Exploring the Dynamic Interaction between Income and Health: Time-Series Evidence from Scandinavia

by

David Fielding[§] and Kalvinder Shields*

[First Draft: August 2001]

Abstract

This paper complements existing cross-section and panel data analyses of the interaction of income and health outcomes by applying a cointegrating VAR model of income and health to time-series data for several Scandinavian countries. The results are consistent with previous cross-section and panel results, but also highlight the complexity and heterogeneity of the dynamic relationships that generate them. In particular, there are substantial differences across countries in the relative importance of different underlying causes of the health-income correlation.

JEL Classification Numbers: I10, O52

[§] Department of Economics, University of Leicester. * Department of Economics, Melbourne University. Address for correspondence: Department of Economics, University of Leicester, Leicester LE1 7RH, UK. E-mail address: DJF14@LE.AC.UK.

1. Introduction

It has long been accepted that economic development is a key determinant of health outcomes in a country (see for example Preston, 1976), and that health is an important part of a country's stock of human capital and hence factor productivity (Fogel, 1994). There is a wide body of microeconometric evidence to support these relationships.[1] More recently, there have been a number of cross-country studies that aim to unpack the observed correlation between income and health outcomes at the macroeconomic level by using panel data techniques. These studies have produced a number of very interesting results, but the limitations of panel data mean that some important questions regarding the interaction of health and income remain unanswered. The aim of this paper is to complement such work – and to answer some of these questions – by employing macroeconomic time-series data for individual countries. By way of introduction, we will first review what has been discovered from panel data analysis, and what issues remain to be addressed.

Two recent papers deal with the problem of consistently estimating the impact of per capita income on health outcomes across countries. Pritchett and Summers (1996) use IV to estimate the impact of log per capita GDP on infant mortality and life expectancy in an international panel data set. Using the life expectancy measure, they find no significant relationship, perhaps because income affects this measure of health with very long lags. But using the infant mortality measure (infant deaths per 1000), they do find some significant effect. The estimated size of the effect varies greatly with the instrument(s) chosen, but is between –0.2 and –1.0; there is similar variation in the level of statistical significance of the estimate. Easterly (1999) uses a similar data set and methodology, but with different instruments, and produces estimates of income elasticity within the same range (and with similar significance levels) for both of the health indicators.[2]

Similarly, Bhargava et al. (2001) apply an IV estimator in order to quantify the impact of life expectancy and child mortality on log per capita GDP.[3] The panel data set is similar to that used in the Pritchett and Summers and Easterly studies. The impact of health on income is found to be statistically significant under all model specifications, although parameter estimates again vary substantially. In the Bhargava paper the regression equations are non-linear and the relationship between income and health is non-monotonic.

---

[1] On micoreconometric evidence for the importance of health as a component of human capital, see Basta et al. (1979), Spurr (1983), Bhargava (1997) and Strauss and Thomas (1998). On microeconometric evidence for the impact of income on nutrition and health, see Behrman and Deolalikar (1988) and Ravallion (1990).

[2] The Pritchett and Summers (1996) instrument set includes the terms of trade, the investment-output ratio and indicators of price distortions; the Easterly (1999) instrument set includes inflation, financial depth and indicators of price distortions.

[3] The dependent variable in the regressions in this paper is per capita GDP growth; but the lagged level of per capita GDP is on the LHS of the regressions, so the model is implicitly one of the level of per capita GDP.

The results in these papers demonstrate that there is a strong relationship in cross-section between a country's health performance and its per capita income, with causality running both ways. However, they leave a number of issues still to be addressed.

First, they deal largely with the factors determining the variations in income and health across countries: the time dimension in the data used is very small, and exhibits little variation, relative to the cross-sectional dimension. The estimated coefficients are to be interpreted as cross-country averages of the impact of one variable on another. The results say nothing about the size of effects within a particular country, or within a particular type of country: as the authors above acknowledge, it is necessary to aggregate over high-income and low-income countries in these panel data sets to generate enough sample variation for interpretable results. Since most of the variation is coming from LDCs, the panel results do not constitute very strong evidence on the relationship between health and income in OECD countries.

Secondly, the restricted time dimension in the panel data – five or ten-yearly intervals over a 30-year period – does not permit very detailed analysis of the dynamics of the relationship between income and health. It is not possible to produce very precise estimates of the length of time it takes for a shock to income (or health) to have its full impact on health (or income). Nor is a panel regression the appropriate vehicle for investigating the nature of shocks to income and health: for example, for determining how much of a typical shock is transitory and how much is permanent.

The main reason that these questions have not been pursued is the lack of reliable time-series data on income and health variables in most countries. Whereas macroeconomic time series are often reported at monthly or quarterly intervals, health data is usually reported annually; so time-series data covering a century or more are required for a reasonable sample size. Nevertheless, there are a few countries for which such data are available. In this paper we will exploit such time series in order to address the issues raised above.

The longest time series available are for three Scandinavian countries (Sweden, Finland and Denmark). In this paper we will explore the relationship between per capita income and health outcomes using time-series data in these countries for the period 1867–1997.[4] We will address the issues raised above by focusing on the following three questions.

(i) What fraction of innovations in health consists of permanent changes in health outcomes, and what fraction consists of transitory effects? In other words,

[4] Norway and Iceland do not appear in this study because of a lack of data. As in many other European countries, the Second World War disrupted the collection of economic data in these two countries, so we do not have continuous time series for more than about 50 years.

what is the degree of persistence in shocks to health? How does this compare with the degree of persistence in shocks to income?

(ii) To what extent, on average, have correlations between income and health over time been due to common shocks: wars, for example, or technological innovations that improve both factor productivity and the efficiency of health provision? (In a cross-section context, this question could only be addressed directly by finding proxies for technology, and it is by no means obvious that such proxies could ever be found. In a time-series context, the shocks can be modelled as stochastic innovations.) To what extent have the correlations been due not to common shocks but to causal associations – in both directions – over time? (In a time-series context, the causal associations can be interpreted in terms of Granger-causality.)

(iii) How does the magnitude of the association between health and income in a time series compare with the magnitude in a cross section? Does the high degree of cross-sectional correlation noted in previous studies correspond to a high degree of inter-temporal correlation?

In dealing with these questions, we will make note of any heterogeneity that appears across the three countries. This will tell us something about the degree of variation there is around the cross-country average effects estimated in existing panel data studies. There are a number of potential sources of heterogeneity. The three countries have very different political histories. Sweden experienced a gradual transition to full democracy over the period 1866-1917, and did not participate in either World War. Denmark experienced a similar gradual transition over the period 1850-1901, but it was occupied by Germany through 1939-45. Finland was occupied by Russia up to 1917, and has fought three wars with Russia since then. Various parts of Finland have been occupied by Russia at various times during the 20th century. The varying magnitude of (and violence associated with) these political changes may translate into larger shocks to economic and demographic variables, and to differently characterised dynamic interactions. Similarly, Finland and Sweden have lower population densities than Denmark, and agricultural products make up a larger fraction of their output and exports. The consequence of these differences in economic structure remains to be seen.

In Section 2 below we present the modelling framework, and in Section 3 variable definitions and data sources. Section 4 contains the results of the modelling exercise, and an interpretative discussion. Section 5 concludes.

