

Lockdown Blues: The Effect of Social Norms on the Psychological Cost of Unemployment During the COVID-19 Pandemic

Dylan Hallahan

Supervisor: Benjamin Schoefer

April 29, 2022

ABSTRACT

This paper explores and demonstrates the ability of social norms to moderate the emotional toll of unemployment. I utilize the traditional method of using the unemployment rate as a proxy for the strength of the social norm to work. This model predicts a higher unemployment rate would alleviate the decline in subjective well-being caused by unemployment. The relationship is tested by regressing subjective well-being on an interaction between individual unemployment and the unemployment rate. This test suggests a higher unemployment rate does mitigate the psychological cost of unemployment both before and during the COVID-19 pandemic. This paper acts as a preliminary analysis of the relationship between unemployment and subjective well-being during the COVID-19 pandemic and provides an analysis of the relationship both before and during the pandemic.

Well, we know when people lose work, they are miserable. So, if we're going to have less work, I'd like to see everybody have a little bit less rather than many people not working at all.

-David Autor

1 Introduction

According to the conventional neoclassical labor supply model, individuals trade their leisure for employment and compensation. If a person becomes unemployed, the lost wages are defined as the pecuniary cost of their unemployment. According to this model, if a person is adequately compensated for their lack of salary, they should be at the same or greater utility level than before they were laid off. Unexpectedly, researchers have found time and again that a loss of employment garners large non-pecuniary or psychological costs of unemployment. These costs seem to be substantial and are often seen to be dramatically higher than any loss in income would indicate. This raises a few important questions that are not only relevant to labor economics but also to public policy, especially as we move towards a future that may include bouts of prolonged unemployment for some portion of our working-age population (Autor 2019).

As it stands, there are many possible explanations as to the exact channels that lead from unemployment to unhappiness. A foundational theory laid out in Jahoda (1981) proposes that work provides several non-monetary benefits and thus contains a “latent function.” This function encapsulates the benefits derived from time structure, social interaction, setting goals and defining purpose, personal status, and enforcing activity. Each of these channels has its own set of evidence and proponents; however, it seems likely that all play some role in the decline of subjective well-being (SWB) associated with suddenly losing one’s job. The majority of contemporary research focuses on novel ways of studying these different forces so that we can better under-

stand the experiences and mental state of the unemployed. Currently, the relative strength of each proposed channel is far from fully understood or agreed upon. This presents a problem, as adequate policy response may vary heavily depending on which channel it is targeting.

COVID-19 presents an interesting opportunity to study this topic, as it brought large shifts in employment and potentially in the non-monetary costs associated with unemployment. This paper utilizes these shifts to study the relationship between the SWB and both the national and state-level unemployment during the COVID-19 pandemic and recession, comparing it to the same relationship pre-pandemic. This is particularly relevant to the branch of study that examines the role of societal norms and pressures on SWB. Historically, the employment status of others is often used as a proxy to understand the relative pressure on any individual generated both by society and one's internal measure of self-worth. This paper utilizes data from the NHIS survey alongside recent data from the household PULSE survey. I use this data to compute both individual and aggregated measures of subjective well-being. The primary result is an individual-level regression that tests the interaction between the unemployment rate and employment status both before and after the pandemic.

This paper contributes to three primary bodies of literature. The first and most direct is the work done examining the effect of the unemployment of others and social norms on unemployment-related SWB loss. These papers look to insight from psychology that suggests individual happiness is often determined through comparison to those in close physical proximity. The intuition is that an individual will be less affected if they live in an area with higher rates of unemployment. These people are using the circumstances of others to determine how they feel about themselves. Papers in the area include Jahoda (1981), Di Tella et al. (2003), Clark (2003), Clark et al. (2010), Shelds et al. (2009), and Knight et al. (2021). These papers will be the primary focus of the literature review, as they suggest a casual relationship that is an important factor in interpreting the findings of this paper. These authors

primarily rely on panel data and focus on regional unemployment as a source of heterogeneity and focus on the effects of immediate social comparison rather than a social comparison to aggregate unemployment.

The second body of literature focuses on the effects of the COVID-19 pandemic on mental health, subjective well-being, and the labor market. Of those that focus on mental health, Bryson (2021), Nordt (2020), Sobotková (2022), and Achdut (2020) are relevant. Nordt (2020) is especially relevant, as it creates a model using historical data to try to estimate suicide rates as the result of unemployment on a global scale. Fisher et al. (2020), is also relevant, as it utilizes state-by-state lockdown timing differences to perform an analysis of the effect of the lockdown on mental health. Closely related is the recent work on SWB such as Katja et al. (2021), which examines the effect of remote work on SWB, and Roberto et al. (2020), which looks at the SWB cost of lockdown isolated from the overall effects of the pandemic. This paper also speaks to the recent work on the somewhat peculiar labor market conditions surrounding the pandemic, as well as gender imbalances in both mental health (Spantig et al., 2021) and labor force participation (Fuller et al., 2020).

The final group of research this paper contributes to are studies on the effect of unemployment as a signal to future employers. Traditionally, papers such as Katz (1991) have proposed a theory that firms don't possess complete information as to the cause of an employment gap, i.e. the employee could have been laid off due to their own incompetence, or due to something out of their control; therefore, post unemployment wages won't differ between the two groups. Of particular interest is Nakamura (2008), which suggests that the proportion of workers laid off due to their individual productivity is lower during recessions and is, therefore, a less valid signal for prospective employers. The trends in SWB observed in this paper could be an extension of reduced fear of becoming the victim of unemployment stigma. This could also explain why many workers were willing to delay reentering the labor force during the pandemic, perhaps for a better opportunity down the line, and also

why some of their unemployment-related SWB declines may have been temporarily elevated. Taken more broadly, it could also help explain an apparent undervaluation of non-employment suggested in Jäger et al. (2020).

2 Literature Review

This literature review will primarily focus on the use of SWB as a tool for measurement, and on papers that relate SWB to the unemployment of others.

2.1 Measurement

Determining a coherent yardstick for subjective well-being can be somewhat more complicated than other more traditional measures such as income. Because it is self-reported, researchers must rely heavily on the survey design, which can often change over time in small ways, such as the wording of a particular question. The most direct way to measure SWB is to only look at surveys that ask participants to rate their level of satisfaction with their own life. The problem is that often surveys don't include this specific question and prefer to ask a slew of questions aimed at assessing mental health and psychological well-being. An example of this type of question from the PULSE survey would be: "Over the last 2 weeks, how often have you been bothered by feeling down, depressed, or hopeless?" So far this paper has used mental health and subjective well-being somewhat interchangeably. The reason for this is that they are often considered to be the same variable in this line of study. One convention in SWB research is to use an index that takes the average of several questions surrounding mental health and uses it as a proxy for SWB. The idea is that psychological well-being (PWB) is a very close approximation of subjective well-being, as they are both built off the more abstract concept of general well-being (Hayes et al., 2013). Research that studies the causation between the two finds a strong causal relationship between PWB and SWB. Despite this, some researchers do still prefer

to differentiate the two, and in surveys where both are independently available, it is not uncommon to regress upon both separately.

Another feature of this field of study is that it can roughly be divided into research that uses either cross-sectional or longitudinal data. Each method offers unique strengths and insights. Cross-sectional data provides the most robust and consistent examples of the difference in SWB between employed and unemployed workers. However, cross-sectional data typically fails to offer any information on causality. One common observation is that there is a strong possibility for reverse causation, i.e. workers with poor mental health and lower SWB may be more likely to be laid off. Here, longitudinal data can give information on SWB for the same individual before and after being laid off, giving us stronger evidence for causation. For a complete understanding of the relationship, both methods should be used in tandem to show both causation and magnitude.

