Environmental Adaptation of Risk Preferences*

Salvatore Di Falco[1] and Ferdinand M. Vieider[2,3]

[1] *Geneva School of Economics and Management, Switzerland, Salvatore.DiFalco@unige.ch*
[2] *RISL$\alpha\beta$, Department of Economics, Ghent University, Belgium, ferdinand.vieider@ugent.be*
[3] *RISL$\alpha\beta$ Africa, University Mohammed VI Polytechnic, Morocco*

March 18, 2021

Abstract

We present incentivized panel data measuring risk preferences from across Ethiopia, and pair them with rainfall data. We use these data to test evolutionary predictions on environmental adaptation of risk preferences. We find rainfall shocks to decrease risk tolerance for the same individuals over time in the short run. We also find that historical rainfall characteristics and geographical features can explain 40% of the variation in preferences across individuals in the long run. The short-term effects are perfectly aligned with the long term effects we document, painting a unified and highly consistent picture. This provides the first real world evidence that preferences may systematically adapt to the environment of the decision maker.

Keywords: risk preferences; panel data; shocks

JEL-classification: C93; D03; D80; O12

*We gratefully acknowlege financial support by the Swiss National Science Foundation (SNF), and by the International Development Research Centre (IDRC). We thank the Environment and Climate Research Center (ECRC) at the Ethiopian Development Research Institute (EDRI) for logistical support. We are grateful to Mohammed Abdellaoui, Bart Cockx, Gerdie Everaert, Nick Netzer, Amma Panin, Arthur Robson, Laura Schechter, Aleksei Tetenov, Peter Wakker, and Jeffrey Wooldridge for helpful discussions. We thank Avichal Mahajan for invaluable research assistance. All errors remain our own.

A great deal can be learned about rational decision making by taking into account [...] *the limitations upon the capacities and complexity of the organism, and by taking account of the fact that the environments to which it must adapt possess properties that permit further simplication* [sic] *of its choice mechanisms.*

Herbert A. Simon (1956), p. 129

1 Introduction

Preferences towards risk and uncertainty play a key role for economic decision making. They contribute not only to the determination of investments, but also to labor market choices, investments into education, and marriage and fertility decisions. As drivers of entrepreneurship, they contribute to shaping the development and growth prospects of entire countries (Galor and Michalopoulos, 2012; Doepke and Zilibotti, 2014). It thus appears desirable to understand what shapes risk preferences. Our ability to explain the variation in preferences over time and across individuals, however, remains limited.

We contribute to this topic by documenting the role of environmental adaptation in shaping risk preferences. Despite recent advances in our knowledge about socio-demographic correlates of risk preferences (Tanaka, Camerer and Nguyen, 2010; Dohmen, Falk, Huffman, Sunde, Schupp and Wagner, 2011; von Gaudecker, van Soest and Wengström, 2011; Choi, Kariv, Müller and Silverman, 2014; Noussair, Trautmann and van de Kuilen, 2014; Falk, Becker, Dohmen, Enke, Huffman and Sunde, 2018), relatively little is still known on what causally determines preferences. We provide direct evidence suggesting that risk preferences systematically adapt to the decision maker's environment.

We organize our results through the lens of evolutionary models of preference adaptation (Robson, 2001a; Netzer, 2009). The scarcity of cognitive resources juxtaposed to infinitely many possible outcomes makes it evolutionarily optimal for neural sensitivity to adapt to the expected consumption opportunities present in a given environment. Such models are underpinned by neurological sensory adaptation mechanisms (Tobler, Fiorillo and Schultz, 2005; Wark, Lundstrom and Fairhall, 2007; Khaw, Glimcher and Louie, 2017), and can be used to derive concrete predictions for our data. We think about adaptation of preferences as being driven by expectations about future outcomes as determined by historical realizations of consumption opportunities. Cognitive limitations imply that only discrete increases in utility can be detected, so that utility takes the form of a step function with the steps corresponding to constant increases in utility. The shape of the utility function will then be determined by the allocation of the consumption thresholds at which the jumps in utility take place.

The key insight emerging from the model is that to avoid costly mistakes, these *perception thresholds* will be allocated where outcomes are most frequent and where mistakes are most costly from an evolutionary point of view. Netzer (2009) showed that such a setup results in a reference-dependent utility function proportional to the distribution of consumption opportunities in the environment, thus providing evolutionary underpinnings for an S-shaped utility function incorporating decreasing sensitivity relative to a reference point (Markowitz, 1952; Kahneman and Tversky, 1979). Sudden shifts in expectations caused by unexpected outcome realizations will result in shifts of the attention thresholds, and thus in shifts of the utility function which are distinct from movements along a pre-existing utility function (see Netzer, 2009, p.947). Section 2 provides a sketch of this model, which guides our empirical analysis.

We test environmental adaptation in preferences using the results of a panel study conducted with subsistence farmers from across the Ethiopian highlands. We obtained detailed, incentivized measures of the risk preferences of 900 respondents living in 20 Woredas (administrative districts). The Woredas were chosen in a stratified design to represent differences in both average rainfall and rainfall variation. We link the preference measures to a database of rainfall combining infrared satellite imagery with data from rain gauges on the ground. This provides an ideal testbed to study the causal determinants of preferences. The Ethiopian highlands are characterized by high environmental variability both over time and across space. Given the dependence of local livelihoods on rain-fed agriculture (Dercon and Christiaensen, 2011; Dercon and Porter, 2014), we can use exogenous variation in rainfall during the growing season to study how environmental shocks change preferences for the same individuals over time.

Our model posits that the effect of rainfall shocks ought to pass through consumption levels. Since agricultural production and hence consumption is closely linked to rainfall in our study setting, we prefer to use the latter as our main variable of interest due to its exogenous nature. To show the suitability of this approach, we also conduct a detailed analysis of how agricultural yields are affected by rainfall shocks (see section S2.2 in the supplementary materials). Using rainfall deviations from the local historical mean during the main growing season, we find that both shortfalls in rain and excess rainfall reduce agricultural yields significantly. We thus use separate indicators for the two types of shocks in our main analysis. In the long term, however, agricultural practices adapt to local circumstances, so that higher average levels of rain are indeed beneficial to agricultural production—a finding that is consistent with the existing evidence for sub-Saharan Africa (Barrios, Bertinelli and Strobl, 2010).

We find that rainfall shocks reduce risk tolerance within subjects over time. We further show that the cross-sectional effects of the observed rainfall shocks go in the opposite direction of the time-changing effects in our panel data. This emphasizes the added value of our findings over a literature that is to date prevalently cross-sectional.[1] The spurious results in our cross-sectional data are driven by large and systematic differences in preferences across geographical regions pre-existing the shocks themselves, which leads to a 'randomization failure' (details in section 5.1). These findings thus contribute to consolidating a literature on the effects of different types of 'shocks' on preferences (Voors, Nillesen, Verwimp, Bulte, Lensink and Van Soest, 2012; Cameron and Shah, 2015; Hanaoka, Shigeoka and Watanabe, 2018; Jakiela and Ozier, 2019), which has arrived at highly contradictory conclusions (Chuang and Schechter, 2015).

In the long run, we find that environmental factors—beyond driving changes in preferences over time—also explain a large part of the variation in risk preferences across space. These effects are highly consistent with the movements we observe over time, painting a coherent picture on how preferences change and adapt to a given environment. We find that historical rainfall variables and geographical characteristics such as altitude explain over 40% of the cross-sectional variation in idiosyncratic preferences, defined as the individual preference component obtained after filtering out changes over time. These findings depart discretely from previous studies, which concluded that observable characteristics of decision makers and their socio-economic surroundings could only explain a small fraction of the variance in risk preferences between individuals (von Gaudecker et al., 2011; L'Haridon and Vieider, 2019).[2]

[1] Our shocks are 'exogenous' in the sense that they cannot be influenced by respondents, thus excluding reverse causality. In cross-sectional analysis, however, these shocks may still be correlated with the error term, thus not meeting the bar of exogeneity according to the econometric definition of the term. This is indeed what drives the difference between the within and between estimators we document.

[2] An exact figure of the variance explained in previous studies is difficult to come by because of the wide variety of techniques used and the inconsistent reporting of variance metrics. von Gaudecker et al. (2011) state that the variation associated with demographic chracteristics, "is small compared to the variance ascribed to unobserved heterogeneity" (p. 666). Sutter, Kocher, Glätzle-Rützler and Trautmann (2013) explain about 4% of the variance in risk preferences in their sample of school children. Noussair et al. (2014) can explain at most 6% of the variance in their estimated risk aversion parameter using a wide array of demographic and economic characteristics. Vieider, Lefebvre, Bouchouicha, Chmura, Hakimov, Krawczyk and Martinsson (2015) report R^2 measures between 0.01 annd 0.07 for their incentivized measures of risk tolerance. At the higher end of the spectrum, Cesarini, Dawes, Johannesson, Lichtenstein and Wallace (2009) attribute 16% of the variance in risk preferences to genetic

A key assumption underlying our long run analysis is that the current place of residence of respondents must correspond to their past place of residence. This would prevent selection effects due to migration. We therefore collected information about respondents' place of birth. We found that about three quarters of the sample was born in the current village of residence. In a robustness analysis using the data on birth place, we do not find any evidence for systematic selection effects. The main patterns we document further remain stable if we restrict our sample to respondents who continue to reside in the village where they were born. The coherence of the long-run effects with the time-changing effects, which warrant a causal interpretation under much milder assumptions (see section 4 for details), provides further support for a tentatively causal interpretation of these findings.

It should be emphasized that our findings cannot be interpreted as simple movements along a utility function, but constitute genuine shifts in preferences. The differences we document across space hold for households at the same wealth level, thus excluding simple movements along a fixed utility function defined over lifetime wealth as postulated by expected utility theory. Measuring utility over significant stake ranges, we unequivocally find the utility function to be characterized by increasing relative risk aversion and constant absolute risk aversion. Given that we find relative risk aversion to *increase* following shocks, however, accounting for our effects through movements along the utility function would require a function characterized by *decreasing* relative risk aversion—the opposite pattern of what we find, and a form which has not received any empirical support (see Wakker, 2010, section 3.5, for a review).

2 Theoretical model

We start by deriving predictions for our data building on the evolutionary insights discussed by Robson (2001a), and their further development by Netzer (2009). Let $y \in \mathcal{Y}$ designate fitness in an evolutionary sense (i.e., number of surviving children), and let $V(y)$ designate utility over such fitness levels. The central assumption of the model is that neural mechanisms are fundamentally evolved to maximize evolutionary fitness. Given the low frequency with which fitness outcomes are observed, however, in practice maximization will take place over per-period consumption levels (Robson, 2001b). Let

factors in a sample of Swedish twins.

$c \in \mathcal{C}$ designate consumption. Since fitness will be a function of consumption, we can now write $y = \phi(c)$, where ϕ is a function mapping consumption into fitness. It follows that $V(y) = V(\phi(c))$, which we will henceforth simply write as $U(c)$.

The fundamental insight underlying the model is that individuals have limited cognitive capacity, which contrasts with potentially infinitely many values of c. This will make it optimal for an organism to adapt to the most likely ranges of outcomes in its environment. This insight builds directly on neurological evidence on how signals about rewards are encoded in the brain and translated into decisions (Tobler et al., 2005; Stauffer, Lak and Schultz, 2014). Given the scarcity of neural resources, it will be evolutionarily optimal for an organism to allocate the finite number of perceptual thresholds at its disposal where they matter most. Utility will then take the form of a step-function with a finite, but potentially large, number N of distinct steps corresponding to *just noticeable differences* in utility (Robson and Whitehead, 2017). Different outcomes located on the same utility step cannot be distinguished from each other, resulting in random choice. Assume without loss of generality that each step corresponds to a utility increment of $1/N$. The shape of the utility function will now be determined by the location of the thresholds in the outcome space $\mathcal{C}$ determining the steps in utility. In particular, Netzer (2009) showed that the thresholds need to closely track the cumulative distribution function of outcomes in the environment in order to minimize the likelihood of costly mistakes. That is, it is optimal to be most sensitive to different outcomes where they occur most frequently, and where the mistakes are most costly from an evolutionary point of view.[3]

The solid blue line in figure 1 shows a highly stylized utility function, $U(c)$, as predicted by the model. The curve is steepest in the regions corresponding to frequent outcome realizations, corresponding to the peak of the probability density function of consumption opportunities in the environment, $g(c)$, where the discrimination thresholds, $\hat{c}_i$, are closest together. It is flatter for very small and very large consumption levels, since these are less frequent and/or good discrimination between outcomes in these regions may be less important. This provides a natural underpinning for modelling decreasing sensitivity relative to a reference point, such as proposed by Markowitz (1952) and incorporated into prospect theory by Kahneman and Tversky (1979).

[3] Utility is proportional to consumption opportunites rather than mimicking them exactly because it is also important how different consumption opportunities map into evolutionary fitness, as captured by the mapping parameter ϕ. This is also why large shortfalls in consumption will elicit a larger reaction than potential excess consumtion. See Netzer (2009) for technical details.

Figure 1: Utility step functions across environments

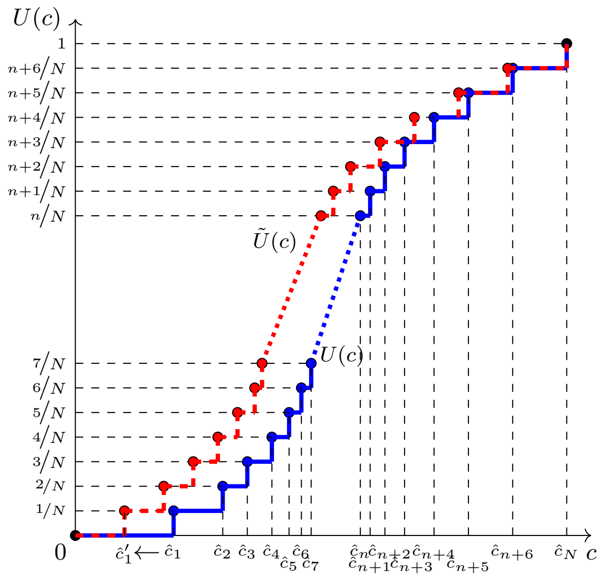

Stylized illustration of the adaptive model. Consumption levels $\hat{c}_i$ and $\hat{c}'_i$ represent perception thresholds. Given that utility differences are equally spaced based on the observation that such differences are 'just noticeable', the spacing of the perception thresholds fully determines the shape of the utility function. The blue, solid line indicates a baseline utility function labelled $U(c)$. The dashed red lines indicate a different function $\tilde{U}(c)$, which is shifted to the left and thus more concave, indicating an increased degree of risk aversion. Dotted lines indicate linear interpolation.

People living in an environment with relatively low overall consumption levels, i.e. represented by a pdf g' shifted to the left relative to g, would be expected to be more risk averse than people living in lusher environments, simply because their attention thresholds, $\hat{c}'_i$, would be geared towards taking optimal decisions for such lower consumption levels, and thus shifted leftwards. By the same token, people living in environments where historical realizations are more dispersed, i.e. where the function g' has fatter tails than g, would also need to allocate more attention to lower consumption realizations, given the potentially deleterious impact of such low consumption on fitness. This is illustrated by the utility function $\tilde{U}(c)$ in figure 1, represented by the red dashed line, which is more concave than the blue function, thus expressing higher levels of risk aversion. This would hold even if both utility functions were characterised by constant absolute risk aversion, setting this account apart from wealth effects under expected utility theory.

The comparative statics just described have at their base a mechanism by which the perceptual thresholds are updated over time to reflect (perceived) shifts in the environment. That is, if the environment shifts from a distribution function $g(c)$, which is known from experience, to a function $g'(c)$, we would also expect the utility function to change from $U(c)$ to $\tilde{U}(c)$. Put differently, beliefs about the environment, and with them the utility function, will adapt over time if the decision maker observes outcomes which differ from the predicted or expected outcomes (Schultz, Dayan and Montague, 1997; Schultz, 2016). Following Robson and Whitehead (2017), we thus allow the consumption thresholds to be time-dependent and to follow a simple updating rule:

$$\hat{c}_i^{t+1} = \hat{c}_i^t + \xi_s[g_s^{t+1}(c) - g_s^t(c)], \tag{1}$$

where ξ is a function governing the extent to which the thresholds react to perceived shifts in the density of consumption. The subscript s indicates that both the adjustments and the perceptions of the density function may be subjective.[4] To the extent that the consumption distribution g^{t+1} is perceived to be lower than g^t, i.e. there is a negative shift in the density function resulting from a rainfall shock, the consumption thresholds will adapt to this by shifting downward. This represents the adaptive mechanism of our model, which translates perceived shifts in the environment into changes in the utility function, and hence in risk preferences.

3 Data and measurements

3.1 Sampling Framework and Descriptives

Sampling. The sampling area comprises the mountainous Ethiopian heartland. Lower-lying regions in the south and east of the country were excluded from the sampling frame because they have different geographical features and farming practices, and because of security concerns. Observations inside these regions derive from 20 different *Woredas* (administrative districts). Figure 2 shows the geographical distribution of sampled households. The sampling area measures 7.5 degrees latitude times 5.2 degrees longitude, corresponding to 581 by 714 kilometers.

The sampling frame was developed to ensure representation at the Woreda level of rainfall patterns in terms of both annual total and variation (see S6 for further details). From each Woreda, 50 households were randomly selected from municipal rosters dating from 2004. Upon our first visit in 2013, we identified 930 households for participation in the study. Some households were absent and could not be reached after several attempts,

[4]The subjectivity of the adjustment function serves to drive home that different people may adjust at different speeds, but also that additional characteristics of the environment may drive the extent of adaptation. We remain agnostic as to what may exactly enter this function, leaving this issue to the empirical analysis.

Figure 2: Geographical location of samples in Ethiopia

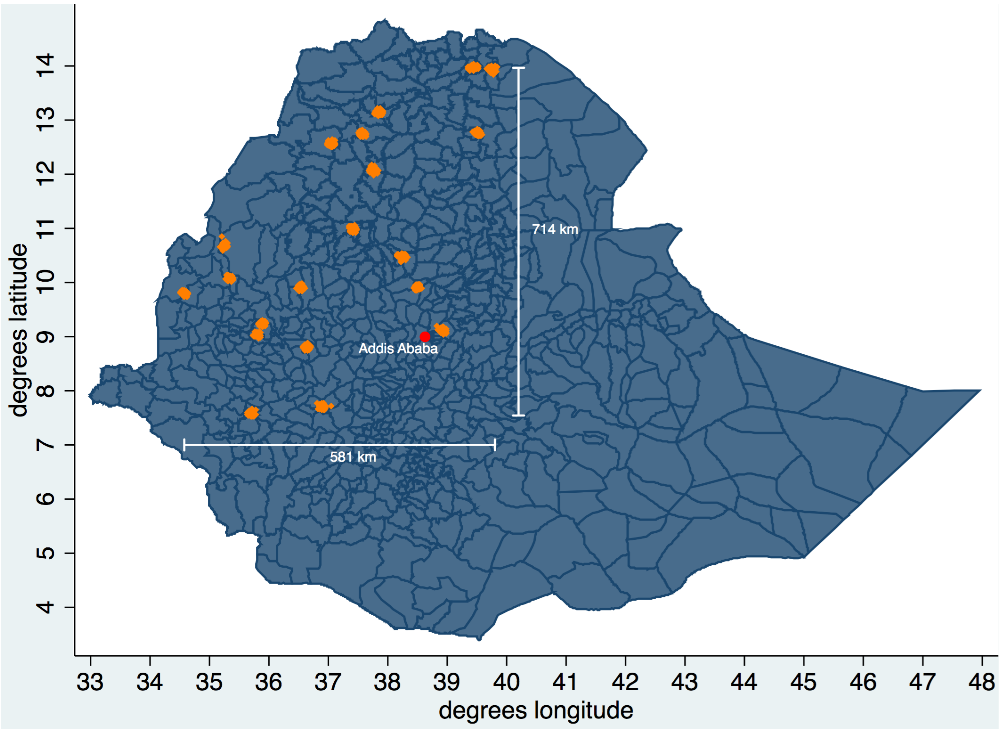

leaving us with a sample of 923 in the first round of the experiment (2013). In the second round (2015), this number was reduced to 910 households, with some households lost due to issues in identifiers that impeded matching, and some because of attrition. In the third round (2017), this sample was further reduced to 861 household. Overall, we end up with a total sample of 906 households, since we can only use households with at least two years in the data in our panel data analysis.[5] We do not find attrition to be explained by any observable characteristics of the Woreda or the household, including altitude, historical mean rainfall and variation in rain, distance to the national and regional capitals, and individual characteristics such as gender, age, or indeed risk tolerance as measured in 2013 (results available upon request).

