Does the Risk of Poverty Reduce Happiness?*

Stefano A. Caria† and Paolo Falco‡

Nov, 2017

Abstract

We investigate the unexplored link between the risk of poverty and happiness in the context of a developing country. Using unique longitudinal data, we estimate workers' vulnerability to income poverty and find a strong negative relationship between vulnerability and life-satisfaction, over and above the positive income effect commonly documented in the literature. The result is robust and cannot be reduced to the effect of two-sided uncertainty. A matched behavioural experiment shows that respondents are significantly loss-averse. We conclude that downside risk is an important determinant of happiness and of economic decisions under uncertainty. Policies that mitigate downward risk may thus have direct impacts on both well-being and efficiency.

JEL codes: D60, I31, I32, D81, O12.
Keywords: poverty, low pay, vulnerability, risk, subjective well-being, happiness, loss-aversion.

*This paper uses data from the Ghana Urban Household Panel Survey (GUHPS), conducted by the Centre for the Study of African Economies as part of ongoing research into African labour markets funded by the ESRC, RECOUP, IDRC, DFID and the Gates Foundation. We are grateful to Marcel Fafchamps, William Maloney and Francis Teal for helpful comments, and to the participants to the CSAE Conference 2011, the IZA-WB Conference 2011 and the RES Conference 2013. We are indebted to Moses Awoonor-Williams who coordinated the data collection and to the enumerators of the CSAE team in Ghana. The usual disclaimer applies.

†Oxford Department of International Development, University of Oxford (stefano.caria@qeh.ox.ac.uk)

‡OECD and University of Oxford, 2 rue André Pascal, 75016, Paris (paolo.falco@oecd.org)

1 Introduction

Economic outcomes are often characterised by pervasive uncertainty. This is particularly true when insurance markets and safety nets are incomplete, saving opportunities limited, and many individuals rely on risky entrepreneurial activities to generate their incomes (Banerjee and Duflo, 2007). Poverty, which constitutes a possible outcome for many, has in turn profound impacts on the quality of the lives people live. Recent studies have shown that low income correlates with lower life satisfaction and with a larger loss in well-being following shocks in other domains of life (Clark et al., 2008; Kahneman and Deaton, 2010). In developing countries, where poverty is widespread, the correlation between economic outcomes and life satisfaction is even stronger (Howell and Howell, 2008).

In this study we investigate the relationship between the risk of income poverty and life satisfaction (interpreted henceforth as 'happiness'[1]), and the link between sensitivity to downward risks and decision-making. In particular, we tackle the following two questions. Is there a connection between happiness and the risk of poverty? And how are people's decisions affected by exposure to such risk?[2] While the connection between life satisfaction and low income has been heavily researched, the one between life satisfaction and the risk of income poverty is still unexplored. This is partly due to the challenges of estimating the probability distribution of income convincingly. At the same time, it appears to be a very important area of research, especially in developing countries, where widespread exposure to uninsured shocks makes the risk of future income poverty pervasive for both poor and

[1] Researchers distinguish two components of happiness (Kahneman and Deaton (2010)). The first component is life satisfaction: the evaluation we make of our own life. The second component is emotional well-being, or the tendency to experience positive or negative affect. In this paper we analyse responses from a survey question on life satisfaction and hence we focus the analysis on the first component. Throughout the text, we will use the terms 'happiness' and 'life satisfaction' interchangeably.

[2] We choose to focus on income, rather than consumption for two main reasons. First, we aim to link directly to the existing literature on subjective well-being, which has widely explored the relationship between happiness and income. Second, in urban contexts like the one we are studying, labour earnings are typically the main source of income and earning shocks are directly transmitted to consumption. Changes in income and consumption are hence likely to be tightly correlated. We fully acknowledge the difficulties that arise from attempting to measure income precisely in a low-income setting with widespread informality, such as Ghana, especially when self-employment is prominent. Section 2 provides a detailed discussion of the methodological challenges involved.

non poor households.[3] Evidence is also missing on the connection between the determinants of happiness and those of decision making: are the same individuals whose happiness is sensitive to downside risk loss-averse in economic decisions? Vulnerability may affect individual behavior in ways that are detrimental to economic efficiency.[4] Such evidence is thus a necessary first step towards a full assessment of the welfare effects of economic vulnerability.

The context of our analysis is the urban labour market in Ghana, a growing African country. Ghana is an interesting setting for our analysis as the country experienced substantial poverty reduction in recent years (Nsowah-Nuamah et al., 2010) while, as our results will suggest, large numbers are still exposed to a significant risk of poverty. Given the novelty of the question and of the testing strategy, our results provide leads that may prove relevant in other contexts as well.

Our estimate of the risk of poverty builds upon the work by Chaudhuri (2003) and Chaudhuri et al. (2002), who propose two indices of vulnerability to poverty that are amenable to empirical estimation based on panel and cross-sectional variation respectively.[5] Using data from the Ghana Urban Household Panel Survey (GUHPS), a long panel dataset gathered by the Centre for the Study of African Economies in urban Ghana, we obtain estimates of the two indices for a representative sample of working age Ghanaian earners. We focus more extensively on the panel measure, since it enables us to estimate individual-specific vulnerability. We further rely on the longitudinal nature of the data to investigate the relationships of interest, between the risk of income poverty and life-satisfaction. Improving upon most of the existing literature on happiness in developing countries, we are able to control for *individual fixed effects in the happiness model*, ruling out potential biases from unobserved personality traits. Previous studies have indeed highlighted the importance of unobserved heterogeneity in happiness regressions (Ferrer-i

[3] For example, in a recent study of seven West-African capitals, Bocquier et al. (2010) construct a multidimensional index of employment vulnerability and find that 85% of private sector workers are vulnerable on the basis of at least one criterion in 2002-2003.

[4] Throughout the analysis the term 'vulnerability' will be used to refer to 'the risk of falling below the income poverty line'.

[5] The two indices are reviewed in a survey article by Ligon and Schechter (2004), who compare the performance of different vulnerability measures through Monte Carlo simulations.

Carbonell and Frijters, 2004; Graham et al., 2004; Powdthavee, 2010).

Our main result is *a strong negative relationship between vulnerability to income poverty and workers' happiness*, over and above the positive income effect commonly documented in the literature. It is both statistically significant and economically meaningful. Reducing the risk of poverty by 20 percentage points (which amounts to entirely offsetting the risk of poverty for the median worker in our sample) has the same effect on well-being as increasing earnings by 50%. When we bootstrap the estimation sequence to account for imprecision in the measure of vulnerability, the results do not change. Upon testing for the role of two-sided uncertainty as opposed to downward income losses, we find that the effect of downward vulnerability on happiness is more evident. These findings become more compelling when we consider the extent of the vulnerability to poverty we uncover. About 35 percent of all workers, and 15 percent of currently non-poor workers, face a risk of poverty of at least 50 percent. Vulnerability decreases the life satisfaction of a large pool of individuals.

In addition, we analyze the choices of a sub-sample of respondents in a set of behavioral games designed to elicit attitudes towards risky prospects. Our maximum likelihood estimates reveal that subjects are characterised on average by a substantial degree of loss-aversion. We are careful not to collapse the distinct notions of experienced utility and decision utility (Kahneman et al., 1997).[6] Our findings from the behavioral experiment show that, besides influencing subjective well-being, downside risk also has an appreaciable impact on economic decisions.

Our work relates to two different strands of the literature. First, we contribute to the study of downside risk in developing countries. This literature has focused on measurement (Chaudhuri, 2003; Ligon and Schechter, 2004), on the persistence of downside shocks (Dercon, 2004; Dercon et al., 2005), and on the strategies employed to minimize and cope with shocks (Rosenzweig and Binswanger, 2003; Dercon, 1996; Fafchamps, 2003; Dercon

[6] Decision utility refers to the weights that people assign to outcomes when making choices. Experienced utility refers to the quality of an experience. The life satisfaction question we employ in this analysis asks respondents to make an evaluation of the latter. The behavioral experiments allow us to make inferences about the former.

and Christiaensen, 2007; Fafchamps, 2009). Most importantly, our analysis contributes to the growing literature on the determinants of happiness. A number of empirical papers have documented a cross-sectional correlation between income and happiness (Kahneman and Deaton, 2010), which does not disappear once individual fixed effects are accounted for (Ferrer-i Carbonell and Frijters, 2004; Powdthavee, 2010). A separate concern has been that of adaptation: the happiness effects of income gains seem transitory and tend to disappear once income reference points have adjusted (Easterlin, 2001; Frey and Stutzer, 2002; Di Tella et al., 2007; Knight and Gunatilaka, 2008). The literature has also explored the effect of social comparisons on well-being (Blanchflower and Oswald, 2002; Kingdon and Knight, 2004; Luttmer, 2005). Our contribution is to highlight the fact that risk and, in particular, the risk of poverty is a major negative determinant of life satisfaction. Morevoer, we show that the same people who manifest loss-sensitivity in life evaluation make economic decisions consistent with loss-aversion.