2.2 Establishing the Non-Pecuniary Cost of Unemployment

When looking at causal evidence to support the SWB cost of unemployment, it is important to note that the current research is heavily biased toward a small set of surveys and countries. In particular is the German Socio-economic Panel (SOEP), the British Household Panel Survey (BHPS), the 2008 UK Household Longitudinal Study (UKHLS), and the Household, Income and Labor Dynamics in Australia (HILDA). These surveys are some of the only longitudinal surveys that include questions that we can use to measure SWB. These surveys are also taken often enough that the effects of recent unemployment can be observed. Because of this, they make up a large portion of available research and may create a biased result if applied to other countries. Due to the importance of societal norms in this research, it is highly likely that these estimates vary across cultures.

With that being said, there is a general consensus that there is a strong causal

link between SWB and unemployment. For the BHPS data, Clark and Oswald (1994), Theodossiou (1998), and Clark (2003) all find a strong negative coefficient. Theodossiou in particular notes a sharp rise in symptoms of anxiety and depression among the unemployed, as well as a stronger effect for men and middle-aged people (ages 23-50). Using the SOEP data, Winkelmann and Winkelmann (1995) and Gerlach and Stephan (1996) use different methods to arrive at the same conclusions on both the negative coefficient and gender and age differences. Of particular importance to this paper is the meta-analytical study conducted by McKee-Ryan et al. (2005) that looked into both cross-sectional ($n = 52$) and longitudinal ($n = 10$) magnitude. They utilized the Cohen's d method to find the standardized mean differences between multiple samples with different scales. Across the cross-sectional papers, they found that unemployed workers had dramatically lower psychological well-being ($d = -.57$), along with lower levels of subjective and objective physical well-being. Of the longitudinal studies, they found a negative correlation ($d = -.39$) associated with unemployment and a corresponding ($d = .89$) correlation associated with reemployment. Cohen's d is the difference in means as measured by the standard deviation, so a $-.57$ is interpreted as the average physiological well-being of an unemployed individual being $.57$ of a standard deviation lower than an employed person. The convention is that anything lower than 0.2, 0.4, and 0.6 are considered small, medium, and large effects respectively. This evidence presents a convincing argument as to the causal relationship between unemployment and a sharp decrease in SWB, psychological well-being, and physical well-being.

This longitudinal research is very helpful when it comes to identifying causality. Intuition suggests a few very reasonable threats to internal validity. The first and most commonly discussed is the possibility of reverse causation. The thought is that poor mental health or other self-reported measures could be causing the unemployment, and therefore causing the unemployed to have lower levels of well-being. It is also possible that the same individuals with low reported well-being will stay unemployed

for longer, and their lack of well-being could result in depression, which could then result in not having the motivation to search for work, undervaluing the wages they could garner by working. This problem is further muddied by the possibility that SWB could be related to an individual's decision to leave the labor market. As it stands, there is no research that looks at the existence of this relationship.

Another issue in studying this relationship is the strong possibility of compositional issues. Unemployed individuals could be predisposed to having a lower level of SWB and in times of high unemployment, individuals who would normally be employed with higher SWB could enter the unemployed pool and drive the number up, giving the impression of a relationship. The longitudinal difference in difference evidence helps dismiss both of these concerns. Not only is the drop in well-being almost instantaneous after unemployment, some research even suggests that the drop occurs at the date they find out they will be unemployed, not the day they are laid off. All of the papers here find no relationship between SWB and an employee's chances of being fired.

Another method of quantifying the real cost of unemployment and its magnitude is to calculate the amount of money that an unemployed person would have to be compensated to bring them to the same level of SBW as before they were laid off. Using this method, Winkelmann and Winkelmann (1998) calculate that income would have to be seven times higher than it was pre-unemployment. Other papers such as Frijters et al. (2004) find a stronger effect for the drop in income relative to unemployment, suggesting only an 82% to 104% increase. Blanchflower and Oswald (2004) use cross-sectional data in the United States to estimate that compensation should be around \$60,000 per year in addition to unemployment benefits. Additional research finds limited adaptation over time, i.e. those who remain unemployed fail to adapt to their circumstances and return to a pre-unemployment level of SWB. In particular, Clark and Georgellis (2013) and Bauer et al. (2015) find no evidence for adaptation in either men or women. This evidence seems strong, however, I

believe there is an issue incorporating the behavior of individuals who leave the labor force. The movement of discouraged workers could present a problem for much of this research as their decision to leave the labor market may be correlated with either higher or lower levels of SWB.

2.3 The Effect of Social Norms and the Unemployment of Others

Of key importance to this paper is past work on the relationship between unemployment-related SWB, the employment of others, and the effect of social comparisons. In this area, there is generally less consensus and no meta-analysis studies providing conclusive evidence. The overarching trouble is that there are two conflicting intuitions that may outweigh the other. The first is that the unemployment rate in an area (large or small) reflects an individual's opportunity for reemployment. If a person becomes unemployed in a region or time with a higher unemployment rate and a poor job market, they will personally have a harder time finding a job and therefore are expected to suffer from more unemployment-related SWB loss. On the other hand, an unemployed person could look at their surroundings, compare themselves to other unemployed people, and use that as a measuring stick to gauge how they should feel about their own unemployment. The second intuition builds on the work of Akerlof (1980), which suggests that a social norm to work plays a large role in labor market incentives and that the unemployment rate can be used as a proxy for the strength of the social norm. These two arguments present a negative and positive correlation respectively.

While not stated in the current literature, I believe there is a third intuition that suggests a reduction in SWB loss when the unemployment rate is high. Nakamura (2008) shows that during recessions, the proportion of workers laid off due to their individual low productivity is decreased. This implies that firms will have a harder

time using a gap in employment to determine the value of a potential employee. My suggestion is that individuals are in some way aware of this reduced stigma. They believe that job loss might not incur the same level of discrimination as it might have if they were let go during a period of low employment. This interaction could lead to a reduction in SWB loss associated with unemployment during recessions.

Although it is not the focus of the McKee-Ryan et al. (2005) meta-analysis, the authors briefly discuss the unemployment rate variable as an additional moderating variable. They use research from Cohn (1978) and Turner (1995) to argue and conclude that the first intuition should win out and the relationship should be negative. While Turner (1995) does present a strong argument, more recent data and research suggest the opposite conclusion. Di Tella et al. (2003) use Eurobarometer data for twelve countries over almost twenty years and focus on the effects of aggregate unemployment rates rather than local rates. Interestingly, they find that unemployment is correlated with lower happiness in every country they studied, except Spain.

More importantly, they find that life satisfaction is negatively correlated for both the employed and the unemployed and that when the employment-unemployment gap shrinks, unemployment rates go up. Clark (2003) uses longitudinal data to come to a similar conclusion with stronger causal evidence. However, he finds the relationship to be dramatically stronger in prime-age working men. Several studies have replicated these methods, some finding support for Clark's conclusion (Shields et al., 2005, Powdthavee, 2007, Shields et al., 2009, Clark et al., 2009, Flint et al., 2013, Howley and Kight, 2021), while others do not (Frijters et al., 2006, Oesch and Lipps, 2013). Clark has continued to refine his method and presents more and more convincing arguments. Howley and Kight (2021) is the most recent, and in my opinion, most convincing paper on the subject due to their robust methodology and longer-term longitudinal data set.

The unfortunate reality is that this relationship may vary over time and cir-

cumstances. It is unclear as to whether national or local employment rates offer substantially different information either as a reemployment signal or a self-blame signal. There is a strong possibility that an individual that was laid off in 2008 with high local and national unemployment might face unexpected and hard-to-measure differences compared to an unemployed factory worker in a high unemployment area with low national unemployment. While it seems clear that there is an element of comparison, it is less clear who exactly people compare themselves to. Overall, it does seem likely that there is a positive correlation between both local and national unemployment and the SWB of both the unemployed and employed, as well as the gap between the two.