Timeline. We use rainfall levels during the main agricultural season, or *Meher*, as our main independent variable.[6] The Meher allows the main staple crops, such as teff, maize, and wheat, to be grown. The main rainfalls tend to start in late June and continue through September. Harvest takes place from October to November (figure 3). There

[5] Note that this figure does not line up exactly with the changes from wave to wave detailed above. This happens because a few of the households 'lost' between 2013 and 2015 do show up again in 2017. We even have some households that could not be found in 2013 showing up in 2015 and 2017.

[6] Temperature varies very little across time close to the equator, and most of the variation in temperature in our samples takes place across Woredas, with only 5% ocurring across time.

Figure 3: Time line of data collection

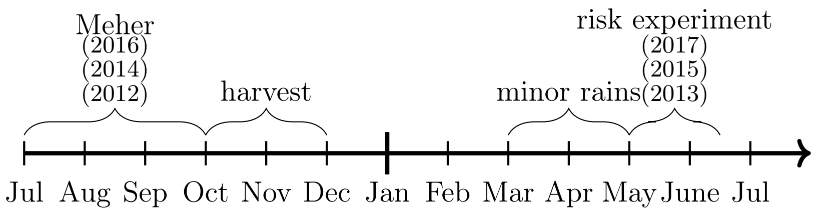

further is a minor rainy season in March to early May, called *Belg*, used mostly for small vegetable crops (onions, peppers, and some pulses). These small crops are mostly used for immediate consumption, and have a minor impact on the overall yearly food production. We thus use rainfall in the minor rainy season as a placebo. We conducted all risk measurements in May and early June—an idle period during which no farming activities take place.

Subject characteristics. We always conduct the experiment with the self-declared household head by means of individual interviews. The mean age of the household heads participating in our study is 49.84 in 2013, and 84% are male. All households live mainly from farming, and 26% declare to have some non-farm income as well. Only 45% of the participants are literate. Households farm an average of 0.41 hectares of land (about 1 acre; SD 0.62 ha).

3.2 Rainfall data

We obtained our rainfall data from the Climate Hazards Group, using the Climate Hazards Infra-Red Precipitation with Station data (Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, Husak, Rowland, Harrison, Hoell et al., 2015). The data combine satellite imagery with station data to produce a grid of rainfall data with a $0.05° \times 0.05°$ resolution (3×3 nautical miles close to the equator). This gives us 343 distinct observations about historical rainfall levels. We refer to these separate locations as *areas*. The data comprise rainfall levels from 1981 to the present. Our main measure of interest is the total rainfall occurring during the main rainy season, or *Meher*, by area. In addition, we use the total rainfall during the minor rainy season, or *Belg*, as a placebo, since only minor crops such as vegetables and some pulses are grown in this period.

We assemble historical measures by area from 1981 to 2010. We define shocks as standardized negative and positive absolute deviations from these means:

$$d_{at} = \frac{\ell_{at} - \mu_a}{sd_a}, \tag{2}$$

where d_{at} indicates the absolute deviation in a given year t in a determined area a, ℓ_{at} indicates the local rainfall level in that year, μ_a is the local average historical rainfall from 1981 to 2010, and sd_a is the historical standard deviation in the same time span. This definition captures the informative value of the rainfall realization relative to historical realizations. The area-specific rainfall measures are matched to individual households using GPS coordinates. That is, all households whose GPS coordinates fall within a given rainfall area are attributed the measures specific to that area.

The assumption underlying the use of this measure is that agricultural practices are adapted to local circumstances, and that what counts as a shock are deviations from typical or 'expected' conditions. In this sense, both shortfalls in rain or droughts, and excess rainfall or floods, may constitute a shock. That said, the effect of shortfalls and excesses in rainfall are likely to work though very different mechanisms. Droughts are widely recognized as being problematic in the context of sub-Saharan agriculture, with rain levels considered to be generally too low (Barrios et al., 2010). Deleterious consequences of droughts are well-documented in the development literature (Rose, 1999; Maccini and Yang, 2009), including for Ethiopia (Dercon and Porter, 2014).

Excess rainfall may be deleterious for different reasons. Large and highly concentrated volumes of rainfall may result in flooding, erosion of fields in Ethiopia's mountainous terrain, washing out of seeds, or rotting of harvests. Using data on maize yields in the US, Li, Guan, Schnitkey, DeLucia and Peng (2019) showed that excessive rainfall can lead to drops in yield comparable to those caused by drought, but that this effect is less uniform and interacts with other characteristics of the terrain and the environment. Borgomeo, Khan, Heino, Zaveri, Kummu, Brown and Jägerskog (2020) show that excessive soil moisture can impact the yields of especially maize and wheat, two important crops in our study area. Derbile, File and Dongzagla (2016) present evidence from Ghana showing that smallholder farmers are adversely affected by both droughts and excess rainfall. Using global data on rainfall and GDP at the level of $0.5° \times 0.5°$ cells and applying fixed effects estimators, Damania, Desbureaux and Zaveri (2020) show that the relationship between rainfall and GDP growth is inverse-U shaped, with GDP growth increasing with rainfall up to a certain level, after which it starts declining. Most of this effect is driven by developing countries in their data.

To gain a better understanding of these issues in our context, we obtained detailed plot-wise data on agricultural production (for details, see section S2 in the supplementary materials). In each survey round, we asked for the land area allocated to different crops in the previous season, as well as the yield for each type of crop. A detailed analysis of these data reveals two main insights. First, crops are adapted to the local conditions, that is, high-yield and high-value crops such as maize and teff—an indigenous grain used to make *injera*, the local bread—are grown mostly in regions with high historical levels of rainfall, and with relatively low historical standard deviations. Other crops such as barley and sorghum tend to be grown in regions with lower levels of historical rainfall. Especially sorghum, a high-yield and relatively high-value crop, seems to be adapted to drier conditions, since its absence in high-rain areas may otherwise seem puzzling.

The second finding concerns the impact of positive and negative deviations in rainfall from historical means on crop yields. Given that crops are adapted to historical conditions, both types of rainfall shocks result in considerable reductions in crop yields. While excess rainfall may well work though different channels than droughts, such as high concentration during critical phases of the crop growth cycle, or interactions with specific terrain characteristics such as proximity to streams, terrain steepness, etc., it tends to have effects that are equally deleterious to those caused by drought. Notice also that while effects on yields and consumption tend to be severe, effects on wealth are much more muted. This is because wealth consists mostly of agricultural land, which technically belongs to the state and cannot be sold in Ethiopia, and the houses built on that same land. Wealth fluctuations, if any, are thus by necessity minor in nature.

Rainfall variability across geographic areas as well as year-on-year within each area is large. Historically, we observe most of the variation in rainfall across Woredas, which accounts for fully 76% of the variation. The variation across areas within a given Woreda, on the other hand, is relatively small at 5% of the total. The remaining 19% of the overall variation takes place within any given area over time. Figure 4 shows the rainfall deviations for the Meher seasons immediately preceding our preference measurements (the deviations lagged once and twice are shown in figures S2 and S3). In 2012 we observe some excess rainfall, although few households experience rainfall more than 1 SD in excess of the historical average. In 2014 we observe extensive droughts, with

Figure 4: Rainfall deviations from historical average

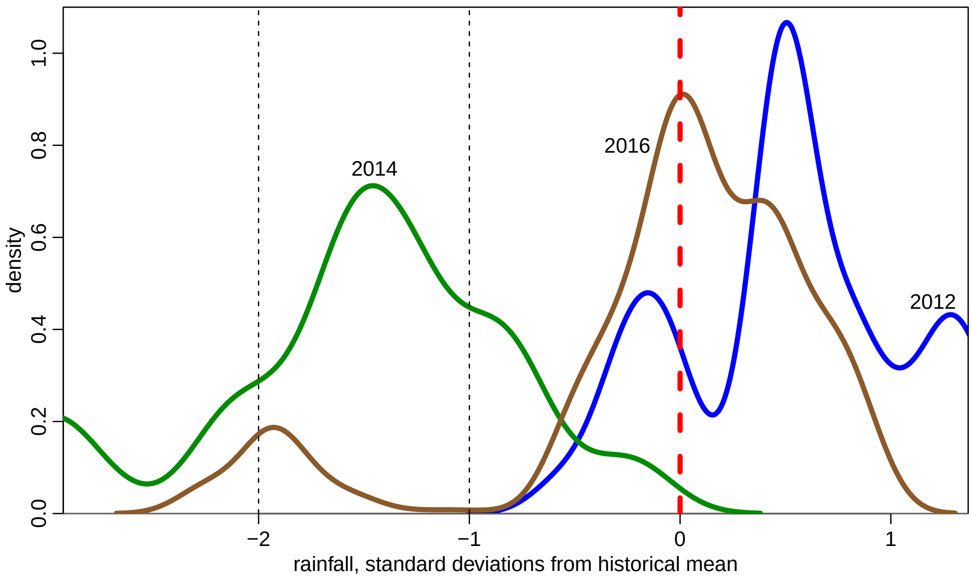

The figure shows the rainfall deviations from the historical mean in each area in standard deviations, separately for each Meher season immediately preceding our experimental measurements. 2012 saw some light excess rainfall, while 2016 was largely normal, except for a few households that experienced a shortfall around 2 SDs. 2014 was characterized by extreme droughts that affected a large part of the sample, albeit to different degrees. The graph is cropped for better display, removing the most extreme 2.5% of the distribution to either side.

a majority of the sample experiencing shortfalls in rain between 1 and 2 SD below the historical average. This constitutes one of the worst droughts in recent memory, on a par with the one of 1982, which triggered extensive famines.[7] Finally, in 2016 we observe largely regular rainfalls, with a minority of respondents experiencing severe droughts. Overall, we thus observe considerable droughts during our study period but only moderate excess rainfall. As a consequence, we expect any effects of shortfalls in rain to be clearly identified, while any effects of excess rainfall will likely be more tricky to detect due to the weak 'treatment'. In addition to the variation over time, we also observe large geographical variation in rainfall patterns in each season. Figure S1 shows maps of our 20 Woredas indicating average rainfall levels for the Meher immediately preceding our risk experiments. By comparing the maps to each other, one can see that the Woredas affected by rain shortfalls and excesses change over time.

[7] While the drought was comparable to the one in 1982, its consequences were not. The consequences of the 1982 drought, with over a million lifes lost over the following years, were so disastrous bacause of the combination with civil war which made the access for help organizations all but impossible (Dercon and Porter, 2014).

3.3 Risk preference data

Elicitation procedure

We elicited 14 certainty equivalents (CEs) for each respondent per round.[8] CEs are well suited for experiments in developing countries, because they are amongst the simplest tasks to measure risk preferences. Physical representations of the choice problems are straightforward. Only monetary amounts vary within a given choice list, while probabilities stay fixed. This makes it easy to lay out money on a table and represent probabilities physically, which is a great advantage given people's familiarity with money. Furthermore, they are easy to manipulate and to use in the construction of nonparametric indices, as well as in the identification of the parameters of preference models, thus explaining their popularity (Bruhin, Fehr-Duda and Epper, 2010; Abdellaoui, Baillon, Placido and Wakker, 2011; Dohmen et al., 2011; Sutter et al., 2013).

The 14 prospects differed both in terms of probabilities of obtaining the high amount or *prize*, and in terms of the amounts themselves. The design follows the one used by Vieider et al. (2015), but only used the gain part of that experiment and known probabilities. We presented prospects with 50-50 probabilities first, namely Birr $\{(30,0);(60,0);(120,0);(180,0);(180,60);(180,120)\}$. These prospects were followed by prospects in order of ascending probability, with $p = i/8, i = 1, ..., 7$, offering either Birr 120 or else 0.[9] The expected earnings for a risk-neutral participant were around €18 (USD 24) in PPP, with the highest prize reaching €30 PPP. These are significant amounts for rural Ethiopian households, a majority of whom live on less than $2 PPP per day. Tasks were kept in a fixed order to facilitate the physical representation of the prospects using colored balls and money, since only either probabilities or outcomes would typically change from one task to the next. A test of order effects conducted with students in Vietnam showed that such a fixed ordering facilitated the task, while not producing different results from a random order (results available upon request). A previous experiment in rural Ethiopia also showed no order effects (Vieider, Beyene, Bluffstone, Dissanayake, Gebreegziabher, Martinsson and Mekonnen, 2018). Importantly, the lists are balanced on average, so

[8] In 2017, we elicited a total of 17 CEs. The three additional CEs were a repetition of CEs already included in the 14 initial ones, and were added to obtain an indicator of the test-retest reliability of our measures. By adding them to the end of the experiment, after the 14 regular measures had been obtained, we avoided tinkering with the main design features of the panel.

[9] Given that all choices were presented physically to the subjects, many of whom were illiterate, we did not have experimental instructions in the traditional sense. We did, however, have scripts that were meant as a reference for the enumerators. These scripts are included in section S5.

that the expected value switching point falls into the middle of the choice list, serving to avoid systematic noise deriving from the administration of unbalanced choice lists (Andersson, Tyran, Wengström and Holm, 2016; Vieider, 2018).

Subjects were asked to choose repeatedly between a prospect and a list of sure amounts ranging between the high and the low amount of the prospect and changing in steps of 3 Birr. Since they have to choose between this invariant lottery and different sure amounts, it is straightforward to find the amount at which subjects want to switch from choosing the prospect to choosing the sure amount. The CE of the prospect is then simply encoded as the average sure amount around the switching point (using an interval regression between the two bounding values does not affect our results in any way). In an initial example, subjects were first offered a choice between a given prospect or zero. They were then offered a choice between that same prospect and the highest outcome of the prospect. This procedure served to test the understanding of the tasks, and to nudge subjects towards switching from the lottery to the sure amount at some point in the list. If this procedure showed that a subject had not understood the task, enumerators were instructed to explain the task again. Single switching was not enforced after this in the elicitation process. However, only in very few instances did subjects want to switch back to the lottery after they had switched to the sure amount. We dropped the five observations where this happened from the data. At the end of the experiment, one choice task was chosen at random to count for real pay—the standard procedure in this kind of elicitation. Subjects also obtained a participation fee of 30 Birr, to compensate them for their time and ensure that nobody left empty-handed.

Descriptive insights on risk tolerance

We start from discussing the stability of risk tolerance over time. To measure the consistency of our measures, we can look at the test-retest reliability—the correlation between identical measures taken in the same year and the same experimental session. We only included such measures in 2017, when we repeated three of the original 14 prospects at the end of the experiment. For the prospect offering the PPP-equivalent of €20 or else 0 with $p = 0.5$, we find a test-retest reliability of 0.713. The other two retests, for the same outcomes obtaining with $p = 0.125$ and $p = 0.875$, respectively, we find correlations of 0.788 and 0.759. These values are close to those observed with students in the

West. Brooks, Peters and Zank (2013) report that about 70% to 73% of repeated choices matched the initial choices, and provide a short review indicating similar findings by others. Abdellaoui, Kemel, Panin and Vieider (2019) report correlations between 0.75 and 0.8 in an experiment using high stakes with Western students. We thus conclude that risk tolerance is reasonably stable in the very short run, indicating the meaningfulness of our measures of risk tolerance.

Table 1: Correlations of risk-tolerance over time

	2013	2015	2017
2013	1		
2015	0.306 ($p < 0.001$)	1	
2017	0.213 ($p < 0.001$)	0.265 ($p < 0.001$)	1

Correlation coefficieints indicate Spearman rank order correlations between means of risk-tolerance per year.

We next examine the inter-temporal correlation of our measures across the years of the survey. Table 1 shows the correlations between the average measures of risk tolerance per year (i.e. taking the average CE across all tasks). The Spearman correlation between the average measure in 2013 and the average CE in 2015 falls slightly above 0.3, with the correlation between 2015 and 2017 falling somewhat below that value. Correlations between 2013 and 2017—with four years intervening between the measurements—are lower, at 0.21. These correlations, while certainly not large, fall towards the upper quartile of the inter-temporal correlations discussed by Chuang and Schechter (2015).[10]

This brings us to a description of the levels of risk tolerance. Figure 5 shows risk tolerance for 50-50 prospects offering a prize of x or else 0. (Figure S4 depicts risk tolerance across probability levels, and indicates the typical pattern of relative risk tolerance declining in probability; see Fehr-Duda and Epper, 2012; L'Haridon and Vieider, 2019). The measure depicted in panel 5(a), shows a clear pattern of decreasing relative risk

[10]The raw correlations just discussed do not yet take the noisiness of the measures into account. The test-retest reliability discussed above allows us to correct the raw correlations described above for attenuation. Denote the raw correlation by $\rho(x, y)$, where x and y can designate different measurements, either using the same task at different periods in time, or using different tasks within the same session. Let x' and y' indicate re-tests—measurements using an identical tasks within the same session. The correlation coefficient corrected for attenuation will now be:

$$\hat{\rho}(x, y) = \frac{\rho(x, y)}{\sqrt{\rho(x, x')\rho(y, y')}} \tag{3}$$

If we thus correct the correlation coefficients reported, the true inter-temporal correlations increase to about 0.4. Notice that these calculations assume that the test-retest reliability is constant across time, since we can only use the values obtained in 2017.

Figure 5: Risk tolerance in 50-50 prospects

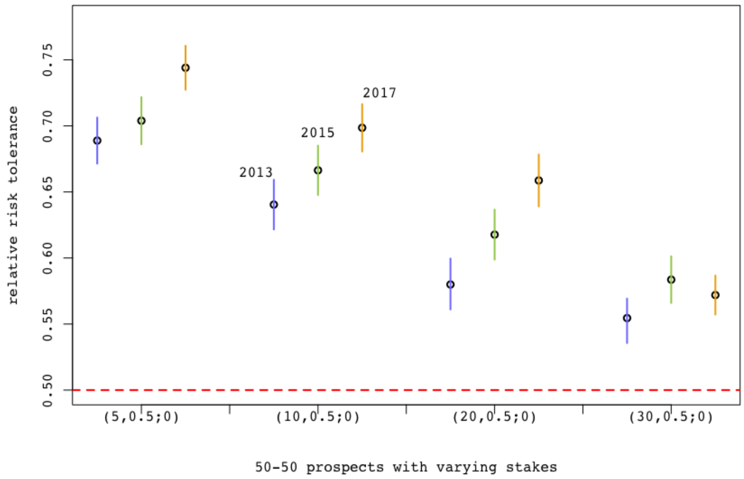

(a) Relative risk tolerance

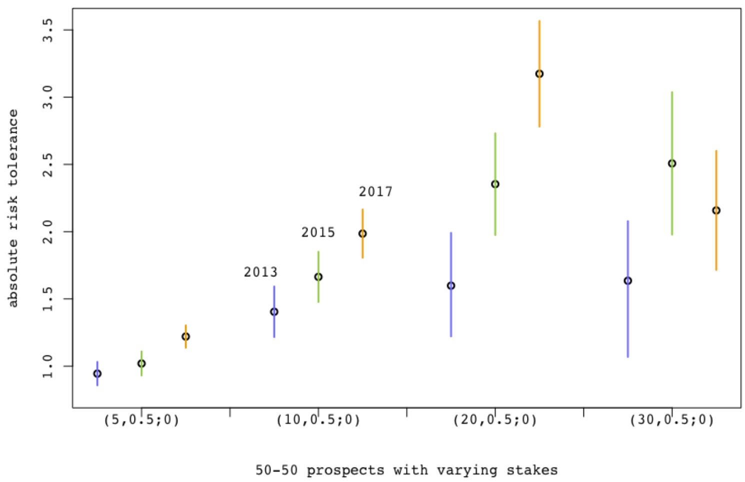

(b) Absolute risk tolerance

The figure shows non-parametric indices of risk tolerance for 50-50 prospects offering a prize x or else 0. We focus on tests of utility as stakes change for a given probability of 0.5, since such tests are valid not only for expected utility theory, but also generalizations such as prospect theory. Relative risk tolerance is defined as $\frac{ce-y}{x-y}$, where x is the higher and y the lower outcome of the prospect. This constitutes an index of risk tolerance *relative* to the outcome range (from a theoretical point of view, one can think of the measure as a decision weight under dual-expected utility; Yaari, 1987). The dashed horizontal line in panel (a) indicates risk neutrality. Absolute risk tolerance is defined as $ce - ev$ (i.e., a negative risk premium), where ev designates the expected value of the prospect. The pattern we find is one of increasing relative risk aversion (IRRA) and constant absolute risk aversion (CARA).

tolerance, or equivalently, increasing relative risk aversion (IRRA). This is highly significant, with each subsequent measure as stakes increase resulting in lower levels of relative risk tolerance, corresponding to the typical pattern found in the literature (Holt and Laury, 2002; Fehr-Duda, Bruhin, Epper and Schubert, 2010; Bouchouicha and Vieider, 2017). Panel 5(b) shows a measure of *absolute* risk tolerance for the same prospects. The pattern here is less clear. While there appears to be a tendency toward increasing absolute risk tolerance, or decreasing absolute risk aversion (DARA), when passing from the smallest prize to the next larger one, this pattern subsides as prizes increase further. Fitting a parametric model to the data, we find that an exponential utility function, incorporating IRRA and constant absolute risk aversion (CARA), fits the data significantly better than a logarithmic function (Scholten and Read, 2014), combining IRRA with DARA (WAIC of 238,664.5 versus 238,740.2 in favour of the exponential function, giving it a weight of 1;[11] this holds both in an expected utility framework, and in generalizations allowing for nonlinear probability weighting—see section S3 for details).

