost of a failure at the square of the per-act risk in Column 7. This still does not lead to statistically significant results for Comparative Static V, but it decreases the precision of the estimates for the linear and quadratic terms in risk.

Because Column 7 captures both of the relevant comparative statics that can be measured in this dataset, I will rely on it as my preferred specification for further analysis.

5.3. Testing for a non-monotone relationship

Looking just at the squared coefficients in Table 4 would suggest that there is some evidence of a non-monotone relationship between per-act risks and risk-taking but that it is not robust to changes in the specification used. However, as shown in Section 4, this method is likely to be misleading if the underlying data-generating process is based on the rational fatalism model of Section 3. Instead, I employ the two-step test described in Section 4.4. First, I run the Lind-Mehlum parametric test for a U-shaped relationship. The p-values for this test are given in the third-to-last row of Table 4. While the unconditional relationship (Column 2) does not display significant non-monotonicity, the test rejects monotonicity for the other three specifications at the 0.1 level. In the preferred specification of Column 7, the p-value is 0.096.

To confirm that the Lind-Mehlum test is not misleading due to poor model fit, I run a partially linear regression of risk-taking on per-act risks, including all the covariates from Column 7 (with the exception of the squared term in risks, which will be captured by the LOWESS curve). This regression is plotted in Figure 5. Although the LOWESS curve has a flattened U shape, no obvious relationship is visually evident from the plotted datapoints. However, the partially linear estimator employed does not accept sampling weights, meaning that these results are potentially biased due to the oversampling of certain groups – unmarried individuals and especially people from rural areas. Because of this, Figures 6 and 7 show the results split out by peri-urban and rural residents. Since rural residents are 95% of the weighted sample, I focus on Figure 7 as a basic approximation of the weighted partially linear regression. This plot shows a more noticeable U-shaped relationship, implying that the weighted relationship

we applied the Lind-Mehlum test to has a reasonably parabolic shape and the test results are probably trustworthy.

In summary, there is suggestive evidence for a non-monotone relationship between risk-taking and per-act risks in this dataset. Moreover, the first-order relationship is consistently positive, which is the opposite of the prediction of the simple linear self-protection model. Strikingly, these results are based on a weighted sample that is 95% rural. This is consistent with the finding of Kaler (2003) that some men in rural areas employ rationally fatalistic reasoning. People near the peri-urban trading center may have more access to HIV testing, which would tend to resolve the uncertainty about their HIV status and eliminate the fatalistic effect.

5.4. Plausibility of Fatalism

While these results imply that risk responses may be fatalistic for some people in Zomba District, it is not clear how plausible this is – could people conceivably think their risk is really high enough that they cross the tipping point value for per-act risks, \tilde{r} ? As a first pass at this question, I compare people's beliefs to a plausible tipping-point: I assume that people use the true risk-aggregation function $\Phi(r, n, m)$. Then the tipping point value of r is $1 - e^{-\frac{1}{m}}$. If people think they have just a single unprotected exposure to a partner who is HIV-positive, their tipping point is about 63%; higher levels of exposure imply a lower tipping point. Looking at the distribution of beliefs in this population, (which is highly-skewed toward 100%), I find that 78.9% of people have beliefs that would be above the tipping point by this standard: if they have any past exposure to someone they are sure is HIV-positive at all, their beliefs about per-act risks are high enough to lead to fatalism.

A more conservative approach is to use people's beliefs about their past partner's HIV status (at the time they had sex) multiplied by their perceived per-unprotected-act risk of

contracting HIV from an infected partner. By this standard, 9.7% of people have risk beliefs that are consistent with fatalism even if they only had unprotected sex with their past partners a *single* time – far lower than the actual figure, since condom use is very rare in Malawi. Thus even by a very conservative standard, one would expect a non-trivial number of people from Zomba District to react fatalistically to HIV transmission risks.

An alternative method of examining the plausibility of fatalism in this sample is to isolate how people behave when they think a sex partner is infected with HIV. For this I rely on a set of questions on the SASB survey that asked respondents about their primary sex partner, which was defined to be the person they had sex with the most times in the past week, or, if they had not had sex in the past week, then the person with whom they had most recently had sex. Although a large share of respondents are married, we did not restrict this partner to be the respondent's spouse (and intentionally did not ask if they were). Respondents were asked about the likelihood that this partner was HIV-positive; 52 of the 447 respondents in the dataset either reported a high likelihood or said that they were sure their partner was infected. Figure 10 presents a scatterplot for these 52 individuals, with their perceived per-act risk of infection from sex with an infected partner on the x-axis and the number of times in the past month that they had unprotected sex with this partner on the y-axis. Immediately notable is the large cluster of beliefs at 100%. Many of these respondents also report positive levels of unprotected sex with their partner. A best-fit line confirms a positive relationship between risk-taking and perceived riskiness. Due to the small sample size, this relationship is not statistically significant. It is nevertheless suggestive of potential fatalism among people who think they are being exposed to HIV.

5.5. Limitations

In Section 4.5, I lay out three potential identification issues with using observational data to analyze risk responses. With the present dataset I am unable to address either of the two sources of reverse causality that are likely to bias these results. I expect that the first reverse causality issue is of fairly limited importance. The HIV epidemic in Malawi is fairly mature and all individuals face very similar true prevalences. Moreover, people's perceived risks have little relationship with the actual risk they face. The actual per-act risk of HIV transmission is about 0.1% (Wawer et al. 2005), far below the average belief of 85%. And the true prevalence of HIV in Malawi's Southern Region is 17.6%, also substantially less than the mean belief for this variable (49%). Since people do not actually understand the true risk, the causal positive effect of behavior on the true risk is unlikely to cause a large bias in my results.