Outside of using the unemployment rate as a proxy for social norms, other methods have been suggested. Within the papers that utilize these methods exists by far the most convincing evidence I have seen for the often-cited “Protestant work ethic.” Van Hoorn and Maseland (2013) argue that if the Protestant work ethic exists that it should reveal itself in a larger unemployment-related SBW loss in historically Protestant countries. They use the World Values Survey and European Value Survey to perform a cross-country analysis. They find that both Protestants and people unemployed in Protestant countries experience more unemployment-related loss to SWB.

Another method uses panel data to observe what happens when individuals who were unemployed retire. This line of thinking considers the possibility that unemployment-related shame and identity issues could be immediately alleviated by a shift in identity. Hetschko et al. (2014) find that there is a substantial increase in SWB for this group. This work has complications, but I believe it offers the best explanation for unemployed workers who choose to leave the labor force. The tendency for the labor force participation rate to fall when unemployment is high is a broad problem for this paper but I believe this research does address some of the concerns.

A final explanation argues that there is no innate social norm to work but a social norm not to live off others. The intuition here is that the SWB loss does not come from the stigma surrounding unemployment but rather from receiving unemployment benefits. Stutzer and Lalive (2004) find a correlation between a larger average unemployment-related SWB loss and the percentage of voters in favor of a reduction in unemployment benefits in Sweden. Chadi (2014) and Hetschko et al. (2020) both use German data to analyze movements between unemployment, subsidized employment, and normal employment. They find that moving from unemployment to subsidized employment does increase SWB, but not as much as moving to normal employment. This leads to the conclusion that the stigma surrounding unemployment is probably multi-faceted and includes stigma around working, receiving transfers, and other possible factors surrounding status.

2.4 Within-Household Spill Over Effects

A final area of research that will be important in contextualizing the findings of this paper is that which looks into the spillover effects of unemployment on those directly around the unemployed individual. It would stand to reason that if an individual was to become unemployed it would affect not only their SWB but also their partner's well-being. Winkelmann and Winkelmann (1995) find that an unemployed husband does reduce the SWB of the wife, but not vice versa. Marcus (2013) is probably the most sophisticated and clear look into the subject. He utilizes matching and a difference in difference framework to replicate the earlier result, although he found a significantly stronger direct effect on men than the indirect effect on women. Bubonya et al. (2017) and Nikolova et al. (2019) also replicate this finding using different surveys.

3 Methods

3.1 Data

This paper utilizes two primary data sets. The first is the NHIS Survey which looks at a variety of health and demographic factors, is conducted annually, and contains relevant data from 2000 to 2019. This first data set is used as a reference point to further motivate a deeper analysis of the second data set.

One primary challenge in this analysis is that the NHIS Survey could not be conducted during the COVID-19 pandemic due to its use of personal home interviews, with this method dating back around fifty years. The absence of this health-related survey data right when it might have been most valuable for policymakers presents a substantial problem. To fill the gap, a wide array of federal agencies collaborated to create the Household PULSE Survey, which will be the primary focus of this analysis. This survey does present some challenges, the first of which is that the format was changed dramatically from the NHIS model to a twenty-minute online survey.

The Census Bureau refers to this data set as experimental and states that it will most likely be discontinued sometime in 2022. The data set does have a somewhat smaller sampling size than the NHIS survey but appears to be an unbiased and reputable source for economic analysis. Due to its updated purpose, the survey also contains questions surrounding COVID-19 that other less agile surveys have failed to include. The largest and most relevant factor is that the survey is collected every few weeks, allowing for a “play-by-play” of conditions throughout the pandemic. It was collected in multiple phases, each containing its own weak variables that represent the order. From April 23rd, 2020 to February 7th, 2022, forty-three waves of data were collected. This data presents a fantastic opportunity to study the relationships previously discussed in this paper.

There are two somewhat important complications. The first is that the survey

does change slightly over time, notably altering the question relating to the cause of unemployment. This creates a new variable for the second half of the surveys that needs to be merged with the variable associated with the first half. Another complication is the lack of survey time consistency. The surveys were collected in waves, with each survey window lasting between one and two weeks, and the time in between separate surveys varying in the same way. This makes more advanced time series analysis and lining the survey up with other indicators, such as unemployment, very challenging.

The independent variable used in this analysis is an average of the variables ANXIOUS, WORRY, DOWN, and INTEREST, all of which are commonly used to evaluate PWB and SWB. The questions all follow a similar tack: “Over the last two weeks, how often have you been bothered by having little interest or pleasure in doing things?” to which the respondent gives an answer ranging from one, which is not at all, to four, which is every day. The average, which I have given the label Well-being (WB), is then transformed to a zero to one scale, where one is the best possible well-being, and zero is the worst. For the measure of unemployment, it is impossible to join a separate data set containing more precise information from the BLS. Instead, both the aggregate and state unemployment rate is calculated by dividing the number of unemployed by the number of employed for each time period and state. While this is less accurate, it seems to give a very close approximation of the unemployment rate calculated elsewhere.

It should also be noted that while there is a colloquial difference between the terms being fired, being laid off, and being furloughed, this paper generally makes no distinction between them due to the lack of a “how did you become unemployed?” question in the data I’ll be using. The PULSE survey assumes that all unemployment during COVID-19 was caused by COVID-19. This seems unlikely, but it forces the exclusion of an examination of the differential effects of the cause of unemployment. Another theoretical problem with this data set is the possibility that a large number of

people were claiming unemployment but should have been better classified as retired or out of the labor force. Given the estimated scale of UI fraud and the spike in retirement during COVID-19 (Castro 2020), this does seem like a real possibility. This may be corrected in the future, but it may just be a continual uncertainty for all research studying employment during this time period.

3.2 Theoretical Framework and Regression

The theoretical framework most relevant both to this paper and to this topic is provided in the Blanchflower and Oswald (2004) paper which also uses cross-sectional data in the United States. They use the following framework to better conceptualize the relationship between reported well-being and utility:

$$r = h(u(y, z, t)) + e$$

Where r is the reported well-being. h is a function relating actual to reported well-being to u which is a measure of a person's true utility, or true well-being. u itself is a function of y , real income, z , which represents a vector of personal characteristics, such as unemployment and gender, and t , which is the time period. e is the error that arises in the individual challenges humans face when asked to rate their own happiness, for example, two people who are equally happy might score themselves differently. Blanchflower and Oswald also use the framework to describe self-reported well-being more broadly as a function of circumstances, aspirations, comparisons with others, and a person's baseline happiness. They are particularly interested in the way that the time period affects r .

Grafting this framework onto the premise of this paper, we might think of the way social norms might affect the value one ascribes to different things in their life. If a number such as the unemployment rate is representative of either social norms or points of comparison, then it should impact the way in which z affects r . Taken to an extreme, the way in which every personal characteristic impacts individual well-

being, either positively or negatively, could be determined by the way they are valued by those surrounding the individual. Even the value a culture places on income could affect the individual motivation to become entrepreneurial or wealthy. The economist Richard Easterlin consistently argues that even the effect that income has on utility and well-being is heavily moderated by social comparisons. He observes the fact that happiness does not trend upwards over time even as incomes rise, giving us the still debated Easterlin Paradox. Another way to interpret this framework is in its similarity to the idea of experienced utility put forth by Kahneman et al. (1997), which is still a foundation for work in behavioral economics. The concept of true utility in this context is also a mirror for the concept of true preferences, which is still the subject of active debate.

