Papers in Evolutionary Economic Geography

12.15

The Resilience of Dutch Regions to Economic Shocks. Measuring the relevance of interactions among firms and workers.

Dario Diodato and Anet Weterings

Utrecht University
Urban & Regional research centre Utrecht

The Resilience of Dutch Regions to Economic Shocks

Measuring the relevance of interactions among firms and workers

Dario Diodato and Anet Weterings

Netherlands Environmental Assessment Agency
The Hague, the Netherlands
P.O. Box 30.314, 2500 GH The Hague, The Netherlands
Email: Dario.Diodato@pbl.nl and Anet.Weterings@pbl.nl

Abstract

Although increasing attention is paid to the resilience of regions to economic shocks, theoretical and empirical insights in the determinants of regional resilience are still limited. This paper aims to make a first step in quantifying regional resilience. Using a model, we explore how three regional factors jointly contribute to the resilience of regions to economic shocks: 1) the network of buyer-supplier relationships within and between regions, 2) the level of relatedness between industries, which facilitates intersectoral labor mobility and, 3) the geographical position of a region which determines the possibilities of commuting for workers. The supply network mainly determines the propagation of the shock, while possibilities for intersectoral and interregional labor mobility affect a regional economy's capacity to recover from the shock. To illustrate the workings of the model, it is applied to the case of the Netherlands using data on buyer-supplier relationships within and between Dutch regions, as well as on intersectoral and interregional labour mobility.

Keywords: regional resilience, input-output network, labor mobility, related labor flows, commuting flows, the Netherlands

JEL classification: J61, O18, R11

1. Introduction

Both the recent economic recession and the burst of the dotcom bubble showed that not every region is affected in the same way by a sudden economic shock (see Holm & Østergaard 2010; Martin 2012). The same goes for local economic shocks, such as the closing down of a large company. Holm et al. (2012) illustrate this for the closing down of four shipyards spread across Denmark. Laid-off employees were confronted with very different possibilities for finding a new job depending on where the shipyard was located. This raises the question how it is possible that some regions are better capable of coping with economic shocks than others, that is, why regions differ in their resilience to shocks.

Although an increasing number of studies explores the idea of economic resilience of regions in economic geography (see for instance the special issue of Cambridge Journal of Regions, Economy and Society, vol. 3, 2010), theoretical and empirical insights in which factors affect a region's resilience to economic shocks are still very limited. Until now, most studies focus on discussing the possible application of the concept of resilience to regional differences in economic growth and the identification of the nature and scale of regional differences to economic shocks (e.g., Simmie & Martin 2010; Martin 2012; Fingleton et al. 2012). Most studies do mention the possible factors which may affect a region's resilience to economic shocks, but only a few also empirically examine the relevance of those factors (Holm & Østergaard 2010; Glaeser et al. 2011). The explanation of regional resilience is a highly complex issue as many different factors play a role and resilience tends to change over time. Nevertheless, for regional policymakers, a better understanding of the determinants of regional resilience is crucial.

For these reasons, the main aim of this paper is to make a first step in quantifying regional resilience. Following Martin (2012), we argue that resilience refers to the capacity of a region to resist a shock, as well as the speed with which it can recover from the shock. The resilience in these two phases, the initial shock and the recovery, depends on different regional attributes. We focus on the sectoral composition of regions, that is, the mix of sectors present in the region. Inspired by the ideas on regional smart specialisation policy (McCann & Ortega-Argilés 2011), we identify three determinants of resilience: embeddedness, relatedness and connectivity. In the first phase, resilience depends on the mix of activities present in a region and the buyer-supplier relationships among these activities. If sectors are embedded within the region through local buyer-supplier relationships, a shock hitting a central actor can propagate through the supply chain and affect the whole local economy. In the second phase, what determines the region's ability to recover from the shock is the speed of adaptation. We associate this to the possibilities that a region offers to laid-off employees to find a new job. Here, the relatedness between industries and connectivity between regions matter. Employees who lost their job due to the shock will be better able to find a new one when other sectors in the region require similar kind of skills as the sector where they used to work, that is, when industries in the region are more related. Alternatively, laid-off workers also have an advantage when neighbouring regions, at commuting distance, offer a range of jobs that match their skills.

Following this reasoning, we have composed a two-stage model to explore how the three different determinants (buyer-supplier relationships, intersectoral and interregional labour mobility) jointly contribute to the resilience of regions to economic shocks. To illustrate the workings of the model, we apply it to the case of the Netherlands. We gathered empirical data for 12 Dutch regions (NUTS-2 level) and 59 sectors (2-digit NACE-codes) on regional buyer-supplier relationships, as well as on intersectoral and interregional job mobility to measure regional differences in resilience. The actual resilience of a region depends on the combined effect of the three determinants. Using the example of the Dutch regions, we show that this is not simply the sum of the three effects taken separately.

The outline of the paper is as follows. The following section describes more extensively how the three determinants are assumed to affect regional resilience and section 3 explains the composition of the model. In the second part of the paper, we apply the model to the case of the Netherlands. Section 4 describes which data has been used and how each determinant is measured. Section 5 shows the model outcomes for different types of shocks and the final section provides the conclusions, discussion and recommendations.

2. Theoretical background

With respect to the definition of resilience in the context of regional economies, Simmie and Martin (2010) and Martin (2012) have identified two different theoretical perspectives: the equilibrist and the evolutionary view. The equilibrist view considers regions to be resilient when they are able to withstand a shock, that is, if the regional system is capable of returning to the same structure and state as prior to the shock. The evolutionary approach, on the other hand, focuses on how well a regional system is capable of adapting its structure in response to shocks. What matters for the long-run success of regions is to what extent the regional system is capable to adapt to changes in competition, technology, market opportunities and pressures and institutions. Consequently, the effect of a shock on a regional system is assumed to depend on the 'resilience building' within the region prior to the shock (Simmie & Martin 2010).

We adopt the evolutionary approach to regional resilience and consider a region to be economically resilient when the region is able to absorb a sudden shock, such as the bursting of a bubble or the closing down of a large company, by changing its structure. In other words, we assume that a region will not return to its pre-shock state but instead structural changes have to take place within the region to be able to recover from the shock.

Many factors can affect the resilience of regions to an economic shock: the sectoral composition, export orientation of firms, innovative propensity of firms, skills of the workforce, level of entrepreneurship within the region, but also institutional arrangements (see Glaeser et al. 2011, Fingleton et al. 2012 and Martin 2012). In this paper, we focus on the sectoral composition of the region (that is, the mix of sectors present in the region). This is because the portfolio of activities reflects the set of skills, know-how, and productive capacity, which is crucial in case a shock forces a reconfiguration of the economy, as we further explain. In his analysis of the concept of resilience, Martin (2012) notices that the effect of an economic downturn on the regional economy consists of two phases. The first phase is the shock itself, while the second phase is the recovery from the shock. In our search for determinants of regional resilience, it is, then, important that we distinguish between these two phases. The question 'what are the determinants of regional resilience?' can be split in two: 1) What makes a region more capable to withstand a shock? 2) What makes a region more capable to adapt and recover from the shock?

With respect to the first phase, we argue that the magnitude of the shock within the region depends on the region's sectoral composition both directly and indirectly. Generally, regions with a more diversified sectoral portfolio are assumed to be less sensitive to economic shocks as the risk of being hit by a shock is spread among those sectors (Frenken et al. 2007). However, the magnitude of the shock can be reinforced when sectors are embedded in the region in the sense that backward linkages are localized (McCann & Ortega-Argilés 2011). When the shock hits a sector that has a central position in the local input-output network, it will indirectly also affect production and employment in the other sectors present in the region. Consequently, regions with a more varied sectoral composition are not necessarily less vulnerable to shocks if those different sectors are regionally embedded through supply relationships.

To determine the ability of a region to recover from the shock (the second phase), we focus on the ability of employees to find a new job without having to move. Most people prefer to live close to family and friends and search for jobs within commuting distance from their current home (for empirical evidence see Dahl & Sorenson 2009). Even when people are confronted with a strong regional shock such as the closing down of a large company, most of them seem to consider residential relocation an undesirable option. Holm et al. (2012) showed that after the closing down of four shipyards in Denmark, more than 85% of the laid-off employees did not move elsewhere. Therefore, we consider regional economies to be resilient if laid-off employees are able to find a new job within commuting distance from their home. Regions are considered to be less resilient when they experience a higher increase in unemployment or a higher share of people that decides to migrate to other regions for work. Both cases are indications that the regional economy is unable to reabsorb the excess labor and, therefore, not able to recover from the shock.

An economic shock will trigger a fall in output demand and employers often response by laying off employees in order to reduce costs and the scale of production (Fingleton et al.

2012). If the shock is large enough, laid-off employees are unlikely to be able to find a new job in the sector in which they used to work. Therefore, in order to find a new job, these former employees have two options: find a job in another sector (intersectoral labor mobility) or by commuting to another region (interregional labor mobility).

If other sectors in the region are less affected by the shock, intersectoral labor mobility is an effective way to absorb the shock. However, the opportunities for an employee to change jobs from one sector to another depends on the relatedness of the sectors with respect to the skills that are required for the work. When the production processes of two sectors require at least to some extent the same skills, the skills that the workers have acquired during their prior job are also useful in the new sector (Neffke & Henning, forthcoming). This increases the likelihood that the former worker will find a new job in the other sector. Consequently, regions with a sectoral portfolio that consists of industries which require similar kind of skills are better able to absorb the shock than regions with a portfolio of unrelated industries. If intersectoral labor mobility takes place, the recovery from the shock will lead to structural changes in the sectoral composition of the region, as employment will shift from one industry to another within the region.

When within the region, the opportunities for finding jobs in a related sector is limited – either because the sectors within the region are highly unrelated or because the related sectors have also been hit hard by the shock – the local labor force may have to adjust to the new situation by looking for a job in neighboring regions. This would imply a cost for the worker, since he or she must dedicate time and money to commuting. However, this solution can provide a satisfactory relief compared to the much more expensive alternative of residential relocation or unemployment.

The possibilities for intersectoral labor mobility mainly depend on the connectivity of the region to other regions. Regions with a more central location or regions neighboring more densely populated regions will offer better opportunities for interregional labor mobility than more peripheral areas. The sectoral portfolio of the region may seem to matter less in this case, however, the ability of the workers to find a new job in a neighboring region does depend on to what extent the two regions have a similar sectoral portfolio. Workers will only be able to find a new job in the neighboring region when employment levels in either the sector in which they used to work or related sectors are affected less by the shock in that region.