The second issue, in contrast, is more important for these data than for data that employ the true HIV prevalence as a proxy for risks. Since people develop their beliefs in part through experience I expect this to introduce a net negative bias in these results. Without an exogenous source of variation in risk beliefs it is not possible to determine the exact nature of the bias in these results, but the negative bias from the second issue is probably more important than the positive bias due to the first one. That would imply that the true relationship would be initially flatter than what is implied by the partial linear regressions in Figures 5, 6, and 7, and then steeper after the tipping point. In other words, this observational analysis is likely to undermeasure the true extent of fatalism in this population. This would not be the case if the belief formation process worked in the opposite way, with more experience leading to higher risk beliefs.

The measurement error issue is more tractable within the current dataset. As discussed in Section 4.5, one simple metric for the measurement error in my risk belief variable is the excess frequency of “50%” responses. Figure 8 shows a histogram of the per-act risk beliefs of my respondents. There is substantial heaping at exactly 50%. Unfortunately, no viable alternative variables are available.¹⁶ The survey questionnaire did not ask a followup question about whether people are simply unsure, so it is not possible to exclude only those people or including a dummy for being unsure. As a crude approximation to that method, however, I replicate the regressions from Section 5.3 and 5.4 using only the subset of people who are not “uncertain”, meaning they did not answer 50% to the per-act risk variable. The results are presented in Table 5 and Figure 9. I find uniformly stronger responses to HIV risks across all specifications, and lower p-values for the Lind-Mehlum test. The preferred specification in Column 7 finds large and statistically significant first- and second-order terms, and the Lind-Mehlum test rejects monotonicity at the 0.05 level. The semi-parametric plot shows a reasonably clear-cut U shape. While this approach is not perfect, it implies that measurement error is substantially attenuating these results, so the true response will be larger in magnitude and more statistically significant than what is seen in Table 4 and Figures 5 through 7.

¹⁶ There is a question about the total prevalence of HIV in the population, but a printing error meant that it was collected only for males. As mentioned above, the data does have the prevalence among past partners, but this has no appreciable relationship with risk-taking, either positive or negative.

6. Conclusion

I develop a theoretical model that generalizes the linear risk-response relationship assumed in the literature to the case where responses may not be linear. I do this by allowing people to employ a subjective risk aggregation function that satisfies several broad conditions about its shape, and show there is a tipping point above which increases in the per-act risk lead to more risk-taking rather than less. This result holds for any valid risk-aggregation function that satisfies a set of simple conditions. Even very simple heuristic risk aggregation functions that require no sophisticated understanding of probability theory also exhibit the tipping point that is central to my results. The rational fatalism model implies that responses to risks will have both a downward-sloping (self-protective) and upward-sloping (fatalistic) region. It advances the previous literature by showing this effect not just for specific simple benefit functions and the true risk aggregation function, but for a wide range of plausible choices. The rational fatalism model also shows that fatalistic responses can occur for interior solutions and not just in situations where people choose to take as many risks as possible. Based on this model, and imposing some assumptions about the benefit from risky acts and the risk aggregation function, I generate simulated data and use it to test the effectiveness of standard econometric approaches to data on risk-taking a per-act risks. I find that the typical specifications can generate misleading inferences: they may fail to identify non-monotonic relationships and will generate estimates of the average response that are attenuated relative to the true value and may also be of the incorrect sign. I develop my own approach based on the Lind and Mehlum (2010) parametric test for non-monotonicity and the Yatchew (1996) partially linear regression technique. I apply these methods to preliminary observational data from Malawi's Zomba District, and find suggestive evidence of non-monotonicity in that region. I also show that depending on how we think about

past exposures, between 9.7% and 84.4% of people in the region have beliefs that are above the tipping point into fatalistic behavior implied by the true risk aggregation function. Looking just at the subset of 52 individuals who think their primary sex partner is HIV-positive, I find that risk-taking increases with perceptions about the per-act risk, although the sample is not large enough to rule out a slope of zero. These data are limited by potential measurement error as well as reverse causality. A first pass at the measurement error issue confirms that it is probably attenuating the measured relationship between risk-taking and per-act risks, and that the true U-shape is even more pronounced than observed in this data.

The results of this paper are subject to several important limitations. First, while the rational fatalism model substantially extends the usual linear risk response relationship used in the literature, it abstracts from the possibility of multiple periods. In particular, agents may be aware of the effect of their current behavior on their future decisions about risky acts. This model is appropriate for individuals with short planning horizons, or as an abstract model of risky behavior among adults where the cost of a failure is simply dying before old age (when no risks will be taken).

Second, while it demonstrates that data generated by the rational fatalism model can display non-monotone responses, and that conventional methods of looking for that non-monotonicity may not perform well, this is done only for a single simulated dataset. A superior approach would be to generate a large number of such datasets under plausible assumptions about parameter values, and explore the general effectiveness of various methods. This would in turn necessitate developing a formal and automated way of testing for non-monotonicity that works for this model, probably by adapting Bowman et al. (1998). An additional difficulty with

doing this is that many combinations of seemingly-plausible parameter values generate only corner solutions, which are not amenable to any kind of test and hence uninformative.

Third, I cannot resolve the reverse-causality issue in my preliminary data, since it lacks an exogenous source of variation in risk beliefs. I argue that correcting it is likely to decrease the extent of self-protection in measured in the data while increasing the extent of fatalism, so that this preliminary data is probably underestimating the extent to which Malawians from that district are fatalistic and not overestimating it. However, it is impossible to know what the true relationship is without an instrument or an experiment. Future work on this topic should focus on conducting an experiment which provides information about HIV transmission risks to people in Malawi and in other places where people substantially overestimate the per-act risk of contracting HIV.