3.3 Empirical Specification

Utilizing this framework, I will test the premise of the paper through three separate regressions. The first is a regression of well-being (WB) on individual employment status (E), the national unemployment rate (NUR), gender (G), and age (Age) as well as a three-way interaction term between the first three variables.

$$\begin{aligned}
 WB &= \beta_0 + \beta_1 E + \beta_2 NUR + \beta_3 G + \beta_4 Age + \beta_5 (NUR * E) \\
 &\Rightarrow +\beta_6 (NUR * G) + \beta_7 (E * G) + \beta_8 (NUR * E * G) + \beta_9 (E * Age) + e
 \end{aligned}
 \tag{1}$$

The second regression will replicate the results of the first using the data set from the COVID-19 period.

$$\begin{aligned}
WB &= \beta_0 + \beta_1 E_{\text{COVID}} + \beta_2 NUR_{\text{COVID}} + \beta_3 G_{\text{COVID}} + \beta_4 (NUR_{\text{COVID}} * E_{\text{COVID}}) \\
&\Rightarrow +\beta_5 (NUR_{\text{COVID}} * G_{\text{COVID}}) + \beta_6 (E_{\text{COVID}} * G_{\text{COVID}}) \\
&\Rightarrow +\beta_7 (NUR_{\text{COVID}} * E_{\text{COVID}} * G_{\text{COVID}}) + \beta_8 Age_{\text{COVID}} + \beta_9 (E_{\text{COVID}} * Age_{\text{COVID}}) + e
\end{aligned}
\tag{2}$$

The final regression will run the same test however this time with the unemployment of the state the individual resides (SUR) in rather than the national unemployment rate.

$$\begin{aligned}
WB &= \beta_0 + \beta_1 E_{\text{COVID}} + \beta_2 SUR_{\text{COVID}} + \beta_3 G_{\text{COVID}} + \beta_4 (SUR_{\text{COVID}} * E_{\text{COVID}}) \\
&\Rightarrow +\beta_5 (SUR_{\text{COVID}} * G_{\text{COVID}}) + \beta_6 (E_{\text{COVID}} * G_{\text{COVID}}) \\
&\Rightarrow +\beta_7 (SUR_{\text{COVID}} * E_{\text{COVID}} * G_{\text{COVID}}) + \beta_8 Age_{\text{COVID}} + \beta_9 (E_{\text{COVID}} * Age_{\text{COVID}}) + e
\end{aligned}
\tag{3}$$

Using this specification I'd like to test four hypotheses, one that are fairly proven and three that are not.

H_1 : WB will be lower for the unemployed than it is for the employed

H_2 : The the effect of unemployment on WB will be lower when the unemployment rate is higher

H_3 : The moderating effect of the unemployment rate will be stronger for men than for women

H_4 : More localized employment data will improve the fit of the model

The first hypothesis is a reflection of the established relationship between unemployment and reduced WB . The third is a test of the less established relationship

between the unemployment rate and the magnitude reduction in WB. If these hypotheses are valid, then I would expect B_1 will be negative and B_3 will be positive. H_3 tests the intuition that if the unemployment rate does moderate unemployment related WB loss then it will have a stronger effect on men due to gender norms. H_4 tests the intuition that individuals will tend see use those in close proximity more often so a more localized employment rate will be a better predictor of WB .

4 Results

The results are broadly split into two sections. The first looks at historical patterns in the relationship between the unemployment rate and subjective well-being loss. The second look at the same relationship after the pandemic and comments on the key differences. The analysis in both parts will start with an exploration of the data through a series of charts and figures that look at aggregates both in relation two each other and then over time. Then it will move to the individual level regressions for both time periods which are the primary focus of the paper.

4.1 Pre-COVID-19

The primary goal of this paper is to test if the national unemployment rate has an effect on the relationship between SWB and unemployment. The clearest way to visualize the relationship is to compose an aggregated well-being that represents the mean for both employed and unemployed, as well as the unemployment rate for each time period. The first and most obvious way to show the simple relationship is to plot $WB_{\text{Difference}}$, which is equal to $WB_{\text{Employed}} - WB_{\text{Unemployed}}$, against the national unemployment rate.

Figure 1




Figure 1 shows a negative correlation between the two variables indicating a higher unemployment rate is correlated with a reduction in the difference between WB_{Employed} and $WB_{\text{Unemployed}}$. This in itself presents interesting evidence that the difference between the two does have a strong underlying relationship, but it doesn't tell us where the difference is coming from. Figure 2 shows the well-being of the unemployed and the employed separately.

Figure 2

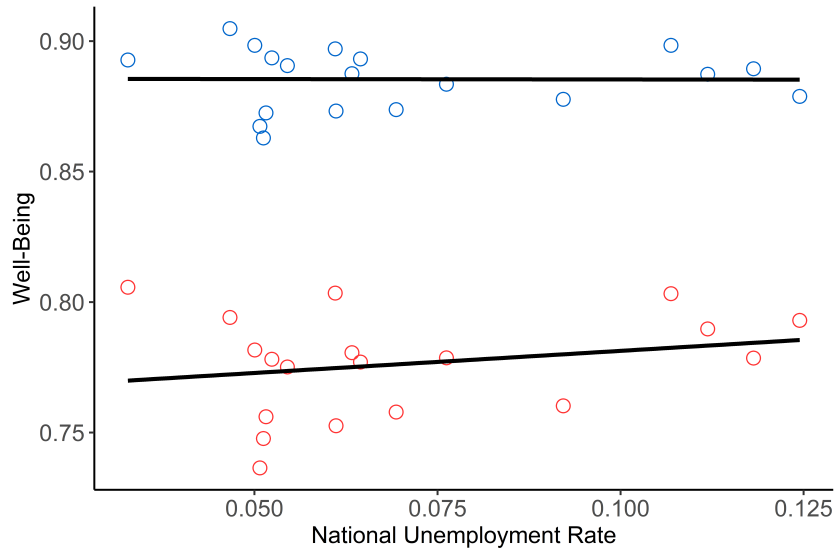


Figure 2 indicates that the well-being of the employed is not correlated with the unemployment rate in this pre-COVID-19 period, so the correlation between $WB_{\text{Difference}}$ and the unemployment rate comes mostly from changes in $WB_{\text{Unemployed}}$. Figure 3 and Figure 4 both show the time-series relationship between $WB_{\text{Difference}}$ and the unemployment rate, and WB_{Employed} and $WB_{\text{Unemployed}}$.

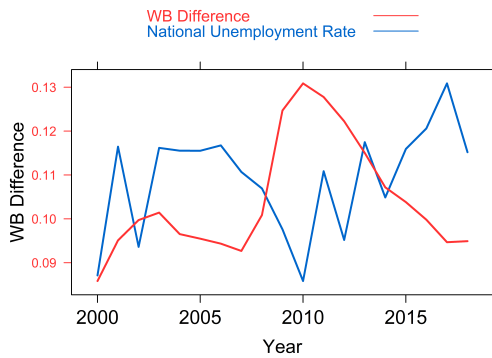


Figure 3

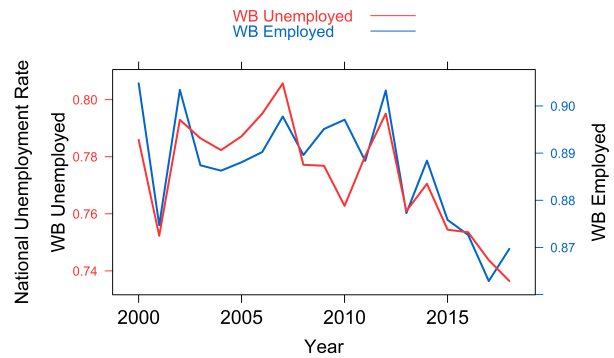


Figure 4

Figure 4 underlines the observation that WB_{Employed} and $WB_{\text{Unemployed}}$ are mostly correlated with each other, suggesting exogenous factors ranging from political events to the zeitgeist probably have an consistent effect on both WB_{Employed} and $WB_{\text{Unemployed}}$. It's very important that the two are so closely correlated, as it shows that events or changes over time have a very similar impact on both the unemployed and the employed. Despite this volatility, the difference between the two remains strongly correlated with unemployment.

