The Value Function with Life Satisfaction Data

Kanninen, Ohto; Mahler, Daniel Gerszon

The Value Function with Life Satisfaction Data[*]

Ohto Kanninen[†] & Daniel Mahler[‡]

March 21, 2017

Abstract

Prospect theory and its empirical applications have shown that in some contexts people make choices based on the value of those choices relative to a reference point. The resulting mapping from outcomes to utility is called the value function and it exhibits loss aversion and diminishing sensitivity. These properties make the value function an S-curve with a kink at the reference point. In this paper, we use the German Socio-Economic Panel ($n > 250,000$) to test whether the properties of the value function extend from narrow gambling choices in experiments to yearly changes in earnings evaluated with life satisfaction. We find that the mapping from changes in earnings to life satisfaction mimics the predicted S-curve remarkably well when the reference point is generated from individuals' past earnings. This finding is robust to a large set of alternative specifications. In congruence with experimental evidence, we find that earnings losses have around 2 times greater impact on life satisfaction than earnings gains. We emphasize that the S-curve we find need not be causal, since the changes in earnings were based on observational data. However, we can rule out that certain other factors produced the observed relationship, including expected utility theory.

[†]Address: Labour Institute for Economic Research, Pitkänsillanranta 3A, 00530, Helsinki, +358 41 513 7175 (ohto.kanninen@labour.fi).

[‡]Address: Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26, 1353 København K, Denmark, +45 2680 3188 (daniel.mahler@econ.ku.dk).

1 Introduction

When expected utility theory describes the behavior of an agent, the agent is called rational. To account for observed systematic departures from such rational behavior, Kahneman and Tversky (1979) proposed prospect theory as an alternative theory. At the core of prospect theory is the value function, which measures the subjective value derived from the argument in question, such as income or consumption, and encompasses three departures from rationality. Firstly, the argument of the value function is defined relative to a reference point. Secondly, it contains loss aversion, implying that it is steeper for losses than gains. Thirdly, it exhibits diminishing sensitivity, meaning that it is concave for gains and convex for losses.

Starting from Kahneman and Tversky (1979) the properties of the value function have been well-documented in experimental settings where subjects are asked to choose between various gambles. Two elements of the value function, reference dependence and loss aversion, have also been shown to influence behavior in many settings outside the lab (see Barberis (2013) for an excellent recent review). However, diminishing sensitivity, the feature of prospect theory that generates an S-curve, has remained elusive in non-experimental settings. Indeed, Barberis (2013) concludes that in empirical applications of prospect theory diminishing sensitivity seems "much less important." Shleifer (2012) notes that the value function occasionally is graphed as a simple piecewise linear function, thus fully neglecting diminishing sensitivity.

In this paper, we test for the existence of an S-curve outside of the lab. We find that the mapping from changes in earnings to life satisfaction follows the hypothesized S-curve remarkably well. In congruence with previous experimental evidence, we find that life satisfaction is associated with earnings losses about twice as steeply as earnings gains. We complement the findings of Kahneman and Tversky (1979) in three important ways. Firstly, we look at experienced utility, as measured by life satisfaction, rather than decision utility. Secondly, we consider changes in labor income rather than narrow gambling situations. Thirdly, and perhaps most importantly, we extend the previous laboratory findings to a non-experimental setting. For a wider applicability of any laboratory finding, this last step is pertinent, but it sets challenges to causal inference.[1] [2]

We measure income by self-reported net monthly earnings in logs. Our baseline reference point is the log of last year's monthly earnings plus the average change in

---

[1] Levitt and List (2007) list multiple occasions where laboratory findings do not extend to the field. This can happen due to the context in which choices are embedded.

[2] A fourth departure from Kahneman and Tversky (1979) is that we look at realized outcomes. Consequently, we are not concerned with uncertainty and probability weighting.

log earnings. We use the log difference as our main explanatory variable, but show that our results are robust to using level differences or proportional changes. Our data comes from 30 years of the German Socio-Economic Panel (GSOEP), which allows our baseline set-up to contain more than 250,000 observations.

We emphasize that we observe the S-curve in non-experimental data. In principle the S-curve could arise spuriously from expected utility theory if there is a particular correlational structure between life satisfaction, income changes and income levels. To test whether this is the case, and inspired by the reference-dependent utility function in Kőszegi and Rabin (2006), we test for the presence of the value function while simultaneously controlling for income levels. We find evidence for both the S-curve and expected utility theory, suggesting that both theories map unto experienced utility. This is in line with prior experimental evidence (Harrison and Rutström, 2009). In particular, we find that if the median earner experiences a standard deviation change in log earnings, the value function has roughly twice the impact on life satisfaction than the standard utility function. For individuals with low earnings, diminishing utility kicks in and the standard utility function plays a relatively larger role.

We study at length whether some other omitted variable could be behind the observed S-curve, such as health status, changes in work hours, job changes, and lagged income. When controlling for these we still find evidence for the S-curve. We also test whether our results could be driven by reversed causality. In our main specification we use individuals' current income. We conduct robustness checks where our income measure is the annual household income in the year prior to the survey. The S-curve remains.

Another potential concern is that diminishing sensitivity arises mechanically due to the boundedness of the life satisfaction scale, which goes from 0 to 10. Individuals who report 10 out of 10 in life satisfaction and get an income increase above their reference point cannot increase their stated life satisfaction further. The reverse applies to individuals who report 0 out of 10. Three points argue against this generating diminishing sensitivity. Firstly, rather few individuals find themselves at the boundary of the life satisfaction scale (4.5% report being 10/10, 0.2% report being 0/10 and only 1.25% being below 3). If we exclude individuals reporting 0, 1, 9, or 10 our results hold - indeed, they become stronger. Secondly, if diminishing sensitivity is purely mechanical, any variable positively associated with life satisfaction should generate S-curves as well. We show that this is not the case. Thirdly, as a robustness check we transform the life satisfaction variable such that responses close to the boundary carry a larger weight. This does not change the results.

In sum, we fail to find an alternative explanation for the S-curve and infer that it is likely that the value function plays a causal role in converting income changes into utility. Experimental or quasi-experimental downward income changes of varying magnitudes remain to be studied to confirm the causality of our observed relationship.

In further robustness checks we vary the independent variable, the income variable, the dataset (with the British Household Panel Survey) and the reference point. We also test the results with various subsamples. The robustness checks suggest that the relationship is not spurious. However, not all the properties of the S-curve hold with all specifications. Generally, diminishing sensitivity is more robust to changes in the specification than loss aversion, which fluctuates between 1 and 4 depending on the specification.

To our knowledge, we are the first to find that the value function follows the predictions from prospect theory outside of lab settings. The S-curve has previously been found using experienced utility in experimental settings (see for example Galanter (1990) and Carter and McBride (2013)). The most extensive study using experienced utility in a non-experimental context is by Vendrik and Woltjer (2007). They look for the S-curve, also using the German Socio-Economic Panel, but where the reference point is defined as the mean income of a reference group. They observe a globally concave value function and thus rule out an S-curve with their reference point. We complement their findings by showing that generating a reference point based on past earnings does generate an S-curve.

A number of papers have documented elements of the value function by using experienced utility outside of experiments. For example, Luttmer (2005) and Ferrer-i Carbonell (2005) find that individuals' well-being depend on the distance between their own income and the income of a reference group. Boyce et al. (2013) study loss aversion using panel data. They observe bigger impacts of losses than gains in income. De Neve et al. (2014) exploit macroeconomic variation in incomes to study loss aversion, which they find support for.

The rest of the paper is organized as follows. Section 2 describes the theory and outlines our main hypotheses. Section 3 describes the data and our empirical specification. Section 4 presents the results. Section 5 concludes.

2   Theory

At the core of prospect theory is the value function, which has three fundamental properties: loss aversion, diminishing sensitivity, and reference dependence.[3] The argument in the value function is changes in income or consumption rather than levels of income or consumption as in the neoclassical approach. Kahneman and Tversky (1979) argued that people tend to be more sensitive to differences between small than large changes and more sensitive to losses than gains. Thus, the value function is concave for gains, convex for losses, and steeper for small losses than gains, i.e. it exhibits diminishing sensitivity and loss aversion. Together with differentiability everywhere except at the reference point, this generates a value function with an S-shape as shown in Figure 1. The value function is not expected to be predictive of losses so large that "ruin or near ruin is a possible outcome" (Kahneman, 2003).

Figure 1: Textbook Version of the Value Function

Notes: From the Wikipedia entry on prospect theory.

Let $V(\Delta y)$ be the value function, where Δy is changes in y, which in the value function can be income, wealth or consumption. Income can be the argument of the value function formulation in its own right or serve as a proxy for consumption. In principle, differences in consumption and income relative to a reference point could have independent effects in the value function. In expected utility theory the argument is always considered to be consumption and income only acts as a proxy for it.

---

[3]Prospect theory also hypothesizes that people tend to overweight small probabilities and underweight large ones. Our focus is on realized incomes for which there is no uncertainty.

We assume that $V(\Delta y)$ is continuous for all Δy, twice differentiable for all $\Delta y \neq 0$, and that $V(0) = 0$. This notation implicitly assumes that the reference point is the status quo, that is, no increase in y. Given that the average increase in real income typically is larger than zero, this may seem like an overly pessimistic reference point. In our empirical specification we will allow for individuals to have a larger than zero increase as the reference point. Hence, strictly speaking our value function will take the form $V(y_{it} - r_{it})$, where r_{it} is the idiosyncratic reference point for individual i at time t. For notational convenience we denote $y_{it} - r_{it} = \Delta y$ in this section.

Based on Kahneman and Tversky (1979), Bowman et al. (1999), and Kőszegi and Rabin (2006), we make three sets of testable hypotheses regarding the properties of the value function:

H1a: $V'(\Delta y) \geq 0$ and $V''(\Delta y) < 0$, for $\Delta y > 0$,

H1b: $V'(\Delta y) \geq 0$ and $V''(\Delta y) > 0$, for $\Delta y < 0$.

H1a and **H1b** define diminishing sensitivity in the positive (**H1a**) and negative (**H1b**) domain.

H2: $\dfrac{\lim\limits_{\Delta y \to 0} V'(-\Delta y)}{\lim\limits_{\Delta y \to 0} V'(\Delta y)} = \delta > 1$, for $\Delta y > 0$,

H2': $\dfrac{|V(-\Delta y)|}{V(\Delta y)} > 1$, for $\Delta y > 0$,

H2": $\dfrac{|V(-\Delta y)|}{V(\Delta y)} = \delta^c > 1$, for $\Delta y > 0$.

H2 defines loss aversion around the reference point, where $\delta > 1$ is the loss aversion parameter. It implies that very small losses have a larger impact than very small gains. **H2'** is a more demanding definition, which assumes loss aversion holds for all values of Δy. **H2"** is even stronger, as it assumes loss aversion is constant in the whole domain of the value function (Tversky and Kahneman, 1991). We denote the parameter of constant loss aversion δ^c, defined by Tversky and Kahneman (1991) and estimated to be around 2.25 over the range of a few hundred dollars in Tversky and Kahneman (1992). Loss aversion for large changes is ambiguously defined and has taken many definitions in the literature. Kahneman (2003) and Köbberling and Wakker (2005) only look at loss aversion for small changes (**H2**), whereas Kahneman and Tversky (1979) and Tversky and Kahneman (1992) assume it holds at changes of all sizes (**H2'**).

