Channeling your Anger: Indirect Effects of Unemployment on Wellbeing

Team 8

April 13, 2018

Abstract

Using micro-data from the European Values Survey over several decades, and European data on crime, we analyze to what extent the effects of aggregate unemployment on individual subjective wellbeing are mediated through the effects on sentiment towards immigrants and crime rates. We add to the existing literature through three key innovations: First, we aggregate several survey indicators of wellbeing into a single measure of subjective wellbeing using principal components analysis. Second, we discuss the scope and requirements for identification of causal indirect effects in wellbeing regressions and provide a simple approach to comparing indirect effect sizes. Finally, we implement an instrumental variables approach to identifying *causal* estimates of the indirect effects of unemployment through the immigration and crime channels using exogenous local labor demand shocks as an instrument. We find that immigration concerns and crime rates are plausible channels of mediated multiplier effects of unemployment. Moreover, we show that unemployment effects mediated through shifts in sentiment towards immigrants and increases in non-violent crime have a greater contribution to the effect of unemployment on wellbeing than effects mediated through increases in violent crime rates.

Contents

1 Introduction

There are several important studies that show a significant effect of unemployment on wellbeing. For instance, Clark and Oswald (1994) found that mental distress was higher among the unemployed than among the employed in Great Britain, but observed that this gap was smaller in the British regions with a higher unemployment rate. Similarly, Tella et al. (2003) show that self-reports of subjective wellbeing co-move with macroeconomic indicators over the business cycle, with unemployment rates in Europe and the United States tending to have a large, negative effect. However, the effect of unemployment on the non-employed and the precise channels through which any spillovers are not yet understood to the same degree.

In this paper, we focus on the channels through which unemployment can affect self-reported subjective wellbeing even for the employed - that is, spillovers. We focus on how unemployment affects sentiments towards immigrants and crime rates and thereby may affect wellbeing. To quantify these effects, we use the most recent waves of the European Value Survey (EVS), which contains high-quality, micro-data on self-reported subjective well-being that covers over 40 countries. In addition, we combine this survey data with a rich set of other data sources, including information on regional crime rates and industrial production, and alternative measures of unemployment and income.

To isolate these channels, we adopt the following strategy. First, we estimate the causal effect of regional unemployment on regional attitudes towards immigrants and regional crime rates. To do so, we construct "local demand shocks" following a long, established literature in labor economics and use this as an instrument for regional unemployment (Bartik, 1991). We next attempt to estimate the effect of the mediation channel on wellbeing by carefully selecting of reasonable controls.

We find that variation in local unemployment is consistently associated with a large rise in negative attitudes towards immigrants. This result is robust to a variety of empirical specifications. Moreover, we find that the effect of unemployment on wellbeing is mediated through this channel. We find that a large effect of regional unemployment on wellbeing is mediated through attitudes towards immigrants. In fact, the effect of regional unemployment on wellbeing mediated through attitudes towards immigrants is 3-5 times as large than the direct effect of regional unemployment on wellbeing. Second, we find that variation in local unemployment is only weakly associated with changes in violent and non-violent crime rates. We find suggestive evidence that the effect of regional unemployment on wellbeing mediated through violent crime is quite small but find that the effect mediated through non-violent crime is also much larger than the direct effect of regional unemployment on wellbeing.

The remainder of the paper proceeds as follows. Section 2 provides a brief framework to discuss the estimation of mediation channels for unemployment through attitudes towards immigrants and crime. Section 3 describes our data and measures of wellbeing, unemployment, attitudes towards immigrants and crime. Section 4 further describes our empirical approach. Section 5 presents the results and Section 6 concludes.

2 An Empirical Model for Mediation Analysis

There may be unobservable omitted variables that bias our estimates of the effect of unemployment on mediating variables (attitudes towards immigrants and crime) and the effect of mediating variables on wellbeing. This challenge is illustrated using a directed acyclic graph

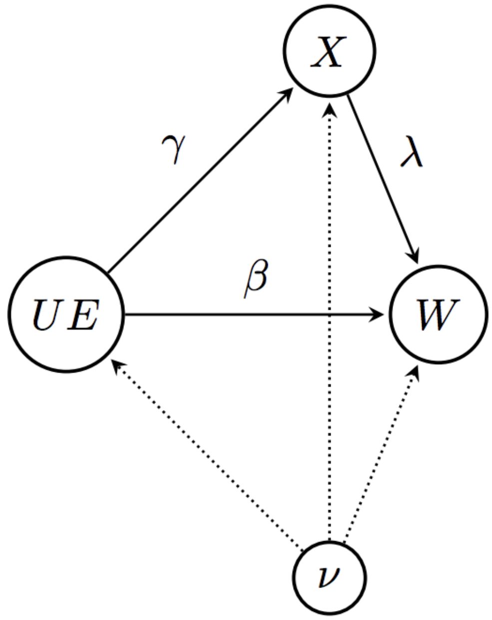

Figure 1: Causal directed acyclic graph (DAG) to represent causal relationship between unemployment UE and wellbeing W, possible mediators X and unobservables ν. See Pearl (1995) for additional details on the use of DAGs to understand causal mechanisms.

(DAG) in Figure 1.[1] We are interested in the causal effect of the mediator X on wellbeing W and the causal effect of unemployment UE on the mediator on X. However, there may be an unobserved confounder ν that simultaneously affects all three variables. As a result, we cannot identify the true causal effects β, λ or γ, unless we can eliminate the confounding path through ν.

In order to identify these coefficients, we either need to make a number of additional assumptions or use an instrument Z that is not affected by the confounder. Figure 2 illustrates the case in which we instrument for unemployment. An instrument Z that is not affected by ν allows us to identify the coefficients β and γ.

However, an instrument for unemployment does not allow us to identify λ. Instead, we could try to limit the effect of the unobserved confounder by noting that once we control for unemployment in a regression of wellbeing on X, we have limited the effect of confounders to the effect of ν on both X and W. We can then proceed to calculate indirect effects of unemployment on wellbeing as the product of γ and λ either by assuming that the causal arrow from ν to X does not exist, or by finding a different instrument for X that allows us to identify λ. That is, based on the causal relationships described in the DAGs, the indirect effect of unemployment on wellbeing mediated by X is defined as

$$\lambda \times \gamma.$$

[1] DAGs are powerful, widely-used tools in computer science and statistics for describing causal relationships in the presence of confounders and mediators. See Pearl (2009) for a classic reference on DAGs and causal mechanisms.

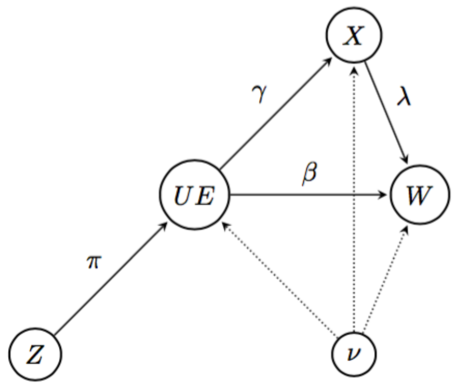

Figure 2: Causal directed acyclic graph (DAG) to represent causal relationship with instrumental variable Z, between unemployment UE and wellbeing W, possible mediators X and unobservables ν.

3 Data

3.1 European Values Study

Our individual-level data come from the European Values Study (EVS). The EVS is a repeated-cross-sectional survey conducted in over 40 European countries that collects information about respondents' beliefs, values, attitudes, and outcomes for key dimensions of their lives.[2] The survey has been conducted four times – in 1981, 1990, 1999, and 2008 – thus providing researchers with a unique way to investigate how Europeans thoughts and outcomes in topics such as family, work, religion, politics, and society have changed over several decades.

