HEALTH ECONOMICS GROUP
Faculty of Medicine and Health
Norwich Medical School

Labour market importance of poor health in the Russian Federation

Yevgeniy Goryakin[1], Marc Suhrcke[1]

[1]Health Economics Group, Norwich Medical School,
University of East Anglia, UK

November 2013 HEG working paper 13-05

Labour market importance of poor heath in the Russian Federation[1]

Yevgeniy Goryakin

Marc Suhrcke

Abstract

This paper examines the economic consequences of poor health in the Russian Federation, a country with exceptionally adverse adult health outcomes. In both baseline OLS models, as well as in models with fixed effects, acute ill-health events in Russia generally have much weaker association with the logarithm of hours worked than with labour force participation. However, this analysis ignores people who reported zero hours worked. After including them into the analysis with the help of a two part model, the magnitude of the effect of several health events, especially more serious ones, on hours worked increases dramatically. In addition, people with poor self-assessed health living in rural areas are less likely to stop working, compared to people living in the cities, and, perhaps surprisingly, women are also less likely to stop working than men. While there is no conclusive explanation, it is potentially due to the existence of certain barriers that prevent people with poor health from withdrawing from the labour force easily in order to take care of their health. If this is the case then better social insurance protection mechanisms, including more comprehensive unemployment and health insurance, may be required to alleviate the economic burden of ill health in Russia.

[1] Comments and suggestions on earlier drafts from William Dow, Richard Scheffler, David Levine, Jenny Liu, Ha Nguyen, Marcello Morciano, Charlotte Davies, Angela Robinson and from seminar participants at the University of California, Berkeley, from participants at the Population Association of America conference in Detroit, USA, as well as at the International Economic Association in Beijing, China, are gratefully acknowledged.

November 2013 HEG working paper 13-05

Introduction

In the past two decades, Russia has experienced a comparatively radical transformation from a socialist to a market economy. While creating economic opportunities for a large number of people, the process of economic disruption associated with the transition has also entailed a heavy and widely documented social and human toll for the Russian population (UNDP, 2005). Compared with many other Eastern European and Former Soviet Union countries, Russia had the largest baseline real GDP per capita pre-transition, but suffered one of the greatest collapses of output. By 1998, its economy had shrank by more than 75% compared to the early 1990s (Stillman, 2006). Unemployment increased throughout the 1990s, starting to fall only around 2000. Notably, the Russian population also experienced large deteriorations in various health indicators. For example, life expectancy at birth decreased by six years between 1990 and 1994, then recovered from 1994 to 1998 by three years, before declining again in synchrony with the large financial crisis starting in 1998 (Stillman 2006).

For comparison, deaths from preventable causes in Britain were higher than in Russia in 1965, but steadily decreased thereafter (Andreev et al., 2003). Out of the group of countries with comparable levels of per capita incomes, until recently Russia had one of the highest male mortality rates, and even did worse than many significantly poorer countries (Suhrcke, 2007). For example, in 2000-2001, the probability of dying for Russian males between ages 15 and 60 was about 42 %, while this probability was about 10% in Japan, 14% in the US, 13 % in Germany, and 22% in Turkey (Suhrcke, 2007).

In contrast to most developing countries, this deterioration in health was mostly attributable to increases in non-communicable diseases and injuries[2] (Bloom and Canning, 2000; Suhrcke, 2007). It also appears that behavioral factors (such as increased consumption of alcohol (Shkolnikov et al., 2001)), rather than declining standards of medical care, were the principal drivers of these trends (Stillman, 2006). For example, worsening diet, increased rates of smoking, alcohol consumption and mental stress have all been implicated in the past (Bloom and Canning, 2000). As the increasing gender gap in life expectancy may be partly attributable to such negative coping behaviors as greater consumption of alcohol and smoking by men (McCartney et al., 2011; Weidner and Cain, 2003), the evidence of such gap in Russia also

[2] Although infectious diseases, and HIV/AIDS and tuberculosis in particular, have also been growing at an alarming rate in Russia.

supports the view that stress and behaviorally-related factors may have played an important role in declining health in Russia. Likewise, Brainerd and Cutler (2004) failed to find evidence that material deprivation could explain an increase in mortality rates in Russia in the 1990s, and concluded that the most significant predictors of mortality were alcohol consumption and stress (Brainerd and Cutler, 2005).

As a large burden of disease in Russia occurs among that part of the working age population between 40 and 55 years old (Stillman, 2006; Suhrcke, 2007), economic consequences of ill health might be considerable (Marquez et al., 2007). Yet, from a theoretical perspective, while poor health is expected to lead to a decline in productivity and therefore to lower wages, the predicted effect of health on labour force participation is ambiguous. This is because the substitution effect – individuals choosing to substitute work for leisure in response to lower hourly wages – would act against the income effect (Currie and Madrian, 1999).

The ambiguity of the association between poor health and hours worked may even be further reinforced by the following: while intuitively, one might expect more serious ill-health events (e.g., myocardial infarctions or strokes) to cause bigger reductions both in terms of the probability of labour force participation and in terms of hours worked, in reality, this may not be the case, in that more seriously sick people drop out of the labour force and report zero hours worked. If this is the case, then the need arises to appropriately adjust for the features of the data in any empirical analysis seeking to allow for extrapolation of the results to the whole sample, regardless of whether they report positive hours worked.

A further hypothesis relates to the heterogeneity in the effect of health on labour supply, specifically on labour force participation: some of the previous literature has noted that the poor may continue working despite having serious health problems (Benjamin et al., 2003), simply because they cannot afford to retire or dedicate their time to treatment rather than work. In this scenario, the economic costs of illness may be underestimated for people in poor socioeconomic circumstances. Therefore, one proposition worth examining is that wealthier people are more likely to drop out of the labour force in response to adverse health events.

A related hypothesis can be derived about it not being wealth per se, but rather access to appropriate medical care and social insurance mechanisms that make it easier for people to stop working in response to adverse health events. If this is the case, then people living in the cities may find it easier to stop working when they are ill. Also, one can expect that the potentially detrimental effect of poor health on labour force participation may be at its strongest near retirement age, when the benefit of working (i.e., continuing to receive the income stream necessary for living expenses) may be outweighed by the greater disutility of work due to illness, and therefore greater preference for leisure. In addition, older people may be more likely to reduce their labour supply if their productivity (and thus remuneration) is more likely to decrease in response to illness, compared to younger age groups. On the other hand, at younger ages, people may have to disregard their deteriorating health simply in order to financially sustain themselves and their families. Finally, one can expect that the effect of poor health on labour supply will be stronger for women then for men across the age distribution, both because the opportunity costs of not working are usually higher for men (who tend to earn more than women), and because men are often the main family breadwinners.