References

- Ahituv, A., Hotz, V. J., & Philipson, T. (1996). The Responsiveness of the Demand for Condoms to the Local Prevalence of AIDS. *The Journal of Human Resources*, 31(4), 869–897.
- Arunachalam, R., & Shah, M. (2009). *Compensated for Life: Sex Work and Disease Risk in Ecuador* (Working Paper). Population Association of America Annual Meeting. Retrieved from <http://paa2009.princeton.edu/abstractViewer.aspx?submissionId=91998>
- Attanasio, O. (2009). Expectations and Perceptions in Developing Countries: Their Measurement and Their Use. *American Economic Review*, 99(2), 87–92.
- Auld, M. C. (2006). Estimating behavioral response to the AIDS epidemic. *Contributions to Economic Analysis & Policy*, 5(1), 12.
- Barnett, T., & Blaikie, P. (1992). *AIDS in Africa: Its present and future impact*. London: The Guilford Press.
- Bowman, A. W., Jones, M. C., & Gijbels, I. (1998). Testing Monotonicity of Regression. *Journal of Computational and Graphical Statistics*, 7(4), 489–500. doi:10.2307/1390678
- Chimbiri, A. M. (2007). The condom is an “intruder” in marriage: Evidence from rural Malawi. *Social Science & Medicine*, 64(5), 1102–1115.
- De Paula, A., Shapira, G., & Todd, P. (2009). How Beliefs about HIV Status Affect Risky Behaviors: Evidence from Malawi, Second Version.
- Delavande, A., & Kohler, H. P. (2009). Subjective expectations in the context of HIV/AIDS in Malawi. *Demographic Research*, 20(31), 817–875.
- Delavande, A., & Kohler, H. P. (2011). *HIV/AIDS-related Expectations and Risky Sexual Behavior in Malawi* (Working Paper).
- Dupas, P. (2011). Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya. *American Economic Journal: Applied Economics*, 3(1), 1–34.
- Fang, H., Keane, M., Khwaja, A., Salm, M., & Silverman, D. (2007). Testing the mechanisms of structural models: The case of the Mickey Mantle effect. *The American economic review*, 97(2), 53–59.
- Geoffard, P.-Y., & Philipson, T. (1996). Rational Epidemics and Their Public Control. *International Economic Review*, 37(3), 603–624. doi:10.2307/2527443
- Godlonton, S., & Thornton, R. (2011). Circumcision, Information, and HIV Prevention.
- Godlonton, Susan, Chinkhumba, J., Hayes, B., & Thornton, R. (2010). *Circumcision, Information, and HIV Prevention Study - Urban Sample Dataset*. University of Michigan.

- Godlonton, Susan, Thornton, R., & Munthali, A. (2009). *Circumcision, Information, and HIV Prevention Study - Rural Sample Dataset*. University of Michigan.
- Juhn, C., Kalemli-Ozcan, S., & Turan, B. (2009). *HIV and fertility in Africa: First evidence from population based surveys* (Working Paper No. Discussion Paper No. 4473). IZA Discussion Paper Series. Institute for the Study of Labor.
- Kaler, A. (2003). "My girlfriends could fill a yanu-yanu bus": Rural Malawian men's claims about their own serostatus.
- Leon, G., & Miguel, E. (2011). Transportation Choices, Fatalism and the Value of Life in Africa (Working Paper). Working Group in African Political Economy. Berkeley: University of California.
- Lillard, L., & Willis, R. (2001). *Cognition and wealth: The importance of probabilistic thinking* (Working Paper No. WP 2001-007) (p. 48109–1220). University of Michigan Retirement Research Center.
- Lind, J. T., & Mehlum, H. (2010). With or Without U? The Appropriate Test for a U-Shaped Relationship. *Oxford Bulletin of Economics and Statistics*, 72(1), 109–118.
- Luke, N., Clark, S., Zulu, E., Westlands, N., & Eliya, K. (2009). The Relationship History Calendar: Improving the Scope and Quality of Data on Youth Sexual Behavior. *Brown University typescript*.
- NAC. (2009). *National HIV/AIDS Prevention Strategy: 2009 to 2013*. Lilongwe: Malawi National AIDS Commission (NAC).
- O'Donoghue, T., & Rabin, M. (2001). Risky Behavior among Youths: Some Issues from Behavioral Economics. *NBER Chapters*, 29–68.
- Oster, E. (2012). HIV and sexual behavior change: Why not Africa? *Journal of Health Economics*, 31(1), 35–49. doi:10.1016/j.jhealeco.2011.12.006
- Padian, N. S., McCoy, S. I., Balkus, J. E., & Wasserheit, J. N. (2010). Weighing the gold in the gold standard: challenges in HIV prevention research. *AIDS*, 24(5), 621–635. doi:10.1097/QAD.0b013e328337798a
- Sterck, O. (2011). *Should prevention campaigns disclose the transmission rate of HIV/AIDS? Theory and evidence from Burundi*. Université catholique de Louvain, Institut de Recherches Economiques et Sociales (IRES). Retrieved from <http://ideas.repec.org/p/ctl/louvir/2010042.html>
- Stoneburner, R. L., & Low-Beer, D. (2004). Population-level HIV declines and behavioral risk avoidance in Uganda. *Science*, 304(5671), 714.
- Viscusi, W. K. (1990). Do smokers underestimate risks? *Journal of Political Economy*, 98(6), 1253–1269.

Wawer, M. J., Gray, R. H., Sewankambo, N. K., Serwadda, D., Li, X., Laeyendecker, O., Kiwanuka, N., et al. (2005). Rates of HIV-1 Transmission per Coital Act, by Stage of HIV-1 Infection, in Rakai, Uganda. *The Journal of Infectious Diseases*, 191, 1403–1409.