The data is nationally representative of the adult population within each country for the given collection year. Each country randomly sampled adults through a multi-stage procedure until some specified quota was met (e.g., each country aimed for 1,200 respondents in 1981). In some years, countries over sampled specific populations, but the EVS provides sample weights to recover nationally representative samples. In most years and countries, the survey was conducted as face-to-face interviews, but there are some instances where countries collected responses using other methods (e.g., phones, online, postcards). See GESIS Leibniz Institute for the Social Sciences (2015) for more details about the country-level quotas over the years and for a discussion of how some specific countries might have deviated from the sampling procedure.

The EVS dataset contains 164,997 observations with almost 1,400 variables for each individual.[3] We restrict attention to the final wave of the EVS. It is the only wave that asked respondents about attitudes towards immigrants and we can only link the survey to crime data

[2]Specifically, the EVS contains responses from Albania, Armenia, Austria, Belarus, Belgium, Bosnia-Herzegovina, Bulgaria, Canada, Cyprus, Northern Cyprus, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany (East/West), Great Britain, Greece, Hungary, Iceland, Ireland, Northern Ireland, Italy, Kosovo, Latvia, Lithuania, Luxembourg, Malta, Republic of Macedonia, Republic of Moldova, Republic of Montenegro, the Netherlands, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Switzerland, Sweden, Turkey, Ukraine, and USA. Not all countries are represented in each year of the survey. See GESIS Leibniz Institute for the Social Sciences (2015) for a more detailed description of which countries were sampled in which years. Note that although Canada and USA are not European, for convenience, we will continue to refer to the sample as "Europe" or "European" throughout.

[3]Note that we do not include 538 potential duplicate observations in our analysis sample.

after 2008. This gives us a final sample size of 66,237 observations. The construction of our measures of wellbeing and and unemployment are described in more detail below.

3.2 Measures of wellbeing

Each individual survey item that elicits respondents to evaluate a particular aspect of their life satisfaction and happiness captures a different dimension of subjective wellbeing. A priori, it is unclear which survey item best captures the "true" subjective wellbeing of the respondent. As a result, we take an eclectic approach and consider a variety of measures. We use four measures of wellbeing that we aggregate into a single, aggregated measure of wellbeing and also consider each measure individuals.

We first consider the survey items that elicit respondents' overall

- life satisfaction (*A170*)
- job satisfaction (*C033*)
- subjective state of health (*A009*)
- feelings of happiness (*A008*)

Each of these survey items appears in each wave of the European Value Survey and so, they provide the most coverage of respondents. Each of these measures have appeared in the literature on subjective wellbeing. For example, Tella et al. (2003) uses a survey measure of general life satisfaction in its analysis of Europe and a survey measure of overall feelings of happiness in its analysis of the United States. Diener (2006) argues that subjective wellbeing also incorporates assessments of job-satisfaction. Van Praag et al. (2003) considers general satisfaction, job satisfaction, financial satisfaction and health satisfaction among several others to define the subjective wellbeing of respondents.

Using these individual survey items, we construct our own measure of subjective wellbeing from the European Values Survey that aggregates responses from four survey items across the four waves. There are a variety of survey items in the European Values Survey that elicit an assessment of the quality of life from respondents. Each of these survey items provides a noisy measure of the respondent's "true" subjective wellbeing that is unobservable and so, we may wish to aggregate these responses into a single measure of subjective wellbeing.

This suggests modeling the subjective wellbeing of a respondent as a latent factor that we extract from survey responses using principal components analysis. In particular, a respondent's subjective wellbeing is the first principal component that underlies their responses to four wellbeing related survey items. Several survey items are missing for respondents either because the respondent did not answer. We use an expectation-minimization (EM) algorithm that simultaneously computes the principal components and imputes missing values following Stock and Watson (2002). In short, the EM algorithm iteratively updates the imputed missing values as the fitted values from the factors and loadings of the previous step. The principal components are then re-estimated using the data with the newly imputed missing values. The algorithm iterates until convergence. The results are summarized in Figure 3 below. As can be seen, each survey item is highly correlated with one another. This gives confidence that the principal components analysis extracts some unobserved, common signal underlying this array of items. The estimated first-principal component mostly loads onto the life satisfaction, and job satisfaction survey items. In addition to considering the aggregated measure, throughout the analysis, we will also consider how each mechanism operates on each individual survey item of subjective wellbeing.

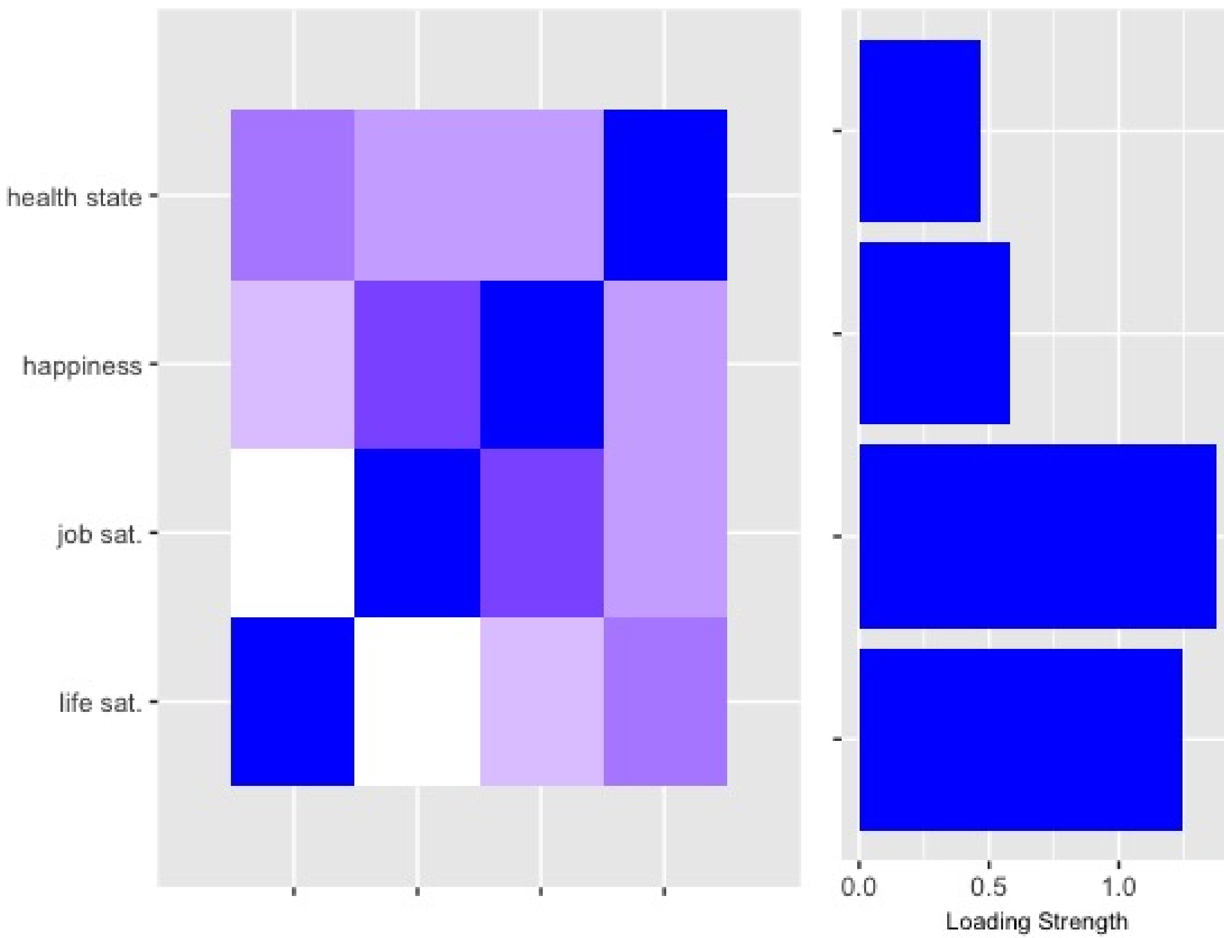

Figure 3: Pairwise correlation heatmap of wellbeing survey items and factor loadings of the first principal component. Correlation heatmap: the darker the shade of blue, the larger the correlation. Factor loadings bar chart: the larger the bar, the more represented the variable is in the first component.