Inspired by Kőszegi and Rabin (2006), we formulate a reference dependent utility function that combines the standard utility function and aspects of the value function as a third hypothesis:[4]

H3: $u(c|r) = m(c) + n(c|r) \approx m(y) + V(\Delta y),$

where r is the reference consumption, $m(c)$ is the standard utility function and $n(c|r)$ is a gain-loss utility function. We approximate the consumption in $m(c)$ with income, y. Kőszegi and Rabin (2006) assume that the decision-maker assesses gain-loss utility in each dimension separately. We focus on just one general dimension of consumption as proxied by earnings. Following this specification, we will test whether both components of the reference-dependent utility function translate into experienced utility.

3 Data & Empirical Specification

3.1 Data

We use the German Socio-Economic Panel (GSOEP), which is a nationally representative household survey conducted yearly since 1984. We use data from 1984–2014. The total sample size meeting our baseline empirical specification contains 41,259 individuals and 251,437 observations. As the measure of value we use life satisfaction, which is the answer to the question *"How satisfied are you with your life, all things considered?"* The answer categories range from 0 (completely dissatisfied) to 10 (completely satisfied). A histogram over the answers to the question is given in Figure 2.

We use self-reported net monthly earnings as the income variable. Only employed individuals can report any earnings. Thus, unemployed and people outside the labor market are excluded from the main analysis. The income variable has been deflated and is expressed in constant 2010 EUR. Income variables that include all respondents and all forms of income exist at the household level. Most of these other income measures reflect the income for the year prior to the interview, and as such need not reflect the income available to the respondent at the time of the interview. We therefore prefer to use contemporaneous labor earnings. Later on we will show that

---

[4]Kőszegi and Rabin (2006) use the terms "consumption utility" and "gain-loss utility." Since we make some departures from their model, we will instead use the terms standard utility function and the value function.

Figure 2: Histogram of Life Satisfaction

Notes: Histogram of answers to the question, *"How satis-fied are you with your life, all things considered?"* 0 means completely dissatisfied and 10 completely satisfied.

the results also hold when using gross or net household income from the year before the interview using the entire sample.

3.2 Empirical Specification

In our baseline specification, we estimate the following function:

$$LifeSat_{it} = \alpha_i + V^-(y_{it} - r_{it})N_{it} + V^+(y_{it} - r_{it})P_{it} + \epsilon_{it},$$

where y_{it} is the log of earned income and r_{it} is the reference level of earnings. In our baseline case we set $r_{it} = y_{it-1} + \mu_t$, where $\mu_t = \frac{1}{n_t}\sum_{j \in n_t}(y_{jt} - y_{jt-1})$. That is, we take the reference point to be last year's earnings plus the average increase in earnings in the sample, all in logs. μ_t ranges from 0.004 to 0.071 over the 30 years of data. The mean over the entire sample is 0.034. We will use several other reference points as robustness checks. $V^-(y_{it} - r_{it})$ and $V^+(y_{it} - r_{it})$ are estimated with a fourth order polynomial or using a restricted cubic spline. We use power functions as robustness checks.[5] P_{it} and N_{it} are indicators of positive and negative difference from the reference point. All estimates include individual fixed effects and at times also year fixed effects (discussion of this to follow). We cluster the standard errors at the individual level and deliberately interpret the life satisfaction answers as cardinal. Using an ordered logit model or a related model would take

---

[5]We use a fourth order polynomial rather than a higher or lower order polynomial, as the fourth order minimizes the 10-fold cross validated root mean square error. We prefer this over the power functions such that we can keep the individuals fixed effects in the regression.

out information that are useful for the present purposes.

The distribution of Δy_{it} is given in Figure 3.[6] 97.3% of the observations are contained in the window. A change of log earnings of -1 is equivalent to having 37% (e^{-1}) of last year's earnings, while a change of log earnings of 1 is equivalent to having 272% (e^1) of last year's earnings. The histogram of changes from the reference point, $y_{it} - r_{it}$, looks almost identical but shifted slightly to the left.

Figure 3: Histogram of Changes in Log Earnings
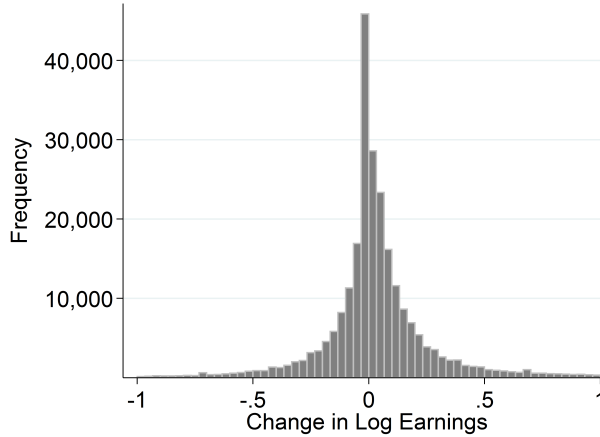
*Notes:* Histogram of changes in log monthly net earnings.

We use log differences rather than relative changes for three reasons. Firstly, in psychology, the Fechner–Weber Law states that subjective sensation is proportional to the logarithm of stimulus intensity (Fechner, 1912). If the value function follows from such stimulus intensity sensation, the log difference is an appropriate measure. Secondly, the distribution of log differences is symmetric, where as that of relative changes is right-skewed, making log differences a convenient measure. Thirdly, using relative changes would mechanically make loss aversion more likely. Using logs will therefore give us a more conservative estimate of loss aversion. Our methodological choice differs from most of the literature. Vendrik and Woltjer (2007) use the relative income gap while Carter and McBride (2013) and Kahneman and Tversky (1979) use absolute changes. We will show later that using relative differences or absolute differences still generates an S-shape but that the nature of loss aversion changes.

---

[6]The figure masks nominal wage rigidities for three reasons: 1) the earnings are deflated, 2) the earnings are self-reported, 3) the earnings are after taxes. If we look at the distribution of nominal self-reported earnings (see Figure A.1), 12% have precisely a zero increase in their earnings, and 30% have a negative change. The 30% figure compares to e.g. Barattieri et al. (2014), who find that in the U.S., 11.5% of quarterly non-zero gross wage changes are negative.

4 Results

To start out with, we graphically inspect whether the relationship between changes from the reference point and life satisfaction follows the pattern we are looking for. In Figure 4 we regress life satisfaction on a restricted cubic spline of changes from the reference point. Panel (a) does not contain year fixed effects while panel (b) does. Both contain individual fixed effects. The splines suggest that we have evidence for diminishing sensitivity and loss aversion in our data. However, when using year fixed effects, positive deviations from the reference point are insignificant. One of the reasons for this could be that year fixed effects take away business cycle variation. If everyone is having a good year, this will be captured by the year fixed effects, diluting the variation that the value function can pick up.

Figure 4: Restricted Cubic Splines

(a) Without Year Fixed Effects   (b) With Year Fixed Effects
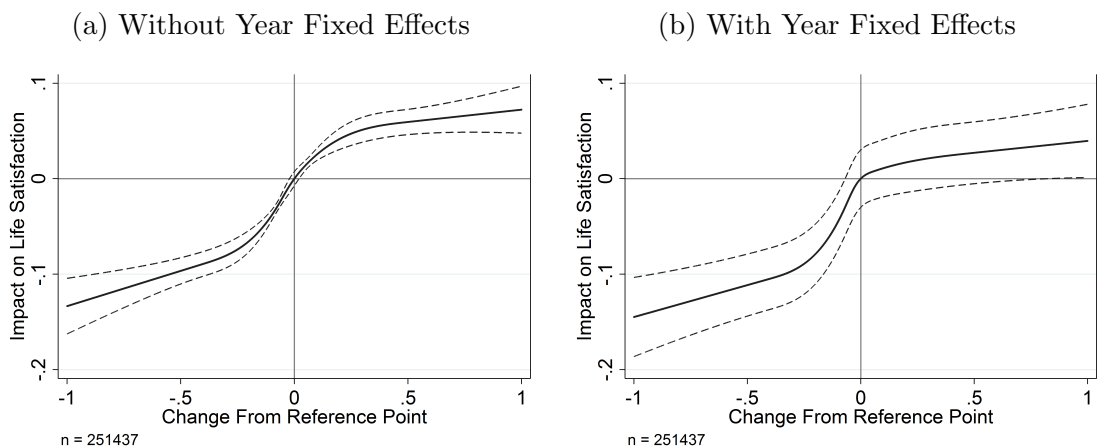
Notes: Panel (a) shows the predicted values from a fixed effects regression of life satisfaction on changes from the reference point. Panel (b) in addition controls for year fixed effects. We use a restricted cubic spline with 5 knots. Dashed lines indicate 95 pct. confidence interval. Standard errors are clustered at the individual level. We use last year's earnings plus the average change in earnings, all measured in logs, as the reference point.

It is important to note that our choice of reference point is mechanically related to the amount of loss aversion in this set-up. If we choose a more ambitious reference point, the vertical line shifts to the right and the amount of loss aversion will be greater, and vice versa. If the reference point is decreased to about a zero increase in nominal wages, positive deviations from the reference point are significant when we use year fixed effects.

The standard deviation of life satisfaction is 1.63 and the standard deviation of changes in log earnings is 0.34. Based on Figure 4(a), a log earnings loss of 1 standard deviation is associated with a decrease in life satisfaction of 0.05 standard deviations

10

(0.085/1.63). A log earnings increase of 1 standard deviation is associated with an increase in life satisfaction corresponding to 0.03 standard deviations (0.054/1.63). Given that much of the variation in life satisfaction is pure noise, we consider these effects to be of relevant size.

4.1 The Reference Point

The previous discussion imposed a very specific reference point on the respondents, the reason being that there is no obvious way to define the reference point for the value function (see Barberis (2013) for a discussion of this). Kahneman (2003) states that the reference point is "usually the status quo." In this setting, it seems a bit implausible that all individuals expect no changes in their real earnings. Kőszegi and Rabin (2006) assume that the reference point is defined as rational expectations about outcomes, something which Abeler et al. (2011) find convincing support for. It is unclear how best to operationalize these expectations, but it is plausible that individuals rationally expect to have an average increase in earnings.

The GSOEP provides an explicit question that tries to get to the reference point in the surveys of 1992, 1997, and 2007. The respondents are asked to state what would be a *somewhat inadequate income*, a *barely adequate income*, a *good income*, and *a very good income* in net monthly terms for the household (see Van Praag and Frijters (1999) for more information on these questions). In addition, in 2002, 2007, and 2012 individuals are asked what would be *the minimum net household income needed to get by*. In Figure 5 we plot a similar graph as before, but now the reference points are these self-reported, subjective income levels. Hence, $r_{it} = y_{it,good}$ etc. We use a measure of households' monthly net income as the income variable to align the income measure with these particular questions.

