Negative Emotions, Income, and Welfare: Causal Estimates from the PSID[☆]

David Clingingsmith^{a,*}

^a*Department of Economics, Case Western Reserve University, Cleveland, Ohio, 44106.*

Abstract

I use instrumental variables to estimate the causal effect of family income on the frequency with which individuals experience negative emotions. Doubling income reduces the experience of negative emotions by 0.26 SD on average. Percentage changes in income have a constant effect on negative emotion for family incomes below \$80,000. Above \$80,000, the effect of percentage changes declines, reaching zero at \$200,000. A dollar change in family income has an eight times larger effect at the 20th percentile of income than the 80th percentile. Effects of income are similar on the high levels of negative emotion characteristic of mental illness, except there is no satiation.

Keywords: income, emotion, affect, well-being, welfare, mental illness

JEL: D6, I3, J3

1. Introduction

Income is widely used as a measure of individual welfare in economics. The theory of revealed preference holds that income produces well-being indirectly by enabling the acquisition of goods and services. Income is thus an indirect measure of welfare. Since the 1970s, a literature has developed that seeks to measure subjective well-being (SWB) more directly through survey questions. An important task of this literature has been to understand the relationship between SWB and income.¹

[☆]Thanks to Carol Graham and two anonymous referees as well as Silke Forbes, Erzo F. P. Luttmer, Matt Sobel, Justin Sydnor, and Mark Votruba for helpful comments and suggestions. Andrew Loucky and Anthony Gatti provided excellent research assistance. The Ohio Supercomputer Center provided a grant of computing resources. The collection of the PSID data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and the National Science Foundation under award number 1157698.

*Corresponding author.

Email address: david.clingingsmith@case.edu (David Clingingsmith)

¹Relevant contributions include Easterlin (1974); Diener et al. (1993); Oswald (1997); Graham and Pettinato (2002); Blanchflower and Oswald (2004); Stevenson and Wolfers (2008); Deaton (2008); and
Preprint submitted to Elsevier *July 6, 2016*

One advantage of direct measures of well-being is that they enable us to address the welfare effects of income redistribution. Neoclassical microeconomic theory, grounded in revealed preference, is unable to do so because it relies on Pareto improvement to judge that welfare is greater in one state of the world than another.² This criterion cannot be used to assess the welfare effects of redistributive policies that result in losses of income to some and gains to others.

In this paper, I estimate the causal effect of income on the frequency with which people experience negative emotions. The experience of negative emotions is one of several complementary measures of SWB. I construct estimates of the marginal effect of income across the income distribution, which enables me to assess the welfare effects of redistributive transfers. At present, Pischke (2011) is the only other work I am aware of that has attempted to identify the causal relationship running from income to SWB.³

Most of the existing literature on income and SWB uses *evaluative* measures, such as life satisfaction or Cantril's ladder, that ask people to assess their well-being on a scale. Negative emotion, by contrast, is an *experience* measure of well-being.⁴ The present paper extends the recent work of Kahneman and Deaton (2010), who study the cross-sectional correlation of income with both experience and evaluative measures of well-being, by using panel data, introducing a richer measure of negative emotion, and providing causal identification.

My measure of negative emotion, the K6 index, aggregates the frequency with which a person has experienced six negative emotions during the past 30 days into a 0-24 scale.⁵ Lower K6 scores indicate greater SWB. The Panel Study of Income Dynamics (PSID)

Kahneman and Deaton (2010).

²A Pareto improvement is a change in an allocation that makes at least one person better off without making anyone worse off.

³Pischke and Schwandt (2012) discusses difficulties in the identification employed in Pischke (2011) that the present paper avoids.

⁴White and Dolan (2009) argue that a comprehensive measure of SWB requires knowing both the experience of positive and negative emotion, absolute and relative measures of life satisfaction, and a measure of life purpose. Others see a subjective inventory as an incomplete account of well-being, and wish to include or substitute other concepts such as capability Sen (e.g. 1993). Kahneman et al. (1997) discusses how experienced well-being is related to the classical utilitarians' notion of utility.

⁵The index was developed to measure serious mental illness in surveys (Kessler et al., 2003).

has collected the K6 from survey respondents in six waves since 2001 and is the main source of data in this paper (PSID, 2013).⁶

Identification of the causal effect of income on SWB is challenging because there are many plausible shocks that are likely to affect both variables. Panel data techniques, such as including a lag of the outcome among the regressors, are unlikely to be flexible enough to fully capture such shocks. I therefore use an instrumental variables approach. Because marginal effects across the distribution of income are of key interest, the instrument should have a strong first stage across its entire support in addition to satisfying the exclusion restriction. Intuitively, we need a source of exogenous variation that moves income substantially for individuals at all income levels. The policy changes that often serve as convincing instruments typically affect only a small segment of the population and will therefore not meet our needs.

I turn to variation in average wages across industries and geographic to build my instrument. I compute average weekly earnings by single-digit industry, geographic region, and year using the Current Population Survey (CPS). Then I assign these average wage estimates to individuals in the PSID based on their industry and geographic region in the *first year* in which they appear in the panel.⁷ This method of linking average wages to PSID individuals ensures the instrument will not be affected by subsequent choices individuals make about their industry or region. Since the initial selection into an industry-region might be related to unobserved characteristics, I include a set of initial-year industry-region dummies as part of the strategy. This means the causal effect of income is identified by changes in average wages in an individual's initial industry-region. The instrument is strongly correlated with income, with first-stage F-statistics above 60.

My preferred IV specification shows that a doubling of prior-year family income leads to an average decline of 0.96 in the K6 score ($p < 0.05$).⁸ The effect is large compared to the average K6 score of 3.3 points and is equal to about 0.26 standard deviations (SD).

⁶Repeated cross-sectional data on the K6 has also been collected in the National Health Interview Survey (NHIS) since 1997.

⁷This approach is closely related to Bartik's (1991) shift-share instrument.

⁸Family income is always expressed in 2013 dollars in this paper.

Nonparametric IV techniques reveal how the marginal effect of family income varies between family incomes of \$10,000 to \$200,000. The marginal effect in percentage terms is around -0.85 for family incomes between \$10,000 and \$80,000. It then begins to fall quickly, reaching -0.18 at \$160,000 and ending near zero at \$200,000, suggesting satiation at that point. In dollar terms, the marginal effect of income falls steadily, though at a decreasing rate, from -0.06 per \$1,000 of income at family income of \$10,000 to zero at family income of \$200,000. In other words, the marginal benefit of income in terms of negative emotional experience is highest for those with the lowest incomes.

We can use these estimates to compute the welfare effects of transfers between individuals at different parts of the distribution. A transfer of \$1,000 to someone at the 20th percentile from someone at the 80th percentile (\$27,000 and \$113,000, respectively) would reduce the experience of negative emotion by the 20th-percentile person *eight times* more than it would increase the experience of negative emotion by the 80th-percentile person. While this estimate of welfare effects comes with a number of caveats, it suggests that redistributive transfers will be welfare improving even if the efficiency losses of making them are quite high.

This paper makes several contributions to the large existing literature on SWB. First, I offer better-identified estimates of the effect of income on SWB than exist in the literature through the combined use of panel data and instrumental variables. Relatedly, I also offer better-identified estimates of the conditional effects of age, gender, and race and offer estimates of the effect of income on serious mental illness (SMI). Second, I estimate marginal effects of income across the distribution and use those estimates to compute the effects of redistributive transfers on welfare. Third, I introduce the K6 index to the literature as a measure of experienced well-being and make comparisons to measure of life satisfaction. I address how my contributions are linked to the literature in more detail in the final section of the paper.

Section 2 describes the datasets and the construction of the instrument. Empirical specifications, a discussion of identification, and explanation of the estimation techniques

used in the analysis are covered in section 3. The results, including the first stage and IV sensitivity analysis, are found in section 4. Section 5 discusses the results and their relationship to the literature.

2. Data

The primary data source for the analysis is the Panel Survey of Income Dynamics (PSID). The PSID is a longitudinal survey of United States residents that began in 1968 with a core set of families. The core families comprised two samples, a nationally representative sample and an oversample of low-income families. Since then, the PSID has expanded to include the descendants of these individuals, their spouses and spouse equivalents, and families. In 1997 and 1999, additional households were added to reflect immigration since 1968. In the analysis to follow, I work with individuals who entered the sample both through the two core samples and the immigrant sample, whether as original members, descendants, or spouses. I use the 2001, 2003, 2007, 2009, 2011, and 2013 waves. The PSID provides weights that can be used to compute statistics that are nationally representative either cross-sectionally or longitudinally.

2.1. Negative Emotion and Life Satisfaction Measures in the PSID

The PSID began including the K6 Index of Non-specific Psychological Distress in the 2001 round. The K6 was originally developed by Ronald Kessler and colleagues (2002; 2003) to enable researchers to measure serious mental illness in surveys. To develop this tool, the group first had a sample of adults answer an extensive set of questions related to their mental state. They then gave a subset of sample members a blinded face-to-face clinical evaluation by a psychiatrist for serious mental illness. Item response theory was used to select a subset of six questions that provided a sensitive measure of serious mental illness as diagnosed by clinical evaluation. This subset of six questions is used to compute the K6 score. As it turned out, the best subset of questions ask about a respondent's recent experience of a familiar set of negative emotions. The same negative

emotions that are a common part of everyday experience are diagnostic of mental illness when experienced excessively.

The six questions asked in constructing the K6 Index are

In the past thirty days, how often did you feel...

1. so sad nothing could cheer you up?
2. nervous?
3. restless or fidgety?
4. hopeless?
5. that everything was an effort?
6. worthless?

There are five possible responses: 1) all of the time, 2) most of the time, 3) some of the time, 4) a little of the time, or 5) none of the time. Points are assigned to each response, ranging from four points for “all of the time” to zero points for “none of the time.” The points from all six questions are summed to give the respondent’s K6 score. Having a high K6 score indicates that a person likely suffers from a serious mental illness.

I use the K6 primarily as a measure of experienced well-being. Ideally, a measure of experienced well-being would include emotions of both a negative and positive valence. The K6 included only negative emotions. Indeed, recent neuroscientific evidence suggests that negative and positive emotions actually belong to separate neural systems and can be experienced simultaneously (Norris et al., 2010). While this is a defect of the K6 for measuring experienced well-being, there is evidence that the correlations of negative and positive emotions with income are similar (Kahneman and Deaton, 2010).

The K6 has been administered in six rounds of the PSID. It was part of each biennial round between 2001 and 2013 with the exception of 2005. A total of 12,521 PSID respondents have completed it at least once. This yields 48,972 K6 observations, or an average of 3.9 observations per respondent.⁹

⁹The average is less than six for several reasons. First, data is collected about both sample members

Figure 1 shows a histogram of K6 scores in the PSID data. The responses are heavily skewed toward the right. The median score is 2 and the mean is 3.2. The median person thus either experiences one of the negative emotions some of the time or two of them a little of the time. About 30% of respondents reported “none of the time” for all six emotions. This shouldn’t be taken to mean that these individuals never experienced negative emotion at all, but rather that such experience was infrequent enough for them to endorse “a little of the time” when the time frame was the past 30 days.

The K6 index has been part of the National Health Interview Survey (NHIS) since 1997. The K6 score distribution is similar in the PSID and NHIS. Weighting 2013 data to be nationally representative, we find that 58.6% have a K6 of 2 or less in the PSID compared to 66.7% in the NHIS. (The difference is primarily due to there being more zeros in the NHIS.) Similarly, 7.6% have a K6 of 10 or greater in the PSID compared to 7.9% in the NHIS.