In sum, we consider three regional factors to shape the reaction of a regional economy to sudden economic downturn: the embeddedness of the supply chain within the region, the relatedness of sectors with respect to the skills required for the work and the connectivity of the region to surrounding regions. The embeddedness of sectors determines the initial propagation, while relatedness and connectivity are associated to regions' capacity of adapting to the shock, by re-absorbing the labor force.

3. A model of adaptive resilience

To study how the embeddedness of the supply chain, relatedness of industries and the connectivity of the region could jointly contribute to regional resilience we combine them into one model. In this way, we can analyze how the different factors affect resilience in a more coherent manner. We construct a two-stage model to capture the two phases of an economic downturn: a) the initial shock and b) the process of recovery.

To model the initial shock (the first phase) we rely on an input-output (IO) model (Miller and Blair, 2009). This model permits to simulate how a shock propagates to the rest of the economy through the backward linkages of the supply chain. IO models are based on some simplifying assumptions, the most relevant of which are: one product type outputted per industry and fixed technical coefficients. Although for some studies this rigidity may be an handicap, it serves well our purpose of analyzing sudden shocks. In fact, an unexpected downturn may cause the immediate cancellation of existing supply contracts, while it might require some time before production techniques adjust.

Looking at the economy, we can think of shocks of various nature: the closing down of a large company, the bursting of a housing bubble, a global downturn or even a natural disaster. In IO output analysis, shocks can be modeled by a drop in final demand. By carefully choosing which product and which regions are affected, the researcher can simulate a variety of shocks. The analysis works this way: in the beginning, an exogenous shock is defined. The final demand of a certain product category, or a set of them, drops by a chosen percentage. The IO model, then, determines how this initial drop is transmitted to different industries, in different regions. The shock is propagated according to the distribution of activities among regions and the supply relations among them (embeddedness).

We define X as the output and d as demand before the shock. The input requirement matrix A determines the needs for intermediates of a sector. More precisely, it tells how much of a product an industry needs to produce 1 unit of output. Given the indices *ro* and *rd* for regions of origin and destination, and given *so* and *sd* for sectors of origin and destination, we have the following equivalence:

$$X_{ro,so} = (I - A_{ro,so}^{rd,sd})^{-1} d_{rd,sd} \tag{1.1}$$

The matrix $(I-A)^{-1}$ is known as the Leontief-inverse. The matrix product between this and the vector d gives the total output needed to satisfy both the final demand and the demand for intermediates needed by the economy. Equation (1.1) can be seen as the state of equilibrium before the shock. The shock is simulated by changing the vector d of final demand and then looking at how this has an effect on total production. If we define the new final demand as d^* and $\tilde{X}$ as the new output, we write:

$$\tilde{X}_{ro,so} = (I - A_{ro,so}^{rd,sd})^{-1} d^*_{rd,sd} \tag{1.2}$$

We note here that since the drop in demand is exogenously defined, and since this has an important effect in the model, it is worth dedicating some attention to this while choosing the entity of the shock. At the end of this section, and in section 5, this is explored more in depth.

After the shock is determined, we want to analyze how the labor market reacts to it, through relatedness and connectivity. The central concept of our model is that workers who lost their job in region *ro* and sector *so* can quickly find a new job if related industries are accessible for them. Hence regions with a higher connectivity and more related industries hold an advantage and are more resilient. Resilience, however, does not only depend on these elements. Also the size of the initial shock has an effect on the likelihood that laid-off employees are reabsorbed. The second part of the model needs to take into account all these elements to assess regional resilience.

In order to bridge the IO analysis and the labor dynamics, we associate a reduction in output to an increase in unemployment in the following fashion:

$$U_{ro,so} = X_{ro,so} - \tilde{X}_{ro,so} \tag{1.3}$$

Where U represents the total number of newly unemployed[1]. In the same way, we associate the reduction in output ($X - \tilde{X}$) with unemployment, we can look at the production (X and $\tilde{X}$) as a proxy for jobs available before and after the shock. Even though other options are possible, it is convenient for our purpose to assume that unemployment before the shock is zero ($U = 0$). Before the shock, there are X jobs and no unemployment, while after we have $\tilde{X}$ jobs and U unemployed[2]. Our aim is to find a way to way to model a sudden increase in unemployment. Thus, the focus is entirely on the newly unemployed workers created by the shock (short-term disruption) rather than on the long-term dynamics. The question that we ask is, given the location and specialization of the newly unemployed and the jobs that resisted the downturn, how is the labor market most likely to re-organize and re-absorb the labor force?

To model the re-absorption process, we borrow the concept of matching functions from labor economics literature. A matching function is a function that uses the number of people looking for a job and the number of jobs available, as inputs, and returns, as output, the number of successful matches (that is, working relations established) per unit

[1] This implies the assumption that different industries (and different regions) use capital and labor in a fixed ratio. Since the assumption is likely to be strict in this context (for instance, the oil industry has a higher capital/labor ratio), we are introducing a bias. However, given the complex dynamics of the model that we present, and given the explorative nature of this paper, we decided not to correct for it. Normative studies should, instead, consider to include a correction.

[2] We notice that another simplifying assumption of the model presented is that, before the shock, workers do not commute. This assumption, together with the assumption of full employment, are not fundamental to the model.

of time (Mortensen and Pissarides, 1999). In our case, we call the matching function m and we write:

$$m = F(U, \tilde{X}) \qquad (1.4)$$

Equation (1.4) tells that the rate at which new jobs are formed depends on both the total number of unemployed (that we proxy with the size of the shock $U = X - \tilde{X}$) and the number of vacancies (that we proxy with the size of output after the shock $\tilde{X}$). It is highlighted that the model, despite using a concept which is central to labor economics, differs from this stream of literature in many respects. Here, the matching function is used as a way of modeling search routines, but we stray from other aspects of a model of labor economics (e.g. breaching rate, search costs, equilibrium unemployment) since they are outside our scope of modeling the advantages of connectivity and relatedness. In comparison to standard matching function, the model we present adds a regional and sectoral dimension to it. How will the newly unemployed in region ro and sector so ($U_{ro,so}$) re-organize themselves, given the survived output ($\tilde{X}_{rd,sd}$)? The matching function, then, becomes:

$$m_{ro,rd}^{so,sd} = F(U_{ro}^{so}, \tilde{X}_{rd}^{sd}) \qquad (1.5)$$

The rate to which successful matching occurs between workers of region ro, sector so and jobs in region rd, sector sd depends on the number of unemployed at the origin U_{ro}^{so} and on the number of jobs at destination $\tilde{X}_{rd}^{sd}$. We note that the matching function, as we wrote it, has an origin-destination component, both for regions and sectors. A successful matching between an unemployed in one region and a job in another will result into commuting[3]. Matching between two different sectors leads to intersectoral labor mobility. It is fundamental in our conceptualization that the likelihood of a matching between region ro and region rd is proportional to the proximity between

[3] Other works have also introduced a spatial dimension in search theory of labor economics with the purpose of modeling commuting (see for instance Rouwendal 2004 or Van Ommeren and Rietveld 2005). In these models, commuting declines with distance because the costs associated to it reduce the utility of agents. The model we present here is simpler in many respects and it is the search behavior itself that leads to an inverse relation between distance and commuting.

these two regions (connectivity). At the same time, it is equally important that the likelihood of a matching between sector so and sector sd is proportional to the relatedness between these two sectors. As a consequence, we need to take these two elements into account when we choose a functional form for (1.5).

Literature in labor economics proposes a variety of functional forms of which the linear, non-linear and Cobb-Douglass matching technologies are the most often applied (Diamond and Maskin 1979, Mortensen and Pissarides 1999). The linear and non-linear functions have the advantage that one can theoretically derive them from micro-founded individual behavior of agents. The Cobb-Douglass function, instead, gives more control to the researcher over its parameters, which in turn helps to deal better with issues such as return to scale.

Although all three options are suitable, we select a non-linear matching function with only one agent (the employee), which is actively engaged in the searching process. The reasons for this choice are related to the mathematical properties of the function. The final measure for resilience is size neutral: doubling the number of unemployed in a region simply doubles the matching rate of the region (no returns to scale)[4]. Furthermore, we prefer it from a theoretical point of view, as we can more convincingly and elegantly include relatedness and connectivity into the matching framework, by deriving the matching function from micro-founded individual behavior of agents.

The matching function is built as follows. In every region ro, sector so there are U_{ro}^{so} newly unemployed, while in every region rd, sector sd there are $\tilde{X}_{rd}^{sd}$ jobs available. The unemployed use their time to look for jobs. We are interested in finding the calling rate, the expected number of calls (e.g., visits to firms' websites, sent applications, job interviews) that an unemployed person makes every unit of time to potential jobs. If we imagine that workers call more often firms in sectors which are related and in regions which are connected, we can assume that the number of calls that the unemployed in

[4] Matching is homogeneous of degree one, while the measure of resilience is homogeneous of degree zero. In literature, it is debated whether returns to scale of matching are constant or increasing (Mortensen and Pissarides, 1999). There is some evidence that the matching process may have increasing returns. The choice for a size neutral function, though, allows us to keep the focus of the paper on evaluating resilience with respect to embeddedness, relatedness and connectivity.

region ro, sector so makes to perspective jobs in region rd, sector sd is proportional to the relatedness between so and sd, and also proportional to the connectivity between ro and rd.

We define connectivity ($c_{ro,rd}$) as the probability that a call is made between ro and rd, and relatedness ($r_{so,sd}$) as the probability that a call is made between so and sd. With $\sum_{ro,rd} c_{ro,rd} = 1$ and $\sum_{so,sd} r_{so,sd} = 1$, we write parameter $f_{ro,rd}^{so,sd}$ as the combined probability.

(1.6) $$f_{ro,rd}^{so,sd} = c_{ro,rd} r_{so,sd}$$

This combined probability can also be seen as a calling rate[5]. The total (expected) number of calls coming from region ro, sector so equals $f_{ro,rd}^{so,sd} U_{ro}^{so}$. The expected number of calls received by a job at destination is $(f_{ro,rd}^{so,sd} U_{ro}^{so}) / X_{rd}^{sd}$. Now we have all the elements to define the matching rate: the total number of calls coming from ro, so that are expected to reach available jobs in rd, sd.

(1.7) $$m_{ro,rd}^{so,sd} = \frac{f_{ro,rd}^{so,sd} U_{ro}^{so} \tilde{X}_{rd}^{sd}}{X_{rd}^{sd}}$$

Finally we can define regional resilience as the number of matching per unemployed person, per unit of time.