Moving on to the regression, the strategy is to take the national employment for the time series and merge it into the individual data set that was used to create the aggregated numbers seen above. The most important test is conducted by measuring the interaction between the national unemployment rate and the dummy for individual unemployment (NUR times Unemployed).

Table 1: Pre COVID-19 Results

	<i>Dependent variable:</i>				
	Well-Being				
	(1)	(2)	(3)	(4)	(5)
NUR		0.002 (0.012)	0.004 (0.012)	0.015 (0.017)	0.014 (0.017)
Unemployed	-0.107*** (0.001)	-0.126*** (0.004)	-0.130*** (0.004)	-0.134*** (0.005)	-0.087*** (0.006)
Age				0.001*** (0.00002)	0.001*** (0.00002)
Gender			0.033*** (0.001)	0.036*** (0.002)	0.036*** (0.002)
Gender:NUR				-0.044* (0.025)	-0.044* (0.025)
Gender:Unemployed			0.011*** (0.003)	0.028*** (0.008)	0.028*** (0.008)
NUR:Unemployed		0.228*** (0.045)	0.220*** (0.045)	0.317*** (0.062)	0.336*** (0.062)
Unemployed:AGE					-0.001*** (0.0001)
Gender:NUR:Unemployed				-0.213** (0.090)	-0.209** (0.090)
Constant	0.886*** (0.0003)	0.885*** (0.001)	0.869*** (0.001)	0.831*** (0.002)	0.827*** (0.002)
Observations	340,586	340,586	340,586	340,586	340,586
R ²	0.021	0.021	0.029	0.034	0.034
Adjusted R ²	0.021	0.021	0.029	0.034	0.034

This regression validates the first, second, and third hypotheses. The coefficient on unemployment is negative and statistically significant, which confirms past research that points to a causal relationship between unemployment and a reduction in WB , and acts as a validity check for the use of this data set and method. The interaction between NUR and Unemployed is positive and statistically significant, suggesting that, as unemployment rises, the WB cost of unemployment falls. This is a numeric reflection of Figure 1.

The final test we can conduct is a triple interaction between Gender, National Unemployment Rate, and Unemployed. This effect is also negative and statistically significant. We can investigate this further by breaking down the interaction between the National Unemployment Rate and the Unemployed variable by gender.

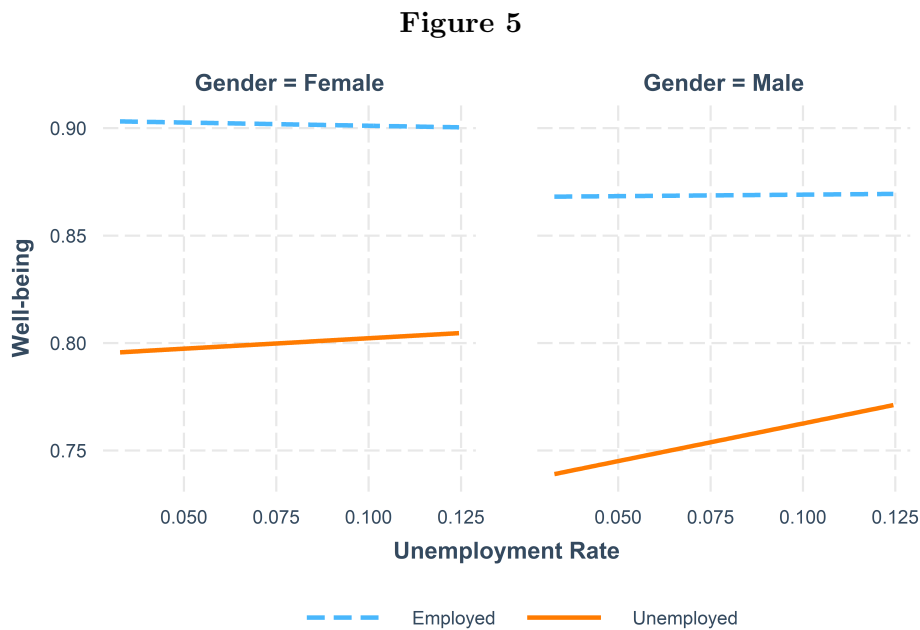



Figure 5 and the interaction terms suggest that the alleviating effect of a higher unemployment rate is stronger for men than it is for women. This adds credence to the intuition that negative feelings and loss of self-worth may be stronger than a loss of nonmonetary employment benefits. If time structure was an important part

of unemployment-related SWB loss, then this interaction would be hard to explain.

4.2 COVID-19

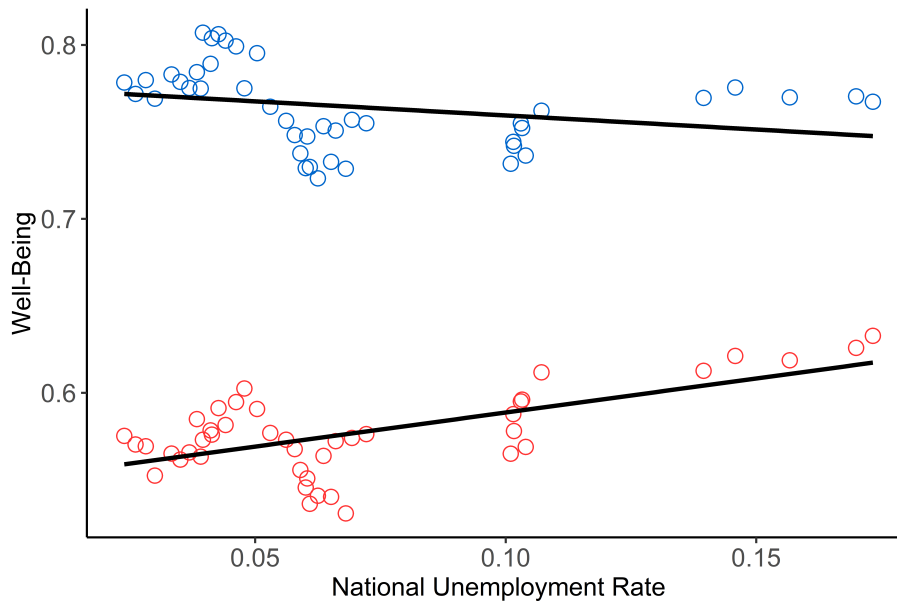
Now I can replicate the method used for the Pre-COVID-19 data set and compare the relationships and see if they still hold. Once again, we'll start with the aggregates for each time period in our sample.

Figure 6



Here, we can once again observe the negative relationship. The correlation does appear to be somewhat stronger than it was before the pandemic, although it is important to note that the difference could be due to having access to far more granular data. We can also look at this same relationship separated by employment status in Figure 7. One of the primary differences between pre-COVID-19 and during COVID-19, which we will explore in more detail later, is that this time we do observe a downward sloping line for the employed. This implies that during COVID-19, employed individuals are better off when the unemployment rate is lower.

Figure 7



We can also see that this data is somewhat grouped into a few separate clumps. It would seem like these clumps might correspond with the different phases of the PULSE data set. This is somewhat true, but may be deceptive. The different phases did not have a substantial amount of time in between them, raising questions as to why there is such a strong group around the 10% unemployment rate. Looking at Figure 8, we can see that the case is a plateau in the fall in the unemployment rate between weeks five and twelve. It is unclear if this is a problem in this data set or the result of comparing weekly data to monthly data.