3.3 Measures of unemployment and income

Our main measure of unemployment is derived from the micro data of the European Value Survey (EVS). In the EVS, each respondent is asked what her employment status is: "full-time", "part-time", "self-employed", "retired", "housewife", "student", "unemployed", or "other". We consider the responses "full-time", "part-time", "self-employed' as being employed and the response "unemployed" as being unemployed. All other responses are considered to be out of the labor force.

We then calculate the aggregate regional unemployment rate by summing all the unemployed and employed individuals to obtain the labor force and calculating the unemployment rate as the share of the labor force that is unemployed. More precisely, in all the regressions at the individual level, we will be using leave-one-out estimates of regional unemployment rates, that is, omitting each individual's own unemployment status from the aggregate included in their regression in order to avoid a mechanical correlation between the aggregate status and the individual's.

As a robustness check, we consider four additional measures of unemployment rates in our empirical approach below:

1. Regional Unemployment (Helliwell and Huang (2014)) – We reconstruct the measure of unemployment described in Helliwell and Huang (2014) using the EVS data. Specifically, this measure considers all respondents who reported their employment status as "other" or had missing data as unemployed. Our aggregation to the regional level is similar to the procedure described above.

2. Regional Unemployment (Everyone) – Similar to our main measure, this measure considers all respondents who stated their work status as "unemployed" as being unemployed. However, this measure now considers all respondents who answered the question as being in the labor force. Our aggregation to the regional level is similar to the procedure described above.

3. Regional Unemployment (C029) – In addition to asking for a respondent's work status, the EVS also directly asks respondents if they are employed or not. The respondents' responses to this variable are encoded in binary variable C029. This measure considers all respondents who answered "yes" to this question as employed and all respondents who answered "no" as unemployed. Missing responses are considered out of the labor force. Our aggregation to the regional level is similar to the procedure described above.

Table 1 reports the pairwise correlation of all the unemployment measures we consider. As can be seen, all of the measures are positively correlated. The measures created from the EVS work status variable are all highly correlated (correlation coefficients > 0.9) and the measure created from the binary EVS variable is slightly less correlated with these EVS work (correlation coefficients between 0.6 and 0.73). The national unemployment measure from the World Bank has correlations with the EVS measure that range from 0.46 to 0.75. Robustness checks that consider each of these unemployment measures are reported in Section 5 below.

Table 1: Correlation Table of Unemployment Measures

Variables	Reg–Main	Reg–Helliwell	Reg–Everyone	Reg–C029
Reg–Main	1.000			
Reg–Helliwell	0.973	1.000		
Reg–Everyone	0.977	0.952	1.000	
Reg–C029	0.723	0.704	0.609	1.000

Notes: This table reports the pairwise correlation between the various unemployment measures we are considering. See the text for a descriptions of each measure.

3.4 Measures of crime

In order to be able to assess the indirect effect of unemployment on group wellbeing through the channel of greater crime rates, we obtain data on crime incidence for a majority of European countries from Eurostat. Eurostat reports crime statistics at a number of regional aggregation levels, with the most detailed being NUTS 3. Where available, we use the NUTS 2 level incidence - where not, we aggregate the respective NUTS 3 regions to construct a NUTS 2 incidence measure. As the last survey wave of the EVS starts in 2008, we use crime incidence for 2008-2010 and sum across the three years to obtain a measure of average crime rates at the same time as our measure of unemployment.

The crime incidence is measured for both violent crimes (robberies, homicides) and non-violent crimes (burglaries, motor vehicle theft). In order to adjust crime incidence to a comparable measure across countries, we express crime incidence as a rate per 100,000 inhabitants per year by dividing the average 2008-2010 incidence by the 2008 population of the region.

To get an idea of how crime incidence varies between different European countries, we have plotted the average incidence of non-violent crime rate per year in 2008-2010 in the European countries for which we have data. Note that non-violent crimes appear to be more prevalent in Southern Europe, Western Europe, and Northern Europe, and less prevalent in Central Europe. This suggests that there seem to be broader regional and perhaps cultural patterns in crime rates that we should be concerned about in our analysis, as discussed in section 2.

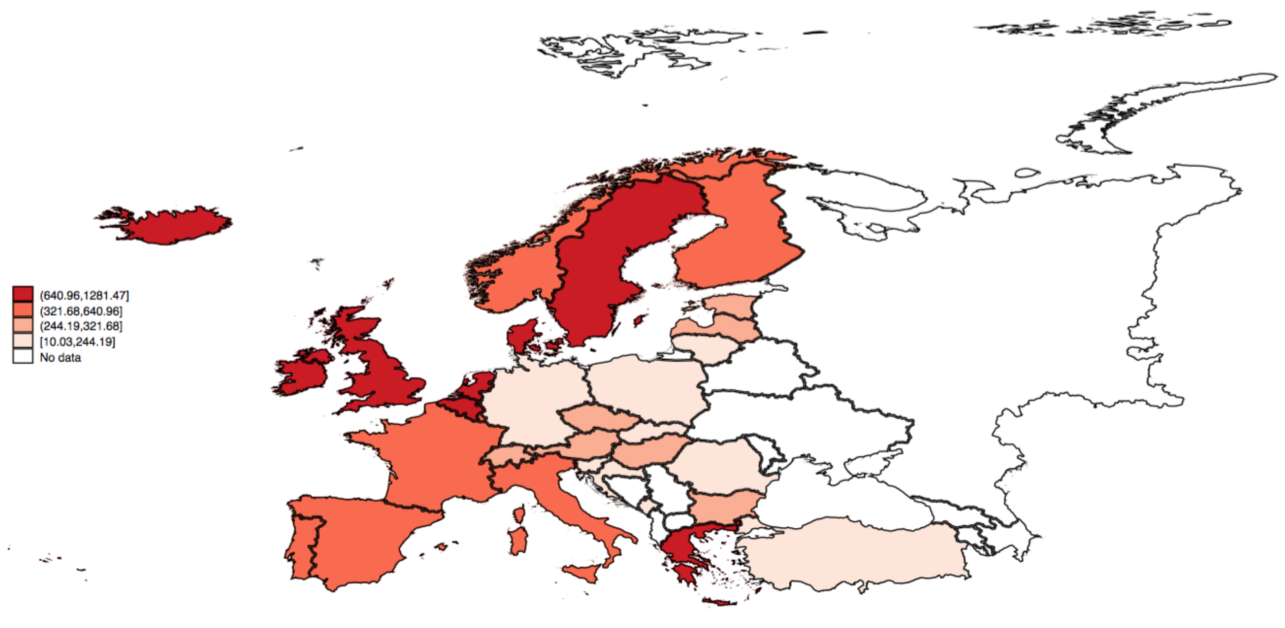

Figure 4: This figure presents the country-level nonviolent crime rates. The figure reports the number of nonviolent crimes per 100,000 people per year between 2008-2010 for each country. The darker a country, the higher the nonviolent crime rate.