(1.8) $$res_{ro} = \frac{\sum_{rd,so,sd} m_{ro,rd}^{so,sd}}{\sum_{so} U_{ro}^{so}}$$

The final measure of resilience is an indication of speed. How fast can a region re-absorb the labor force, which was laid-off during a shock? Although this framework could be

[5] If one multiplies the probability with the appropriate scalar the two concepts can be matched in this case. In alternative, one can redefine the unit of time appropriately. Though, doing this is irrelevant for our purpose of exploring regional differences in resilience as the final indicator of resilience would just end up rescaled.

suitable to study the long-term dynamics of the labor market, we do not explore this option. This is because we are interested in the possibilities that workers have to recover after a sizeable sudden shock. In our view, this concerns the short-period and it is satisfactorily captured by the measure in (1.8). Extensions of the model could expand the analysis to the dynamics in the long-run. The model has four important properties worth discussing.

First, in the IO part of the model (embeddedness), none of the regions and sectors are – in general – more resistant than others (no a priori advantage). This is because the initial shock is exogenously defined by the researcher. If we model a generalized shock (decrease of demand in all products by 5%), every sector and every region will decrease its output by the same percentage. This will make all of them equally resistant to the initial shocks. If, instead, we model a non-neutral shock (e.g., a decrease of demand in food by 5%) than the shock will have the biggest impact in those regions that have a relatively high share of production dedicated to the food supply chain. There are no regions that are in general more resistant to shocks, but only regions that are resistant to certain kind of shocks and regions that are resistant to other kind of shocks.

Second, relatedness and connectivity, on the contrary, do give certain regions and sectors some a priori advantages. Better connected regions and a more related sector can re-absorb shocks more quickly than isolated sectors in remote regions. This property is not a direct consequence of the model, but a feature that we have chosen to include. An important assumption is implied here. The overall searching effort (e.g. the total number of applications) is smaller for workers in remote regions and in poorly related sectors. This is in contrast with the assumption of a fixed budget constraint normally adopted in labor economics, but it is required in our model to differentiate regional resilience. A way of interpreting this assumption is that workers are discouraged to send applications in regions that are too far or in sectors that are unrelated. Thus, workers which benefit from more connectivity and relatedness send more applications and are more likely to quickly find a job. The same assumption implies that a generalized improvement of connectivity and relatedness in the nation can be seen as an improvement in the resilience of the whole nation, as matching happens faster.

Third, the measure of resilience we propose is characterized, to some extent, by hysteretic behavior, similar as suggested by Martin (2012). The speed of recovery is relatively stable to moderate shocks, but a strong downturn may seriously compromise regional resilience. In fact, a large and uneven shock may also have serious effects in regions that are very well connected and that are specialized in related activities.

Fourth, the matching process works as a gravity model, where the force (matching) between origin and destination depends on the unemployed at origin and the jobs at destination. Nevertheless, the measure of resilience is size neutral. The matching function is divided by U_{ro}^{so}, so there is virtually no mass effect with respect to the origin. As for the destination, resilience depends on the share of survived jobs on total jobs before the shock ($\tilde{X}_{rd}^{sd} / X_{rd}^{sd}$) and it also carries no mass effect. With this property, our measure of resilience does not grant an a priori advantage to larger regions.

The model we just presented is not suited for easy labeling. On the one hand, it has some characteristics of mainstream equilibrium models. At the starting point the economy can be seen in perfect equilibrium, with no unemployment and market clearing. Then, by assuming an exogenous shock, we observe the changes in the economy (this modus operandi is known as comparative statics or, in some context, sensitivity analysis). On the other hand, the adjustment process follows an evolutionary dynamic. Laid-off workers are matched with new jobs according to search routines, which depends on their prior work experience and current residential location, not by following rational behavior (e.g. utility maximization). If one had to study the evolution of the adjustment in the long period, could explore considerations about long-run equilibria, stability of the system, breaking points and so forth. However, the main interest of this paper is the adjustment itself, rather than the equilibrium after the adjustment. This is because, in our conceptualization, regional resilience depends precisely on that: regions that are not able to re-employ laid-off workers within a reasonable time will inevitably experience a decline. Hence, the measure of resilience we propose in this paper is based on the speed of adjustment, which makes the model primarily an evolutionary one.

4. Data and measurements

To better display how the model works, we apply it to the case of the Netherlands. To do so, data is fed into the model to simulate the shock and the recovery process and, subsequently, to assess the resilience of Dutch regions. For an empirical application of the model there are three determinants, for which we need data. First, we need to have a working input-output model to measure the level of embeddedness of sectors within regions. This requires data on sectoral-regional demand and output, as well as data on intermediate supply interactions. Second, micro-data on intersectoral labor mobility is required to assess the actual possibilities of workers to find a job in another industry (relatedness). Third, it is necessary to have data on commuting to measure what is the willingness to travel for work daily and at what distance (connectivity). In this section, we explain what data and which methodology we have used to measure these three determinants of resilience. For the analysis of resilience of Dutch regions we distinguish among 12 regions (NUTS-2 level) and 59 goods and services (NACE Rev. 1.1., 2-digit classification).

4.1 Embeddedness of input-output linkages

To assess the embeddedness of input-output linkages within the 12 Dutch regions, we use an input-output (IO) model that has been developed by the PBL Netherlands Environmental Assessment Agency for the year 2000. More precisely, the PBL IO model concerns 59 NACE Rev. 1.1 2-digits product categories and 256 NUTS-2 European regions. Information on regional production, regional consumption, national use tables, national import tables, international trade, regional freight and international business flights, have been combined with the purpose to infer the most likely structure of supply relationships in Europe. In the ideal case, we would use data on the actual network of supply relationships among the 59 sectors in 12 regions and, possibly, also with the rest of the world. However, since this kind of data are not collected, the IO model from PBL is the best option available.

The construction of the PBL data was carried out as follows. First, international trade of goods (Feenstra et al., 2005) and services (Eurostat, 2009) was made consistent with the national accounts. Second, national use and supply tables were regionalized using regional statistics on production and consumption (Cambridge Econometrics, 2008). Third, interregional trade was inferred using data on interregional freight (Ministerie van Infrastructuur en Milieu, 2007) for goods and interregional business fights (MIDT, 2010) for services. Fourth, this first estimate of trade was constrained to be consistent with the regionalized use and supply accounts. Next, the regional tables were made suitable to an IO framework by a diagonalization procedure (one output per industry, see Miller and Blair 2009). And finally, the tables were given the required extra sectoral dimension. Trade was split proportionally, by sector of destination or final demand.

The final result is a matrix of technical coefficients (the A-matrix) 256x59 sectors of origins, supplying 256x59 sectors of destination. In this research, we do not need this level of detail, as we focus only on the twelve Dutch regions. The remaining 244 regions have been grouped together into one single foreign region. Hence, the A-matrix has 13x59 sectors of origin and 13x59 sectors of destination. To simplify, final demand has been made neutral with respect to region of destination. It is a 13x59 column, which tells what is the demand for a product of a sector in a region, but it does not distinguish where the final consumer is located. So when we discuss a drop in final demand by 5%, we do not specify where this demand is coming from. For the purpose of the IO analysis of this paper, this is a sufficient simplification, as we can still select the production of which region or sector is affected by the initial shock.

4.2 Relatedness between industries

To assess the possibilities of workers to find a job in another industry, we use data on labor flows between the 59 sectors in the Netherlands from 2001 to 2004. Labor flows indicate how likely it is that a person who loses a job in one industry is able to find a new job in another industry. To measure this likelihood, we apply the methodology developed by Neffke and Henning (forthcoming) who use labor flows to infer to what degree two industries use the same skills. They call this 'relatedness' between industries and originally used the measure to analyse knowledge transfer between industries. We calculate relatedness in the same way but, instead, use it to assess the probability of intersectoral labor mobility within the Netherlands.

The information on labor flows is composed using register data from Statistics Netherlands that contains information on all jobs in the Netherlands from 2001 until 2004. To avoid any disturbance of short-term jobs, we have selected those jobs which are not registered at a temporary employment agency, of which the part time factor is larger than 0.5 and where someone has worked more than 3 months during one year (52.9% of all jobs). All individual labor market moves, that is, changes in employment by an employee from one year to the next by moving to another firm unit, have been aggregated to the industry level (4-digit NACE codes). The 4-digit industries that employ fewer than 250 individuals per year on average because their job flows are too small. The resulting database contains the job flows between in total 437 4-digit industries.

Besides relatedness, labor flows between industries depend on several other characteristics of industries such as the size of the industry, its growth rate and wage levels. Therefore, to be able to assess whether a flow between two industries is exceptionally large, a baseline has been composed that takes those industry characteristics into consideration. Following Neffke and Henning (2009), we estimated a zero-inflated negative binomial model with the observed labor flows as dependent variable and the three industry characteristics of both the industry of origin and destination as independent variables (see Appendix 1 for details). Based on the predictions obtained by the model ($\hat{F}_{so,sd}$), the estimated skill-relatedness from sector *so* to sector *sd*, $r_{so,sd}$ is defined as the ratio of observed to predicted flows:

(1.9) $$r_{so,sd} = \frac{F_{so,sd}}{\hat{F}_{so,sd}}$$

Where

$F_{so,sd}$ = observed labor flow from sector *so* to sector *sd*

$\hat{F}_{so,sd}$ = predicted labor flow from sector *so* to sector *sd*

When the index equals 1, the two industries involved are unrelated, values lower than 1 indicate dissimilarity in required skills and values higher than 1 relatedness.

As this relatedness index is sensitive to differences in the size of the industry, confidence intervals were constructed using a binomial test to determine which flows are significantly lower or higher than the expected flows with a p-value of 5% (for further details about the method applied see Neffke & Henning 2009). For all flows which were not significantly higher or lower than expected, the value of the predicted flow has been replaced by the value of the observed flow, leading to a score of 1 on the relatedness index for these industries.

To link the data on the relatedness between industries to the input-output relationships which are only available on the 2-digit NACE code level, we aggregated the relatedness index from the 4-digit to the 2-digit level. This is done by summing the observed and predicted flows with replacing for the latter group of flows the predicted value by the observed value if the flow was found to be insignificant on the 4-digit level. After this aggregation, we calculated the relatedness index by dividing the observed flow over the predicted flow on the 2-digit level.

