Written Exam at the
Department of Economics, summer 2020
Economic Growth
Final Exam, June 2, 9am-noon

3-hour open book exam. Answers only in English.

This exam question consists of 5 pages in total.

Upload your answer as one single PDF-file. The PDF must be named with exam number only (e.g. '127.pdf') and uploaded to Digital Exam.

This exam has been changed from a written Peter Bangsvej exam to a take-home exam with helping aids. Notice that any communication with fellow students or others about the exam questions during the exam is considered to be cheating and will be reported. It is also considering cheating to let other students use your product.

Be careful not to cheat at exams! You cheat at an exam, if during the exam, you:

- Make use of exam aids that are not allowed
- Communicate with or otherwise receive help from other people
- Copy other people's texts without making use of quotation marks and source referencing, so that it may appear to be your own text
- Use the ideas or thoughts of others without making use of source referencing, so it may appear to be your own idea or your thoughts
- Or if you otherwise violate the rules that apply to the exam

1 Short essay questions

Question 1.a

Suppose that AI suddenly makes it profitable to replace some tasks previously undertaken by skilled workers. Use the task-based model with three skill levels to explain what happens to the ranges of tasks performed by the three types of workers, and their relative wage rate. What is the intuition?

Answer: The increased use of AI is analogous to an increase in the supply of skilled workers. The implication is that the range of tasks performed by skill workers/AI increases, whereas the ranges of tasks performed by the other skill levels become smaller. Because AI works as a supply shock, it will decrease relative wages for skilled workers. The effect on wages of medium skilled workers relative to low skill workers is uncertain.

Question 1.b

In their paper "Beyond GDP? Welfare across Countries and Time", Jones and Klenow assume that life expectancy is determined at country level, and does not vary across individuals within each country. Would relaxing that assumption lead to higher or lower differences in welfare across countries? Would the US look better or worse compared to countries in Western Europe? Discuss.

Answer: Life expectancy will most likely correlate with consumption possibilities because richer people are better able to afford both health care and a healthy lifestyle. That would exacerbate inequalities within countries, thereby reducing overall welfare. By implication, more unequal countries in terms of consumption possibilities will now look worse. Developing countries are typically more unequal that developed countries, so the change to how life expectancy is modeled will increase the welfare gap between rich and poor countries. By the same argument, the US would look relatively worse than Western Europe. One could also note that the correlation between life expectancy and income presumably is smaller in countries with universal health care, which again would favor Western Europe in comparisons with the US or with developing countries.

Question 1.c

Suppose that the economy contains a continuum of firms with identical O-ring production functions, and a continuum of workers with different levels of human capital. Assume further

that economic growth manifests itself through increased complexity of the products of each firm (i.e., in a higher n). What happens to wage inequality in the economy? Explain.

Answer: A feature of O-ring production function is that skill levels are equalized within firms, meaning that some firms will only use high quality workers, whereas others will only use low quality workers. A higher n means that more can go wrong, and because of complementarity across workers within firms, that implies that the value of employing low-skill workers decreases compared to that of high-skill workers. Wage inequality will therefore increase.

2 Trade and the labor market

Consider the regression i Autor, Dorn, and Hanson (2015):

$$\Delta Y_{jkt} = \gamma_t + \beta_1 \Delta IPW_{jt}^{China-US} + \beta_2 RSH_{jt} + X'_{jt}\beta_2 + \delta_k + e_{jkt} \tag{1}$$

Y_{jkt} is the outcome of interest, such as the employment-to-population ratio, in commuting zone j, census region k, at time t. $IPW^{China-US}$ captures import from China, RSH_{jt} is the share of the workforce employed in routine jobs, and X_{jt} is a vector of control variables. The $\Delta's$ signify changes over the sample period 1990-2007. First differencing takes care of time-invariant characteristics of the commuting zones. γ_t represents time fixed effects, and δ_k represents possible differential time trends in the census regions (which are more aggregate entities than commuting zones). The parameters of interest are $\beta_1$and β_2, which capture the effects of trade and computerization, respectively.

Question 2.a

Suppose now $IPW_t^{China-US}$is a measure of actual imports from China into commuting zone j. In this case, explain why $\beta_1$might be a biased estimate of the true causal effect of import competition from China. Explain in which direction potential sources of endogeneity might bias the estimate.

Answer: A positive demand shock in commuting zone j will increase both employment and imports, leading to a positive bias in β_1. A positive productivity shock (supply shock) will increase employment, but decrease imports from China because domestic goods are now more competitive.