In this paper, we examine in detail one dimension of the economic consequences of the exceptionally poor adult health in Russia, i.e. its effects on two labour market outcomes- labour force participation and hours worked, using a rich set of micro data collected over 14 years. In the following section, we will describe the data and variables in more detail, together with the discussion of specific empirical strategies required for identification of the parameters of interest. Finally, we will present results and discuss them.

Data

In this paper, we use individual, household and community-level data from rounds 6-17 of Russia Longitudinal Monitoring Survey- Higher School of Economics (RLMS-HSE) dataset collected in 1995-2008 by the University of North Carolina Population Center. RLMS-HSE is a household-based survey, designed to be nationally representative. Data has been collected in a repeated survey of household dwelling units since 1992 (see table 1), although the first part of the survey, collected until 1995, is too different to be included in this analysis[3]. Multistage probability sampling was used to select survey participants. Geographically-determined strata were chosen first, followed by primary sampling units (PSU) further selected from the strata. Finally, within each PSU, the population was stratified for greater estimation efficiency, and the sample size was determined. The population was restricted to adults in their predictive years, i.e.

[3] More information about this is on the survey website: http://www.cpc.unc.edu/projects/rlms-hse/project/sampling

aged 18-65 (inclusive). The minimum number of eligible respondents per round was 5,794, and the maximum was 8,767. The total number of observations in the whole dataset was 50,072.

We have two dependent variables measuring labour supply: a binary indicator for current labour force participation status, and a natural logarithm of hours worked in the last 30 days, for those reporting positive hours worked. The former variable has a value of one if a respondent says that he or she currently works, is on paid leave or unpaid leave. The value of zero is assigned if the person says he or she doesn't work. The hours worked variable was formed using the answers to the following question:

> "How many hours did you actually work at your primary [secondary] place of work in the last 30 days?"

After summing reported hours worked in a primary and secondary place, we took a natural log of this variable, as the untransformed outcome variable is highly skewed.

With health being a multi-dimensional concept, we used its various definitions in our specifications. Specifically, we consider the following measures of health:

1) Self-reported health (SAH). Respondents answered a question asking them to evaluate their health according to five categories, ranging from very good to very poor. We created a binary variable, assigning it a value of one if respondents rated their health as being poor or very poor, and zero otherwise. The main reason for this transformation is because we would like to take advantage of the panel feature of the data, by estimating individual fixed effects models.

One of the problems with this measure is that it may depend not only on underlying health, but also on socioeconomic status: if richer, or more educated people are more likely to contact a medical care system and get diagnosed with previously unknown conditions, they may actually give lower marks to their health than the reference group (Currie and Madrian, 1999). In this case, the estimated relationship between SAH and labour market outcomes may be subject to a downward bias, unless socioeconomic controls are included. On the other hand, more pessimistic people may underestimate their health, and may also have lower labour supply (thus leading to an upward bias). However, including individual fixed effects should reduce if not totally eliminate this source of this bias, as long as it is time-invariant. In addition, given that SAH usually measures underlying true health with error, and additionally assuming that this error is subject to classical error in the variable assumption (i.e., not correlated with the unobserved true health variable) (Wooldridge, 2002), the parameter on the relationship between SAH and labour market outcome variables may be downward-biased.

2) We also included the effect of diagnosed, although self-reported conditions. Specifically, the acute disease indicators were derived from the answers to the following questions:

-"Has a doctor ever diagnosed you as having had a stroke--blood hemorrhage in the brain?"

-"Have you ever been diagnosed with a "myocardial infarction"?"

For several self-reported chronic conditions we specifically used dummies with the value of one assigned for people who answered that they have liver, lung, kidney, heart diseases, or diabetes. We expect parameters estimated for these variables to be smaller in size than for strokes and myocardial infarctions. It's important to emphasize here that we do not believe that these conditions are not serious. Nevertheless, they may be more difficult to diagnose (in contrast to strokes and myocardial infarctions), and resulting greater measurement error may lead to the downward bias in the parameter estimates for these variables (Wooldridge, 2002). In addition, these chronic conditions (especially heart disease) may indeed have a debilitating effect on the ability to work, although the likelihood of this increases with age. However, since we restricted our sample to those between 18 and 65 years, we hypothesize that it is unlikely that the effect of these conditions will be as strong as for strokes and myocardial infarctions, especially given that the proportion of people self-reporting chronic conditions is quite large (see table 1)

We also included a number of theoretically relevant control variables in the model, including age, residence (urban/rural), marital status, education, family size, wealth, household access (yes or no) to water, sewer, heating, hot water, as well as year and regional dummies. The full list of variables and their description is provided in Appendix B.

Estimation

To start with, we are interested in estimating the parameters in the following model:

(1) $Y_{it} = \beta_1 H_{it} + \beta_2 X_{it} + u_{it}$

Where Y_{it} is a binary indicator for labour force participation or a variable measuring log hours worked for those reporting positive hours for person i at time t; H_{it} is a variable (or a vector) measuring health; X_{it} is a vector of exogenous sociodemographic controls likely to be correlated with health and labour supply[4], such as age, education, marital status, wealth, urban/rural residence, household size, access to water, sewer, heating, hot water, as well as region of living; and u_{it} is an error term.

In general, two major issues are likely to plague the validity of the estimating parameters of model 1. First, health may be correlated with the error term u_{it}. For example, even conditional on including a range of covariates contained in the vector $X_{it,}$ health may still be correlated with certain unobserved determinants of Y_{it}. Some of them (such as individual ability (Strauss and Thomas, 2007)) may be time invariant, while others may change over time. Second, health may be correlated with unobserved country-wide economic shocks, which may affect employment levels. To deal with these two concerns, we are going to estimate parameters using the following model:

(2) $Y_{it} = \beta_1 H_{it} + \beta_2 X_{it} + \alpha_i + \delta_t + \varepsilon_{it}$

Where α_i is a time-invariant endowment of person i possibly correlated with health (e.g., ability); δ_t is the time effect, and ε_{it} is idiosyncratic error term, assumed to be an independent and identically distributed (iid). We estimate model (2) by taking advantage of the panel nature of our data, i.e. by including individual level fixed effects, as well as time effects. This allows us to control for the important source of unobserved heterogeneity in health, possibly at the expense of the loss of precision, especially if health is substantially serially correlated. Although health may not necessarily be correlated with the time-invariant α_i, allowing a more efficient estimation approach with random effects, we have decided to only present results for the OLS and individual fixed effects models.