This resulted in a matrix for all 3,481 (=59x59) industry combinations that exist at the 2-digit level of the NACE classification. This matrix describes a network indicating how related each industry is to all other industries (see Appendix 1). A normalization procedure is applied to use this matrix in the resilience framework. First, we want the score of relatedness to be in a finite range. We transform the score in the following fashion:

$$r_{so,sd}^{new} = \frac{r_{so,sd}}{r_{so,sd} + 1} \tag{1.10}$$

In this way the score of relatedness ranges between 0 and 1. Second, we are interested only in those industries which are positively related (statistically). We, therefore, only include values larger than 0.5 in the matrix. Third, we set the diagonal equal to 1, so, intraindustry labor flows have the maximum score of relatedness. Lastly, all the elements of the matrix are divided by the matrix total, so that $\sum_{so,sd} r_{so,sd}^{new} = 1$.

4.3 Connectivity

To obtain insights in over what distances employees in the Netherlands are willing to travel for work daily (connectivity), we use data on commuting flows between the region where people live and where they work in the Netherlands in 2008. Similar as for the intersectoral job flows, this data is also taken from the register of Statistics Netherlands providing information on all jobs in the Netherlands. An implication of using job data as the source is that people who have more than one job are counted more than once. Furthermore, the work location is not known on the establishment level and therefore people who work for a company with more than one establishment are assigned to the establishment nearest to their home. This may lead to an underestimation of the actual commuting distance. However, other possible sources for commuting data have gathered information through surveys which is likely to cause even larger biases.

To get from commuting patterns to connectivity, some steps have to be made. The number of commuters between two regions depends on the number people living in the region of origin, the number of jobs in the region of destination and the distance between the two regions. As we only want to capture the distance over which people are willing to commute, we need to correct for the size effect. It is hypothesized that number of workers W in the region of origin (residential region), number of jobs J in the region of destination (work region) and distance d between these two regions contribute to the formation of commuting pattern cp in a similar fashion as Newton's law of gravity:

$$cp_{ro,rd} = \beta_0 \frac{W_{ro}^{\beta_1} J_{rd}^{\beta_2}}{d_{ro,rd}^{\beta_3}} \tag{1.11}$$

We have constructed a dataset with all possible pairs of the 12 Dutch regions, collecting data on workers by region of residence, job positions by region in which the work is located and distance. To account for infrastructure endowment and congestion, we use generalized travel costs as an indicator for the distance between the regions. The measure is the sum of the costs to travel from one region to another and the costs of the time spent travelling weighted by the value of time for commuting (Significance 2009). Commuters can travel between regions using different types of modes (car, train, bus-metro-tram). For each combination, we have selected the cheapest option. With this dataset assembled, we can estimate the parameters (*betas*) of the model in (1.11). To do so, we need to choose an estimation technique. Similar as intersectoral job flows, commuting flows are non-negative and integer-valued. However, for none of the combinations of the 12 provinces, the number of commuters is zero and, therefore, we use a negative binomial model instead of a zero-inflated negative binomial model (for an overview of the results see Appendix 2).

Next, we calculate the predicted commuting flows for all region combinations using the estimated parameters (*betas*) for the three independent variables by replacing the actual number of employed for both the region of origin and the region of destination with the variable average to standardize the effect of the mass of the region:

$$c_{ro,rd} = \hat{\beta}_0 \frac{\overline{W}_{ro}^{\hat{\beta}_1} \overline{J}_{rd}^{\hat{\beta}_2}}{d_{ro,rd}^{\hat{\beta}_3}} \tag{1.12}$$

This results in a matrix of 144 (12x12) regional combinations. If we divide each combination by the matrix total, we obtain our final measure of connectivity: a network which describes the likelihood that employees would travel to another province to find a new job.

5. The resilience of Dutch regions and sectors

Using the model we presented in section 3, and with the aid of the data described in the previous section, we evaluate the resilience of the 12 regions and 59 sectors in the Netherlands. At the first stage of the model of resilience, there is a shock. Following the input-output framework we use at this stage, we conceptualized the shock as a sudden reduction of final demand. The type of shock is exogenous to the model and it needs to be defined by the researcher, who has the liberty to choose the entity of the shock, and in which sectors and regions the drop in demand occurs. To illustrate how sensitive the models is to the choice of the shock, the results of the model are presented in two parts.

First, we describe the results of the two-stage model of resilience for a shock in which every sector, in every Dutch region, experiences a fall in sales of final products, as large as 5%, while the demand from the rest of the world remains stable. Since all Dutch regions and sectors are affected in the same way, the model results provide insights in the general differences between the Dutch regions with respect to the three determinants of resilience: embeddedness, relatedness and connectivity[6].

In the second part of section 5, we deepen our understanding of the stability of the model and examine under which circumstances this stability may be compromised by testing how different types of shocks affect the resilience of Dutch regions. We will define shocks in a number of ways, from the ones specific to a certain product category to downturns related to just one region. The first simulation described in section 5.1 then serves as a benchmark or reference point to which we compare the results of the other simulations.

5.1. Resilience of Dutch regions

In the benchmark case we hypothesize a uniform contraction of demand. The production of every good and service produced for final consumption in the Netherlands is reduced by 5%. It is recalled that we grouped the rest of the world into only one geographical area ('region 13'), with all of the 59 sectors identified. These foreign sectors both use Dutch inputs and serve Dutch industries. At the same time, Dutch industries are serving each other, in a complex web of regional-sectoral industrial relations. With the help of the IO model, we evaluate the drop in production by region-sector. We then calculate the regional contraction as the proportion between the output after and before the shock.

[6] The most generic shock possible is a drop in demand of the same percentage both within and outside the Netherlands. However, in this case, the model would only return a scaled down version of the economy before the shock and consequently does not provide any insights in the ability of Dutch regions to withstand a shock.

Figure 1. Effect of the initial shock: resilience due to embeddedness. The y-axis indicates the value of the new output, as a percentage of the output before the shock.

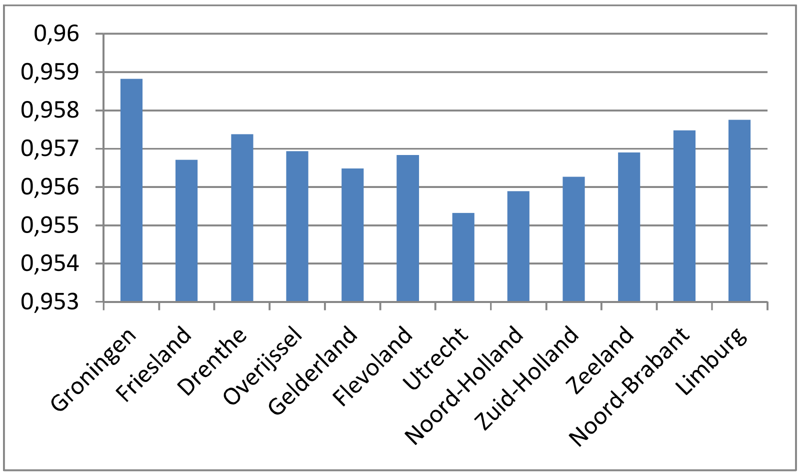

Figure 1 is a simple representation of the effect of the input-output propagation for each of the 12 Dutch regions. It shows how the - originally even - drop in internal demand affects the regions in different ways. Utrecht, Noord-Holland, and Zuid-Holland experience the greatest output reduction. This can be largely explained by their above-average specialization in services as these activities are highly embedded[7]. Services sell most of their final output to local consumers and a significant part of their inputs is constituted by other services, also supplied by local enterprises. Manufacturing activities, on the contrary, only partially rely on local markets, as their production can be more easily exported (World Bank 2009). The shock that we study in this simulation, does not directly affect the economy outside the Dutch border. This implies that the regions which are more manufacturing oriented, and subsequently more export oriented, are more resilient to this internal shock, as they can keep selling their products to foreign markets, while this possibility is partially precluded to service oriented regional economies. This explains why manufacturing regions, such as Noord-Brabant or Limburg, absorb the

[7] Large shares of business and financial services are concentrated in the region of Amsterdam (Noord-Holland). Also Utrecht hosts a number of firms in the financial sector, and, in addition to that, has a strong concentration of computing services. Zuid-Holland is the region where Rotterdam and The Hague are located and this region's specialization is more mixed, due to the different focus of the two cities. Nonetheless, the service sector is still more dominant within the sectoral composition of this region, especially compared to more rural areas of the Netherlands.

shock more successfully (see Figure 1). The same holds for the resource sector, which has high added value and is export-oriented. This also explains why Groningen, that is specialized in the extraction of natural gas, is the most resilient region, with respect to this first stage of the shock. The effect of the shock on the other Dutch regions is between the two extremes, reflecting the moderate specialization of these regions in export-oriented goods.

This simulation shows that regions with a more embedded supply network are the least capable to withstand a shock in domestic demand. This does not imply that these regions are also the least resistant to any other type of shock, actually quite the contrary. Global downturns are more likely to hit export-oriented regions, while embedded regions are less affected thanks to the support of internal demand. In section 5.2, we will show how a drop in external demand is reflected into the Dutch economy.

After this initial shock, we turn to the second stage of the model, the speed of recovery, and focus on the differences between regions in the rate at which laid-off employees find a new job. For this part, we use the matching model described in section 3 which measures the number of work contracts signed per employee per unit of time. The speed of recovery depends first on the number of people that lost their jobs and the number of jobs that survived the initial shock. We take these numbers from the IO part of the model. So in this simulation, Utrecht, Zuid-Holland and Noord-Holland start with the highest number of unemployed and Groningen, Noord-Brabant and Limburg with the lowest.

The other two determinants which are central in the recovery process are the relatedness between sectors and the connectivity between regions. Connectivity gives an advantage to workers from central regions, because it enlarges their possibilities to find a new job in neighboring regions. Relatedness benefits regions that are specialized in activities, which use similar skills. Better yet, the most resilient regions are specialized in activities which use similar skills, but that are not too strongly embedded through input-output relationships. In this way, the decline in one industry does not indirectly lower the availability of possible alternative jobs for workers.

Given the complexity of the matching process and to clearly distinguish between the effect of connectivity and relatedness on regional differences in the speed of recovery, we have split the description of the recovery process in three parts. First, we look at the bare effect of the relatedness between industries. To do so, we assume that employees only search for a new job within the region where they used to work, but – following the relatedness between sectors – that they do dedicate some time in looking for jobs in different sectors. Next, we focus on the effect of connectivity. To provide insights in the relevance of this regional factor, we assume that intersectoral labor mobility does not take place. Workers can only look for jobs in their sector of origin, but they can do it in different regions. Finally, we study the two effects combined. This will constitute our final measure of resilience, the speed of recovery of different regions.