Table 2: Woreda-level descriptives of risk tolerance and environmental characteristics

Woreda	mean CE	SD CE	hist. rain	hist. rain SD	altitude	altitude SD
Atsbi Wonberta	11.38	6.97	323.21	100.59	2773.97	49.67
Bambasi	15.42	7.59	706.60	67.54	1425.16	14.43
Bereh Aleltu	14.83	7.53	685.29	80.16	2503.05	66.77
Bichena	12.58	7.90	616.45	100.11	2405.28	66.85
Chilga	12.33	6.61	751.72	84.66	2178.99	50.40
Debark	13.78	7.06	801.85	107.67	2826.96	73.29
Endamehoni	11.76	6.96	432.42	146.45	2445.90	55.40
Gesha Daka	17.48	7.41	722.58	58.01	2200.65	110.45
Gimbi	13.52	7.27	998.21	90.50	1807.23	39.00
Haru	14.67	7.11	1037.59	99.24	1852.28	134.35
Hawzein	11.58	7.13	368.10	82.54	2193.17	61.17
Hidabu Abote	14.88	7.46	703.60	81.41	2331.65	201.80
Kersa	15.05	8.02	736.02	79.08	1812.77	37.57
Libo Kemkem	12.92	6.96	803.17	111.01	1846.18	22.02
Limu	14.88	8.05	903.68	55.60	2213.83	45.25
Nunu Kumba	15.15	7.69	943.51	89.95	2302.51	93.91
Quarit	12.60	8.19	805.42	76.88	2189.72	71.10
Sirba Abay	16.20	7.56	662.50	70.46	866.42	29.45
Wogera	13.75	7.10	851.99	109.32	2844.13	71.77
Wonbera	12.96	8.11	676.49	67.25	2392.69	85.94
Total	13.83	7.42	725.25	88.85	2178.62	67.83

All numbers reported in the table represent Woreda-level averages. CEs are measures in PPP-Euros. Altitude is measured by GPS, and reported in metres above the sea. The variable 'hist. rain' represents the average historical rainfall in the Woreda during the Meher season between 1981 and 2010. The variable 'hist. rain SD' represents the standard deviations between yearly Meher seasons over the same period.

Finally, we present some district-level descriptives. Table 2 shows the average CE by Woreda over all tasks across the three waves of data collection, jointly with some other mean characteristics of the Woreda, such as the average altitude above the sea and its standard deviation; the mean historical rainfall levels during the Meher season and their

[11] *WAIC* stands for Watanabe-Akaike Information Criterion; see e.g. Gelman, Hwang and Vehtari (2014a) or McElreath (2016), chapter 6, for a discussion. The weight attributed to a model can be intuitively interpreted as a probability that the specific model is the best amongst the tested models.

standard deviations; and the average distance from the capital. The mean CE varies considerably across Woredas, ranging from a low of 11.38 in Atsbi Wonberta to a high of 17.48 in Gesha Daka. There is substantial variability of CEs within each Woreda.

Figure 6: Worelda level correlations of risk tolerance with environmental characteristics

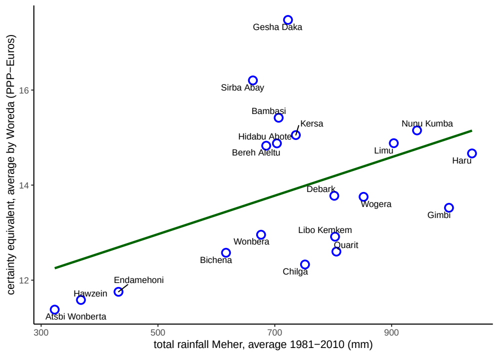

(a) RT and average historical rainfall

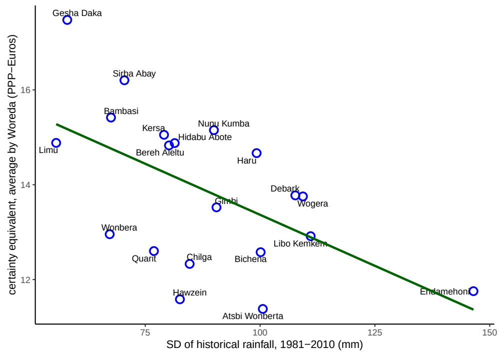

(b) RT and historical rainfall variation

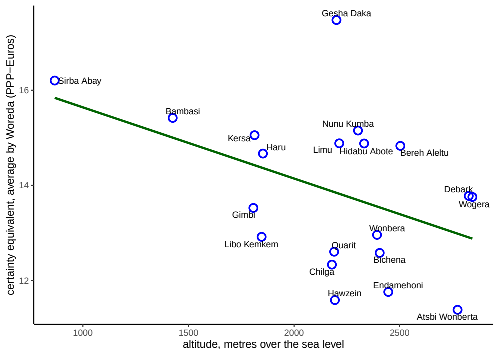

(c) RT and altitude

.

Figure 6 shows raw correlations between the Woreda-level average CEs and some of the main rainfall and geographic characteristics in the table, also averaged at the Woreda level. Panel 6(a) reveals a positive correlation between risk tolerance and the historical rainfall level during the Meher ($r = 0.469$, $p = 0.037$, Pearson correlation). Panel 6(b) indicates that Woredas with higher variability in historical rainfall tend to be less risk tolerant ($r = -0.562$, $p = 0.009$). Risk tolerance is also lower in Woredas at higher altitudes (panel 6(c); $r = -0.436$, $p = 0.055$). Of course, these graphs only give us a first indication of these correlations, since measures are aggregated at the Woreda level and we only examine one characteristic at a time. We will return to these issues in section 5.3.

4 Econometric Analysis

We analyze our data by means of a within estimator in combination with an error structure that explicitly models our sampling framework. As famously shown by Mundlak (1978), in balanced panels the within estimator yields results identical to those of individual fixed effects implemented through dummy variables. This implementation thus allows us to rigorously document the effects of time-changing characteristics, and to document the effects of time-invariant environmental features in one and the same regression. We augment the usual error structure by two additional error terms. An error term subordinate to the individual-year residual allows us to use several measurements per individual and year. An additional error at the level of the Woreda explicitly allows for spatial covariation in the individual-level residuals. This has the effect of clustering the errors at the level of stratification (Cameron and Miller, 2015, p. 318). It further has substantive implications for the inferences we draw on the effect of time-invariant environmental characteristics.

Our dependent variable consists of a measure of relative risk tolerance, $rt = \frac{ce-y}{x-y}$, where ce indicates the certainty equivalent, and x and y are the high and low outcome of the prospect, respectively. This is a measure of risk tolerance relative to the outcome range of the prospect, and can be thought of as a decision weight in the context of Yaari's (1987) Dual Expected Utility model. This measure is convenient in terms of interpretation, and should not distract from the fact that our setup is model-free. We explicitly model heterogeneity between prospects, i, for a given subject, s, sampled from a district or Woreda, w, in a given year, t:

$$rt_{wsti} = \alpha_t + (S_{st} - \overline{S}_s)\beta_1 + \overline{S}_s\beta_2 + X_s\gamma + \omega_w + \nu_{ws} + \eta_{wst} + \epsilon_{wsti}. \tag{4}$$

The model consists of a regression part, $\alpha_t + (S_{st} - \overline{S}_s)\beta_1 + \overline{S}_s\beta_2 + X_s\gamma$, and of a composite error term, $\omega_w + \nu_{ws} + \eta_{wst} + \epsilon_{wsti}$. The matrix $\overline{S}_s = \frac{1}{T}\sum_{t=1}^{T}(S_{st})$ contains the intertemporal means of our time-changing rainfall shocks, d_{at}, and time-changing controls, so that $(S_{st} - \overline{S}_s)$ contains per period deviations from the intertemporal means. Our primary interest is for the coefficient vector β_1, which contains the within-estimates, i.e. it captures how preferences change over time for the same individuals following shocks. The coefficients β_2 capture the between effects, i.e. the effects of the average shocks across the three years in the cross-section, which do not warrant a causal interpretation (see below). The constant α is subscripted by t to indicate that we allow for time fixed effects. The matrix X_s contains fixed characteristics of the environment of a given subject s, with γ the vector of coefficients. One of the great advantages of the within estimator we use is indeed that we can document the effect of the environmental characteristics in X in the same regression used to document the effect of time-changing characteristics, without compromising on the identification of the time-changing effects.

This brings us to the composite error term, $\omega_w + \nu_{ws} + \eta_{wst} + \epsilon_{wsti}$. The part $\nu_{ws} + \eta_{wst}$ constitutes the standard error structure used in conjunction with the within estimator with one observation per period in the data, with $\eta_{wst} \sim \mathcal{N}(0, \sigma_t^2)$ and $\nu_{ws} \sim \mathcal{N}(0, \sigma_s^2)$, where σ_t and σ_s indicate the standard deviations at the time and subject level respectively (see e.g. Wooldridge, 2015, section 14-2a and onwards , or Allison, 2009, p. 23). We augment this basic error structure with two additional terms. The term $\epsilon_{wsti} \sim \mathcal{N}(0, \sigma_r^2)$ represents residuals at the level of the measurement, i, with σ_r^2 the residual variance. This allows us to use all measurements obtained for a given individual in a given year, rather than having to average over the different measurements. The term $\omega_w \sim \mathcal{N}(0, \sigma_w^2)$ represents an additional hierarchy at the level of the Woreda, with σ_w^2 the Woreda-level variance. We insert this term to explicitly model our sampling framework. This clusters the standard errors at the Woreda level, thus providing conservative estimates of the standard errors (Cameron and Miller, 2015; Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014b).

For a causal interpretation of the time changing effects in S_{it}, we assume η_{wst} to be independent of the rainfall shocks in S_{st}, conditional on any time-changing controls. This corresponds to the standard assumption about the within estimator (as well as about the fixed effects model in general), and is usually considered to be a relatively mild assumption. The time average of the shocks in $\overline{S}_s$ are needed to define our within estimator, and do not warrant a causal interpretation. For any of the time-invariant regressors in X_s to have a causal interpretation, we need to assume its conditional independence from the composite error term $\omega_w + \nu_{ws}$. This is a much stronger assumption, which is not generally warranted. Therefore, the causal interpretation we suggest for some of the longterm environmental characteristics in X_s is based on supplementary arguments of plausibility and coherence with the time-changing effects and the predictions of our model, rather than being established purely econometrically.

While the clustering is an important second-order effect of the Woreda-level error term, ω_w, it also has substantive implications for the analysis of time-invariant environmental characteristics, since it nests subject-level residuals, ν_{ws}, in Woreda-level residuals, ω_w, instead of the intercept, α_t. Figure 7 plots the idiosyncratic preferences estimated based on equation 4 empty of covariates (i.e. the residuals $\alpha + \omega_w + \nu_{ws}$), against the residuals obtained from an otherwise equivalent model that drops the error term ω_w. The residuals estimated in the model without ω_w are pooled towards a global mean given by the intercept α, shown as a dashed vertical line in the graph. This means, inter alia, that estimates beyond 0.8 are entirely discounted as being unlikely (see Gelman and Pardoe, 2006, for technical details). This does not happen when the Woreda-level residuals are introduced, simply because respondents in the entire Woreda exhibit similarly high levels of risk tolerance. The individual-level residuals, ν_{ws}, are now pooled towards the Woreda-level residuals, ω_w, indicated by the dashed horizontal lines. This explicitly models the expectation that individuals within one and the same Woreda will be more similar to each other than individuals in different Woredas. It is important to note that this modeling choice follows from our stratification strategy, which thus also provides the justification for it.

Figure 7: Plot of idiosyncratic risk tolerance in the model with and without Woreda-level residuals

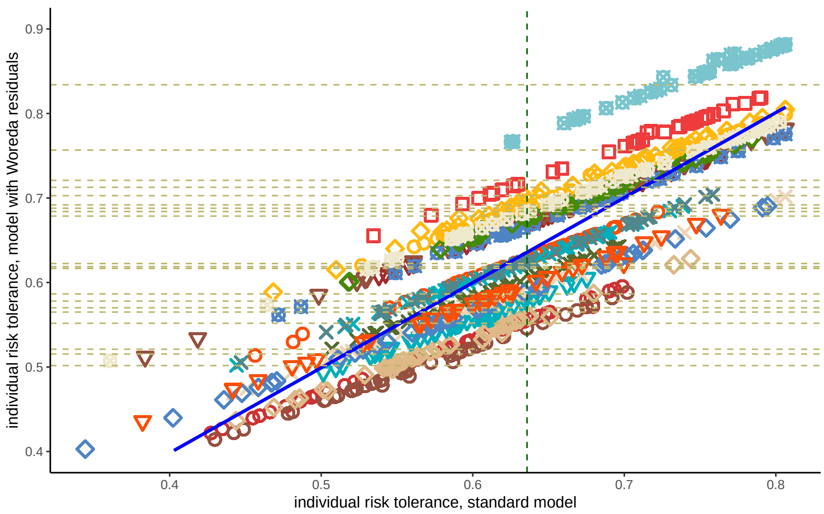

The hierarchical model further allows us to document the co-variation in preferences across the various levels of analysis (Gelman and Hill, 2006; McElreath, 2016). A useful metric to achieve this is the intra-class correlation (ICC), defined as the proportion of variance captured at a given level relative to the overall level of variance. For instance, the ICC across time is defined as the level of variance across time relative to the sum of all four variance terms, $\rho(t) = \frac{\sigma_t^2}{\sigma_r^2+\sigma_t^2+\sigma_s^2+\sigma_w^2}$. An intuitive interpretation of the ICC of a given level is that it captures the correlation between randomly drawn observations at that level (see Snijders and Bosker, 2012, section 3.3), in this particular example, the resemblance of two measurements obtained from the same individual in a given year. The ICC at the Woreda level obtained while abstracting from the variation across measurements and across time, $\rho(w|\sigma_r = 0, \sigma_t = 0) = \frac{\sigma_w^2}{\sigma_s^2+\sigma_w^2}$, then serves to quantify the geographic correlation of preferences within one and the same Woreda. The latter plays an important role in our data, given the geographical similarity of environmental circumstances and the ensuing spatial correlation in preference patterns.

5 Results

Co-variation in preferences across time and space

Some interesting insights can be obtained from a variance decomposition of our measures. The highest levels of variance registers at the residual level, with $\sigma_r^2 = 0.043$, and over time, with $\sigma_t^2 = 0.043$. Variance across subjects and Woredas is considerably lower at $\sigma_s^2 = 0.007$ and $\sigma_w^2 = 0.008$, respectively. At the level of measurements in a single time period, we find an ICC of $\rho(t) = 0.42$, indicating that two random measurements taken for the same random individual in a random year show a correlation of 0.42 on average. This is consistent with the large variation in preferences across prospects we find, and corresponds to typical correlations observed in experiments with students.

There are two further measures we want to look at. The first is the ICC at the individual level, ie. aggregating across individuals and Woredas while assuming the residual variance to be equal to 0, $\rho(s, w|\sigma_r^2 \equiv 0)$. We can interpret this as a correlation of the means per measurement period over time, which we quantify at $\rho(s, w|\sigma_r^2 \equiv 0) = 0.261$. This figure indeed corresponds closely to the average raw correlations between the measures, shown in table 1. Put differently, 74% of the total variation between aggregated measures takes place across time. The upshot of this finding is that it should come as no surprise that cross-sectional analysis performs poorly at identifying correlates of risk tolerance—cross-sections measured in different years do look very different from each other, and should thus be expected to yield different results in regressions.

The final comparison we are interested in concerns the similarity of individuals within the same Woreda, i.e. the spatial co-variation of preferences at the level of sampling stratification. It is most meaningful to examine this measure while setting the inter-temporal variance to 0, $\sigma_t^2 \equiv 0$. What is left are then the risk preferences across individuals once inter-temporal fluctuations are averaged out, which we will refer to as *idiosyncratic preferences*. Under this assumption we find an ICC of $\rho(w|\sigma_r \equiv 0, \sigma_t \equiv 0) = 0.568$. In other words, the idiosyncratic risk tolerance of two random individuals from a random Woreda shows a correlation of 0.568. To put this figure into perspective, we can compare it to the variance captured at the country level in cross-country comparisons of risk tolerance. Falk et al. (2018), Bouchouicha and Vieider (2019), and L'Haridon and Vieider (2019) present concordant evidence that the variation captured at the country level is about 10% of the total (none of these studies quantified variation across regions within countries). This suggests that the environment of a respondent plays a major role in the determination of her risk preferences. It also illustrates the importance of explicitly accounting for spacial co-variation in preferences at the Woreda level in the econometric analysis, since neglecting such co-variation would result in biased estimations.

5.1 Within- versus between-effects of rainfall shocks

We next detail the effects obtaining from a longitudinal versus cross-sectional examination of our data. In particular, we deploy the standard within-between estimator suggested by Wooldridge (2015) as an alternative to the Hausman test—if the within and between estimators coincide, then the conclusions from a longitudinal and cross-sectional analysis of the data will yield the same results, and one could apply a random effects model. If the two estimators differ, then the residuals are not independent from the predictors, imposing the use of a within or fixed effects estimator.

Figure 8 shows the different estimators, separately for positive rainfall deviations from historical average values (floods) and for negative rainfall deviations (droughts). The within estimator clearly shows a negative effect of rainfall shocks on risk tolerance of both droughts and floods. In both cases, the between estimator significantly differs from the within estimator. For rain shortfalls, it indeed goes in the opposite direction, which may lead one to (wrongly) conclude that rainfall shortfalls increase risk tolerance based on the cross-sectional evidence. This shows the dangers of drawing inferences from cross-sections even in contexts where the shocks are exogenous, and random ex

Figure 8: Within versus between effects of rainfall shocks

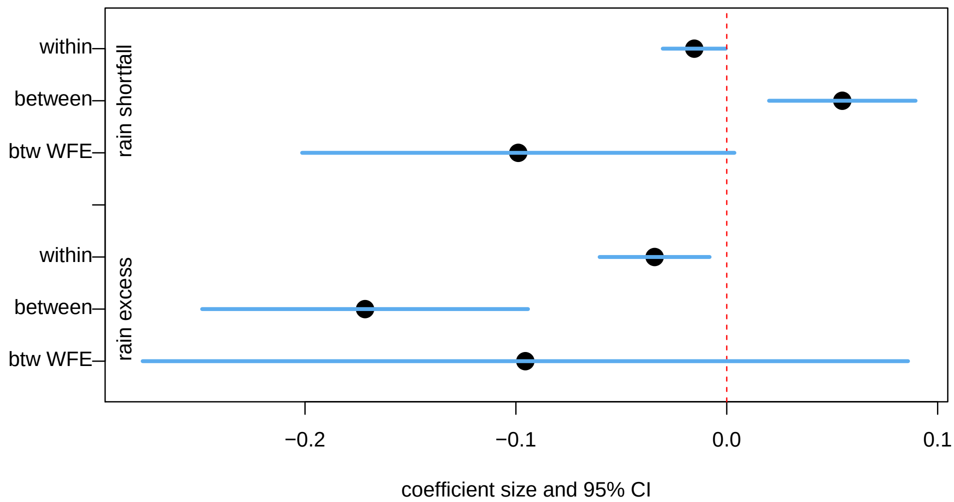

Graph of regression coefficients with 95% confidence interval. The shorthand 'btw WFE' stands for the between estimator with Woreda fixed effects.

ante, as they are in our case. Once we add Woreda fixed effects the standard errors of the between estimator explode. This is unsurprising, since very little rainfall variation is observed within Woredas in any given year. While the between effects are no longer different from the within effects, they are also no longer different from zero. This would again yield very different—and ultimately misleading—inferences.

An implication of these differences in the within and between estimators is that cross-sectional analysis of our data would necessarily result in biased conclusions. This insight acquires special significance in our context, since the great majority of previous studies investigating the effect of shocks on preferences have used cross-sectional data. One may also wonder about the underlying reason for these divergent effects. After all, our rainfall shocks are exogenous and random *ex ante*. Potential confusion may arise from different usages of the term 'exogenous'. Applied researchers often take that term to indicate that the predictor of interest is unaffected by any actions the study participants may undertake. This excludes reverse causality, and our measures unambiguously fulfil that criterion. In econometrics, however, the term 'exogenous' is used to signify 'uncorrelated with the error term'. Exogeneity of the first type is thus not sufficient to guarantee exogeneity of the econometric type. The significantly positive between estimator for rain shortfalls then just indicates that these shocks by chance prevalently hit relatively risk tolerant districts during our study period (and vice versa for floods).