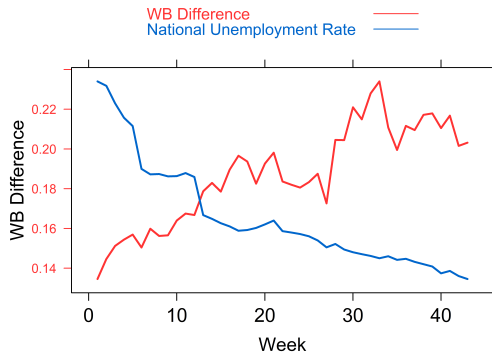


Figure 8

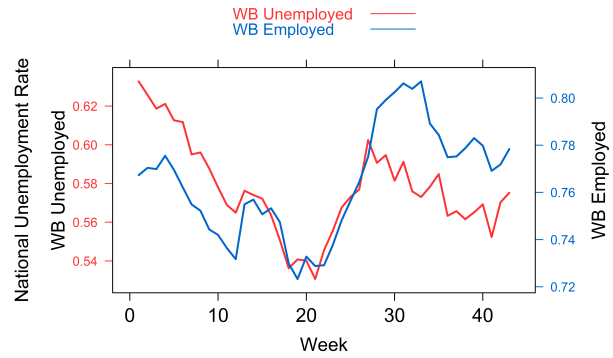


Figure 9

Looking at Figure 9, we can once again see that WB_{Employed} and $WB_{\text{Unemployed}}$ are highly correlated with each other. This once again leads to the conclusion that there are many exogenous factors that influence general well-being. Due to the granular nature of this data, there are some additional complications. For example, there is little to no evidence on whether SWB is in any way seasonal. We can also speculate on the cause of the sharp decline in SWB around the twentieth week. This roughly corresponds with the 2020 presidential election. Obviously, it would be hard to draw a causal link between the two, but it may warrant further research.

Now we can move on to the regression. This time we can observe the same negative coefficient on Unemployment as well as the positive coefficient on the interaction between the Unemployment and National Unemployment Rate variables. From this, we can once again infer that unemployment is costly from a well-being perspective. The interaction term NUR and Unemployed also suggests that the alleviating effects of high unemployment were consistent with pre-pandemic findings.

Table 2: COVID-19 NUR Results

<i>Dependent variable:</i>					
Well-Being					
	(1)	(2)	(3)	(4)	(5)
NUR		-0.061*** (0.004)	-0.063*** (0.004)	-0.053*** (0.005)	-0.053*** (0.005)
Unemployed	-0.184*** (0.001)	-0.235*** (0.001)	-0.220*** (0.001)	-0.225*** (0.002)	-0.249*** (0.002)
Gender			0.066*** (0.0003)	0.063*** (0.001)	0.064*** (0.001)
Age				0.003*** (0.00001)	0.003*** (0.00001)
Gender:NUR				-0.012 (0.008)	-0.012 (0.008)
Gender:Unemployed			-0.038*** (0.001)	-0.041*** (0.003)	-0.041*** (0.003)
NUR:Unemployed		0.550*** (0.013)	0.562*** (0.012)	0.579*** (0.016)	0.584*** (0.016)
Unemployed:Age					0.0005*** (0.00004)
Gender:NUR:Unemployed				0.008 (0.025)	0.004 (0.025)
Constant	0.771*** (0.0002)	0.775*** (0.0004)	0.748*** (0.0004)	0.615*** (0.001)	0.617*** (0.001)
Observations	2,173,614	2,173,614	2,173,614	2,173,614	2,173,614
R ²	0.050	0.051	0.068	0.092	0.092
Adjusted R ²	0.050	0.051	0.068	0.092	0.092

The primary coefficient of interest, NUR:Unemployed, suggests the difference between the Well-Being of an employed and unemployed person is around 21% when the unemployment rate is high (4%) and around 17% when the unemployment rate is high (12%).

This data set also gives access to information on an individual's state. This allows us to perform the same regression as above, but now with the State Unemployment Rate rather than National Unemployment Rate. This time, the model does have a better fit and the interaction between gender, the unemployment rate, and individual unemployment are no longer statistically significant. This suggests that using a more localized measure of unemployment does yield slightly stronger results. This result does support the fourth hypothesis however it is clear that both NUR and SUR can play an explanatory role.

Table 3: COVID-19 SUR Results

<i>Dependent variable:</i>					
Well-Being					
	(1)	(2)	(3)	(4)	(5)
SUR		-0.087*** (0.003)	-0.091*** (0.003)	-0.091*** (0.003)	-0.091*** (0.003)
Unemployed	-0.184*** (0.001)	-0.222*** (0.001)	-0.206*** (0.001)	-0.211*** (0.001)	-0.233*** (0.002)
Gender			0.066*** (0.0003)	0.064*** (0.001)	0.064*** (0.001)
Age				0.003*** (0.00001)	0.003*** (0.00001)
Gender: SUR				-0.012** (0.005)	-0.011** (0.005)
Gender: Unemployed			-0.039*** (0.001)	-0.043*** (0.002)	-0.044*** (0.002)
SUR: Unemployed		0.284*** (0.007)	0.291*** (0.007)	0.298*** (0.009)	0.300*** (0.009)
Unemployed: Age					0.0004*** (0.00004)
Gender: SUR: Unemployed				0.019 (0.014)	0.017 (0.014)
Constant	0.771*** (0.0002)	0.780*** (0.0003)	0.753*** (0.0003)	0.620*** (0.001)	0.623*** (0.001)
Observations	2,173,614	2,173,614	2,173,614	2,173,614	2,173,614
R ²	0.050	0.051	0.069	0.092	0.092
Adjusted R ²	0.050	0.051	0.069	0.092	0.092

One of the primary differences between the pre-COVID-19 and COVID-19 periods is the statistically significant negative coefficient on the Unemployment Rate. This implies that in the COVID-19 period the employed were actually made worse off during times of high unemployment. This finding is actually more in line with prior research. In particular, Howley and Kight (2021) find a similar relationship, although they offer no strong explanation. It does present a problem if the cost of unemployment goes down as unemployment rises, in which case we might expect the benefit of unemployment to rise rather than fall. In my view, there are two possible explanations. The first is that spillover effects play a large role. In times of high unemployment, an employed individual is more likely to have an unemployed spouse which in turn would reduce their own SWB. Looking at prior research as well as Figure 10, it seems as though this explanation would probably only make sense if the effect of the unemployment rate on the employed was stronger for women than it is for men. This finding would mirror Howley and Kight (2021), who found that employed men were completely unaffected by the unemployment rate, while both employed and unemployed women were negatively affected, with the effect being stronger for the former. Once again I can visualize the three-way interaction term and we can look at differences in the relationship by gender.

Figure 10



Contrary to prior research, this paper finds both men and women are affected by the unemployment rate in similar ways. This points to a second, and in my view more likely, explanation. This interaction is due to exogenous forces specific to the COVID-19 pandemic and the associated lockdown. In particular, this model has a very possible and unique empirical threat. It is very likely that the lockdown itself had a strong effect both on the unemployment rate and on SWB. Recent research suggests that lockdown has had a strong effect on general well-being and mental health (Kien Le My Nguyen, 2021). However, there is not yet clear evidence of this effect being different for different groups, such as gender or employment. The lockdown itself may have had a strong negative effect on the SWB of the employed. Anecdotally, many individuals I know in customer-facing roles said that their job was made significantly harder due to their newfound responsibilities such as enforcing mask and vaccine requirements. In 2021, the FAA reported a record number of unruly passenger complaints around 70% of which related to mask requirements.