Figure 2 shows the regional differences in the speed of recovery when workers search for a new job according to the relatedness between the industry where they used to work and other industries, but the only look for new jobs in their own region. It is clear that the regions Flevoland, Utrecht Noord-Holland and Zuid-Holland are very resilient in this respect. Despite the fact that the simulated shock affected these regions hard due to their specialization in services, many new matches are formed thanks to the presence of more related sectors in these regions. In fact, the higher relatedness of sectors is likely to follow from the dominance of the service sector in these regions. In the network of relatedness, services are well linked, most likely because the competences needed for these activities are more interchangeable than the ones required for manufacturing or the primary sector (see Appendix 1). Nevertheless, the relatedness network also shows that there are sections of manufacturing that are also rather dense. This may explain the relatively high resilience of manufacturing regions like Noord-Brabant, Limburg, Gelderland and Overijssel. The regions which are more specialized in agriculture or resource extraction show up as the least resilient.

Figure 2 – Speed of recovery: resilience due to relatedness. The y-axis indicates the number of contracts signed, per unemployed person, per unit of time.

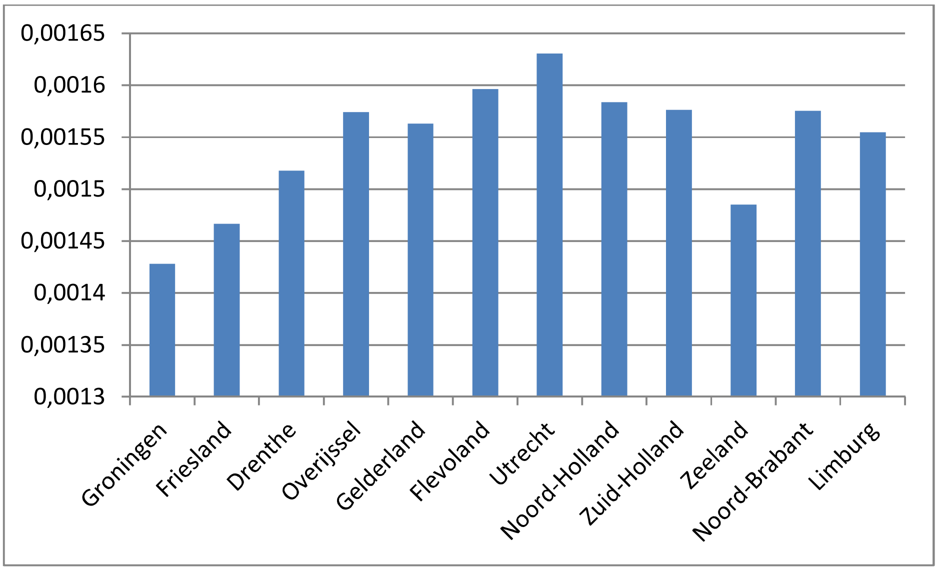

Figure 3 shows the regional differences in the rate of recovery, that is the number of signed contracts, per worker, per unit of time, according to the connectivity between regions. Unemployed look for jobs only in the industry in which they used to work, but these can be located in both their own region and regions within commuting distance. It comes to our attention that the bar-graph roughly draws a wave, displaying low values at the extremities and high values in the center. Since the regions in analysis are listed, more or less, from north to south, we can interpret this wave-shaped results as the manifestation of the effect of centrality on resilience. The regions of Utrecht and Noord-Holland show the highest recovery rate. In these regions, many new matches are formed thanks to their central position in the Netherlands. The same goes for Flevoland, a very well connected region just next to Noord-Holland, and – to a certain extent – for Gelderland and Zuid-Holland, which are relative gainers in the recovery process. In contrast, peripheral regions are relative losers. The three regions in the south, Limburg, Zeeland and Noord-Brabant, resisted well the initial shock, but can recover less quickly because they are less centrally located[8]. Consequently, the unemployed, although lower in number, have more difficulties to find a job than their central counterparts. The regions in the north, Groningen, Drenthe and Friesland, do clearly better. Thanks to lower travel costs, especially among each other, their speed of recovery due to their connectivity is comparable to some of the central regions. However, it is important to remember that the simulated shock of a drop of 5% in internal demand affected these regions less than the more central regions. The question is whether the connectivity of the northern regions is strong enough to recover in a similar rate as the central locations in case of more extreme shocks.

Figure 3 – Speed of recovery: resilience due to connectivity. The y-axis indicates the number of contracts signed, per unemployed person, per unit of time.

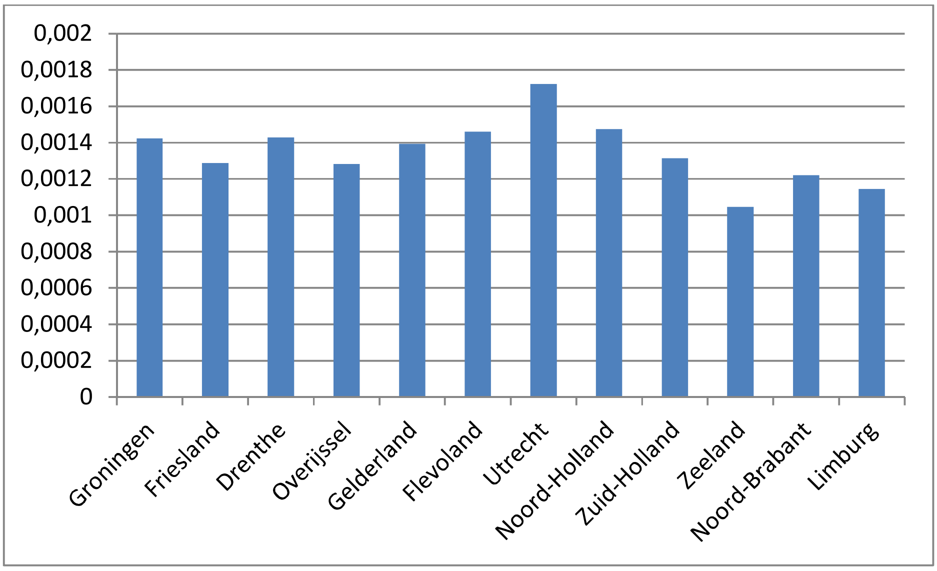

Lastly, we combine the effect of connectivity and relatedness and assume that laid-off employees search for new jobs both in other regions and in other sectors. Figure 4 shows the regional differences in resilience of Dutch regions. Note that the combined framework is not a simple average of the previous two. In line of principle, a region can

[8] It should be noted that our model underestimates the connectivity of the southern regions as no data is available on the number of people commuting from these regions to bordering regions in Germany and Belgium. Especially for the southern regions, their proximity to the larger cities of Aachen and Antwerp may offer opportunities for laid-off employees to find new job. Hence, for future research, it is important to further explore how this location may affect their resilience.

be badly connected, and score low with respect to relatedness of activities in its own territory. But when both commuting and intersectoral mobility are permitted, the activities of this region may be more related to the activities of its neighboring regions. Hence, the region will show a higher speed of recovery in the combined framework than in both of the individual ones.

Figure 4 – Combined effect: Regional resilience. The y-axis indicates the number of contracts signed, per unemployed person, per unit of time.

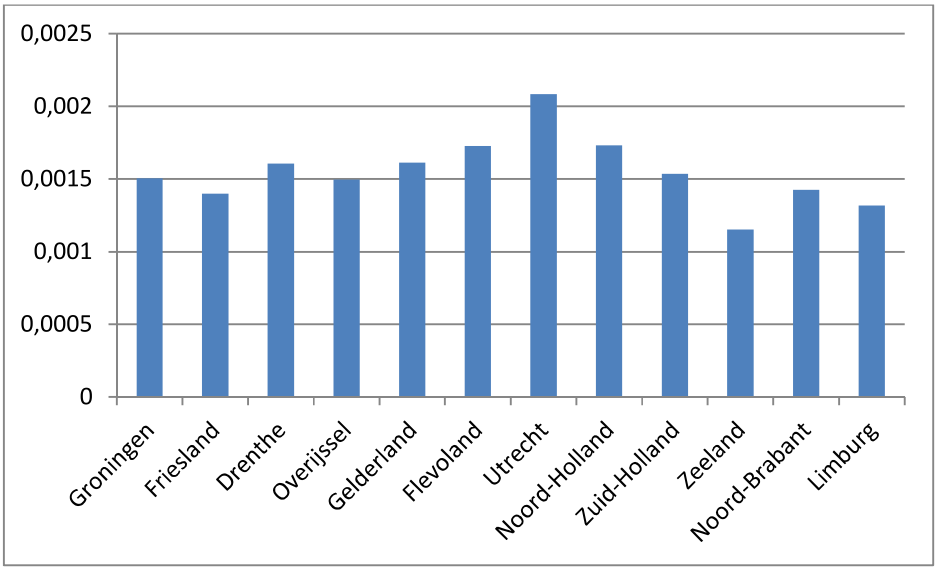

For the Dutch regions, however, we do not observe such large changes in resilience in the combined framework. The resilience of the Dutch regions shown in Figure 4 is in line with what we observed for the effect of connectivity and relatedness separately. The regions that had the highest rate of recovery in the two separate runs, Utrecht, Noord-Holland and Flevoland, still have the highest recovery rate in the combined framework. Zuid-Holland and Gelderland are also still quite resilient compared to other regions. Groningen, Drenthe and Friesland are peripheral regions, which benefitted from their relatively high level of connectivity, but that scored poor on relatedness. This lack of specialization in related activities lowers their overall recovery rate. The province of Noord-Brabant, instead, mainly benefits from its specialization in related activities and, shows better resilience if we evaluate the score taking into account relatedness. Limburg and Zeeland have the slowest recovery speed. Limburg was able to resist the shock quite well due to its specialization in manufacturing, but its location in the far south of the country lowers the ability of the region to recover from the shock and compromises the overall resilience of the region.

5.2. Sensitivity to shocks

In this section, we test the sensitivity of the model to the type of shock that is chosen by the researcher. It has been argued that shocks propagate in different ways according to the nature of the shock itself. In section 5.1, we already showed how a drop in internal demand affects the Dutch regions differently, according to the differences in specialization of regions. Tables 1 and 2 depict the results of the simulation with other types of shocks. The first row in Table 1 is the counterpart of the benchmark simulation that is extensively described in section 5.1: we simulate a contraction in the global economy, while we assume that demand for Dutch final products remains unchanged. In this situation, the regions with the most embedded supply network are the ones who resist the shock the best, due to their relatively low proportion of exported output. However, the change of the shock only marginally affects the overall resilience of the regions, that is, their recovery speed. If we rank regions by resilience, we observe that there is almost no difference between the case of a shock in internal demand and the one of a global downturn. Only Groningen and Overijssel change positions, while their level of resilience stays rather stable.