Consider this highly stylized example. Assume a country is divided into risk seekers, living in the south, and risk averters, living in the north. Further assume that shocks truly have no effect on risk preferences. If the particular shocks we observe in a given year hit prevalently in the south, we may conclude from cross-sectional analysis that shocks *increase* risk tolerance. If, on the other hand, the shocks we observe hit prevalently in the north, we might be tempted to conclude that shocks *decrease* risk tolerance. Even if exogeneity and randomness hold for our measures of shocks, the conclusions drawn would be mistaken in both cases, since preferences were not uniformly distributed across regions ex ante. Unless we observe and correctly measure differences in preferences pre-existing the shocks themselves, these differences will be subsumed in an error term that is now correlated with the predictor variables. Given the 'exogenous'—in the applied rather than econometric usage of the term—nature of rainfall, it is all too easy to mistake such spurious correlations for causal effects. While this is a highly stylized example relying on there being only two regions with different preferences, the example readily generalizes to much larger numbers of distinct regions. Only once the number of regions goes to infinity while the shocks stay random can we be sure that this problem will no longer occur—a case approaching the gold standard of individual randomization.

5.2 Rainfall shocks reduce risk tolerance

We now describe the effects of shocks on risk tolerance over time. Table 3 shows the regressions of risk-tolerance on rainfall deviations (placebo regressions using the minor rains can be found in section S4). We present reduced form regressions, regressing risk tolerance directly on rainfall deviations and not including any economic controls. While our model postulates that the effect ought to pass through consumption, the latter is likely to be endogenously determined, thus raising the spectre of reverse causality from risk-tolerance to consumption. Section S2 in the supplementary materials presents a detailed, parcel-wise analysis of agricultural yields, and shows that the effects of rainfall shocks on yields provide additional evidence for the coherence of the effects presented below with the mechanism postulated by our model.

The effects shown in table 3 are difficult to interpret due to the polynomial expressions. To overcome this shortcoming, figure 9 shows the total effect of rainfall shortfalls.

Table 3: Regression of risk-tolerance on rainfall shocks (within effects)

dep. var: risk-tolerance	(1)	(2)	(3)	(4)	(5)	(6)	(7)
rain shortfall	-0.011*	-0.087***	-0.108***	-0.128***	-0.209***	-0.207***	-0.218***
	(0.006)	(0.019)	(0.022)	(0.030)	(0.039)	(0.039)	(0.044)
rain excess	-0.032**	-0.215***	-0.218***	-0.221***	-0.291***	-0.287***	-0.402***
	(0.013)	(0.043)	(0.045)	(0.046)	(0.061)	(0.061)	(0.071)
rain shortfall sq.		0.028***	0.038***	0.036***	0.058***	0.057***	0.059***
		(0.007)	(0.008)	(0.009)	(0.012)	(0.012)	(0.012)
rain excess sq.		0.132***	0.139***	0.147***	0.200***	0.198***	0.265***
		(0.032)	(0.033)	(0.034)	(0.041)	(0.041)	(0.045)
rain shortfall lag 1			-0.005	-0.057***	-0.037*	-0.038*	-0.044**
			(0.014)	(0.020)	(0.021)	(0.021)	(0.021)
rain excess lag 1			0.028	-0.038	-0.023	-0.024	0.011
			(0.033)	(0.038)	(0.040)	(0.040)	(0.045)
rain shortfall lag 1 sq.			0.008**	0.015***	0.007	0.008	0.010**
			(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
rain excess lag 1 sq.			-0.005	0.023	0.014	0.015	0.013
			(0.016)	(0.018)	(0.019)	(0.019)	(0.022)
altitude * shortfall							-0.042***
							(0.011)
pos. skewness * shortfall							-0.015**
							(0.008)
animals * shortfall							0.042**
							(0.021)
animals * excess							0.075*
							(0.039)
year fixed effects	NO	NO	NO	YES	YES	YES	YES
rain dev. lag 2	NO	NO	NO	NO	YES	YES	YES
controls	NO	NO	NO	NO	NO	YES	YES
Nr. Households	906	906	906	906	906	906	906
Observations	39420	39420	39420	39420	39420	39420	39420
R^2 over time	0.004	0.018	0.027	0.034	0.045	0.048	0.083

All coefficients shown in the table refer to within effects. Between effects are not shown to save space. Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). The shorthand 'sq.' indicates the square of the deviations. Equation (1) only contains linear effects of positive and negative rainfall deviations. Equation (2) adds the squares of these measures. Subsequent regressions add the same rainfall variables lagged once (3), year fixed effects (4), and the lag 2 rainfall variables (5). Equation (6) includes the following time-varying controls: number of animals; access to irrigation (dummy); and non-farm income. Equation (7) contains interaction effects of rain shortfall and excess with standardized measures of altitude, positive skewness in the historical rainfall distribution, animals held, and the area of the land farmed (only significant effects reported because of space constraints). The inter-temporal R^2 is calculated as $1 - \sigma^2_{mo}/\sigma^2_{mj}$, where mo indicates the model empty of covariates, and mj refers to the model in the regression (see Snijders and Bosker, 2012, for a textbook treatment).

The grey lines represent the total sampling uncertainty surrounding the mean parameter estimates. For an average prospect offering either €20 PPP or else nothing, a rainfall shock of 1.5 SDs reduces the CE by €3.59 PPP—an economically sizeable effect.[12] For the largest shortfalls beyond 2 SDs, the curve starts bending slightly upward again. However, these effects are mostly driven by a handful of outliers experiencing a particularly severe drought, and should thus be interpreted with caution. Figure 10 shows the

[12] To calculate the economic effects for an average prospect, we use of the observation that our index of relative risk tolerance can be interpreted as a decision weight under dual expected utility. We can then simply calculate the change in the decision weight for a given rainfall shock from the coefficients in the table, and mutliply this change with the prize of the prospect to obtain the change in CE.

effect of rainfall lagged by one year. The effects are consistent with those of a drought immediately preceding the measurement, but weaker and less precisely identified.

Figure 9: The effect of rainfall shortfalls

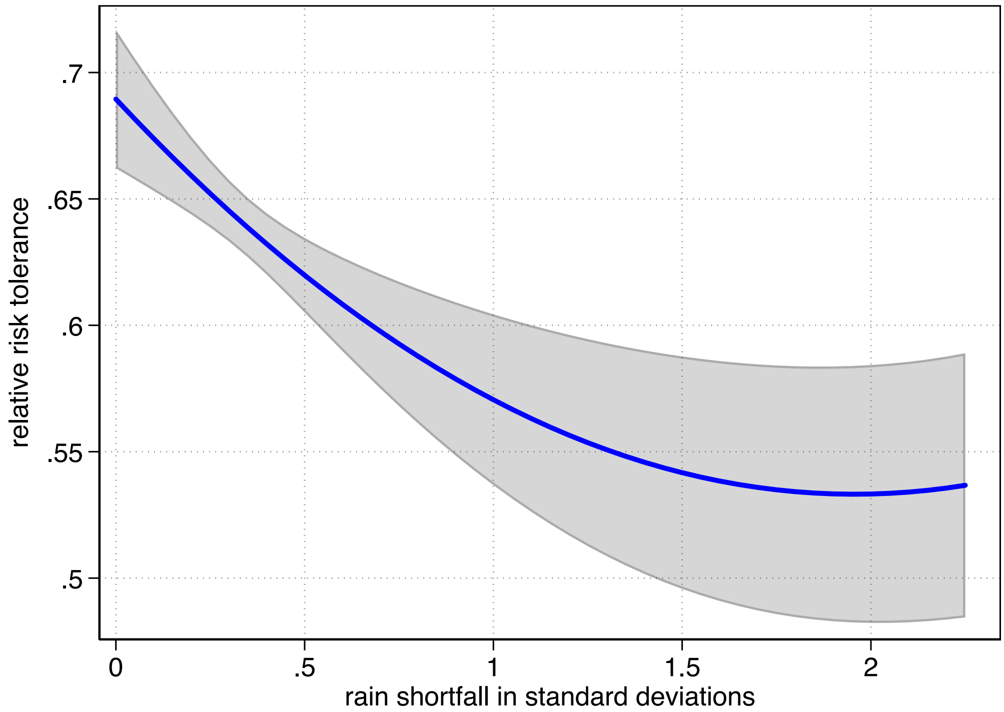

Graph of overall effect of rainfall shortfalls ('droughts') on relative risk tolerance. The solid blue line represents the mean effect of the polynomial. The grey area represents the the 95% prediction interval. Outliers based on the most extreme 5% in rainfall deviations are not shown in the graph.

Figure 12 shows the equivalent effects for excess rainfall. Two main differences stand out. One, the effect shows a much more pronounced U-shape than for shortfalls. Two, there is considerably more uncertainty surrounding the estimates. These differences can be traced partially to the fact that we simply observe much less excess rainfall than shortfalls during our study period. This phenomenon registers both at the extensive margin and at the intensive margin, with fewer subjects being affected by excess rainfall. An additional issue is that excess rainfall may affect people very differently depending on where and how it occurs. For instance, some additional rainfall may be beneficial if it occurs in places with low average rainfall. Furthermore, a given amount of additional rainfall distributed evenly over the Meher may have very different effects from the same rainfall occurring in one or two days during critical phases of the planting period. We thus argue that excess rainfall is inherently different—and more complex—than shortfalls.

The treatment effects differ by household and environmental characteristics, as high-

Figure 10: The effect of rainfall shortfalls, lag 1

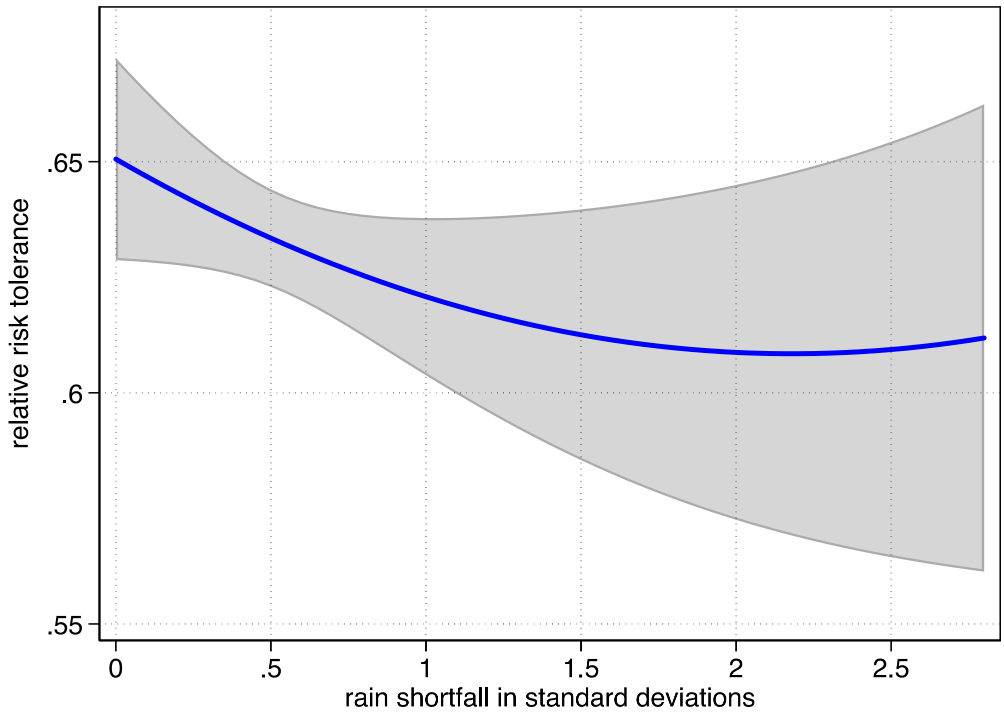

Graph of overall effect of rainfall excesses ('floods') on relative risk tolerance. The solid blue line represents the mean effect of the polynomial. The grey area represents the 95% prediction interval. Outliers based on the most extreme 5% of the sample are not shown.

lighted by the interaction effects in regression (7). In particular, rain shortfalls have a stronger negative impact at high altitude, while excess rainfall has less of a negative impact at high altitude. Households with animal holdings tend to generally react less to rainfall shocks, which is consistent with animal herding reacting less strongly to rainfall variations, and especially to excess rainfall, but also with animals being used as a buffer stock to be used against rainfall shocks. Finally, we find that in places having a positive skewness in historical rainfall a shortfall in rain has a considerably larger impact than in places with negative skewness in historical rain (the opposite holds true for excess rainfall, but it is very imprecisely estimated). This supports our intuition that what counts is the deviation in a given rainfall realization from to the historical distribution.

Before concluding this section, we examine the amount of inter-temporal variance explained by our model. Regression (6) in table 3 explains 4.8% of the inter-temporal variance. Adding the heterogenous treatment effects in regression (7) brings this figure to 8.3%. Much of the variation in preferences over time thus remains unexplained. This suggests that aggregating over the unexplained inter-temporal fluctuation in risk

Figure 11: The effect of rainfall excesses

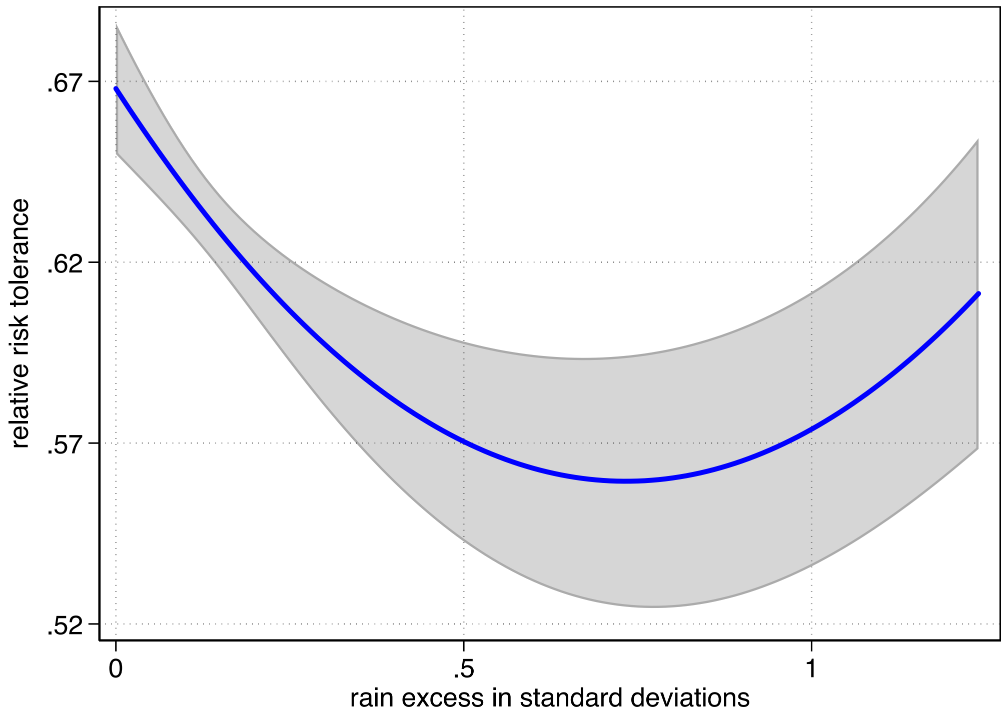

Graph of overall effect of rainfall excesses ('floods') on relative risk tolerance. The solid blue line represents the mean effect of the polynomial. The grey area represent the 95% prediction interval. Outliers based on the most extreme 5% of the sample are not shown.

tolerance will be at least as important as filtering out any effects of observable variables when it comes to stabilizing the estimates of idiosyncratic risk tolerance.

5.3 Environmental determinants of risk tolerance

We now examine correlates of idiosyncratic risk tolerance. We encode idiosyncratic risk tolerance as the means of the individual-level intercepts, defined as the sum of the individual-level and Woreda-level residuals, $\alpha_{2013}+\omega_w+\nu_{ws}$. This means that we capture individual-level preferences purified of the effects of shocks as captured in our model, and averaged over time. The basis for our analysis of idiosyncratic risk tolerance is formed by regression (6) in table 3. Using regression (7) instead does not affect our conclusions in any substantive way, but the effects become more difficult to interpret due to the presence of interaction terms including environmental characteristics in that regression.

We start from a graphical analysis of the raw correlations between idiosyncratic preferences and environmental characteristics. Figure 12(a) shows the correlation between the mean (panel 12(a)) and the standard deviation (panel 12(b)) of historical rainfall in

Figure 12: Coorelations between historical rainfall indicators and idiosyncratic risk tolerance

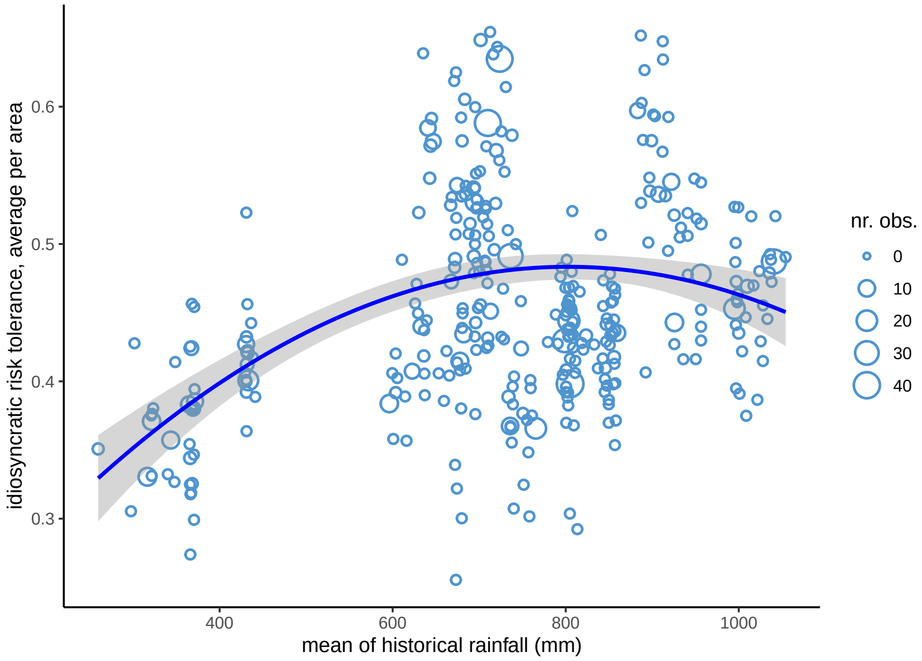

(a) Historical mean of rainfall and risk tolerance

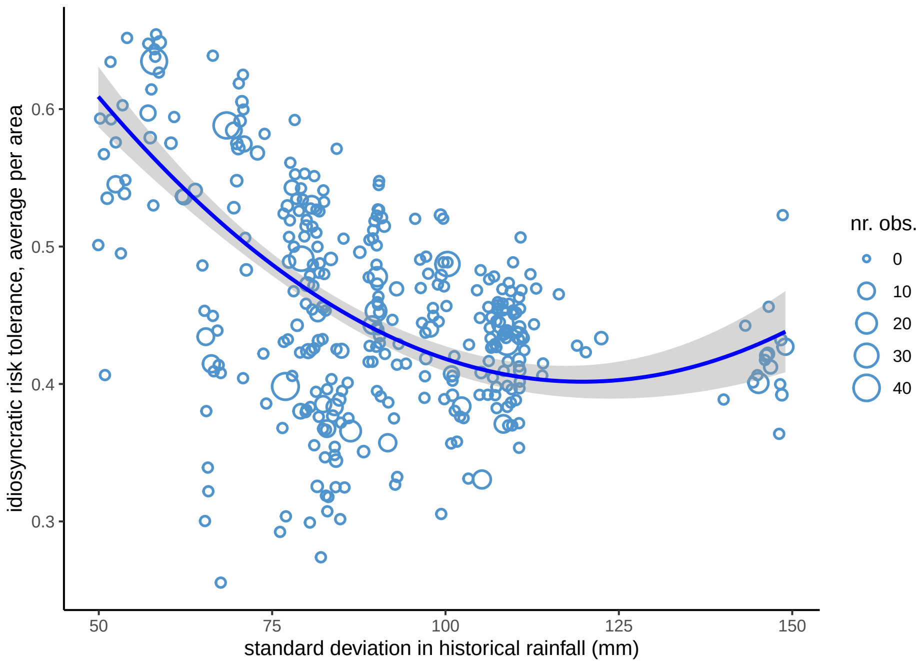

(b) Historical SD of rainfall and risk tolerance

Graph of idiosyncratic risk tolerance against the mean and standard deviation (SD) in historical rainfall. Since historical rainfall data differ by rainfall areas rather than individuals, we show idiosyncratic risk tolerance aggregated by area and weighed by the number of observations contained in each point.

a given area and idiosyncratic risk tolerance. Risk tolerance is increasing in the mean of historical rainfall. At the same time, risk tolerance decreases strongly in rainfall SD, at

Figure 13: Correlations between altitude and idiosyncratic risk tolerance

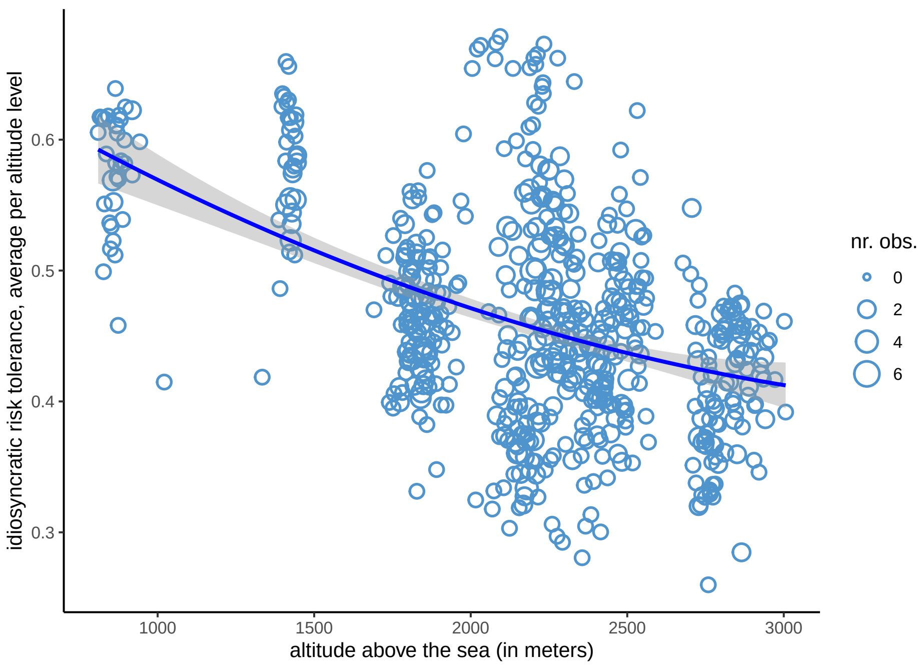

Graph of idiosyncratic risk tolerance against altitude as measured by the GPS coordinates. Idiosyncratic risk preferences is aggregated by the independent observations at exactly the same altitude, and weighted by the number of observations.

a decreasing rate. In addition to rainfall levels, geographical features may also impact preferences. Given the mountainous geography of Ethiopia, one would expect that the altitude at which a farm is located will impact productivity, since temperatures decline quickly with altitude, and because higher altitudes create vulnerability because of exposure to wind, quick draining of soils, difficulty in ploughing due to the steepness of fields, etc. (Diamond, 2005). Figure 13 shows the correlation between idiosyncratic risk tolerance and altitude. Risk tolerance steeply declines with altitude, as expected.