This intuition might play a strong role in the downward sloping curve for the employed, but it is hard to imagine any substantial way in which the lockdown made the experiences of the unemployed better. I would expect many of the unique facets of the lockdown, such as isolation, would affect the employed and the unemployed roughly the same. The only explanation that comes to mind is that the enhanced unemployment benefits and stimulus checks during the beginning of the pandemic had a strong effect on SWB. If this was the case, we might expect to observe a spike in SWB for both the employed and the unemployed around the disbursement of the checks. We might also expect a sharp drop in the SWB of the unemployed immediately following the end of the enhanced unemployment insurance. Looking at Figure 9, neither of these events seem to be visible in the data. In this time series, we can also see that the relationship continued after all aid had ended as well. The strong correlation between the SWB of the employed and unemployed implies that exogenous shocks affect the SWB of the unemployed and the employed equally. In general, the unique circumstances of the COVID-19 pandemic and associated lockdown do present a very real threat to the empirical model. With that being said, the data does strongly suggest that a higher unemployment rate has an alleviating effect on unemployment-related SWB loss.

4.3 Further Analysis

These are several things that could have strengthened this analysis, and several unanswered questions that this data could help answer. The best way this analysis could be improved is with a more complete list of control variables. In many ways, this paper serves more as a preliminary analysis, but the data contains many other variables which might be of further interest such as age and race.

The data also has other information that could be used to conduct a more precise analysis such as “recently unemployed,” and “unemployment insurance” variables. More research would look into the differences between the recently unemployed

and those who had been unemployed for a longer period. It also would also examine the covariance between unemployment insurance and individual employment as a check. At a later date when more data is released from the longitudinal data mentioned earlier in this paper, it would also be prudent to replicate past methodologies and perform a test that directly compares the pre-COVID-19 and COVID-19 periods. During this analysis, I generally avoided a direct comparison, as the data sets did vary, particularly in their frequencies and collection method. This paper instead focused on how this COVID-19 data-set suggests a continuation of past results and used the pre-COVID-19 results as a way to ground the analysis. In the future, a longitudinal study will provide a more robust answer to the questions posed here.

Another interesting avenue to explore would be to look into a difference in difference analysis at the aggregate state level, rather than the individual level. This might allow a more complete look at the relationship between the change in the state-level unemployment rate, and the change in the difference in SWB between the employed and the unemployed over time. As an aside, I'd also be very interested to see if the way a state voted in the 2020 election had an impact on the SWB of its constituents. It would also be interesting to see if a state's political affiliation has an effect on the SWB loss associated with unemployment. This line of research might support the findings of Stutzer and Lalive (2004) that showed a connection between the percentage of voters in favor of welfare reform and unemployment-related SWB loss.

5 Discussion

This paper tests the relationship between both the national and state-level unemployment rate and the SWB loss associated with unemployment throughout the COVID-19 pandemic, the primary hypothesis is that the unemployment rate and a reduced social norm to work reduced the non-monetary costs of unemployment.

This paper acts as a preliminary exploration into the relationship between social norms, unemployment, and SWB. The main finding was that a higher unemployment rate seems to continually alleviate the emotional suffering caused by unemployment. This supports the hypothesis that the reduction in self-blame caused by comparison outweighs the added anxiety of being laid off in a weak job market. In my view, it also points to a broader explanation that the loss in SWB is primarily due to internal feelings of shame and loss of status, rather than external factors such as changes in social interaction and time structure. It should be noted that the cross-sectional nature of this paper leaves many questions unanswered, and should ideally be followed up with a longitudinal expiration when data is made available.

This research has two broad implications. The first is the effect of the aforementioned relationship and its impact on the labor markets. In my view, work on subjective well-being can help tie together different areas of economic research and offer empirical evidence to some unanswered questions. In the 2016 paper “Where Have All the Workers Gone?,” Kruger (2016) examines the secular decline in labor force participation and utilizes SWB data to re-enforce his claims. He points out that individuals who dropped out of the labor force appear to be doing better than individuals who remain unemployed. This is consistent with the evidence presented here. If dropping out of the labor force alleviates some of the negative emotions associated with unemployment, this might add an incentive to do so. He then goes on to show the causal relationship between the rise of the opiate epidemic and the decline in labor force participation. However, he does not explicitly state that there could be a relationship between the two. It could be the case that a decline in well-paying jobs that would traditionally employ high school-educated men could have led to worse mental health, which in turn could have had causal effect on rates of opiate addiction. Another feature he points out is men’s apparent resistance to transitioning to traditionally female-dominated professions, such as nursing. This points to an underlying concept: when determining the desirability of a job, an individual is rarely only

looking at hours and salary. Harder to measure factors such as how dignified or how compatible the job is with one's identity are likely of key importance.

These insights should be used when evaluating labor market policies. As has been pointed out by Kruger, the decades-long decline in the labor force participation of prime-age workers should be a real policy concern. However, I believe the results presented here and in related literature suggest that a large portion of the response must include changes in how society views unemployment. While policies such as adult reeducation and opiate addiction programs will likely be very useful, it seems unlikely that they will be successful at alleviating the social stigma placed on those without work. In my view, these programs need to be coupled with a general reduction in the stigma associated with periods of unemployment. One possible solution could be a policy aimed at making it illegal for corporations to discriminate against employees of an employment gap. So far the effects of government mandated employment programs on SWB are unclear. However it seems likely that engaging unemployed workers in meaningful, productive and dignified work could go a long way in alleviating the identify issues associated with unemployment.

Another more narrow implication of this work is that it could explain some of the lag between job growth and employment growth experienced through the COVID-19 pandemic. Particularly in the first quarter of 2021, there were a handful of months where it was expected that large gains in new job openings would immediately translate into strong employment growth. In April 2021, employment growth was around 75% of market estimates (Cox 2021). While I haven't found any formal explanations as to why market expectations were so off, it seems the consensus at the time was that the labor market was far tighter than expected, especially in the face of a relatively high employment rate. In my view, this was probably caused by an increase in the relative value of non-employment. Federal aid and stimulus payments likely had a large impact on this prospective holdout, however, I believe this research suggests that a reduction in the social norm to work also played a role.

References

- [1] Netta Achdut and Tehila Refaeli. “Unemployment and psychological distress among young people during the COVID-19 pandemic: psychological resources and risk factors”. In: *International journal of environmental research and public health* 17.19 (2020), p. 7163.
- [2] George A Akerlof. “A theory of social custom, of which unemployment may be one consequence”. In: *The quarterly journal of economics* 94.4 (1980), pp. 749–775.
- [3] David Autor. *Work of the Past, Work of the Future*. Tech. rep. National Bureau of Economic Research, 2019.
- [4] Jan Michael Bauer et al. “Effects of different life events on life satisfaction in the Russian Longitudinal Monitoring Survey”. In: *Economics Letters* 129 (2015), pp. 91–94.
- [5] David G Blanchflower and Alex Bryson. *Biden, Covid and mental health in America*. Tech. rep. National Bureau of Economic Research, 2021.
- [6] David G Blanchflower and Andrew J Oswald. “Well-being over time in Britain and the USA”. In: *Journal of public economics* 88.7-8 (2004), pp. 1359–1386.
- [7] Franco Bonomi Bezzo, Laura Silva, and Maarten Van Ham. “The combined effect of Covid-19 and neighbourhood deprivation on two dimensions of subjective well-being: Empirical evidence from England”. In: *PloS one* 16.7 (2021), e0255156.
- [8] Melisa Bubonya, Deborah A Cobb-Clark, and Mark Wooden. “Job loss and the mental health of spouses and adolescent children”. In: *IZA Journal of Labor Economics* 6.1 (2017), pp. 1–27.
- [9] Adrian Chadi. “Regional unemployment and norm-induced effects on life satisfaction”. In: *Empirical Economics* 46.3 (2014), pp. 1111–1141.