3.5 Measures of sentiment towards immigrant

In order to be able to assess the indirect effect of unemployment on group wellbeing through the channel of sentiment towards immigrants, we construct a measure using data from the European Values Survey (EVS). We generate an aggregate measure of an individual's sentiment towards immigrants by averaging responses to three survey questions:

1. "Immigrants take away jobs from [respondent's nationality]"

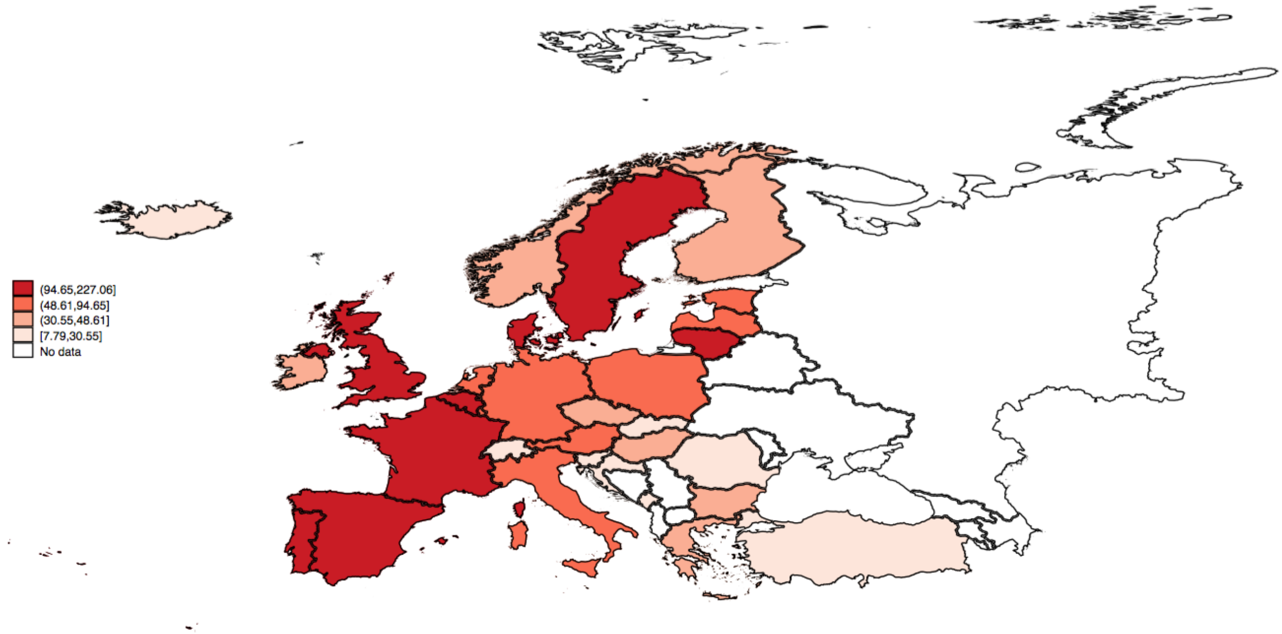

Figure 5: This figure presents the country-level violent crime rates. The figure reports the number of violent crimes per 100,000 people per year between 2008-2010 for each country. The darker a country, the higher the violent crime rate.

2. "Immigrants increase crime problems"
3. "Immigrants are a strain on welfare system"

All responses are reported on a 10-point Likert-type scale where a 1 is the most negative and a 10 is the most positive. Before averaging the responses, we rescale the distribution for each question by the standard deviation observed in the analysis sample.

Figure 6 plots the average value of negative sentiment towards immigrants that we observe for each country in the EVS. Note that the aggregate measure has been rescaled so that a positive value represents *more* negative sentiment. It appears that negative sentiment towards immigrants is more prevalent in Central Europe and some Eastern European countries, and less prevalent in Western Europe. Similar to crime rates, this suggests that there seem to be broader regional and perhaps cultural patterns in attitudes towards immigrants that we should be concerned about in our analysis, as we discussed in section 2.

3.6 Industry employment by region

In order to identify exogenous regional shocks to employment (for the last survey wave only), we construct regional-level employment by industry for the years preceding the 2008 survey. We obtain data at the NUTS 2 regions level for all available industries at the NACE Level 1 classification[4] from Eurostat. The included industries are: Mining and quarrying; Manufacturing; Electricity, gas, and water supply; construction; Wholesale and retail trade; repair of motor vehicles, motorcycles and personal and household goods; Hotels and restaurants; Transport, storage and communication; Real estate, renting and business activities.

[4]NACE refers to the "Statistical Classification of Economic Activities in the European Community". Level 1 of the industry classification includes 21 industries.

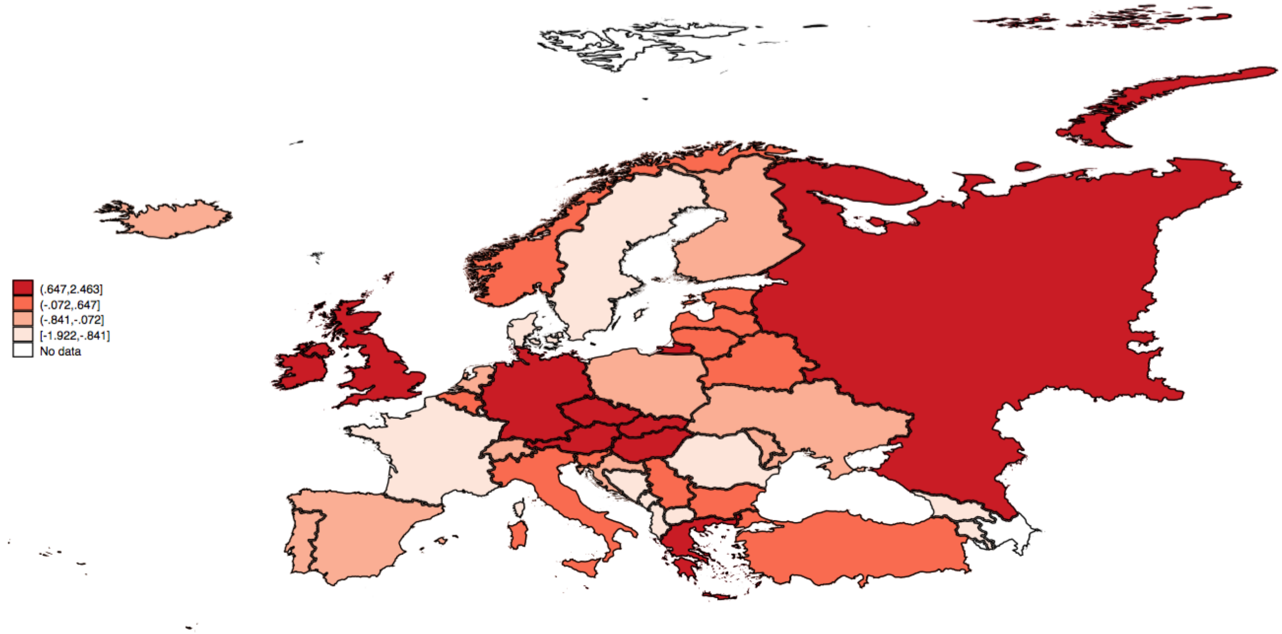

Figure 6: This figure presents the average negative sentiment towards immigrants in each country (standardized such that the country-level variable is mean zero and standard deviation one across all countries). Refer to the text for how the individual-level variable was constructed from the survey data. The darker a country, the more negative sentiment the country has towards immigrants on average.

For each industry, we obtain the number of people employed for each year from which we subsequently construct country totals for the included industries. After dropping missing observations, we end up with 314 and 320 NUTS 2 regions, respectively, for which we can construct 2002-2007 or 2006-2007 industrial employment growth shocks using the methodology detailed in the section on Bartik shocks below.

4 Empirical Approach

4.1 Estimation Procedures

The measures of wellbeing in our study are binary or ordinal in nature (for example, life satisfaction on a 10-point scale). The simple, textbook approach when working with data that have this structure is to use maximum likelihood techniques for estimation that explicitly model the structured nature of the data-generating process. Canonical examples are Logit or ordered Probit. These techniques are extremely common in the literature on happiness.

A non-linear model such as ordered Probit, which seems natural in this context, suffers from the *incidental parameters problem*. It produces inconsistent estimates of the coefficients on dummy variables such as country fixed effects. Moreover, it produces estimates that suffers from large bias in small samples (Neyman and Scott, 1948). These poor properties affect the quality of the estimates for the parameters of interest. In particular, the poor properties in the estimates of country fixed effects in an order Probit model of wellbeing will produce estimates of the coefficients on unemployment that have poor properties in small and large samples (biased and inconsistent) (Newey, 2017). Furthermore, the commonly viewed theoretical advantages of nonlinear methods only hold if we were estimating the "correct" nonlinear model. Identifying this requires very strong additional assumptions on the error term in the underlying latent model. Finally, heteroskedasticity and non-independence of errors are very challenging to deal with in non-linear models and in their presence, the general performance of MLE estimators is not well understood. As a result, Logit or order Probit models are extremely poor estimators in this context.