Question 2.b

How does Autor, Dorn and Hanson (2015) try to solve such endogeneity problems? What are the identifying assumptions? Are the problems you have mentioned in your answer to the previous question solved? Explain.

Answer: They use a shift-share instrument based on local employment shares across industries, and national trends in imports from China of goods produced by those industries. The identifying assumptions are that commuting zones are too small to affect national trends, and that the initial industry-mix is exogenous once fixed effects are accounted for by first-differencing. Under these assumptions, the endogeneity biases mentioned above will be solved. One could imagine cases in which both assumptions are violated. One would be spatial correlation of demand shocks, although this possibility is partially alleviated by including δ_k. Another would be that if national regulation reduced productivity across the board in some industries (a negative supply shock), then employment and production in those industries will decline if the elasticity of demand is higher than one. Imports from China and elsewhere might partially fill the shortfall, but employment would have fallen regardless.

Question 2.c

The 2SLS results of Autor, Dorn and Hanson (2015) unambiguously show that import competition from China have reduced the employment rate in affected commuting zones. Explain why the effect on the national employment rate nevertheless are ambiguous. Would you consider an increase or a decrease in the national employment rate a more likely consequence of Chinese import competition? Does your answer depend on the time horizon?

Answer: Autor, Dorn, and Hanson (2015) estimate local effects of Chinese imports in commuting zones directly affected by import competition **relative** to less affected commuting zones. But if cheaper Chinese goods increases demand for all other domestic goods, then labor demand in commuting zones not directly in competition with China might increase. The net effect on employment depends on how big the savings are, on what they are spent on, and on how many of the affected workers who find a new job. In the short term, it seems likely that the net effect on employment is negative, as at least some of the spending previously used on domestic goods are now spent on foreign goods. But as some of the affected workers begin to find new employment, the net effect might be zero or even positive.

Question 2.d

Suppose that you re-estimate the regression using data from 1990-2017 instead. Would you expect the estimate of $\beta_1$to change? If so, in what direction? Explain.

Answer: Job losses from Chinese import competition might be persistent, but we should nevertheless expect the unemployed to gradually find new jobs (or retire). By implication, most job losses captured by ΔY_{jkt} should be expected to be of a relatively recent date. In the short sample, most of the increase in Chinese import competition happened after China's entry into the WTO in 2001, i.e., towards the end of the sample. The results of Autor, Dorn, and Hanson (2015) should consequently be viewed as reflecting short-term effects. In the longer sample, 1990-2017, many of the workers who lost their job to Chinese competition in the 2000s probably are employed again, or retired. The estimated β_1 should consequently be expected to be numerically smaller in the long sample, although it is still likely to be negative and significant.

3 Population growth and decline

Global population growth has been slowing down for decades, and the number of people in the world may even start start to decline sometime in the future (see Figure 1 at the end of the exam question). In this part of the exam, you are asked to discuss potential consequences of this trend in light of the models considered in the course.

Question 3.a

What is the likely consequence of a smaller global population (i.e., negative population growth) for global productivity growth according to Paul Romer's R&D based endogenous growth model ("growth through expanding variety")? Explain why it makes a difference whether declining population manifests itself in fewer skilled or unskilled workers.

Answer: In Paul Romer's model there is an assymetric effect of changing the level of the skilled or unskilled labour supply. If the former changes, growth declines, whereas a decline in the latter (in the standard formulation of the model) leaves the growth rate unaffected. The reason is the following. If the total supply of skilled labor expands it admits more skilled works in the production of final goods as well as in the production of new ideas (Research labor). More human capital in the final goods sector will increase demand of new varieties and thus of new ideas, which expands the demand for (skilled) R&D labor. Accordingly, more

"H" will serve to expand growth, which means the model would predict that fewer skilled laborers will reduce productivity growth. The key difference to unskilled labor is that "L" is not used in R&D activtities, but solely in final goods production. More L will, much like a greater supply of H, expand demand for new varies, and thus new ideas. This mechanism will work to increase growth. At the same time, however, more L will make skilled labor more productive in the final goods sector, prompting a reallocation of H away from research and into final goods production. This mechanism works to lower growth. In the standard formulation of the model the two countervailing effects off set, leaving growth unaffected. Accordingly, if we use the Romer model as a guide a declining world population will only work to lower growth if it translates into a lower supply of skilled workers, which need not be the case.

Question 3.b

How and why would your answer change if you use Charles I. Jones modification of the Romer model to form your expectations? Which of the two frameworks appears empirically most relevant?