Comparing models (1) and (2), we see that we have made a restriction that the original error u_{it} does not contain any time-varying unobserved heterogeneity. This may be viewed as a

[4] In the discussion that follows, we will use the term "labour supply" to describe both labour force participation and hours worked for those reporting positive hours worked.

weakness of our approach, although we try to deal with it by including a large range of controls in vector X_{it}. A second problem is that health may also be simultaneously determined with labour force participation in model (2), which neither the control variables nor individual fixed effects can address. This may happen, for example, in the context of the so-called justification hypothesis, when a person explains reduction in labour supply by reporting worse health status (Haveman et al., 1994; Kerkhofs et al., 1999). In addition, loss of a job may affect self-reported health directly (as a stressor) or indirectly through a reduction in health-related consumption. Having said that, the reverse effect running from labour supply to health is unlikely to be of significant concern when we measure health with chronic or acute health conditions (Currie and Madrian, 1999). It may, however, pose a more serious problem for the SAH variable. All these issues may be potentially addressed by instrumental variable estimation. Unfortunately, no theoretically and practically convincing instruments were found in RLMS-HSE for this or other health measures.

We will model labour force participation following equation (2) with a linear probability specification. Although this approach has some drawbacks (e.g. heteroscedastic disturbances and out of bound predictions), they are relatively easy to deal with, given that we are more interested in estimation than prediction, and with the estimation standard errors robust to heteroscedasticity and clustering being straightforward (Scheffler et al., 2007). The benefit of this approach is the ease of interpreting coefficients, and especially computational advantages for the specifications that include individual fixed effects.

Next, we consider how poor health may affect the number of hours worked. Since the distribution of this outcome is highly skewed, it will be log transformed. Although formulation (2) will allow us to model the effect of poor health on the logarithm of hours worked for those reporting positive hours worked, it will not show the combined impact of health on *both* participation (i.e., extensive margin of labour supply) and hours worked (i.e., intensive margin).

To deal with this, we will treat the number of hours worked as a corner solution outcome (Wooldridge, 2002), where zero represents the extensive margin side, while positive hours worked- the intensive margin. Since the conditional expectation of hours worked will be a nonlinear function of the covariates of interest, the marginal effects of health on the logarithm of hours worked can in principle be estimated by Tobit, however this approach imposes the restriction that the effects of health on participation and hours worked should have the same sign.

Instead, we will recover the marginal effects using the two-part model approach (Dow and Norton, 2003) described in more detail below.

Specifically, distribution of hours worked can be presented in the following form, describing a two-stage decision process (Cameron and Trivedi, 2005):

(3) $f(y|x) = \begin{cases} \Pr[d=0 \mid x] & \text{if } y=0 \\ \Pr[d=1 \mid x] f(y \mid d=1, x) & \text{if } y>0 \end{cases}$

Equation (3) describes a conditional density function for the random variable y of interest (in our case the hours worked by the individual), which are only positive for those who work (i.e., d=1). This density is equal to the probability that a person does not participate in the labour force if hours worked are zero (y= 0), or it equals the product of the probability that a person participates in the labour force[5] times the truncated density of y conditional on labour force participation and other factors X, for those with non-negative reported hours worked (y>0).

A natural way to model participation d in equation (3) is by assuming that it is determined by a latent variable framework, and thus a logit or probit functional form is often used (Cameron and Trivedi, 2005). We will run a probit model of the labour force participation on all covariates of interest, using the entire sample of respondents:

(4) $Pr[y>0 \mid X] = \Phi(X'b_{1;}\, \varepsilon_1)$

Where y is hours worked and X is a vector of all covariates, including health. Note that y>0 implies d=1 in equation (3), i.e. a person participating in the labour force. Using estimated parameters b_1, we will predict labour force participation for the whole sample. After that, using the sample of those only reporting positive hours worked, we will estimate parameters in the following specification:

(4) $Ln(y) = X'b_2 + \varepsilon_{2,} \quad for\ y>0$

[5] Strictly speaking, while d is indeed a probability of working, the hours-worked-variable y will be conditional on reporting positive hours worked, which is slightly different from labour force participation. Specifically, 95% of those who work, do report positive hours worked in the sample, while about 5% do not. The reason for this discrepancy is that some people may be on paid/maternity leaves, for example.

This will allow us to predict E[y|y>0, X] for the whole sample, using only the parameter b_2 estimated in the positive hours sample of equation (3).

Next, we will estimate person-specific marginal effects of changes in health on total hours worked using the following formula (Dow and Norton, 2003), taking account of the fact that the outcome is in logs, rather than levels:

(5) $dY/dX = b_2 * \Phi(X'b_1) * exp(X'b_2 + 0.5\sigma^2) + b_1 * \varphi(X'b_1) * exp(X'b_2 + 0.5\sigma^2)$

We will then take an average of these person-specific marginal effects for each health variable of interest, bootstrapping standard errors.

For comparison, we are also interested in how the marginal effects of health on hours worked differ when no extrapolation to the whole sample is made. To be able to draw an appropriate comparison, note again that our outcome variable is log-transformed, and, hence, we need to be careful in translating the parameter β_1 from equation (2), where the outcome is log of hours worked, into a marginal effect of health on hours worked. One way to do it is by estimating the difference: $E(Y_{it} \mid H_{it} = 1; X_{it}) - E(Y_{it} \mid H_{it} = 0; X_{it})$, by predicting both components for each individual, and then taking the difference in mean predictions (Trivedi, 2010):

For simplicity, let us assume that model (1) is expressed in a shortened version:

(6) $Ln(Y_{it}) = \beta_1 H_{it} + u_{it}$

Where variable H_{it} is health status, for example having a certain condition (diagnosed or self-assessed), and u_{it} is an error term. The example can be easily generalized to the case when other control variables are also included.

Exponentiating both sides, and taking the expectation, we obtain:

(7) $E(Y_{it} \mid H_{it}) = \exp(\beta_1 X_{it}) \times E(\exp(u_{it}))$

From this, we can see that $E(Y_{it} \mid X_{it}) \neq \exp(\beta_1 X_{it})$. Since our goal is to estimate $E(Y_{it} \mid H_{it} = 1; X_{it}) - E(Y_{it} \mid H_{it} = 0; X_{it})$, we need to make an additional correction for the error term, as described below.