Given the poor behavior of non-linear estimators, we proceed using linear techniques such as ordinary least squares throughout our analysis. While the OLS estimator does not exploit the non-linear functional form of the conditional expectation function, it yields an *unbiased* and *consistent* estimate of the *best linear approximation* to the conditional expectation function. Especially in specifications with a large number of dummy variables, this approximation tends to quite be accurate. For example, Stock and Watson (2015) argue that in most applications with a binary dependent variable, OLS provides very similar estimates to Logit and Probit. Taken together, linear techniques make it clear what population object we are estimating – the best linear approximation to the population conditional expectation – and return estimates that have desirable properties in both finite and large samples. Given the concerns outlined above, when the true data-generating process is unknown, estimation with *OLS is a more cautious approach than Logit or Ordered Probit Models.* We proceed with this in mind.

Consider the following population regressions:

$$W_{ijc} = \lambda X_{ijc} + \beta_1 U_{ijc} + \beta_2 \bar{U}_{jc} + \kappa_w Z_{ijc} + \alpha_c + \epsilon_{ijc} \tag{1}$$

$$X_{ijc} = \gamma \bar{U}_{jc} + \kappa \bar{Z}_{jc} + v_i \tag{2}$$

where the subscripts denote individual i, in region j of country c. W_{ijc} is subjective wellbeing, X_{jic} is the mediator of interest, U_{ijc} is an individual-level unemployment indicator, $\bar{U}_{jc}$ is regional-level unemployment, Z_{ijc} is a set additional individual-level controls, $\bar{Z}_{jc}$ is a set of regional-level controls and α_c is country-fixed effects. Note that when we consider attitudes towards immigrants, the mediator of interest is at the individual-level i.e. individual-level attitudes towards immigrants. When we consider crime as the mediator of interest, X_{ijc} is at the regional-level. Following the framework outlined in Section 2, we interpret λ as the causal effect of the regional-level mediator of interest on individual subjective wellbeing and γ as the causal effect of regional-level unemployment on the regional-level mediator of interest. We then define the mediated effect of regional-level unemployment on individual subjective wellbeing through the regional level mediator as

$$\lambda \times \gamma$$

Put in another way, a one unit rise in regional-level unemployment on average leads to a $\lambda \times \gamma$ unit change in individual subjective wellbeing through the mediator of interest. Therefore, our empirical strategy focuses on producing quality estimates of λ and γ.

Note that throughout the analysis we do not use individual income nor individual income as a control variable. In particular, in the regression of unemployment, regional income is a "bad control." That is, regional income is an outcome of the change in the regional unemployment rate and so, holding regional income fixed would eliminate some of the variation of interest (Angrist and Pischke, 2008). Similarly, in the regression of subjective wellbeing on the mediator and unemployment, individual income and regional income are similarly bad controls as they may eliminate variation of interest.

4.2 Estimating γ: Bartik instruments

First, we describe our strategy for identifying γ. As noted above in Section 2, we may worry that unemployment and its mediator are simultaneously caused by some unobservable shocks. If this is true, we are not identifying the true population coefficients in a regression of the mediator on unemployment. One strategy to overcome this is to use an instrument to identify variation in unemployment that is plausibly exogenous. That is, we wish to identify an instrument Z that does directly affect the mediator while being a valid instrument for unemployment. This instrument can then be used in a standard two-stage least squares (TSLS) estimator.

One plausible instrument for unemployment consists of exogenous shocks to local labor demand. We follow a long and active literature (Bartik, 1991; Boustan, 2010; Boustan et al., 2010; Notowidigdo, 2011; Bartik, 2015; Maggio and Kermani, 2016; Diamond, 2016; Goldsmith-Pinkham et al., 2018) in instrumenting for changes in local labor demand using the variation in local industry employment due to national industry growth dynamics interacting with the pre-existing local industry structure. Due to data availability constraints, we can only construct this instrument for the 2008 survey wave.

To construct this instrument, we multiply regional industry exposure in a baseline year (either 2002 or 2006) by the national leave-one-out growth rate experienced by each industry. That is, using the 2006 baseline year example, for each geography g, industry i and year t, we construct the so-called "Bartik" employment shock as

$$B_{g,t} = \sum_{i=1}^{k} \phi_{g,2006,i} \cdot \left(\frac{\nu_{-g,t,i} - \nu_{-g,t-1,i}}{\nu_{-g,t-1,i}} \right),$$

where $\phi_{g,2006,i} = \frac{e_{g,2006,i}}{e_{g,2006}}$ is the industry i employment share in geography g in the baseline year. Moreover, $\nu_{-g,t,i} = \frac{\sum_{c \neq g}^{n} e_{c,t,i}}{\sum_{i=1}^{k} \sum_{c \neq g}^{n} e_{c,t,i}}$ is the leave-one-out estimate - omitting region g from both the numerator and the denominator - of the national employment share in industry i in year t. Consequently, the Bartik shock input consists of the product of the importance of industry i in the regional economy and the national leave-one-out growth rate of the share of employment in that industry. The Bartik shock for geography g in year t is then simply the sum of these national shocks, weighted by the local industry structure, across all industries. For instance, if the local industry structure remained the same as in 2006 and regional industries grew at exactly the rate of their leave-one-out equivalents at the national level, then the Bartik shock would be precisely equal to total regional employment growth in the industries contained in the sample. After computing Bartik shocks and taking into account the limited availability of the industry data, we are left with data for the last wave from 218 NUTS-2 regions to be included in the IV analysis.

For the instrumental variables analysis, we cluster our standard errors at the regional-level. We do so because the treatment that we instrument for is at the regional-level (regional unemployment) and so, the error structure of individuals within a given region is correlated. This is in-line with the recommendation of Abadie et al. (2017) - the IV design is analogous to a randomized experiment with a stratified treatment design.

4.3 Estimating λ: Regression analysis

Next, we describe our strategy for estimating λ. Once again, we are concerned that there exists some unobserved confounder that simultaneously affects the regional-level mediator of interest and subjective wellbeing. Ideally, we would identify plausibly exogenous variation in the mediator of interest and use this to estimate λ. As in the previous sub-section, one strategy is to find an instrument that is correlated with the mediator and satisfies an exclusion restriction on subjective wellbeing. However, this is a particularly challenging context to find an instrument.

For instance, suppose that the mediator of interest is regional-level crime rates. One common instrument for crime rates in the literature is local religiosity. For example, the literature on religiosity and economic and social behavior highlight differences between Protestants and Catholics that plausibly translate into investment in crime-precaution and criminality: In this literature, religious organizations are one type of institution that cultivate and incentivize certain types of behavior which develop into "cultural norms" (Weber, 1976; Foucault, 1977). Individuals in more hierarchical religions, like Catholicism have been found to have lower levels of trust in cross-country and micro-level studies (Porta et al., 1996) and to be associated with lower government performance (La Porta et al., 1999). At the same time, it is not plausible that religiosity satisfies the exclusion restriction in the second stage — variation in religiosity is likely correlated with subjective wellbeing reports.

There is a similar difficulty when considering attitudes to immigrants as a mediator. For example, Brunner and Kuhn (2014) estimates the causal effect of local immigrant density on attitudes to immigrants using immigrant share of the local labor market as an instrument. While this instrument may satisfy the exclusion restriction on subjective wellbeing, it also may be a weak instrument. In particular, in our application, we would use immigrant share of the local labor market as an instrument for attitudes towards immigrants. That is, our first-stage would be the reduced-form of Brunner and Kuhn (2014). Weak instruments pose several well-known challenges for estimation and inference (Staiger and Stock, 1997). Moreover, results on the asymptotic distribution of the two-stage least squares estimator in the presence of weak instruments have been shown to be poor approximations in finite samples (Young, 2017).

