STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
659 Degree project in economics
Spring 2020

The Generalist Trap – A Theoretical Exploration of Specialisation in Labour Markets

Henry Chen (24254) and Marcus Hagman (24124)

Abstract: It has previously been acknowledged that the increase in productivity resulting from time allocated to a task is a key driver of specialisation in economies. This paper presents three models exploring the topic of specialisation by assuming that the allocation of an agent's time influences production both directly and indirectly, the latter through an experience-based learning effect. The first model presented focuses on a single agent, free to allocate their time as they wish on a continuous scale. In this model, specialising leads to productivity gains but increases exposure to transaction cost. According to the model, lifetime utility can only be maximised by either pure generalists or pure specialists, of which the preferable strategy is determined by transaction costs. When the effect of a transition from generalist to specialist is explored, the model demonstrates that the decision entails a short-term loss, but also a long-term potential gain in utility. The subsequent models incorporate game theory, allowing focus on the entire economy. Both these models demonstrate that even in situations where everyone benefits from a specialised economy, an individual may benefit from remaining a generalist. We call this phenomenon *the generalist trap*, defined as a situation where specialisation on a societal level is the desired outcome, but barriers of various forms prevent this from being achieved.

Keywords: time allocation, specialisation, generalisation, labour market, human capital, learning by doing, game theory

JEL: J220, J240, C720, C780

Supervisor:	Karl Wärneryd
Date submitted:	13 May 2020
Date examined:	27 May 2020
Discussants:	Henrik Taro Hansson and Olof Tydén
Examiner:	Johanna Wallenius

We would like to extend our sincere gratitude to Karl Wärneryd, our supervisor, and Murat Üngör at the University of Otago for generously offering their advice and expertise.

Contents

1 Purpose

The importance that specialisation, or division of labour, has had for the growth of economies has been observed for hundreds of years. In The Wealth of Nations, Adam Smith (1776) states that "The greatest improvement in the productive powers of labour, and the greater part of the skill, dexterity, and judgement with which it is anywhere directed, or applied, seem to have been the effects of the division of labour." Smith goes on to present an example of a pin factory. Whilst a single worker producing pins from start to finish could make no more than 20 pins in one day, a team of ten workers, each assigned their role in the production, could make upwards of 48,000 pins in a single day – a 240-fold increase in efficiency.

Implicit in Smith's argument is that the driver for specialisation of individuals is the simple fact that performing a task repeatedly leads to an increase in productivity. This effect is often referred to as *learning by doing* in the literature. Though the connection between experience-based learning and specialisation has long been acknowledged, the consequences of the connection remain relatively unexplored. Seeking to provide theoretical contribution to this issue, the research question of this paper is as follows:

How does experience-based learning affect the decision of economic agents to specialise and what does this imply regarding the labour market as a whole?

To answer this question, we develop a set of theoretical models based on the key assumption that an agent's allocation of time influences production both directly through a production function and indirectly through a skill acquisition function. The models provide insight not only into the circumstances wherein agents should specialise or not, but potentially also with regard to the behaviour of labour markets on a greater scale. The aim, in other words, is not to conclude whether being a generalist or a specialist is preferable in general, but to map out the factors that have an influence on this dilemma, as well as determine the implications for the economy.

Two of the models we will develop are based on game theory, as the outcome associated with a strategy is dependent on the strategies used by other agents in the population. This provides insight into which levels of specialisation the economy is likely to converge to. The models display situations where the economy stabilises at points that are sub-optimal for the economy as a whole, from which we can draw conclusions about the efficiency of the market. We also explore what the implications of a transition to an increased degree of specialisation are for the agent that makes the shift.

The conclusions of any model are sensitive to the assumptions on which the model is built. The different models we develop are based on different sets of assumptions. We chose this approach, rather than focusing on a single model, to show how the predictions of the model depend on the assumptions made with regard to the nature of the market.

2 Current State of Knowledge

This section contains relevant findings from previous research to provide background for later sections. The chosen literature is divided into four categories: division of labour, international economics, empirical findings, and evolutionary biology.