The results of this analysis bear important policy implications that may generalise well-beyond the African context. Our findings provide clear motivation for policy interventions to reduce people's exposure to (downside) risk. They also suggest that non-Rawlsian models of growth, whereby 'someone may be left behind', may fail to enhance general welfare despite rising average incomes, if the risk of falling behind is sufficiently widespread. Lastly, loss-aversion motivates individuals to forgo economic opportunities that are profitable in expectation but may involve outcomes below the reference point. A reduction in vulnerability may result in efficiency gains too.

The article is structured as follows. Section 2 introduces the data we use in the analysis. Section 3 outlines the empirical strategy. First, it explains the methodology to estimate income vulnerability; second, it outlines the happiness model. Section 4 presents and discusses the results. Section 5 concludes.

2 Data

Our analysis is based on data from the Ghana Urban Household Panel Survey (GUHPS), which was conducted by the Centre for the Study of African Economies (CSAE) in the cities of Accra, Kumasi, Takoradi and Cape-Coast starting in 2004.[7] Respondents were drawn by stratified random sampling of urban households from the Population and Housing Census of 2000. The survey was designed to cover all household members of working age at the time of the interview. This paper focuses on the period between 2004 and 2009, when the survey was repeated every year with the sole exception of 2007.[8] Panel datasets of this length are unusual in developing countries, and are particularly uncommon in Africa.[9]

A module on subjective well-being – designed in accordance with the existing literature – was added to the survey in 2005. Our analysis will focus on the answers to the following two questions: (a) *'All things considered, how satisfied are you with your life as a whole these days?'* (b) *'All things considered, how satisfied are you with your current work?'*. In both cases, the options given to respondents were: *'1.Very Dissatisfied, 2. Dissatisfied, 3. Neither Satisfied Nor Dissatisfied, 4. Satisfied, 5. Very Satisfied'*. Figure 1 depicts the distribution of answers. Responses appear to be skewed towards positive values. For our quantitative analysis, we attribute numerical values on a scale from 1 to 5 to these answers, where 1 corresponds to 'Very Dissatisfied' and 5 to 'Very Satisfied'. Despite early criticism of their ability to accurately capture well-being (e.g. Mullainathan and Bertrand (2001)), these measures have been consistently used throughout the literature. Moreover, psychologists have recently been able to validate the use of these questions, by showing their correlation with other measures of well-being, such as smiling more frequently (Gra-

[7] The first wave was collected between the end of 2003 and the beginning of 2004, but for simplicity we will refer to it as 2004.

[8] Subsequent survey waves were unavailable to us at the time of writing. Moreover, extending the analysis beyond 2009 would pose severe challenges due to a change in the survey question on life satisfaction introduced in 2010.

[9] The panel is unbalanced, but attrition is not an absorbing state, in the sense that respondents who are not interviewed in a given wave are kept in the sample and re-interviewed in subsequent years. Out of the initial sample of respondents interviewed in 2004, over 92% were re-interviewed at least once more in the following years and about 65% are observed in at least 3 waves. Random attrition would decrease precision and pose a classical problem of attenuation bias in the happiness model. In this case, the large effects we estimate would be a lower bound of the true impacts. On the other hand, if the people whose happiness is least affected by the risk of poverty are most likely to drop out, our estimates may be biased upwards. As a robustness check, the analysis presented below has been repeated on the strictly balanced sample of individuals who are interviewed in all waves and we find that this does not affect the results. The point-estimates of the coefficients discussed below do not change substantially, despite a drop in precision due to the fall in sample-size. Previous studies on this data have found no evidence of selective attrition biasing the results (see Falco et al. (2011) and Falco et al. (2015)).

ham et al., 2004; Layard, 2005; Oswald and Wu, 2010).

< Figure 1 here >

A selection of key summary statistics for our sample of interest is presented in Table 1. The average worker is 35.3 years old and has 8 years of formal education. Most of them are self-employed (largely in the informal sector), as it is typical in the labour markets of many developing countries.[10] For the self-employed income is calculated as revenue minus cost (i.e. profit).[11] In low-income settings, such as urban Ghana, it is difficult to obtain precise estimates of such profit and we are aware of the challenge. A number of published studies, however, have successfully relied on profit information from GUHPS (see Falco and Haywood (2016) and Falco et al. (2015) for the most recent examples). Moreover, a study by Fafchamps et al. (2012) lends additional support to the data. They use Personal Digital Assistant (PDA) technology to cross-check profit figures in a sample of micro-entrepreneurs in urban Ghana which is very similar to the GUHPS sample. They find that only a small proportion of profit-calculations is imprecise.[12]

< Table 1 here >

[10] The sample described in Table 1 is confined to workers for whom self-reported subjective well-being data is available. Furthermore, it only includes paid workers, for whom income is observed and we are therefore able to construct a measure of income vulnerability.

[11] The resulting profits are attributed to the owner of the business. The remaining household members who work in the business are categorized as paid or unpaid family labour, depending on whether they receive an income. This may lead to some imprecision as it may be difficult to clearly identify business ownership in some cases. However, the fact that the large majority of the self-employed in our sample (about 80%) does not employ any other person (including household members, either paid or unpaid) makes the potential magnitude of the problem relatively small. When a self-employed works by himself/herself (i.e. hires no other workers), the value of the variable 'employees' is set to 1. Consequently, its log in the regressions below equals 0. The same approach was followed by Falco et al. (2011).

[12] There might be additional reasons of concern beyond the precision of profit calculations. For instance, in the absence of accurate bookkeeping, micro-entrepreneurs may simply recall the 'wrong' revenue and cost figures, which might equally affect standard (paper-based) and PDA-based surveys if the latter are unable to spot inconsistencies across revenue, costs and profit figures. Most micro-enterprise surveys suffer from such problems and the fact that GUHPS data has already been successfully employed in several published studies lends credibility to our analysis.

3 Empirical methodology

3.1 Constructing an indicator of vulnerability to income poverty

This section outlines the methodology to construct the vulnerability indicators used in the remainder of the analysis. For a detailed discussion of the relative merits of different vulnerability indices, the reader is referred to the survey paper by Ligon and Schechter (2004). The analysis in this article will draw on the two measures proposed by Chaudhuri (2003) and Chaudhuri et al. (2002). The former relies on time-series variation in individual earnings and suits particularly well the characteristics of our dataset, where subjective well-being is recorded for the same individuals over a number of consecutive years (in addition to income and other worker characteristics from which we can model vulnerability). The latter method attempts to model cross-sectional variation in earnings and to infer from it the degree of individual vulnerability. Ligon and Schechter (2004) compare the performance of these two (and several other) vulnerability indices via Monte Carlo simulations and their conclusion is in favour of the panel approach as the best performing indicator of actual vulnerability.[13] In light of their findings and of the data at our disposal, our main focus will be on the panel method. The online appendix presents alternative estimates based on the cross-sectional approach and offers insights on the differences between the two.

Following Chaudhuri (2003), the income vulnerability of a worker at time t is defined as the probability that the worker's income will fall below a certain threshold (z) next period. This differs from a simple measure of volatility. The distribution from which a given worker's income is drawn can have very large volatility, but if most of the variation occurs above the threshold z, this worker has a low probability of earning less than z and hence has low vulnerability.