- [10] Fang Fang Chen et al. “Two concepts or two approaches? A bifactor analysis of psychological and subjective well-being”. In: *Journal of happiness studies* 14.3 (2013), pp. 1033–1068.
- [11] Andrew Clark, Andreas Knabe, and Steffen Rätzel. “Boon or bane? Others’ unemployment, well-being and job insecurity”. In: *Labour Economics* 17.1 (2010), pp. 52–61.
- [12] Andrew E Clark. “Unemployment as a social norm: Psychological evidence from panel data”. In: *Journal of labor economics* 21.2 (2003), pp. 323–351.
- [13] Andrew E Clark and Yannis Georgellis. “Back to baseline in Britain: adaptation in the British household panel survey”. In: *Economica* 80.319 (2013), pp. 496–512.
- [14] Andrew E Clark and Andrew J Oswald. “Satisfaction and comparison income”. In: *Journal of public economics* 61.3 (1996), pp. 359–381.
- [15] Richard M Cohn. “The effect of employment status change on self-attitudes”. In: *Social psychology* (1978), pp. 81–93.
- [16] Jane RW Fisher et al. “Mental health of people in Australia in the first month of COVID-19 restrictions: a national survey”. In: *Medical journal of Australia* 213.10 (2020), pp. 458–464.
- [17] Ellen Flint et al. “Do local unemployment rates modify the effect of individual labour market status on psychological distress?” In: *Health & place* 23 (2013), pp. 1–8.
- [18] Roberto Foa, Sam Gilbert, and Mark Otto Fabian. “COVID-19 and subjective well-being: Separating the effects of lockdowns from the pandemic”. In: *Available at SSRN 3674080* (2020).

- [19] Paul Frijters, John P Haisken-DeNew, and Michael A Shields. “Money does matter! Evidence from increasing real income and life satisfaction in East Germany following reunification”. In: *American Economic Review* 94.3 (2004), pp. 730–740.
- [20] Paul Frijters et al. “Can the large swings in Russian life satisfaction be explained by ups and downs in real incomes?” In: *scandinavian Journal of Economics* 108.3 (2006), pp. 433–458.
- [21] Knut Gerlach and Gesine Stephan. “A paper on unhappiness and unemployment in Germany”. In: *Economics Letters* 52.3 (1996), pp. 325–330.
- [22] Robert Gibbons and Lawrence F Katz. “Layoffs and lemons”. In: *Journal of labor Economics* 9.4 (1991), pp. 351–380.
- [23] Clemens Hetschko, Andreas Knabe, and Ronnie Schöb. “Changing identity: Retiring from unemployment”. In: *The Economic Journal* 124.575 (2014), pp. 149–166.
- [24] Clemens Hetschko, Ronnie Schöb, and Tobias Wolf. “Income support, employment transitions and well-being”. In: *Labour economics* 66 (2020), p. 101887.
- [25] Peter Howley and Sarah Knight. “Staying down with the Joneses: Differences in the psychological cost of unemployment across neighbourhoods”. In: *Work, Employment and Society* (2021), p. 09500170211003483.
- [26] Simon Jäger et al. “Wages and the Value of Nonemployment”. In: *The Quarterly Journal of Economics* 135.4 (2020), pp. 1905–1963.
- [27] Marie Jahoda. “Work, employment, and unemployment: Values, theories, and approaches in social research.” In: *American psychologist* 36.2 (1981), p. 184.
- [28] Daniel Kahneman, Peter P Wakker, and Rakesh Sarin. “Back to Bentham? Explorations of experienced utility”. In: *The quarterly journal of economics* 112.2 (1997), pp. 375–406.

- [29] Wolfram Kawohl and Carlos Nordt. “COVID-19, unemployment, and suicide”. In: *The Lancet Psychiatry* 7.5 (2020), pp. 389–390.
- [30] Kien Le and My Nguyen. “The psychological consequences of COVID-19 lockdowns”. In: *International Review of Applied Economics* 35.2 (2021), pp. 147–163.
- [31] Jan Marcus. “The effect of unemployment on the mental health of spouses—Evidence from plant closures in Germany”. In: *Journal of health economics* 32.3 (2013), pp. 546–558.
- [32] Frances McKee-Ryan et al. “Psychological and physical well-being during unemployment: a meta-analytic study.” In: *Journal of applied psychology* 90.1 (2005), p. 53.
- [33] Katja Möhring et al. “The COVID-19 pandemic and subjective well-being: longitudinal evidence on satisfaction with work and family”. In: *European Societies* 23.sup1 (2021), S601–S617.
- [34] Emi Nakamura. “Layoffs and lemons over the business cycle”. In: *Economics Letters* 99.1 (2008), pp. 55–58.
- [35] Milena Nikolova and Sinem H Ayhan. “Your spouse is fired! How much do you care?” In: *Journal of Population Economics* 32.3 (2019), pp. 799–844.
- [36] Daniel Oesch and Oliver Lipps. “Does unemployment hurt less if there is more of it around? A panel analysis of life satisfaction in Germany and Switzerland”. In: *European Sociological Review* 29.5 (2013), pp. 955–967.
- [37] Nattavudh Powdthavee. “Are there geographical variations in the psychological cost of unemployment in South Africa?” In: *Social indicators research* 80.3 (2007), pp. 629–652.
- [38] Yue Qian and Sylvia Fuller. “COVID-19 and the gender employment gap among parents of young children”. In: *Canadian public policy* 46.S2 (2020), S89–S101.

- [39] Michael A Shields and Stephen Wheatley Price. “Exploring the economic and social determinants of psychological well-being and perceived social support in England”. In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 168.3 (2005), pp. 513–537.
- [40] Michael A Shields, Stephen Wheatley Price, and Mark Wooden. “Life satisfaction and the economic and social characteristics of neighbourhoods”. In: *Journal of Population Economics* 22.2 (2009), pp. 421–443.
- [41] Eva Sobotková et al. “Analysis of the current situation in international show jumping and assessment of the influence of the proportion of Thoroughbred in the pedigree, horse demographics and sport season on the performance of horses”. In: *Journal of Veterinary Behavior* 47 (2022), pp. 12–21.
- [42] Alois Stutzer and Rafael Lalive. “The role of social work norms in job searching and subjective well-being”. In: *Journal of the European Economic Association* 2.4 (2004), pp. 696–719.
- [43] Rafael Di Tella, Robert J MacCulloch, and Andrew J Oswald. “The macroeconomics of happiness”. In: *Review of Economics and Statistics* 85.4 (2003), pp. 809–827.
- [44] Ioannis Theodossiou. “The effects of low-pay and unemployment on psychological well-being: a logistic regression approach”. In: *Journal of health economics* 17.1 (1998), pp. 85–104.
- [45] J Blake Turner. “Economic context and the health effects of unemployment”. In: *Journal of health and social behavior* (1995), pp. 213–229.
- [46] André Van Hoorn and Robbert Maseland. “Does a Protestant work ethic exist? Evidence from the well-being effect of unemployment”. In: *Journal of Economic Behavior & Organization* 91 (2013), pp. 1–12.

- [47] Liliana Winkelmann and Rainer Winkelmann. “Happiness and unemployment: a panel data analysis for Germany”. In: *Applied Economics Quarterly* 41.4 (1995), pp. 293–307.
- [48] Liliana Winkelmann and Rainer Winkelmann. “Why are the unemployed so unhappy? Evidence from panel data”. In: *Economica* 65.257 (1998), pp. 1–15.