It seems to be the case that a reference point just around a *barely adequate income* and *a minimum needed income* generates an S-curve with the amount of loss aversion found in experimental settings. On average, the income individuals consider to be barely adequate is log 0.14 lower than their actual income, while they consider a minimum needed income to be log 0.23 lower.

Note also that the self-reported numbers seem to be a more precise measure of the reference point. The effects are about five times as large with the subjective reference points compared to the reference points based on lagged earnings. Thus, lagged earnings could be considered a noisy measure of the reference point. One potential reason for this is that lagged earnings are subject to some degree of hedonic adaptation (see Frederick and Loewenstein, 1999; Loewenstein and Ubel, 2008; Kimball et al., 2015). One of the findings in the literature on hedonic adaptation

Figure 5: Using Subjective Reference Points

(a) Somewhat Inadequate Income

(b) Barely Adequate Income

(c) Minimum Needed Income
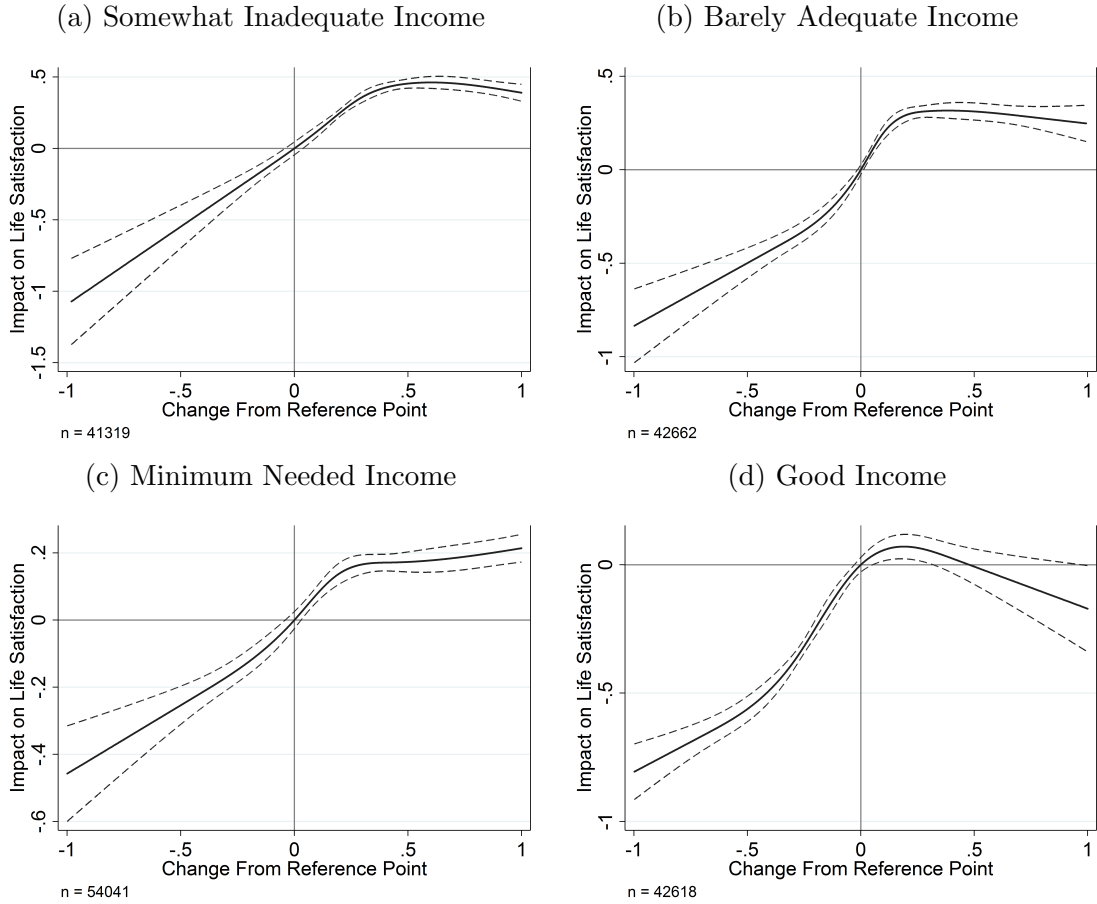
(d) Good Income

Notes: The figure shows the predicted values from regressing life satisfaction on the distance between net monthly household income and the net monthly household income deemed necessary by the respondent to have "a good income" etc. All variables are measured in logs.

is that people expect life events to have a larger effect on their well-being than actually is the case (see Riis et al., 2005; Ubel et al., 2005). People are not usually interviewed the moment they hear about their raise in salary. By the time they are asked about their life satisfaction, they may have partly adapted to the higher or lower earnings compared to the previous year. As they are asked to state subjective reference earnings at the same time as they report their income, no such adaptation takes place and the measured effect is larger.

Fundamentally, it is problematic to use subjective measures of the reference point, since they are likely to have correlated measurement errors with life satisfaction. An advantage of generating reference points based on lagged earnings is that we have more than ten times the sample size. Moreover, when using lagged earnings our results still seem to be driven by individuals that experience earnings changes above or below their subjective expectations. To illustrate this, we exploit a ques-

12

tion where individuals are asked how likely it is that they will get a pay raise within the next two years. We binarize the answers according to the stated likelihood. For our results to be driven by unexpected earnings changes, in the positive domain, the S-curve should be driven by individuals who received an unexpected pay raise. Conversely, in the negative domain the S-curve should be driven by individuals who were least expecting a pay cut. Figure 6 shows that this is indeed the case.

Figure 6: The Role of Pay Raise Expectations
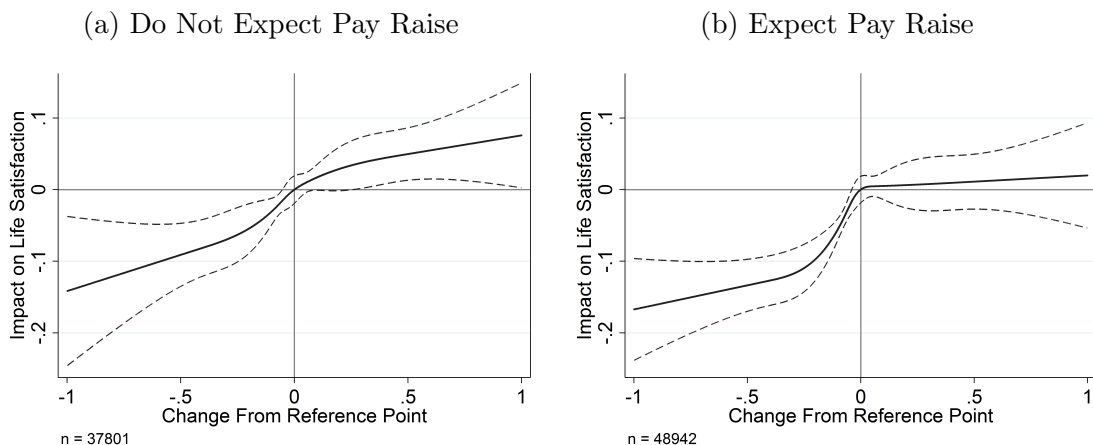
(a) Do Not Expect Pay Raise (b) Expect Pay Raise

Notes: The figure shows the predicted values from regressing life satisfaction on changes from the reference point. Panel (a) includes individuals who expect a pay raise within the next two years while panel (b) includes individuals who do not expect a pay raise.

For these reasons, our main specification will continue to use reference points based on the previous year's earnings. We will further study the robustness of the results to other choices of reference points later in the analysis.

4.2   Testing the Hypotheses

To test the relationship between earnings changes and life satisfaction more formally, we explore which of our hypotheses that have empirical leverage. We want to emphasize that our setting is not experimental. Thus, the tests we present here are not definite and further research is needed. However, we diligently study whether confounding variables or reversed causality could drive the results.

First, we test **H1**, i.e. diminishing sensitivity. To make these results presentable in a table and allow for kink at the reference point, we run the same regressions as presented in Figure 4, but instead of a spline we fit second and fourth order polynomials on the positive and negative domain separately. We do so both with and without year fixed effects. The results are presented in Table 1.

13

Table 1: Regression Results

		Without year fixed effects				With year fixed effects			
		(1)		(2)		(3)		(4)	
		2nd order		4th order		2nd order		4th order	
		Coef	SE	Coef	SE	Coef	SE	Coef	SE
	y-r	0.15^{***}	(0.02)	0.20^{***}	(0.05)	0.09^{***}	(0.02)	0.10^{**}	(0.05)
Gain	$(y\text{-}r)^2$	-0.05^{***}	(0.01)	-0.15^{*}	(0.08)	-0.03^{**}	(0.01)	-0.06	(0.08)
domain	$(y\text{-}r)^3$			0.04	(0.03)			0.01	(0.03)
	$(y\text{-}r)^4$			-0.00	(0.00)			-0.00	(0.00)
	y-r	0.25^{***}	(0.03)	0.37^{***}	(0.05)	0.27^{***}	(0.03)	0.41^{***}	(0.05)
Loss	$(y\text{-}r)^2$	0.07^{***}	(0.02)	0.35^{***}	(0.09)	0.08^{***}	(0.02)	0.39^{***}	(0.09)
domain	$(y\text{-}r)^3$			0.15^{***}	(0.05)			0.16^{***}	(0.05)
	$(y\text{-}r)^4$			0.02^{***}	(0.01)			0.02^{***}	(0.01)
Within r^2			0.002		0.002		0.016		0.016
AIC			760,992		760,976		757,325		757,311

Notes: * p<0.10, ** p<0.05, *** p<0.01. All models contain individual fixed effects and individual-level clustered standard errors. The reference point, r, is set to last year's monthly earnings plus the average yearly change in monthly earnings. $n = 251,437$.

The first model uses a piece-wise quadratic relationship between changes in earnings from the reference point and life satisfaction. All four coefficients are highly significant and have the necessary signs to generate diminishing returns. The same applies if we include year fixed effects. This is in line with **H1**. In the second and fourth specification we use a 4th order polynomial. In the loss domain all higher order terms are highly significant whereas there is less significance in the positive domain. Our preferred specification is to use a fourth order polynomial without year fixed effects. We prefer not using year fixed effects, since we believe they take out too much of the variation we want to exploit. We think business cycle variation should be captured by the value function rather than fixed away.

Given the evidence for diminishing sensitivity in the previous specifications we next test **H2**, that is, whether loss aversion is present. To test the hypotheses we need to calculate various statistics that are a function of both the level and slope of the predicted S-curve. To generate confidence bands around these test statistics we bootstrap 1000 resamples at individual level clusters. We use the percentile method to derive the confidence bands. Panel (a) of Figure 7 shows the generated S-curve with bootstrapped standard errors. It is equivalent to the predictions from the fourth order polynomial without year fixed effects presented in Table 1. Panel (b) of Figure 7 shows the corresponding slope of the curve. The slope at the reference point is significantly larger in the negative domain than in the positive domain. This is in line with **H2**.