Let ν_{it} be the inverse of vulnerability, that is, i's probability at t of earning an income

[13] This is a sensible conclusion, considering the likely presence of unobserved individual fixed effects that cannot be controlled for in a cross sectional model of earnings and could therefore mislead the analysis of vulnerability. An interesting alley for future research would be to explore how the results of this analysis would change upon using other vulnerability measures, including Ligon and Shechter's own index of vulnerability (see Ligon and Schechter (2003)).

above z at $t+1$:

$$\nu_{i,t} = Pr(y_{i,t+1} > z) \tag{1}$$

When z is the poverty line, ν_{it} is the probability at t that household i will not be income poor at $t+1$. Following standard Mincerian earnings analysis, we assume that income is generated by the following process:

$$ln(y_{i,t}) = \delta X_{i,t} + \eta_i + \tau_t + e_{i,t} \tag{2}$$

where X_{it} is a bundle of observable characteristics, η_i is an individual unobservable fixed effect, τ_t captures time-effects that are common across workers (e.g. aggregate income growth factors and common shocks) and e_{it} is a stochastic component.

Second, we assume the variance of e_{it} to be a function of worker and household characteristics.

$$ln(\sigma^2_{lny_{i,t}}) = \theta K_{i,t} + \xi_i \tag{3}$$

where ξ_i is an individual fixed-effect in the model of income variance and $K_{i,t}$ may or may not contain additional worker characteristics, outside the set $X_{i,t}$. In Chaudhuri (2003) $K_{i,t}$ and $X_{i,t}$ coincide.

The variance of the stochastic component can be modeled empirically using the log of first-stage residuals from the earnings model:

$$ln(\hat{e}^2_{i,t}) = \theta K_{i,t} + \xi_i + \omega_{i,t} \tag{4}$$

given that:

$$\frac{1}{T}\sum_{t=1}^{T} \hat{e}^2_{i,t} \rightarrow_p \sigma^2_{lny_{i,t}} \tag{5}$$

Assuming income to be (log)normally distributed and Φ to be the cumulative distribution function of the log normal distribution, we can now compute the probability of not earning less than z at $t+1$ for every worker i using the following expression.[14]

$$\hat{\nu}_{it} = \hat{Pr}(ln(y_{it}) > ln(z)|X_{it,}K_{it}, \hat{\delta}, \hat{\theta}) = 1 - \Phi\left[\frac{ln(z) - \hat{\mu}_{it}}{\hat{\sigma}_{it}}\right] \tag{6}$$

Where $\hat{\mu}_{it}$ denotes the predicted value of (log) income, and $\hat{\sigma}_{it}$ the predicted variance.

This measure does not differentiate transitory from permanent shocks. When saving, credit and insurance options are limited, both transitory and permanent shocks can have a significant impact on consumption. However, in the presence of an effective saving technology, the two may have different impacts on welfare since transitory fluctuations can be smoothed out through precautionary savings. We acknowledge that separating transitory and permanent shocks would be useful, especially in the largest urban areas where financial markets are likely to be more accessible. Unfortunately, attempting to explicitly separate the permanent from the transitory component of income variation (e.g. following the approach by Meghir and Pistaferri (2004)) would pose major challenges given the length of the panel at our disposal. This remains an open alley for future research.

3.2 Empirical model of happiness

Having constructed a measure of income vulnerability, we can now explore its relationship with subjective well-being. The following equation describes our workhorse model of happiness:

$$h_{i,t} = \beta y_{i,t} + \gamma \nu_{i,t} + \delta Z_{i,t} + \kappa_i + \epsilon_{i,t} \tag{7}$$

where h_{it} is worker i's level of life satisfaction in period t, y_{it} is income at time t and ν_{it} is the inverse of vulnerability in the same period; Z_{it} is a vector of worker characteristics that are expected to be correlated with life-satisfaction. κ_i is an unobserved happiness fixed-effect that accounts for unobserved traits that make an individual naturally more (or less) prone to be satisfied with his/her life (e.g. optimism). Our main hypothesis is that

[14] The reader should note that, differently from the definition in (1), our estimates of vulnerability will be obtained as the probability of falling below the poverty line given worker characteristics at t, rather than $t+1$. This choice was made based on the idea that workers are most likely to assess their future prospects on the basis of their current characteristics, some of which might themselves be stochastic and subject to unpredictability.

β and γ are positive (once again, note that $\nu_{i,t}$ is the inverse of vulnerability and, hence a 'good' in this specification): *increasing income and decreasing vulnerability enhances life satisfaction*. In order to test it, we will attempt to overcome several identification challenges.

First, a number of time-varying and time-invariant determinants of happiness may be correlated with income and vulnerability. If omitted from the analysis, those variables may bias the results. Among the time-invariant factors, one can think of *personality traits* and endowments of *social and human capital* (which may have direct impacts on both job prospects and life-satisfaction). More extroverted and optimistic individuals, for instance, may be both 'naturally' satisfied with their life and more likely to find good, secure employment, or, equally plausibly, more willing to face the risks and uncertainty of entrepreneurship. The same may hold for educated or well-connected people. Among the time varying unobservables, *working conditions* are a first, obvious source of bias. Powdthavee (2010) argues that income gains are often correlated with deterioration in the conditions of work and the latter may have an important influence on life satisfaction. Vulnerability might also be correlated with working conditions, though we have no strong a-priori evidence of the sign of such correlation. *Relative income* is a third potentially confounding factor. Extensive empirical evidence has been generated showing that relative income is correlated with the life-satisfaction of individuals in both developed and developing countries (Blanchflower and Oswald, 2002; Luttmer, 2005; Kingdon and Knight, 2004) and it is natural to assume that relative income will be correlated with absolute income and vulnerability. We will attempt to account for these potential sources of bias by including in the model controls for working-conditions (proxied by a measure of satisfaction with work) and for a worker's position in the income distribution. Most importantly, thanks to our panel dataset we will be able to *control for all time-invariant individual characteristics correlated with happiness (e.g. personality traits)*.

The second challenge is methodological: life-satisfaction is generally recorded in datasets like GUHPS as a categorical variable. Modeling it as a discrete (ordered) outcome would, therefore, appear to be the most appropriate approach. However, such approach would not easily lend itself to controlling for those time invariant unobservables that we have argued are of great relevance in the determination of life satisfaction. To address this issue Ferrer-i Carbonell and Frijters (2004) develop a conditional estimator for the fixed effects logit model. Their findings show that *'it makes virtually no difference whether one assumes ordinality or cardinality of happiness answers, whilst allowing for fixed effects does change results substantially'* (Ferrer-i Carbonell and Frijters, 2004).[15] It therefore seems justifiable to assume cardinality of the life satisfaction indicator and use the corresponding estimators.

Third, issues of reverse causality may arise in the analysis. High levels of life satisfaction may help individuals earn higher incomes or reduce their income vulnerability (Graham et al. (2004), De Neve and Oswald (2012)). Such effects may again bias the estimated coefficients β and γ. In order to fully address this problem, we would be required to specify an FE-IV regression approach. However, doubts are often raised about the validity of the instruments proposed by the authors who have attempted the IV or FE-IV approach for income, such as Knight and Gunatilaka (2008) and Powdthavee (2010).[16] Hence, we do not attempt to instrument vulnerability, while fully acknowledging the possibility that these concerns might be important.

Finally, the vulnerability index is a non-linear function of the first two moments of the earnings distribution, which are both modeled as functions of household and individual characteristics in the first stage of the estimation. It follows that the happiness model (where we include both income and vulnerability on the right-hand side) contains *two* functions of those characteristics among the regressors. Separate identification of these two functions implicitly relies on assumptions regarding the relationship between income and well-being. Existing studies have often imposed *linearity* on the relationship and, for comparability, we choose the same approach.[17]

[15] Ferrer-i Carbonell and Frijters (2004) base their analysis on data from the German Socio-Economic Panel. The life satisfaction question in their data uses the same phrasing as GUHPS. The only difference is that it is posed on a 1-10 scale, while in Ghana the scale is from 1 to 5. This should not pose a challenge for comparability. The results reported in our paper confirm the main conclusion of Ferrer-i Carbonell and Frijters (2004) on the importance of controlling for individual fixed effects in empirical models of life satisfaction.

[16] Furthermore, the vulnerability variable has been constructed as a deterministic function of the predicted values of an earnings model, which would complicate an IV strategy.

[17] Fafchamps and Shilpi (2008, 2009) report non-parametric results that show a linear relationship between consumption expenditures and subjective satisfaction with consumption levels, lending empirical support to this modelling choice.