A similar picture is found when other generalized shocks are simulated. The three other rows in Table 1 represent the effects of a 5% shrinkage in demand for, respectively, the primary sector, manufacturing and financial and business services. In all three cases, our overall assessment of the resilience of Dutch regions remains rather stable. Despite the fact that the initial shock, due to supply relationships within the region, spreads differently in each of the three cases, the relatedness and connectivity of regions imply that the same regions recover more quickly in all scenarios. We do notice, nonetheless, that when the fall in demand concerns only the financial and business sector, Utrecht and Noord-Holland are more strongly affected in their recovery rate. Utrecht almost loses its first position to Flevoland, while Noord-Holland is overtaken by Groningen.

Table 1. Sensitivity analysis - generalized shocks

<table>
<tr><td>Propagation (input-output)</td><td>Resilience</td></tr>
<tr><td colspan="2">*5% foreign demand*</td></tr>
<tr><td>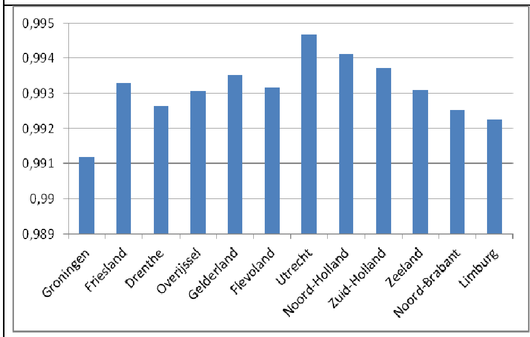
</td><td>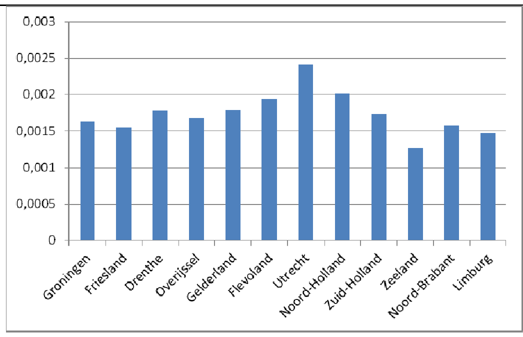
</td></tr>
<tr><td colspan="2">*5% primary sector*</td></tr>
<tr><td>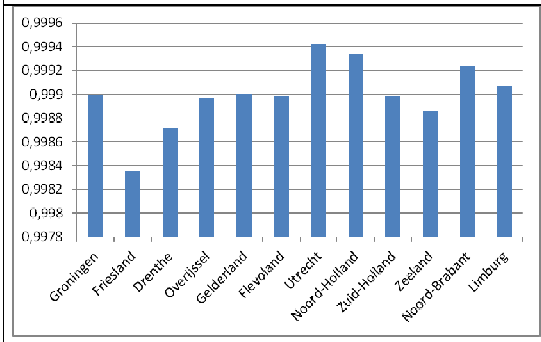
</td><td>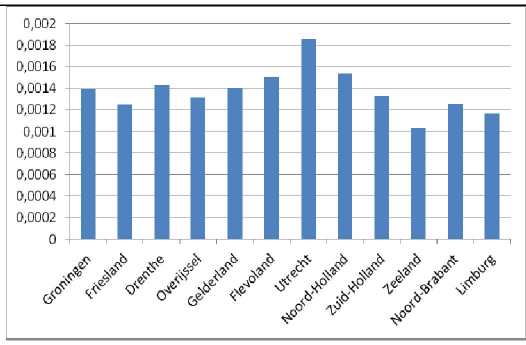
</td></tr>
<tr><td colspan="2">*5% manufacturing*</td></tr>
<tr><td>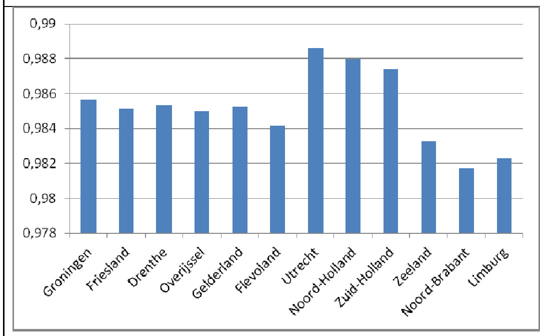
</td><td>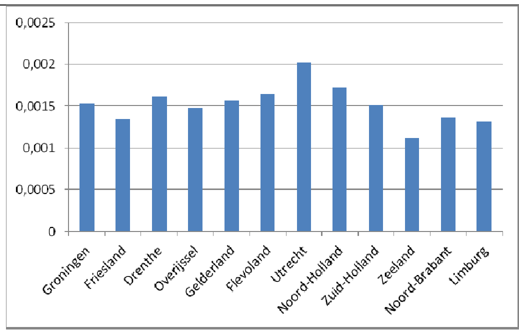
</td></tr>
<tr><td colspan="2">*5% finance and business services*</td></tr>
<tr><td>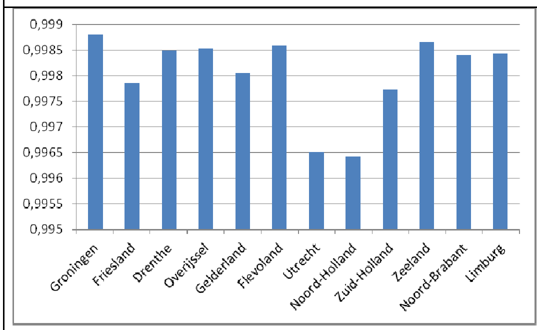
</td><td>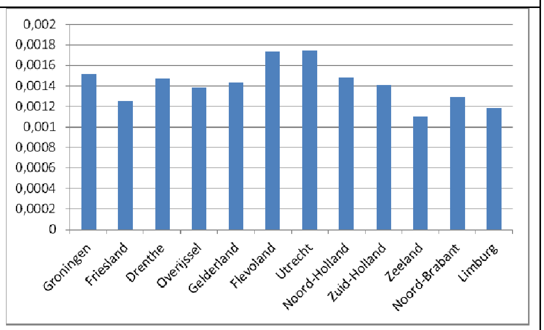
</td></tr>
</table>

In Table 2, the results of simulations of more extreme shocks are shown. In this case, we pinpoint four regions and test how the speed of recovery is modified if a drastic 10% decrease in demand hits just one of them at a time. The more embedded regions Noord-Holland and Zuid-Holland are more affected by the IO propagation than the more export-oriented regions of Noord-Brabant and Groningen. However, while Groningen is indeed much less affected by the shock in local demand (initial hit causes a shrinkage of 5.93%), the impact of the shock in Noord-Brabant (7.26%) is not much lower than in the two services-specialised regions (7.53% and 7.46% respectively). With respect to the adaptive resilience, we notice that, even in these extreme conditions, the regional differences in the recovery rate are fairly stable. Only the drop in demand in Groningen leads to a notably different recovery pattern. Especially the two regions bordering Groningen, Friesland and Drenthe, recover much quicker after a shock in Groningen than after a shock in, for instance, Noord-Brabant. To better understand what causes this difference, we have taken a more in-depth look at the resilience of Friesland to a shock in Groningen and one in Noord-Brabant. Even though Noord-Brabant is located further away from Friesland than Groningen, the shock in Groningen does not cause a much stronger reduction in output in Friesland than the shock in Noord-Brabant (respectively 69 and 65 Million Euros). However, the two shocks do affect very different activities in Friesland. The shock in Noord-Brabant mainly lowers the output of manufacturing activities and hardly affects services in Friesland, which reflects the fact that manufacturing products can more easily be bought and sold over longer distances than services. The shock in Groningen, on the contrary, does cause a drop in output of services, in particular in other business services. Almost 22% of the total output reduction in Friesland due to the shock in Groningen concerns the output from these service activities, compared to only 4.6% in case of a shock in Noord-Brabant. However, those services concern activities such as accountancy, consultancy and advertising which require quite generic skills and, therefore, are highly related to other sectors. Consequently, laid-off employees in these activities can more easily find a new job in other sectors than laid-off employees in most manufacturing activities. This explains why Friesland recovers much quicker from a shock in Groningen than from a shock in Noord-Brabant.

Table 2. Sensitivity analysis - localized shocks

<table>
<tr><th>Propagation (input-output)</th><th>Resilience</th></tr>
<tr><td colspan="2">10% Noord-Holland</td></tr>
<tr><td>
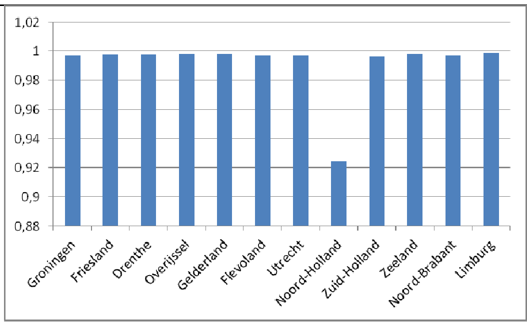

</td><td>
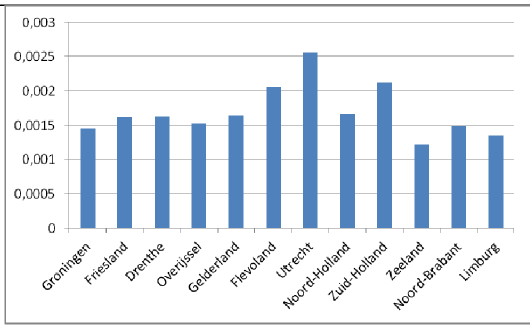

</td></tr>
<tr><td colspan="2">10% Zuid-Holland</td></tr>
<tr><td>
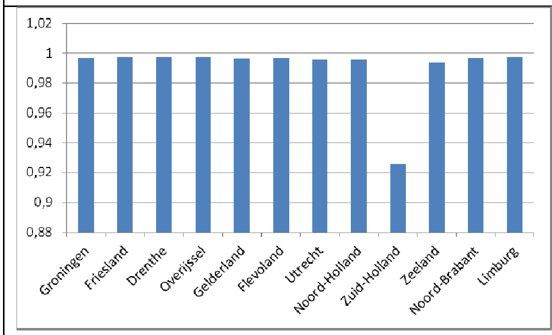