As a result, we instead estimate λ using simple regression analysis varying the set of controls. That is, we estimate

$$W_{ijc} = \lambda \bar{X}_{jc} + \beta_1 U_{ijc} + \beta_2 \bar{U}_{jc} + \kappa_w Z_{ijc} + \alpha_c + \epsilon_{ijc}$$

including control variables that may be plausible confounders. It is important to note that this does not identify the causal effect of the mediator on subjective well-being. However, given the data, this is the best we can do and it provides a rough sense of the magnitude of the mechanism.

4.4 Control variables

The choice of control variables in this empirical design needs to balance two different concerns: on the one hand, there is the risk of "bad controls" (Angrist and Pischke, 2008), that is, the issue of controlling for variables that are also outcomes of changes in regressor of interest, such that holding them constant would in fact be eliminating some of the variation of interest.

On the other hand, there is a concern that a given regressor is endogenous with regard to wellbeing. In that case we might be able to eliminate most of the risk of unobservables causing the regressor and wellbeing jointly by controlling for all of the characteristics that might be jointly causing both wellbeing and the regressor.

As a result of these considerations, we consider a specification that includes standard variables identified in the literature on subjective wellbeing (see, e.g. Tella et al. (2003)) as affecting wellbeing and which are also plausibly not causally dependent on the regressor of interest.

For our attitudes towards immigrants regressions, the resulting list includes the sex interacted with educational level, age and age squared of the respondent, as well as whether they are married, widowed, separated or divorced. Moreover, we also include separate variables for having any children at all and the total number of children (also interacted by sex), as the extensive and intensive margins of having children may have differential effects as noted for example by Aaronson et al. (2014) in the context of the fertility transition in the U.S. South.[5] However, since the outcome variable in our crime rate regressions only has regional-level variation, all of the individual-level controls listed above can no longer be used. In these regressions, we instead control for country-level fixed effects and log population of the regions. These controls are included in both the OLS and Bartik shock specifications.

5 Results

5.1 Effects of unemployment on attitudes towards immigrants

When we estimate the effects of unemployment on intermediate social outcomes γ as described in Equation 2 for the channel of concerns about immigration, we obtain the results shown in Table 2 and 3.

Table 2: Effect of Regional Unemployment on Immigration Sentiment - OLS

	(1) Immigr. Sentiment	(2) Steal Jobs	(3) Increase Crime	(4) Strain Welfare
Regional Unemp. Rate	-0.2341*	-0.1392	-0.1689	-0.4121**
	(0.1362)	(0.1396)	(0.1504)	(0.1621)
Employment Status FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Gender FE	Yes	Yes	Yes	Yes
Gender x Demographics	Yes	Yes	Yes	Yes
R-squared	0.1068	0.1228	0.0926	0.1101
Observations	62866	61946	61083	60046

* p<0.10, ** p<0.05, *** p<0.01

Notes: Standard errors clustered at the region level.

First, consider the simple OLS results in Table 2: We find a consistently significant association of higher regional unemployment with a more negative attitude towards immigrants. As the dependent variable does not have a natural scale, the size of these coefficients is somewhat difficult to interpret, but will be discussed in more detail below in our analysis of indirect effect sizes. When we use the exogenous labor demand shocks to identify the local average treatment effect of unemployment in an instrumental variables setting for the same outcome variables, we obtain the coefficients shown in table 3. Note that the IV coefficients are bigger by a factor of 10-40 than the OLS estimate. This change in size combines two different effects: On the one

[5] The corresponding variable names in the EVS data are the following: Education level (X025); Sex (X001); Age and Age Squared (X003); Married, Widowed, Divorced, Separated (X007); Number of Children and Any Children (X011).

hand, the labor demand shock instrument is only available for a smaller population of regions, so the results reflect a subset of the regions included in the OLS analysis. Moreover, the effect is identified from the "compliers" in unemployment as a result of a labor demand shock, that is, from the marginal people gaining or losing a job as a result of the shock. We might expect that population to have stronger shifts in sentiments than those who are long-term unemployed, for instance.

However, the overall message is quite consistent: Regional unemployment causes a significant worsening in the population's attitude towards immigrants, which may therefore represent one potential mediating channel for group effects of individual unemployment.

Table 3: Effect of Regional Unemployment on Immigration Sentiment - IV

	(1) Immigr. Sentiment	(2) Steal Jobs	(3) Increase Crime	(4) Strain Welfare
Regional Unemp. Rate	-2.7164	-3.4926*	-0.3309	-4.1607**
	(1.8317)	(2.0316)	(1.8731)	(1.9944)
Employment Status FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Gender FE	Yes	Yes	Yes	Yes
Gender x Demographics	Yes	Yes	Yes	Yes
First Stage F-Statistic	67.8233	68.2375	68.4526	67.5034
Observations	21247	20989	20829	20376

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Notes: Standard errors clustered at the region level.

5.2 Effects of unemployment on crime

When we consider regional crime rates as the mediator of interest, our estimates of γ are summarized in Tables 4 and 5. Since the crime rate observations are at the regional level, the analysis is collapsed so that regressions are run at the region (NUTS-2) level. Both tables control for country fixed effects and log population of the regions (to capture differences between city and rural areas). We find that the regional unemployment rate has large and relatively precisely estimated positive effects on crime. In particular, a one percentage point increase in unemployment increases the annual rate of violent crimes by 152.7 instances per 100,000, the annual rate of non-violent crimes by 471.7, and the annual rate of homicides by 1.05 instances per 100,000.

Unfortunately when collapsed to the region level and weighting all regions equally, the Bartik shock instruments produce only a weak first stage on regional unemployment. This yields imprecise IV point estimates, as can be seen in Table 5.

5.3 Mediation results

Table 6 in the main text present the results of our specifications to estimate λ for attitudes towards immigrants and crime rates. Table 6 estimates λ in a simple OLS specification of individual-level aggregated wellbeing on regional unemployment rates, individual attitudes towards immigrants and crime rates. We are interested in the coefficients on attitudes towards

Table 4: Effect of Regional Unemployment on Crime - OLS

	(1) Violent Rate	(2) Non-violent Rate	(3) Homicide Rate
Regional Unemployment Rate	152.7164***	471.6787***	1.0472***
	(53.9822)	(172.5909)	(0.2452)
Country FE	Yes	Yes	Yes
Population (Log)	Yes	Yes	Yes
R-squared	0.4468	0.7329	0.8395
Observations	195	195	195

* p<0.10, ** p<0.05, *** p<0.01

Notes: Standard errors clustered at the region level.

Table 5: Effect of Regional Unemployment on Crime - IV

	(1) Violent Rate	(2) Non-violent Rate	(3) Homicide Rate
Regional Unemployment Rate	-561.0029	3145.5390	1.5408
	(1028.3612)	(3057.4143)	(2.8010)
Country FE	Yes	Yes	Yes
Population (Log)	Yes	Yes	Yes
First Stage F-Statistic	1.4122	1.4122	1.4122
Observations	151	151	151

* p<0.10, ** p<0.05, *** p<0.01

Notes: Standard errors clustered at the region level.

immigrants, the violent crime rate and the non-violent crime rate. A one standard deviation increase in positive attitudes towards immigrants leads to a 0.06 unit increase in the aggregate well-being measure. The point estimate on violent crime is not significant and should interpreted as a null result. Finally, the non-violent crime rate appears to have a significant negative effect on aggregate wellbeing (a one standard deviation increase in the violent crime rate reduces wellbeing by 0.15 units). Table 9 in the appendix uses the Bartik instrument from earlier as an instrument for unemployment in the same specification. In both of these estimation procedures, the estimated coefficients on these variables do not vary substantially across specifications. For simplicity, we proceed with the reported OLS estimates

Table 6: Direct Effects on Wellbeing ($\hat{\lambda}$) - OLS

	(1) Wellbeing PC	(2) Wellbeing PC
Regional Unemployment Rate	-0.2976*	0.0324
	(0.1761)	(0.1908)
Immigrant Sentiment		0.0615***
		(0.0115)
Violent Rate		-0.0001
		(0.0002)
Non-violent Rate		-0.0004***
		(0.0001)
Employment Status FE	Yes	Yes
Country FE	Yes	Yes
Gender FE	Yes	Yes
Gender x Demographics	Yes	Yes
R-squared	0.1044	0.1184
Observations	103883	29474

* p<0.10, ** p<0.05, *** p<0.01

Notes: Standard errors clustered at the region level.