4 Results

We present here three sets of results. First, we discuss our estimates of income vulnerability. Second, we present a number of regressions of happiness on vulnerability, which constitute the central results of our analysis. This section also offers a test to distinguish between the effect of vulnerability and that of two-sided uncertainty. Third, we show the results of a complementary analysis of workers' attitudes to gains and losses based on data from a field experiment.

4.1 Vulnerability Estimates

Table 2 shows the results from estimating the earnings and variance models used to predict vulnerability later in the analysis. The first feature of the results is that while the earnings model (col 1) shows a relatively high predictive power, trying to predict the variance of earnings proves to be a much more challenging exercise (col 2 - 5).[18]

Upon experimenting with different earnings specifications, we concluded that the best model is one that controls for individual fixed effects and for a set of key time-varying covariates (col 1), the choice of which is grounded in a long-established literature on Mincerian earnings regressions (see Rankin et al. (2010) for an application on Ghana using the GUHPS dataset). The results confirm a number of standard patterns observed in related studies of earnings in Sub-Saharan Africa. First, we find a statistically significant effect of firm-size on earnings (captured by positive coefficients on the log of firm-size for wage-employees and on the log of the number of hired employees for the self-employed). Second, we detect a sizable civil service premium and a positive premium for longer tenure in the job. Third, while the linear effect of age cannot be estimated when time-trends are also controlled for, we find some indication of the typical concavity of the age-earnings profile (though the coefficient on age squared is not statistically significant). Since the estimation in col 1 is carried out with controls for individual fixed effects, it is not possible to separately identify the coefficients on time-invariant characteristics such as education and gender.

[18] This is to be expected, partly because a fraction of what appears to be true variation in earnings may in fact be random measurement error.

We estimate several variance models. We first estimate a model with the same covariates used in the earnings regression and no fixed effects. This model is only able to explain a small portion of the variation in the data. We then estimate a model that includes additional covariates that were not part of the earnings model. Motivated by the observation that social networks can provide an important buffer against negative income shocks, such additional variables include the respondent's ethnic background and marital status. This strategy allows us to marginally increase the fit of the model, but requires us to make the assumption that ethnicity and marital status only determine the variance of income and not its mean. Our preferred model is thus the model reported in column 5, where we include the same covariates used in the earnings regression as well as allow for individual fixed effects to determine the variance of income.[19] This strategy is close to the one employed in Chaudhuri (2003) and the measure of vulnerability used in the rest of the paper will be based on this method.

< Table 2 here >

Given our estimates of the earnings model, before we can calculate vulnerability we need to define a low-earnings threshold (alternatively referred to as 'poverty line'), z. Figure 2 shows the percentage of people who are below different income thresholds, while figure 3 shows the resulting cumulative distribution of the vulnerability index. Our choice for the remainder of the paper will be to set z = 10 GHC per month, which approximately translates into 40 USD.[20] When we experimented with alternative lines in the vicinity of

[19] Ethnicity does not vary over time and we also have very little variation in marital status. This gives us a second reason to omit ethnicity and marital status in the variance model with fixed effects.

[20] All income figures have been deflated and the entire analysis will be conducted at constant prices with 1997 as the base year. The reader should be alert to the fact that in 2007 the Ghana Cedi was converted into the New Ghana Cedi at a rate of 10,000 Ghana Cedis to 1 New Ghana Cedi. For simplicity, all the income figures used in the analysis have been converted into New Ghana Cedis. The exchange rate with the USD is not adjusted for PPP.

this value, the main patterns in our results did not change.[21] For $z = 40$ USD the risk of poverty we estimate is substantial for large portions of our sample. The central line in Figure 3 shows that, on average, 35 percent of workers in a given year have a chance of 50 percent or more of being poor in the next period. When we disaggregate the estimation between currently poor and non-poor workers (Figure 4), we find that even among the latter a sizable group faces a large probability of becoming poor.

< Figure 2 here >

< Figure 3 here >

< Figure 4 here >

4.2 Happiness

This section will present the results from estimating our happiness model. Figure 1(b), in the previous section, plots the histogram of happiness responses after splitting the sample by low/high income relative to the poverty line. The histogram shows prima-facie evidence of the link between income and happiness that we are attempting to formally test, with people who are above the low-income threshold more likely to report to be 'satisfied' with their life.

Table 3 reports the results from estimating the workhorse model of happiness (equation 7) using first OLS (col 1 - 2) and then controlling for fixed effects (col 3 - 4). Our first

[21] The poverty line we have chosen (10 Ghana Cedis per month) translates into nearly 170 PPP-adjusted USD. Assuming the typical earner in our sample sustains a family of 5 people with one additional adult (older than 14) and three children (or alternatively a family of 4 with 2 additional adults and one child, based on a standard equivalence scale), this translates into about 1.7 dollars per capita a day, which is close to standard definitions of poverty. This figure is also consistent with the responses to a question added to GUHPS in 2009, which asked respondents to report the income level they deemed sufficient to cater for (i) basic needs and (ii) a comfortable life. We interpret the answers as a direct, albeit crude, measure of workers' reference points, and variation below the reference point can be considered as downside risk. Upon plotting those answers (available upon request), we found that for the vast majority of the sample (more than 90%), our chosen poverty threshold lies below both measures of minimum desirable income. This lends strong support to the assertion that our low-income range is within the domain of poverty as perceived by urban Ghanaians. Nonetheless, we are aware that, while we have good reasons to set the poverty line at 10 Ghana Cedis per month, other thresholds could be used. We thus experimented with poverty lines that range from about one third to about three times our preferred value and found no major changes in our main results.

result is a positive and significant effect of absolute income on life-satisfaction, in line with the existing literature (e.g. De Neve and Cooper (1998)). This relationship is evident in the OLS regressions and, rather strikingly, it does not change significantly once we control for fixed effects. It appears, therefore, that time-invariant unobservables correlated with earnings are not biasing the estimated effect of income on happiness. Interestingly, the size of the estimated coefficient on the log of income in col 3, 0.017, is remarkably close to that estimated by Powdthavee (2010) using data from the British Households Panel Survey and a fixed-effect estimator, 0.019.

On the other hand Table 3 shows that individual fixed effects play an important role in the relationship between vulnerability and life-satisfaction. Once we control for them, we find *a strong negative relationship between vulnerability and happiness, over and above the income effect* just described (recall that in the regression tables this is reported as a *positive* relationship between the *inverse of vulnerability* and happiness). This is the key result in the paper. It is both statistically significant and economically meaningful. Reducing the risk of poverty by 20 percentage points (which amounts to entirely offsetting the risk of poverty for the median worker in our sample) has the same effect on well-being as increasing earnings by 50%.

The change in the coefficient on the inverse of vulnerability, which we obtain upon controlling for fixed effects, is an indication that individuals who tend to be more satisfied with their lives (high κ_i) face more downside risk.[22] This could be the case if a common personality trait determines life satisfaction as well as occupational choice. For example, individuals who value autonomy may be happier as well as more likely to work in sectors such as self-employment, where autonomy comes at the cost of a higher risk of low earnings (Benz and Frey, 2008; Falco et al., 2015). We investigate this hypothesis by computing the correlation between the estimated fixed effects from the happiness model in (7) and the probability of being self-employed (which in Ghana, as shown by Falco (2014), is associated with higher earnings risk than salaried work). Consistently with our hypothesis, we find a positive and highly significant correlation (not shown for conciseness).

[22] The fixed effects appear to correlate negatively with the 'inverse' of vulnerability. Hence, they correlate positively with downside risk.

Our estimation also includes controls for work satisfaction (proxying changes in working conditions), income quartile, age and its square, and marital status. Work satisfaction is closely correlated with life satisfaction and shows by far the biggest positive coefficient in the life satisfaction regression.[23] The income quartile dummies allow us to control for the position of respondents in the income distribution, which has been shown to be a significant preditor of well-being (Clark et al., 2008). Their inclusion in the regression model does not affect our main results (if anything, the coefficient on vulnerability increases slightly). This should reassure the reader that changes in relative income are not confounding our estimates of the relationship between income vulnerability and happiness.[24]

Finally, the vulnerability index has been constructed using estimates from a first-stage model of earnings. Hence, it carries a degree of statistical imprecision that could pose a challenge to the significance of our estimates in the second-stage model of happiness. In order to check the robustness of our results to such concern, we bootstrap the entire estimation sequence (including the first stage to construct the vulnerability index), sampling with replacement to obtain 200 replications of the original sample. The results are summarised in Figure 5 and 6, where we plot the distribution of the bootstrapped coefficients on earnings and on the inverse of vulnerability (from the happiness model), and they are consistent with the discussion so far.