We now enter all of these measures jointly into a regression framework. The regressions are shown in table 4. The regressions simply add environmental characteristics to regression (6) in table 3, using the specification set out in equation 4. The time-changing part shown in table 3 is not displayed again in order to save space, but remains unaffected. Regression (1) includes only the historical rainfall mean from 1981 to 2010. Regression (2) adds the historical standard deviation for the same period. Regression (3) further adds the square of the standard deviation, and regression (4) the altitude,. All effects have the signs we would expect based on the figures above, and all of them are statistically significant. Regression (5) further controls for the age and gender of the respondent, and for the land area farmed by the households. None of these variables are significant, and we omit them from the table.

Table 4: Risk-tolerance and environmental factors

	(1)	(2)	(3)	(4)	(5)
mean of historical rain	0.027**	0.025**	0.023**	0.017**	0.017**
	(0.012)	(0.011)	(0.009)	(0.008)	(0.008)
SD historical rain		-0.216***	-1.407***	-1.415***	-1.415***
		(0.078)	(0.423)	(0.379)	(0.379)
SD hist. rain sq.			0.614***	0.643***	0.643***
			(0.215)	(0.192)	(0.192)
altitude				-0.050*	-0.050*
				(0.027)	(0.027)
animals (intertemp. mean)	0.040*	0.041*	0.039*	0.043*	0.043*
	(0.024)	(0.024)	(0.023)	(0.024)	(0.024)
Observations	39420	39420	39420	39420	39420
controls	NO	NO	NO	NO	YES
Nr. respondents	906	906	906	906	906
Observations	39420	39420	39420	39420	39420
R^2 across respondents	0.137	0.260	0.368	0.409	0.409

The results reported are based on equation (6) in table 3, with the cross-sectional variables added to that specification. Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). The shorthand 'sq.' indicates the square of a variable. Regressions (1) to (4) control only for individual-level factors used as inter-temporal means of time-changing effects, including the inter-temporal mean of animals, the inter-temporal mean of irrigation, and the inter-temporal mean of non-farm income. Regression (5) introduces additional controls such as gender and age of the respondent, and area of land farmed.

The environmental variables used in the regressions explain a large part of the variance in preferences across respondents. The historical mean and SD in rainfall alone explain 26% of the variance in idiosyncratic risk tolerance, increasing to 36% when the square of the SD is added. Further adding altitude we reach a figure of 40% of the overall variance across respondents that is explained by characteristics of the respondents' environment. This figure is one order of magnitude larger than typical values reported in the literature (see footnote 2). The reason for this superior performance is twofold. One, our subjects are highly dependent on the environment for their subsistence, and no institutionalized safety nets exist. We would thus expect environmental effects to be particularly strong. Two, filtering out the effects of shocks and aggregating across time stabilizes our individual-level preference measures. Indeed, given the large variation over time, it is unsurprising that regressions using cross-sectional data perform poorly—a cross-section measured in one year looks very different from the same cross-section measured in a different year.

5.4 Robustness to selection effects

While the time-changing effects we documented using the panel structure permit an unambiguously causal interpretation under relatively mild assumptions, such an interpretation is not as straightforward for the long-term effects documented in the last section. The biggest challenge to such a causal account—albeit not the only one, see discussion in section 4—derives from systematic selection effects. Assume that all Woredas exhibit equal levels of risk tolerance initially, but that over time the most risk tolerant individuals leave the 'bad' Woredas, i.e. the Woredas at high altitude and with high levels of rainfall variation, and migrate to Woredas with more favourable conditions. As a result, Woredas with lower rainfall variation and at lower altitudes may then show higher levels of risk tolerance, just as we observe. We do not consider such an account to be plausible for several reasons. For one, migration between Woredas is difficult in Ethiopia, because of the already-mentioned restrictions to the ownership of land. Furthermore, the between-Woreda patterns we document line up nicely with the patterns we find over time for the same individuals using our within estimator, and for which a causal interpretation seems clearly warranted.

Luckily, we do not have to rely on plausibility alone. We have data on whether a participant was born in the village where he now lives, and on whether a participant was in that village at the age of 18 if he had migrated in his youth. Overall, 64% of respondents were born in the village where they now reside, and 74% lived in their current village of residence at 18 years old. Note that these figures are likely to overestimate migration between Woredas, since many of the respondents not born in their village of residence are likely to have migrated from neighbouring villages in the same Woreda.[13] That said, if indeed there is a significant number of migrants from a different Woreda—and if the hypothesis set out above holds true—then we ought to find migrants to exhibit above-average levels of risk tolerance compared to locals.

We test this hypothesis in the regressions reported in table 5. Regression (1) regresses risk tolerance on a dummy indicating whether the respondent was born in the village, including all the time-varying variables from regression (6)from table 3. Regression (4) does the same using the dummy indicating whether somebody was in the village

[13]It was unfortunately not possible to obtain data on between Woreda migration. Given that such migration is considered illegal in Ethiopia—and given that the mere question about the Woreda of origin would raise suspisions in the current climate of ethnical and regional tensions—our team of enumerators did not feel comfortable to even try and field this question.

Table 5: Effect of migration status on risk tolerance

	(1)	(2)	(3)	(4)	(5)	(6)
born in village	0.003	-0.000	0.073			
	(0.011)	(0.011)	(0.079)			
in village at 18 years				0.010	0.006	0.100
				(0.012)	(0.012)	(0.085)
mean of historical rain		0.017**	0.023***		0.017**	0.022**
		(0.008)	(0.009)		(0.008)	(0.009)
SD of historical rain		-1.415***	-1.420***		-1.418***	-1.488***
		(0.380)	(0.374)		(0.381)	(0.353)
SD of historical rain sq.		0.643***	0.635***		0.644***	0.684***
		(0.192)	(0.189)		(0.192)	(0.178)
altitude		-0.050*	-0.045		-0.049*	-0.193
		(0.027)	(0.030)		(0.027)	(0.118)
born in village * mean rain			-0.010*			
			(0.006)			
born in village * SD rain			0.032			
			(0.054)			
born in village * altitude			-0.012			
			(0.025)			
in village at 18 * mean rain						-0.008
						(0.006)
in village at 18 * SD rain						-0.008
						(0.059)
in village at 18 * altitude						-0.012
						(0.026)
time-changing variables	YES	YES	YES	YES	YES	YES
household-level controls	NO	YES	YES	NO	YES	YES
Nr. respondents	906	906	906	906	906	906
Observations	39420	39420	39420	39327	39327	39327

The results reported are based on equation (6) in table 3, with the cross-sectional variables added to that specification. Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). The shorthand 'sq.' indicates the square of a variable. Distances to Addis Ababa, the capital, is calculated using geodesic distance. Equations (1) to (5) control only for individual-level factors used as inter-temporal means of time-changing effects, including the inter-temporal mean of animals, the inter-temporal mean of irrigation, and the inter-temporal mean of non-farm income. Regression (6) introduces additional controls such as gender and age of the respondent, and area of land farmed.

at the age of 18 instead. Neither in one case nor the other does the dummy show a significant result, and the coefficients are very small. Regressions (2) and (5) add the environmental characteristics from table 4. Once again, people born in their current village of residence or having resided there since the age of 18 are in no way different from the rest of the population. The effects of the environmental characteristics documented above meanwhile do not change. We thus conclude that the hypothesis on selection effects set out above does not find any support in our data.

One may be concerned about attenuation in our results due to the potential mix of migrants from different Woredas and neighbouring villages in our data, even though the extremely small coefficients make it unlikely that our null result is driven purely by such attenuation. For instance, a more complex hypothesis would hold that there are different

Table 6: Risk-tolerance and environmental factors, respondents born in village of residence

	(1)	(2)	(3)	(4)	(5)
mean historical rain	0.030**	0.029**	0.025**	0.018*	0.019*
	(0.014)	(0.012)	(0.010)	(0.010)	(0.010)
SD historical rain		-0.248***	-1.328***	-1.305***	-1.344***
		(0.087)	(0.500)	(0.455)	(0.467)
SD historical rain sq.			0.556**	0.579**	0.592**
			(0.253)	(0.229)	(0.235)
altitude				-0.068**	-0.073**
				(0.034)	(0.035)
animals (intertemp. mean)	0.045	0.046	0.041	0.049	0.048
	(0.037)	(0.036)	(0.036)	(0.037)	(0.038)
controls	NO	NO	NO	NO	YES
Nr. respondents	573	573	573	573	573
Observations	24989	24989	24989	24989	24989
R^2 across respondents	0.098	0.228	0.311	0.352	0.347

The results reported are based on equation (6) in table 3, with the cross-sectional variables added to that specification. Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). The shorthand 'sq.' indicates the square of a variable. Equations (1) to (4) control only for individual-level factors used as inter-temporal means of time-changing effects, including the inter-temporal mean of animals, the inter-temporal mean of irrigation, and the inter-temporal mean of non-farm income. Regression (5) introduces additional controls such as gender and age of the respondent, and area of land farmed.

types of migration involved. Regular within-Woreda migration happens everywhere, and may not be linked to risk tolerance. However, at the same time, between-Woreda migration may be systematically linked to risk tolerance. The absence of aggregate effects in regressions (1), (2), (4), and (5) already suggests that—if such migration indeed exists—it would be relatively moderate compared to the first type. To nevertheless test this hypothesis more in depth, regressions (3) and (6) include interaction effects between the dummy indicating whether somebody was in the village since birth or by the age of 18, and the mean of historical rainfall, its standard deviation, and altitude. We find no support for the systematic selection hypothesis. The only significant term is the interaction between having been born in a village and the mean of historical rain, which is significant at the 10% level. While the direction of the effect is consistent with people not born in the village being more risk tolerant in areas characterized by higher historical rainfall levels, the size of the effect is tiny compared to the differences between Woredas we document, even before applying any statistical corrections for multiple testing.

As yet another additional test to rule out selection as an explanation for the patterns we find, we run our previous analysis based purely on respondents born in the villages where they currently reside (this is more conservative than using respondents who were in the village at 18, which does not change our conclusions). Table 6 replicates the analysis in table 4 based only on the sample of individuals born in the villages where they currently reside. The effects of the historical mean of rainfall, its variation, and altitude reported using the full sample are replicated. If anything, these effects result slightly reinforced, showing that migration and systematic selection effects cannot explain the effects we document. The one major exception to this general conclusion is constituted by the number of animals, which is no longer significant in this regression. While we cannot fully exclude migration in previous generations to have caused some selection, we consider strong effects of such long-ago selection effects implausible, given the substantial changeability of preferences over time we documented above. We thus conclude that selection effects are unlikely as an explanation for the long-term patterns we document, so that we are inclined to interpret them as plausibly causal.

6 Discussion

We have documented large differences in preferences across different environments. We are inclined to interpret these environmental effects as plausibly causal. For one, the effects of the historical mean and standard deviation are fully consistent with the changes over time we documented, for which a causal interpretation seems warranted. Both the time-changing effects and the cross-sectional effects of the long-term rainfall and environmental characteristics line up perfectly with the ones predicted by our model. Our data on agricultural yields further reinforce this narrative, by showing the negative impact of rainfall shocks on agricultural yields in our sample.

Selection effects do not provide a plausible explanation for our findings on long-run determinants of preferences. The Ethiopian constitution mandates that land belongs exclusively to the state. Committees allocate use rights to households. A key condition for the allocation of land is that the household members remain residents of the same *Kebele*, an administrative level subordinated to the Woreda (Rahmato, 2008). This allocation system creates a disincentive for migration, which is consistent with empirical evidence on economic migration in Ethiopia (De Brauw and Mueller, 2012). While rural to urban migration *does* exist, both the observed scale of migration and the urbanization rate in Ethiopia are too small to account for the large differences between Woredas we document. Finally, accounts based on selection preceding the current land distribution, several decades or even centuries ago, and subsequent transmission of preferences through the generations seem difficult to reconcile with our finding of systematic changes in preferences over time following shocks.

We find that over 40% of the variation in risk tolerance in our sample is explained by environmental factors. This high figure needs to be put into perspective. One of the reasons for the high value is that the preferences of individuals are highly correlated geographically in our data, with over 50% of the overall variance between individuals occurring between environments, rather than between individuals within one and the same environment. Cesarini et al. (2009) documented the genetic heritability of risk preferences and explained 16% of the variation—one of the highest proportions in the literature to our knowledge (though see also Zhong, Chew, Set, Zhang, Xue, Sham, Ebstein and Israel, 2009). It should, however, be clear that the relative role of genetic and environmental factors will itself not be constant across environments. Indeed, we would expect environmental factors to play less of a role in relatively more homogenous Western populations (see Ridley, 2003, for a book-length discussion). This was indeed one of the reasons for carrying out the experiment in Ethiopia—to maximize our chances of detecting environmental influences.

The effects we presented cannot be explained by mere movements along a fixed, innate utility function defined over any level of wealth one might possibly face over one's lifetime. Assume for a moment that we only observe wealth effects (i.e. movements along a pre-existing utility function defined over lifetime wealth). Since we use a measure of relative risk tolerance to analyze our results, we find shocks—presumably resulting in a decrease in wealth—to increase relative risk aversion. This means that movements along the utility require a function characterized by *decreasing* relative risk aversion in order to account for our findings. Measuring utility over considerable stakes, we found utility to be characterized by *increasing relative risk aversion*—the exact opposite pattern, and the prevalent finding in the empirical literature (Fehr-Duda et al., 2010; Wakker, 2010; Bouchouicha and Vieider, 2017). This results in a contradiction. Controlling for wealth further does not impact our results. We thus conclude that moves along a pre-determined utility function cannot organize our results.

This raises the interesting question of what can account for the effects we observe. We consider the evolutionary models of Robson (2001a) and Netzer (2009) to be the most promising possibility. If one allows for cognitive limits to the extent with which humans can detect changes in utility, then it is evolutionarily optimal for utility to adapt to the probability distribution of consumption opportunities present in the environment, so as to allocate scare cognitive resources where they are most useful. Preferences would then be expected to be updated following changes in perceptions of the likelihood of different consumption opportunities (Robson and Whitehead, 2017). This could then account for the shifts in preferences we observe over time. It could also account for the effects of fixed environmental characteristics. For instance, high variability in rainfall ought to result in a utility function that is geared towards avoiding mistakes where they are most costly, i.e. toward avoiding catastrophically low consumption outcomes. Utility would then systematically adapt to the environment, with lower average consumption opportunities, as well as higher variability in consumption over time, resulting in increased risk aversion.

The account just presented may seem at odds with the high levels of risk tolerance we observe in Ethiopia in general—a finding that is consistent with the evidence from comparative studies showing high risk tolerance especially in Africa (Vieider et al., 2015; Falk et al., 2018; Vieider et al., 2018; Bouchouicha and Vieider, 2019). However, this is only an apparent contradiction, and different mechanisms may be at work with and between countries (Bouchouicha and Vieider, 2019). It may further be possible to organize these results even within the framework of our model. In rich countries, important decisions are relatively rare. Most decisions are taken over small to moderate amounts. To wit, people tend to overinsure moderate losses (Sydnor, 2010). The evolutionary model would thus predict the utility function to rise steeply initially, and then to level off gradually up to very high outcomes, resulting in a highly concave function. The situation is exactly the opposite in developing countries. Important decisions need to be taken frequently, following the agricultural cycle. With consumption close to the subsistence level, such decisions are of vital importance, and mistakes can be very costly indeed. The model can thus reconcile the apparently contradictory findings of increased risk tolerance in poorer countries, and decreased risk tolerance in more vulnerable regions within those same countries.

7 Conclusion

We presented unique incentivized panel data on risk preferences from Ethiopia, and paired those data with detailed historical data on rainfall levels. This allowed us to investigate the effects of rainfall shocks on risk preferences. Using a within estimator to emulate individual fixed effects, we found rainfall deviations to reduce risk-tolerance. We also showed how an analysis of cross-sectional data would have led to the exact opposite conclusion, showing our contribution over a literature that has used mainly cross-sectional data. Looking at historical rainfall metrics and fixed geographical characteristics, we uncovered effects that are highly consistent with the time-changing ones. In particular, we found idiosyncratic risk tolerance to decrease strongly in the standard deviation of historical rainfall, in altitude, and in distance to markets. We excluded selection effects as a plausible explanation for these effects. Overall, our results thus indicate that preferences systematically adapt to the environment faced by the decision maker.