2. The Modelling Framework

2.1 The underlying econometric model

One major advantage of adopting a time-series approach to modelling income and health is that we can obtain consistent estimates of the interactions between the two, without explicitly modelling their response to exogenous factors for which data might be lacking, by invoking the Wold Decomposition Theorem. Consider the case in which we have one measure of per capita income and one measure of health for $i = 1, 2, \ldots, m$ different countries. (In Section 4 below, $m = 3$.) It is possible that these $2m$ variables are all inter-related: shocks to income can be passed from one country to another through trade, and shocks to health through contagion. These inter-relations can be captured through a VAR model. Assuming that the variables are all difference-stationary, it is always possible to obtain an MA representation of their growth rates in the following form (Wold, 1983):

$$\Delta z_t = \mu + A(L)\varepsilon_t . \tag{1}$$

In this expression $\Delta z_t = (\Delta y_t, \Delta p_t)'$ represents a stacked $(2m \times 1)$ vector where $\Delta y_t = (\Delta y_{1t}, \Delta y_{2t}, \ldots\ldots, \Delta y_{mt})'$ is an $(m \times 1)$ vector containing values on income growth in each of the m regions in time t, $\Delta p_t = (\Delta p_{1t}, \Delta p_{2t}, \ldots\ldots, \Delta p_{mt})'$ is an $(m \times 1)$ vector containing values on the growth in the health indicator in each of the m regions in time t and Δ is the difference operator. $\mu = (\mu^y, \mu^p)'$ is a $(2m \times 1)$ vector, where $\mu^y = (\mu_1^y, \mu_2^y, \ldots\ldots, \mu_m^y)'$ and $\mu^p = (\mu_1^p, \mu_2^p, \ldots\ldots, \mu_m^p)'$ are both $(m \times 1)$ vectors and contain the mean values of income and health, respectively, in countries $i = 1, \ldots, m$. Similarly, $\varepsilon_t = (\varepsilon_{it}^y, \varepsilon_{it}^p)'$ for $i = 1, \ldots, m$ is a $(2m \times 1)$ vector of mean zero, serially uncorrelated innovations experienced by income growth (ε_{it}^y) and health improvements (ε_{it}^p) in country i at time t, with a covariance matrix Ω. In this multivariate model, $A(L)$ is a matrix polynomial given by

$$A(L) = \sum_{k=0}^{\infty} A_k L^k , \quad A_o = I_{2m} \tag{2}$$

and the (i,j)-th element of $A(L)$ is the lag polynomial $a_{ij}(L)$. Hence, for instance, in addition to the effects of current and past values of innovations on Δz_t in region i itself, income growth (health improvements) in country i may also be affected by past values of shocks to country j to income growth or health improvements. Moreover, there may be a systematic association between the occurrence of shocks in country i and those taking place elsewhere (captured by the non-zero off-diagonal elements of $\mathbf{W}$).

Expression (1) has a fundamental moving average representation, and, in general, this can be approximated by a finite order VAR model of the form

$$B(L)?z_t = ?z_t + B_1?z_{t-1} + B_2?z_{t-2} + \ldots\ldots\ldots + B_q?z_{t-q} = \boldsymbol{m}^* + \boldsymbol{e}_t, \tag{3}$$

where B_s, (s=1,2,... q) are (2m x 2m) matrices of coefficients, and the (i,j)-th element of B_s, denoted b_{yipjs}, relates to the coefficient on health improvements (denoted by p) in region j, lagged by s periods, in the equation explaining income growth (denoted by y) in region i. In this finite order VAR model, income growth in country i is explained by q lagged values of income growth in region i, q lagged values of health improvements in country i, plus q lagged values of income growth and health improvements in all other countries, and a random innovation, e_{it}^{y}; i.e.

$$?y_{it} = \mu_i^{y*} + \sum_{s=1}^{q} b_{yiyis}?y_{i,t-s} + \sum_{s=1}^{q} b_{yipis}?p_{i,t-s} + \sum_{s=1}^{q}\sum_{j\neq i} b_{yiyjs}?y_{j,t-s} + \sum_{s=1}^{q}\sum_{j\neq i} b_{yipjs}?p_{j,t-s} + e_{it}^{y} \tag{4}$$

If important interactions exist between the levels of z_t, the existing modelling framework can be readily adapted to allow for the presence of cointegrating relationships in the form of restrictions on the MA representation in (1). The error-correction form of (3) can be expressed as,

$$?z_t = \boldsymbol{m}^* + \sum_{s=1}^{q} \mathbf{G}_s?z_{t-s} - \mathbf{P}z_{t-q-1} + \boldsymbol{e}_t, \tag{5}$$

where □ is a (2m × 2m) reduced rank matrix determining the extent to which the system is cointegrated. The identification of the cointegrating vectors is discussed in Section 4 below.

2.2 Interpreting and measuring the persistence of shocks to income and health

The multi-country, multivariate VAR model presented above provides a flexible framework within which an analysis of income and health determination can be carried out. Most existing applications of this sort of VAR modelling framework are in the area of macro-econometrics.

Within this framework there are several ways of identifying the consequences of shocks to the system. In many macro-econometric applications the modeller has the confidence to impose a priori restrictions on the system, and to translate the estimated reduced-form shocks into a set of structural innovations, as in Blanchard and Quah (1989). But in some applications (as in our own) there is no theoretical ground for such restrictions, and other authors (for example Giacometti and Pinelli, 1999) chose not to impose a particular set of theoretical long-run restrictions on their model. Instead they explore the dynamics of their model through impulse response analysis. However, the application of impulse response analysis is not theoretically innocuous. The impulses to which the system's response is measured are orthogonalizations of the estimated reduced form innovations. These orthogonalizations (for example, Choleski decomposition) are not invariant to the ordering of the variables in the system. Implicit in the ordering is a theory about how the variables interact: in effect, a set of short-run restrictions.

We wish to avoid such restrictions, since our intention is to provide insights into the dynamic interaction of health and income variables rather than to identify an underlying structural model. Our analysis of the dynamics is conducted by constructing measures of generalized impulse responses.

Of particular interest are the long-run responses of the variables in z_t to shocks, and the dynamics of adjustment to the long run. Lee and Pesaran (1993) and Pesaran and Shin (1996) provide a framework for identifying the effects of specified types of shock. We can investigate the evolution of individual variables in response to shocks without resorting to a priori restrictions, by using generalized impulse response analysis. We will next provide a brief description of the measurement of the impact of shocks, showing how they may be used to construct measures of interest.

Specifically, referring to the multivariate, multi-country model described in equations (1-3), if e_r is a $(2m \times 1)$ selection vector with unity in its r^{th} element, and zeros elsewhere, then the generalized impulse response of any one variable k in the system to a "typical" shock to income ($j = 1$ £ r £ m) or health ($j = m + 1 \leq r \leq 2m$) in a particular country at time horizon N is given by:

$$G(j,k,N) = \left(\frac{e_k{}' A(N) \Omega e_j}{s_{jj}} \right) \tag{6}$$

where s_{jj}, the square root of the diagonal element of W for the j^{th} variable (i.e., the standard error of the j^{th} equation), is the magnitude of the initial shock. In equation (6) $A(0) = I_{2m}$, so

$G(j,j,0) = \sigma_{jj}$. As $N \to \infty$ we have a measure of the permanent effect of the shock. The persistence measures incorporate all of the interactions between variables in the system, insofar as they affect income or health in a country.

The $G(j,k,N)$ terms are in effect conditional expectation measures for the k^{th} variable at horizon N, given a "typical" shock to the j^{th} variable in the current period ($t = 0$). This typical shock is not necessarily orthogonal to other shocks in the system, because the off-diagonal elements of W are not necessarily equal to zero. This distinguishes generalized impulse responses from traditional impulse response measures. As a consequence, a $G(j,k,0)$ term will not necessarily be equal to zero, even when $j \neq k$: it reflects the degree of correlation of j and k, conditional on the history of the system.

The impulse response measures described above can be used to address questions (i-iii) in section 1 in the following ways.

(i) The ratio $G(j,j,\infty)/\sigma_{jj}$ – the magnitude of the asymptotic response of variable j in a particular country to a typical shock to that same variable, scaled by the magnitude of this initial shock – is a convenient measure of the degree of persistence of the shock. For difference-stationary variables we will expect this measure to be strictly positive, but we have no theoretical prior about its absolute size. As the measure approaches unity, we have a situation in which 100% of a typical shock to j persists in the long run. A high value for the persistence measure implies that most of a typical shock constitutes a permanent change in the variable. In the case of health, this means that any recent improvement (or deterioration) in health outcomes can be expected to last. When health improves (or deteriorates), it is largely a result of permanent changes in the socio-economic environment (for example, changes in health technology). Conversely, a small value of the persistence measure indicates that changes in health outcomes are largely – though not entirely – transitory phenomena. To put it another way, the measure of persistence indicates the magnitude of secular movements in standards of health relative to the magnitude of transitory fluctuations in standards. Comparing the persistence measure for health in each of the three countries with the persistence measure for income will provide some evidence on whether the degree of persistence in health is greater or smaller than the degree of persistence in macroeconomic variables.