< Table 3 here >

< Figure 5 here >

< Figure 6 here >

4.3 Alternative explanations

The measure of vulnerability employed in this paper focuses on the notion of *exposure to downside risk* and we interpret our findings as showing that the risk of income poverty has

[23] As a robustness check, we tried to exclude work-satisfaction from the estimation and the results did not change significantly (we only detected a slight increase in the effect of vulnerability).

[24] In addition to what is reported in the table, we experimented with finer quantile disaggregation (quintiles and deciles) and the main results did not change.

a significant impact on well-being. An alternative explanation could be that individuals dislike income volatility per se (rather than exposure to downside risk). In this section we attempt to disentangle the two hypotheses by replacing the vulnerability measure used so far with *two-sided measures* of earnings volatility.

First, we use the raw squared residual $\hat{e}_{it}^2$ from a first stage earnings regression with fixed effects as a proxy for income volatility and find no significant relationship with happiness (table 4), despite the sign of the estimated effect is always negative (as we would expect if workers are risk-averse).[25] The lack of statistical significance might be due to the fact that ex-post realizations of the shock are a noisy proxy of the expected degree of vulnerability workers perceive (and are affected by). A way to circumvent the problem is to model the variance of these residuals, as we already did in section 3.1, and use the predicted value as a measure of expected variance. Upon doing that, we document once again a negative effect of volatility on life satisfaction that is not statistically significant (the results are not reported for conciseness, but available upon request). Overall, this evidence points to the conclusion that vulnerability to downside income risk, as analysed in the previous section, plays a more clear-cut role in the determination of well-being than two-sided volatility.

< Table 4 here >

4.4 Choice among risky prospects

Our final piece of evidence comes from a behavioural experiment that studies individual choices between risky prospects when downside risk is present and when it is absent. Our objective is to investigate whether downside risk impacts the economic decisions of the respondents in our sample. We do so by estimating the level of loss-aversion implied by the respondents' choices in a series of lottery games. This complements our analysis of life satisfaction and highlights the role of downside risk in the domain of decision-making (Kahneman et al., 1997).

The experiment, extensively described in Falco (2014), was run in 2007, with a random sub-sample of 307 GUHPS respondents. It consisted of 21 choices between pairs of mon-

[25] Using experimental data, Falco (2014) shows that the majority of GUHPS respondents are indeed risk-averse.

etary lotteries. Each ’game’ was framed as a choice between two opaque urns containing marbles of different colours and, correspondingly, different monetary values.[26] After being shown the composition of each urn, respondents were asked to choose the one from which they would prefer to draw a marble. Prior to making their choices, they were informed that at the end of the game one of their 21 preferred lotteries would be randomly selected and played out. The winnings of that game would then be paid to the respondent. Monetary incentives of this kind are used to induce truthful revelation of preferences.

Choices were framed in terms of losses and gains with respect to the reference point of no gain over the initial endowment. This standard manipulation is ubiquitous in the experimental literature on loss-aversion.

Using data from the experiment, we perform a maximum likelihood estimation of the following utility function, which incorporates a loss-aversion parameter, λ. [27]

$$u(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x^{\beta}) & \text{if } x < 0 \end{cases} \tag{8}$$

This is a standard parametrization of utility functions in the prospect theory literature (Wakker, 2010). An estimate of λ greater than one is evidence of loss-aversion: losses are felt more than gains. In line with prospect theory, we further assume that prospects are evaluated as a weighted sum of the utilities of the various outcomes, where the weights are transformations of the actual probabilities given by the following probability weighting function:

$$\begin{aligned} \omega(p) &= \tfrac{p^{\gamma}}{(p^{\gamma}+(1-p)^{\gamma})^{\frac{1}{\gamma}}} \quad \text{if } x \geq 0 \\ \omega(p) &= \tfrac{p^{\phi}}{(p^{\phi}+(1-p)^{\phi})^{\frac{1}{\phi}}} \quad \text{if } x < 0 \end{aligned} \tag{9}$$

Given these assumptions about the form of the utility function and the weighting of

[26] A detailed description of the experimental setup is contained in Barr (2007).

[27] Following prospect theory, this utility function has separate parameters to define the curvature in the loss and in the gain domain. We assume the reference point is 0 and define x as the gain over the initial endowment. This is consistent with the framing of the lotteries.

probabilities, we can calculate the difference in utility between the two lotteries in each pair. For each two lotteries R and L we obtain:

$$\nabla PU = \frac{\sum_R \omega(p_R)u(x_R) - \sum_L \omega(p_L)u(x_L)}{\mu} \tag{10}$$

∇PU includes an error term capturing the possibility that individuals make mistakes when assessing the utility they would derive from a lottery. We model this error term following the recommendations of the literature (Hey and Orme, 1994; Andersen et al., 2010).The choice of a lottery is then modelled as a stochastic function of ∇PU. The log-likelihood is hence given by:

$$lnL(\alpha, \beta, \lambda, \gamma, \phi, \mu; y, X) = \sum_i [(ln\Phi(\nabla PU)|y_i = R) + (ln\Phi(1 - \nabla PU)|y_i = L) \tag{11}$$

Details on the estimation procedure are further outlined in Harrison (2008) and Falco (2014).

We first estimate the parameters of utility function (8) pooling the choices of all respondents. We cluster standard errors at the individual level, as recommended by Andersen et al. (2010). The estimate of loss-aversion we obtain is 1.77, which is in line with previous experimental findings (Booij et al., 2010; Wakker, 2010). Using a standard test, we can reject the null of $\lambda = 1$ at a 1-percent significance level.

< Table 5 here >

Furthermore, we attempt to estimate a coefficient of loss-aversion for each individual in the sample. Our maximum likelihood routine converges for 266 respondents. However, for 45 of them we obtain estimates of λ above 10, which are inconsistent with the upper bounds reported in other studies (Booij et al., 2010; Wakker, 2010). We exclude these from the analysis. Out of the remaining observations, we estimate λ above 1, indicating loss-aversion, for 55 percent of individuals. The precision of these individual estimates is however low, so we are able to reject the null hypothesis of $\lambda = 1$ for only 22 percent of the respondents. Figure 7 shows the distribution of estimated loss-aversion coefficients.

< Figure 7 here >

Overall, the results in this section complement the discussion above by showing that the workers in our sample are characterised by significant loss-aversion and, therefore, downside risk will not only affect their well-being but also their decisions. Thus, labour markets that expose workers to large downside risks not only harm workers' welfare, but may also lead people to forgo profitable (but risky) investment opportunities (Blumberg and Kremer, 2014).

An interesting extension to the analysis would be to investigate whether individuals with higher loss-aversion experience higher losses in well-being as a result of vulnerability. While it is possible, this may not necessarily be the case since loss-aversion describes the weight that people give to outcomes below a reference point *when making decisions*, and this 'decision utility' may not correlate perfectly with 'experienced utility' (Kahneman et al., 1997). However, as the individual estimates of loss-aversion are rather noisy, we are not in a position to convincingly investigate in our data whether income vulnerability causes larger well-being costs for loss-averse individuals.[28]

5 Conclusions

This article investigates the relationship between income and well-being in a growing developing country, with a focus on the previously unexplored link between *the risk of income poverty and happiness*. Using unique longitudinal data from a representative household survey from urban Ghana, we are able to measure the probability of income poverty at the individual level and explore its relationship with life-satisfaction. Our results are compelling.

[28] There is a second reason not to attempt this analysis. In the experiment we measure loss-aversion with respect to a reference point created by a fixed endowment, which represents the 'status quo' before the decision. This is standard practice in the literature that studies loss-aversion. In the previous sections, on the other hand, we have studied the risk of falling below a poverty line. The poverty line will differ from the status quo for many individuals. Thus, while we study attitudes towards downside risk in both cases, the two measures we use define 'downside risk' differently. In general, the literature on loss-aversion has not yet fully understood how individuals choose reference points and we flag this as an area for future research (Kőszegi and Rabin, 2006). We thank an anonymous referee for raising this point.