Kessler and colleagues found that threshold score of 13 balanced false positives and false negatives when using the K6 to predict whether a representative sample of US individuals has a clinically diagnosed serious mental illness (SMI) (Kessler et al., 2003). About 3.2% of the K6 scores in the PSID are greater than 13. However, as we will see, very few individuals with higher family incomes have K6 scores that are greater than 13. The linear probability model I use to compute nonparametric IV estimates does not behave well when dependent variables have conditional means close to zero. I have therefore elected to use a threshold of 10 for the serious mental illness (SMI) indicator for this paper. Kessler et. al.’s (2003) analysis suggests that this will increase the overall rate of false positives from 36% to 48%. In other words, about $\frac{2}{3}$ of those classified as having an SMI using the threshold of 13 actually have one, whereas about $\frac{1}{2}$ do using the threshold of 10.

The PSID began to collect a measure of life satisfaction in the 2009 wave. At the

and their spouses, while the K6 questions are only asked of the actual survey respondent. Which spouse responds varies from wave to wave. Second, children enter the sample as potential respondents when they set up their own households.

beginning of the interview, respondents are now asked “Please think about your life as a whole. How satisfied are you with it? Are you completely satisfied, very satisfied, somewhat satisfied, not very satisfied, or not at all satisfied?” Questions of this type are frequently used in the literature on SWB, and their inclusion in the PSID enable us to better understand how the measure of negative emotion related to this literature.

Mean responses to the life satisfaction question are shown in Figure 2, panel A. The modal respondent is very satisfied with their life as a whole, and less than 10% report being not very or not at all satisfied. Life satisfaction is weakly correlated with the K6 index. The Spearman correlation coefficient between the two variables is 0.32, which is similar to what Kahneman and Deaton (2010) found for similar measures in Gallup data. Panel B shows the relationship between life satisfaction and the K6 index. Those who are more satisfied tend to have lower K6 scores, as we would expect. Those who are not very or not at all satisfied show more variance in their K6 scores. Only some of those who are not very or not at all satisfied with their lives experience high levels of negative emotion. However, among those who are very or completely satisfied with life, almost none experience high levels of negative emotion. Noting that the life satisfaction measure is of a broader nature than the K6, one interpretation of this data is that a low level of negative emotion is necessary but not sufficient to produce a high level of life satisfaction.

The PSID collects extensive data on income. I use total family income in the *prior year* as my main income measure. This measure includes income from wages, assets, net business profits, transfers (such as welfare benefits and unemployment compensation), and social security for all family members. It excludes non-cash benefits such as food stamps and is pretax. Total family income better captures the total resources available to an individual than labor market income because income tends to be shared, at least to a degree, among family members. Using total family income for the *prior year* has econometric advantages that I discuss in section 3. I also collect full set of demographic data on the respondents that are suitable for use as exogenous controls. I adjust total

family income to 2013 dollars using the chained CPI.

The other important source of data for the analysis besides the PSID is the Current Population Survey (CPS). I use the CPS Merged Outgoing Rotation Group (MORG) files to construct estimates of mean weekly earnings by industrial sector, geographic division, and year.¹⁰ I use cross-walks to recode industries into the two-digit sector definitions from the 1990 Census to produce a consistent industry coding across the 2000-2012 CPS and 2001-2013 PSID. These mean weekly earnings estimates are matched to individuals in the PSID based on their industry and geographic region in a base year and are then used as an instrument for income.

2.2. Summary Statistics

I work with a subsample of the PSID that contains all adults aged 23 and older who responded to the K6 question at least once. I further limit the sample to individuals who were employed in at least one year in which they answered the K6 questions to facilitate the construction of an instrumental variable. These restrictions reduce the available observations to 38,675. I further limit the subsample to those observations in which the complete set of the demographic information I use to construct control variables was recorded, yielding a subsample of 38,251 total observations on 10,156 individuals.

I present some summary statistics for the PSID subsample in the paper in Table 1. The table shows observation-level means and standard deviations. While the respondents are broadly similar to the U.S. non-institutional adult population, there are some important differences arising both from the inclusion of a low-income oversample in the construction of the PSID and from some respondents entering the sample by cohabiting with an original sample member of descendent. Compared to the 2013 American Community Survey (ACS), the PSID respondents are slightly older (median age 40 vs 38 years), more likely to be female (60% vs 51%), much more likely to be African-American (34% vs 13%), somewhat less likely to be married or cohabiting (52% vs 58%), more likely

¹⁰The merged files are created and maintained by the National Bureau of Economic Research. They may be obtained from <http://www.nber.org/data/morg.html>.

to be employed (78% vs 58%), similarly educated (57% vs 59% have more than high school), and slightly better off (median household income of \$53,412 vs \$52,250) (U.S. Census Bureau, 2013). When we use the PSID sample weights to compare the PSID to the ACS, the mean characteristics are more similar. However, the weights exclude many respondents, and since my instrumental variables approach requires statistical power, I primarily use unweighted regressions.

3. Empirical Methods

The primary tool I use in the analysis is a panel regression. The general specification is

$$R_{it} = \alpha_t + \beta Y_{i,t-1} + X'_{it}\psi + W'_{it}\theta + \varepsilon_{it}. \quad (3.1)$$

The dependent variable R_{it} is the K6 score, the SMI indicator, or the life satisfaction score. The independent variable of interest is the base 2 log of total family income in the prior year $Y_{i,t-1}$. By using the base 2 log, the coefficient β measures the effect of a doubling of prior-year family income. The vector X_{it} contains control variables. Each year gets its own intercept α_t . The vector W_{it} contains variables that depend on the panel estimation technique used, such as fixed effects or lags of the dependent variable.

Prior-year family income offers some advantages over current-year income in this application. Since the prior year has ended, measured income has less measurement error than what would be inferred from current earnings. There is also some evidence that the experience of well-being may adapt after a period of time to an income change (e.g. Frederick and Loewenstein, 1999). If the frequency of negative emotional experience responds to income changes mainly in the short run, it is possible that the short-run effects could be large while the medium run effects are zero. Using prior-year income as our income measure means that β measures a medium-run effect.

3.1. Identification

A number of obstacles stand in the way of interpreting an estimate of β as the causal effect of prior-year income on our measures of negative emotion. Both simultaneity and omitted variable bias are a concern in this setting. We know that mood disorders, which feature a high frequency of negative emotions, are associated with lower productivity when comparing workers with similar jobs (Wang et al., 2004). This probably also leads to lower wages. Individuals with higher initial levels of happiness have been shown to earn more later on (Diener et al., 2002; Graham et al., 2004). However, a large literature beginning with Alloy and Abramson (1979) shows that individuals with higher levels of negative emotion are less prone to a variety of cognitive biases (Keltner and Lerner, 2011). These mechanisms suggest that $Y_{i,t-1}$ and ε_{it} will be correlated, though it is unclear whether on balance the correlation will be positive or negative.

Some of the identification issues described above can be addressed using panel data methods, though only partially. Fixed effect and lagged dependent variable (LDV) estimation can control for persistent or time-varying individual differences in the propensity to experience negative emotions. However, there are likely to be confounds that panel methods cannot address in the data we have. For example, a long-term depression beginning in $t - 1$ that is caused by the loss of a loved one may both depress family income at $t - 1$ and increase negative affect at t . This depression is not captured by either the $t - 2$ lag or the fixed effect. It would bias an estimate of β downwards.

The solution is to develop an instrumental variables approach to identification. The first-stage equation for an instrumental variable $Z_{i,t-1}$ in this application is

$$Y_{i,t-1} = \alpha_t + \gamma Z_{i,t-1} + X'_{it}\psi + W'_{it}\theta + \nu_{it}. \quad (3.2)$$

Two important requirements for instrumental variable $Z_{i,t-1}$ are that it be strongly correlated with endogenous explanatory variable $Y_{i,t-1}$ and that it be conditionally uncorrelated with the with the estimating equation error term ε_{it} . This latter condition,

the exclusion restriction, is usually represented as $E(Z_{i,t-1}\varepsilon_{it}) = 0$.

Finding an appropriate instrument for prior-year family income is challenging. As Pischke (2011) points out, many of the changes in state or federal policies that provide convincing exclusion restrictions in applied economics will not have a sufficiently strong first stage in a relatively small dataset such as the PSID. Further, such institutional variation in income tends to have a limited support across the income distribution. Instrumental variables estimates are local average treatment effects (LATE), meaning that they represent the responses of individuals whose variable of interest changes as a result of the instrument (Angrist and Pischke, 2008). A key question for this paper is whether the semi-elasticity of negative emotion to income is different for higher and lower incomes. An instrument that has a broad support is necessary to answer this question.

I take advantage of cross-industry variation in average wages as a source of identification. While some variation in wages across industries can be explained by employee characteristics, a large literature in labor economics argues that rents are just as, if not more, important (e.g. Dickens and Katz, 1987; Katz and Summers, 1989; Holzer et al., 1991; Hiscox, 2002). This suggests that the instrument will be strong even controlling for productivity-related characteristics, and that when so controlling the identifying variation will come from rents. I compute industry average wages by geographic region using the merged outgoing rotation group files of the Current Population Survey (CPS MORG).

Cross-industry variation has been used before to attempt identification of the effect of income in cross-sectional data (Luttmer, 2005; Pischke, 2011). In the panel setting I use, it is important to realize that individuals may pursue higher income by switching industries. Indeed, Pischke and Schwandt (2012) note that industry affiliation is correlated with fixed background variables such as height and mother's education, suggesting industry choice is not random. Further, it is reasonable to expect industry switching to be affected by emotional state. Individuals who experience more negative emotion might be less likely to take advantage of opportunities to switch. This means that average wages in an individual's current industry may be confounded by industry switching.

I take advantage of the panel structure of the dataset to address this limitation of cross-industry wage variation. My approach is related to Bartik’s (1991) shift-share instrument, which has seen a number of applications in recent years (Autor and Duggan, 2003; Luttmer, 2005; Card, 2009; Lewis, 2011; Clingingsmith, 2014). Let j index an individual’s current industry and g be the geographic region in which they reside. For each individual, I select the industry and region from the first year they appear in the subsample. I call these the base industry $\bar{j}$ and base region $\bar{r}$. I then assign an industry-region average wage to this individual in each PSID year using the base industry and region.

$$Z_{ijr,t-1} = \bar{w}_{\bar{j}\bar{r},t-1} \quad (3.3)$$

This variation still reflects an individual’s initial industry and region, which might be correlated with productivity-related characteristics. When I use the instrument, I therefore control for each base region by base industry cell.

One can nevertheless conceive of reasons why the exclusion restriction for this instrument may not hold exactly. A shock to base region-industry average wages at $t - 1$, over and above its effect on family income at $t - 1$, may also contain information about future prospects. If this information became known to an individual and has an emotional impact that endures to time t , the exclusion restriction would be violated.

There is some evidence that SWB is negatively affected by the level of income of a relevant comparison group (Veenhoven, 1991; Diener et al., 1993; Blanchflower and Oswald, 2004; Luttmer, 2005; Clark et al., 2008). If individuals perceive their base region-industry as a peer group against which to compare themselves and if the effect on levels of income also applies to changes in income, which is the variation we are using to identify the effect of income, the exclusion restriction could be violated.

To reduce concerns that region-industry average wages might have a direct effect, I use broad definitions of both region and industry. I use nine census divisions for geographic categories and nine industries for industrial categories. Individuals’ knowledge of trends in average wages most likely come from more local comparisons. Further, our controls

for initial industry-region effects mitigates direct effects from differences in one’s level of income to the region-industry mean. Further, since the comparison group usually used in the relative income literature is geographic, I estimate models with state-year effects to control for changes in comparison group income.

Finally, it is reasonable to suppose that any true direct effect of the instrument on negative emotion conditional on controls will be much smaller than the effect of family income. In Appendix A1 I conduct a sensitivity analysis that shows the IV estimates are robust to moderate violations of the exclusion restriction.

3.2. *Nonparametric Methods*

I employ nonparametric regression techniques at a number of points in the paper. I use both kernel and B-spline regressions. Kernel regressions are a familiar applied technique (see Cameron and Trivedi, 2005). B-spline regressions are more exotic: they use polynomials defined over a partition of the variable of interest to approximate the shape of its relationship with the outcome. Varying the degree of the polynomial and the fineness of the partition allow for any smooth function to be approximated. The partition points are referred to as knots. I provide a brief introduction to B-splines in section 3.2 of the online appendix.