(ii) The quantity $G(j,j+m,0)$ for $j = 1, \ldots, m$ measures the immediate change in health in each country that accompanies a typical shock to income there. If shocks to income and health are highly correlated, then this quantity will be larger.

Similarly, $G(j+m,j,0)$ measures the immediate change in income in each country that accompanies a typical shock to health there. If $G(j,j+m,0)$ and $G(j+m,j,0)$ are large relative to the corresponding asymptotic quantities $G(j,j+m,\infty)$ and $G(j+m,j,8)$, then we can conclude that the observed correlation between income and health is largely due to common shocks to the two variables, rather than a Granger-causal link between one variable and the other. However, if the asymptotic quantities are relatively large, then there is evidence that the observed correlation between income and health is at least partly the result of a dynamic interaction between the variables in the system.

(iii) The absolute values of the asymptotic quantities $G(j,j+m,\infty)$ and $G(j+m,j,8)$ indicate the magnitude of the long-run impact of a shock to one variable on the level of another. Although these quantities are not directly analogous to cross-sectional regression or correlation coefficients, they can nevertheless be cautiously compared with the results of previous cross-sectional studies, in order to give a sense of the magnitude of the inter-temporal association between health and wealth relative to its cross-sectional counterpart.

The discussion immediately above focuses on $G(\cdot,\cdot,0)$ and $G(\cdot,\cdot,8)$ measures: on immediate and asymptotic responses. However, the profile of the transition between the two, at finite positive values of N, can also provide useful information about the dynamic interaction between health and income. For this reason, the tables in section 4 listing persistence measures at $N = 0$ and N ? 8 will be accompanied by figures depicting the shape of the transition path between the two.

Moreover, the discussion thus far focuses on income and health interactions within a particular country. Although the persistence measures discussed above implicitly incorporate the interaction of variables across countries, more can be done to exploit the fact that we are estimating health and income equations for several countries within a single system. In particular, in the three-country case, we can calculate impulse responses following a universal shock to income by using the selection vector $e_j = (1\ 1\ 1\ 0\ 0\ 0)'$, and to a universal shock to health by using the selection vector $e_j = (0\ 0\ 0\ 1\ 1\ 1)'$. (In such cases our scaling factors, corresponding to the denominator in equation (6), will be $\Sigma_{j=1}^{j=m} s_{jj}$ and $\Sigma_{j=m+1}^{j=2m} s_{jj}$ respectively.) Similarly, system-wide responses to shocks can measured by using the selection vectors $e_k = (1\ 1\ 1\ 0\ 0\ 0)'$ for income and $e_k = (0\ 0\ 0\ 1\ 1\ 1)'$ for health. In this way, we can look at system-wide quantities corresponding to all the country-specific quantities discussed in (i-iii) above. So we will be able to complement our cross-country comparison of persistence measures with a description of pan-Scandinavian persistence measures.

3. Definition and Properties of the Data

3.1 Data sources and definition

The three countries on which this paper focuses (Sweden, Finland and Denmark) report consistent national accounts and demographic data, covering a limited range of variables, from as early as the mid-19th century.[5] The unusually long period of data coverage allows us to take a time-series approach to modelling the interaction of income and health. Our indicator of per capita income each year (y_t) will be the logarithm of real GDP at market prices minus the logarithm of the population. These two series are taken from the World Bank's World Development Indicators for the period 1960–96, and from Mitchell (1981) for earlier years. Consistent data are available for all three countries are available from 1867 at the latest, giving a sample size of 130.

The only health indicators that are available for such a long period of time are the crude death rate and the infant mortality rate (i.e., the proportion of infants dying within one year of birth). The disadvantage of using the crude death rate is that it is likely to respond to income with very long lags, since the death rate for the oldest cohort in the population is the consequence of health inputs over the last 70–80 years. Pritchett and Summers (1996) find no significant relationship between life expectancy and income in their panel data set, and the problem is likely to be even more severe when estimating a dynamic model of income and health using time-series data. We will therefore work with the infant mortality measures that are available. In order to ensure that the variable we are using (p_t) is not bounded, we will use a logistic transformation of the infant mortality series. That is:

$$p_t = \log(m_t) - \log(1 - m_t) \qquad (7)$$

where m_t is the ratio of infant deaths to the number of live births for each year. The six time series that will appear in our VAR (y_t and p_t for each country) are depicted in Figures 1–3. In the tables that follow, income and infant mortality in Sweden alone are designated as y_{st} and p_{st}; the corresponding designations for the Finnish and Danish variables are y_{ft}, p_{ft}, y_{dt} and p_{dt}.

[Figures 1–3 here]

3.2 Time-series properties of the data

Before estimating the VAR, it is necessary to ascertain the order of integration of the variables of interest. Augmented Dickey-Fuller (ADF) test statistics (not reported) confirm

---

[5] Comparable data for most other OECD countries begins much later. In the USA, for example, it begins only after the First World War.

the impression of Figures 1-3, that the null of difference-stationarity cannot be rejected for any of the variables. It is therefore appropriate to search for cointegration between the six variables, to see if the matrix **P** in equation (5) above has some strictly positive rank. There could be as many as five cointegrating vectors (stationary linear combinations) with six difference-stationary variables. With appropriate identifying restrictions, these could be interpreted as long-run relationships between the six variables that define the steady state of the system.

In samples as small as ours, tests for multivariate cointegration have low power, and anyway there is no obviously intuitive interpretation of a cointegrating vector in more than two of our variables. So we test for the existence of up to five bivariate cointegrating vectors: two linking together income across the three countries, two linking together infant mortality, and one linking income and infant mortality in one country (and therefore, by substitution of the other four vectors, linking income and infant mortality in the other two). Using the method of Engle and Granger (1987), which involves applying an ADF test to the residuals from a bivariate static regression in levels, we test for cointegration for the following pairs: $\{y_{st}, y_{dt}\}$, $\{y_{dt}, y_{ft}\}$, $\{p_{st}, p_{dt}\}$, $\{p_{dt}, p_{ft}\}$ and $\{y_{ft}, p_{ft}\}$. The null of no cointegration cannot be rejected for the first of these pairs, even at the 10% level; but the null can be rejected for the other four. (This result is not sensitive to the five pairs chosen to test for cointegration: whatever pairs are chosen, it always appears that all the variables except y_{st} are cointegrated with each other.)

Our VAR will therefore include four cointegrating vectors. The parameters in these vectors – the elements of **P** in equation (5) – are reported in Table 1. These are equal to the parameter estimates from the Engle-Granger regressions, except in the case of $\{p_{st}, p_{dt}\}$ where the intuitively appealing restriction of the parameter to unity cannot be rejected, even at the 10% level. Table 1 allocates a number to each vector ("cv1" to "cv4") for reference in Table 2, which reports the full VAR estimates.

[Tables 1-2 here]

4. Results

4.1 Characteristics of the fitted VAR

We do not report the unrestricted estimates of equation (5), but rather a restricted model that includes just that set of lags of each variable in each equation that minimizes the Schwartz-Bayesian Information Criterion. χ^2 test statistics for the validity of the restrictions on a second-order VAR are reported in Table 2: in no equation are the restrictions rejected. The cointegrating vectors are jointly significant at the 5% level in each equation except that for

Δy_{st}, confirming the cointegration of five out of the six variables. Coefficients on the cointegrating vectors are consistent with the long-run stability of the system.

The R^2 and σ_{jj} values reported in Table 2 indicate that there is some heterogeneity across the equations in terms of goodness-of-fit and the size of a typical innovation. Shocks to infant mortality appear to be larger on average than shocks to income. Shocks in Finland are larger than shocks in Denmark, which are larger than shocks in Sweden. The LM test statistics for residual autocorrelation and for heteroscedasticity that are reported in the table are not significant at the 5% level.