To compute the mediated effect of unemployment on wellbeing mediated through attitudes towards immigrants and crime rates, we follow the discussion in Section 2. In particular, the mediated effect is defined as

$$\hat{\lambda} \times \hat{\gamma}$$

$\hat{\lambda}$ is drawn from Column (2) of Table 6. $\hat{\gamma}$ is pulled from Column (1) of Table 3 for attitudes towards immigrants. Because the IV estimates for the effects of unemployment on crime are highly imprecise, we instead use the OLS estimate for $\hat{\gamma}$. That is, we use Columns (1) and (2) of Table 4 for the violent and non-violent crime rates respectively. The mediation results are summarized in Table 7.

We find that the effect of regional unemployment on wellbeing intermediated through attitudes towards immigrants and non-violent crimes is larger than the direct effect of regional unemployment on wellbeing. In particular, we find that a 1% increase in the regional unemployment rate leads to a 0.17 unit decline in wellbeing through a rise in negative attitudes towards immigrants. Similarly, a 1% rise in the unemployment rate leads to a 0.17 unit decline in wellbeing through a rise in non-violent crime. These effects are large and translate into a

Table 7: Mediated Effects of Unemployment on Wellbeing

Mediator (X):	(1) Immigrant concern	(2) Violent crime	(3) Non-violent crime
Unemp. effect on mediator ($\hat{\gamma}$)	-2.72	152.7164	471.6787
Mediator effect on wellbeing ($\hat{\lambda}$)	0.0615	-0.0001	-0.0004
Indirect effect estimate ($\hat{\lambda} \cdot \hat{\gamma}$)	-0.17	-0.02	-0.19

Notes: For crime the effects of unemployment on the mediator are obtained from the OLS regressions, as the IV estimates were too noisy, whereas for the $\hat{\gamma}$ for immigrant concern we use the IV estimate of the coefficient.

roughly 0.1 standard deviation decline in wellbeing through the attitudes towards immigrants and non-violent crime channels.

5.4 Robustness check: Definitions of unemployment

In order to show the robustness of our results, we repeat the calculation of the multipliers through the immigration sentiment and crime channels using the different definitions of the unemployment rate defined in Section 3 above. Table 8 shows these calculations for each of the 4 unemployment measures, with the results being largely stable across different definitions.

Table 8: Indirect Effect Estimates ($\hat{\lambda} \cdot \hat{\gamma}$) for Various Unemployment Measures

	(1) Main	(2) Helliwell	(3) Everyone	(4) C029
Immigrant concern	-0.17	-0.24	-0.48	-0.15
Violent crime	-0.02	-0.002	-0.01	-0.01
Non-violent crime	-0.19	-0.04	-0.04	-0.11

Notes: This table reports the indirect estimate for each of the unemployment measures described in Section 3. Each column reports estimates a different unemployment measure. See Table 7 for a description of how they were calculated.

5.5 Robustness check: Direct effects on sub-components of wellbeing

Finally, we consider how the direct effects of regional unemployment, attitudes toward immigrants and crime rate vary across different sub-components of wellbeing. In particular, we consider how the direct effect of each variable on life satisfaction, job satisfaction, happiness and subjective state of health. To do so, we plot the coefficients for each variable on the aggregated wellbeing measure along with the coefficients for each variable on each component multiplied by that component's loading in the aggregated wellbeing measure. The results are presented in Figures 7 to 10 in the appendix. As can be seen, the direct effect of regional unemployment, attitudes to immigrants and crime rates each move each component measure of wellbeing in the same direction as the aggregate measure of wellbeing.

6 Conclusion

Using micro-data from the European Values Survey, we analyzed the two channels through which unemployment affects subjective well-being. In particular, we consider how unemployment affects attitudes towards immigrants and crime and thereby affects subjective wellbeing.

We find that immigration concerns and crime rates are plausible channels of mediated multiplier effects of unemployment. Moreover, we show that unemployment effects mediated through shifts in sentiment towards immigrants and increases in non-violent crime have a greater contribution to the effect of unemployment on wellbeing than effects mediated through increases in violent crime rates.

These results suggest that indirect effects analyses such as the approach presented in this paper may have important implications for policymakers: On the one hand, when mitigating the negative effects of unemployment, policymakers need to take into account the costs imposed on the employed population through indirect effects. On the other hand, there may be substantial benefits to be had from identifying the most important channels mediating the wellbeing effects of unemployment and then targeting interventions towards these channels.

Building on these results, future research should explore further potential *causal* channels through which unemployment may affect wellbeing by identifying more plausible instruments and other quasi-experimental tools through which these mediated wellbeing effects can be reliably identified.

References

Aaronson, D., F. Lange, and B. Mazumder (2014): "Fertility transitions along the extensive and intensive margins," *American Economic Review*, 104, 3701–24.

Abadie, A., S. Athey, G. W. Imbens, and J. Wooldridge (2017): "When Should You Adjust Standard Errors for Clustering?" *NBER Working Paper No. 24003*.

Angrist, J. D. and J.-S. Pischke (2008): *Mostly harmless econometrics: An empiricist's companion*, Princeton University Press.

Bartik, T. J. (1991): "Boon or Boondoggle? The debate over state and local economic development policies," .

——— (2015): "How Effects of Local Labor Demand Shocks Vary with the Initial Local Unemployment Rate," *Growth and Change*, 46, 529–557.

Boustan, L. P. (2010): "Was postwar suburbanization white flight? Evidence from the black migration," *The Quarterly Journal of Economics*, 125, 417–443.

Boustan, L. P., P. V. Fishback, and S. Kantor (2010): "The effect of internal migration on local labor markets: American cities during the great depression," *Journal of Labor Economics*, 28, 719–746.

Brunner, B. and A. Kuhn (2014): "Cultural Distance and Attitudes towards Immigrants: Evidence from Swiss Voting Results," Tech. rep., IZA.

Clark, A. E. and A. J. Oswald (1994): "Unhappiness and unemployment," *The Economic Journal*, 104, 648–659.

Diamond, R. (2016): "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000," *The American Economic Review*, 106, 479–524.

Diener, E. (2006): "Guidelines for National Indicators of Subjective Well-Being and Ill-Being," *Journal of Happiness Studies*, 7, 397–404.

Foucault, M. (1977): *Discipline and Punish: The Birth of the Prison*, Peregrine Books, Vintage Books.

GESIS Leibniz Institute for the Social Sciences (2015): "EVS 1981 - 2008 Variable Report: Longitudinal data Files," .

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2018): "Bartik Instruments: What, When, Why and How?" Tech. rep., National Bureau of Economic Research.

Helliwell, J. and H. Huang (2014): "NEW MEASURES OF THE COSTS OF UNEMPLOYMENT: EVIDENCE FROM THE SUBJECTIVE WELL-BEING OF 3.3 MILLION AMERICANS," *Economic Inquiry*, 52.

La Porta, R., F. Lopez-de Silanes, A. Shleifer, and R. Vishny (1999): "The quality of government," *Journal of Law, Economics, and organization*, 15, 222–279.

Maggio, M. D. and A. Kermani (2016): "The importance of unemployment insurance as an automatic stabilizer," Tech. rep., National Bureau of Economic Research.

Newey, W. (2017): "Course materials for 14.386 - New Econometric Methods," *MIT OpenCourseWare*.