First, our analysis reveals a substantial risk of poverty for a large share of the population. Second, we find *a strong negative relationship between vulnerability to poverty and life-satisfaction*, over and above the positive income effect commonly documented in the literature. The result is both statistically significant and economically meaningful. Reducing vulnerability by 20 percentage points has the same effect on well-being as increasing income by 50%. Interestingly, we find that failing to control for individual fixed effects leads to significant bias and misleading conclusions. Further, we attempt to disentangle the effect of *downside risk* on happiness from the effect of *two-sided uncertainty*. We find that the former has the clearest impact on subjective well-being. Finally, in a matched behavioural experiment which elicits respondents' attitudes towards risky prospects we find evidence of significant levels of loss-aversion among our respondents. This suggests that the effect of downside risk is not limited to life evaluation, but extends to decision making.

Our results highlight the importance of social protection as an instrument to shield workers from the risk of falling into poverty. Loss-averse agents will be more willing to undertake productive investments when safety nets and insurance minimize this risk. Importantly, our evidence suggests that these policies will directly improve life satisfaction, and should be considered alongside interventions that target other determinants of well-being, such as working conditions and work-life balance. More broadly, our results suggest that the well-being gains that developing countries can obtain by raising average incomes may be limited if the risk of falling into poverty remains high for a large share of the population.

References

Steffen Andersen, Glenn W. Harrison, Morten Igel Lau, and Elisabet E. Rutstrom. Behavioral econometrics for psychologists. *Journal of Economic Psychology*, 31(4):553–576, August 2010.

Abhijit V Banerjee and Esther Duflo. The economic lives of the poor. *The journal of economic perspectives*, 21(1):141, 2007.

A. Barr. Attitudes to Risk in Ghana: Field Manual. *Unpublished Manuscript*, 2007.

Matthias Benz and Bruno S. Frey. The value of doing what you like: Evidence from the self-employed in 23 countries. *Journal of Economic Behavior & Organization*, 68(3-4): 445–455, December 2008.

D.G. Blanchflower and A.J. Oswald. Well-being over time in britain and the usa. *Journal of Public Economics*, 88 (7-8):1359–1386, 2002.

Seth Blumberg and Michael Kremer. Dynamic loss aversion, growth, and development. *Working Paper*, 2014.

Philippe Bocquier, Christophe J. Nordman, and Aude Vescovo. Employment vulnerability and earnings in urban west africa. *World Development*, 38(9):1297–1314, September 2010.

Adam Booij, Bernard Praag, and Gijs Kuilen. A parametric analysis of prospect theory's functionals for the general population. *Theory and Decision*, 68(1):115–148, February 2010.

S. Chaudhuri. Assessing vulnerability to poverty: concepts, empirical methods and illustrative examples. Columbia University, unpublished manuscript, 2003.

S. Chaudhuri, J. Jalan, and A. Suryahadi. Assessing household vulnerability to poverty from cross-sectional data: A methodology and estimates from indonesia. Columbia University, Department of Economics Discussion Paper 0102-52, 2002.

Andrew E Clark, Paul Frijters, and Michael A Shields. Relative income, happiness, and utility: An explanation for the easterlin paradox and other puzzles. *Journal of Economic Literature*, pages 95–144, 2008.

Jan-Emmanuel De Neve and Andrew J. Oswald. Estimating the influence of life satisfaction and positive affect on later income using sibling fixed-effects. CEP Discussion Papers dp1176, Centre for Economic Performance, LSE, November 2012.

K.M. De Neve and H. Cooper. The happy personality: a meta analysis of 127 personality traits of subjective well being. *Psychological Bulletin*, 125:197–229, 1998.

S. Dercon. Risk, crop choice and savings: evidence from tanzania. *Economic Development and Cultural Change*, 44(3):485–514, 1996.

S. Dercon. Growth and shocks: evidence from rural ethiopia. *Journal of Development Economics*, 74(2):309–329, August 2004.

S. Dercon and L. Christiaensen. Consumption risk, technology adoption and poverty traps: evidence from ethiopia. CSAE Working Paper, 2007.

S. Dercon, J. Hoddinott, and T. Woldehanna. Shocks and consumption in 15 ethiopian villages, 1999-2004. *Special Issue on Risk, Poverty and Vulnerability in Africa, Journal of African Economies*, 14(4):559–585, 2005.

R. Di Tella, J. Haisken-De New, and R. MacCulloch. Happiness adaptation to income and to status in an individual panel. NBER working paper 13159, 2007.

R. Easterlin. Income and happiness: towards a unified theory. *Economic Journal*, 111 (473):465–484, July 2001.

M. Fafchamps. *Rural Poverty, Risk, and Development*. Elgar Publishing, December 2003.

M. Fafchamps. Vulnerability, risk management, and agricultural development. 2009. paper presented at the AERC Conference on Agriculture and Development held in Mombasa, Kenya, on May 28-30, 2009.

Marcel Fafchamps and Forhad Shilpi. Subjective welfare, isolation, and relative consumption. *Journal of Development Economics*, 86(1):43 – 60, 2008. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2007.08.004.

Marcel Fafchamps and Forhad Shilpi. Isolation and subjective welfare: Evidence from south asia. *Economic Development and Cultural Change*, 57(4):641–683, 07 2009.

Marcel Fafchamps, David McKenzie, Simon Quinn, and Christopher Woodruff. Using PDA consistency checks to increase the precision of profits and sales measurement in panels. *Journal of Development Economics*, 98(1):51–57, 2012.

Paolo Falco. Does risk matter for occupational choices? Experimental evidence from an African labour market. *Labour Economics*, 28(C):96–109, 2014.

Paolo Falco and Luke Haywood. Entrepreneurship versus joblessness: Explaining the rise in self-employment. *Journal of Development Economics*, 118(C):245–265, 2016.

Paolo Falco, Andrew Kerr, Neil Rankin, Justin Sandefur, and Francis Teal. The returns to formality and informality in urban africa. *Labour Economics*, 18(S1):S23–S31, 2011.

Paolo Falco, William F. Maloney, Bob Rijkers, and Mauricio Sarrias. Heterogeneity in subjective wellbeing: An application to occupational allocation in Africa. *Journal of Economic Behavior & Organization*, 111(C):137–153, 2015.

A. Ferrer-i Carbonell and P. Frijters. How important is methodology for the estimates of the determinants of happiness? *Economic Journal*, 114(497):641–659, 2004.

B. Frey and A. Stutzer. What can economists learn from happiness research? *Journal of Economic Literature*, 40(2):(402:435), 2002.

Carol Graham, Andrew Eggers, and Sandip Sukhtankar. Does happiness pay?: An exploration based on panel data from russia. *Journal of Economic Behavior and Organization*, 55(3):319 – 342, 2004.

G. W. Harrison. Maximum likelihood estimation of utility functions using stata. *University of Central Florida, Working Paper 06-12*, 2008.

John D Hey and Chris Orme. Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6):1291–1326, November 1994.

Ryan T Howell and Colleen J Howell. The relation of economic status to subjective well-being in developing countries: A meta-analysis. *Psychological bulletin*, 134(4):536, 2008.

Daniel Kahneman and Angus Deaton. High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, 107(38): 16489–16493, 2010.

Daniel Kahneman, Peter P Wakker, and Rakesh Sarin. Back to bentham? explorations of experienced utility. *The Quarterly Journal of Economics*, 112(2):375–405, May 1997.

G. Kingdon and J. Knight. Community, comparisons and subjective well-being in a divided society. CSAE Working Paper, 2004.

John Knight and Ramani Gunatilaka. Aspirations, adaptation and subjective well-being of rural-urban migrants in china. January 2008. University of Oxford, Economics Working Paper Series, No. 381.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, pages 1133–1165, 2006.

R. Layard. *Happiness: lessons from a new science*. Penguin, 2005.

Ethan Ligon and Laura Schechter. Measuring vulnerability. *Economic Journal*, 113(486): C95–C102, March 2003.

Ethan Ligon and Laura Schechter. Evaluating different approaches to estimating vulnerability. June 2004. The World Bank, Social Protection Discussion Paper No. 0410.

E.F.P. Luttmer. Neighbours as negatives: relative earnings and well-being. *Quarterly Journal of Economics*, 120:963–1002, 2005.

Costas Meghir and Luigi Pistaferri. Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32, 01 2004.