Given the non-linearity of some panel data regressions of income on health (for example, those in Bhargava et al., 2001), we also test for the validity of the functional form of each equation using a RESET test. The test statistics are reported in Table 3: in no case can the null of validity be rejected at the 5% level.

[Table 3 here]

The individual coefficients in Table 2 do not have a straightforward individual interpretation, so the next section explores the characteristics of the estimated model by reporting and discussing generalized impulse response measures and impulse response profiles.

4.2 Generalized impulse responses in the system

In Sections 1 and 2 above, we raised three general questions that could be addressed through the estimation of our cointegrating VAR. In this section, we will use our estimates to address each in turn.

(i) The first question concerns the degree of persistence in shocks to each variable, as measured by the quantity $G(j,j,\infty)/\sigma_{jj}$.[6] The left hand side of Table 4 reports these quantities for each variable in each of the three countries, plus the corresponding quantities for a universal shock to each variable in the system. The table indicates a substantial degree of heterogeneity across the three countries. 89% of a shock to Swedish income and 52% of a shock to Swedish infant mortality persists in the long run. At the opposite extreme, the corresponding figures for Finland are 24% and 3%. The figures for Danish shocks and for universal shocks are between these two extremes. In all cases, the degree of persistence in income is greater than the degree of persistence in infant mortality. Sweden has experienced relatively small shocks compared with Denmark and Finland, but the degree of persistence in these shocks has been a little greater than in Denmark, and much greater than in Finland.

[6] In Tables 4-5 we approximate the infinite horizon as 40 years. The generalised impulse response profiles in Figures 4-7 indicate that this is quite a close approximation.

The unscaled persistence measures $G(y,y,\infty)$ and $G(p,p,\infty)$ on the left hand side of Table 5 indicate that the greater degree of persistence in Sweden outweighs the fact that the initial shocks are smaller than in the other two countries: out of all of the three countries, Sweden has the largest values of $G(y,y,\infty)$ and $G(p,p,\infty)$. The standard errors in Table 5 indicate that the difference between the two extremes (Sweden and Finland) is significant at the 5% level. The generalized persistence profiles in Figures 4-7 indicate that most of these differences are evident within 10 years of the shock. There is little movement in most of the profiles after the first 10 years.

[Tables 4-5 and Figures 4-7 here]

Within the VAR framework that we are using it is not possible to identify directly the reasons for this heterogeneity. However, it is not surprising that the most populous country (Sweden) experiences the smallest shocks and the least populous country (Finland) experiences the largest shocks: we should expect a higher variance in a smaller population, unless shocks to individuals are perfectly correlated. Moreover, there are clear reasons why Sweden should exhibit the most persistence: unlike the other two countries, it has not been involved in any major international conflict, nor has it been occupied by a foreign power. The relatively adverse conditions during international conflicts are likely to represent a less persistent type of shock than technological innovations and changes in productivity. But the magnitude of the differences between the countries is still remarkable.

It is also noteworthy that across all the countries shocks to infant mortality are larger than shocks to income, but less persistent. To use some macroeconomic jargon, health "business cycles" are shorter but more extreme than economic ones. Macroeconomic stabilization policy typically focuses on the size cyclical movements in income and inflation. But macroeconomic variables may in fact exhibit less extreme cycles than health indicators. Our results support Sen (1998), who argues that health indicators are not necessarily less sluggish than macroeconomic ones, and therefore no less appropriate as measures of economic performance, even in the short run.

(ii) The second question concerns the relative importance of two reasons for a correlation between income and health: on the one hand, large common shocks to the two variables, and on the other, substantial dynamic interaction between them. Tables 4-5 and Figures 4-7 show that here, too, there is substantial heterogeneity across the three countries.

At one extreme lies Denmark. In Denmark there is no significant correlation between innovations in per capita income and innovations in infant mortality. (In fact the point estimate of the correlation coefficient is positive, but only a fraction one standard deviation from zero. Table 6 lists the values of all the correlation coefficients for the innovations in the six equations.) This is illustrated by the bottom two graphs in Figure 7, in which the impulse response profiles G(y,p,N) and G(p,y,N) begin very close to the zero line. However, the unconditional sample correlation between the two variables is -0.98. This negative correlation is explained by the dynamics illustrated in the two graphs. A positive shock to y_d will in subsequent periods lead to a lower p_d; similarly, a positive shock to p_d will in subsequent periods lead to a lower y_d. At the 20-year horizon, both effects are just about significant at the 5% level, although at the infinite horizon the standard error on G(p,y,8) is slightly too high to register statistical significance using conventional confidence intervals. The asymptotic effects and their associated standard errors are listed on the righthand side of Table 5.

[Table 6 here]

At the opposite extreme lies Finland. In Finland there is a negative correlation between innovations in per capita income and innovations in infant mortality, and G(y,p,N) and G(p,y,N) are significantly below zero for very small values of N, as shown in the bottom two graphs in Figure 6. However, these negative effects persist very little, and for all values of N greater than eight they are insignificantly different from zero. The dynamics of y_f and p_f dampen down the effects of common shocks pushing income up and infant mortality down. This is emphasised by the figures on the righthand side of Table 4, which show the ratios of the first point on each impulse response profile to the last point. In the case of Finland, the asymptotic measures are only a fraction of the size of the initial effects. The negative unconditional correlation coefficient between the two variables (-0.99) is a largely a consequence of the common shocks.

The figures for shocks to Sweden, and for shocks to all of the countries, represent an intermediate case, as shown in Figures 4-5 and in Table 4. There is a negative correlation between innovations in per capita income and innovations in infant mortality, and the dynamic interaction between the two variables magnifies this effect. The estimated magnification effects for Sweden (shown in Table 4) are 1.33 for G(y,p,N) and 2.76 for G(p,y,N). The figures for shocks to all countries are a little larger.

So there is no straightforward answer to the question of how much the negative correlation between per capita income and infant mortality is due to common shocks, and how much it is due to one variable Granger-causing the other. Even within Scandinavia, there is considerable heterogeneity in the extent to which one factor or the other is more important. This suggests that the results of papers based on IV estimates of the interaction between income and health using panel data are represent average effects around which there is a great deal of international variance.

(iii) Any comparison of panel regression results with those presented here must therefore be interpreted with a great deal of caution. Nevertheless, the asymptotic generalized impulse response measures for the effect of shocks to per capita income on infant mortality are broadly in line with the Pritchett and Summers (1996) and Easterly (1999) results discussed in Section 1. For a comparison with the elasticities estimated in these papers, we can construct asymptotic generalized impulse response measures corresponding to the G (y,p,8) figures in Table 5 (these capture the long-run effects on infant mortality of a typical shock to income), but with unit shocks to per capita income instead of one standard error shocks. The resulting figures are −0.43 for Sweden, −0.35 for Finland, −0.61 for Denmark and −0.75 for a shock to all countries.[7]

5. Conclusion

Using a cointegrating VAR framework, we have been able to identify the characteristics of the dynamic interaction between per capita income and infant mortality in three Scandinavian countries: Sweden, Finland and Denmark. Although there is a negative association between the two variables in all three countries, there is a considerable degree of heterogeneity in the dynamics that underlie this association. In Finland, and to a lesser extent in Sweden, the negative correlation is largely a result of common contemporaneous shocks to both income and health. However, in Denmark the shocks are orthogonal, and the negative correlation is entirely due to the fact that each variable is Granger-caused by the other. Moreover, there is substantial heterogeneity in the characteristics of typical shocks: Swedish shocks are the smallest, but have the greatest degree of persistence; Finnish shocks are the largest, but have the least persistence. The such a large degree of heterogeneity should be manifested even within Scandinavia suggests that the results of cross-country panel data studies of the interaction of income and health represent average effects around which there is likely to be a great deal of variance. This is an important caveat if these studies are used as an input in policy decisions in individual countries.