Yatchew, A. (1997). An elementary estimator of the partial linear model. *Economics Letters*, 57(2), 135–143. doi:10.1016/S0165-1765(97)00218-8

Table 1: Results of optimal risk-taking on per-act risk, simulated data

	(1)	(2)	(3)	(4)	(5)
Per-Act Risk	2.911*** (0.171)	2.573*** (0.648)	2.870*** (0.105)	2.150*** (0.432)	-16.49*** (0.906)
(Per-Act Risk)^2		0.425 (0.785)		0.905* (0.488)	43.84*** (1.954)
Cost per act			-50.07*** (1.012)	-50.10*** (1.012)	-47.72*** (1.561)
Constant	4.114*** (0.0734)	4.165*** (0.122)	11.63*** (0.169)	11.74*** (0.197)	12.71*** (0.289)
p-value(Lind-Mehlum U-Shape Test for n and X)	-	1.000	-	1.000	<0.001
Adjusted R-squared	0.0979	0.0978	0.789	0.789	0.797
n	3000	3000	3000	3000	1547

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Summary Statistics for Demographics

	Mean	Std. Dev.	Min.	Max.	N
Demographics					
Peri-Urban	0.05	0.22	0	1	447
Female	0.59	0.49	0	1	447
Married	0.88	0.32	0	1	447
Age	29.46	8.25	17	49	447
Years of education	5.52	3.03	0	13	444
Literate	0.60	0.49	0	1	447
# of female children	1.47	1.31	0	7	438
# of male children	1.55	1.21	0	6	445
Desired # of additional children	1.16	1.45	0	20	447
Attractiveness [1-5] [*]	3.44	0.90	1	5	442
Household spending in past month, PPP USD ^{**}	142	232	0	3852	447
Household income in past month, PPP USD ^{**}	247	561	0	19369	447
Religion					
Catholic	0.23	0.42	0	1	447
CCAP	0.13	0.34	0	1	447
Pentecostal	0.13	0.33	0	1	447
Church of Christ	0.11	0.31	0	1	447
Anglican	0.06	0.24	0	1	447
Other Christian [†]	0.14	0.35	0	1	447
Muslim	0.20	0.40	0	1	447
Ethnic Group					
Lomwe	0.55	0.50	0	1	440
Yao	0.22	0.41	0	1	440
Chewa	0.16	0.37	0	1	440
Other [‡]	0.07	0.26	0	1	440

Data taken from Kerwin et al. (2011). Means and standard deviations constructed using sample weights. Summary statistics for other regression controls available from the author upon request.

* Enumerators (of the same sex as the respondents) rated respondents on how attractive they are, from 1 to 5.

** Constructed by dividing self-reported values in Kwacha by the World Bank ICP PPP exchange rate for 2011, which is 39.46.

† Other Christian includes indigenous Christian churches (2.0%), Baptist (1.6%), 7th-Day Adventist (0.2%), and miscellaneous (10.3%).

‡ Other includes Mang'angja (2.7%), Nyanja (2.6%), Ngoni (1.3%), Sena (0.4%), and trace numbers of Tumbuka and Shona.

Table 3: Summary Statistics for Subjective Beliefs and Sexual Behaviors

	Mean	Std. Dev.	Min.	Max.	N
Subjective Beliefs					
Transmission Risks					
Share of people resp. finds attractive who are HIV-positive	0.49	0.21	0	1	414
Prob. of HIV Xmission per unprotected sex act w/infected partner)	0.85	0.26	0.01	1	423
Average Prob. of HIV Xmission per act, attractive people	0.43	0.24	0	1	413
Cost of HIV Infection					
Years of life remaining	29.96	17.48	0	82	436
Years from HIV infection to death, without ARVs	4.36	4.09	1	80	423
Years from HIV infection to death, with ARVs	15.85	12.19	2	100	423
Share of HIV-positive people who would receive ARVs	0.62	0.29	0	1	421
Expected years from HIV infection to death	11.43	8.50	1.12	100	421
Expected life years lost if resp. gets HIV	18.86	16.40	0	78	410
Sexual Behaviors					
Total sex acts in past week	2.11	2.68	0	13	447
Unprotected sex acts in past week	1.93	2.67	0	13	447

Data taken from Kerwin et al. (2011). Means and standard deviations constructed using sample weights.

Table 4: Regressions of risk-taking on perceived per-act risks, full sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Average Prob. of HIV Xmission per act, attractive people	2.833**	-1.046	1.520**	-3.732	1.524	-5.800**	-4.277
	(1.304)	(3.508)	(0.759)	(2.588)	(1.133)	(2.494)	(3.274)
(Average Prob. of HIV Xmission per act, attractive people)^2		4.185		5.598**		7.442***	5.818
		(4.222)		(2.638)		(2.592)	(3.577)
Expected life years lost if resp. gets HIV					-0.00804	-0.0156	-0.00419
					(0.0221)	(0.0203)	(0.0260)
(Cost of HIV)*(Prob. HIV per Act)					-0.0267	-0.0129	-0.0841
					(0.0415)	(0.0374)	(0.127)
(Cost of HIV)*(Prob. HIV per Act)^2							0.0793
							(0.128)
Constant	0.753	1.401***	-1.850	-0.551	-1.054	0.506	0.349
	(0.547)	(0.453)	(3.255)	(3.337)	(3.896)	(3.912)	(3.942)
Controls Used [†]			X	X	X	X	X
Std. Errors Clustered by Area	X	X					
p-value(Lind-Mehlum U-Shape Test for n and X)	-	0.385	-	0.075	-	0.010	0.096
Adjusted R ²	0.0633	0.0780	0.524	0.535	0.528	0.546	0.545
N	413	413	380	380	367	367	367

Heteroskedasticity-robust standard errors in parentheses. All regressions run using the sampling weights as pweights.

[†] Controls include gender, age, age squared, education, education squared, number of male and female children, desired future children, marital status, literacy, attractiveness, attractiveness squared, and peri-urban location, and fixed effects for area, ethnicity, religion, media exposure (TV/radio/newspaper) and enumerator. Coefficients on controls suppressed due to space considerations but available from author upon request.