Nonparametric IV has a complexity not present in the parametric case. An inverse problem must be solved and it may be “ill-posed,” which means that the resulting estimates may be highly sensitive to small changes in the data. This problem and its solution through a technique called *regularization* are discussed in Horowitz (2011). I employ the R package `crs` to estimate IV spline regressions (Racine and Nie, 2012). Estimation implements the Landweber-Fridman regularization method.

4. Results

We begin the empirical analysis with an examination of how the K6 score distribution varies with prior-year family income. Table 2 divides respondents into thirteen family

income bins and shows the percentage of respondents in each income bin whose K6 scores fall into each of nine score bins. For example, the top left cell indicates that 25% of respondents with prior-year family income between \$5,000 and \$9,000 had a K6 score of zero or 1. The rightmost columns shows the K6 mean for each income bin and the fraction who have a positive SMI indicator.

As the income level increases, we see that the K6 score distribution shifts to the left and that the highest score bins in particular become more sparsely populated. The share of respondents with K6 scores of three or less grows from 42% to 75% over the range of income. More dramatically, the population share whose K6 score is 10 or greater, consistent with serious mental illness (SMI), falls from 20% at the lowest income level to only 2% at the highest.

Figure 3 presents local-linear kernel regressions that show how the K6 score, SMI indicator, and life satisfaction score vary with prior-year family income. Time effects have been removed from all variables. The horizontal axis ranges from approximately the 5th to 95th percentiles of family income. Panel A plots the mean K6 score, which falls from just over 5 for family incomes of \$10,000 to about 2.2 for family incomes of \$200,000. This decline in the mean corresponds to about 0.75 unconditional standard deviations. The slope of the graph is roughly linear for incomes below \$80,000 and then becomes more shallow. Panel B shows the conditional percentage of respondents with a K6 indicating SMI ($K6 \geq 10$). The relationship between SMI and prior-year family income is clearly concave, with income having the strongest correlation at low levels. The conditional probability of having an SMI falls from about 17.5% for family income of \$10,000 to 2.5% for family income of \$200,000. Panel C shows that the conditional mean of life satisfaction rises from about 3.5 to 4.1, or about 0.7 SD. In contrast with the K6 and SMI, the life satisfaction curve has a slight S-shape, with the steepest slope for incomes between \$40,000 to \$100,000.

4.1. Pooled Cross Section and Controls

I next explore how conditioning on a broad set of demographic variables affects the relationship between the measures of SWB and family income. I use pooled estimations of 3.1 that set $W_{it} = 0$.

The bivariate semi-elasticity of the K6 score to prior-year family income is -0.62, which means a doubling of income is associated with a 0.62 point reduction in the K6 (Table 3, column 1). Recall for comparison that the mean K6 is 3.3 and the SD is 3.8. A doubling of income leads to a 3.3 percentage point decline in the SMI indicator (column 2) and an increase in life satisfaction of 0.12.

I introduce controls for gender, race, and age in columns 4 through 6. It is reasonable to consider these variables exogenous because they are predetermined. There are only a slight changes in the coefficients of interest for all three outcomes. The regression reveals that women both experience more negative emotion than men but also have higher life satisfaction. Coefficients on the race variables are statistically insignificant for the K6 and life satisfaction, though African Americans have a substantially higher rate of SMI. The regressions also reveal that there is a strong tendency for negative emotion to decline with age beginning at around 39. Life satisfaction exhibits a U-shaped pattern in age, with those 48 to 56 years old experiencing the lowest life satisfaction and those over 65 the highest. The age and gender effects for life satisfaction are consistent with those found in other U.S. and international data (Oswald, 1997; Graham and Pettinato, 2002; Blanchflower and Oswald, 2004).

I next introduce measures of education level, unemployment, marital status, and the number of children into the regression (columns 7-9). These variables are more problematic as controls because they may themselves be a function of both an individual's state of well-being and their income. I therefore will not include them in later IV regressions to which I attach a causal interpretation because doing so could introduce bias.

The first thing to note is that introducing these variables reduces the size of the prior-year family income coefficient by about a third for the K6 and SMI and about half for

life satisfaction. The gender, race, and age effects are broadly similar to those in columns 4-6. However, because these controls are potentially affected by income we cannot say whether the change in the coefficient on income is due to an improvement in specification or to and increase in selection bias.

Compared to those with less than a high-school education, the more highly educated experience less negative emotion conditional on age, income, and race. The effects of education on life satisfaction are less clear cut. Unemployment and marriage also have very strong relationships with both negative emotion and life satisfaction. The unemployed experience more negative emotion and are less satisfied with life, while marriage has the opposite effects. Those who are divorced and not remarried have similar negative emotion and life satisfaction to the never-married. Having a single child has no significant effect on negative emotion or life satisfaction. Having two or more increases life satisfaction, while negative emotion is reduced by having exactly two but not more.

Note that even when the full set of covariates are included, the R^2 of the negative emotion and life satisfaction regressions are only 8% and that of the SMI regression in only 5%. This means that more than 90% of variation in SWB is unexplained.

From this analysis, we can conclude that income continues to have a robust correlation with negative emotion when we control for other factors thought to influence well-being. Gender, age, education, employment, marital status, and family size all have sizeable relationships with negative emotion conditional on income.

4.2. Panel Regressions

The PSID subsample contains an average of 3.9 observations per respondent. I take advantage of repeated measures to estimate fixed effect and lagged dependent variable regression models as discussed in section 3.1. To make comparison between the K6 and life satisfaction regressions easier, I convert both of these variables to z-scores by demeaning them and dividing by their standard deviations.

Table 4 presents pooled, fixed effect, lagged dependent variable estimates in separate panels. Columns show estimates for both the K6 z-score, SMI indicator, and life

satisfaction z-scores without controls and then with the exogenous and full control sets.

The pooled OLS estimates show us that the effect size for negative emotion and life satisfaction are similar (Table 4, panel A). A doubling of real family income reduces negative emotion by about 0.16 SD and increases life satisfaction by about 0.14 SD.

Fixed effects estimates are substantially smaller than the pooled OLS estimate (Table 4, panel B). With or without the exogenous controls, a doubling of family income reduces the K6 by 0.03 SD and raises life satisfaction by 0.04 SD. These estimates are only a quarter as large as the pooled estimates. This suggests that a substantial part of the cross-section correlation between income and SWB may be due to persistent individual differences that are related to income. The overall pattern is similar for the SMI measure.

To the extent that the persistent individual differences eliminated by the fixed effects are confounders, fixed-effects estimates will be less biased than pooled estimates. An example of a confounding fixed factor is long-run physical health. Physical health is positively related to productivity and negatively related to negative emotion, and excluding it from the cross-sectional regression biases down the coefficients. We therefore expect the fixed effect estimates to be less negative. On the other hand, fixed effects estimates are more susceptible to attenuation bias due to measurement error than the pooled estimate. Family income tends to be fairly persistent and is difficult to compute accurately in a survey setting, which means individual-level changes may be noisy. Lastly, the fixed effects estimates only make use of within-respondent variation in family income, while most of the actual variation is across individuals. Some variation removed by fixed effects could be due to the persistent causal effect of income on negative emotion.

Including a lag of the dependent variable (LDV) rather than a fixed effect is more appropriate when the important omitted variables are time-varying. I show LDV estimates in Table 4, panel C. Lags are available only for the 2003 and 2009-2013 waves, so this specification is estimated on a smaller sample. The LDV semi-elasticities are about twice as large those estimated with fixed effects but half the size of pooled estimates. With or without exogenous controls, doubling family income reduces negative emotion

by 0.07 SD and increases life satisfaction by 0.08 SD. Again, the pattern is similar for the SMI measure. The lagged terms have coefficients of about 0.5 for both the K6 and life satisfaction, suggesting strong persistence in SWB. Note that adding the lag increases the R^2 of the regressions by about fivefold.

4.2.1. Panel Marginal Effects by Family Income

Average effects such as those reported in table 4 are of limited use in understanding the welfare effects of redistributive changes in family income. For that purpose, we need to understand how the marginal effects of income vary in the level of income. To conduct this analysis, I use kernel-weighted regression to calculate the marginal effect β in equation 3.1 at different levels of family income (see Cameron and Trivedi, 2005, section 9.5).

Figure 4 shows the marginal effects for the K6 z-score in panel A and the life satisfaction z-score in Panel B. Each panel shows effects for a 100% increase in family income in the right column and for a \$1,000 increase in income in the left column. I add the effect of dollar changes in income because these are most relevant for thinking about redistribution.

The effects of family income on negative emotion are fairly constant in percentage terms up to family income of about \$80,000, when they begin to fall to zero (pooled and LDV specifications). Since a larger dollar amount of income makes a percentage change at higher levels of income, we see that the effects of dollar changes in income fall much more steeply, particularly between family income of \$10,000 and \$20,000.

In contrast, the effects of percentage changes in family income on life satisfaction are strongest in the middle of the distribution, for incomes between \$40,000 and \$80,000. This means that the level effects fall much more slowly with income. For example, an additional \$1,000 has about double the effect on life satisfaction for someone at \$10,000 in income than someone at \$60,000. The difference in effect on negative emotion is closer to five times. One way to characterize this difference is that income has a strong effect on both experienced and evaluated well-being at low levels of income, but that the effect

on experienced well-being falls more dramatically as income rises.

4.3. IV Regressions

The panel estimates presented in section 4.2 are potentially confounded by unobservables that are neither fixed over time nor adequately controlled for by the lag of the K6 score. For this reason, I will now turn to estimates produced using the instrumental variables approach discussed in section 3.1. Recall that the instrument $Z_{i,t-1}$ is prior-year mean weekly earnings in the census division and industry in which each PSID respondent worked in their first year in the subsample, conditional on a dummy set for division-industry in that first year.

4.3.1. First Stage Estimates

Table 5 shows first-stage estimates for the pooled, fixed effect, and LDV specifications. The first stage is quite strong in the pooled specification. The elasticity of prior-year family income to the instrument is 0.61 without controls and around 0.45 with them (columns 1 and 2). The partial F-statistics are 101 and 63, respectively. Because these regressions include base year division-industry dummy variables, the identifying variation is coming from changes in division-industry averages over time. There is no first stage in the fixed effect specification (columns 3 and 4). The first stage in the LDV specification is quite weak, with F-statistics of 7 without and 3 with controls (columns 5 and 6). In using the instrument, I therefore employ the pooled specification.

Recall that the life satisfaction variable was only collected in the final three of the six PSID waves I am using. The first of these waves was 2009. I created a version of the instrument for the 2009-2013 data using 2009 as the initial year. There is no first stage for any of the three specifications in this case (see Table A2 in the online appendix). While there is a strong unconditional correlation between the instrument and prior-year family income, almost all of it is cross-sectional, and controlling for the initial industry-division cell absorbs that variation. The IV analysis thus proceeds using the K6 variable only.

The IV regression method produces an estimate of the local average treatment effect, or LATE (Angrist and Pischke, 2008). Loosely, the estimate $\hat{\beta}$ is a weighted average of individual effects β_i where the weights depend on the degree to which the instrument has an impact on prior-year family income for individual i . The estimate is thus “local” to those individuals whose incomes are shifted by the instrument. Suppose that the division-industry instrument $Z_{i,t-1}$ is correlated with prior-year family income in low-wage industries but not for high-wage industries. The IV estimates produced under this condition would be local to the subpopulation in low-wage industries.

We can gain some insight into the LATE in this application by analyzing the correlation of the instrument with prior-year family income across the instrument’s support. I apply the nonparametric spline method described in section 3.2 to the first-stage equation. The instrument enters the first stage as an unknown function $g(Z_{i,t-1}, C_{it})$. Figure 5 panel A plots predicted values from a spline estimate of expected prior-year family income against the instrument that uses five knots and is degree three. The correlation between prior-year family income and the instrument is positive over the entire support, but is clearly strongest for average earnings below about \$800 per week. The LATE estimated by IV is therefore more representative of individuals in the lower-wage division-industries.