The other main result from the VAR model is that shocks to infant mortality are larger but less persistent than shocks to per capita income. This stylized fact is true of all the countries. This suggests that health measures are not a “sluggish” indicator of economic performance: if

[7] The caveat is that we have used a logistic transformation of the infant mortality series, but the other two papers have not. We do not attempt a comparison with the results of the paper by Bhargava et al. (2001) discussed in Section 1 because the results of that paper are embodied in a model with a non-monotonic relationship between income and health.

anything, they are less sluggish than macroeconomic variables. It also suggests that policy-makers could raise social welfare by extending the focus of stabilization policy beyond the narrowly economic domain.

References

Basta, S., Soekirman, M., Karyadi, D. and Scrimshaw, N. (1979) "Iron Deficiency Anemia and the Productivity of Adult Males in Indonesia, American Journal of Clinical Nutrition, 32, 916–25.

Behrman and Deolalikar (1988) "Health and Nutrition", in Chenery, H. and Srinivasan, T. (eds.), Handbook of Development Economics (Volume 1), Amsterdam: North-Holland.

Bhargava, A. (1997) "Nutritional Status and the Allocation of Time in Rwandese Households", Review of Economic Studies, 77, 277–95.

Bhargava, A., Jamison, D., Lau, L. and Murray, J. (2001) "Modeling the Effects of Health on Economic Growth", Journal of Health Economics, 20, 423–40.

Blanchard, O. and Quah, D. (1989) "Dynamic Effects of Aggregate Demand and Supply Disturbances", American Economic Review, 79, P655–73.

Lee, K. and Pesaran, M. (1993) "Persistence Profiles and Business Fluctuations in a Disaggregated Model of UK Output Growths", Richerche Economiche, 47, 293–322.

Easterly, W. (1999) "Life During Growth", manuscript, World Bank.

Engle, R. and Granger, C. (1987) "Cointegration and Error Correction: Representation, Estimation and Testing", Econometrica, 55, pp. 251–76.

Giacometti, R. and Pinelli, D. (1999) "Asymmetric Shocks and Long-Run Economic Performances Across Italian Regions", unpublished manuscript, University of Bergamo, Bergamo, Italy.

Mitchell, B. (1981) European Historical Statistics 1750–1975 (Second Revised Edition), London: Macmillan

Pesaran, H. and Shin, Y. (1996) "Cointegration and the Speed of Convergence to Equilibrium", Journal of Econometrics, 71, 117–43.

Pritchett, L. and Summers, L. (1996) "Wealthier is Healthier", Journal of Human Resources, 31, 841–68.

Ravallion, M. (1990) "Income Effects on Undernutrition", Economic Development and Cultural Change, 38, 489–516.

Sen, A. (1998) "Mortality as an Indicator of Economic Success and Failure", Economic Journal, 108, 1–25.

Spurr, G. (1983) "Nutritional Status and Physical Work Capacity", Yearbook of Physical Anthropology, 26, 1–35.

Strauss, J. and Thomas, D. (1998) "Health, Nutrition and Economic Development", Journal of Economic Literature, 36, 766–817.

Wold, H. (1983) A Study in the Analysis of Stationary Time Series, Uppsala: Almqvist and Wiksells.

Table 1: The Cointegrating Vectors

$p_s - 1.000 \cdot p_d$ (cv1)	$p_d - 0.889 \cdot p_f$ (cv2)
$y_d - 0.800 \cdot y_f$ (cv3)	$y_f + 0.703 \cdot p_f$ (cv4)

Table 2: The Fitted VAR Model

Dy_s	coeff.	s.e.	t-ratio
cons	0.0206	0.0033	6.1714
$\Delta y_f(1)$	0.1936	0.0543	3.5626
$\Delta p_s(1)$	0.0773	0.0334	2.3108
$\Delta p_d(1)$	-0.0529	0.0306	-1.7274
$\Delta y_s(2)$	-0.2322	0.0926	-2.5070
$\Delta y_f(2)$	0.1314	0.0600	2.1898

Dp_s	coeff.	s.e.	t-ratio
cons	-0.0512	0.0087	-5.9235
$\Delta p_s(1)$	-0.1983	0.1008	-1.9673
$\Delta p_f(1)$	0.1414	0.0607	2.3304
$\Delta p_d(1)$	-0.2921	0.0910	-3.2088
$\Delta y_s(2)$	0.7573	0.2231	3.3936
$\Delta p_f(2)$	0.1282	0.0534	2.3994
cv1(1)	-0.1782	0.0686	-2.5957
cv2(1)	0.0321	0.0593	0.5423
cv3(1)	-0.0810	0.0839	-0.9651
cv4(1)	-0.2118	0.0694	-3.0511

Dy_f	coeff.	s.e.	t-ratio
cons	0.0137	0.0046	2.9698
$\Delta y_s(1)$	0.2783	0.1400	1.9872
$\Delta y_f(1)$	0.3493	0.0947	3.6889
$\Delta y_d(1)$	-0.1781	0.1144	-1.5568
$\Delta p_s(1)$	0.1494	0.0576	2.5940
$\Delta p_d(1)$	-0.1188	0.0524	-2.2660
cv1(1)	-0.0697	0.0384	-1.8159
cv2(1)	-0.1101	0.0324	-3.3990
cv3(1)	0.0920	0.0475	1.9359
cv4(1)	-0.1543	0.0382	-4.0424

Dp_f	coeff.	s.e.	t-ratio
cons	-0.0700	0.0143	-4.8929
$\Delta y_s(1)$	0.5554	0.3465	1.6027
$\Delta y_f(1)$	-0.5493	0.2404	-2.2846
$\Delta p_f(1)$	-0.1981	0.0841	-2.3560
$\Delta p_d(1)$	-0.3917	0.1344	-2.9147
$\Delta y_s(2)$	0.6255	0.3344	1.8707
$\Delta p_d(2)$	-0.2664	0.1181	-2.2551
cv1(1)	-0.1988	0.0926	-2.1461
cv2(1)	0.3936	0.0910	4.3257
cv3(1)	-0.1713	0.1224	-1.3988
cv4(1)	-0.2053	0.0983	-2.0892

Dy_d	coeff.	s.e.	t-ratio
cons	0.0163	0.0046	3.5797
$\Delta y_s(1)$	0.2241	0.1221	1.8355
$\Delta y_f(1)$	0.1128	0.0826	1.3651
$\Delta y_d(1)$	-0.0389	0.0989	-0.3934
$\Delta p_s(1)$	0.1738	0.0550	3.1575
$\Delta p_d(1)$	-0.1899	0.0516	-3.6802
$\Delta y_d(2)$	-0.1989	0.0862	-2.3086
$\Delta p_s(2)$	0.0770	0.0483	1.5927
$\Delta p_d(2)$	-0.0821	0.0467	-1.7579
cv1(1)	-0.1053	0.0363	-2.8975
cv2(1)	-0.0220	0.0293	-0.7521
cv3(1)	-0.0759	0.0424	-1.7910
cv4(1)	-0.1204	0.0347	-3.4695

Dp_d	coeff.	s.e.	t-ratio
cons	-0.045	0.0096	-4.6968
$\Delta y_s(1)$	0.9552	0.2490	3.8358
$\Delta y_d(1)$	-0.2263	0.2084	-1.0863
$\Delta p_f(1)$	0.1968	0.0681	2.8915
$\Delta p_d(1)$	-0.5195	0.1010	-5.1409
$\Delta p_s(2)$	0.1470	0.1036	1.4182
$\Delta p_f(2)$	0.1411	0.0621	2.2741
$\Delta p_d(2)$	-0.1950	0.1014	-1.9235
cv1(1)	0.0429	0.0697	0.6158
cv2(1)	0.0061	0.0668	0.0917
cv3(1)	-0.0353	0.0911	-0.3870
cv4(1)	-0.1765	0.0724	-2.4390

equation	R^2 value	s value	T1§	T2§	T3§
Δy_s	0.1764	0.0274	$\chi^2(11) = 7.776$	$F(1,122) = 0.241$	$F(1,127) = 1.401$
Δp_s	0.3160	0.0706	$\chi^2(7) = 3.744$	$F(1,118) = 1.600$	$F(1,127) = 0.072$
Δy_f	0.2959	0.0398	$\chi^2(7) = 4.035$	$F(1,118) = 0.683$	$F(1,127) = 0.026$
Δp_f	0.4060	0.1037	$\chi^2(6) = 3.295$	$F(1,117) = 0.203$	$F(1,127) = 0.858$
Δy_d	0.2407	0.0344	$\chi^2(4) = 0.871$	$F(1,115) = 0.350$	$F(1,127) = 0.943$
Δp_d	0.3787	0.0760	$\chi^2(5) = 4.060$	$F(1,116) = 3.035$	$F(1,127) = 0.583$

§ T1: test for validity of the restrictions imposed on each equation. T2: test for residual autocorrelation. T3: test for heteroscedasticity.