If u~N(0, δ^2), then E(exp(e))=exp(0.5δ^2) (where δ^2 is estimator of model regression error in equation (6)), and therefore the marginal effect of having illness on hours worked can be estimated as follows (Trivedi, 2010):

$$(8)\ E(Y_{it} \mid H_{it} = 1; X_{it}) - E(Y_{it} \mid H_{it} = 0; X_{it}) =$$

$$\exp(0.5\delta^2) \times (\exp(Ln(Y_{it}) \mid H_{it} = 1; X_{it}) - \exp(Ln(Y_{it}) \mid H_{it} = 0; X_{it}))$$

To emphasize, the Heckman estimator is an appropriate one when the selection problem exists. For example, when no wages are reported, it does not mean that people work for zero wages, they are simply unobserved for them, in which case the conditional expectation function of interest is for potential wages. On the other hand, in this paper we are interested in the marginal effect of poor health on the observed hours of work. In contrast to missing wage, zero hours are actual outcome, ie the outcome of utility maximizing behavior of individuals (Wooldridge 2002). In other words, we are not trying to solve a selection problem (which we believe does not exist), but rather try to account for the non-linear conditional expectation function of the outcome variable with the flexible two part model.

Results

Descriptive statistics

We present descriptive statistics for the pooled sample of respondents 18 years of age (the age at which Russian commonly begin to enter labour force) and older in Table 1[6]. We see that the mean age of the sample participants did not fluctuate much. On the other hand, the proportion of males has been generally fluctuating over time around 43-45% range, with no clear trend. It is also interesting that the proportion of people living in the cities had been on the decline prior to 2001, after which it rose again and stabilized at higher levels. This may suggest that as the economic conditions started to improve, people started to move back to the cities. Unsurprisingly, the percentage of the sample population currently working appears to be strongly driven by the macroeconomic conditions, as this indicator reached its trough at the time of the 1998-1999 economic crisis.

[6] To correct for the non-independence in the data collection, we adjusted the estimated standard errors for the effect of clustering at the commune level.

As far as the specific heath indicators are concerned, we can see that the proportion reporting their health as poor or very poor has been on a steady decline over the observed period, although it started to decrease to a greater extent after 2000. The proportion of working age people who have had myocardial infarctions or strokes (and survived them) has been increasing over the years, although it is difficult to say whether this is due to objectively worse cardiovascular health, or to better survival of patients experiencing such episodes. The fact that the proportion of working age people with diabetes and heart disease has been overall on the increase suggests that the former explanation may play an important role. Finally, the proportion of patients with chronic self-reported liver or lung disease has slightly declined over the observed period.

Finally, in Figure 1, we see that self-reported health is strongly related to the probability of working: among working age adults with very poor health, about 80% are not working. On the other hand, although the proportion of people working steadily increases with better health, the gap between those working and not working in very good health is small. However, this is not really surprising if one notes that about 63% of the respondents who have rated their health as very good are younger than 30 years. A substantial proportion of people in this age group (about 16%) are students, and an additional 15% classify themselves as temporarily not working.

Currently working

In Table 2, we examine the association between having several health indicators and two individual labour market outcomes- currently working and the logarithm of the number of hours worked in the last 30 days. For the currently working outcome, this association is strong and significant for most health variables in both OLS specification (column 1) and the Probit one (column 2, where marginal effects are shown), with results being very similar. We find that being in poor self-reported health is associated with an about 20% lower probability of working. The association of currently working with two indicators of acute health events - heart attacks and strokes - is somewhat smaller, but note that here we are controlling for SAH and chronic heart disease. When these two variables are removed (not shown here), heart attacks are associated with a 16%, and strokes with a 22% lower probability of work. A stronger effect from ever having strokes can probably be explained by a more lasting impact on the ability to work for stroke survivors, or (speculatively) from greater selective mortality among miocardial infarction sufferers. Unsurprisingly, most other chronic conditions are weakly, or not at all related to the probability of working. Two exceptions are heart or lung diseases, related to correspondingly an 8% and 5% lower probability of work.

All control variables have the predicted signs (results available on request). Thus, males, married people, those with more education and with more wealth are all more likely to be employed than the reference groups. There is a nonlinear relationship between the probability of work and family size: it attains its maximum at the size of 3-5 members, and then falls continuously with every additional member. As the family size increases, the probability of working continuously falls.

There is a significant drop in the size of the parameter estimates after the inclusion of individual fixed effects for poor health, diabetes, lung and heart disease, suggesting there was a sizeable upward bias in these variables. On the other hand, the size of the effect became stronger for liver disease, suggesting that OLS estimates may have been downward-biased for this variable. Finally, there was little change for myocardial infarctions and strokes, suggesting that the relationship between these two variables and working may be little affected by individual-level bias.

Hours worked

In column 4, we see that poor health is significantly associated with fewer hours worked. Thus, having this status is linked to about 2% fewer hours worked in both the OLS and IFE models, although the IFE parameter is not significant. Rather surprisingly, this association is only significant in both OLS and IFE for only one variable-heart disease. Also note that removing poor health and heart disease dummies made no difference to the parameters on myocardial infarction and strokes, in contrast to the case when the outcome variable was ‘currently working’. Nevertheless, we need to bear in mind that this association may be underestimated for the sample reporting positive hours only, and in our case, all those not working were excluded since we took a natural log of our outcome variable. We will deal with this issue further below.

Heterogeneity of health effects

An interesting issue that is dealt with relatively rarely in the empirical literature is how the estimated association between health and labour supply varies across a distribution of a particular population characteristics. In this section, we graphically examine how the SAH indicator is related to the probability of labour force participation by age, wealth, residence and gender. In this case, it is more instructive to consider the estimates for the whole sample than only for those who reported positive hours worked; therefore we chose the labour force participation indicator rather than log hours worked as our outcome variable.

In Figure 2, we show the distribution of the coefficients from the linear probability regression of labour force participation on poor health, stratified by age and place of residence. For the urban subsample in particular, there appears to be some evidence that as age increases, the relationship between poor health and labour force participation first gets stronger, and then weakens after retirement age. The weakening effect after the age of 60 can be due to the possibility that it is mostly the least well-off people who continue working after retirement, while for the rest, the link between working and health no longer exists. Also, as expected, for most parts of the age distribution, the effect of health on labour force participation for the subsample of urban dwellers is stronger. The difference increases in middle ages and continues until official retirement age (55 years for women; 60 years for men). This could be either because middle-aged people living in the urban areas are wealthier, or because they have better access to social and insurance services that allow them not to work upon suffering adverse health events.

From Figure 3 we see that there is considerable difference in the effect of poor health by education status in most age groups, suggesting that it is people with *less* education (and therefore with a lower socioeconomic status) who are more likely to stop working in response to being in poor health. Therefore, the previous finding of a stronger effect of health on labour force participation in urban than in rural areas does not appear to be due to greater wealth accumulated by the city dwellers. A more plausible explanation appears to be the availability of certain urban-specific factors (e.g., formal or informal social protection mechanisms) that may make it easier to stop working in the cities when adverse health events occur.