</td><td>
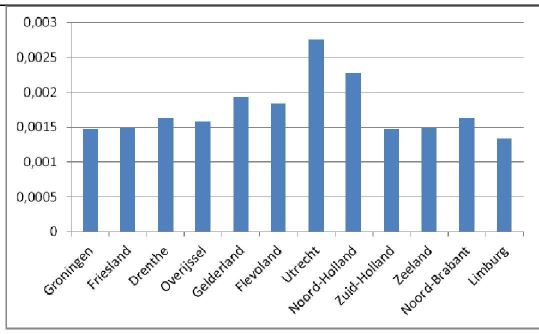

</td></tr>
<tr><td colspan="2">10% Noord-Brabant</td></tr>
<tr><td>
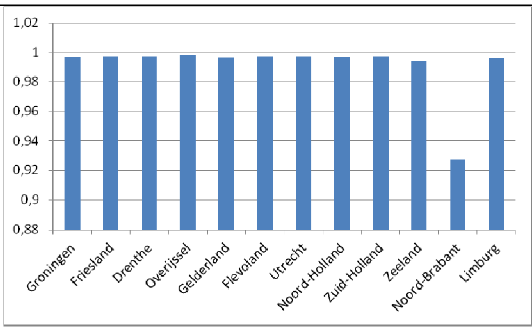

</td><td>
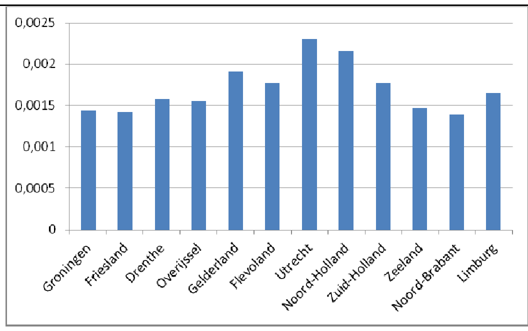

</td></tr>
<tr><td colspan="2">10% Groningen</td></tr>
<tr><td>
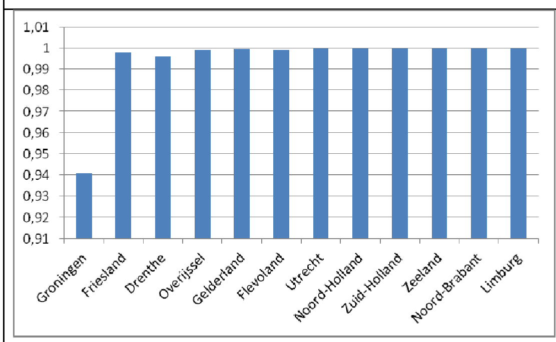

</td><td>
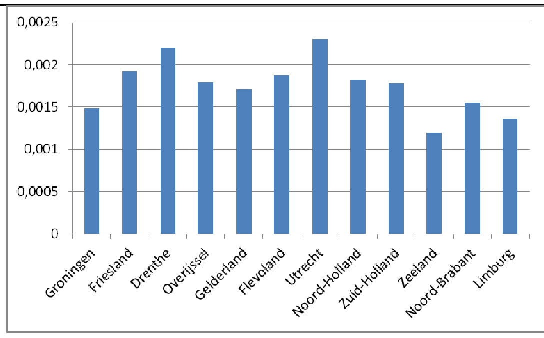

</td></tr>
</table>

6. Conclusions and discussion

In this paper, we attempt to conceptualize and measure what may cause some regions to be more resilient to economic shocks than other regions. Starting from the existing literature on regional resilience, we suggest that different determinants of resilience are important in the two different phases of the downturn. In the first phase, the initial hit, the extent of the damage depends on the how the shock is propagated through the regional input-output structure (embeddedness). Once a shock has hit a regional economy, we suggest that the second phase, the recovery, depends on how fast laid-off workers can be reabsorbed in the labor market. We link this ability to the degree of intersectoral labor mobility among existing activities in this region (relatedness) and to the geographical position of the region (connectivity).

To better understand the contribution of embeddedness, relatedness and connectivity to regional resilience, we combine the three determinants into one model. Using data on NUTS-2 regions in the Netherlands, we illustrate the behavior of the model and show what are the implications for the resilience of Dutch regions of such a conceptualization.

With respect to the initial hit, we show that a highly embedded region, that is a region with a large portion of the supply chain located within its territory, is more vulnerable to internal shocks, but less vulnerable to external shrinkages in demand. The opposite goes for export-oriented regions. If a foreign economy is hit by a recession, export-oriented regions are more exposed to the shock. Because of this characteristic of the model, and because we have no a priori reason to believe that an internal shock is more or less common than an external one, we conclude that there is no conformation of the regional supply chain that makes regions inherently more or less resilient than others.

Different is the case of the recovery phase, after the shock. The actual value of resilience depends on the combined effects of embeddedness, relatedness and connectivity and, this is not simply the sum of the three effects separately. However, generally, the adaptive resilience of a region is higher if it is specialized in sectors that are not strongly intertwined in the supply network, but that do require similar kind of skills and, therefore, are rather related to one another through labor mobility. Furthermore, the resilience of a region is also higher if is well connected to other regions, but how much the region benefits from this depends on how related and how embedded the sectors of the region are to those in nearby regions.

The sensitivity analysis shows that some regions have inherently higher adaptive resilience than others, irrespective of the type of economic shock occurs (e.g., internal drop of demand, global downturn, sector-specific shock). Nevertheless, if the initial shock is very large and localized, the capacity of a region to adapt and absorb the shock can be compromised. From this perspective we notice that, if the regional adaptation process to shocks works in the way we conceptualize in this paper, resilience shows hysteretic behavior.

Extensions of this study and future research can go in several directions. First, there is room for improvement for the model. The measure of resilience described in this paper is completely neutral with respect to size. This choice was made because it allows a better focus on the contribution of embeddedness, relatedness and connectivity to resilience. Nevertheless, some further work is required to understand the role of size on regional resilience. This could be done by experimenting with different matching functions. The Cobb-Douglas function, for instance, gives more freedom in this respect. Second, we illustrate the behavior of the model using actual data, but we do not test the model's capacity of predicting resilience. Previous literature on regional resilience has attempted to empirically measure regional resilience (Fingleton et al. 2012), while this research measures possible determinants, a logical following step would be to verify the explanatory power of these determinants into predicting resilience. Third, the theoretical ground of our measure of resilience is that laid-off workers have to find a job. If they do not succeed in this quickly enough, they may decide to move out of the region, contributing to the decline of the regional economy. But the fact that workers will move if they do not find a job within reasonable time is only assumed. We do not actually model what this time would be. Until when (or up to what costs) are the unemployed able to withstand the situation and when are they starting to move out of the region? Extensions of this model can take this into account. Next, from an empirical point of view, it is interesting to use the model to explore differences in resilience for high and low-educated employees. Since the latter group of workers is less likely to commute over larger distances (Schwanen et al., 2001) and less likely to move to another region (Van Ham, 2005), a shock that hits these workers may impact the regional resilience quite differently than a shock impacting high-educated employees. Lastly, the model could be also enriched by elements which would narrow the gap with mainstream economic theory. Future research, for instance, could attempt to explore the role of wages.

References

Cambridge Econometrics (2008). Regional production, investment and consumption in Europe for year 2000. Data acquired in 2008.

Dahl, M.S., Sorenson O. (2009) The embedded entrepreneur. *European Management Review* 6, pp. 172-181.

Diamond, P.A., Maskin, E. (1979) An equilibrium analysis of search and breach of contract, I: steady states. Bell Journal of Economics, The RAND corporation 10(1), pp. 282-316

Eurostat (2009), International trade in services. Data for years from 2000 to 2004, downloaded in 2009.

Feenstra, R.C., Lipsey, R.E., Ma, H.D.A.C., Mo, H. (2005) *World trade flows: 1962-2000*. NBER working paper No. 11040, January 2005.

Fingleton, B., Garretsen, H., Martin, R. (2012) Recessionary shocks and regional employment: evidence on the resilience of UK regions. *Journal of Regional Science* 52(1), pp. 109-133.

Frenken, K., van Oort F.G., Verburg, T. (2007) Related variety, unrelated variety and regional economic growth. *Regional Studies* 41(5), pp. 685-697.

Glaeser, E.L., Ponzetto, G.A.M., Tobio, K. (2011) *Cities, skills and regional change*. Discussion paper number 2191 HIER. Cambridge, Ma.: Harvard Institute of Economic Research.

Holm, J.R., Østergaard, C.R. (2010) *Sources of regional resilience in the Danish ICT sector*. DRUID working paper No. 10-28. Copenhagen/Aalborg: Copenhagen Business Schoool, Department of Industrial Economics and Strategy/Aalborg University, Department of Business Studies.

Holm, J.R., Østergaard, C.R., Olesen, T.R. (2012) *The unleashing of creative potential from exiting firms… not really?* Paper presented at DRUID 2012, 19-21 June, Copenhagen.

Martin, R. (2012) Regional economic resilience, hysteresis and recessionary shocks. *Journal of Economic Geography* 12(1), pp. 1-32.

McCann, P., Ortega-Argilés, R. (2011) *Smart specialization, regional growth and application to EU cohesion policy*. Economic Geography working paper 2011. Groningen: Faculty of Spatial Sciences, University of Groningen.

MIDT (2010). Regional flights, business and first class. Data for november 2000.

Miller, R.E., Blair P.D. (2009) *Input-Output Analysis 2nd edition.* Cambridge: Cambridge University Press.

Ministerie Infrastructuur en Milieu/Ministerie van verkeer en waterstaat, AVV (2007) *Basisbestanden goederenvervoer 2004*. NEA: Rijswijk.

Mortensen, D.T., Pissarides, C.A. (1999) New developments in models of search in the labor market. In: Ashenfelter, O., Card, D. (eds.) *Handbook of Labor Economics, vol. 3*, Amsterdam: Elsevier, pp. 2567-2627.

Neffke, F., Henning, M. (2009) *Skill-relatedness and firm diversification*. Working paper #0906 of Papers on Economics and Evolution. Jena: Max Planck Institute of Economics, Evolutionary Economics Group.

Neffke, F., Henning, M. (forthcoming) Skill-relatedness and firm diversification. *Strategic Management Journal*.

Rouwendal, J. (2004) Search Theory and Commuting Behavior. *Growth and Change* 35(3), pp. 391-418.