Table 3: Tests for Functional Form Misspecification

	y equation	p equation
sweden:	$F(1,122) = 2.2645[0.87]$	$F(1,115) = 1.3520[0.75]$
finland:	$F(1,118) = 0.6829[0.59]$	$F(1,118) = 0.0689[0.21]$
denmark:	$F(1,118) = 0.0255[0.13]$	$F(1,117) = 2.8990[0.91]$

Table 4: Asymptotic Generalized Impulse Responses Scaled by the Size of Initial Shocks

	y on y	p on p	y on p	p on y
	$G(y,y,8)/s_{yy}$	$G(p,p,8)/s_{pp}$	$G(y,p,8)/G(y,p,0)$	$G(p,y,8)/G(p,y,0)$
system	0.591924	0.387790	2.650678	4.026683
sweden	0.890943	0.522224	1.334088	2.756764
finland	0.237127	0.031769	0.783585	0.412821
denmark	0.389766	0.228618	-3.23411	-3.73881

Table 5: Unscaled Asymptotic Generalized Impulse Responses with Standard Errors

	coeff.	s.e.		coeff.	s.e
G(y,y,8) measures			G(y,p,8) measures		
y on y (system)	0.048417	0.013755	p on p (system)	0.073895	0.015442
y on y (sweden)	0.025587	0.002236	p on p (sweden)	0.038009	0.010974
y on y (finland)	0.009731	0.007738	p on p (finland)	0.003383	0.010856
y on y (denmark)	0.013650	0.006585	p on p (denmark)	0.017762	0.012091
G(p,p,8) measures			G(p,y,8) measures		
y on p (system)	-0.061562	0.027435	p on y (system)	-0.040142	0.007793
y on p (sweden)	-0.012371	0.007151	p on y (sweden)	-0.010087	0.002709
y on p (finland)	-0.014273	0.011099	p on y (finland)	-0.002898	0.007609
y on p (denmark)	-0.021316	0.011251	p on y (denmark)	-0.011108	0.007024

Table 6: Innovation Correlations

	y_s	y_f	y_d	p_s	p_f	p_d
y_s	1.000000	0.419785	0.388139	-0.127402	-0.066738	0.007296
y_f	0.419785	1.000000	0.412337	0.012716	-0.171054	-0.063713
y_d	0.388139	0.412337	1.000000	-0.074494	-0.153903	0.084839
p_s	-0.127402	0.012716	-0.074494	1.000000	0.335972	0.453681
p_f	-0.066738	-0.171054	-0.153903	0.335972	1.000000	0.199424
p_d	0.007296	-0.063713	0.084839	0.453681	0.199424	1.000000

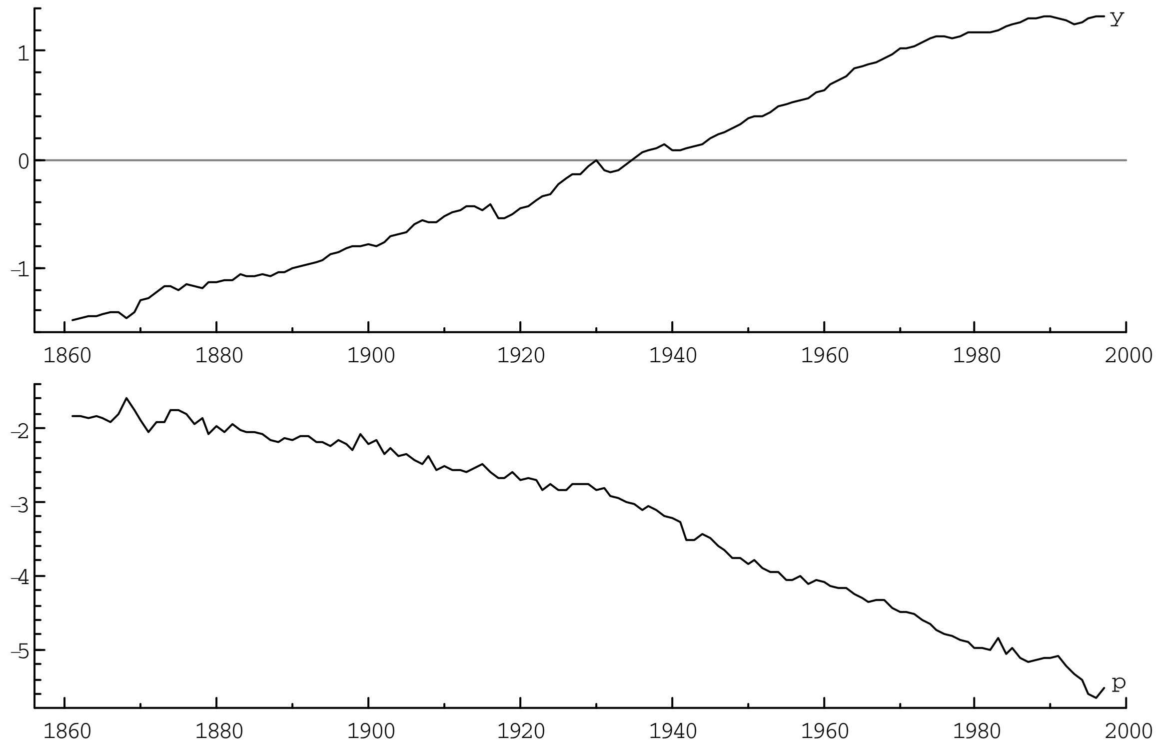

Figure 1: Time-series for Sweden

(y = log of real per capita GDP and p = logistic of infant mortality rate)