Finally, a surprising finding in Figure 4 is that men are considerably more likely to stop working when they experience health problems than women. A possible part of the explanation here is that men may generally define “poor health” differently from women, in that for men only particularly serious conditions may be seen as a sign of “poor health”[7]. Nevertheless, the magnitude of the difference is quite unexpected for most age groups. One potential explanation is that women may feel more responsible for their families, compared to men, and thus may continue working despite being in poor health (Ashwin and Bowers, 1997; Fong, 1993). Using another dataset, a similar finding was made in another paper (Goryakin et al., 2013).

Extrapolating to the whole sample

The results listed in the previous section suggest that several health conditions, in particular stroke, heart disease and poor SAH, may have a considerably stronger effect on labour force participation than on the log of hours worked. While this is in line with our prior expectation, we should be making a fair comparison between the effects of health conditions on the intensive and extensive margins of labour supply. To achieve this goal, we need to take into account the fact that for people who do not work, the log-transformed values of their hours worked in the last 30 days (ie those presented in column 5 of Table 2) will lead to missing observations. As discussed in the empirical section, we prefer to correct for this by means of a two part model. In Table 3, we present the two part model results (column 1), comparing them to the log-linear model for the sub-sample of those who reported positive hours worked, with marginal effects derived (column 2) using the correction described in equation (8).

The two part model estimation showed that the largest effect on the overall number of hours worked in the last 30 days was found for heart disease and strokes, which were associated with respectively about 20 and 17 fewer hours worked in the last 30 days than for the reference group. Similarly, the effects of poor health, myocardial infarctions, diabetes, and lung disease on hours worked was found to be substantial in column 1, which is different from the evidence previously presented for the outcome variable ‘log hours worked’ for the sample of those reporting positive hours (Table 2), as well as marginal effects from the log-transformed model estimated using model (8) (column 2, Table 3). Thus, only poor health, heart and lung diseases

[7] In support of this theory, the RLMS-HSE data suggests that men are less likely to self-report poor health than women, even though they are generally unhealthier by other, more objective indicators (e.g., life expectancy at birth, rates of cardiovascular mortality).

were significantly associated with the loss of working hours in the last 30 days, although the magnitude of the effect was much smaller than for the corresponding parameters in column (1). This again highlights the issue of the differential effects of diseases on the extensive and intensive margins of labour supply, depending on their severity, and the importance of taking account of both for an objective picture. For two chronic diseases (liver and kidney disease), the two-part model does not make as much difference as for the others conditions, as none of their parameters are significant in either model. This is intuitively logical, because the latter appear to affect the total monthly hours worked primarily through labour force participation, rather than through hours worked for the sample reporting positive hours worked.

Discussion

As expected, the negative association between health and labour force participation was strong (both in OLS and IFE models) for several conditions we believed to be potentially more serious, and/or more likely to be diagnosed, i.e. indicators for poor health, strokes and myocardial infarctions. Interestingly, contrary to our initial expectation, one chronic disease in particular-heart disease- had a consistently strong negative association with labour force participation across specifications, suggesting that even though it is self-reported, any potential downward bias resulting from the measurement error in misdiagnosis is likely to be outweighed by its seriousness. For other chronic conditions in IFE specification, the association was either smaller in magnitude (liver disease), insignificant (diabetes or lung disease), or had a wrong sign (kidney disease). However, overall, 4 out of 8 conditions had a significant negative association with working in both OLS and IFE specifications. These findings are in notable contrast to the estimated relationships between health and the logarithm of hours worked in the last 30 days: only one condition (heart disease) had a significant negative association with log hours worked in both OLS and IFE specifications.

Does this imply that health measures do not have a significant effect on the intensive margin of labour supply? As estimates from the two-part model demonstrated, the answer is ‘No’. Indeed, we found that the largest effect on the overall number of hours worked was for poor health, strokes and heart disease. Interestingly, the effect was also significant and quite large in size for diabetes, suggesting that it may have a stronger effect on the intensive margin of labour supply than on the labour force participation. Although the effect was significant for myocardial infarction, it was relatively modest in size. This may be due to the selective mortality effect among the most serious cases, rather than dropping out of the labour market. The lesson here is that one needs to carefully choose the appropriate estimators when dealing with samples with a lot of missing values.

One potential concern in this paper is the use of a two part model to model the effect of poor health on the logarithm of hours worked. It may be argued, for example, that even though we are interested in estimating marginal effects on the actual number of hours worked, rather than on the potential number of hours worked if people were in the labour force, one can still argue that there is a selection issue to deal with. While similar concerns have been considered elsewhere (Dow and Norton, 2003; Salas and Raftery, 2001), we nevertheless conducted a formal statistical test for whether the selection issue exists (results available upon request). Specifically, we assumed the following exclusion restriction: the number of household members could affect the number of hours worked only through their effect on the probability of working. If this assumption is valid, then the parameter on the inverse Mills ratio would identify whether any selection problem exists. In this instance, we found that although the number of household members was strongly and significantly related to the probability of work, the Mills ratio parameter (lambda) was not significant.

How do our results compare with the existing literature on the effect of ill health on labour market outcomes in Russia? Overall, there were only a few similar studies conducted on data collected in that region. First, using RLMS survey data collected in 1997-2004, Abegunde et al (2008) estimated the effect of chronic and non-chronic illness on a range of labour market outcomes, although labour force participation per se was not one of them. Nevertheless, they did find there was significant association between illness and lost days from work for heads of households (the association appeared stronger with acute compared than with chronic diseases). A study using data collected in 10 post Soviet countries, (Goryakin et al., 2013) found that poor health was significantly related to about a 15% lower probability of work in the community fixed effects specification, this magnitude increased to about 36% in the instrumental variable model (note how this compares to 20% we found in our OLS model in this paper). Finally, Suhrcke et al (2007) found that in Russia, self-assessed good health was mostly unrelated to log of hours worked per week. However, this finding could also have been due to its focus only on those who reported positive hours worked, rather than the true lack of effect. Note that we found a similarly weak association between almost all measures of health and the logarithm of hours worked.

Frequently in the empirical literature, conclusions about the labour market consequences of poor health are based on only one or two of such outcomes, which may give an incomplete picture. For example, one can find comparatively little published research on the association of health with hours worked. However, if this result is only obtained for the sample of those reporting only positive hours worked, and those reporting zero hours worked are ignored, this can give a misleading picture. By also considering the health effect on labour force participation in the context of a two-part model, we are more likely to accurately assess the overall effect of health on labour supply.