Sendhil Mullainathan and Marianne Bertrand. Do people mean what they say? implications for subjective survey data. *American Economic Review*, 91(2):67–72, May 2001.

Nicholas Nsowah-Nuamah, Francis Teal, and Moses Awoonor-Williams. Jobs, skills and incomes in ghana: How was poverty halved? *CSAE Working Paper Series*, 2010-01, 2010.

Andrew J. Oswald and Stephen Wu. Objective confirmation of subjective measures of human well-being: Evidence from the usa. IZA Discussion Papers 4695, Institute for the Study of Labor (IZA), January 2010.

N. Powdthavee. How much does money really matter? estimating the causal effects of income on happiness. *Empirical Economics*, 39(1):77–92, 2010.

N. Rankin, J. Sandefur, and J. Teal. Learning and Earning in Africa: Where are the Returns to Education High? *CSAE Working Paper Series*, 2010-02, 2010.

M. Rosenzweig and H. Binswanger. Wealth, weather risk and the composition and profitability of agricultural investments. *Economic Journal*, 103(416):56–78, 2003.

Peter P. Wakker. *Prospect Theory*. Cambridge University Press, 2010.

Tables

Table 1: Summary Statistics, 2005†

Variable	Mean	Std. Dev.
Age	35.27	10.73
Educ (ys)	8.038	4.032
Male	0.446	0.497
Private Sect. Wage	0.265	0.442
Public Sect. Wage	0.074	0.262
Ln(employees)	0.133	0.418
Ln(firmsize)	0.842	1.551
Years since started current job	10.625	9.888
Married	0.53	0.499
Ga-Dangme	0.145	0.352
Ewe	0.06	0.237
Mole-Dagbani and Hausa	0.096	0.295
Other ethnicity	0.112	0.316
Obs.††		740

†Restricted to observations in the happiness model, pooling survey waves.
††Only 738 observations include ethnic origin.

Table 2: Estimation of Vulnerability

Dep. Var.	y	$\sigma^2(K)$	$\sigma^2(X)$	$\sigma^2(X^2)$	$\sigma^2(X, FE)$
	(1)	(2)	(3)	(4)	(5)
Age		.035	.019		-.0005
		(.028)	(.026)		(.001)
Age2	-.0003	-.0004	-.0002	-5.49e-08	-.0005
	(.0004)	(.0003)	(.0003)	(4.76e-08)	(.001)
Educ		-.028	-.055		
		(.034)	(.032)*		
Educ2		.003	.004	7.17e-06	
		(.002)	(.002)*	(8.15e-06)	
Male		.134	.133	.127	
		(.095)	(.092)	(.092)	
Private Sect. Wage	-.147	-.991	-.944	-.967	-.234
	(.074)**	(.149)***	(.145)***	(.121)***	(.177)
Public Sect. Wage	.196	-1.342	-1.319	-1.313	-.120
	(.112)*	(.184)***	(.180)***	(.177)***	(.269)
Ln(employees)	.187	.008	.040	.044	-.030
	(.051)***	(.115)	(.113)	(.057)	(.121)
Ln(firmsize)	.055	.005	-.013	-.0003	.062
	(.021)**	(.042)	(.042)	(.008)	(.051)
Yrs since curr. job start	.005	.003	.004	.0002	.005
	(.003)*	(.006)	(.006)	(.0002)	(.007)
Married		-.052			
		(.095)			
Eth.: Ga-Dangme		-.037			
		(.119)			
Eth.: Ewe		.392			
		(.171)**			
Eth.: Mole Dag. - Hausa		.578			
		(.155)***			
Other ethnicity		-.181			
		(.162)			
Const.	2.588	-3.041	-2.642	-2.514	-3.293
	(1.346)*	(.519)***	(.491)***	(.190)***	(1.294)**
Indiv. Fixed Effects	Yes				Yes
Time Dum.	Yes	Yes	Yes	Yes	Yes
Obs.†	3659	3014	3110	3110	3110
R^2	.685	.073	.065	.064	.627

Notes: y: (log) real monthly earnings; $\sigma^2(K)$: (log) variance of y modelled as a function of K; $\sigma^2(X)$: (log) variance of y modelled as a function of X; $\sigma^2(X^2)$: (log) variance of y modelled as a function of X^2; $\sigma^2(X, FE)$: (log) variance of y modelled as a function of X incl. indiv. fixed effects (this specification is used to compute vulnerability subsequently in the paper); X is the set of key regressors in the income model, K is an augmented set of regressors to include potential determinants of the variance; omitted occupational category = self-employed; omitted ethnicity = Akan; *Public Sect. Wage* includes all salaried workers in the public sector, including civil servants and workers in public enterprises. †NxT; Col (2) to (5) have fewer observations than (1) since the variance analysis must be confined to respondents who appear multiple times in the panel. Col (2) has fewer observations than (3) to (5) due to missing values in the marriage and ethnicity variables. *Confidence:* *** $\leftrightarrow$ 99%, ** $\leftrightarrow$ 95%, * $\leftrightarrow$ 90%; Robust standard errors in parentheses.

Table 3: Happiness and Vulnerability

	Pooled OLS		*Fixed Effects*	
	(1)	(2)	(3)	(4)
(1-Vul)	-.006	.005	.087	.117
	(.017)	(.018)	(.036)**	(.044)***
LnRealEarn	.013	.033	.017	.051
	(.006)**	(.012)***	(.008)**	(.017)***
LnWorkSatis	.618	.615	.588	.587
	(.014)***	(.014)***	(.025)***	(.025)***
Married		.027		.023
		(.010)**		(.020)
Age		-.006		-.019
		(.003)*		(.014)
Age2		.00007		.0002
		(.00004)*		(.0002)
EarnQuart=2		-.0009		-.026
		(.018)		(.025)
EarnQuart=3		-.048		-.069
		(.023)**		(.034)**
EarnQuart=4		-.060		-.099
		(.032)*		(.048)**
Const.	.426	.489	.388	.784
	(.019)***	(.059)***	(.036)***	(.273)***
Obs.†	2507	2507	2507	2507
R^2	.45	.454	.422	.425

Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%; Robust standard errors in parentheses; †NxT; The number of observations is lower than in the earnings model in table 2 due to missing values in the life-satisfaction variable (most notably due to the fact that there was no question on life satisfaction in the very first wave of GUHPS (2004)). We keep 2004 data in the earnings regressions to ensure we model the earnings process as precisely as possible.

Table 4: Vulnerable to Downside Risk or Averse to Uncertainty? *(Residual)*

	OLS1	OLS2	FE1	FE2
	(1)	(2)	(3)	(4)
$\hat{e}^2_{lny}$	-.003	-.004	-.009	-.011
	(.007)	(.007)	(.010)	(.010)
LnRealEarn	.012	.034	.021	.051
	(.005)**	(.011)***	(.008)***	(.017)***
LnWorkSatis	.618	.615	.588	.588
	(.014)***	(.014)***	(.025)***	(.025)***
Married		.027		.018
		(.010)**		(.020)
Age		-.006		-.007
		(.003)*		(.013)
Age2		.00007		.0001
		(.00004)*		(.0002)
EarnQuart=2		-.003		-.027
		(.018)		(.026)
EarnQuart=3		-.050		-.068
		(.023)**		(.035)*
EarnQuart=4		-.061		-.096
		(.032)*		(.049)**
Const.	.428	.487	.438	.524
	(.019)***	(.059)***	(.031)***	(.247)**
Obs.	2507	2507	2507	2507
R^2	.45	.454	.42	.423

Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%; Robust standard errors in parentheses; †NxT.