2.1 Division of Labour

A theoretical basis for explaining specialisation of individuals remains less explored than specialisation of countries on a greater scale, even though the very existence of modern economics can be largely attributed to it (Rosen, 1983). The most important efforts to better understand what drives division of labour are outlined below.

Maurice Kilberg and Leon Wester (1966) explore in their article *An Economic Model for the Division of Labor* the direct labour cost-implications of specialisation in manufacturing plants, or more specifically the division of labour at assembly lines. They find an optimum cycle-time, which is the amount of time a product spends at each operator's work station. The cycle-time is also used as a measure of division of labour, as a shorter cycle-time implies further division. They find that specialisation is favourable up to a certain point, but that further dividing work would incur higher learning costs that outweigh benefits in the form of decreased imbalance- and non-productive time costs. Hence, division of labour may not simply be a function of the size of the market. Larger production volumes, contrary to Adam Smith's beliefs, do not necessarily justify more extensive division of labour or specialisation.

Specialisation is also strongly linked to economic inequality between social groups. Joseph Henrich and Robert Boyd (2008) develop in their article *Division of Labor, Economic Specialization, and the Evolution of Social Stratification* an evolutionary game-theoretical model based on the assumption that people belong to social groups, and that group members tend to learn from each other. A number of qualitative conclusions are drawn from this model. First, economic specialisation in terms of people occupying different roles can lead to both additional production and stratified inequality, if the different sub-populations remain culturally isolated from each other. Thus, one may decrease social stratification and inequality by increasing the mixing of different social groups. Henrich and Boyd explain that there exists a point where either stratification or an egalitarian equilibrium is stable. This point is referred to as the *stratification threshold* and is dependent on the mixing rate. Secondly, all other things being equal, more stratification leads to higher average production compared to egalitarian societies. Finally, by increasing the surplus available, stratification tends to increase, meaning that the degree of specialisation increases.

Sherwin Rosen (1983) describes that *indivisibilities* constitute a limit for specialisation, or human capital accumulation. Much like it might be unwise to buy a hammer to drive down a single nail, dividing one's labour too far is unfavourable. This implies that specialisation and trade will maximise both social and private returns in human capital accumulation, as labour will always be divided to its greatest extent according to demand. Therefore, specialisation is an efficient use of resources even if all people are identical. The trend we observe in greater specialisation along with economic progress can also be partly explained by increasing indivisibilities. The development of new technologies and knowledge increases the amount and complexity of skills that are available to be acquired, meaning that further specialisation is made possible. Rosen expands, ”This is one reason why the rate of return to education does not fall with economic development and why education is a more desirable investment in advanced economies than in undeveloped ones.” Like Henrich and Boyd, Rosen also concludes that productivity is higher on average when specialisation is feasible.

2.2 International Economics

Many models have been developed around the specialisation of countries. The field of international economics is based on a central idea of comparative advantage, from which countries that trade derive benefit. One could say that the importance of international economics in recent decades, and with it specialisation, is unprecedented as nations are increasingly closely linked and international trade is becoming more and more complex (Palan, 2010).

In the nineteenth century, David Ricardo noted that two countries producing goods at different opportunity costs benefit from trade even if one country has an absolute advantage in both goods (Ricardo, 1817). In other words, international trade can be explained by differences in opportunity costs. Additionally, Ricardo's model illustrates that not only every country gains from trade, but every individual is also made better off, as international trade is predicted not to affect the income distribution. Ricardo's findings have inspired numerous other iterations of his model. However, the model that Ricardo developed was relatively primitive, as labour is considered the only factor of production.

In the real world, trade may also be partially explained by differences in countries' resources. Heckscher and Ohlin showed that different access to resources is a driver of trade (Ohlin, 1933). A country's comparative advantage is influenced by its *relative abundance* of factors of production. This means that for a two-factor model with capital and labour, a country that is *labour-abundant* would benefit from trade with a *capital-abundant* country.