The result that the instrument is strongest for lower-wage division industries might cause concern if those with high family incomes tend to be employed in higher-wage division-industries. Figure 5 panel A plots predicted values from a “reverse” first stage, in which the instrument is the dependent variable and family income the independent variable. This shows that there is little relationship between the conditional mean wage in a person’s division-industry and their family income. Individuals with high levels of family income do not seem to work in systematically higher wage division-industries.

4.3.2. IV Estimates

I compute IV estimates using two-stage least squares. Estimates of the effect of prior-year family income on the K6 score and SMI indicator are shown in Table 6.

The IV estimates show that a doubling of family income leads to a one point decline in the K6 score, both with and without controls (columns 1 to 3). These effects are large and statistically significant. The SD of the K6 score is 3.8, so a doubling of income reduces the K6 by 0.26 SD. While this change in family income is large in terms of the variations experienced by a typical individual, there is a ninefold difference in income between the 10th and 90th percentiles of the family income distribution. Figure 3 shows us that the mean K6 score falls by about three points as family income rises from \$10,000 to \$160,000. The IV estimates suggest that the change in income alone would produce a four point reduction in the K6 score, meaning that other factors correlated with income offset some of its effects.

The IV estimate is not affected by the introduction of exogenous controls in column 2. Estimates of gender, race, and age effects in OLS regressions that control for family income are all likely to be biased because these factors are correlated with the unobservable determinants of income. Column 2 therefore also identifies these effects conditional on family income. Women have a higher level of negative emotion than men by about 0.4 points. There are no significant effects of race. Compared to the omitted 23-30 year old age group, we see that there are no significant age effects for those between 31 and 57 years old, and then a substantial reduction in negative emotion for the 57-64 and 65+ age groups. Age effects are thus less pronounced in the IV than the OLS.

The estimates are unchanged when we introduce state-year effects in column 3. SWB has been shown to be affected by one's income relative to a comparison group, which is usually taken to be a locality such as a state or county (Blanchflower and Oswald, 2004; Luttmer, 2005). The smallest geographic unit available in my sample is the state. Controlling for state-year effects rules out the possibility that the identifying variation is confounded with time-varying state-level unobservables such as average income that may effect negative emotion as described in the literature. Note that the first-stage is not weakened by the addition of state-year effects.

The effect of a doubling of family income on the likelihood of SMI is 5.3 percentage

points without controls (column 4) and 6.5 percentage points with controls (column 5). The effects are statistically significant and nearly as large as the 8% mean likelihood of exceeding the SMI threshold. Figure 3 shows the percentage depressed falls by about 15 points as income doubles four times from \$10,000 to \$160,000. The IV estimates suggest that income alone would produce an even larger decrease.

Women have a higher incidence of SMI by about 2 percentage points, while African Americans have a lower rate by five percentage points. The negative effect for African Americans is particularly interesting because they have a higher unconditional rate of SMI than other groups. Compared to the 23-30 year old age group, SMI incidence increases over the life span, reaching a peak in the 48-56 year old age group, and then falling. This is a similar pattern to the OLS.

Overall, the IV estimates are about 1.6 times larger in magnitude than the OLS estimates. This means that either there is significant bias in the OLS estimates for which the IV corrects or IV bias is itself significant. A sensitivity analysis suggests that these estimates are not in fact very sensitive to moderate violations of the exclusion restriction (see section A1 of the online appendix). If the direct effect of the instrument were 20% as large as the effect of family income, for example, the implied true parameter estimate is -0.7 rather than the -1.0 shown in Table 6.

4.3.3. IV Marginal Effects by Family Income

I next investigate how the response of negative emotion to family income varies with its level, employing the nonparametric IV techniques discussed in section 3.2. We can think of the resulting estimates as LATEs at each point. This means that in interpreting these estimates, we must keep in mind that there may be shifts in the population affected by the instrument with family income. A causal interpretation of the estimate assumes that such changes are small. We saw earlier that the instrument and prior-year family income are strongly and consistently correlated over the ranges of the instrument. This should give us some confidence that the LATEs represent a broad section of the population at each level of income.

I estimate the IV B-spline regression with the 1% tails of the data trimmed. With the K6 score as the dependent variable and the degree and number of internal knots allowed to vary between 0 and 10, cross-validation selected a 3-degree spline with no internal knots. Since we are interested in the elasticities at different points in the distribution, I increased the number of internal knots to 3 for the final estimates. A 3-degree spline was selected by cross-validation from degrees 1 through 10 holding the number of internal knots to 3. For the SMI indicator, cross-validation selected a 5-degree spline with three internal knots.

Figure 6 shows predicted values and marginal effects of income for the K6 score in panel A. The predicted conditional K6 mean falls by about three points as income ranges from \$10,000 to \$200,000. In percentage terms, income has a large negative effect on negative emotion between family incomes of \$20,000 and \$80,000, where it is around -0.8. It then begins to fall quickly, reaching -0.18 at \$160,000 and ending near zero at \$200,000. In other words, the IV spline estimates reveal a satiation point of around \$200,000 above which additional family income has no effect on the mean experience of negative emotion. The marginal effect of \$1,000 in income falls more linearly from -0.06 at family income of \$10,000 to indistinguishable from zero at \$160,000. This clearly shows that in terms of negative emotion a marginal dollar has its greatest effect for those with the lowest incomes.

Predicted values and marginal effects for the SMI indicator are shown in Figure 6, panel B. The predicted incidence of SMI falls from about 0.2 to 0.02 as total family income ranges from \$10,000 to \$200,000. The effect of a percentage change in income falls steadily from around -0.06 to -0.01. Even more so than for the K6 score, income has its largest effect at lower level of income. A \$1,000 change in income reduces the probability of SMI by -0.004 at family income of \$10,000, but this effect falls steeply, hitting just -0.0007 at \$40,000 before levelling off. Income is thus particularly protective against SMI for those at the bottom of the distribution.

4.4. Welfare Effects of Redistribution

The marginal effects of dollar changes in family income allow us to compute the welfare effects of transfers across the income distribution (for more on the assumptions required see section A2 of the online appendix).

Consider a transfer of \$1,000 from someone at the 80th percentile of the distribution of family income to someone at the 20th percentile. These correspond to incomes of \$113,000 and \$28,000, respectively. These income levels differ by a factor of four. This transfer raises the family income of the 10th percentile person by 3.6% while reducing the family income of the 80th percentile person by 0.9%.

The marginal impact of \$1,000 in income on the K6 score is -0.0310 at family income of \$28,000 and -0.0038 at family income of \$113,000. The ratio of these two effects is 8.2. Marginal changes in income are therefore approximately 8.2 times as beneficial in terms of reducing negative emotional experience for an individual at the 20th percentile of income as one at the 80th percentile. The ratio for the SMI indicator at the same percentiles is 10.

These simple calculations cannot take into account the effects of making such transfers on productivity, the possibility that the effect on other aspects of well-being, such as positive emotion, are different from those on negative income, or issues of loss aversion and the effects of earned versus unearned income on well-being. Our preferred social welfare function may also place greater weight on those with higher incomes, though the opposite is more usual. These limitations should be taken as caveats to interpreting these calculations as true welfare effects of a redistributive transfer. However, any countervailing effects would need to be very large to offset the eightfold difference in the marginal effect of a dollar between the 20th and 80th percentiles. Simply put, this exercise provides a steep hill to climb for the case against redistributive taxation.

5. Discussion

I summarize the main findings as follows. Family income reduces both the mean experience of negative emotions and the high frequency of negative emotions characteristic of mental illness. A doubling of family income reduces negative emotion as measured by the K6 score by 0.96 points, or 0.26 SD, and reduces the incidence of SMI by 6 percentage points. The marginal effect of income on negative emotion is much larger when income is low than when it is high. The marginal effect of a dollar of family income is about 8 times higher at the 20th percentile of family income than the 80th, suggesting that redistributive transfers would increase welfare even if the costs of doing so were substantial. Income has a roughly constant percentage effect on negative emotion for incomes between \$10,000 and \$80,000, beyond which it begins to decline. Satiation of negative emotion appears at family income of \$200,000.

Many studies have estimated the relationship between income and SWB.¹¹ This literature has explored both cross-country variation in average SWB and average income and, as in the present study, within-country variation in SWB by income level. Most of the within-country studies have relied on cross-sectional data and have looked at the correlation of SWB and income, often conditioning on covariates. While a few studies have used panel data (Graham et al., 2004; Layard et al., 2008; Di Tella et al., 2010) or instrumental variables (Pischke, 2011), mine is the first to combine these techniques to produce well-identified causal estimates of the effect of income on SWB. This is an important step as economic policy primarily affects well-being through the level and distribution of income, and the problems of reverse causality and omitted variables are widely acknowledged in the literature.

My estimates suggest the estimates of the effect of log income on well-being reported in the literature may be too small. For example, Stevenson and Wolfers (2008) used

¹¹See Easterlin (1974); Diener et al. (1993); Oswald (1997); Graham and Pettinato (2002); Easterlin (2003); Graham et al. (2004); Blanchflower and Oswald (2004); Deaton (2008); Layard et al. (2008); Stevenson and Wolfers (2008); Daly and Wilson (2009); Diener et al. (2010); Easterlin et al. (2010); Kahneman and Deaton (2010); Di Tella et al. (2010); Binder and Coad (2011); Pischke (2011); and Stevenson and Wolfers (2013).

pooled General Social Survey data to estimate that a doubling of income increases life satisfaction by 0.22 SD. While my pooled OLS estimates show a 0.14 SD increase in life satisfaction and a 0.16 SD decrease in negative emotion, the IV estimates show a 0.26 SD decrease in negative emotion. While the existing literature primarily uses evaluative measures of SWB such as life satisfaction, the average effects of income on life satisfaction and negative emotion are similar even if the pattern across levels of income is different. It is reasonable to suppose that the bias follows the same pattern for both variables.

A smaller literature has sought to use SWB measures to understand how the marginal effects of income on welfare vary with income. Layard et al. (2008) assumed a constant-elasticity utility function and, using European and British data, estimated the elasticity of life satisfaction to income of 1.26. An elasticity greater than one implies that the marginal effect of log income falls with income. Binder and Coad (2011) make a similar point using quantile regression on British data. I casually identify a similar pattern using a more flexible estimation method and highlight the implications of my estimates for redistributive taxation.

I find that on average the marginal benefit of income to people at the 20th percentile is eight times higher than those at the 80th percentile. The literature has tended to overlook the vast difference in the marginal benefit of dollar changes in income as it has focused on questions such as satiation and differences between within and across country estimates. These estimates suggest that redistributive transfers can have large net welfare benefits even if the costs of doing so in terms of efficiency are quite large. A further implication is that policy that focuses on the average growth may not be welfare maximizing. It is well known that economic growth in the United States since 1980 has primarily benefited the top of the income distribution. My estimates suggest that policies that reduce overall growth but promote a more equitable distribution can be welfare enhancing.

Most of the existing literature uses evaluative measures of SWB such as life satisfaction, which I analysed in section 4, and Cantril's "ladder of life," which ask respondents to rank their current situation between the worst and best possible lives they could lead

(Graham, 2011). A smaller literature has begun to supplement these evaluative measures with measures of emotional/affective experience (White and Dolan, 2009; Kahneman and Deaton, 2010). My causal estimate of the effect of income on negative emotion support those Kahneman and Deaton (2010) reported in their cross-sectional analysis of Gallup data. Kahneman and Deaton also examined measure of positive affect, and found effects mirroring those of negative affect, which suggests that the restriction of the K6 index to negative emotions may not be that limiting. They further found that marginal effects of percentage changes in income on emotion fall to zero at high levels of income. The advantages of the present study with respect to Kahneman and Deaton are causal identification and a more comprehensive set of questions about negative emotion.