References

Abdellaoui, Mohammed, Aurélien Baillon, Lætitia Placido, and Peter P. Wakker (2011) 'The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation.' *American Economic Review* 101, 695–723

Abdellaoui, Mohammed, Emmanuel Kemel, Amma Panin, and Ferdinand M. Vieider (2019) 'Measuring risk and time preferences in an integrated framework.' *Games and Economic Behavior* 115, 459–469

Allison, Paul D (2009) *Fixed effects regression models,* vol. 160 (SAGE publications)

Andersson, Ola, Jean-Robert Tyran, Erik Wengström, and Håkan J. Holm (2016) 'Risk Aversion Relates to Cognitive Ability: Preferences or Noise?' *Journal of the European Economic Association* 14(5), 1129–1154

Barrios, Salvador, Luisito Bertinelli, and Eric Strobl (2010) 'Trends in rainfall and economic growth in Africa: A neglected cause of the African growth tragedy.' *The Review of Economics and Statistics* 92(2), 350–366

Borgomeo, Edoardo, Hassaan F Khan, Matias Heino, Esha Zaveri, Matti Kummu, Casey Brown, and Anders Jägerskog (2020) 'Impact of green water anomalies on global rainfed crop yields.' *Environmental Research Letters* 15(12), 1–11

Bouchouicha, Ranoua, and Ferdinand M. Vieider (2017) 'Accommodating stake effects under prospect theory.' *Journal of Risk and Uncertainty* 55(1), 1–28

_ (2019) 'Growth, Entrepreneurship, and Risk-Tolerance: A Risk-Income Paradox.' *Journal of Economic Growth* 24(3), 257–282

Brooks, Peter, Simon Peters, and Horst Zank (2013) 'Risk behavior for gain, loss, and mixed prospects.' *Theory and Decision* 77(2), 153–182

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010) 'Risk and Rationality: Uncovering Heterogeneity in Probability Distortion.' *Econometrica* 78(4), 1375–1412

Cameron, A Colin, and Douglas L Miller (2015) 'A practitioner's guide to cluster-robust inference.' *Journal of Human Resources* 50(2), 317–372

Cameron, Lisa, and Manisha Shah (2015) 'Risk-taking behavior in the wake of natural disasters.' *Journal of Human Resources* 50(2), 484–515

Cesarini, David, Christopher T. Dawes, Magnus Johannesson, Paul Lichtenstein, and Björn Wallace (2009) 'Genetic Variation in Preferences for Giving and Risk Taking.' *Quarterly Journal of Economics* 124(2), 809–842

Choi, Syngjoo, Shachar Kariv, Wieland Müller, and Dan Silverman (2014) 'Who Is (More) Rational?' *American Economic Review* 104(6), 1518–1550

Chuang, Yating, and Laura Schechter (2015) 'Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results.' *Journal of Development Economics* 117, 151–170

Damania, Richard, Sebastien Desbureaux, and Esha Zaveri (2020) 'Does rainfall matter for economic growth? evidence from global sub-national data (1990–2014).' *Journal of Environmental Economics and Management* 102, 1–9

De Brauw, Alan, and Valerie Mueller (2012) 'Do limitations in land rights transferability influence mobility rates in Ethiopia?' *Journal of African Economies* 21(4), 548–579

Derbile, Emmanuel K, Dramani JM File, and Alfred Dongzagla (2016) 'The double tragedy of agriculture vulnerability to climate variability in africa: How vulnerable is smallholder agriculture to rainfall variability in ghana?' *Jàmbá: Journal of Disaster Risk Studies*

Dercon, Stefan, and Catherine Porter (2014) 'Live Aid Revisited: Long-Term Impacts of the 1984 Ethiopian Famine on Children.' *Journal of the European Economic Association* 12(4), 927–948

Dercon, Stefan, and Luc Christiaensen (2011) 'Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia.' *Journal of Development Economics* 96(2), 159–173

Diamond, Jared (2005) *Collapse: How societies choose to fail or succeed* (Penguin)

Doepke, Matthias, and Fabrizio Zilibotti (2014) 'Culture, Entrepreneurship, and Growth.' In 'Handbook of Economic Growth,' vol. 2

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner (2011) 'Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences.' *Journal of the European Economic Association* 9(3), 522–550

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde (2018) 'Global evidence on economic preferences.' *The Quarterly Journal of Economics* (4), 1645–1692

Fehr-Duda, Helga, Adrian Bruhin, Thomas F. Epper, and Renate Schubert (2010) 'Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size.' *Journal of Risk and Uncertainty* 40(2), 147–180

Fehr-Duda, Helga, and Thomas Epper (2012) 'Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences.' *Annual Review of Economics* 4(1), 567–593

Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shraddhanand Shukla, Gregory Husak, James Rowland, Laura Harrison, Andrew Hoell et al. (2015) 'The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes.' *Scientific data* 2, 15–66

Galor, Oded, and Stelios Michalopoulos (2012) 'Evolution and the Growth Process: Natural Selection of Entrepreneurial Traits.' *Journal of Economic Theory* 147(2), 759–780

Gelman, Andrew, and Iain Pardoe (2006) 'Bayesian measures of explained variance and pooling in multilevel (hierarchical) models.' *Technometrics* 48(2), 241–251

Gelman, Andrew, and Jennifer Hill (2006) *Data analysis using regression and multilevel/hierarchical models* (Cambridge university press)

Gelman, Andrew, Jessica Hwang, and Aki Vehtari (2014a) 'Understanding predictive information criteria for Bayesian models.' *Statistics and Computing* 24(6), 997–1016

Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014b) *Bayesian data analysis,* vol. 2 (CRC press Boca Raton, FL)

Hanaoka, Chie, Hitoshi Shigeoka, and Yasutora Watanabe (2018) 'Do risk preferences change? Evidence from panel data before and after the great east Japan earthquake.' *American Economic Journal: Applied Economics* 10(2), 298–330

Holt, Charles A., and Susan K. Laury (2002) 'Risk Aversion and Incentive Effects.' *American Economic Review* 92(5), 1644–1655

Jakiela, Pamela, and Owen W. Ozier (2019) 'The Impact of Violence on Individual Risk Preferences: Evidence from a Natural Experiment.' *Review of Economics and Statistics* 101(3), 257–282

Kahneman, Daniel, and Amos Tversky (1979) 'Prospect Theory: An Analysis of Decision under Risk.' *Econometrica* 47(2), 263 – 291

Khaw, Mel W, Paul W Glimcher, and Kenway Louie (2017) 'Normalized value coding explains dynamic adaptation in the human valuation process.' *Proceedings of the National Academy of Sciences* 114(48), 12696–12701

L'Haridon, Olivier, and Ferdinand M. Vieider (2019) 'All over the map: A worldwide comparison of risk preferences.' *Quantitative Economics* 10, 185–215

Li, Yan, Kaiyu Guan, Gary D Schnitkey, Evan DeLucia, and Bin Peng (2019) 'Excessive rainfall leads to maize yield loss of a comparable magnitude to extreme drought in the united states.' *Global change biology* 25(7), 2325–2337

Maccini, Sharon, and Dean Yang (2009) 'Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall.' *American Economic Review* 99(3), 1006–1026

Markowitz, Harry (1952) 'The Utility of Wealth.' *Journal of Political Economy* 60(2), 151–158

McElreath, Richard (2016) *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (Academic Press)

Mundlak, Yair (1978) 'On the pooling of time series and cross section data.' *Econometrica* pp. 69–85

Netzer, Nick (2009) 'Evolution of time preferences and attitudes toward risk.' *American Economic Review* 99(3), 937–55

Noussair, Charles N., Stefan Trautmann, and Gijs van de Kuilen (2014) 'Higher Order Risk Attitudes, Demographics, and Financial Decisions.' *Review of Economic Studies* 21(1), 325–355

Prelec, Drazen (1998) 'The Probability Weighting Function.' *Econometrica* 66, 497–527

Rahmato, Dessalegn (2008) 'Ethiopia: Agriculture policy review.' *Digest of Ethiopia's National Policies, Strategies, and Programs*

Ridley, Matt (2003) *Nature via nurture: Genes, experience, and what makes us human* (New York: Harper)

Robson, Arthur J (2001a) 'The biological basis of economic behavior.' *Journal of Economic Literature* 39(1), 11–33

_ (2001b) 'Why would nature give individuals utility functions?' *Journal of Political Economy* 109(4), 900–914

Robson, Arthur J, and Lorne A. Whitehead (2017) 'Adaptive hedonic utility.' *Working Paper*

Rose, Elaina (1999) 'Consumption Smoothing and Excess Female Mortality in Rural India.' *Review of Economics and Statistics* 81(1), 41–49

Scholten, Marc, and Daniel Read (2014) 'Prospect theory and the "forgotten" fourfold pattern of risk preferences.' *Journal of Risk and Uncertainty* 48(1), 67–83

Schultz, Wolfram (2016) 'Dopamine reward prediction error coding.' *Dialogues in clinical neuroscience* 18(1), 23

Schultz, Wolfram, Peter Dayan, and P Read Montague (1997) 'A neural substrate of prediction and reward.' *Science* 275(5306), 1593–1599

Simon, Herbert A (1956) 'Rational choice and the structure of the environment.' *Psychological review* 63(2), 129

Snijders, Tom A. B., and Roel J. Bosker (2012) *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling,* 2nd ed. (London: Sage Publications)

Stauffer, William R, Armin Lak, and Wolfram Schultz (2014) 'Dopamine reward prediction error responses reflect marginal utility.' *Current Biology* 24(21), 2491–2500

Sutter, Matthias, Martin G Kocher, Daniela Glätzle-Rützler, and Stefan T Trautmann (2013) 'Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior.' *American Economic Review* 103(1), 510–531

Sydnor, Justin (2010) '(Over)insuring Modest Risks.' *American Economic Journal: Applied Economics* 2(4), 177–199

Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen (2010) 'Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam.' *American Economic Review* 100(1), 557–571

Tobler, Philippe N, Christopher D Fiorillo, and Wolfram Schultz (2005) 'Adaptive coding of reward value by dopamine neurons.' *Science* 307(5715), 1642–1645

Vieider, Ferdinand M (2018) 'Violence and risk preference: Experimental evidence from afghanistan: Comment.' *American Economic Review* 108(8), 2366–82

Vieider, Ferdinand M., Abebe Beyene, Randall A. Bluffstone, Sahan Dissanayake, Zenebe Gebreegziabher, Peter Martinsson, and Alemu Mekonnen (2018) 'Measuring risk preferences in rural Ethiopia.' *Economic Development and Cultural Change* 66(3), 417–446

Vieider, Ferdinand M., Mathieu Lefebvre, Ranoua Bouchouicha, Thorsten Chmura, Rustamdjan Hakimov, Michal Krawczyk, and Peter Martinsson (2015) 'Common components of risk and uncertainty attitudes across contexts and domains: Evidence from 30 countries.' *Journal of the European Economic Association* 13(3), 421–452

von Gaudecker, Hans-Martin, Arthur van Soest, and Erik Wengström (2011) 'Heterogeneity in Risky Choice Behaviour in a Broad Population.' *American Economic Review* 101(2), 664–694

Voors, Maarten J., Eleonora E.M. Nillesen, Philip Verwimp, Erwin H. Bulte, Robert Lensink, and Daan P. Van Soest (2012) 'Violent Conflict and Behavior: A Field Experiment in Burundi.' *American Economic Review* 102(2), 941–964

Wakker, Peter P. (2010) *Prospect Theory for Risk and Ambiguity* (Cambridge: Cambridge University Press)

Wark, Barry, Brian Nils Lundstrom, and Adrienne Fairhall (2007) 'Sensory adaptation.' *Current opinion in neurobiology* 17(4), 423–429

Wooldridge, Jeffrey M (2015) *Introductory econometrics: A modern approach* (Nelson Education)

Yaari, Menahem E. (1987) 'The Dual Theory of Choice under Risk.' *Econometrica* 55(1), 95–115

Zhong, Songfa, Soo Hong Chew, Eric Set, Junsen Zhang, Hong Xue, Pak C. Sham, Richard P. Ebstein, and Salomon Israel (2009) 'The Heritability of Attitude Toward Economic Risk.' *Twin Research and Human Genetics* 12(1), 103–107

SUPPLEMENTARY MATERIALS (For online publication)

Environmental Forces Shape Risk Preferences

Salvatore di Falco and Ferdinand M. Vieider

Correspondence to: fvieider@gmail.com

S1 Additional descriptives rainfall data

Figure S1 shows the geographical distribution of shocks during the three Meher seasons immediately preceding our risk measurements. The year 2012 especially sees excess rainfall, which at times is more than 100mm above the historical average. Such excess rainfall is concentrated especially in two Woredas in the north, as well as in two smaller Woredas in the centre of the country. In 2014, we witness extensive droughts. These droughts are especially severe in all the central Woredas, with only two Woredas not being affected at all, and a dew being affected by relatively mild droughts (between 50mm and 100mm less rain than the historical average). Finally, 2016 is a largely normal year, with moderate to severe droughts in two Woredasm and excess rainfall in three Woredas, with one Woreda experiencing an excess above 100mm.

Figure S2 displays the rainfall deviations from the historical means, lagged by one year (i.e. for 2011, 2013, and 2015). 2011 was a largely regular year, with some relatively large outliers in terms of both excess and shortfalls in rain. 2013 was marked by a very wide distribution, with regular rainfall, excess and shortfalls all important for parts of our sample. 2015 follows in the footsteps of the very dry 2014, being characterized by large shortfalls in rain for a large part of our sample.

Figure S3 shows the equivalent figure for the measures lagged by 2 years, i.e. for 2010, 2012, and 2014. We have already discussed 2012 and 2014 in the main text, with 2010 thus providing the only new information. The rains in 2010 almost entirely fall into the region of $\pm$ 100 mm of the historical mean. Overall, this picture tells us that lag 2 will be poorly identified in our regressions, given how 2012 and 2014 are already included in the unlagged predictors, and 2010 does not contain large outliers in rainfall.

Figure S1: Maps of geographical distribution of rainfall, deviation from historical trends, 2012-2016

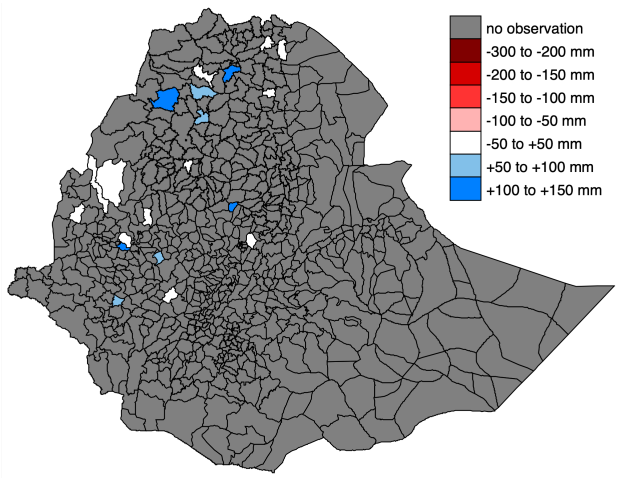

(a) 2012

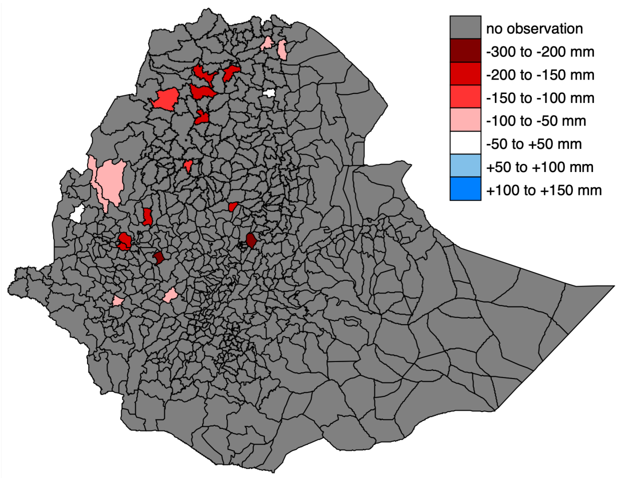

(b) 2014

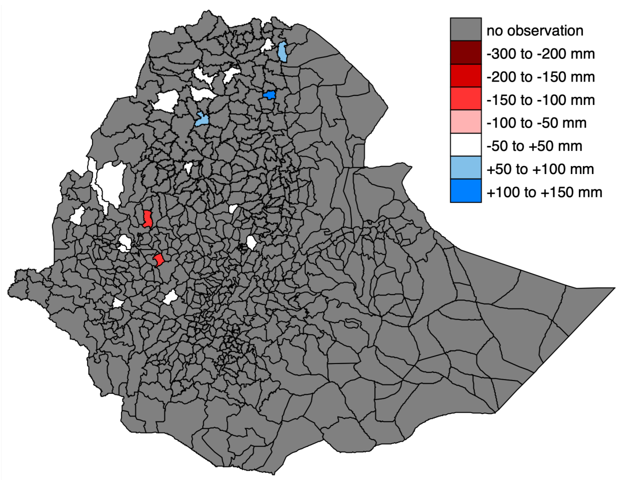

(c) 2016

Figure S2: Average absolute rainfall deviations

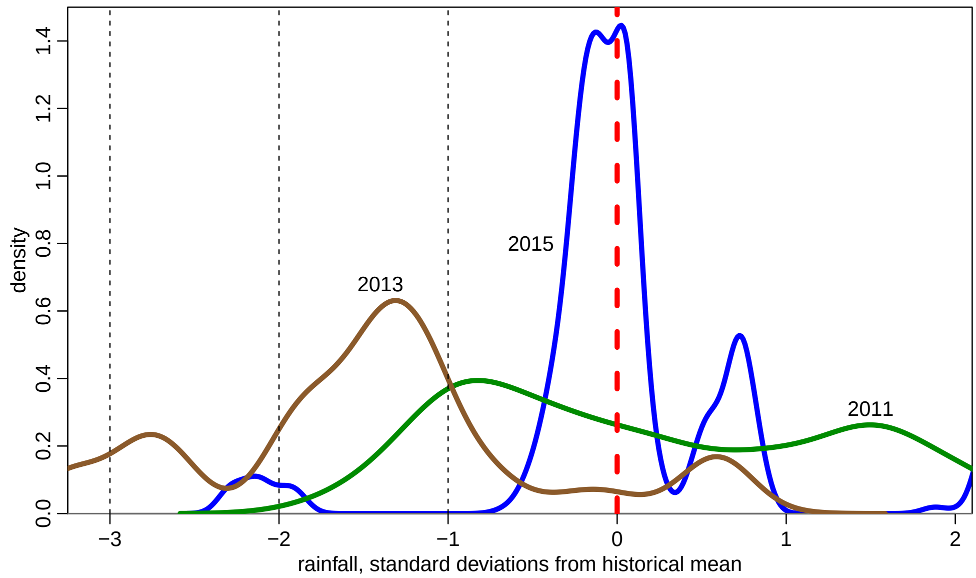

Figure S3: Average absolute rainfall deviations

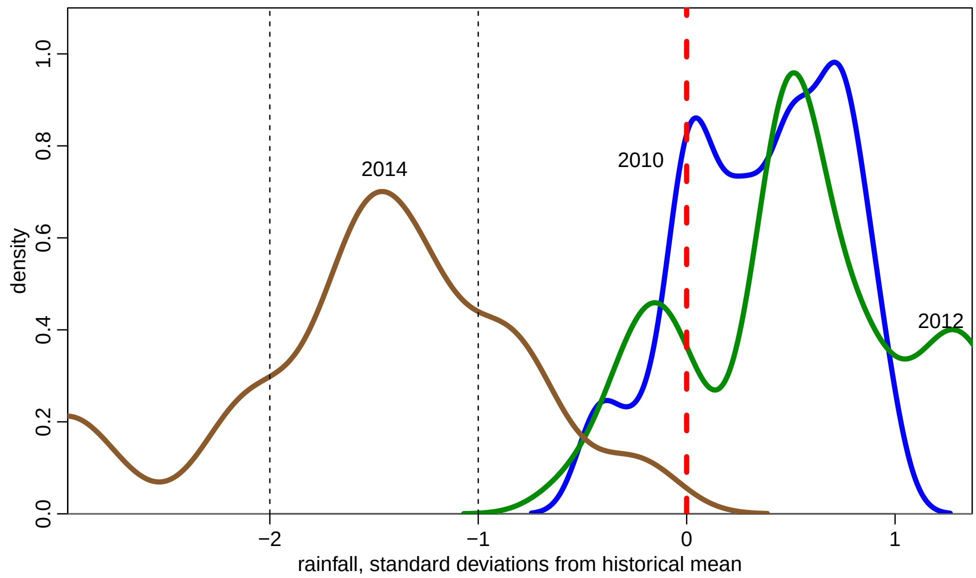

S2 Rainfall shocks and agricultureal yield

S2.1 Descriptive analysis

Our stylized model predicts that rainfall shocks will have an effect on risk tolerance through their impact on consumption, and particularly, through shortfalls in consumption relatively to the historical average. While we use rainfall as an exogenous proxy for consumption shortfalls throughout in the paper, lest we contaminate any causal interpretation of our results by the use of measures such as consumption which may well be endogenous to risk tolerance, it is nevertheless useful to document the effect of rainfall shocks on agricultural yields. This, indeed, serves to back up the premise on which our approach is built—that agriculture is adapted to the local circumstances—and to justify the particular rainfall measures we use.

Table S1: Descriptive data by crop type

cropgrown	freq	plot size (ha)	yield (kg)	kg/ha	altitude (m)	mean rain (mm)	SD rain (mm)
Teff	3263	.389	338	1030	2200	734	94
Maize	2937	.299	452	1812	1981	767	90
Wheat	2564	.273	347	1548	2501	632	107
Barley	2064	.249	312	1525	2499	631	104
Sorghum	777	.595	592	1272	1682	679	87
Potato	735	.171	509	4088	2454	769	92
Fababean	509	.256	226	1016	2420	718	92
Haricot beans	504	.288	221	984	2356	727	92
Millet	492	.249	304	1416	2071	648	93
Chickpeas	359	.298	276	1159	2122	728	104
Field pea	307	.294	204	886	2489	663	96
Cowpea	305	.283	212	860	2483	691	93
Grasspea	239	.242	240	1098	2264	692	98
Noug	188	.442	281	776	1820	779	84
Pepper	129	.167	195	1625	1840	778	97
Sesame	123	.876	655	1000	1055	650	76
Onion	116	.182	761	4779	2000	744	106
Garlic	106	.163	349	2523	2046	741	111

The table lists the 18 most important crops by frequency with which they occur.