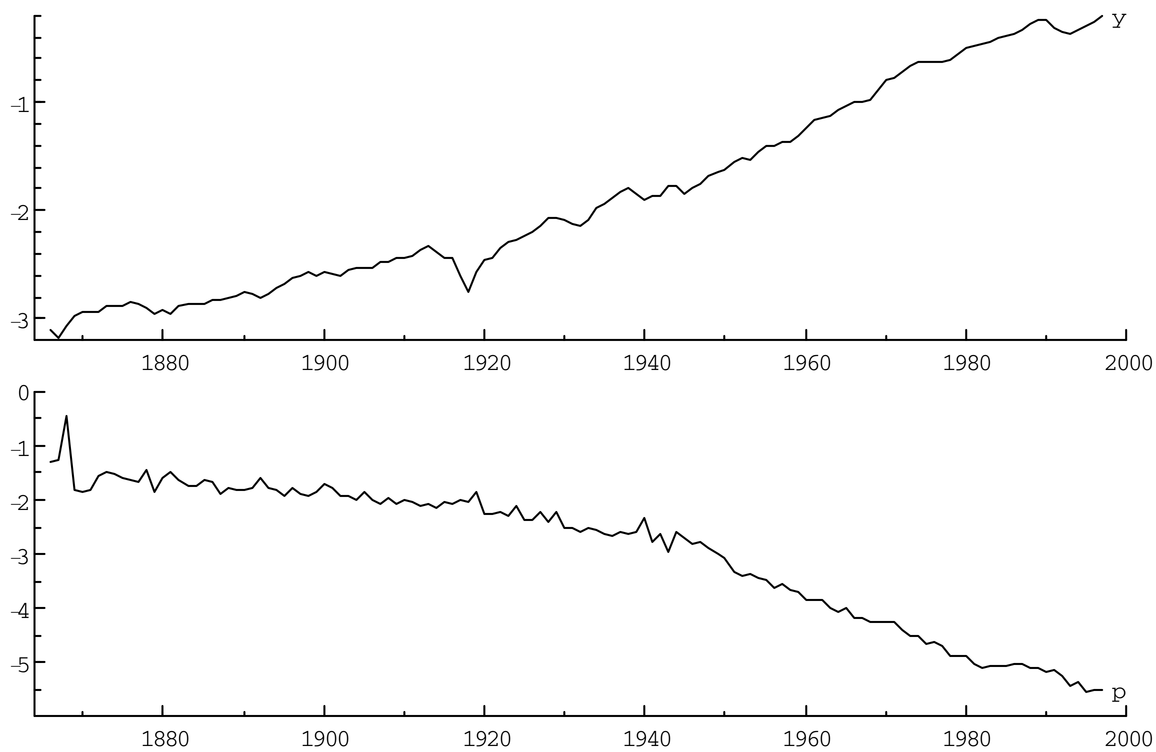

Figure 2: Time-series for Finland

(y = log of real per capita GDP and p = logistic of infant mortality rate)

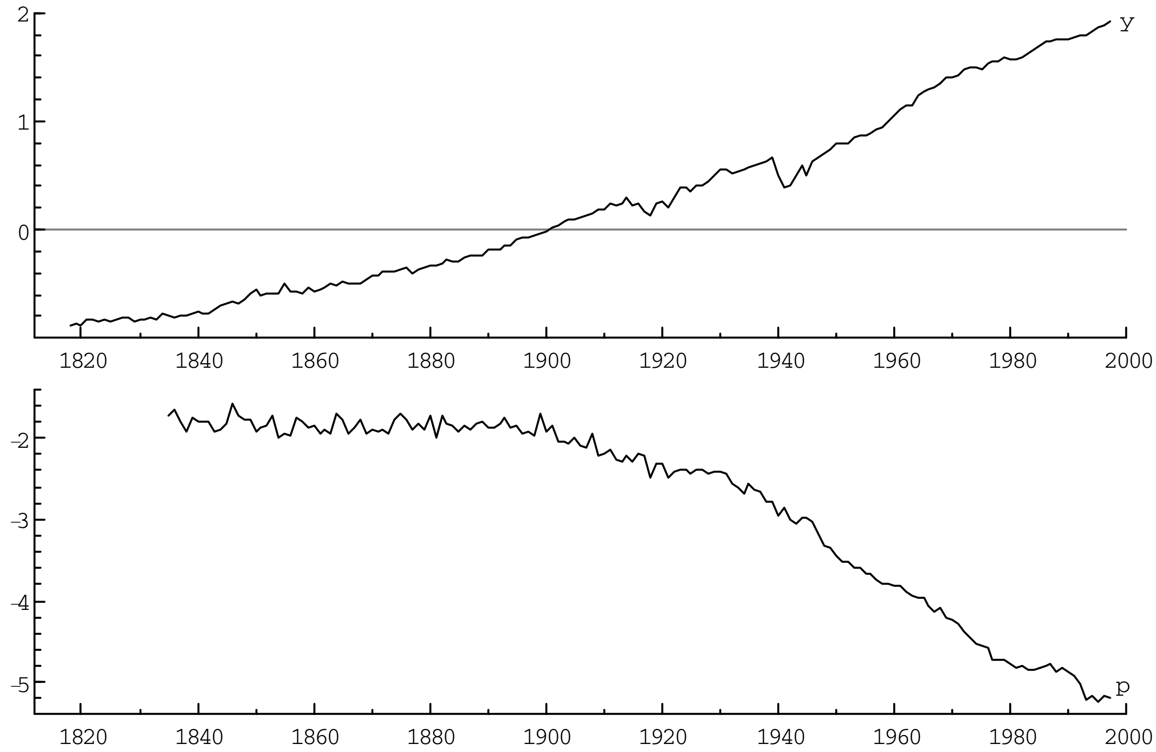

Figure 3: Time-series for Denmark

(y = log of real per capita GDP and p = logistic of infant mortality rate)

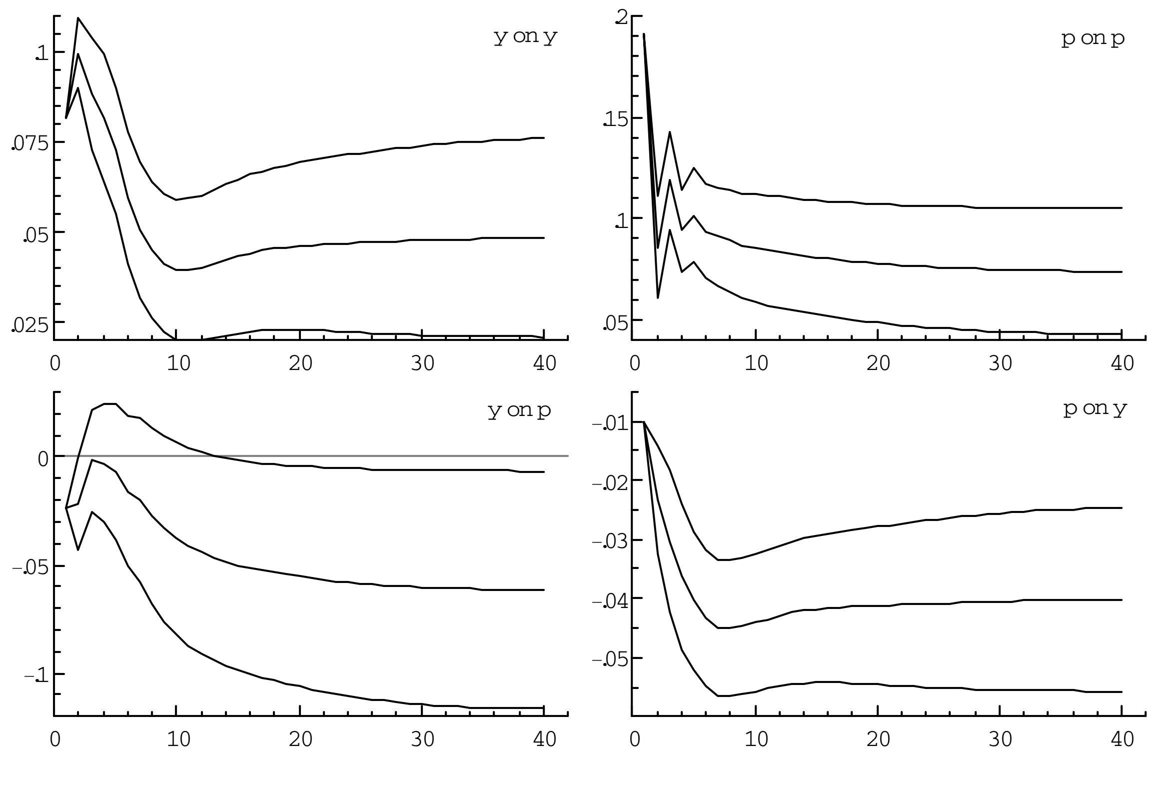

Figure 4: Generalized Impulse Response Profiles for the Whole System ± 2 Standard Errors