Table 5: Maximum likelihood estimates of utility function (8)

Parameter	Estimate	Robust standard error	z	$P > \|z\|$
α	0.42	0.017	24.24	0.00
β	1.90	0.074	25.77	0.00
λ	1.77	0.169	10.45	0.00

Figures

Figure 1: Distribution of Life-Satisfaction

(a) Across the population

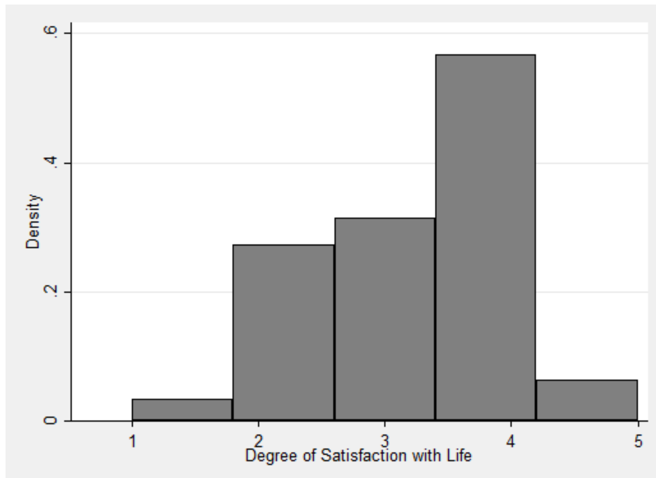

(b) By level of income (employed people only)

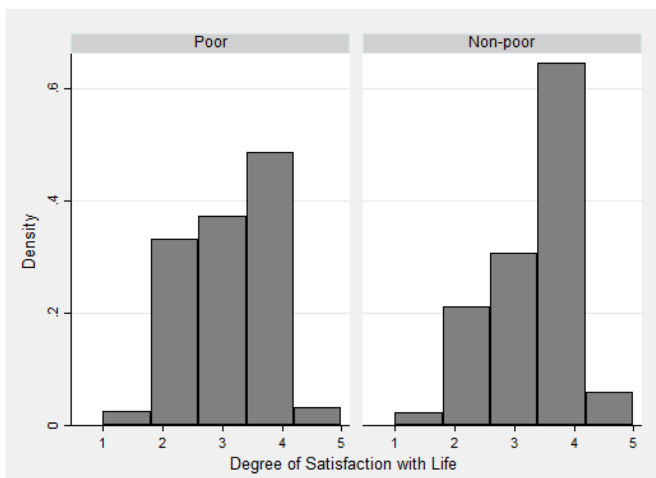

Figure 2: Percentage of employed with $y < z$

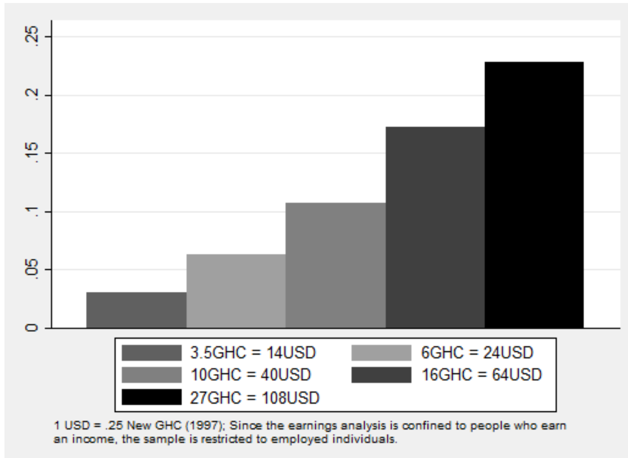

Figure 3: Cumul. Dist. of Vulnerability for different poverty lines (z_t)

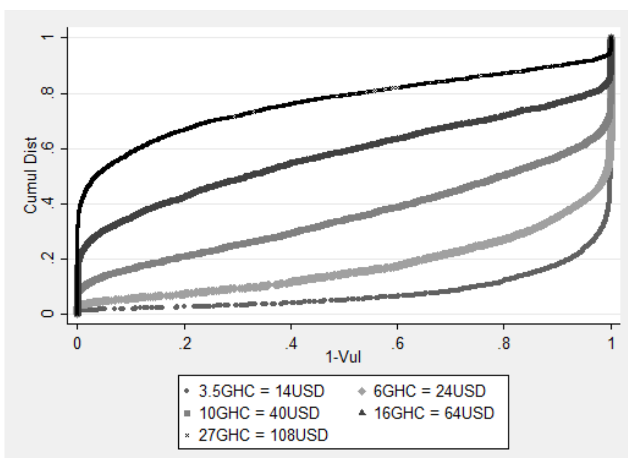

Figure 4: Cumul. Dist. of Vulnerability by current poverty status

z_t = 10 (1997) GhCedis

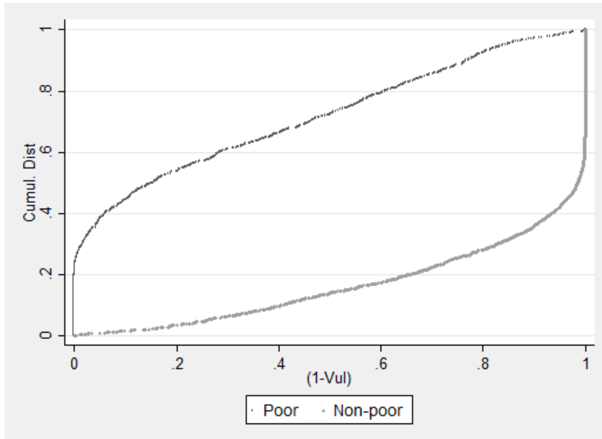

Figure 5: Bootstrapped distribution of the coeff on *LnRealEarn*

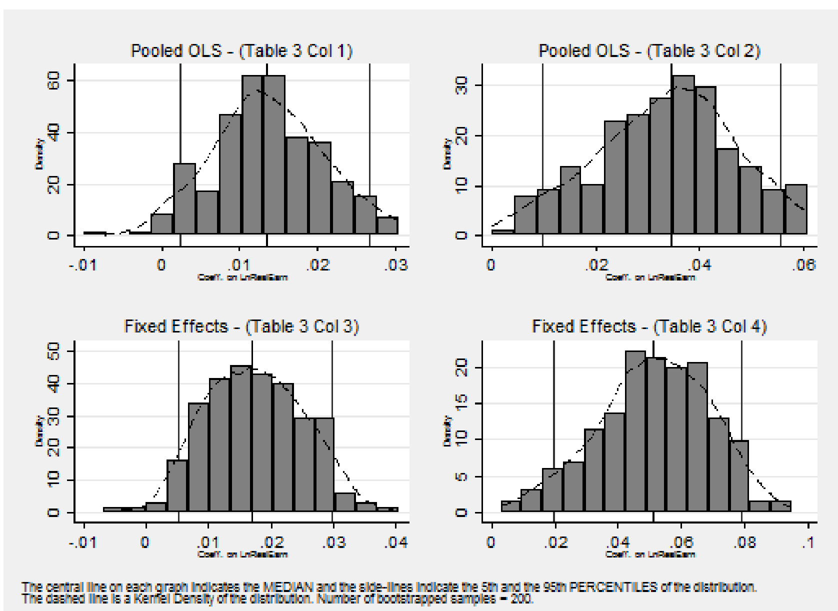

Figure 6: Bootstrapped distribution of the coeff. on *(1-Vul)*

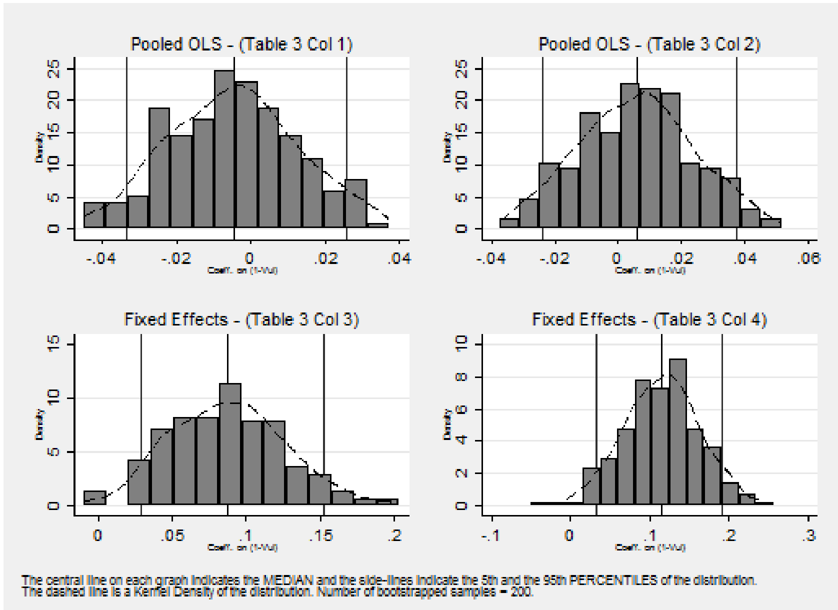

Figure 7: Distribution of λ