*** p<0.01, ** p<0.05, * p<0.1

Table 5: Regressions of risk-taking on perceived per-act risks, excluding “uncertain” respondents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Average Prob. of HIV Xmission per act, attractive people	2.685*	-3.134	0.910	-6.182**	1.673*	-7.589***	-6.788*
	(1.355)	(3.888)	(0.588)	(2.709)	(0.887)	(2.723)	(3.450)
(Average Prob. of HIV Xmission per act, attractive people)^2		6.126		7.516***		9.511***	8.672**
		(4.524)		(2.846)		(2.758)	(3.610)
Expected life years lost if resp. gets HIV					0.0109	0.00537	0.0107
					(0.0157)	(0.0134)	(0.0197)
(Cost of HIV)*(Prob. HIV per Act)					-0.0175	-0.00118	-0.0369
					(0.0336)	(0.0276)	(0.110)
(Cost of HIV)*(Prob. HIV per Act)^2							0.0385
							(0.112)
Constant	0.698	1.596***	-6.952*	-4.511	-7.416*	-4.654	-4.921
	(0.544)	(0.456)	(3.560)	(3.341)	(3.881)	(3.459)	(3.568)
Controls Used [†]			X	X	X	X	X
Std. Errors Clustered by Area	X	X					
p-value(Lind-Mehlum U-Shape Test for n and X)	-	0.217	-	0.0117	-	0.0029	0.0252
Adjusted R ²	0.0783	0.119	0.661	0.681	0.676	0.706	0.705
N	312	312	293	293	281	281	281

Heteroskedasticity-robust standard errors in parentheses. All regressions run using the sampling weights as pweights.

[†] Controls include gender, age, age squared, education, education squared, number of male and female children, desired future children, marital status, literacy, attractiveness, attractiveness squared, and peri-urban location, and fixed effects for area, ethnicity, religion, media exposure (TV/radio/newspaper) and enumerator. Coefficients on controls suppressed due to space considerations but available from author upon request.

*** p<0.01, ** p<0.05, * p<0.1

Figure 1: Shapes of Risk-Aggregation Functions for Low and High Values of Per-Act Risk

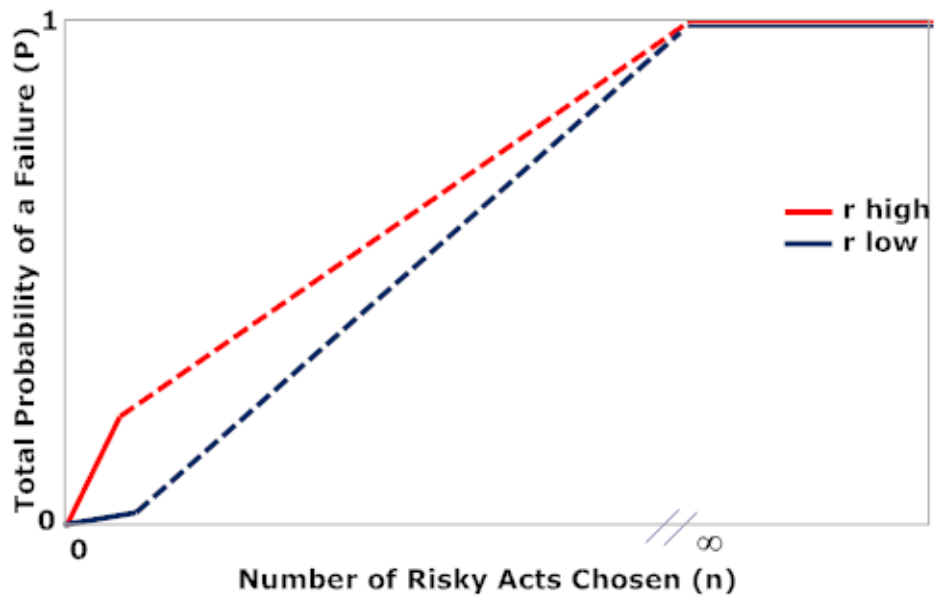


Figure 2: Scatterplot of Optimal Risk-Taking by Per-Act Risk, Simulated Data

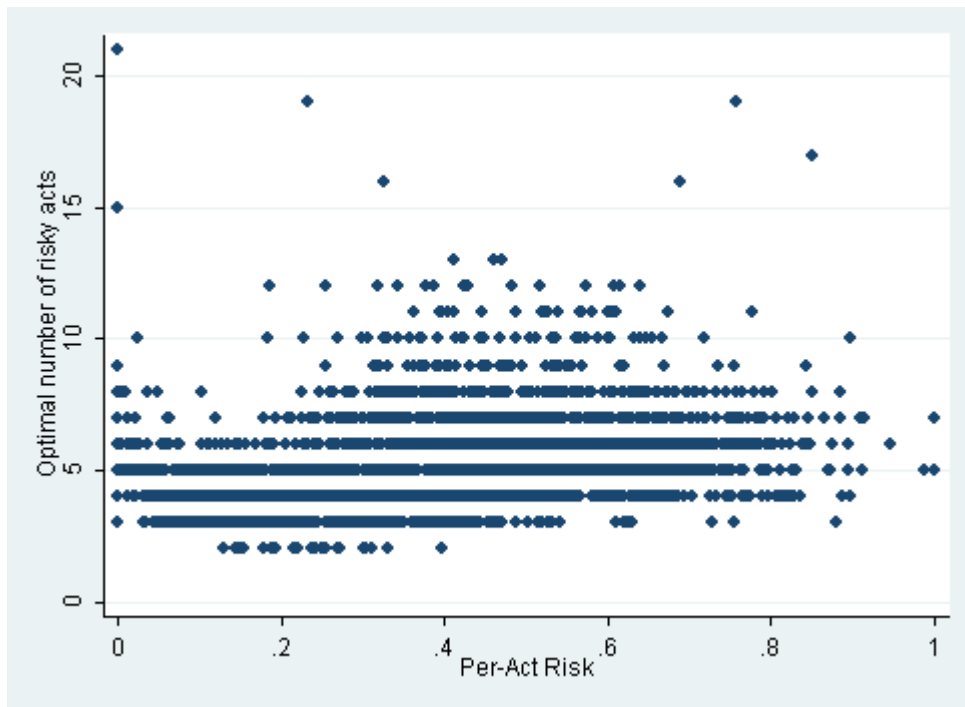


Figure 3: Partially Linear Regression of Optimal Risk-Taking on Per-Act Risk, Simulated Data