Figure 7: Bootstrapped Hypothesis Testing

(a) H1: Level

(b) H2: Slope

(c) H2': Level Difference

(d) H2": Level Ratio
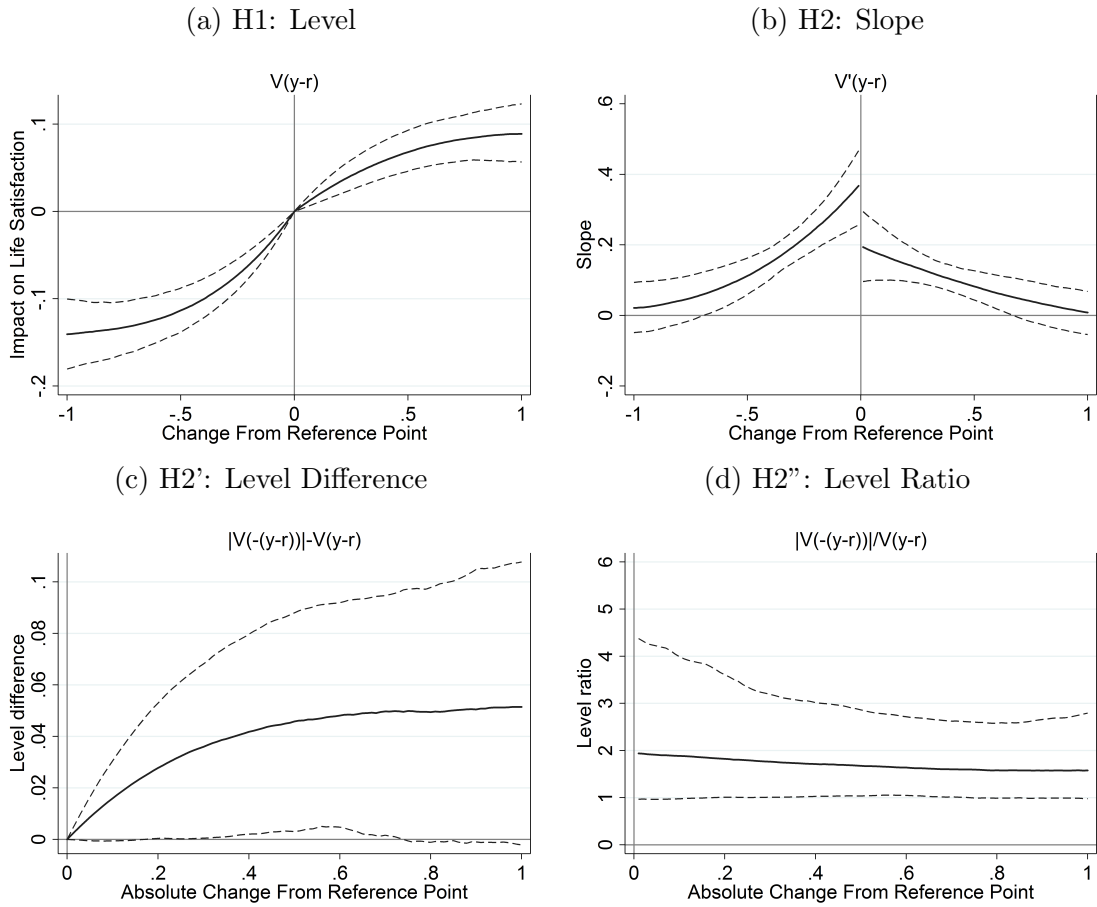
Notes: The graphs display the nature of loss aversion by computing differences and ratios in the impact of losses and gains. Sample size in all graphs is 251,437. Confidence bands are generated by 1000 bootstraps resamples at individual level clusters using the percentile method.

We test **H2'** in panel (c). The graph shows whether a decrease of a given size from the reference point has a larger impact on life satisfaction than a similar increase. This is the case for all changes from the reference point, which is in line **H2'**. In panel (d) of Figure 7 we look at whether loss aversion is constant. We test whether regardless of the size of the change from the reference point, a loss of a given size is worse by a constant multiplier than an equivalent gain. This appears to be the case with a coefficient of around 2. The number is decreasing slightly over the interval but with fairly large confidence intervals. We want to stress this result changes slightly if we use other specifications, such as splines or other polynomials. With some specifications loss aversion is slightly increasing, with others it is slightly decreasing. It always fluctuates around 2, but with rather wide confidence bands. This is consistent with the 2.25 reported by Tversky and Kahneman (1992).

In sum, we cannot falsify any of the different definitions of loss aversion. To

assure that these results are not dependent on the specific functional form we chose, we test the hypotheses using two alternatives to a 4th order polynomial. We use the power function as used in Carter and McBride (2013) and a more flexible power function as used in Vendrik and Woltjer (2007).

$$Power\ function: V(y_{it} - r_{it}) = \begin{cases} \beta_1(y_{it} - r_{it})^{\alpha_1} & : y_{it} - r_{it} > 0 \\ \beta_2(-(y_{it} - r_{it}))^{\alpha_2} & : y_{it} - r_{it} < 0 \end{cases}$$

$$Flexible\ power\ function: V(y_{it} - r_{it}) = \begin{cases} \beta_1 \frac{(1+y_{it}-r_{it})^{1-\alpha_1}-1}{1-\alpha_1} & : y_{it} - r_{it} > 0 \\ \beta_2 \frac{(1+y_{it}-r_{it})^{1-\alpha_2}-1}{1-\alpha_2} & : y_{it} - r_{it} < 0 \end{cases}$$

The results (without individual fixed effects to account for the non-linear specification) are shown in Figures A.2 and A.3. With these specifications, we are never able to negate any of the prior findings. Due to wide confidence intervals, however, we are unable to confirm all of the prior results, particularly whether loss aversion is constant.

Finally, we test **H3**; whether individuals' life satisfaction is increasing both in the positive distance from their reference point, and in the absolute amount of earnings that they have. This test is important to make sure the standard utility function is not behind the S-curve we observe, which in principle could be possible. If individuals only gain utility from their level of income, and changes in income are positively correlated with income levels, the S-curve could arise solely from expected utility theory. We again use a fourth order polynomial above and below the reference point. Now we also control for log earnings. We use a fourth order polynomial to be flexible about how the level of earnings is transmitted into life satisfaction. The corresponding predictions are graphed in Figure 8.

We appear to have evidence for the S-curve and a standard utility function in line with **H3**. This finding suggests that expected utility theory and prospect theory both play an independent role in transmitting income into life satisfaction. We get slightly less loss aversion when controlling for the standard utility function. This is likely because large negative income changes are associated with very low current levels of income, which now are captured by the diminishing marginal utility from expected utility theory.

In order to put the magnitudes in perspective, consider a person with the median earnings (1515 EUR). We can calculate the change in life satisfaction this person would experience from respectively prospect theory and expected utility theory (EUT) if he gained or lost a standard deviation of log earnings (0.34). If the person lost 0.34 log earnings he would experience a decrease in life satisfaction from

16

Figure 8: Testing H3

(a) Prospect Theory

(b) Expected Utility Theory
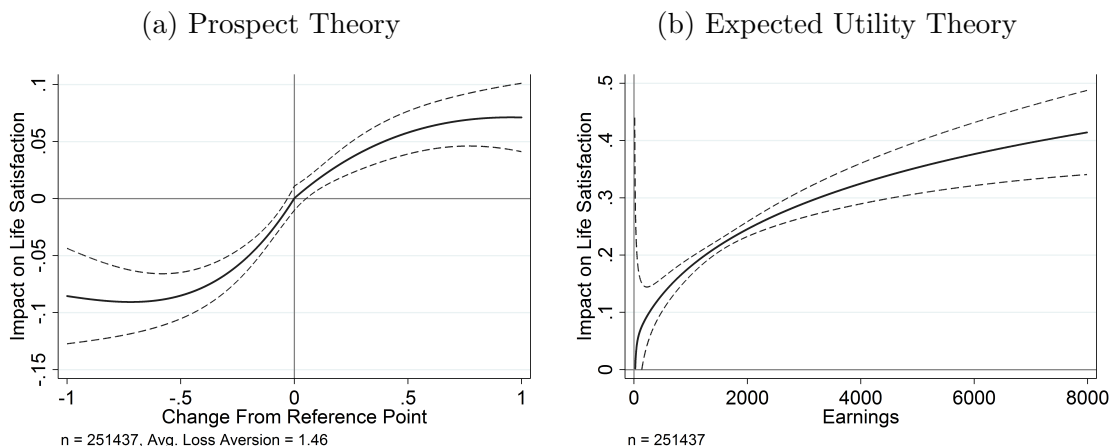
Notes: Predictions from a fixed effects regression of life satisfaction on a fourth order polynomial of log earnings and a piecewise fourth order polynomial of changes from the reference point. Both figures cluster standard errors at the individual level. *Avg. loss aversion* $= \frac{1}{100} \sum_{y_{it}-r_{it}=0.01}^{1} \frac{|V(-(y_{it}-r_{it}))|}{V(y_{it}-r_{it})}$, where $y_{it} - r_{it} = \{0.01, 0.02, ..., 1\}$.

prospect theory of 0.072 and a loss from EUT of about 0.031. If he gained 0.34 log earnings he would experience an increase in life satisfaction of 0.046 from prospect theory and of 0.035 from EUT. Hence, prospect theory seems to have a bigger impact on life satisfaction than expected utility theory for the median earner. For poorer individuals, diminishing utility kicks in and expected utility has a larger impact on life satisfaction.

4.3 Sources of Variation in Income Changes

Since our income changes are not randomized, it is relevant to analyze where they come from. Broadly speaking, our results can be driven by two sources. Firstly, they can be driven by income changes we hope are at play, such as changes in work hours, changes in tax codes, promotions, and job changes. These sources of variation can pose a potential threat if they have a direct impact on life satisfaction. Secondly, our results can be driven by any covariate that is positively associated with both life satisfaction and changes in income or vice versa. This could include lagged income, age, health status, family status etc. In this section we will argue that the first set of variables indeed seem to be a powerful source of variation, but that our S-curve survives even when we control for their direct effect on life satisfaction and for other covariates.

In Table 2 we regress changes in income on a set of variables belonging to each of the two potential sources. With regards to the first class, we control for changes in work hours, voluntary job changes, and involuntary job changes (voluntary: own

17

resignation, employee requested transfer within company, end of self-employment, involuntary: terminated by employer, temporary contract expired, company transferred employee, company closed down). With regards to the second class, we control for changes in health status (proxied by yearly doctor visits), changes in marital status, change in number of kids, and lagged income.

Table 2: Predictors of Income Changes

	Coef	SE
Change in Weekly Work Hours	0.0071***	(0.0002)
Voluntary Job Change	0.0239***	(0.0055)
Involuntary Job Change	−0.0538***	(0.0078)
Lagged Log Income	−0.5028***	(0.0058)
Change in Annual Doctor Visits	−0.0002***	(0.0001)
Kids	−0.0173***	(0.0016)
Single	−0.1183***	(0.0059)
Observations	209,543	
Within r^2	0.309	

Notes: * p<0.10, ** p<0.05, *** p<0.01. Fixed effects regression of changes in income from reference point on a number of potential mediators. Standard errors are clustered at the individual level.