Kahneman and Deaton also found that life satisfaction was most responsive to income in the middle of the distribution with no satiation. Stevenson and Wolfers (2008) similarly report no satiation in life satisfaction. My correlational results for life satisfaction support these findings in the pooled specification, but show satiation at high incomes in the fixed effect and LDV specifications. In contrast, satiation appears in all specifications with negative emotions as the dependent variable, including the causally identified IV estimation.

Overall this pattern of findings supports the contention that different measures of SWB have systematically different relations to income (Graham et al., 2010), which means that they cannot be seen as substitutes for one another in all applications. There are at least three potential reasons why we might see satiation appear in measures of experienced well-being (i.e. affect) rather than evaluative well-being. First, evaluative measures are inherently comparative, and those with high levels of well-being and income may believe that they could do even better with more income, even though experienced well-being is satiated. Under this view, the comparative nature of the measure appears as a kind of bias. Second, evaluative measures may capture aspects of well-being that are not related to emotional states in a straightforward way, such as fulfilment or flourishing, and that are helped by income. Third, the effect of income on life satisfaction may not

be identified, as the finding of satiation in LDV and fixed effects specifications suggests.

A side-benefit of identifying the effect of income on negative emotion is that effects of gender, race, and age conditional on income are also identified. The effects of these characteristics on well-being have also been of interest in the literature on SWB (Oswald, 1997; Frey and Stutzer, 2002; Graham and Pettinato, 2002; Blanchflower and Oswald, 2004; Graham, 2011). Women have a higher level of negative emotion than men by about a tenth of a SD, which is similar to effects reported in Graham and Pettinato (2002). There are no significant differences by race. Age effects are less pronounced than typically reported in the literature. Compared to 23-30 year olds, there is no upward trend in negative emotion with age until age 57. Then, negative emotion begins to fall. Those aged 65 and above have nearly a third of a SD lower negative emotion, which is comparable in size to a 125% increase in income.

Finally, this paper also provides much improved casual evidence that low income is a major cause of serious mental illness. The association between income and mental illness and the problem with reverse causality is well-known in the literature (Lorant et al., 2003; Johnson and Krueger, 2006; McMillan et al., 2010; National Center for Health Statistics, 2011; Sareen et al., 2011). This is the first study, to my knowledge, that addresses the problem of identification of the causal effect of income on the incidence of mood disorders in a large U.S. sample.¹² This study shows that K6 levels exceeding the threshold for serious mental illness is concentrated among those with low income and that the incidence of depression is very responsive to additional income, particularly for those with family incomes below \$40,000. While transfers of income therefore seem a means for preventing serious mental illness, further insights may come from studying the mechanisms through which income insulates people from such conditions. There may be unobserved factors that mediate the relationship between income and SMI at low levels of income, such as access to medical care, that could be targeted directly to reduce SMI.

¹²A randomized evaluation of family income supplements in India found that family income reduced the incidence of mental illness in children (Costello et al., 2010).

References

- Alloy, L. B., Abramson, L. Y., 1979. Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General* 108 (4), 441–485.
- Angrist, J. D., Pischke, J.-s., 2008. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Autor, D. H., Duggan, M., 2003. The rise in the disability rolls and the decline in unemployment. *Quarterly Journal Of Economics* 118 (1), 157–206.
- Bartik, T. J., 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.
- Binder, M., Coad, A., Aug. 2011. From Average Joe’s happiness to Miserable Jane and Cheerful John: using quantile regressions to analyze the full subjective well-being distribution. *Journal of Economic Behavior & Organization* 79 (3), 275–290.
- Blanchflower, D. G., Oswald, A. J., Jul. 2004. Well-being over time in Britain and the USA. *Journal of Public Economics* 88 (7-8), 1359–1386.
- Cameron, A. C., Trivedi, P., 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press.
- Card, D., May 2009. Immigration and Inequality. *American Economic Review* 99 (2), 1–21.
- Clark, A. E., Frijters, P., Shields, M. A., 2008. Relative Income , Happiness , Utility: An for the Explanation Easterlin and Other Puzzles. *Journal of Economic Literature* 46 (1), 95–144.
- Clingingsmith, D., 2014. Industrialization and Bilingualism in India. *Journal of Human Resources* 49 (1), 73–109.
- Conley, T. G., Hansen, C. B., Rossi, P. E., 2012. Plausibly Exogenous. *Review of Economics and Statistics* 94 (1), 260–272.
- Costello, E. J., Erkanli, A., Copeland, W., Angold, A., May 2010. Association of family income supplements in adolescence with development of psychiatric and substance use disorders in adulthood among an american indian population. *JAMA : The Journal of the American Medical Association* 303 (19), 1954–60.
- Daly, M. C., Wilson, D. J., Apr. 2009. Happiness, Unhappiness, and Suicide: An Empirical Assessment. *Journal of the European Economic Association* 7 (2-3), 539–549.
- Deaton, A., Mar. 2008. Income, Health, and Well-Being around the World: Evidence from the Gallup World Poll. *Journal of Economic Perspectives* 22 (2), 53–72.
- Di Tella, R., Haisken-De New, J., MacCulloch, R., Dec. 2010. Happiness Adaptation to Income and to Status in an Individual Panel. *Journal of Economic Behavior & Organization* 76 (3), 834–852.
- Dickens, W. T., Katz, L. F., 1987. Inter-Industry Wage Differences and Industry Characteristics. In: Lang, K., Leonard, J. (Eds.), *Unemployment and the Structure of Labor Markets*. Basil Blackwell,

- New York, pp. 48–89.
- Diener, E., Ng, W., Harter, J., Arora, R., 2010. Wealth and happiness across the world: Material prosperity predicts life evaluation, whereas psychosocial prosperity predicts positive feeling. *Journal of Personality and Social Psychology* 99 (1), 52–61.
- Diener, E., Nickerson, C., Lucas, R., Sandvik, E., 2002. Dispositional Affect and Job Outcomes. *Social Indicators Research* 40 (April), 229–259.
- Diener, E., Sandvik, E., Seidlitz, L., Diener, M., Mar. 1993. The relationship between income and subjective well-being: Relative or absolute? *Social Indicators Research* 28 (3), 195–223.
- Easterlin, R., 1974. Does Economic Growth Improve the Human Lot? In: David, P., Reder, M. (Eds.), *Nations and Households in Economic Growth*. Academic Press.
- Easterlin, R. A., Sep. 2003. Explaining happiness. *Proceedings of the National Academy of Sciences* 100 (19), 11176–11183.
- Easterlin, R. A., McVey, L. A., Switek, M., Sawangfa, O., Zweig, J. S., Dec. 2010. The happiness-income paradox revisited. *Proceedings of the National Academy of Sciences* 107 (52), 22463–22468.
- Frederick, S., Loewenstein, G., 1999. Hedonic Adaptation. In: Kahneman, D., Diener, E., Schwarz, N. (Eds.), *Well-being: The Foundations of Hedonic Psychology*. New York, pp. 302–329.
- Frey, B. S., Stutzer, A., 2002. *Happiness and economics: how the economy and institutions affect well-being*. Princeton University Press, Princeton, N.J.
- Graham, C., 2011. *The pursuit of happiness: an economy of well-being*. A Brookings focus book. Brookings Institution Press.
- Graham, C., Chattopadhyay, S., Picon, M., Feb. 2010. The Easterlin and Other Paradoxes: Why Both Sides of the Debate May Be Correct. In: Diener, E., Kahneman, D., Helliwell, J. (Eds.), *International Differences in Well-Being*. Oxford University Press, pp. 247–278.
- Graham, C., Eggers, A., Sukhtankar, S., Nov. 2004. Does happiness pay? *Journal of Economic Behavior & Organization* 55 (3), 319–342.
- Graham, C., Pettinato, S., 2002. *Happiness and Hardship: Opportunity and Insecurity in New Market Economies*. Brookings Institution Press.
- Hiscox, M. J., 2002. Factor Mobility and Interindustry Technological Change : Evidence on Wage and Profit Dispersion Across U . S . Industries , 1820-1990. *The Journal of Economic History* 62 (2), 383–416.
- Holzer, H., Katz, L., Krueger, A., 1991. Job Queues and Wages. *The Quarterly Journal of Economics* 106 (3), 739–768.
- Horowitz, J., 2011. Applied Nonparametric Instrumental Variables Estimation. *Econometrica* 79 (2), 347–394.
- Johnson, W., Krueger, R. F., Apr. 2006. How Money Buys Happiness: Genetic and Environmental

- Processes Linking Finances and Life Satisfaction. *Journal of Personality and Social Psychology* 90 (4), 680–91.
- Kahneman, D., Deaton, A., Sep. 2010. High Income Improves Evaluation of Life but not Emotional Well-Being. *Proceedings of the National Academy of Sciences of the United States of America* 107 (38), 16489–93.
- Kahneman, D., Wakker, P. P., Sarin, R., May 1997. Back to Bentham? Explorations of Experienced Utility. *The Quarterly Journal of Economics* 112 (2), 375–406.
- Katz, L., Summers, L., 1989. Industry rents: Evidence and implications. *Brookings Papers on Economic Activity* 1989 (1989), 209–290.
- Keltner, D., Lerner, J., 2011. Emotion. In: *The Handbook of Social Psychology*, 5th Edition. McGraw-Hill, New York, pp. 312–347.
- Kessler, R., Andrews, G., Colpe, L., Hiripi, E., Mroczek, D., Normand, S.-L., Walters, E., Zaslavsky, A., Aug. 2002. Short screening scales to monitor population prevalences and trends in non-specific psychological distress. *Psychological Medicine* 32 (6), 959–976.
- Kessler, R. C., Barker, P. R., Colpe, L. J., Epstein, J. F., Gfroerer, J. C., Hiripi, E., Howes, M. J., Normand, S.-L. T., Manderscheid, R. W., Walters, E. E., Zaslavsky, A. M., Feb. 2003. Screening for serious mental illness in the general population. *Archives of general psychiatry* 60 (2), 184–9.
- Layard, R., Mayraz, G., Nickell, S., 2008. The marginal utility of income. *Journal of Public Economics* 92 (8-9), 1846–1857.
- Lewis, E., May 2011. Immigration, Skill Mix, and Capital-Skill Complementarity. *Quarterly Journal of Economics* 126 (2).
- Lorant, V., Delige, D., Eaton, W., Robert, A., Philippot, P., Ansseau, M., Jan. 2003. Socioeconomic inequalities in depression: A meta-analysis. *American Journal of Epidemiology* 157 (2), 98–112.
- Luttmer, E. F. P., Aug. 2005. Neighbors as Negatives: Relative Earnings and Well-Being*. *Quarterly Journal of Economics* 120 (3), 963–1002.
- Ma, S., Racine, J. S., Yang, L., Aug. 2015. Spline Regression in the Presence of Categorical Predictors. *Journal of Applied Econometrics* 30 (5), 705–717.
- McMillan, K. a., Enns, M. W., Asmundson, G. J. G., Sareen, J., Sep. 2010. The association between income and distress, mental disorders, and suicidal ideation and attempts: findings from the Collaborative Psychiatric Epidemiology Surveys. *The Journal of Clinical Psychiatry* 71 (9), 1168–75.
- National Center for Health Statistics, 2011. Health, united states, 2011. Tech. rep.
- Norris, C. J., Gollan, J., Berntson, G. G., Cacioppo, J. T., Jul. 2010. The current status of research on the structure of evaluative space. *Biological psychology* 84 (3), 422–36.
- Oswald, A. J., Nov. 1997. Happiness and economic performance. *The Economic Journal* 107 (445), 1815–1831.