When considering these results, it's important to keep in mind that the effect of illness on individual or household welfare may depend not only on how labour supply responds to disease at the individual level, but also on intra-household allocations on labour, on whether the sick people are in wage or salaried employment, as well as on the features of the social protection systems. In addition, poor health may be related to the loss of non-medical consumption, both due to greater spending on medical care, as well as because of the loss of income (Wagstaff, 2007).[8]

The welfare burden therefore may be borne by a range of players, including the individuals that have fallen sick, their household members, their employers, or the state. For example, there could be little observed relationship between health, income and labour supply both on individual and family levels. However this does not necessarily mean that such health events are costless if the family has to cut back on their non-medical consumption to cover the increased medical costs (especially given the relatively high share of out of pocket spending on drugs in Russia), or if they have to sell off their assets if there are no appropriate insurance mechanisms in place (Benjamin et al., 2003). Alternatively, those suffering from disease(s) may have to continue working despite having poor health, which suggests an additional cost of poor health not easily captured by traditional approaches. This was illustrated with the evidence presented on Figure 2-Figure 4, where we showed that the negative association between heath

[8] In addition, poor heath may directly reduce the utility of consumption.

and labour force participation was considerably weaker in rural areas than in the cities, as well as among women.

To deal with all these issues in one paper would be too ambitious. Rather, our goal is much more modest, but also better defined- we are interested in how the labour supply of individuals responds to poor health. Despite this limitation in scope, our estimation of the effect of being ill on the probability of work, as well as on the number of hours worked, may prove useful for policy analysis, including the estimation of the indirect cost of illness for cost-effectiveness modeling. The main take away message of this research is that having poor health is economically costly to the sick, although the distribution of the burden depends on the character of illness (e.g., how serious is the condition), as well as on whether ill people can "afford" to withdraw from the labour force, because of the nature of where they live and or who they are.

Another important limitation of our research is that some of the variables (e.g. self reported health or hours worked) may be measured with error, which may lead to the downward bias in the estimations, or to less precision. Finally, the estimation may have suffered from residual endogeneity. For example, there could be an issue of reverse feedback from labour market outcomes such as income and labour supply to health. However, we were not able to deal with this due to the lack of good instruments.

Most of the existing literature on the link between health and labour market outcomes has focused on either high income or low income countries, giving middle income countries like Russia relatively little attention- a gap that we hopefully partly addressed with the research presented in this paper. The context of Russia in particular, with its fast-paced economic reforms over the last two decades, including large-scale privatization of state-owned enterprises, provides a particularly rich ground for such research. In the future, more work on the effect of poor health in Russia on other related outcomes- including medical and non-medical consumption, intra-household allocations of labour supply, as well as individual and household-level income, will provide a useful extension to the research presented in this paper.

Table 1. Summary statistics for the RLMS-HSE (1995-2008) sample.

	1995	1996	1998	2000	2001	2002	2003	2004	2005	2006	2007	2008
Poor Health	12.2%	12.0%	12.3%	11.7%	10.7%	10.0%	10.5%	9.5%	8.6%	8.9%	8.4%	8.8%
Myocardial infarction (MI)	1.6%	2.1%	2.4%	2.5%	2.4%	2.5%	2.7%	2.8%	2.7%	2.4%	2.6%	2.6%
Stroke	0.8%	1.0%	1.2%	1.4%	1.6%	1.7%	1.7%	1.7%	1.9%	1.6%	1.9%	2.1%
Diabetes	2.6%	3.3%	3.4%	4.3%	4.7%	4.9%	5.5%	5.3%	5.7%	4.8%	5.7%	6.0%
Heart disease	-	-	-	12.2%	14.9%	17.0%	18.3%	18.8%	19.4%	17.0%	18.7%	19.8%
Liver disease	-	-	-	8.2%	8.9%	8.9%	9.2%	7.8%	7.1%	6.9%	6.6%	7.1%
Lung disease	-	-	-	5.6%	5.5%	5.4%	5.3%	4.9%	4.5%	4.6%	4.3%	4.4%
Currently working	67.8%	66.0%	61.8%	61.8%	62.7%	63.0%	64.1%	63.9%	64.6%	65.7%	67.0%	68.0%
Age	40.5	40.5	40.7	40.9	40.5	40.6	40.5	40.5	40.1	39.9	40.2	40.4
Male	45.3%	45.0%	45.1%	43.8%	44.0%	44.4%	44.3%	44.4%	45.4%	44.8%	44.8%	44.3%
Urban	69.8%	67.9%	66.9%	65.0%	68.3%	68.3%	67.5%	67.7%	67.1%	68.8%	68.2%	68.4%

Notes: All the statistics are adjusted for the sampling weights and clustering at PSU level. Unless otherwise noted, all statistics reported for respondents between ages 18 and 65.

Table 2. Association between health and individual labour market outcomes

	(1) Currently working OLS	(2) Currently working Probit	(3) Currently working IFE	(4) Log hours worked (30 d) OLS	(5) Log hours worked (30 d) IFE
Poor health	-0.196***	-0.21**	-0.055***	-0.019*	-0.012
	(0.012)	(0.01)	(0.008)	(0.011)	(0.011)
MI	-0.056**	-0.06**	-0.060**	0.001	-0.066**
	(0.027)	(0.03)	(0.029)	(0.026)	(0.030)
Stroke	-0.120***	-0.139**	-0.135***	0.031	-0.011
	(0.028)	(0.03)	(0.035)	(0.029)	(0.050)
Diabetes	-0.043***	-0.047**	-0.015	-0.025	-0.066***
	(0.017)	(0.02)	(0.013)	(0.025)	(0.019)
Heart	-0.097***	-0.106**	-0.038***	-0.034***	-0.021*
	(0.012)	(0.01)	(0.007)	(0.012)	(0.012)
Liver	0.012	0.011	-0.015**	0.001	-0.006
	(0.013)	(0.01)	(0.007)	(0.010)	(0.014)
Kidney	0.010	0.011	0.027***	-0.001	0.018
	(0.010)	(0.011)	(0.010)	(0.009)	(0.013)
Lung	-0.038***	-0.04**	0.002	-0.023*	-0.019
	(0.014)	(0.02)	(0.009)	(0.013)	(0.019)
Region dummies	Yes	Yes	No	Yes	No
Observations	60,267	60,267	60,267	38,669	38,669
R-squared	0.142	0.12	0.034	0.037	0.006

Notes: Community cluster-robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
OLS: ordinary least squares. IFE: individual fixed effects.
In addition, all specifications contain round dummies, as well as control variables: dummies for age, being male, married, living in urban areas, having high school and university diplomas, 4 indicators for wealth, corresponding to relevant quintiles occupied by households (adjusted for regional poverty level), four dummies for household size, as well as controls for availability of water, cold water, sewer and heating in the households. Sample restricted to adults between ages 18 and 65. In column 2 (Probit model), marginal effects are shown.