Answer: Charles I. Jones modifies the Romer model by assuming diminishing returns to knowledge in the R&D sector. In the terminology of Jones this effect is referred to as the "fishing out effect", and implies that R&D productivity gradually declines as more knowledge is discovered. This implies, in turn, that in order to sustain growth in knowledge, more and more (skilled) people needs to be allocated to the process of discovery. More and more people needs to be engaged with R&D. Hence, if global population growth declines (especially if it turns negative) this will (dramatically) reduce growth in productivity.
There is by now evidence which in fact suggests that R&D productivity has been declining over the last several decades in the USA. This is true for specific businnesses such as microchip manufacture, where a constant – amazing - annual productivity of growth about 35% since the early 70s only has occured by increasing R&D input by a factor of 70. This suggests (dramatically) declining R&D productivity, consistent with the fishing-out effect in the Jones model, but not with the assumption of constant R&D productivity in the Romer model. Similar findings emerge when you consider the US economy in its entirety. Hence, the Jones model may be the most empirically relevant of the two, if this evidence is a guide, and by extension its predictions regarding the future of productivity growth.

Question 3.c

The real rate of interest on safe assets, such as government bonds, has declined over the last three decades. More specifically, what is called the natural real rate of interest has declined suggesting long-run factors are of importance to the developments. Some growth models predict that declining (but still positive) population growth may be an explanation, others predict that population growth is unlikely to be a major cause. Please name two models of each variety (four in total). Briefly explain why population growth matters/does not matter in each case.

Answer: Faster population growth raises the interest rate (i.e, population growth decline is a potential explatation for r* fact). This prediction will follow from a Standard Solow model (rule of thumb behavior vis-a-vis savings), or a Diamond model (life cycle savings), where productivity growth is exogenous. In the realm of endogenous growth models, the "Jones-model" (or semi-endogenous growth model) predicts that faster population growth leads to faster productivity growth, which should lead to a higher real rate of interest whether the savings motive is as decribed in Solow, Diamond or Ramsey-Cass-Koopmans-based growth models. Hence, to the extent population growth leads to faster productivity growth, a slow rate of population growth should work to lower the long-run real rate. Either if the savings motive is as decribed in the Solow or Diamond models, while growth is approximately exogenous, or, if Jones is right about the population/productivity growth nexus regardless of whether we are in a "Solow/Diamond or RCK world".
There is no effect from population growth in a Ramsey-Cass-Koopmans model, in which the savings motive is bequest and growth is exogenous. Here population growth has no effect on r*. In any endogenous growth model featuring positive scale effects will suggest that declining population growth only should imply a reduction in the speed of increase in r, not lead to a decline.

Question 3.d

Which of the models discussed in the previous question would you trust when it comes to explaining the declining real interest rate? Explain why? (This is an open question with no definitive answer. You answer will be evaluated based on the quality of your arguments).

Answer: Which to trust. If one is persuaded by Jones arguments (see question above), it would appear that population is a great candidate explanation for (some of) the decline in r* since it would be a robust prediction regardless of the precise savings motive (which is

Figure 1: Global population growth

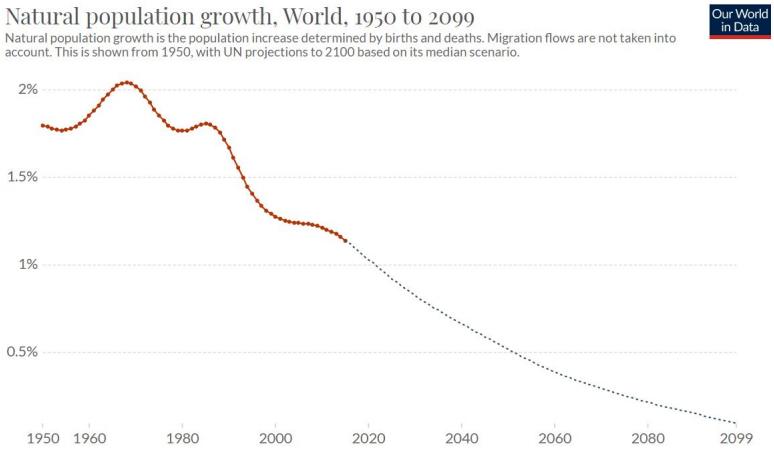

not known). Also, world population growth has in fact been declining over the period in question, which would seem to be a pre-requisite as the decline in r is pervasive across the world. (See Figure 1).