Changes in work hours positively predict changes in income. As work hours are negatively correlated with life satisfaction, if anything, the direct impact of work hours has mitigated the S-curve from our main analysis. Positive changes in health and voluntary job changes positively predict changes in income. This is potentially problematic for our S-curve, since both are positively correlated with life satisfaction. Hence, we want to control for these to assure that our S-curve survives. Lagged income seems to have a lot of explanatory power, whereby individuals with initially high incomes experience lower income changes. Like work hours, this is not problematic for our S-curve as income and life satisfaction are positively correlated. Individuals that become single or get kids experience lower changes in income. It is not entirely clear how this could impact the observed S-curve.

In order to check if our S-curve is driven by any of these variables, we run our baseline regression of life satisfaction on a spline of changes in income while controlling for all of the mediating factors. The resulting predictions are shown in Panel (a) of Figure 9. The S-curve is almost unchanged, although we have less evidence for loss aversion.

The impact of work hours on income changes seem to be a particularly important

Figure 9: Predictions with Controls and Leisure Satisfaction

(a) Predictions with Controls

(b) Leisure Satisfaction

n = 209543, Avg. Loss Aversion = 1.32

n = 227863
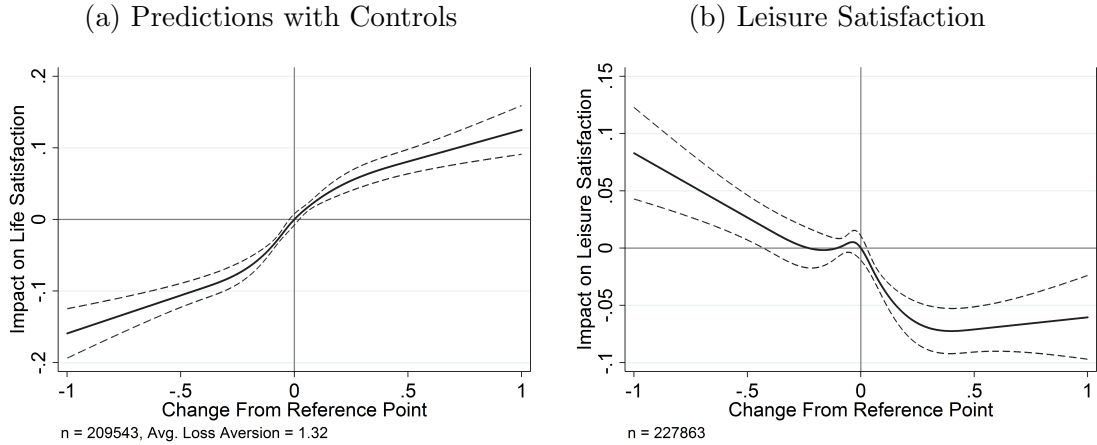
Notes: Panel (a) shows predictions from a fixed effects regression of life satisfaction on a spline of changes from the reference point. The regression controls for changes in work hours, voluntary and involuntary job changes, lagged log income, changes in yearly doctor visits, partnership status, and having kids. Panel (b) shows predictions from a fixed effects regression of leisure satisfaction (answers to how satisfied individuals are with their leisure on a scale from 0-10) on a spline of changes from the reference point.

driving force. This is apparent if we substitute life satisfaction in our baseline model with leisure satisfaction. As shown in Panel (b) of Figure 9, distance from the reference point is *negatively* associated with leisure satisfaction. This can be explained if great income changes come from individuals who increase their work hours accordingly. Again, if anything, this has decreased the strength of the S-curve we found in the main analysis.

4.4 Robustness Checks

The results until now were based on a number of assumptions with regards to the dependent variable, the income measure, the reference point etc. In this section we perform a number of alternative specifications to clarify when the main results are robust. Table 3 shows an overview of the robustness checks we make. Figures supporting all of these results are presented in the Appendix. We list whether we have evidence for diminishing sensitivity (DS) and loss aversion (LA) separately. We consider there to be support for loss aversion if $LA = \frac{1}{100} \sum_{y_{it}-r_{it}=0.01}^{1} \frac{|V(-(y_{it}-r_{it}))|}{V(y_{it}-r_{it})} > 1.5$, where $y_{it} - r_{it} = \{0.01, 0.02, ..., 1\}$. In all cases we use restricted cubic splines with five knots and control for individual fixed effects.

Table 3: Robustness Checks

Robustness check	Main specification	Alternative specification	DS	LA
Dependent variable	Life satisfaction	Personal income satisfaction	Yes	No
		Household income satisfaction	Yes	Yes
		Work satisfaction	Yes	No
		Happiness	No	No
Income measure	Self-reported monthly net earnings (CPI deflated)	Nominal monthly net earnings	Yes	Yes
		Monthly gross earnings	Yes	No
		Monthly gross HH income	Yes	Yes
		Annual net HH income	Yes	Yes
		Annual gross HH income	Yes	Yes
		Annual gross HH earnings	Yes	Yes
Survey	GSOEP	BHPS, monthly net earnings	Yes	No
		BHPS, monthly gross earnings	Yes	No
		BHPS, annual gross earnings	Yes	No
		BHPS, weekly net HH income	Yes	Yes
Reference point	Lagged earnings plus average yearly change in earnings	Predicted log earnings from AR(1)	Yes	Yes
		Predicted log earnings from AR(2)	Yes	Yes
		Peer mean by state	No	No
		Peer mean by region and educ	No	No
		Peer mean by region and sex	No	No
		Peer mean by region, age, sex, and educ	No	No
Transformation of dependent variable	Cardinal from 0-10	log(lifesat/(10-lifesat)) transformation	Yes	Yes
		Remove individuals at boundary	Yes	Yes
Independent variable	Difference from reference point	Income	No	-
		Log Income	No	-
		Annual doctor visits	No	-
		Annual sick days	No	-
Distance measure	Differences in log earnings	Percentage change in earnings	Yes	Yes
		Absolute change in earnings	Yes	No
Sample	All with non-missing values	Excl. 5% with highest/lowest earnings	Yes	Yes
		Excl. 5% with largest earnings changes	Yes	Yes
		Bottom 50%	Yes	Yes
		Top 50%	Yes	No
		Men	Yes	No
		Women	Yes	Yes
		West Germany	Yes	Yes
		East Germany	Yes	Yes
		1984-1999	Yes	No
		2000-2014	Yes	Yes

Notes: All results are based on fixed effects regressions with clustered standard errors at the individual level. A restricted cubic spline with 5 knots is used in all cases. *DS* indicates whether there is evidence for diminishing sensitivity, *LA* whether there is evidence for loss aversion (loss aversion factor greater than 1.5).

Dependent Variable

First we test if the S-curve holds with other subjective well-being variables than life satisfaction. To this end we utilize a battery of variables in GSOEP where individuals are asked how satisfied they are with certain domains of their life on a scale from 0–10. We use questions on satisfaction with job, personal income, and household income. In addition, we use a question on how often individuals have felt happy in the past four weeks. The answer categories to this question are very rarely, rarely, occasionally, often, and very often. The results when using these measures as the dependent variable are given in Figure A.4.

The results with the domain satisfaction variables are quite similar to using life satisfaction. In fact, the domain satisfaction results seem to generate more narrow confidence bans. This is not surprising since earnings changes are more important when people evaluate satisfaction with their income or job rather than with life as a whole. We find no evidence for loss aversion when using personal income satisfaction or work satisfaction. This could be because these measures omit the impact of income changes on leisure. Alternatively, it could be because individuals need larger earnings changes to be satisfied with their job than they need to be satisfied with their life in general.

When the question on happiness is used we get quite wide confidence bans suggesting that this variable is more noisy. Using happiness as the dependent variable generates neither diminishing sensitivity nor loss aversion.

Income Variable

Next, we test the robustness of our result to using other income measures. The main results were based on monthly self-reported net earnings expressed in constant 2010 EUR. We try six other income measures. We keep on using the yearly mean change in the income variable as the reference point. First, we use the same self-reported income but expressed in nominal values. As shown in Figure A.5 we still find support for the S-curve. Next, we use self-reported gross earnings rather than net earnings. We again have evidence for diminishing sensitivity, but this time loss aversion is minimal. We also try using self-reported monthly net household income. This variable includes all income sources so now we also include unemployed and people outside the labor market. With this specification we have evidence for both loss aversion and diminishing sensitivity.

Next, we use income variables from the Cross National Equivalent File, which contains yearly income measures at the household level. We use both gross earnings and gross/net income. Since the incomes are measured in the year prior to the

self-reported life satisfaction, we exclude all interviews conducted after March (31% of the sample), such that the income measure still is somewhat timely. Due to this timing effect, these variables make a great test for whether our main results were driven by reversed causality. Results using these income measures are given in Figure A.5. We find a perfect S-shape using all of the three measures, but with generally larger loss aversion. This could be because individuals have smaller expectations for household income than their personal income and hence use a lower reference point.

Survey

We try to see if our results replicate using the British Household Panel Survey (BHPS). The BHPS is an annual survey that ran from 1991 to 2008. The sample we use contains 143,000 observations spread across 24,000 individuals. The life satisfaction question is phrased slightly differently in BHPS. Respondents are asked *"How dissatisfied or satisfied are you with your life overall?"* on a scale from 1 (not satisfied at all) to 7 (completely satisfied). The BHPS has several income variables. Here we use four different measures. Results are displayed in Figure A.6.

First we use monthly net earnings, which is similar to our baseline specification with the German Socio-Economic Panel. We have evidence for diminishing sensitivity but none for loss aversion. This does not change if we use gross earnings rather than net monthly earnings. Next we try to use annual gross household earnings, as this was the variable in the GSOEP that generated the largest amount of loss aversion. We still have evidence for diminishing sensitivity but not for loss aversion. Finally, we try to use the measure that probably is closest to what individual feel at the moment they answer the survey: weekly net household income. Here we have evidence for both diminishing sensitivity and loss aversion.

In general, however, the results using the BHPS display less evidence for loss aversion and is more fragile to extreme observations. One explanation for the lack of loss aversion may be that Brits have higher expectations than Germans. Indeed, if Brits expect their earnings increase to be around the 80th percentile (such that 4 in 5 get lower earnings increase) then loss aversion re-emerges. It may also be, of course, that loss aversion simply is non-existing in the British sample.

Reference Point

We also vary the reference point. First we predict each individual's income using AR(1) and AR(2) models. We consider the predicted income as the reference point. Hence, income increases larger than predicted are considered positive deviations from the reference point and vice versa. Findings are given in Figure A.7. In both

models we find evidence for both diminishing sensitivity and loss aversion. These results are consistent with the idea that the reference is the rational expectation of future earnings.