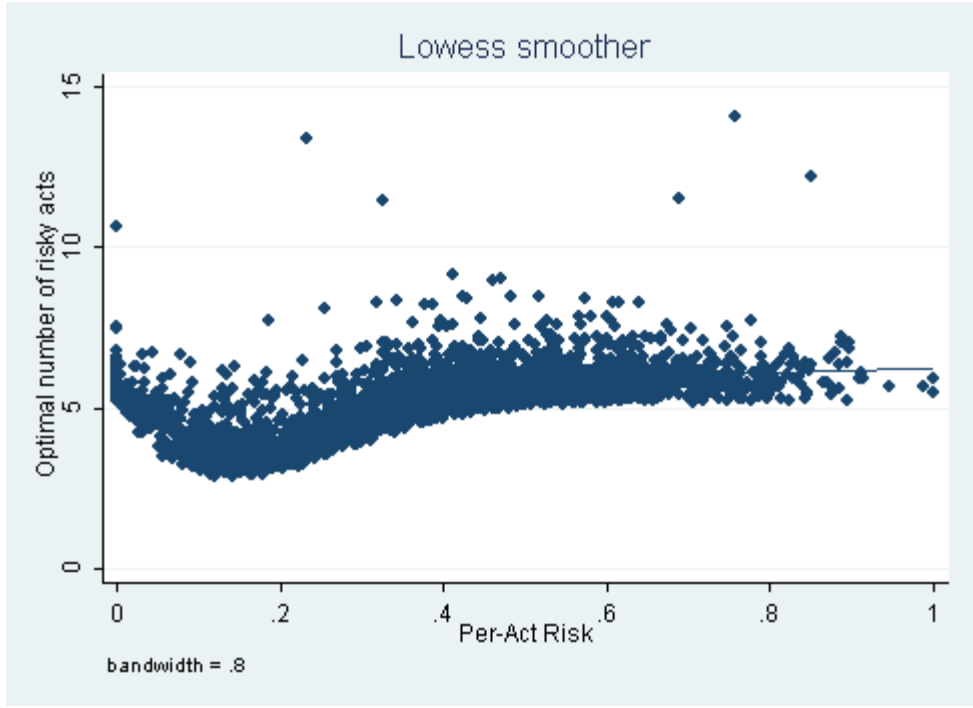


Figure 4: Scatterplot of Unprotected Sex by Perceived HIV Transmission Risk from Attractive People, Unweighted

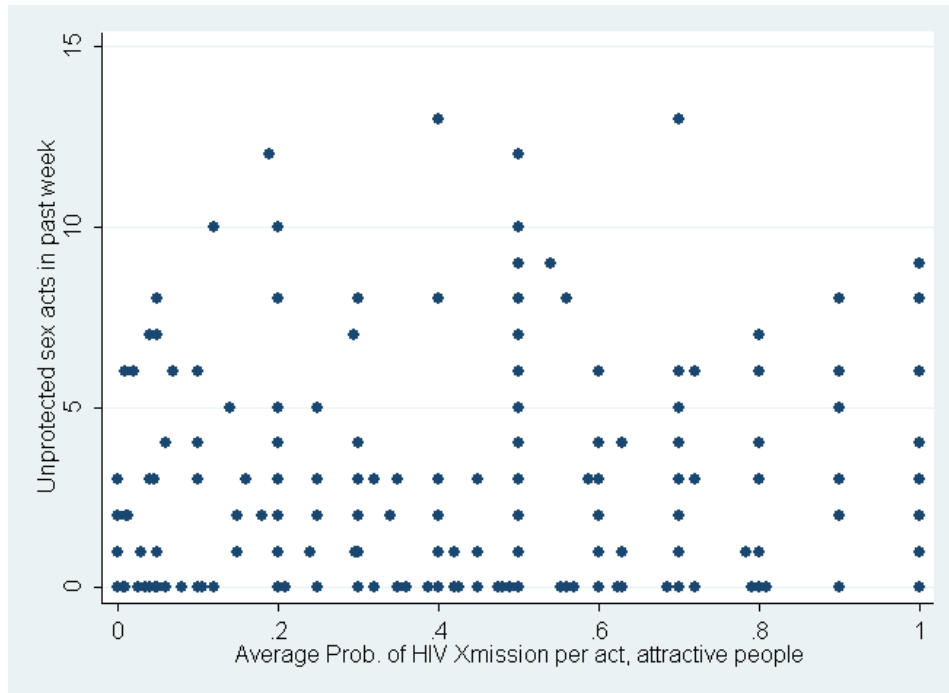


Figure 5: Partially Linear Regression of Unprotected Sex on Perceived HIV Transmission Risk from Attractive People, All Respondents, Unweighted




Figure 6: Partially Linear Regression of Unprotected Sex on Perceived HIV Transmission Risk from Attractive People, Peri-Urban Respondents, Unweighted