Schwanen, T., Dieleman, F.M., Dijst, M. (2001) Travel behaviour in Dutch monocentric and policentric urban systems. *Journal of Transport Geography* 9, pp. 173-186.

Significance (2009) *Gegeneraliseerde reiskosten op gemeenteniveau, memo 24 november 2009*. Den Haag: Significance.

Simmie, J., Martin, R. (2010) The economic resilience of regions: towards an evolutionary approach. *Cambridge Journal of Regions, Economy and Society* 2010 3, pp. 27-43.

Van Ham, M. (2005) Job-related migration in the Netherlands. The role of geographical access to employment in a polynucleated urban region. *Journal of Housing and the Built Environment* 20, pp. 107-127.

Van Ommeren, J., Rietveld, P. (2005) The commuting time paradox. *Journal of Urban Economics* 58, pp. 437-454.

World Bank (2009), World Development Indicators. Washington, DC: World Bank.

Appendix 1. Relatedness between industries using labor flows

Estimating predicted labor flows

To isolate the effect of relatedness, we corrected for three other industry characteristics that also affect the size of the labor flow: size, growth and wage level. We estimate the following model:

$$\begin{aligned} F_{so,sd} = \exp(\alpha + \beta_1 * \ln(Empl_{so}) + \beta_2 * \ln(Empl_{sd}) + \beta_3 * \ln(Wage_{so}) + \beta_4 * \ln(Wage_{sd}) \\ + \beta_5 * \ln(Wage_{so}) + \beta_6 * \ln(Growth_{sd}) + \beta_7 * \ln(Growth_{so})) \end{aligned} \tag{1.13}$$

With the dependent variable $F_{so,sd}$ being the labor flow between sector of origin *so* and sector of destination *sd.* The independent variables have been measured as follows:

$Empl_{so}$: sum of employment in sector of origin *so* across 2001, 2002, and 2003

$Empl_{sd}$: sum of employment in sector of destination *sd* across 2002, 2003, and 2004

$Wage_{so}$: average wage in sector of origin *so* across 2001, 2002, and 2003

$Wage_{sd}$: average wage in sector of destination *sd* across 2002, 2003, and 2004

$$Growth_{so} = \frac{Empl_{so}(2003) - Empl_{so}(2001)}{Empl_{so}(2003) + Empl_{so}(2001)}$$

$$Growth_{sd} = \frac{Empl_{sd}(2004) - Empl_{sd}(2002)}{Empl_{sd}(2004) + Empl_{sd}(2002)}$$

Following Neffke and Henning (2009) we use a zero-inflated negative binomial specification to determine the predicted flow based on the industry-level characteristics with all 437x437 4-digit NACE level industry combinations as cases. This model specification is the most suitable because labor flows are non-negative and integer-valued and characterised by an overabundance of zeros as in many cases no flows exist between industries (57% of all industry combinations). To improve the efficiency of the estimates, the data has been pooled by summing labor flows and employment data across all available years. We include the level of employment in the sector of origin and destination in the regime selection equation, and, in the count data equation, a log-transformation of employment and the wage level in both the sector of origin and destination and the growth of both sectors. Table A1 shows the results.

The sizes of the origin and destination industry have a positive effect on the size of the labor flows. The negative effect of the level of employment of both the region of origin and destination in the regime selection equation further confirms this effect, as the regime selection is farmed in such a way that the probability of that the probability of observing a flow of 0 is coded as a 1. Therefore, a negative effect in this equation indicates a positive effect of the variable on the dependent variable.

Table A1. Results of the zero-inflated negative binomial regression of labor flows

	Parameter estimate	**Standard error**	**p-value**
Count data equation			
Log(Empl_so)	1.902***	0.030	0.000
Log(Empl_sd)	1.896***	0.030	0.000
Growth_so	-1.355***	0.186	0.000
Growth_sd	0.187	0.143	0.190
Log(Wage_so)	-1.294***	0.144	0.000
Log(Wage_sd)	-0.916***	0.129	0.000
Constant	-10.682***	0.528	0.000
Regime selection equation			
Empl_so	-0.000***	0.000	0.000
Empl_sd	-0.000***	0.000	0.000
Constant	0.640***	0.097	0.000
Over-dispersion parameter			
Log(Alpha)	0.826***	0.018	0.000
Wald Chi^2	5084.46***		
Log likelihood	-314369		
N observations	190,969		
N observations flow=0	108,900		

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses.

Faster growing industries have a lower likelihood of labor outflow. The effect of growth in the industry of destination is positive indicating that these industries experience higher labor inflows, however, the effect is not statistically significant. The model also shows that the wage level of both the industry of origin and destination on the size of the labor flow is statistically significant, but negative in both cases. In other words, industries with higher wage levels are less likely to have a higher labor outflow, but also likely to have a lower inflow. We do observe that the limiting effect on the outflow is higher. This effect may be due to the fact that we selected all labor flows and did not limit our model to those employees with wages higher than the industry's median level in the model.

Network of relatedness between industries on 2-digit NACE codes

Figure A.1 shows the network of relatedness between the 59 industries that has been composed with information on labor flows in the Netherlands between 2001-2004 and using the method developed by Neffke and Henning (2009). The figure shows all links with a score of at least 0.65 on the transformed relatedness index and that are significant at the 5% level. We only show these links to improve the visibility of the network. The colour of the nodes indicates to which broader sector an industry belongs.

Figure A.1 Relatedness between 2-digit industries in the Netherlands

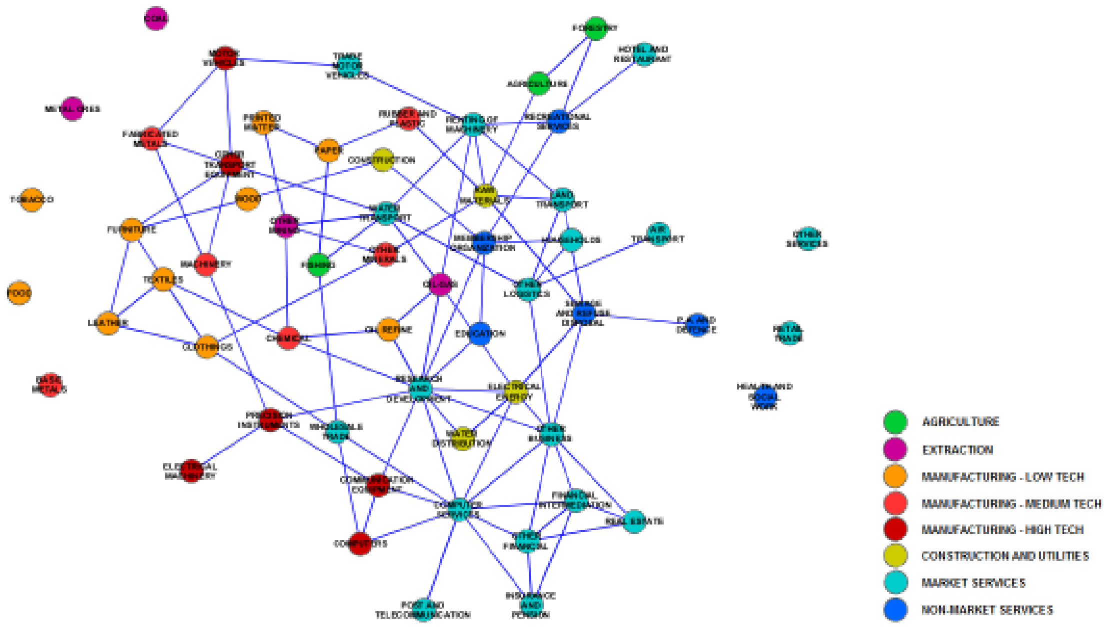

In general, nodes of the same colour cluster together, indicating that more related industries tend to be part of the same broader sector. However, the network also shows that workers are likely to change jobs between broader sectors, even at the 2-digit level. Some industries are quite isolated with only one or even no links to other industries. These industries require certain skills that cannot be easily applied in other industries and most labor mobility occurs between firms active in the same industry.

Appendix 2. Results of negative binomial regression of commuting flows

Statistics Netherlands only provides data on commuting flows between municipalities with a cut-off point of 100. All flows between 1 and 100 are set to 0. To still include information on those flows in the model, we have replaced all 0s by 10 in the matrix of commuting flows on the municipality level, assuming that most of these flows will be quite small. Next, we aggregated this data to the province level. This resulted in the dependent variable of the model: the number of commuting flows between provinces that function as residential location and provinces where the jobs are located in 2008. We estimate the following model:

$$cp_{ro,rd} = \exp(\alpha + \beta_1 * \ln(Travc_{ro,rd}) + \beta_2 * \ln(Empl_{ro}) + \beta_3 * \ln(Empl_{rd})) \quad (1.14)$$

The three independent variables have been measured as follows:

$Travc_{ro,rd}$: the sum of the costs to travel between regions and the costs of the time spent weighted by the value of time for commuting, using the mode of transportation with the lowest cost (2005). The costs of travelling within provinces has been set to the average costs of commuting within all provinces (7.78);

$Empl_{ro}$: number of workers by region of residence (excluding self-employed) in 2008;

$Empl_{rd}$: number of job positions by region in which the work is located in 2008.

Similar as for intersectoral labor flows, these interregional labor flows are non-negative and integer-valued. However, this data is not characterised by an overabundance of zeros as between all provinces commuting flows have been observed. Therefore, we use a negative binomial specification for this estimation. All three independent variables have been log-transformed. Table A2 shows the results of the model.

Table A2. Results negative binomial regression of commuting flows

	Parameter estimate	**Standard error**	**p-value**
Log(Travcosts_ro,rd)	-3.144***	0.088	0.000
Log(Empl_ro)	0. 502***	0.052	0.000
Log(Empl_rd)	0.869***	0.046	0.000
Constant	0.887	1.086	0.414
Over-dispersion parameter			
Log(Alpha)	-1.516***	0.110	0.000
Wald Chi^2	2756.94***		
Log likelihood	-1382.858		
N observations	144		

*** $p<0.01$, ** $p<0.05$, * $p<0.1$. Robust standard errors in parentheses.

All three variables have a statistically significant effect on commuting flows and the sign of the effect is as expected. The sizes of the region of origin and destination both have a positive effect on the size of the commuting flow which indicates that commuting flows are larger between regions where more employed live and where more jobs are available. Travel costs has a negative effect; the more expensive it is to travel to another province, the smaller the commuting flow between provinces.