We also use different mean peer earnings specifications as a reference point. This is what Vendrik and Woltjer (2007) used in their analysis. We construct four different mean peer earnings measures: i) by state, ii) by region and education level (three categories), iii) by region and gender, and iv) by region, education, gender, and age (10 year rolling). Although life satisfaction is increasing in the difference from the mean peer earnings, the relationship between life satisfaction and these reference points shows little evidence for diminishing sensitivity and none for loss aversion. Similar to Vendrik and Woltjer (2007), we therefore find no evidence for the predictions from prospect theory when using peer earnings as the reference point. This could either be because prospect theory does not apply to this reference point, or because the reference groups we can create based on the survey do not capture who individuals actually compare themselves with. It is plausible that individuals use specific colleagues, friends or family members as their reference group, which these broad measures have a hard time capturing.

Transformation of Dependent Variable

A possible concern is that diminishing sensitivity arises mechanically due to floor and ceiling effects. No matter the income change, individuals cannot report life satisfaction levels below 0 or greater than 10. We test whether this is driving the results in three ways. Firstly, we try to transform the dependent variable such that our new dependent variable equals $log(\frac{lifesat}{10-lifesat})$. With this transformation there is further between life satisfaction levels at the boundaries and closer between life satisfaction levels at the center. Hence, more weight is attached to changes close to the boundaries. As shown in Figure A.8, this does not change the results.

Secondly, we try to deal with the boundedness concern more directly by deleting observations where life satisfaction is reported to be 0, 1, 9, or 10. As also shown in in Figure A.8, this only serves to make our findings stronger.

Independent Variable

The third way in which we see if our results are driven by floor and ceiling effects is to regress life satisfaction on other variables, which we know from prior research is positively correlated with life satisfaction. If the boundedness is a concern, we should see S-curves also in these regressions. We already saw that this was not the case when we used peer income as the reference point. To study this further

we regress life satisfaction separately on income, log income, annual doctor visits, and annual sick days (in the latter two, we look for reversed S-shapes). Results are displayed in Figure A.9. In neither case do we see an S-shape. This points against floor and ceiling effects driving diminishing sensitivity.

Distance Measure

Next, we use percentage changes in earnings and absolute changes in earnings rather than differences in logs. We use the yearly mean percentage change in the sample (after discarding changes larger than 1000% income growth) as the reference change. For the absolute changes, we similarly use the mean absolute change (after discarding absolute changes greater than 10,000 EUR). In both cases the S-curve remains. This is presented in Figure A.10. The nature of the loss aversion changes, though.

Sample

Finally, we check if the results are driven by the tails of the distributions. We first look at outliers by removing the 5% most extreme values in terms of, respectively, earnings levels and changes in earnings. As shown in Figure A.11 in both cases the S-shape is unaltered. We also try to divide the sample according to whether individuals are in the top or bottom half of the income distribution. This helps explain if the results are driven by people that had very low or very high initial incomes. Overall, we get S-curves in both cases but with loss aversion differing quite a bit between the two groups.

Next, we divide the sample by gender, region (east/west) and survey year (before/after 2000). As shown in Figure A.12 we find strongest loss aversion for women, East Germany, and the later years of the sample.

In sum, we find that our results our robust to most specifications. However, the amount of loss aversion present changes quite a bit depending on the exact specification.

5   Conclusion

Prospect theory is one of the most canonical results in behavioral economics in the past century. The theory holds that individuals derive value not from their absolute level of income, but rather from changes in income with respect to a reference point. Two of the ingredients of the theory are that individuals display loss aversion (losses are valued more heavily than gains) and diminishing sensitivity (large changes from

the reference point have diminishing marginal impacts), generating an S-curve with a kink at the reference point.

In this paper, we provided a comprehensive test for these two phenomena using data on life satisfaction from the German Socio-Economic Panel. Rather than considering choices under risk, we looked at non-experimental realized outcomes. We used experienced utility rather than decision utility as the measure of value and data on last year's earnings to generate a reference point. Our most naive results, simply regressing life satisfaction on a spline of changes from the reference point, revealed an S-curve strikingly similar to experimental evidence. The size of our data set (250,000+ observations) allowed us to test various specific definitions of loss aversion. We found that earnings losses have about 2 times greater impact on life satisfaction than earnings gains. Our main results are robust to a number of alterations, including using other subjective well-being variables, applying different definitions of income, and using other reference points. The S-curves remains even with a large set of controls and alternative definitions of the reference point.

If we in addition control for the absolute level of earnings (rather than only earnings changes), the S-curves maintains together with a standard utility function, suggesting that both the income level and changes in income matter. For the median earner the value function has slightly greater influence on life satisfaction than the standard utility function.

To our knowledge, we are the first to find that the predictions from prospect theory with respect to loss aversion and diminishing sensitivity hold with life satisfaction outside of an experiment. This result gives support to applying the value function in policy analyses as a significant source of well-being.

References

Johannes Abeler, Armin Falk, Lorenz Goette, and David Huffman. 2011. Reference points and effort provision. *The American Economic Review* 101, 2 (2011), 470–492.

Alessandro Barattieri, Susanto Basu, and Peter Gottschalk. 2014. Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics* 6, 1 (2014), 70–101.

Nicholas C Barberis. 2013. Thirty years of prospect theory in economics: A review and assessment. *The Journal of Economic Perspectives* 27, 1 (2013), 173–195.

David Bowman, Deborah Minehart, and Matthew Rabin. 1999. Loss aversion in a consumption–savings model. *Journal of Economic Behavior & Organization* 38, 2 (1999), 155–178.

Christopher J Boyce, Alex M Wood, James Banks, Andrew E Clark, and Gordon DA Brown. 2013. Money, Well-Being, and Loss Aversion: Does an Income Loss Have a Greater Effect on Well-Being Than an Equivalent Income Gain? *Psychological science* (2013), 2557–2562.

Steven Carter and Michael McBride. 2013. Experienced utility versus decision utility: Putting the 'S' in satisfaction. *The Journal of Socio-Economics* 42 (2013), 13–23.

Jan-Emmanuel De Neve, George W Ward, Femke De Keulenaer, Bert Van Landeghem, Georgios Kavetsos, and Michael Irwin Norton. 2014. Individual Experience of Positive and Negative Growth is Asymmetric: Global Evidence from Subjective Well-being Data. *Available at SSRN 2506600* (2014).

Gustav Theodor Fechner. 1912. Elements of psychophysics. Sections VII and XVI. *Herbert Sidney Langfeld (trad.), en Green CD, Classics in the History of Psychology. An internet resource, York University, Toronto [http://psychclassics.yorku.ca/Fechner/]. Consultado el* 16 (1912).

Ada Ferrer-i Carbonell. 2005. Income and well-being: An empirical analysis of the comparison income effect. *Journal of Public Economics* 89, 5 (2005), 997–1019.

Shane Frederick and George Loewenstein. 1999. Hedonic adaptation. In *Well-being: Foundations of hedonic psychology*, Daniel Kahneman, Edward Diener, and Norbert Schwarz (Eds.). Russell Sage Foundation.

Eugene Galanter. 1990. Utility functions for nonmonetary events. *The American Journal of Psychology* (1990), 449–470.

Glenn W Harrison and E Elisabet Rutström. 2009. Expected utility theory and prospect theory: One wedding and a decent funeral. *Experimental Economics* 12, 2 (2009), 133–158.

Daniel Kahneman. 2003. A perspective on judgment and choice: Mapping bounded rationality. *American psychologist* 58, 9 (2003), 697.

Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* (1979), 263–291.

Miles Kimball, Ryan Nunn, and Dan Silverman. 2015. Accounting for Adaptation in the Economics of Happiness. *NBER Working Paper* (2015).

Veronika Köbberling and Peter P Wakker. 2005. An index of loss aversion. *Journal of Economic Theory* 122, 1 (2005), 119–131.

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

Steven D Levitt and John A List. 2007. What do laboratory experiments measuring social preferences reveal about the real world? *The Journal of Economic Perspectives* 21, 2 (2007), 153–174.

George Loewenstein and Peter A Ubel. 2008. Hedonic adaptation and the role of decision and experience utility in public policy. *Journal of Public Economics* 92, 8 (2008), 1795–1810.

Erzo FP Luttmer. 2005. Neighbors as Negatives: Relative Earnings and Well-Being. *The Quarterly Journal of Economics* (2005), 963–1002.

Jason Riis, George Loewenstein, Jonathan Baron, Christopher Jepson, Angela Fagerlin, and Peter A Ubel. 2005. Ignorance of hedonic adaptation to hemodialysis: A study using ecological momentary assessment. *Journal of Experimental Psychology: General* 134, 1 (2005), 3.

Andrei Shleifer. 2012. Psychologists at the gate: A review of Daniel Kahneman's thinking, fast and slow. *Journal of Economic Literature* 50, 4 (2012), 1080–1091.

Amos Tversky and Daniel Kahneman. 1991. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics* (1991), 1039–1061.

Amos Tversky and Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5, 4 (1992), 297–323.

Peter A Ubel, George Loewenstein, and Christopher Jepson. 2005. Disability and sunshine: Can hedonic predictions be improved by drawing attention to focusing illusions or emotional adaptation? *Journal of Experimental Psychology: Applied* 11, 2 (2005), 111–123.

Bernard MS Van Praag and Paul Frijters. 1999. The Measurement of Welfare and Well-Being: The Leyden Approach. *Well-Being: Foundations of Hedonic Psychology* (1999), 413.

Maarten CM Vendrik and Geert B Woltjer. 2007. Happiness and loss aversion: Is utility concave or convex in relative income? *Journal of Public Economics* 91, 7 (2007), 1423–1448.

A   Supplementary Figures

Figure A.1: Histogram of Changes in Nominal Log Earnings

Notes: Histogram of changes in nominal monthly net earnings. Despite the income variable being self-reported and net of taxes, we see substantial nominal wage rigidities.

Figure A.2: Hypothesis Testing with Power Function

(a) H1: Level
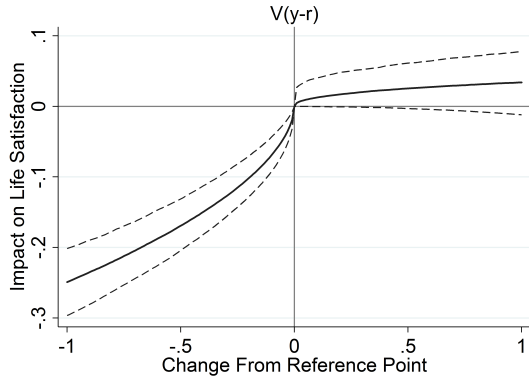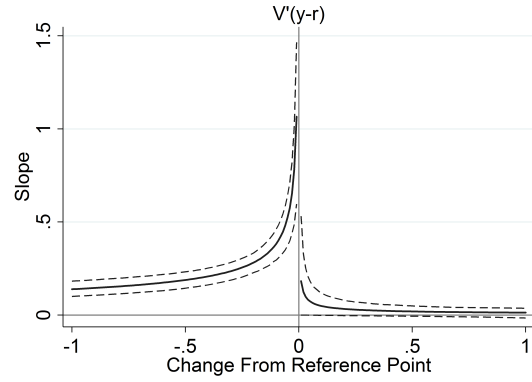
(b) H2: Slope

(c) H2': Level Difference

(d) H2": Level Ratio
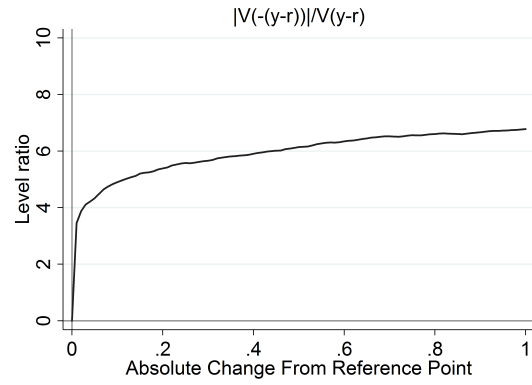
Notes: Hypothesis testing based on the assumption that the value function takes the form:

$$V(y_{it} - r_{it}) = \begin{cases} \beta_1(y_{it} - r_{it})^{\alpha_1} & : y_{it} - r_{it} > 0 \\ \beta_2(-(y_{it} - r_{it}))^{\alpha_2} & : y_{it} - r_{it} < 0 \end{cases}$$

Confidence bans are generated through bootstrapping 1,000 resamples at individual level clusters. Sample size in all graphs is 251,437. In panel (d) the confidence ban is outside the window of the graph. The parameter estimates are as follows:

Parameter	Coef	SE
α_1	0.417	(0.467)
α_2	0.549***	(0.082)
β_1	0.032*	(0.019)
β_2	−0.251***	(0.020)

Notes: * p<0.10, ** p<0.05, *** p<0.01.