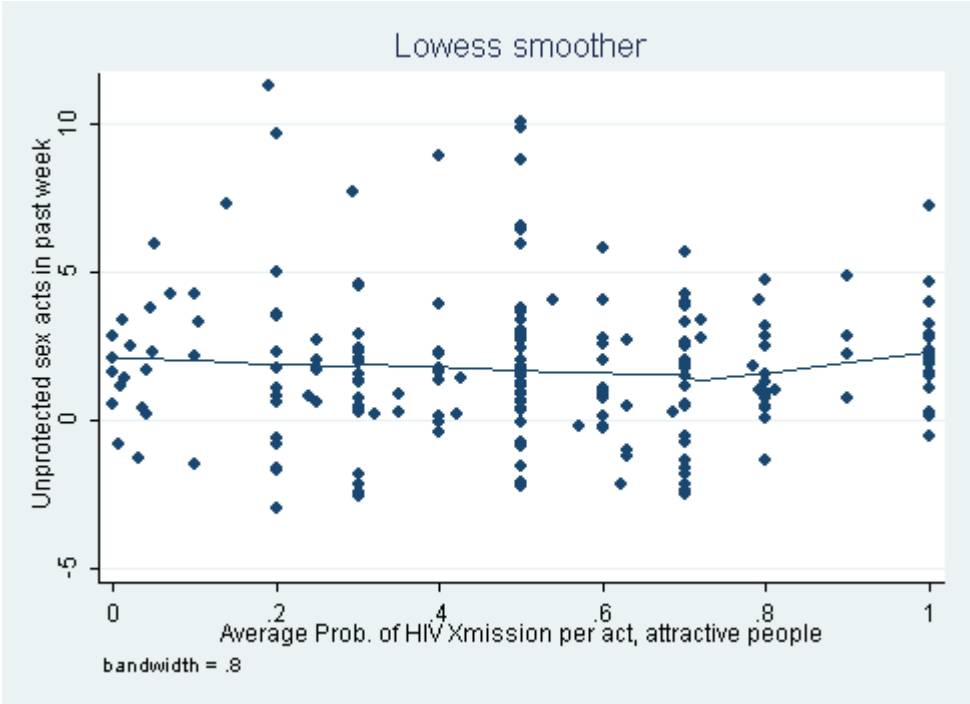


Figure 7: Partially Linear Regression of Unprotected Sex on Perceived HIV Transmission Risk from Attractive People, Rural Respondents, Unweighted




Figure 8: Histogram of Perceived HIV Transmission Risk from Attractive People

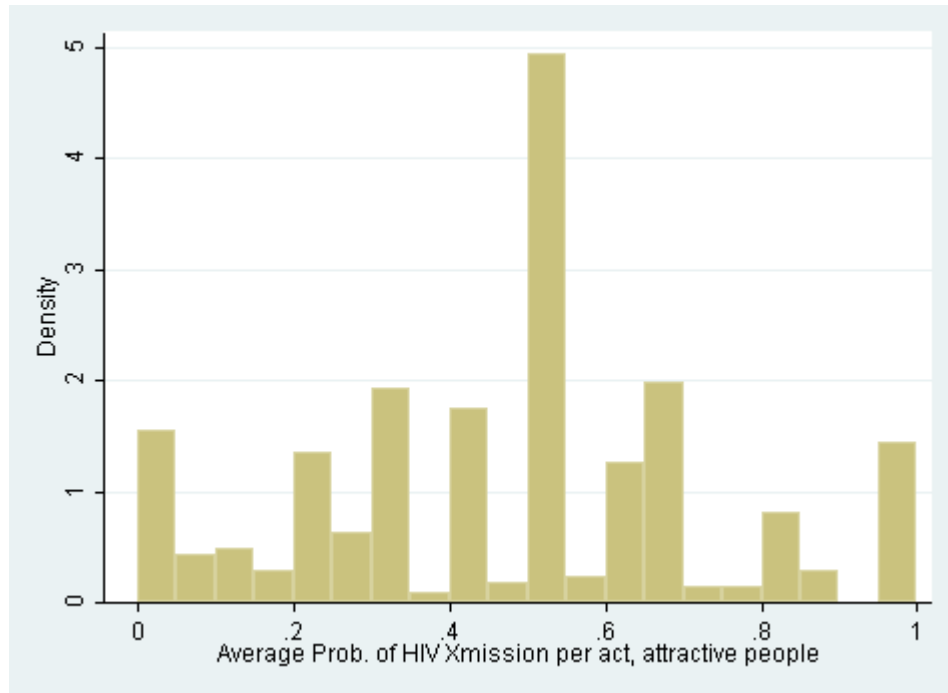


Figure 9: Partially Linear Regression of Unprotected Sex on Perceived HIV Transmission Risk from Attractive People, Excluding “Uncertain” Individuals, Unweighted