- Pischke, J.-s., 2011. Money and Happiness: Evidence from the Industry Wage Structure.
- Pischke, J.-S., Schwandt, H., 2012. A cautionary note on using industry affiliation to predict income. IZA Discussion Paper Series (6840).
- PSID, 2013. Panel study of income dynamics public use dataset.
- Racine, J. S., 2012. A primer on regression splines. Tech. rep.
- Racine, J. S., Nie, Z., 2012. crs: Categorical Regression Spline.
URL <https://github.com/JeffreyRacine/R-Package-crs/>
- Sareen, J., Afifi, T. O., McMillan, K. A., Asmundson, G. J. G., Apr. 2011. Relationship between household income and mental disorders: findings from a population-based longitudinal study. Archives of general psychiatry 68 (4), 419–27.
- Sen, A., 1993. Capability and Well-being. In: Sen, A., Nussbaum, M. (Eds.), The Quality of Life. Clarendon Press, Oxford.
- Stevenson, B., Wolfers, J., 2008. Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox. Brookings Papers on Economic Activity (1), 1–87.
- Stevenson, B., Wolfers, J., 2013. Subjective Well-Being and Income: Is There Any Evidence of Satiation? NBER Working Paper Series (18992).
- U.S. Census Bureau, 2013. American Community Survey 2013, 1-year Estimates, Table S0201.
URL http://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=ACS_13_1YR_S0501&prodType=table
- Veenhoven, R., 1991. Is happiness relative? Social Indicators Research 24 (1), 1–34.
- Wang, P. S., Beck, A. L., Berglund, P., McKenas, D. K., Pronk, N. P., Simon, G. E., Kessler, R. C., Oct. 2004. Effects of major depression on moment-in-time work performance. The American journal of psychiatry 161 (10), 1885–91.
- White, M. P., Dolan, P., Aug. 2009. Accounting for the Richness of Daily Activities. Psychological Science 20 (8), 1000–1008.

Table 1: Summary Statistics for PSID Dataset

	Mean	Median	Standard Deviation		
			Overall	Between	Within
Family Income (2013 Dollars)	\$80,656	\$60,260	\$109,682	\$85,057	\$60,448
K6 Score	3.25	2	3.82	3.28	2.31
Life Satisfaction*	3.79	4	0.86	0.73	0.48
Years of Education	13.53	13	2.38	2.31	0.59
Age	42.62	40	13.12	12.83	3.47
Number of Children	0.89	0	1.18	1.08	0.56
K6 $\geq$ 10	0.08				
Employed	0.79				
Female	0.60				
Married	0.52				
Divorced	0.17				
Black	0.34				
White	0.61				

Notes: There are a total of 38,251 observations on 10,189 individuals in the dataset. The dataset covers the years 2001, 2003, 2007, 2009, 2011, and 2013.

* Life satisfaction was collected in 2009, 2011, and 2013 waves only. There are 21,752 observations on 8,699 individuals for this variable.

Table 2: Percentage of Respondents in K6 Score Categories by Prior-Year Family Income

Prior-Year Family Income	K6 Score Category									K6 Mean	K6 $\geq$ 10 (SMI)
	0-1	2-3	4-5	6-7	8-9	10-12	13-15	16-18	19-24		
5K-9K	25	17	17	11	9	10	4	3	3	5.6	0.20
10K-14K	31	15	15	14	8	9	4	3	2	5.0	0.15
15K-20K	32	18	16	12	8	8	3	2	1	4.5	0.14
21K-28K	33	19	17	12	8	6	3	1	1	4.2	0.11
29K-39K	37	20	17	11	6	5	2	1	1	3.7	0.09
40K-51K	38	21	16	10	6	5	2	1	1	3.5	0.09
52K-64K	43	21	16	9	5	4	1	1	0	3.1	0.07
65K-79K	43	23	16	8	4	3	1	0	0	2.8	0.05
80K-100K	45	24	15	7	4	2	1	0	0	2.7	0.05
101K-126K	47	25	15	8	3	2	1	0	0	2.4	0.03
127K-160K	47	27	13	7	3	2	1	0	0	2.3	0.03
161K-200K	50	26	14	6	2	1	1	0	0	2.2	0.02
201K-360K	50	25	14	7	2	1	0	0	0	2.2	0.02

Notes: Table reports the percentage of respondents in each K6 score category by range of prior year family income. Category percentages sum to 100 by row absent rounding errors. Time effects were removed from both variables. N=38,251.

Table 3: Correlates of Subjective Well-Being

	K6			SML			Life Sat.			K6			SML			Life Sat.		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)									
Log2 Real Family Income $t-1$	-0.621*** (0.023)	-0.033*** (0.002)	0.120*** (0.005)	-0.605*** (0.025)	-0.035*** (0.002)	0.118*** (0.006)	-0.412*** (0.028)	-0.024*** (0.002)	0.055*** (0.007)									
Female				0.472*** (0.061)	0.025*** (0.004)	0.059*** (0.015)	0.519*** (0.061)	0.027*** (0.004)	0.049*** (0.015)									
Black				-0.103 (0.138)	-0.032*** (0.009)	-0.056 (0.039)	-0.149 (0.140)	-0.030*** (0.010)	0.013 (0.038)									
White				0.145 (0.130)	-0.017* (0.009)	0.008 (0.038)	0.246* (0.131)	-0.008 (0.009)	0.014 (0.037)									
Age 31-38				-0.070 (0.068)	0.017*** (0.005)	-0.037* (0.020)	-0.018 (0.071)	0.017*** (0.005)	-0.090*** (0.020)									
Age 39-47				-0.279*** (0.078)	0.021*** (0.005)	-0.039* (0.022)	-0.254*** (0.080)	0.019*** (0.005)	-0.091*** (0.022)									
Age 48-56				-0.428*** (0.080)	0.022*** (0.005)	-0.066*** (0.023)	-0.433*** (0.087)	0.020*** (0.006)	-0.107*** (0.024)									
Age 57-64				-0.751*** (0.096)	0.010 (0.006)	-0.020 (0.025)	-0.725*** (0.105)	0.010 (0.007)	-0.069** (0.027)									
Age 65+				-1.394*** (0.116)	-0.021*** (0.007)	0.134*** (0.031)	-1.362*** (0.126)	-0.020** (0.008)	0.086** (0.034)									
High School Graduate																		
Some College																		
College Graduate																		
Some Graduate School																		
Unemployed																		
Married																		
Divorced																		
One Child																		
Two Children																		
Three or More Children																		
Adjusted R^2	0.052	0.030	0.044	0.064	0.035	0.049	0.077	0.045	0.084									
Observations	38,251	38,251	21,752	38,251	38,251	21,752	38,251	38,251	21,752									

Notes: Family income adjusted to 2013 dollars using chain-weighted CPI. All regressions include year dummies. Standard errors allow for arbitrary correlation of unobservables at the individual level.

Table 4: Subjective Well-Being and Prior-Year Family Income: Panel Regressions

Panel A: Pooled									
K6 Z-Score			SMI (K6 $\geq$ 10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.162*** (0.006)	-0.158*** (0.007)	-0.108*** (0.002)	-0.033*** (0.002)	-0.035*** (0.002)	-0.024*** (0.002)	0.140*** (0.006)	0.137*** (0.007)	0.067*** (0.008)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.052	0.064	0.077	0.030	0.035	0.045	0.044	0.049	0.084
Observations	38,251	38,251	38,251	38,251	38,251	38,251	21,752	21,752	21,752

Panel B: Fixed Effects									
K6 Z-Score			SMI (K6 $\geq$ 10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.034*** (0.006)	-0.034*** (0.007)	-0.026*** (0.007)	-0.007*** (0.002)	-0.006*** (0.002)	-0.006*** (0.002)	0.042*** (0.010)	0.041*** (0.010)	0.029*** (0.010)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.004	0.004	0.008	0.001	0.001	0.002	0.009	0.009	0.016
Observations	38,251	38,251	38,251	38,251	38,251	38,251	21,752	21,752	21,752

Panel C: Lagged Dependent Variable									
K6 Z-Score			SMI (K6 $\geq$ 10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.071*** (0.004)	-0.074*** (0.005)	-0.053*** (0.006)	-0.025*** (0.002)	-0.027*** (0.002)	-0.019*** (0.002)	0.082*** (0.006)	0.083*** (0.006)	0.050*** (0.007)
Lagged DV	0.530*** (0.010)	0.524*** (0.010)	0.519*** (0.010)	0.451*** (0.020)	0.446*** (0.020)	0.439*** (0.020)	0.459*** (0.009)	0.456*** (0.009)	0.436*** (0.009)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.320	0.323	0.327	0.124	0.126	0.132	0.255	0.256	0.268
Observations	23,187	23,187	23,187	23,187	23,187	23,187	12,910	12,910	12,910

Notes: Family income adjusted to 2013 dollars using chain-weighted CPI. Exogenous controls are gender, two race dummies, and five age dummies. Additional controls are four education level dummies, three family size dummies, unemployed, married, and divorced. All columns include year dummies. Fixed effects estimates are computed using the within estimator. Standard errors corrected for clustering at the individual level.

Table 5: First Stage Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
Log ₂ Division-Industry	0.616***	0.449***	0.007	-0.008	0.537***	0.333*
Mean Weekly Income $t-1$	(0.061)	(0.057)	(0.050)	(0.048)	(0.207)	(0.195)
Specification	Pooled	Pooled	FE	FE	LDV	LDV
Controls	None	Exog.	None	Exog.	None	Exog.
First-Stage F-Stat	101.34	62.90	0.02	0.02	6.72	2.91
Observations	38,251	38,251	38,251	38,251	23,260	23,260

Notes: All specifications include year dummies. Pooled and LDV regressions include cell dummies for base division and industry. Fixed effects estimates are computed using the within estimator. Standard errors corrected for clustering at the individual level.

Table 6: Negative Emotion and Family Income: IV Estimates

	K6 Score			SMI (K6 $\geq$ 10)		
	(1)	(2)	(3)	(4)	(5)	(6)
Log ₂ Real Family Income _{<i>t</i>-1}	-1.020*** (0.293)	-0.955** (0.389)	-1.048** (0.410)	-0.053*** (0.020)	-0.065** (0.027)	-0.069** (0.028)
Female		0.434*** (0.120)	0.402*** (0.121)		0.019** (0.008)	0.018** (0.008)
Black		-0.318 (0.262)	-0.264 (0.271)		-0.049*** (0.018)	-0.045** (0.019)
White		0.274 (0.183)	0.311 (0.195)		-0.006 (0.013)	-0.002 (0.014)
Age 31-38		0.110 (0.203)	0.140 (0.210)		0.031** (0.014)	0.033** (0.015)
Age 39-47		-0.046 (0.260)	0.013 (0.270)		0.040** (0.018)	0.042** (0.019)
Age 48-56		-0.166 (0.278)	-0.114 (0.290)		0.043** (0.019)	0.046** (0.020)
Age 57-64		-0.505** (0.240)	-0.481* (0.248)		0.029* (0.017)	0.029* (0.017)
Age 65+		-1.292*** (0.133)	-1.315*** (0.136)		-0.014 (0.009)	-0.015* (0.009)
Region-Industry Dummies	Y	Y	Y	Y	Y	Y
State-Year Dummies	N	N	Y	N	N	Y
First-Stage F-Stat	101.41	65.82	59.33	101.41	65.82	59.33
Observations	38,251	38,251	38,251	38,251	38,251	38,251

Notes: Family income adjusted to 2013 dollars using chain-weighted CPI. Regressions also include year dummies. The instrument for prior-year family income is the prior-year average wage in the industry and geographic division in which individual was employed in their base year. Regressions include cell dummies for base division and base industry of employment. Standard errors allow for arbitrary correlation of unobservables at the individual level.

Figure 1: Histogram of K6 Scores for PSID Subsample

Notes: This histogram shows K6 scores for the 10,156 individuals in the PSID subsample who have completed the instrument at least once. N=38,251.

Figure 2: Life Satisfaction and K6 for 2009-2013 PSID Subsample

Notes: This figure shows responses for the 8,699 individuals in the 2009-2013 PSID subsample who have responded to the life satisfaction and K6 questions. N=21,752.

Figure 3: Negative Affect and Family Income:
Bivariate Local-Linear Regression Plots

42

Notes: Year effects have been removed from variables. Graphs present local-linear regressions using the Epanechnikov kernel. Dashed lines show 95% confidence intervals computed using a non-parametric cluster bootstrap. For panels A and B, $N=38,251$. For panel C $N=21,752$.