30

Figure A.3: Hypothesis Testing with Flexible Power Function

(a) H1: Level

(b) H2: Slope

(c) H2': Level Difference

(d) H2": Level Ratio
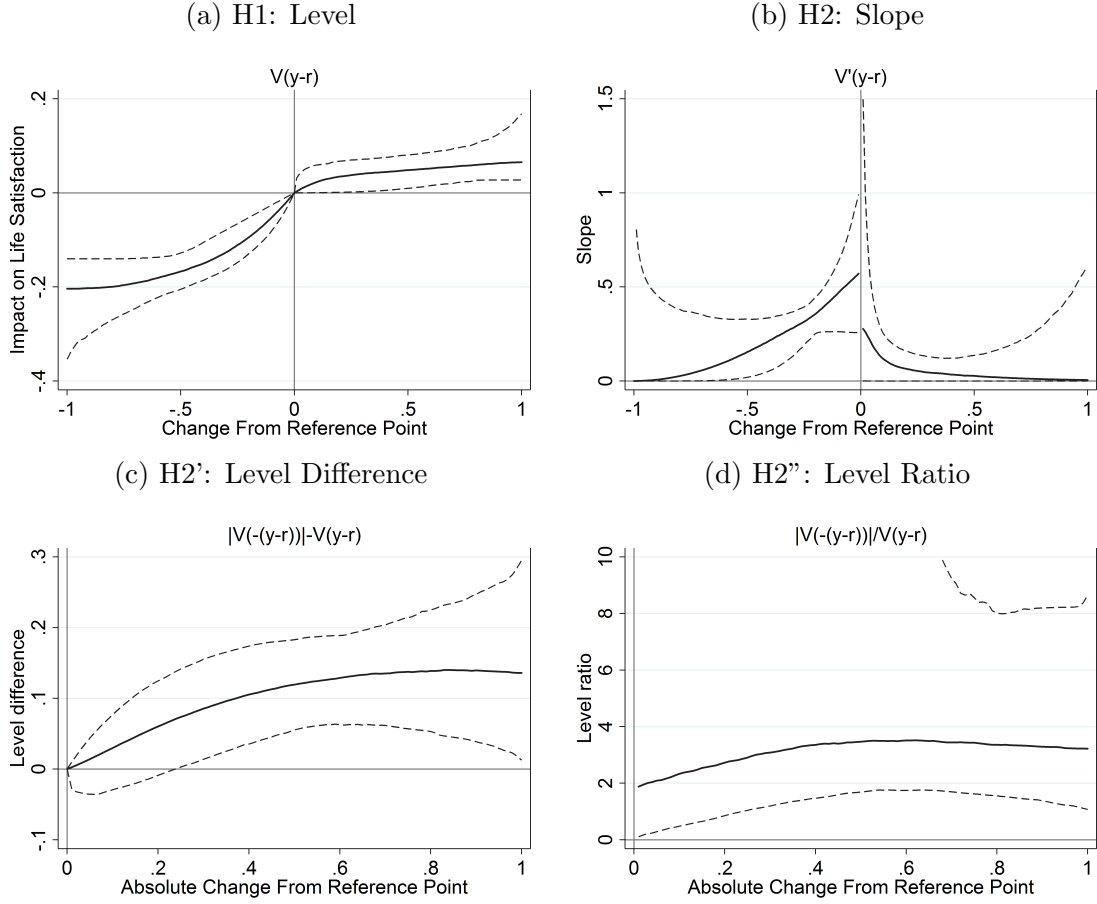
Notes: Hypothesis testing based on the assumption that the value function takes the following form:

$$V(y_{it} - r_{it}) = \begin{cases} \beta_1 \frac{(1+y_{it}-r_{it})^{1-\alpha_1}-1}{1-\alpha_1} & : y_{it} - r_{it} > 0 \\ \beta_2 \frac{(1+y_{it}-r_{it})^{1-\alpha_2}-1}{1-\alpha_2} & : y_{it} - r_{it} < 0 \end{cases}$$

Confidence bans are generated through bootstrapping 1000 resamples at individual level clusters. In panel (d) the upper confidence ban is partly outside the window of the graph. The parameter estimates are as follows:

Parameter	Coef	SE
α_1	7.250	(5.156)
α_2	−1.803**	(0.720)
β_1	0.339	(0.210)
β_2	0.559***	(0.106)

Notes: * p<0.10, ** p<0.05, *** p<0.01.

Figure A.4: Changing the Dependent Variable

(a) Personal Income Satisfaction

n = 122872, Avg. Loss Aversion = 1.17

(b) Household Income Satisfaction

n = 249660, Avg. Loss Aversion = 1.97

(c) Work Satisfaction

n = 246632, Avg. Loss Aversion = .58

(d) Happiness

n = 91985
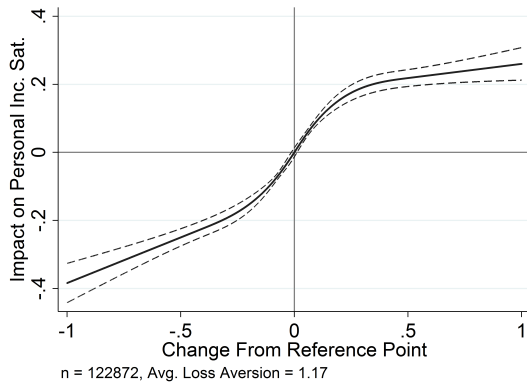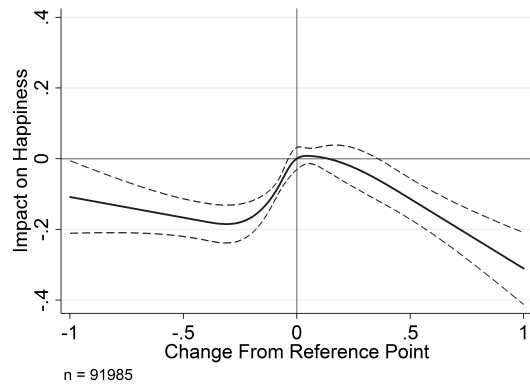
Notes: All variables range from 0-10. Happiness is the answer to how often individuals have felt happy in the past four weeks. The answer categories to this question are very rarely, rarely, occasionally, often, and very often (coded 0, 2.5, 5, 7.5, 10, such that the range is comparable to the domain satisfaction questions).

Figure A.5: Changing the Income Measure

(a) Nominal Monthly Net Earnings

n = 256179, Avg. Loss Aversion = 1.66

(b) Monthly Gross Earnings

n = 251507, Avg. Loss Aversion = 1.15

(c) Monthly Gross HH Income

n = 438419, Avg. Loss Aversion = 1.75

(d) Annual Net HH Income

n = 312307, Avg. Loss Aversion = 4.06

(e) Annual Gross HH Income

n = 296362, Avg. Loss Aversion = 2.77

(f) Annual Gross HH Earnings

n = 238176, Avg. Loss Aversion = 5.36

Notes: Predicted values from fixed effects regressions of life satisfaction on changes in log income from the reference point. The bottom four figures use the entire GSOEP sample including unemployed and people outside the labor market. Panel (c), (d), and (e) include all income sources, not only earnings. The three annual figures contain income in the year prior to the survey. Individuals surveyed after March are excluded to align the timing of these measures with the life satisfaction answers.

Figure A.6: Changing the Survey: BHPS

(a) Monthly Net Earnings

n = 49114, Avg. Loss Aversion = .04

(b) Monthly Gross Earnings

n = 49479, Avg. Loss Aversion = .26

(c) Annual Gross HH Earnings

n = 85276, Avg. Loss Aversion = .64

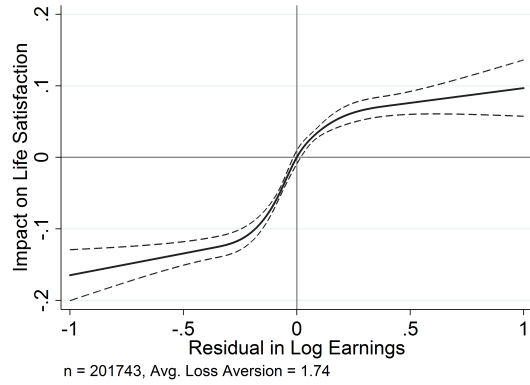(d) Weekly Net HH Income

n = 117475, Avg. Loss Aversion = 2.92

Notes: Predicted values from fixed effects regression of life satisfaction on changes from the reference point. The four income variables are from the British Household Panel Survey.

Figure A.7: Changing the Reference Point

(a) AR(1)

(b) AR(2)

n = 251437, Avg. Loss Aversion = 1.58

n = 201743, Avg. Loss Aversion = 1.74

(c) Peer earnings (by state)

(d) Peer earnings (by region & educ.)

n = 322138, Avg. Loss Aversion = .77

n = 322138, Avg. Loss Aversion = .98
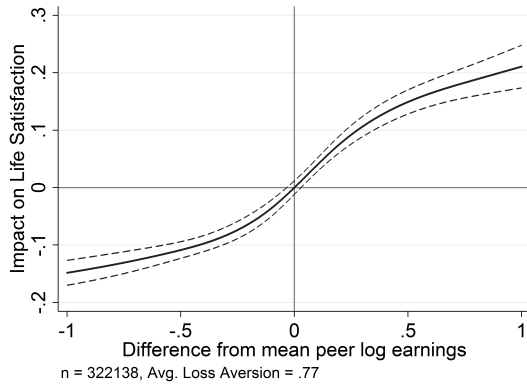
(e) Peer earnings (by region & sex)

(f) Peer earnings (by region/age/sex/educ.)

n = 322138, Avg. Loss Aversion = 1.02

n = 319855, Avg. Loss Aversion = 1.11

Notes: Predicted values from regressing life satisfaction on changes from the reference point using a variety of different reference points. The bottom four panels use the mean log earnings of a peer group as a reference point. Education is split into three categories and age groups are rolling 10 year intervals.