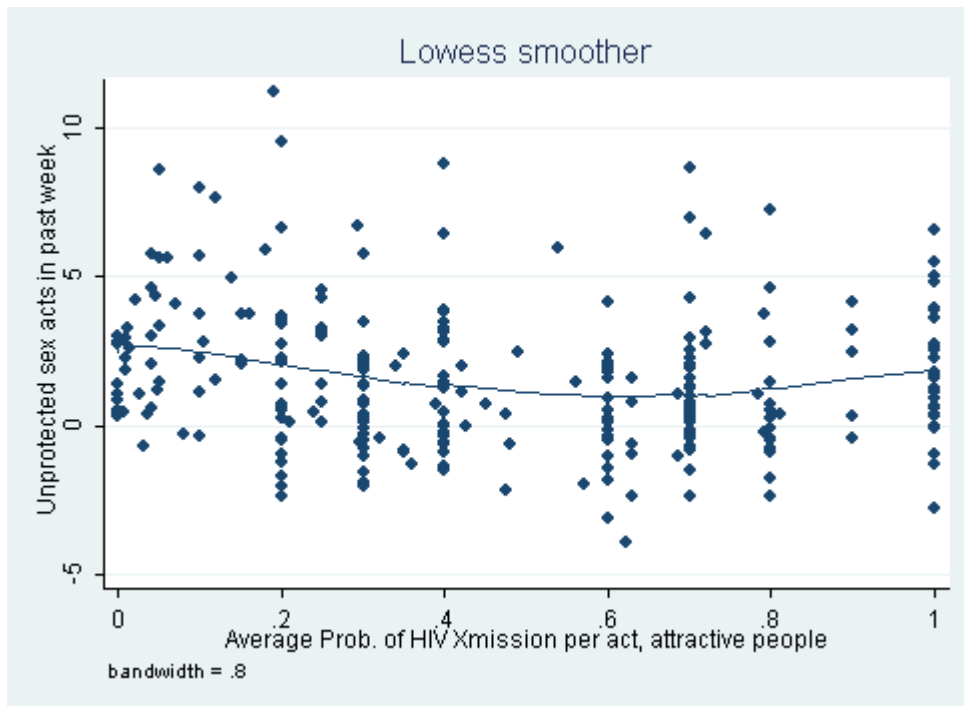
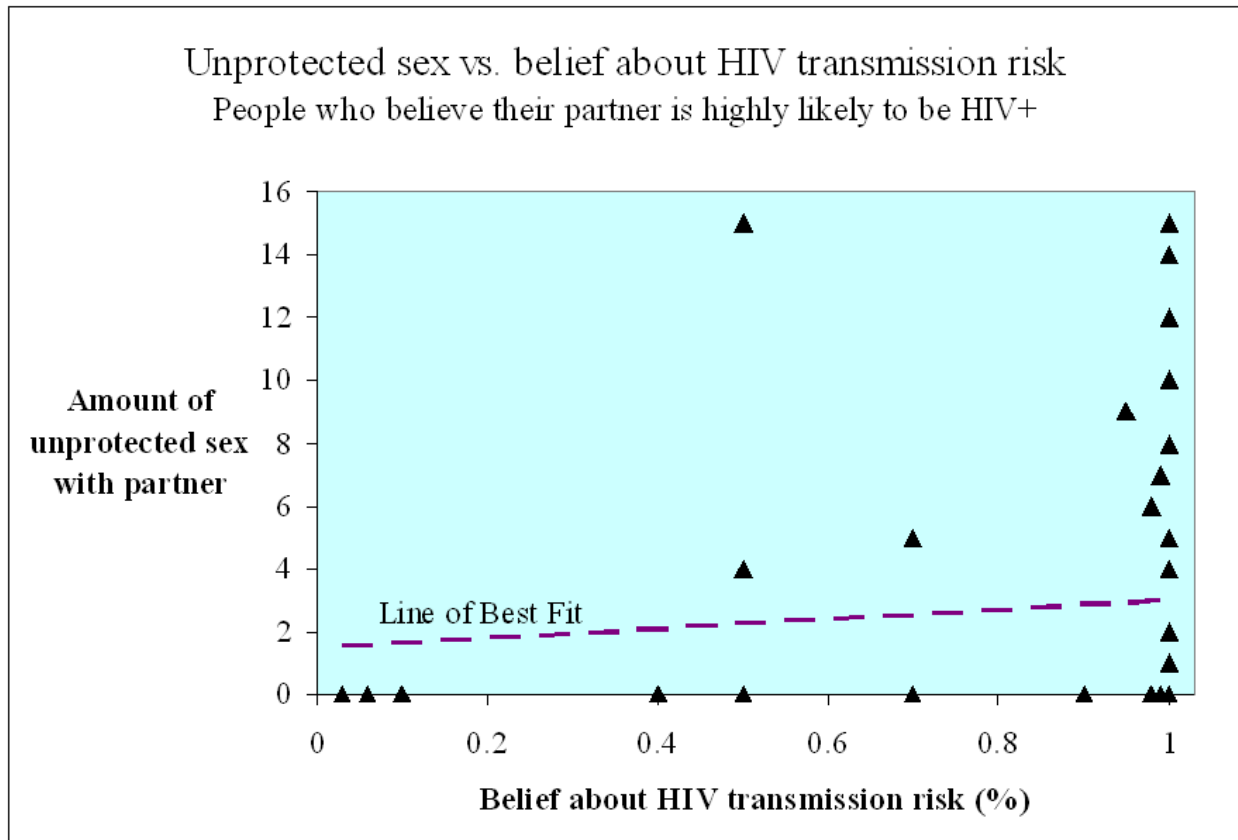


Figure 10: Risk-Taking with Primary Partner for Respondents who Think Partner is Probably HIV+



Appendix A: Additional Comparative Statics

This appendix derives two additional comparative statics based on similar methods to those used in Section 3.3: the response of the optimal value of n to the per-act price p and the cost of the bad outcome c . These two comparative statics are both negative in sign. The second is directly analogous to the Oster (2012) result for a linear risk aggregation function P .

$$\text{III. } \frac{\partial n^*}{\partial p} = -\frac{\frac{\partial G}{\partial p}}{\frac{\partial G}{\partial n^*}} = \frac{1}{v''(n^*) - P_{22}(r, n^* + m) * c}$$
$$\text{IV. } \frac{\partial n^*}{\partial c} = -\frac{\frac{\partial G}{\partial c}}{\frac{\partial G}{\partial n^*}} = \frac{P_2(r, n^* + m) * c}{v''(n^*) - P_{22}(r, n^* + m) * c}$$

Note that the denominator of each expression is simply the second-order condition for an internal optimum and is therefore strictly negative, and that P_2 is negative by assumption.. This immediately yields the following intuitive results:

Proposition B1

$$\frac{\partial n^*}{\partial p} < 0 \text{ (Comparative Static I)}$$

The higher the price of a risky act, the less of it people will do. Conversely, the more negative the price - that is, the more that one is paid to take risks - the more one is willing to do.

Proposition B2

$$\frac{\partial n^*}{\partial c} \leq 0 \text{ (Comparative Static II)}$$

Raising the cost of a failure weakly decreases the number of risks taken, and strictly decreases the number of risks chosen as long as $P_2(r, n + m) \neq 0$.