Figure 4: Family Income Marginal Effects on Subjective Well-Being

Figure 5: First Stage: Spline Regression Plots

Notes: Graphs show results from non-parametric spline estimation of the first stage (panel a) and reverse first stage (panel b). The splines have five segments and is of degree 3. The specification includes dummies for base division by base industry cells. Control variables include gender, race, and age dummies. The grey region shows a 95% confidence intervals estimated by bootstrapping on cluster resamples at the individual level. N=38,251.

Figure 6: Negative Emotion and Family Income: IV Spline Regression Plots

Notes: Graphs show results from IV spline regressions. Regression for the K6 score were degree 3 with four internal knots. Regression for the SMI indicator was degree 5 with four internal knots. Control variables include race, gender, age, and division-industry dummies. The 1% and 99% tails of the distribution were trimmed for the estimation. $N=37,609$. Plots range over total family income from \$10,000 to \$200,000 (2013 Dollars). 95% confidence intervals are shown in grey.

Online Appendix

A1. Sensitivity to Violations of the Exclusion Restriction

A1.1. Methods

The effect on the estimate of β of a failure of the exclusion restriction can be seen by considering the following simplified IV setup that adds $Z_{i,t-1}$ to the estimating equation

$$R_{it} = \beta Y_{i,t-1} + \gamma Z_{i,t-1} + \kappa_{it} \quad (\text{A1})$$

$$Y_{i,t-1} = \Gamma Z_{i,t-1} + \iota_{it}. \quad (\text{A2})$$

The exclusion restriction is satisfied if and only if $\gamma = 0$, which we usually impose by omitting $Z_{i,t-1}$ from the estimating equation. If $\gamma \neq 0$, then 2SLS estimation of the simplified system yields an estimate of β that is biased as follows

$$\hat{\beta} \xrightarrow{p} \beta + \frac{\gamma}{\Gamma}. \quad (\text{A3})$$

When the exclusion restriction does not hold exactly, the biased introduced is inversely proportional to the correlation between the instrument and the endogenous variable Γ and directly proportional to the correlation between the instrument and the outcome γ (Conley et al., 2012). In an assessment of the quality of potential instruments, then, the strength of the first stage and the plausibility that the exclusion restriction holds can be viewed as substitutes. A small exclusion restriction violation with a weaker instrument will produce the same bias, *ceteris paribus*, as a larger exclusion restriction violation from a stronger instrument.

I will now lay out the methods I will use to analyze the sensitivity of 2SLS estimates to violations of the exclusion restriction. We will work with a specification that augments equation 3.1 with the instrument

$$R_{it} = \alpha + \tau_t + \beta Y_{i,t-1} + X'_{it}\psi + W'_{it}\theta + \gamma Z_{i,t-1} + \varepsilon_{it}. \quad (\text{A4})$$

The exclusion restriction holds when $\gamma = 0$. This formulation provides a way of exploring the sensitivity of 2SLS estimates of β . For example, we can see what happens to $\hat{\beta}$ when γ takes on any fixed value γ_0 by simply substituting γ_0 for γ , rearranging so the dependent variable is $R_{it} - \gamma_0 Z_{i,t-1}$, and estimating the resulting equation using 2SLS.

Conley et al. (2012) builds on this idea to produce methods for analyzing the sensitivity of 2SLS estimates based on prior beliefs about the distribution of γ . Consider the Gaussian prior $\gamma \sim N(\mu_\gamma, \Omega_\gamma)$. Let $\mathbf{X}$ be the $N \times k$ matrix of all RHS variables, θ be a $1 \times k$ vector of coefficients, and $\mathbf{Z}$ be the $N \times k$ matrix of instruments. Under the assumption that our uncertainty about γ is of the same order as sampling uncertainty, we can use the approximation

$$\hat{\theta}_{2SLS} \stackrel{\text{approx}}{\sim} N(\theta + \mu_\gamma \mathbf{A}, V_{2SLS} + \mathbf{A} \Omega_\gamma \mathbf{A}'), \text{ where} \quad (\text{A5})$$

$$\mathbf{A} = (\mathbf{X}' \mathbf{Z} (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{X})^{-1} (\mathbf{X}' \mathbf{Z}) \quad (\text{A6})$$

to construct point estimates and confidence intervals given a choice of μ_γ and Ω_γ .

In our application, we expect that $|\gamma|$ will be smaller than $|\beta|$ because real resources should be more important for emotional experience than information about average wages. We can capture this expectation about relative size in a relative prior of the form

$$\gamma \sim N(\zeta \beta, \delta^2 \beta^2). \quad (\text{A7})$$

In this prior, ζ and δ describe beliefs about γ as a fraction of the true value of β . The parameter ζ indicates our expectation about the mean size and sign of γ relative to the effect of interest β . When $\zeta > 0$, γ and β have the same sign. Our expectation about the standard deviation of γ relative to the size of β is given by the parameter δ .

We will consider two applications of the relative prior and the approximations given in equation A5. First, suppose that our belief about γ is mean zero but has some uncertainty. In terms of the prior, this means setting $\zeta = 0$ but δ to some positive number. We can use the approximation to compute adjusted confidence intervals for $\hat{\beta}$

for any value of δ . We will present a figure showing how adjusted confidence intervals vary with δ and show how much uncertainty the estimates can handle and still remain statistically different from zero.

Alternatively, suppose we have a specific belief about the relative size of γ and β . For example, we may think γ is 20% as large as β and the same sign. This would imply $\zeta = 0.2$. As a second application, we will consider a degenerate parameterization of the prior that takes $\delta = 0$. For any given value of ζ , we can then compute the implied true β as

$$\beta = \frac{\hat{\beta}_{2SLS}}{1 + \zeta \mathbf{A}_{\beta\beta}}, \quad (\text{A8})$$

where $A_{\beta\beta}$ is the diagonal element of $\mathbf{A}$ corresponding to our coefficient of interest. We will present a figure showing how the implied $\hat{\beta}$ and confidence intervals vary with ζ . We will use this approximation to present a figure showing how the implied β varies with ζ .

A1.2. Application

We will now analyze the sensitivity of the IV estimates to plausible violations of the exclusion restriction using the method described above. The analysis employs the augmented specification, equation A4, in which the instrument $Z_{i,t-1}$ enters with parameter γ .

The coefficient γ is the effect on negative emotion of the instrument, prior-year division-industry average wages, conditional on actual prior-year family income and controls for base division-industry cell and other exogenous demographic characteristics. Two reasons for γ to not be zero are if division-industry average wages conveys information about future income prospects or if people consider a division-industry to be a peer group and are concerned with relative income.

If division-industry income contains information about prospects, we might expect γ to have the same negative sign as prior-year family income. However, it is also possible that our conditioning on actual prior-year income could lead to a change to the opposite sign. Higher average wages in a person's industry might be a source of worry if it indicates

a worsened competitive position for one's firm. If division-industry income conditional on own income indicates one's relative position, we would also expect γ to have the opposite sign as β . Whatever the true sign of γ , it is reasonable to expect that the magnitude of its effect will be smaller than that of prior-year income itself, or $|\gamma| < |\beta|$. Family income measures real resources that a respondent has at his or her disposal, while division-industry average income is information about the economic environment that, while highly correlated with family income, is likely not even known to the respondent with any degree of accuracy.

I first present estimates of how mean-zero uncertainty about γ affects confidence intervals for the IV estimate of β . I use the prior that $\gamma \sim N(\zeta\beta, \delta^2\beta^2)$ with $\zeta = 0$. The parameter δ scales the standard deviation of the distribution of γ relative to the size of the coefficient estimate β . So when $\delta = 0.2$, the standard deviation of γ is taken to be 20% as big as β .

Figure A1, panel A plots 95% confidence intervals for $\hat{\beta}$ by values of δ ranging from 0 to 0.4. Because $\zeta = 0$, these confidence intervals are centered on the 2SLS estimates. The figures shows that as long as $\delta \lesssim 0.15$, the 95% confidence interval for $\hat{\beta}$ excludes zero. Panel B shows how quantiles of the distribution of γ vary with δ to aid understanding of what a particular value of δ implies. If we were to assume $\delta = 0.15$, that would mean the 75th percentile of the distribution of γ is 0.10 times the size of β , so that in half of the draws from this distribution γ is larger in magnitude than 0.10β . Further, a δ of 0.15 implies there is only a 5% chance that γ is larger in absolute value than 0.29β . We can therefore have a fairly substantial degree of uncertainty about γ relative to β before the confidence intervals become too wide to reject the null of no effect.

However, if one were of the opinion that there is more than a 5% chance that γ is larger than 0.29β , one would choose a larger value of δ than 0.15 for one's prior, in which case the 95% confidence interval for β would include zero.

We next examine what various fixed beliefs about the relative size of γ and β imply for the bias of our estimates. We take $\delta = 0$ and compute the implied parameter estimates

for various values of the relative bias $\zeta = \frac{\gamma}{\beta}$. Figure A2 plots the implied estimate $\hat{\beta}$, along with bootstrapped 95% confidence intervals, against fixed values of $\zeta_0 \in [-0.3, 0.3]$.

Recall that the IV estimate of β is -1.0 with controls. If $\zeta = 0.2$, then the implied true parameter estimate is instead -0.7. If $\zeta = -0.2$, then the implied true parameter is -1.8. So, as long as the magnitude of the direct effect of the instrument is less than 20% as large as the direct effect of family income, we can bound our estimate of β are between -1.8 and -0.7. As a way of gauging the magnitude of this range, it may be instructive to compare to the 95% confidence interval of the IV estimator, which is [-1.8, -0.2].

In sum, this analysis shows the IV estimates are quite robust to violations of the exclusion restriction that would be large in magnitude relative to what we would expect given the way the instrument is constructed.

A2. Measuring the Welfare Effect of Redistribution

We can apply a direct welfare measure such as the self-reported frequency of emotion to estimate the effect of a transfer between individuals with different levels of income. Here I present the assumptions must be made in order to use the empirical analysis presented in this paper for this purpose.

Let use consider an economy with N individuals. Each individual's true welfare as a function of their income is given by $\omega_i(y_i)$. Social welfare is simply the unweighted sum of individual welfare

$$W = \sum_{i=1}^N \omega_i(y_i). \quad (\text{A1})$$

We will consider the effect on social welfare of a transfer from a person with high income to a person with low income. To do this we will make use of the expected value of individual welfare conditional on income being y , which is

$$E(\omega_i(y_i)|y_i = y). \quad (\text{A2})$$

Our goal is to estimate the expected change in total welfare $E(\Delta W)$ from the transfer

of a small amount of income δ between two individuals with different levels of income $y_H > y_L$. We write this expected change as

$$E(\Delta W) = [E(\omega_i(y_i)|y_i = y_L + \delta) - E(\omega_i(y_i)|y_i = y_L)] \\ - [E(\omega_i(y_i)|y_i = y_H) - E(\omega_i(y_i - \delta)|y_i = y_H + \delta)]. \quad (\text{A3})$$

The top line shows the expected difference in welfare for the low-income individual from receiving the transfer while the bottom line shows the expected difference in welfare from giving the transfer. If $E(\Delta W) > 0$ the transfer improves expected total welfare.

Suppose that we have data on a self-reported measure of emotional experience over a period of time n_i and income y_i . We can econometrically estimate the average level of emotion conditional on a given level of income. We will call this estimate $\widehat{n(y)}$. I will now carefully outline the assumptions involved in modifying equation A3 so that we can use $\widehat{n(y)}$ to estimate welfare effects.

Keeping the framework general, we first rewrite our individual welfare function as

$$\omega_i(y_i) = f_i(n_i(y_i)). \quad (\text{A4})$$

This equation captures two kinds of idiosyncrasy in how our self-reported measure n_i is related to the true welfare response of individuals to income.