Table 3. Effect of health variables on the predicted number of hours worked (last 30 days) for the whole sample

	2-Part model (1)	Log-level model (2)
Poor health	-10.745**	-3.428*
	(5.222)	(1.969)
MI	-8.438*	0.129
	(5.073)	(4.936)
Stroke	-17.690***	5.530
	(6.700)	(5.400)
Diabetes	-10.093***	-4.458
	(3.696)	(4.151)
Heart	-19.982***	-6.003***
	(2.280)	(2.044)
Liver	2.043	0.169
	(2.384)	(1.709)
Kidney	1.479	-0.266
	(2.043)	(1.698)
Lung	-8.784***	-4.095*
	(2.319)	(2.399)
Observations	60,268	38,669

Notes: All specifications include contemporaneous controls, regional and round dummies (see table 2). Sample restricted to adults between ages 18 and 65. In the first column, dependent variable is a predicted number of hours worked for the whole sample, using formulas from Dow and Norton (2003). For participation equation, the probit model was run. In the second column, marginal effects of health on hours worked, derived from the log-level model are presented. In first and second column, bootstrapped standard error (using 200 replications) are provided.

Figure 1. Association between self-reported health and labour force participation

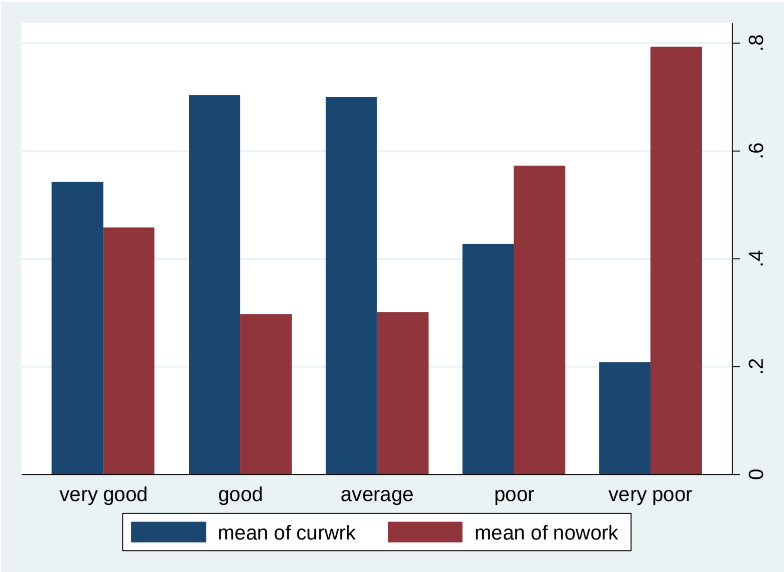

Source: RLMS-HSE dataset. Sample of adults aged 18-65, inclusive.

Figure 2. Coefficients from a regression of labour force participation on poor health dummy, by age and place of residence

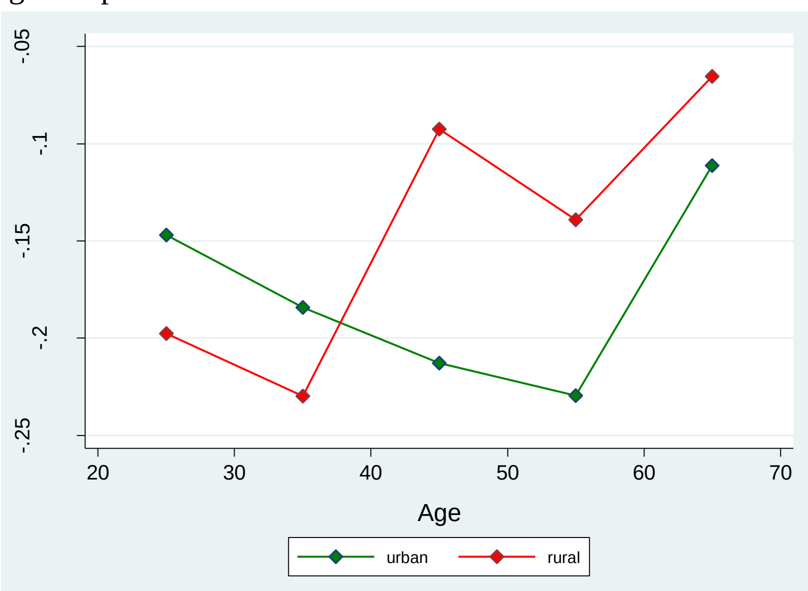

Source: RLMS-HSE dataset. Sample of adults aged 18-65, inclusive.

Figure 3. Coefficients from a regression of labour force participation on poor health dummy, by age and education

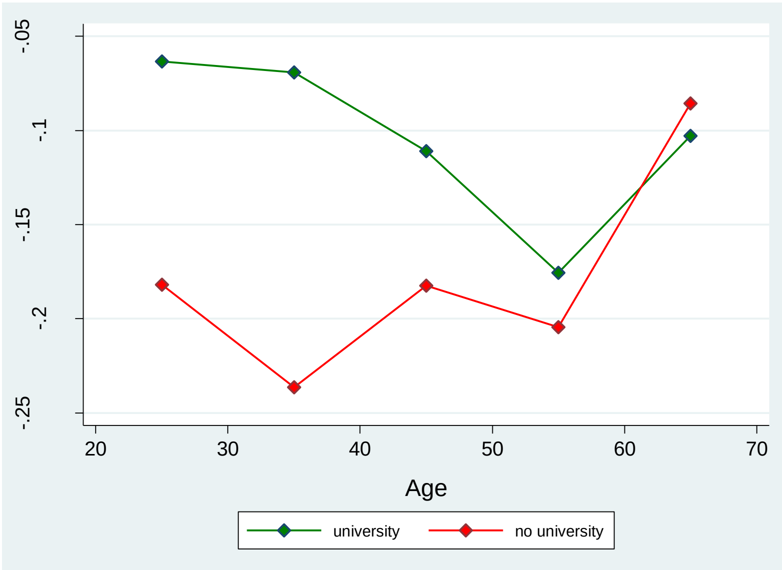

Source: RLMS-HSE dataset. Sample of adults aged 18-65, inclusive. Poor refers to being in the bottom quartile of the wealth distribution; rich refers to being in the top quartile.