Neyman, J. and E. L. Scott (1948): "Consistent Estimates Based On Partially Consistent Observations," *Econometrica*, 16, 1–32.

Notowidigdo, M. J. (2011): "The incidence of local labor demand shocks," Tech. rep., National Bureau of Economic Research.

Pearl, J. (1995): "Causal Diagrams for Empirical Research," *Biometrika*, 82.

——— (2009): *Causality: Models, Reasoning and Inference*.

Porta, R. L., F. Lopez-De-Silane, A. Shleifer, and R. W. Vishny (1996): "Trust in large organizations," Tech. rep., National Bureau of Economic Research.

Staiger, D. and J. Stock (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65.

Stock, J. and M. Watson (2002): "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147–162.

——— (2015): *Introduction to Econometrics*.

Tella, R. D., R. J. MacCulloch, and A. J. Oswald (2003): "The macroeconomics of happiness," *Review of Economics and Statistics*, 85, 809–827.

Van Praag, B. M., P. Frijters, and A. Ferrer-i Carbonell (2003): "The anatomy of subjective well-being," *Journal of Economic Behavior & Organization*, 51, 29–49.

Weber, E. (1976): *Peasants Into Frenchmen: The Modernization of Rural France, 1870-1914*, History e-book project, Stanford University Press.

Young, A. (2017): "Consistency without Inference: Instrumental Variables in Practical Application," Tech. rep.

A Additional Tables & Figures

Table 9: Direct Effects on Wellbeing - IV

	(1) PCA_wellbeing	(2) PCA_wellbeing
Regional Unemployment Rate	2.1453	0.0606
	(2.2624)	(1.4385)
Immigrant Sentiment		0.0535***
		(0.0155)
Violent Rate		0.0001
		(0.0003)
Non-violent Rate		-0.0005***
		(0.0002)
Employment Status FE	Yes	Yes
Country FE	Yes	Yes
Gender FE	Yes	Yes
Gender x Demographics	Yes	Yes
First Stage F-Statistic	68.6805	5.3082
Observations	21548	17350

* p<0.10, ** p<0.05, *** p<0.01

Notes: Standard errors clustered at the region level.

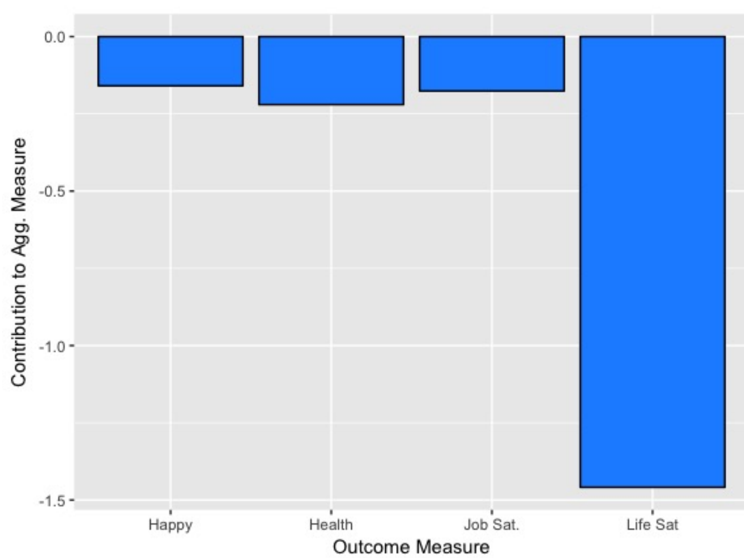

Figure 7: This figures present the direct effects of regional unemployment times the loadings on the different measures of well-being.

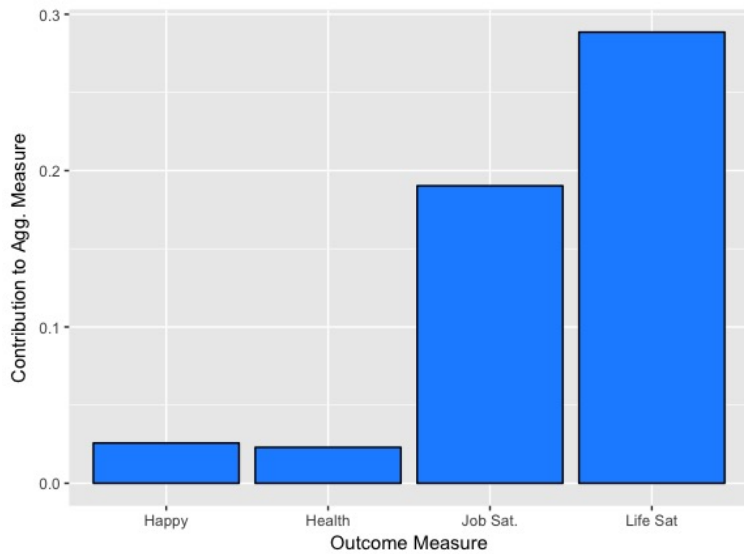

Figure 8: This figures present the direct effects of attitudes towards immigrants times the loadings on the different measures of well-being.

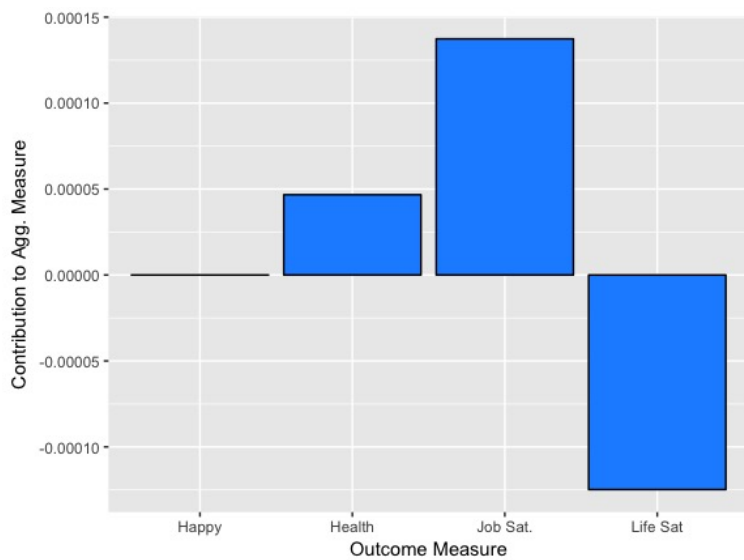

Figure 9: This figures present the direct effects of violent crime times the loadings on the different measures of well-being.

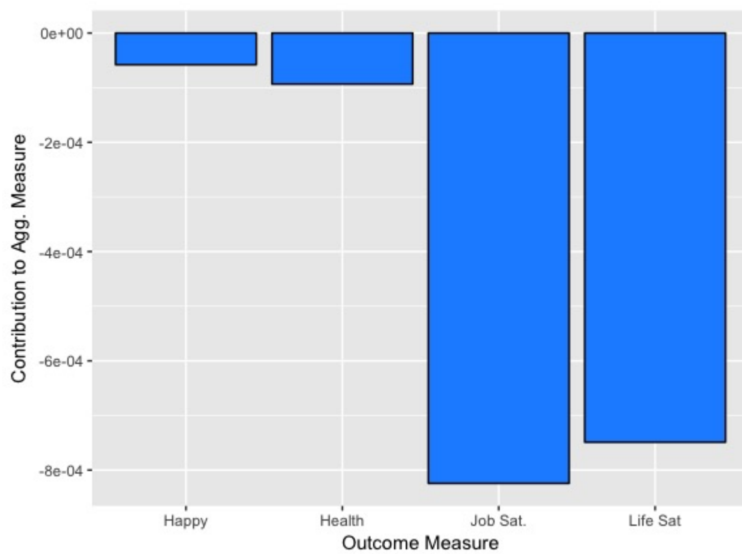

Figure 10: This figures present the direct effects of non-violent crime times the loadings on the different measures of well-being.