For each household, we have parcel-wise data on the crop grown, the area dedicated to that crop (measured in hectares, with one acre corresponding to approximately 0.44 hectares), and the yield measured in kilograms. The households in our sample cultivate a large variety of crops, including various grains, pulses, and vegetables. Table S1 lists the most important ones in terms of frequency with which they are cultivated (out of 3572 possible occurrences). Teff is the most cultivated crop, and it is grown in 91% of the household-year combinations. This is followed by maize, wheat, and barley, with other crops being grown less frequently. In terms of plot size dedicated to it, teff is second only to sorghum, which is however cultivated by much fewer households. In terms of yield per area, however, teff performs rather poorly, with sorghum, potato, and maize taking the top prize. This patterns is not surprising. Teff is indeed an essential ingredient for making Injera—the local bread made out of fermented teff dough—which is an essential part of any meal in Ethiopia.

Table S1 also provides a first indication that the crops planted may depend system-

Table S2: Regression analysis of agricultural land assigned to the top 6 crops

dep var: land area dedicated to crop	maize	wheat	teff	barley	sorghum	patato
mean historical rainfall	0.009*	-0.001	0.034***	-0.008***	-0.048***	0.021**
	(0.005)	(0.003)	(0.006)	(0.003)	(0.013)	(0.010)
SD historical rainfall	-0.151***	-0.090***	-0.276***	-0.061***	-0.089	-0.132***
	(0.029)	(0.020)	(0.040)	(0.016)	(0.096)	(0.032)
altitude	-0.018***	0.008***	0.015***	0.002	-0.035***	0.002
	(0.002)	(0.001)	(0.004)	(0.002)	(0.012)	(0.001)
total land cultivated	0.070***	0.093***	0.166***	0.111***	0.258**	0.073**
	(0.015)	(0.011)	(0.036)	(0.015)	(0.125)	(0.032)
Observations	2951	2565	3272	2070	780	733

Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). The historical rainfall variables are measured in 100s of millimetres. Altitude is measures in 100s of metres, and land area is measured in hectares (1 acre is approximately 0.44 ha). The regressions include error terms at the household and parcel levels.

atically on historical rainfall means and standard deviations, as well as on altitude. To further investigate the geographical distribution of various crops, table S2 shows regressions of the plot size allocated to a certain crop on the mean and standard deviation of historical rainfall and on altitude, while controlling for the total land cultivated by the household, for the six most important crops by frequency. Some clear patterns emerge. Teff and patatos are more likely to be grown—a larger overall land area is allocated to them—in areas with high historical rainfall levels. So is maize, though the effect is only marginally significant, and the coefficient is small. Barley and sorghum, on the other hand, are less likely to be grown in areas with historically high average rainfall levels. This suggests that the latter two crops are rather adapted to drier climates. It also suggests that they may suffer from excessive rainfall, since they both exhibit high yields per area, and are relatively high value crops.

We also observe some systematic associations with the standard deviation of historical rainfall and with altitude. A large standard deviation of historical rain reduces the land allocation for all major crops, likely because the risk from planting such crops becomes too large. Maize and sorghum are less likely to be plated at higher altitudes, whereas teff and wheat are more likely to be planted a high altitudes. All in all, this clearly shows that agriculture is adapted to the local circumstances.

S2.2 Rainfall shocks and crop yields

In this section we regress overall yields per household on rainfall shocks, to determine the effect of positive versus negative deviations from historical averages. Table S3 shows fixed effects regressions of agricultural yields aggregated across all crops on shortfalls and excesses in rainfall relatively to the historical mean, defined as described in equation 2 in the main text. All regressions control for plot size. Regression (1) shows a linear specification regressing yields on shortfalls in rain and excesses in rain. Both show a clear and highly significant negative effect. Regression (2) adds the squared rainfall deviation terms, showing a pattern of decreasing sensitivity to rainfall deviations, replicating the patterns found for risk tolerance in the main text. Regression (3) further adds the mean of historical rainfall, and altitude. Overall yields increase in the historical rainfall mean, indicating that more rainfall is indeed better if farmers are given the time to adapt to it. Finally, higher altitude tends to have a strong negative impact on yield levels. These results remain unaffected by the inclusion of crop fixed effects in regression (4).

Table S3: Regression of agricultural yields on rainfall deviations

dep var: yield in kg	(1)	(2)	(3)	(4)
rain shortfall	-31.536***	-125.768***	-125.174***	-114.100***
	(11.758)	(32.957)	(32.969)	(31.560)
rain excess	-79.657***	-179.958***	-178.413***	-154.927***
	(14.074)	(58.370)	(58.446)	(59.890)
rain shortfall sq.		36.868***	36.674***	33.783***
		(12.062)	(12.059)	(11.768)
rain excess sq.		54.366	53.138	43.946
		(36.137)	(36.178)	(37.127)
mean historical rainfall			8.262**	7.562*
			(3.991)	(4.067)
altitude			-11.369***	-11.117***
			(2.111)	(2.258)
plot size	YES	YES	YES	YES
crop fixed effects	NO	NO	NO	YES
Observations	16879	16879	16879	16404

Standard errors reported in parentheses. Stars signal significance at the 10% level (*), 5% level (**), and 1% level (***). Regressions of yields, in kg and aggregated across all types of crops, on rainfall shocks. Fixed effects are implemented by means of the Mundlak within estimator—see main text for an in-depth discussion. Rain excesses and shortfalls are recorded relative to the historical mean, as defined in equation 2 inn the main text.

S3 Utility fit to risk data

Figure S4 shows how relative risk tolerance changes across the probability range. Two findings stand out. One, we again find very high levels of risk tolerance. Indeed, we find significant risk seeking for all probability levels but the highest two, and only for the highest probability level do we find significant risk aversion. Two, the pattern is clearly one of likelihood-insensitivity. Both these findings line up perfectly with the comparative evidence for students across 30 countries presented by L'Haridon and Vieider (2019), who document that i) risk tolerance systematically decreases in GDP, i.e. developing countries tend to be much more risk tolerant than developed countries; and ii) likelihood insensitivity for gains is universal. This makes it clear that any model ought to capture changes in preferences over outcomes as well as over stakes.

Figure S4: Relative risk tolerance across probabilities

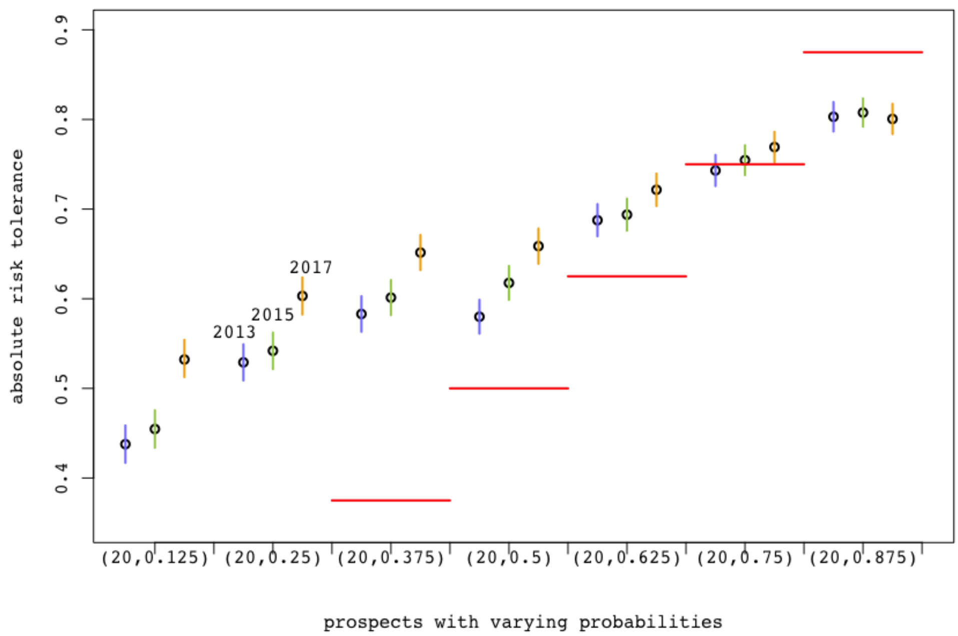

In order to discriminate between decreasing absolute risk aversion and constant absolute risk aversion, we determine the best-fitting utility function to our data. Taking the most general approach, we can describe the indifference between a certain amount and a prospect as follows:

$$u(ce) = w(p)u(x) + (1 - w(p))u(y), \tag{5}$$

where u is a utility function and w a probability weighting function with the typical characteristics. We then estimate the relationship by using techniques akin to those used by L'Haridon and Vieider (2019) by either setting $w(p) = p$, thus assuming an expected utility framework, or by giving $w(p)$ a functional form. We always use a two-parameter formulation, and using either the 2-parameter version of Prelec (1998) of a neo-additive utility function (Abdellaoui et al., 2011) fits the data equally well in combination with either utility function we test.

The utility function is the main part of interest. In particular, we pitch an exponential utility function, reflecting IRRA and CARA, against a logarithmic utility function, reflecting IRRA and DARA. The two functions take the following form:

$$u(x) = \frac{1 - exp(-\rho x)}{\rho} \tag{6}$$

$$u(x) = \frac{ln(1 + \rho x)}{\rho}, \tag{7}$$

where ρ is the coefficient of risk aversion. The exponential function fits the data significantly better under RDU (WAIC of 238,663 vs. 238,740, weight equal to 1; results under EUT are very similar).

S4 Placebo regression using minor rains (Belg)

Table S4 shows the placebo regressions, using the total rainfall measured over the minor rainy season, or *Belg*. Regression (1) regresses risk tolerance on the negative and positive deviations plus their squares in the Belg season only. Regression (2) adds the same measures for the Meher. Regression (3) includes rainfall during the Belg lagged once. And regression (4) once again adds the rainfall measures for the main rainy season. A clear picture emerges. The measures for the minor rainy season are generally not significant, while the measures for the major rainy season show the same significance levels as reported in the main text.

Table S4: Regression of risk-tolerance on rainfall: Placebo regressions

dep. var.: risk-tolerance	(1)	(2)	(3)	(4)	(5)	(6)
Belg rain shortfall	0.004	0.166***	0.010	0.012	0.036	0.040
	(0.018)	(0.056)	(0.066)	(0.066)	(0.070)	(0.070)
Belg rain excess	0.009	0.092***	0.052*	0.036	0.017	0.018
	(0.006)	(0.023)	(0.028)	(0.031)	(0.035)	(0.035)
Belg rain shortfall sq.		-0.101**	0.017	0.023	0.094	0.089
		(0.045)	(0.053)	(0.053)	(0.064)	(0.064)
Belg rain excess sq.		-0.019***	-0.009	-0.007	0.001	0.001
		(0.006)	(0.007)	(0.007)	(0.009)	(0.009)
rain shortfall	-0.016**	-0.097***	-0.116***	-0.148***	-0.181***	-0.178***
	(0.007)	(0.022)	(0.025)	(0.032)	(0.046)	(0.046)
rain excess	-0.023	-0.169***	-0.204***	-0.209***	-0.353***	-0.347***
	(0.015)	(0.047)	(0.050)	(0.053)	(0.079)	(0.079)
rain shortfall sq.		0.029***	0.045***	0.049***	0.056***	0.055***
		(0.007)	(0.009)	(0.010)	(0.014)	(0.014)
rain excess sq.		0.113***	0.167***	0.171***	0.306***	0.302***
		(0.034)	(0.037)	(0.039)	(0.053)	(0.053)

Belg indicates the minor rainfalls used as a placebo. The other variables indicate the major rains, or Meher. Only unlagged variables reported for parsimony. The minor rains do not show any stable significant effects. The effects of the major rains, on the other hand, emerge unscathed. The regressions mirror those in table 3 in the main text, except for regression (7) with interaction effects, which is omitted from this table.

S5 Instructions for enumerators

INSTRUCTIONS

In the present experiment, you will be asked to choose repeatedly between a fixed amount of money and a lottery. The lottery will always give you a chance to win one of two amounts of money. Figure 1 shows a typical choice task. You are asked repeatedly to choose between playing the lottery and obtaining a sure amount of money. **For each row**, you are asked to indicate whether you would prefer to play the lottery or to obtain the sure amount of money by ticking the preferred option.

The urn indicated in the figure contains eight numbered balls. One ball will be extracted from the urn to determine your payoffs in case you should play the lottery. In the lottery displayed, if ball 1 , 2 , 3, or 4 is extracted, you obtain 60 Birr; if ball 5 , 6 , 7 , 8 is extracted, you obtain nothing. Please pay close attention to the amounts to be won as well as the number of balls associated with each outcome, since they change across decisions.

Fig. 1: Example of a typical decision task

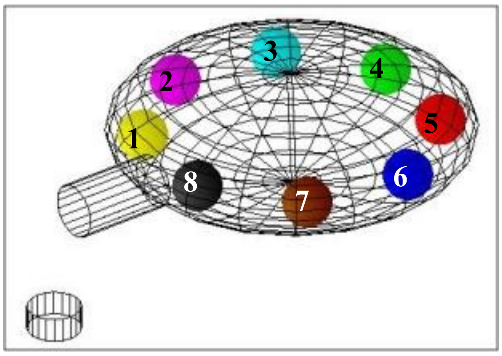

Win **60 Birr** if one of the following balls is extracted:

1 2 3 4

Win **0 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure amount	
O	O	0 Birr for sure
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure

We are interested in the amount for which you will switch from preferring the lottery to preferring the sure amount. Most likely, you will prefer the lottery over a sure amount of 0, and at a certain point switch to the sure amount as the latter increases. Most likely, you would also prefer the sure amount of 60 Birr over the lottery giving you at most 60 Birr, but with a chance of obtaining 0. If you do not want the lottery at all when a positive sure amount is available, you can choose to get the sure amount in the first row and then continue with the sure amount for all choices. Where you will switch from the lottery to the sure amount depends entirely on your preferences—there are no right or wrong answers.

You will be asked to take 17 decisions, for each one of which you will need to decide between a lottery and a series of sure amounts as exemplified in figure 1 above. **Please pay close attention to the amounts to be won as well as the number of balls associated with each outcome!** Indeed, both the higher and lower amount, as well as the number of balls associated to the higher outcome, change between decision problems. Since your final payoff depends on these decisions, it is crucial for you to pay close attention to these features.

Payoff determination
After you have taken all the decisions, one of your decisions will be randomly drawn for real pay, i.e. **the amounts indicated in the decision problem will be paid out for real**. First, one of the 17 decision tasks is drawn at random, using a chance device with equal probability for each decision task to be extracted. For the extracted decision task, one of your decisions, corresponding to one row for which you had to indicate your preference between the sure amount and the lottery, will then be drawn at random with **equal probability for each row**. If for the row that is drawn you have indicated that you prefer the sure amount of money, you will simply be paid that amount.

In case you have chosen the lottery for the randomly determined row, then that lottery will be played according to the probabilities indicated. You will then be paid the outcome corresponding to the ball you drew.

Decision 1

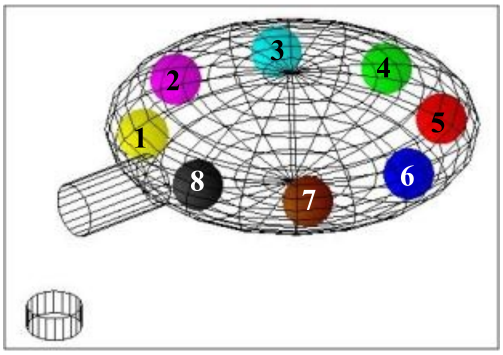

	Lottery	Sure	
	O	O	3 Birr for sure
	O	O	6 Birr for sure
	O	O	9 Birr for sure
	O	O	12 Birr for sure
	O	O	15 Birr for sure
	O	O	18 Birr for sure
	O	O	21 Birr for sure
	O	O	24 Birr for sure
Win **30 Birr** if one of the following balls is extracted:	O	O	27 Birr for sure
Win **0 Birr** if one of the following balls is extracted:			
5 6 7 8			

Decision 2

	Lottery	Sure	
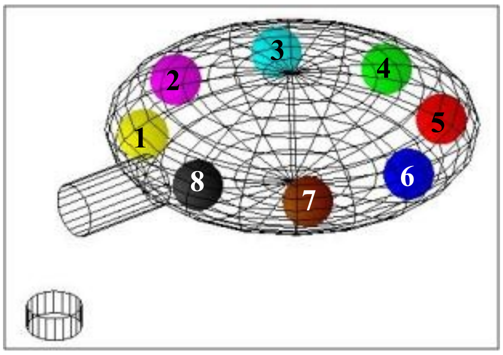	O	O	3 Birr for sure
	O	O	6 Birr for sure
	O	O	9 Birr for sure
	O	O	12 Birr for sure
	O	O	15 Birr for sure
	O	O	18 Birr for sure
	O	O	21 Birr for sure
	O	O	24 Birr for sure
	O	O	27 Birr for sure
	O	O	30 Birr for sure
	O	O	33 Birr for sure
	O	O	36 Birr for sure
	O	O	39 Birr for sure
Win **60 Birr** if one of the following balls is extracted:	O	O	42 Birr for sure
1 2 3 4	O	O	45 Birr for sure
	O	O	48 Birr for sure
Win **0 Birr** if one of the following balls is extracted:	O	O	51 Birr for sure
5 6 7 8	O	O	54 Birr for sure
	O	O	57 Birr for sure

Decision 3

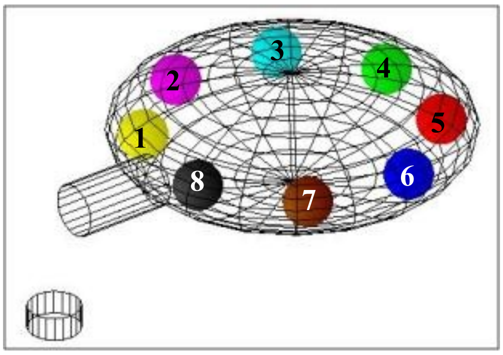

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4

Win **0 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 4

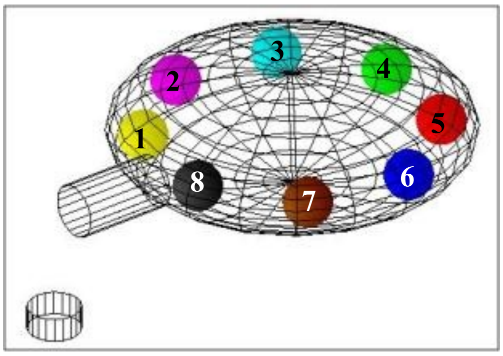

Win **180 Birr** if one of the following balls is extracted:

1 2 3 4

Win **0 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure	
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure
O	O	120 Birr for sure
O	O	123 Birr for sure
O	O	126 Birr for sure
O	O	129 Birr for sure
O	O	132 Birr for sure
O	O	135 Birr for sure
O	O	138 Birr for sure
O	O	141 Birr for sure
O	O	144 Birr for sure
O	O	147 Birr for sure
O	O	150 Birr for sure

Decision 5

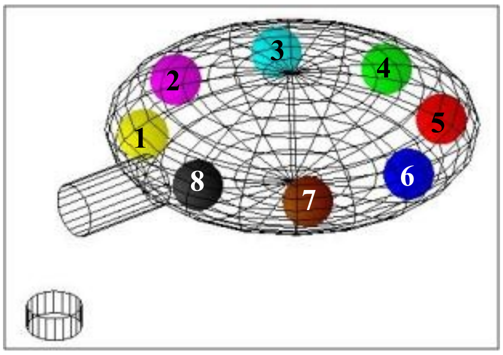

Win **180 Birr** if one of the following balls is extracted:

1 2 3 4

Win **60 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure	
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure
O	O	120 Birr for sure
O	O	123 Birr for sure
O	O	126 Birr for sure
O	O	129 Birr for sure
O	O	132 Birr for sure
O	O	135 Birr for sure
O	O	138 Birr for sure
O	O	141 Birr for sure
O	O	144 Birr for sure
O	O	147 Birr for sure
O	O	150 Birr for sure
O	O	153 Birr for sure
O	O	156 Birr for sure
O	O	159 Birr for sure
O	O	162 Birr for sure
O	O	165 Birr for sure
O	O	168 Birr for sure
O	O	171 Birr for sure
O	O	174 Birr for sure
O	O	177 Birr for sure

Decision 6

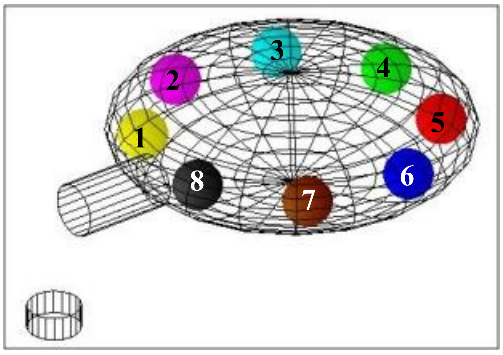

Win **180 Birr** if one of the following balls is extracted:

1 2 3 4

Win **120 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure	
O	O	123 Birr for sure
O	O	126 Birr for sure
O	O	129 Birr for sure
O	O	132 Birr for sure
O	O	135 Birr for sure
O	O	138 Birr for sure
O	O	141 Birr for sure
O	O	144 Birr for sure
O	O	147 Birr for sure
O	O	150 Birr for sure
O	O	153 Birr for sure
O	O	156 Birr for sure
O	O	159 Birr for sure
O	O	162 Birr for sure
O	O	165 Birr for sure
O	O	168 Birr for sure
O	O	171 Birr for sure
O	O	174 Birr for sure
O	O	177 Birr for sure

Decision 7

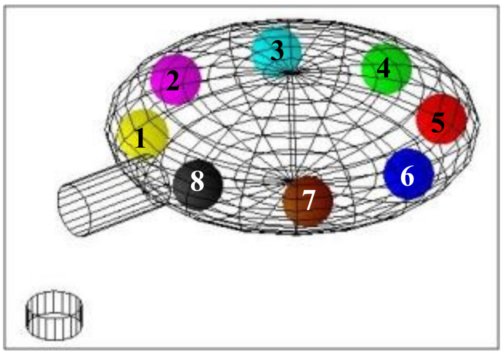

Win **120 Birr** if one of the following balls is extracted:

1

Win **0 Birr** if one of the following balls is extracted:

2 3 4 5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 8

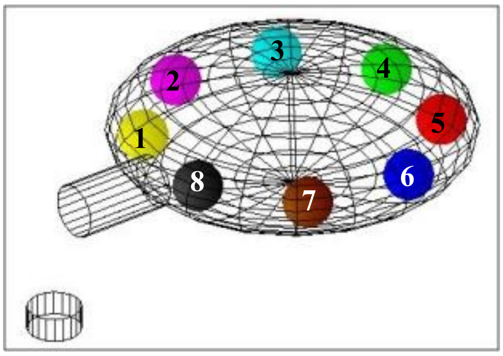

Win **120 Birr** if one of the following balls is extracted:

1

Win **30 Birr** if one of the following balls is extracted:

2 3 4 5 6 7 8

Lottery	Sure	
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 9

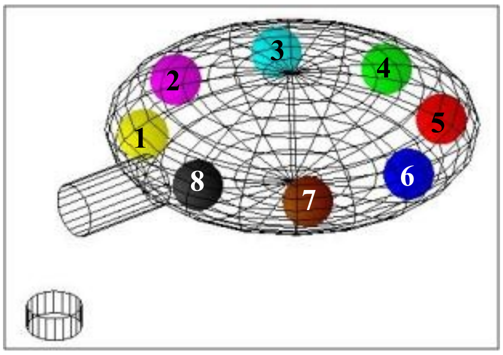

Win **120 Birr** if one of the following balls is extracted:

1 2

Win **0 Birr** if one of the following balls is extracted:

3 4 5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 10

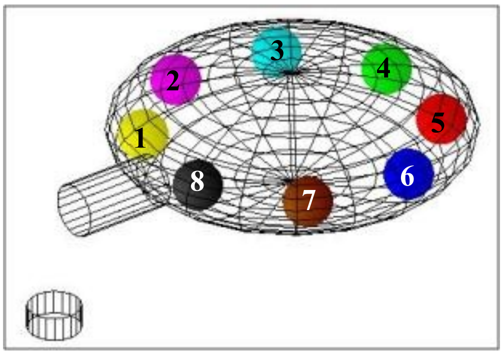

Win **120 Birr** if one of the following balls is extracted:

1 2 3

Win **0 Birr** if one of the following balls is extracted:

4 5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 11

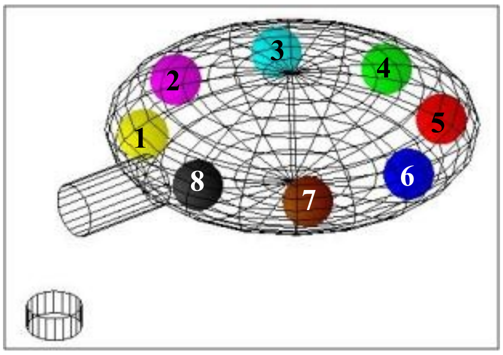

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4 5

Win **0 Birr** if one of the following balls is extracted:

6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 12

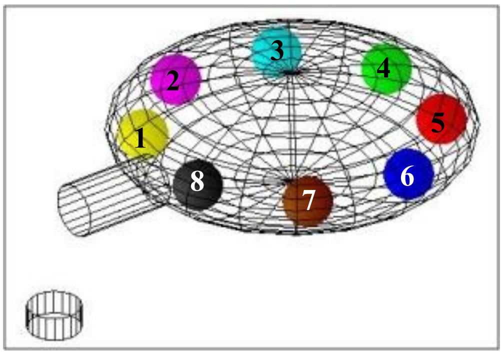

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4 5 6

Win **0 Birr** if one of the following balls is extracted:

7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 13

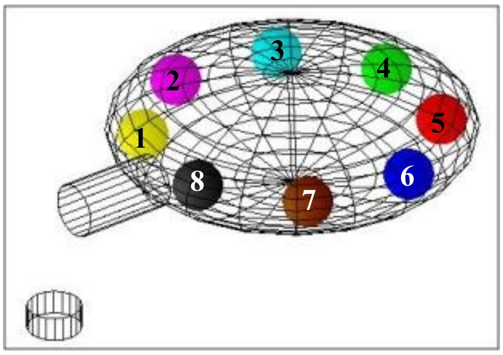

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4 5 6 7

Win **0 Birr** if one of the following balls is extracted:

8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 14

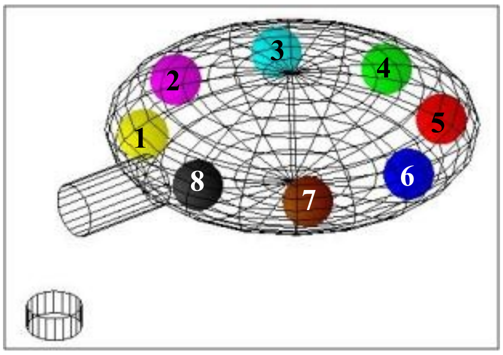

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4 5 6 7

Win **30 Birr** if one of the following balls is extracted:

8

Lottery	Sure	
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 15

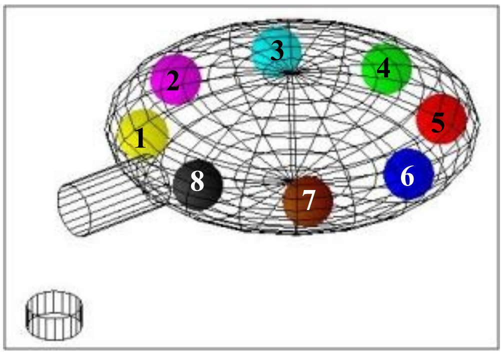

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4

Win **0 Birr** if one of the following balls is extracted:

5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 16

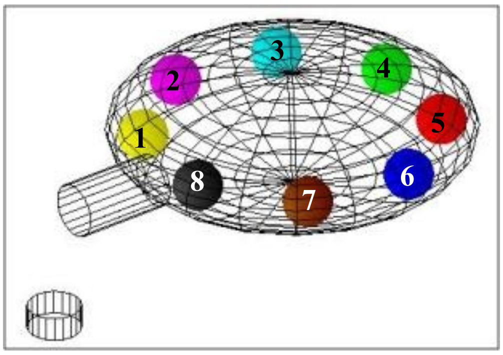

Win **120 Birr** if one of the following balls is extracted:

1

Win **0 Birr** if one of the following balls is extracted:

2 3 4 5 6 7 8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

Decision 17

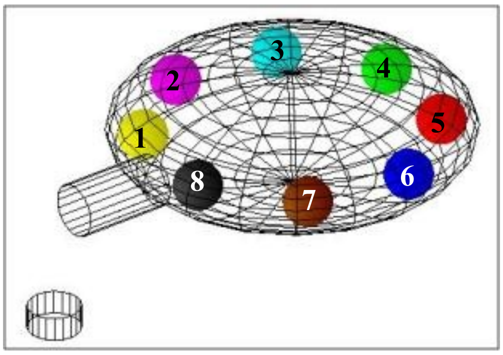

Win **120 Birr** if one of the following balls is extracted:

1 2 3 4 5 6 7

Win **0 Birr** if one of the following balls is extracted:

8

Lottery	Sure	
O	O	3 Birr for sure
O	O	6 Birr for sure
O	O	9 Birr for sure
O	O	12 Birr for sure
O	O	15 Birr for sure
O	O	18 Birr for sure
O	O	21 Birr for sure
O	O	24 Birr for sure
O	O	27 Birr for sure
O	O	30 Birr for sure
O	O	33 Birr for sure
O	O	36 Birr for sure
O	O	39 Birr for sure
O	O	42 Birr for sure
O	O	45 Birr for sure
O	O	48 Birr for sure
O	O	51 Birr for sure
O	O	54 Birr for sure
O	O	57 Birr for sure
O	O	60 Birr for sure
O	O	63 Birr for sure
O	O	66 Birr for sure
O	O	69 Birr for sure
O	O	72 Birr for sure
O	O	75 Birr for sure
O	O	78 Birr for sure
O	O	81 Birr for sure
O	O	84 Birr for sure
O	O	87 Birr for sure
O	O	90 Birr for sure
O	O	93 Birr for sure
O	O	96 Birr for sure
O	O	99 Birr for sure
O	O	102 Birr for sure
O	O	105 Birr for sure
O	O	108 Birr for sure
O	O	111 Birr for sure
O	O	114 Birr for sure
O	O	117 Birr for sure

S6 Details sampling frame

Sampling Frame for BMZ Project "Food and Water Security Under Global Change: Developing Adaptive Capacity with a Focus on Rural Africa" in Ethiopia

Timothy Sulser
27 February 2006

The household sampling frame in Ethiopia was developed to ensure representation at the woreda level of rainfall patterns in terms of both annual total and variation; the four classes of traditionally defined agro-ecological zones (AEZs) found in the basin; vulnerability of food production systems through the proxy of frequency of food aid in the past ten years; and irrigation prevalence. All data used in this sample frame is from the forthcoming *Atlas of the Ethiopian Rural Economy* (Benson et al., in press).

Each woreda was classified according to the following criteria:

Agroecological Zone (traditional typology)	
1	Kolla (blue)
2	Woina Dega (green)
3	Dega (red)
4	Bereha (grey)

Irrigation (percent of cultivated land under irrigation)	
1	no data (lightest blue)
2	0 up to 2
3	2 up to 4
4	4 up to 8
5	8 and greater (darkest blue)

Average Annual Rainfall (total in mm)	
1	0 up to 854 (lightest blue)
2	854 up to 1133
3	1133 up to 1413
4	1413 up to 1692
5	1692 and greater (darkest blue)

Rainfall Variability (coefficient of variation for annual rainfall)	
1	0 up to 62.405 (lightest blue)
2	62.405 up to 80.691
3	80.691 up to 98.976
4	98.976 up to 117.262
5	117.262 and greater (darkest blue)

Vulnerability (number of years food aid received in past 10 years)

1 0 up to 2 (lightest red)
2 2 up to 4
3 4 up to 6
4 6 up to 8
5 8 and greater (darkest red)

Twenty woredas were selected such that across each of the above dimensions the proportion falling into each class for the sample matched as closely as possible the proportions for each class in the entire Nile basin. The selected woredas are indicated in Figure 1 and Table 1. From each of these woredas, 50 households will be randomly selected from municipal rosters to ensure adequate representativeness of the 1000 household sample. Figures 2 through 6 on the following pages present thematic maps for each of the sampling dimensions for the Nile basin.

Table 1. Key to woredas in sample.

1	Hawzen
	Atsbi
2	Wenberta
3	Endamehoni
4	Debark
5	Sanja
6	Wegera
7	Kemkem
8	Enemay
9	Quarit
10	Gimbi
11	Haru
12	Limu
13	Nunu Kumba
14	Kersa
15	Hidabu Abote
16	Bereh Aleltu
17	Wembera
18	Bambasi
19	Sirba Abay
20	GeshaDaka

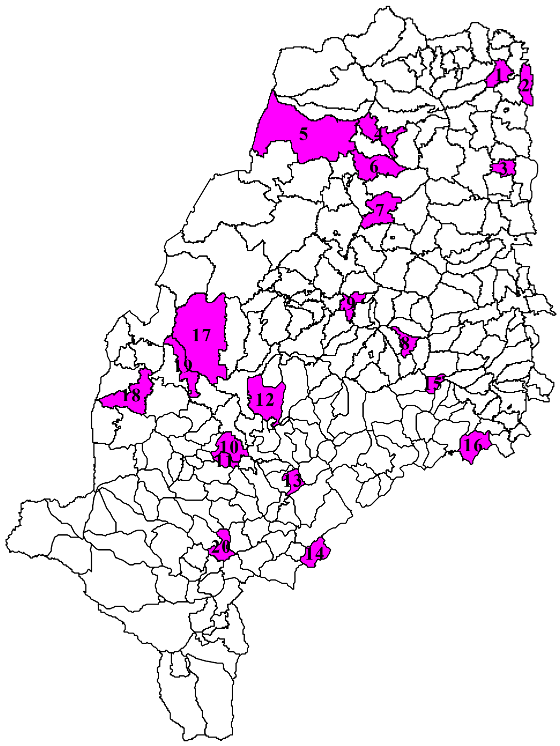

Figure 1. Map of woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names).

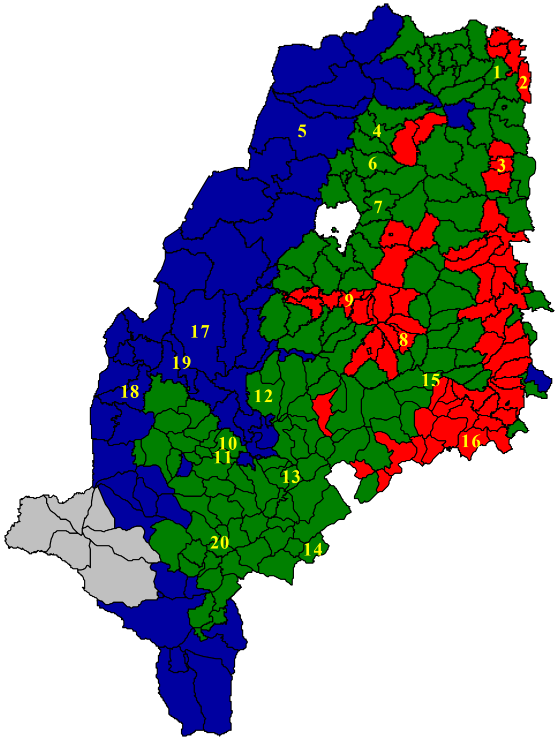

Figure 2. Thematic map of traditional agro-ecological zones and woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names and above classification for details).

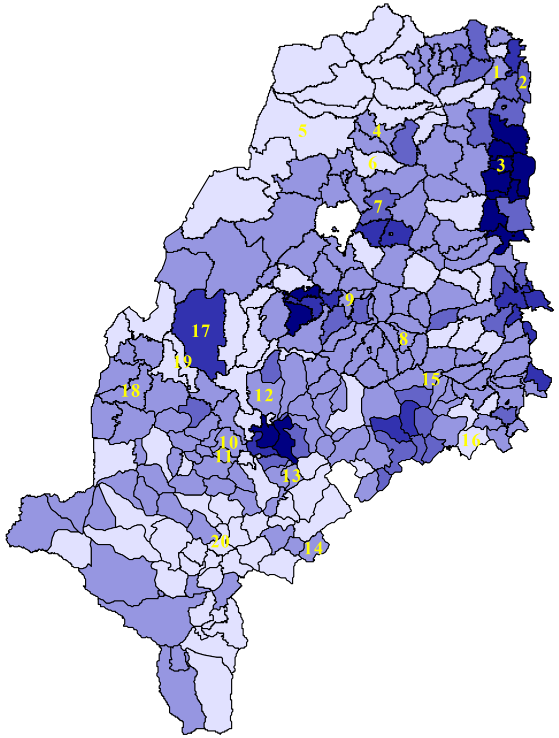

Figure 3. Thematic map of irrigation prevalence and woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names and above classification for details).

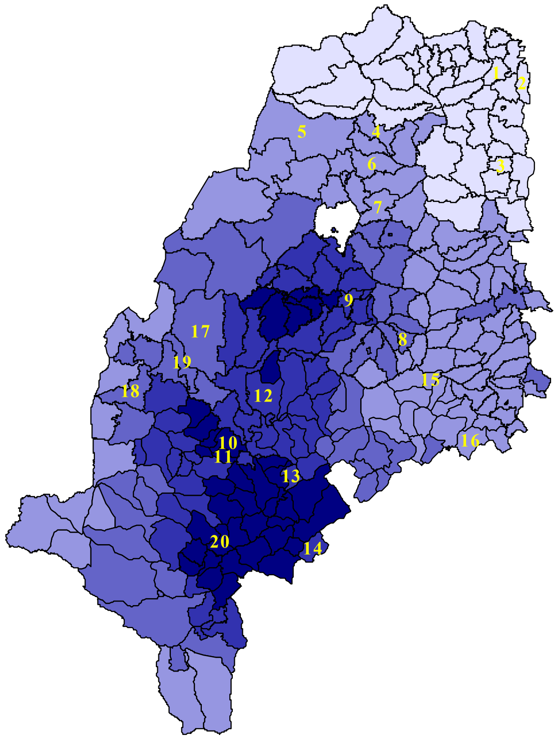

Figure 4. Thematic map of average total annual rainfall and woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names and above classification for details).

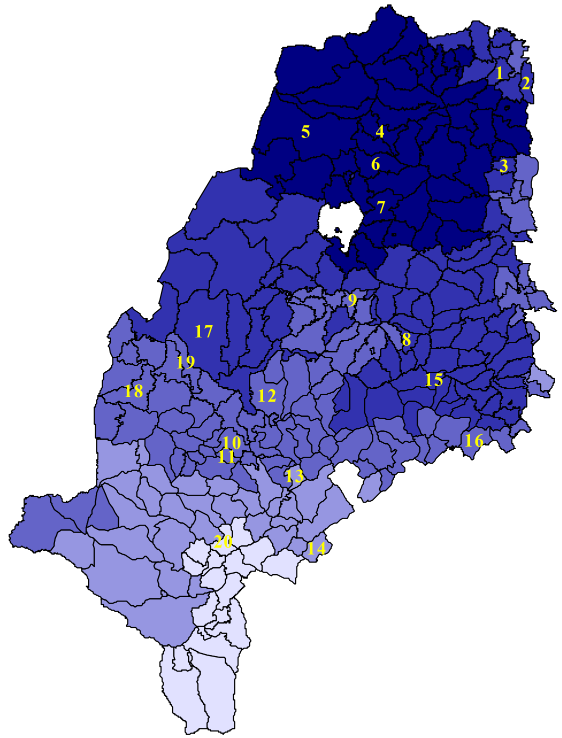

Figure 5. Thematic map of annual rainfall variation and woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names and above classification for details).

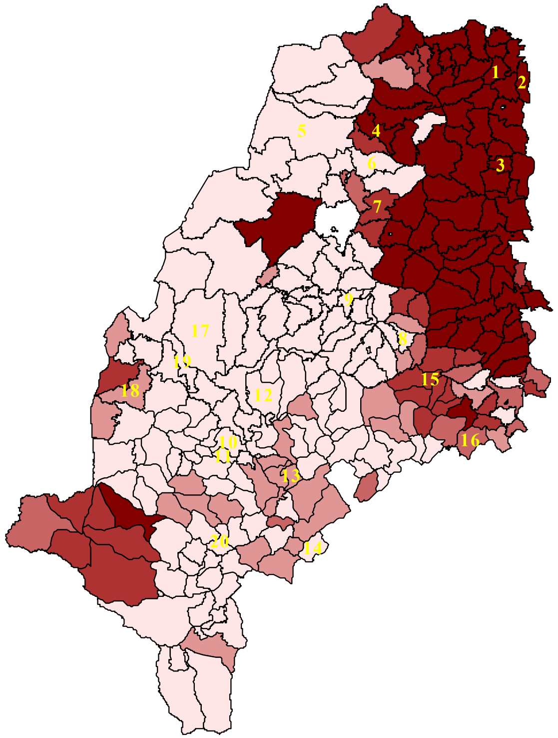

Figure 6. Thematic map of vulnerability and woredas selected for sample in Nile Basin of Ethiopia (see Table 1 for woreda names and above classification for details).