First, the individual-specific function $f_i(n_i)$ allows the effect of self-reported emotion n_i on welfare to be different across individuals. In other words, reporting frequent feelings of sadness might indicate more harm to the true welfare of some individuals than others.

Second, the individual-specific functions $n_i(y_i)$ capture the idea that there may be diversity in self-reported emotional responses to income. Two people who have the same income $\bar{y}$ may report different emotional responses to having their income raised to $\bar{y} + \delta$.

The first assumption we will make is that utility is comparable across individuals. Fully comparable utility means that everyone has the same mapping of self-reported

emotion into well-being:

$$f_i(n_i) = f(n_i). \quad (\text{A5})$$

For our application to small transfers, we can be a little more general and allow for an additive heterogeneity of response $\phi_i(n_i(y_i))$ as long as $E(\phi_i(n_i(y_i))|y) = 0$. In other words, there can be heterogeneity in how self-reported emotion maps into true welfare as long as it is mean zero conditional on income. Our expression for individual utility then becomes

$$\omega_i(y_i) = f(n_i(y_i)) + \phi_i(n_i(y_i)). \quad (\text{A6})$$

Note that there remain plausible forms of heterogeneity that are ruled out by the requirement that $E(\phi_i(n_i(y_i))|y) = 0$. For example, if income does indeed reduce the occurrence of negative emotion, it is possible that those individuals who are more sensitive to negative emotion will invest more in increasing their earnings. This type of heterogeneity is ruled out by this setup.

A challenge in using equation A6 is that we do not have any means of estimating the functional form of $f(\cdot)$. Does the effect of a marginal change in n_i vary with the level of the score, or is it linear? Answering this question involves knowing the “true” measure of welfare. While it is possible neurobiology may one day provide guidance, at the moment we simply do not know. We will therefore make the further simplifying assumption that $f(n_i)$ is a linear function

$$\omega_i(y_i) = n_i(y_i) + \phi_i(n_i(y_i)). \quad (\text{A7})$$

Recall that the data we will be using involves not the individual-specific mappings $n_i(y_i)$ but estimates of the conditional mean $n(y)$. The conditional mean may be written as $n(y) = E(n_i(y_i)|y_i = y)$. We can use this to write

$$\begin{aligned} n_i(y_i) &= E(n_i(y_i)|y_i = y) + [n_i(y_i) - E(n_i(y_i)|y_i = y)] \\ &= E(n_i(y_i)|y_i = y) + \zeta_i(y_i) \end{aligned} \quad (\text{A8})$$

The term $\zeta_i(y_i)$ is the idiosyncratic deviation of each individual's emotional response to income from the conditional mean. Note that by definition $E(\zeta_i(y_i)|y_i = y) = 0$. Substituting equation A8 into equation A7, taking the expectation with respect to income, and simplifying gives us

$$E(\omega_i(y_i)|y_i = y) = n(y). \quad (\text{A9})$$

This shows that under the assumptions we have made, our data on the conditional mean of negative emotion provides a measure of expected welfare conditional on income. We can combine equation A9 and equation A3 to yield

$$\Delta W = [n(y_L + \delta) - n(y_L)] - [n(y_H) - n(y_H - \delta)] \quad (\text{A10})$$

Using equation A10, we can compute the welfare effects of a small redistributive transfer directly from our estimated conditional mean function $\widehat{n(y)}$.

A3. B-Spline Regression

This section briefly explains the B-spline nonparametric regression techniques I use in estimating marginal effects of negative emotion for different levels of income.

Consider the following model of the relationship between an outcome R_{it} , a continuous explanatory variable of interest Y_{it} , and a vector of indicator control variables C_{it} :

$$R_{it} = g(Y_{it}, C_{it}) + \varepsilon_{it}. \quad (\text{A1})$$

In this equation, Y_{it} and C_{it} enter nonparametrically through the unknown function $g(\cdot, \cdot)$. There are a variety of ways to estimate the shape of $g(\cdot, \cdot)$, including local-linear regression, spline regression, and series regression. I use spline regression because it is computationally more tractable in this setting.

A spline $S_{d,n}(y)$ is a numeric function of a continuous variable y composed of polynomial segments of degree d and a number of knots n which occupy distinct locations

$t_i \leq t_{i+1}$ within the domain of the function. The polynomial coefficients may change between knots. The spline is useful in nonparametric estimation because it can approximate functions of an arbitrarily complex shape.

A B-spline (or basis spline) is a spline that has minimal support given d and n . Any spline of degree d with knots n can be expressed as a linear combination of B-splines

$$S_{d,n}(y) = \sum_i \alpha_i B_{i,d}(y), \quad (\text{A2})$$

where i indexes the basis functions, of which there are $d + n + 1$. We can fit a B-spline to data by first computing the $B_{i,d}(y)$ and then using OLS to estimate the $\hat{\alpha}_i$.

The B-spline basis functions are defined recursively (see e.g. Racine, 2012). To define a B-spline, we begin by defining an augmented knot set. To the set of internal knots $(t_1, \dots, t_n)$, append upper and lower boundary knots t_0 and t_{n+1} . This gives the knot set $(t_0, \dots, t_{n+1})$. Next, append a further $d - 1$ copies of the boundary knots $t_{-d} = \dots = t_0$ and $t_{n+1} = \dots = t_{n+d+1}$ to form the augmented knot set $(t_{-d}, \dots, t_0, \dots, t_{n+1}, \dots, t_{n+d+1})$. We index the augmented knot set with $i = 0 \dots n + 2d + 1$.

For each augmented knot i and spline degree d , the basis functions $B_{i,d}(y)$ are defined by

$$B_{i,0}(y) = \begin{cases} 1 & t_i \leq x < t_{i+1} \\ 0 & \text{otherwise} \end{cases}$$

$$B_{i,j+1}(y) = \alpha_{i,j+1}(y) B_{i,j}(y) + [1 - \alpha_{i+1,j+1}(y)] B_{i+1,j}(y), \quad (\text{A3})$$

where

$$\alpha_{i,j}(y) = \begin{cases} \frac{x - t_i}{t_{i+j} - t_i} & \text{if } t_{i+j} \neq t_i \\ 0 & \text{otherwise} \end{cases}$$

Consider the simple example of a quadratic spline ($d = 2$) with no internal knots

($n = 0$) and the domain $[-1, 1]$. The augmented knot set in this case is $(-1, -1, 1, 1)$. The three basis functions are $B_{0,2}(y) = (1 - x^2)$, $B_{1,2}(y) = 2(1 - x)x$, and $B_{2,2} = x^2$.

Our function $g(Y_{it}, C_{it})$ contains both a continuous variable Y_{it} and a vector of categorical variables C_{it} . There are methods for combining spline estimation of Y_{it} with kernel estimation of the categorical variables in a fully flexible way (Ma et al., 2015). However, given the large number of indicator variables in this application, the fully flexible approach proves too computationally intensive. To make the estimation tractable, I allow the indicator variables to enter additively

$$R_{it} = \sum_i \beta_i B_{i,d}(Y_{it}) + C'_{it} \gamma + \varepsilon_{it}. \quad (\text{A4})$$

This equation can be estimated by OLS for any choice of spline degree d and set of knots $(t_1, \dots, t_n)$. The degree and knots may simply be specified to fulfill certain criteria or may be chosen by a model selection procedure such as cross-validation. Cross-validation estimate the model for each unique combination of degrees and knots within specified domains and take that combination that provides the best fit. I use a mix of these procedures.

Table A1: Subjective Well-Being and Prior-Year Family Income: Weighted Panel Regressions

Panel A: Pooled									
K6 Z-Score			SMI (K6≥10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.164*** (0.009)	-0.162*** (0.009)	-0.119*** (0.010)	-0.033*** (0.002)	-0.035*** (0.002)	-0.025*** (0.003)	0.156*** (0.008)	0.162*** (0.009)	0.092*** (0.010)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.053	0.069	0.080	0.032	0.038	0.048	0.053	0.062	0.100
Observations	31,872	31,872	31,872	31,872	31,872	31,872	18,531	18,531	18,531

Panel B: Fixed Effects									
K6 Z-Score			SMI (K6≥10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.043*** (0.008)	-0.043*** (0.008)	-0.037*** (0.008)	-0.008*** (0.002)	-0.008*** (0.002)	-0.008*** (0.003)	0.044*** (0.013)	0.044*** (0.013)	0.032*** (0.013)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.005	0.006	0.010	0.001	0.001	0.003	0.010	0.010	0.019
Observations	31,872	31,872	31,872	31,872	31,872	31,872	18,531	18,531	18,531

Panel C: Lagged Dependent Variable									
K6 Z-Score			SMI (K6≥10)			Life Satisfaction Z-Score			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Log2 Real Family Income $t-1$	-0.079*** (0.006)	-0.082*** (0.007)	-0.066*** (0.008)	-0.027*** (0.002)	-0.028*** (0.002)	-0.021*** (0.003)	0.092*** (0.007)	0.096*** (0.008)	0.067*** (0.009)
Lagged DV	0.540*** (0.012)	0.532*** (0.012)	0.528*** (0.013)	0.446*** (0.027)	0.440*** (0.027)	0.432*** (0.027)	0.470*** (0.012)	0.465*** (0.012)	0.442*** (0.013)
Controls	None	Exog.	All	None	Exog.	All	None	Exog.	All
Adjusted R^2	0.330	0.332	0.335	0.118	0.121	0.127	0.271	0.274	0.287
Observations	19,701	19,701	19,701	19,701	19,701	19,701	11,145	11,145	11,145

Notes: Family income adjusted to 2013 dollars using chain-weighted CPI. Regressions weighted using PSID individual longitudinal weights. Averages of individual weights used for fixed effects regressions. Exogenous controls are gender, two race dummies, and five age dummies. Additional controls are four education level dummies, three family size dummies, unemployed, married, and divorced. All columns include year dummies. Fixed effects estimates are computed using the within estimator. Standard errors corrected for clustering at the individual level.

Table A2: First Stage Estimates: 2009-2013 Data

	(1)	(2)	(3)	(4)	(5)	(6)
Log ₂ Division-Industry	0.353	0.246	0.261	0.247	-0.142	-0.125
Mean Weekly Income $t-1$	(0.258)	(0.247)	(0.227)	(0.226)	(0.410)	(0.398)
Specification	Pooled	Pooled	FE	FE	LDV	LDV
Controls	None	Exog.	None	Exog.	None	Exog.
First-Stage F-Stat	1.87	0.99	1.33	1.19	0.12	0.10
Observations	20,093	20,093	20,093	20,093	11,793	11,793

Notes: All specifications include year dummies. Pooled and LDV regressions include cell dummies for base division and industry. Fixed effects estimates are computed using the within estimator. Standard errors corrected for clustering at the individual level.

Figure A1: Sensitivity to Exclusion Restriction: β versus δ_0

A. Prior-Adjusted 95% Confidence Intervals for $\hat{\beta}$ by Values of δ

B. Quantiles of the Distribution of γ in Terms of β by Values of δ

δ	Quantile			
	0.667	0.750	0.900	0.975
0.10	0.04β	0.07β	0.12β	0.20β
0.15	0.06β	0.10β	0.19β	0.29β
0.20	0.09β	0.13β	0.26β	0.39β
0.30	0.14β	0.20β	0.38β	0.59β

Notes: Panel A shows plots of 95% cluster-robust confidence intervals for 2SLS estimates of $\hat{\beta}$ adjusted for prior beliefs about γ where $\gamma \sim N(0, \delta^2 \beta^2)$. Estimations include controls for gender, race, age, year, and base industry-by-division cell dummies. Panel B shows the sizes of γ in terms of β that correspond to different quantiles of the prior for values of δ .

Figure A2: Sensitivity to Exclusion Restriction: β versus $\zeta_0 = \frac{\gamma}{\beta}$

Notes: Solid lines show 2SLS estimates conditional on fixed values for $\zeta_0 = \frac{\gamma}{\hat{\beta}_{2SLS}}$. Dashed lines show bootstrapped 95% confidence intervals with resampling over individuals.