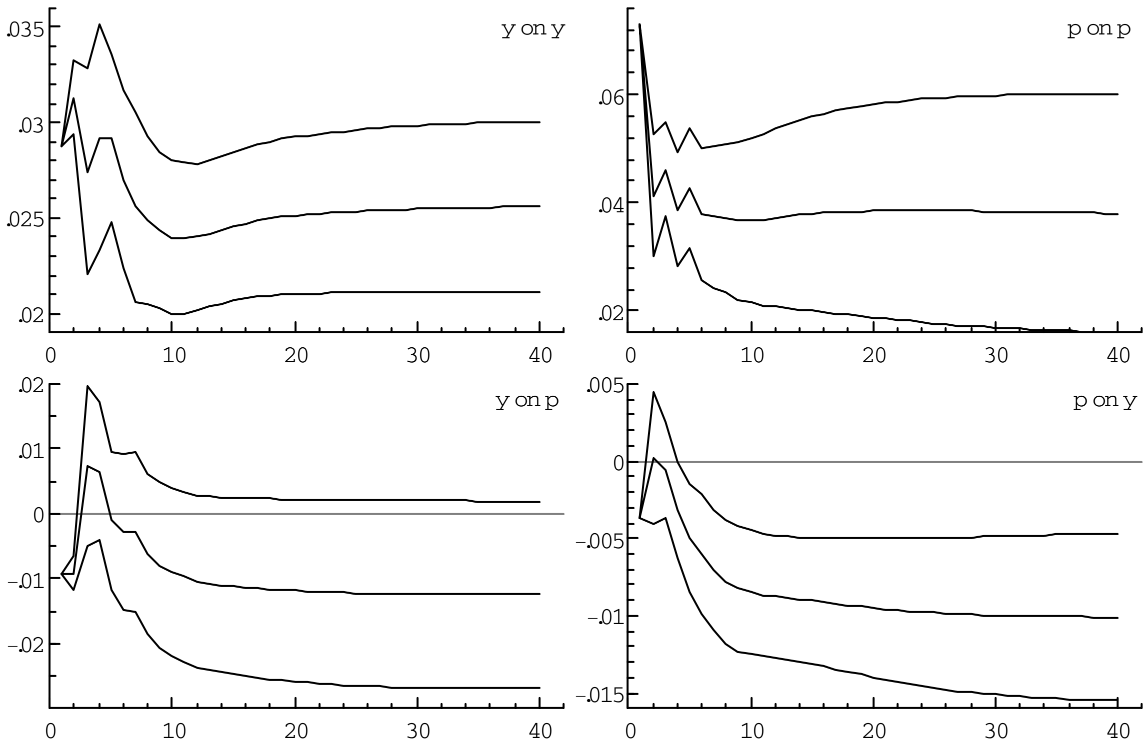

Figure 5: Generalized Impulse Response Profiles for Sweden ± 2 Standard Errors

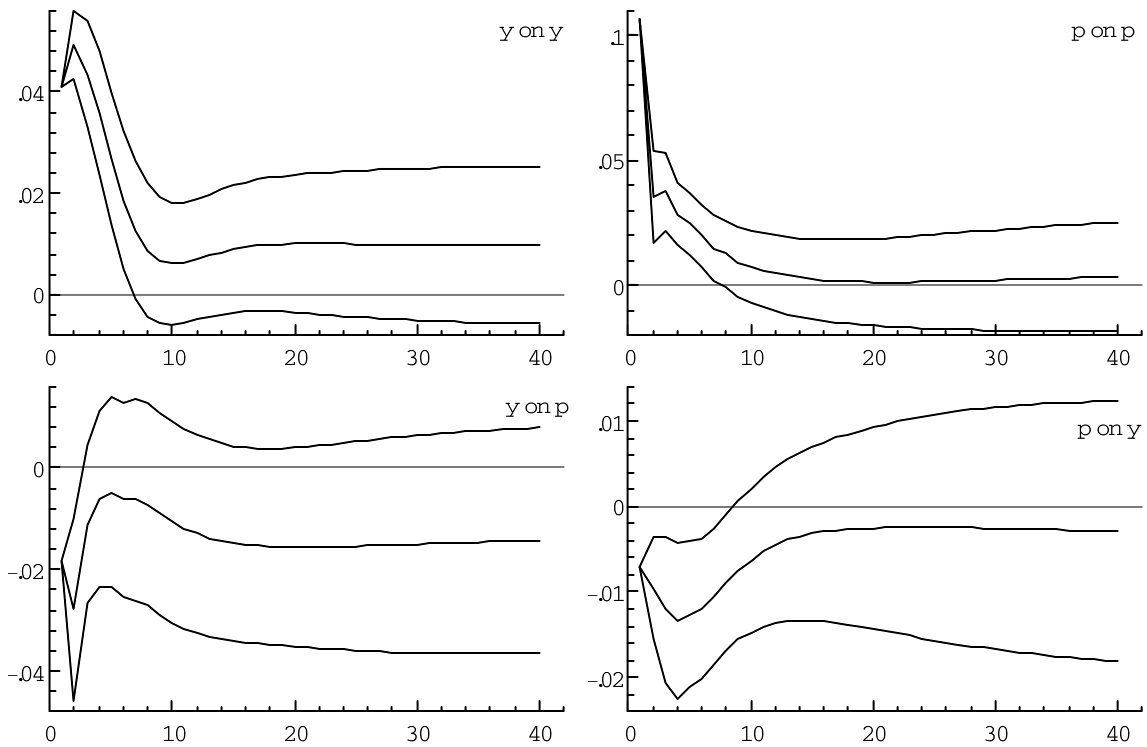

Figure 6: Generalized Impulse Response Profiles for Finland ± 2 Standard Errors

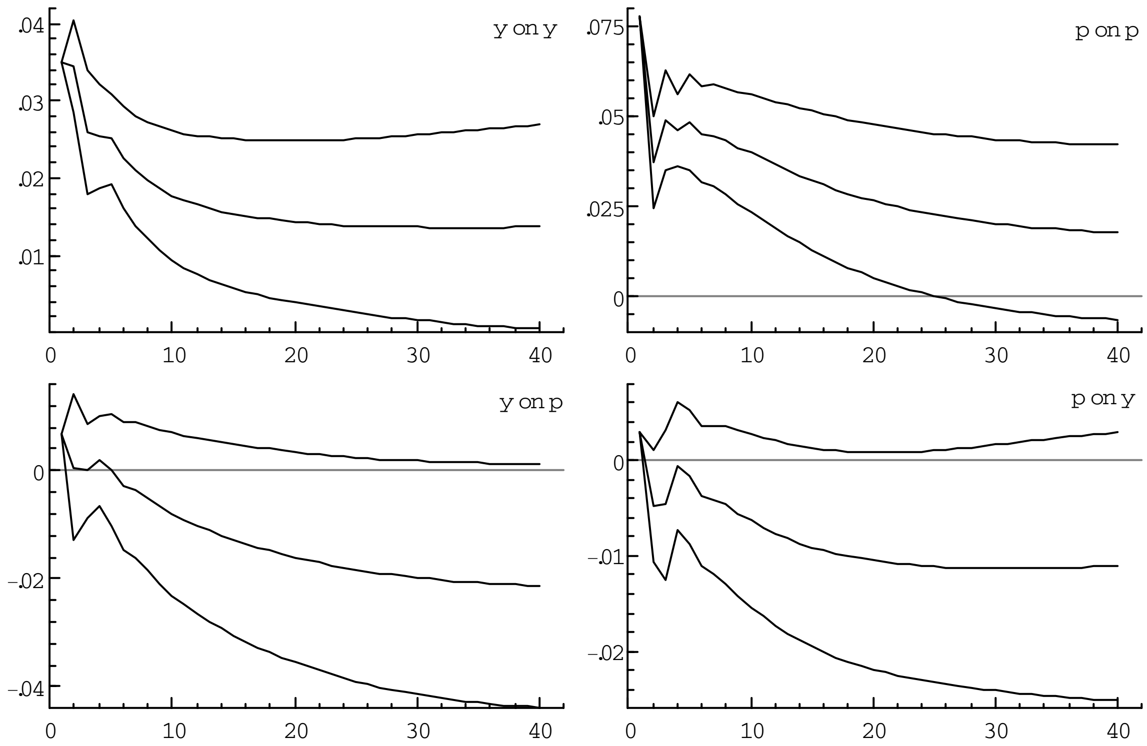

Figure 7: Generalized Impulse Response Profiles for Denmark ± 2 Standard Errors