35

Figure A.8: Transformations of the Dependent Variable

(a) y=log(lifesat/(10-lifesat))   (b) Removing 0-1 & 9-10 in Life Sat.

n = 239617, Avg. Loss Aversion = 1.65   n = 207376, Avg. Loss Aversion = 3.91
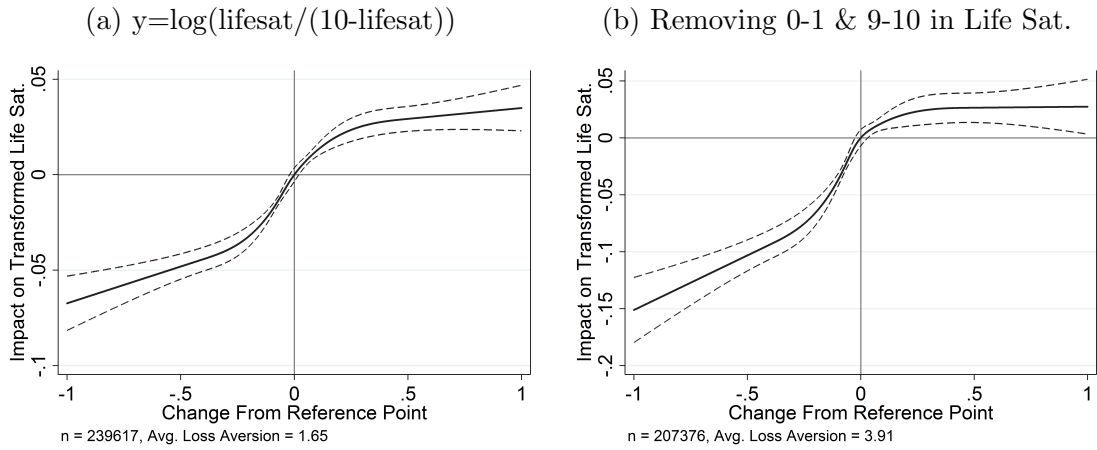
Notes: Panel (a) transforms the left-hand side variable to equal $log(\frac{lifesat}{10-lifesat})$ (individuals with 0 or 10 in life satisfaction are discarded). Panel (b) removes individuals who report 0, 1, 9, or 10 in life satisfaction.

Figure A.9: Changing the Independent Variable

(a) Income   (b) Log Income

n = 255760   n = 255760

(c) Annual Doctor Visits   (d) Annual Sick Days

n = 242894   n = 229304
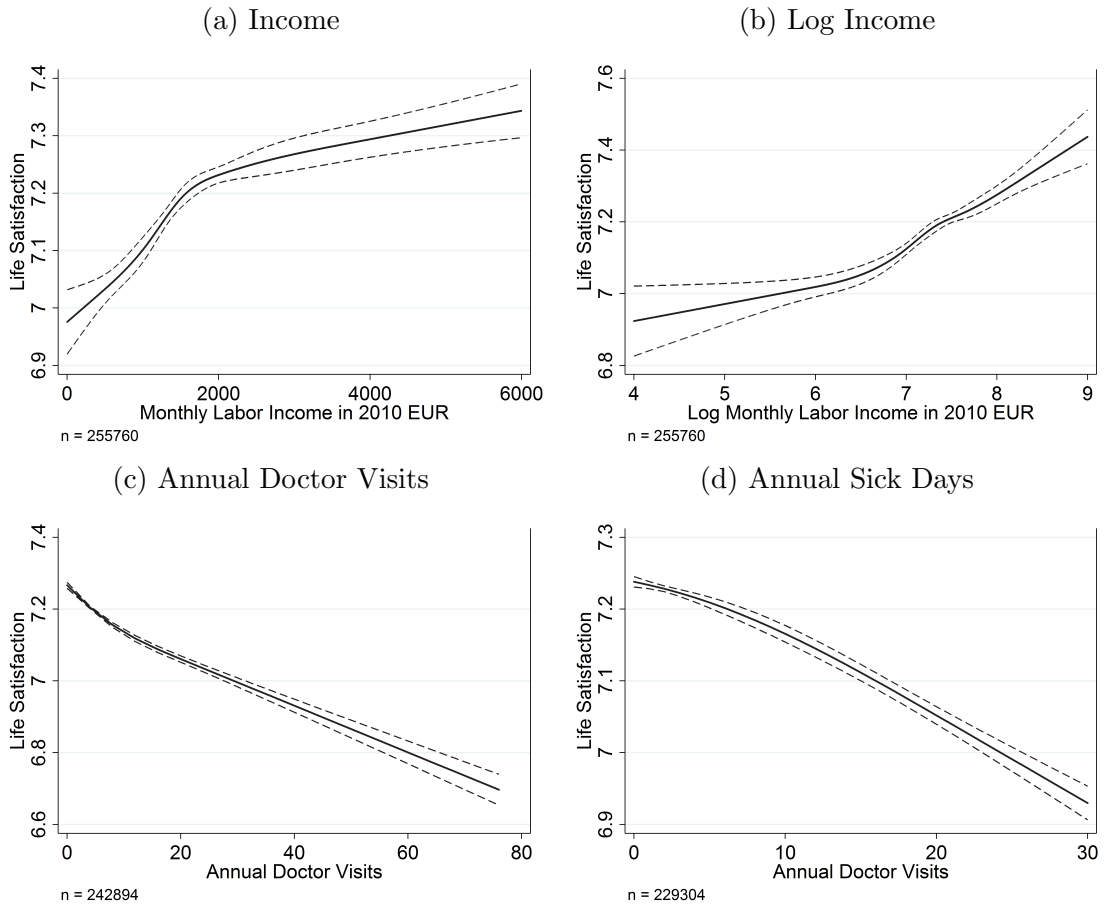
Notes: All panels regress life satisfaction on a spline of the variable in question using individual fixed effects.

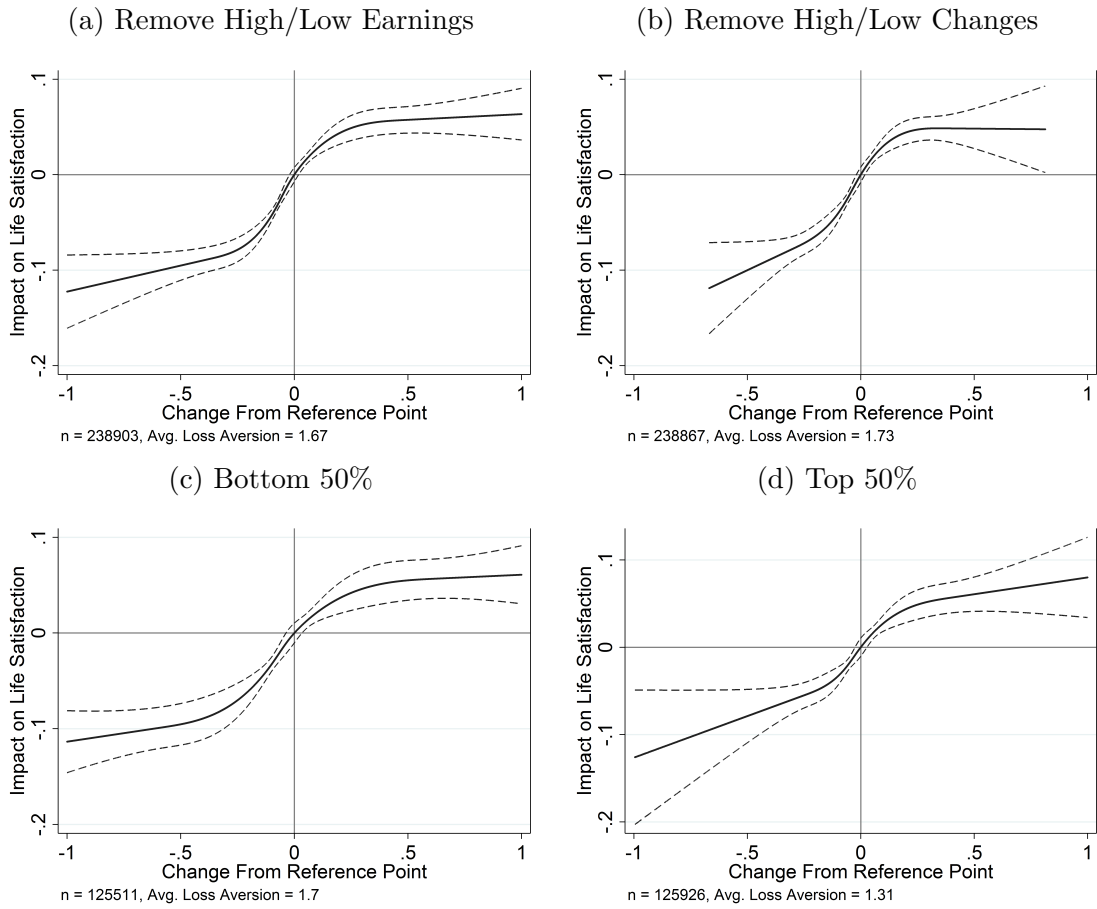Figure A.10: Different Distances from Reference Point

(a) Growth in Earnings

(b) Absolute Change in Earnings

Notes: Predicted values from regressing life satisfaction on difference distances from the reference point. As distance measures we use the percentage change from the year before and the absolute change from the year before. The income variable is monthly net earnings in 2010 EUR. The reference point is defined as, respectively, the average yearly growth and average yearly change in earnings after discarding extreme values.

Figure A.11: Changing the Sample 1/2

(a) Remove High/Low Earnings

n = 238903, Avg. Loss Aversion = 1.67

(b) Remove High/Low Changes

n = 238867, Avg. Loss Aversion = 1.73

(c) Bottom 50%

n = 125511, Avg. Loss Aversion = 1.7

(d) Top 50%

n = 125926, Avg. Loss Aversion = 1.31

Notes: Predicted values using only parts of the sample. Panel (a) excludes individuals below the 2.5 percentile or above the 97.5 percentile of the income distribution. Panel (b) excludes individuals below the 2.5 percentile or above the 97.5 percentile of the 'change in log income' distribution. The bottom two graphs use only the lower or upper half of the respondents by their income.

Figure A.12: Changing the Sample 2/2

(a) Men

(b) Women

n = 139671, Avg. Loss Aversion = 1.37

n = 111766, Avg. Loss Aversion = 2.12

(c) West Germany

(d) East Germany

n = 202779, Avg. Loss Aversion = 1.53

n = 48658, Avg. Loss Aversion = 1.94

(e) 1984-1999

(f) 2000-2014

n = 88839, Avg. Loss Aversion = 1.12

n = 162598, Avg. Loss Aversion = 2.31

Notes: Predicted values using only parts of the sample. All regressions include individual fixed effects and cluster standard errors at the individual level.