Figure 4. Coefficients from a regression of labour force participation on poor health dummy, by age and gender

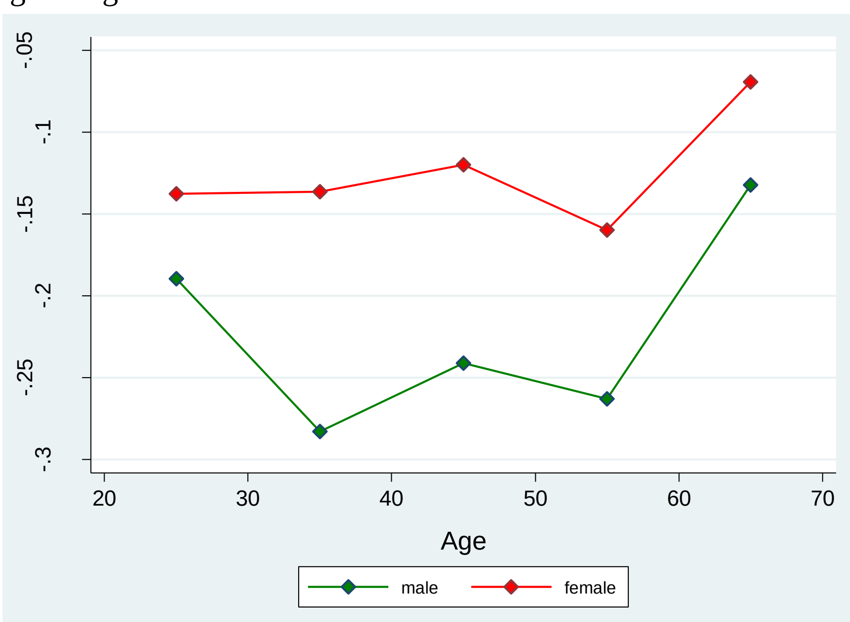

Source: RLMS-HSE dataset. Sample of adults aged 18-65, inclusive.

Appendix B. Variable definitions

Variable	Description
Working	Proportion reporting as working, on paid or unpaid leave at the time of the interview
Log hours worked	Log hours worked in the last 30 days
Poor health	Proportion reporting poor health
Heart	Have heart disease? (self-reported)
Kidney	Have kidney disease? (self-reported)
Liver	Have liver disease? (self-reported)
Lung	Have lung disease? (self-reported)
Household size	4 dummies: 2; 3/5; 6/8; 9/13 members
Age	Age
Married	Proportion in registered marriage, or living together as partners
Male	Proportion of males
Urban	Proportion living in urban areas
Hsdiploma	Proportion having high school diploma as highest degree
University	Proportion having university diploma as highest degree
Water	Have cold water supply in house?
Sewer	Have sewer in house?
Heating	Have heating in house?
Hot water	Have hot water in house?

Source: RLMS-HSE dataset, rounds 6-17. Sample restricted to people aged 18-65, inclusive.

References

Abegunde, D.O., Stanciole, A.E., 2008. The economic impact of chronic diseases: how do households respond to shocks? Evidence from Russia. Social Science & Medicine 66, 2296-2307.
Andreev, E.M., Nolte, E., Shkolnikov, V.M., Varavikova, E., McKee, M., 2003. The evolving pattern of avoidable mortality in Russia. International Journal of Epidemiology, pp. 437-446.
Ashwin, S., Bowers, E., 1997. Do Russian women want to work? Post-Soviet Women: From the Baltic to Central Asia. Ed. Mary Buckley. Cambridge: Cambridge UP, 21-37.
Benjamin, D., Brandt, L., Fan, J.Z., William Davidson, I., 2003. Ceaseless Toil?: Health and Labor Supply of the Elderly in Rural China. William Davidson Institute.
Bloom, D.E., Canning, D., 2000. The Health and Wealth of Nations. Science 287, 1207-1209.
Brainerd, E., Cutler, D.M., 2005. Autopsy on an Empire: Understanding Mortality in Russia and the Former Soviet Union. Journal of Economic Perspectives 19, 107-130.
Cameron, A.C., Trivedi, P.K., 2005. Microeconometrics: Methods and Applications. Cambridge University Press.
Currie, J., Madrian, B.C., 1999. Health, Health Insurance and the Labor Market. Handbook of Labor Economics 3, 3309-3416.
Dow, W.H., Norton, E.C., 2003. Choosing between and interpreting the Heckit and two-part models for corner solutions. Health Services and Outcomes Research Methodology 4, 5-18.
Fong, M.S., 1993. The role of women in rebuilding the Russian economy. World Bank Publications.
Goryakin, Y., Rocco, L., Suhrcke, M., Roberts, B., McKee, M., 2013. The effect of health on labour supply in nine former Soviet Union countries. The European Journal of Health Economics, 1-12.
Haveman, R., Wolfe, B., Kreider, B., Stone, M., 1994. Market work, wages, and men's health. J Health Econ 13, 163-182.
Kerkhofs, M., Lindeboom, M., Theeuwes, J., 1999. Retirement, financial incentives and health. Labour Economics 6, 203-227.
Marquez, P., Suhrcke, M., McKee, M., Rocco, L., 2007. Adult health in the Russian Federation: more than just a health problem. Health Affairs 26, 1040-1051.
McCartney, G., Mahmood, L., Leyland, A.H., Batty, G.D., Hunt, K., 2011. Contribution of smoking-related and alcohol-related deaths to the gender gap in mortality: evidence from 30 European countries. Tobacco control 20, 166-168.
Salas, C., Raftery, J.P., 2001. Econometric issues in testing the age neutrality of health care expenditure. Health Economics 10, 669-671.
Scheffler, R.M., Brown, T.T., Rice, J.K., 2007. The role of social capital in reducing non-specific psychological distress: The importance of controlling for omitted variable bias. Social Science & Medicine 65, 842-854.
Shkolnikov, V., McKee, M., Leon, D.A., 2001. Changes in life expectancy in Russia in the mid-1990s. The Lancet 357, 917-921.
Stillman, S., 2006. Health and nutrition in Eastern Europe and the former Soviet Union during the decade of transition: A review of the literature. Economics and Human Biology 4, 104-146.
Strauss, J., Thomas, D., 2007. Health over the life course. Handbook of development economics 4, 3375-3474.

Suhrcke, M., 2007. Economic Consequences of Noncommunicable Diseases and Injuries in the Russian Federation. WHO Regional Office for Europe.
Suhrcke, M., Rocco, M., McKee, M., Mazzuco, S., Urban, D., Steinherr, A., 2007. Economic consequences of noncommunicable diseases and injuries in the Russian Federation. WHO Regional Office for Europe.
Trivedi, C.C.A.P.K., 2010. Microeconometrics Using Stata. Stata Press.
UNDP, 2005. Russia in 2015: development goals and policy priorities.
Wagstaff, A., 2007. The economic consequences of health shocks: evidence from Vietnam. Journal of Health Economics 26, 82-100.
Weidner, G., Cain, V.S., 2003. The gender gap in heart disease: lessons from Eastern Europe. American journal of public health 93, 768-770.
Wooldridge, J.M., 2002. Econometric analysis of cross section and panel data. The MIT press.