Paul Samuelson and Ronald Jones developed a version of Ricardo's model based on *specific factors* – factors of production which are immobile between industries (Jones, 1971). The distinction between specific factors and mobile factors is not always straightforward. For example, worker mobility varies with the the job occupation and the worker's personality type. Bruce Fallick (1993) measured that a displaced worker in the United States has the same probability of being employed within four years as a similar worker who was not displaced. Compared to an office building that has a lifetime of 30 to 50 years, it can be concluded that labour is a less specific factor than other types of capital. Moreover, one may ponder whether generalists and specialists respectively, are viewed as specific or non-specific factors.

Furthermore, economies of scale have been shown to be yet another driver of trade (Helpman, 1981). The presence of economies of scale may lead to specialisation even when differences between countries and learning by doing is absent. Specialised economies would be able to achieve higher productivity and trade with each other to consume the full range of goods. When analysing economies of scale, it is important to recognise that smaller firms can form clusters that drive scale economies on an aggregated level (Krugman et al., 2015). This implies that size of the market (smaller firms can collectively provide a larger market) affects the overall specialisation of firms and nations.

The role of human capital in economic growth has been investigated by many. Daron Acemoglu (2009) refers to human capital as ”all attributes of workers that potentially increase their productivity in all or some productive tasks”. The Ben-Porath model is described as much of the basis of modern labour economics, and Acemoglu summarises its importance: ”First, it emphasizes that schooling is not the only way in which individuals can invest in human capital, and there is a continuity between schooling and other investments in human capital. Second, it suggests that in societies where schooling investments are high we may also expect higher levels of on-the-job investments in human capital. Thus there may be systematic mismeasurement of the amount or quality of human capital across societies." The Ben-Porath model also shows that human capital investments are typically higher early in ones life. One may ask what portion of human capital investments is dedicated to specialisation and generalisation respectively, as human capital may depreciate as a result of new machines or techniques that erode existing human capital of a worker, and the rate of depreciation may vary depending on the human capital's proportion of specialisation and generalisation.

2.3 Empirical Findings

There has been a lack of research exploring how the division of labour has developed over time. However, specialisation of nations and regions over time has been measured, and specialisation indices have been developed for this purpose. Additionally, *skill premiums*, ratio of skilled to unskilled workers' wages, have been measured for longer periods of time.

Stephen Redding (2001) proposes an empirical framework using data on 20 industries in 7 OECD countries since 1970. Specialisation is derived from the neoclassical trade theory, where a country's extent of specialisation in a specific industry is measured by its share of the country's GDP. Redding finds that there exists substantial mobility in patterns of specialisation. In the United States, an industry has a probability of 0.20 to 0.39 transiting out of its initial quintile of the distribution of GDP shares after 5 years. Redding expands, "The highest levels of mobility are found in Japan and Sweden; Canada, the United Kingdom, and United States have intermediate levels; Denmark and Finland display the least." Moreover, no evidence of increased specialisation, or the extent of production that is concentrated in few industries, is found in the analysed countries. Finland and Denmark even show decreasing levels of specialisation over the time period. Furthermore, over time horizons of 10 years and above, common cross-country effects are influential on specialisation development, as well as country-specific changes in factor endowments.

Due to the effects of globalisation on the specialisation of countries, the importance of measuring the competitiveness of countries in relation to their extent of specialisation has arisen (Palan, 2010). Despite the necessity, there does not seem to be an agreement on which index best captures specialisation. Nicole Palan divides indices into two groups:

measuring absolute specialisation and relative specialisation. Palan defines measurement of absolute specialisation as ”a country would be considered specialised if a small number of industries exhibit high shares of the overall employment of the country”. For example, Italy would be considered specialised in textiles, while Scandinavian countries would be considered specialised in the production of pulp and paper. Palan defines measurement of relative specialisation as ”the deviation of a country’s industry structure from the average industry structure of the reference group of countries”. For example, Finland is relatively more specialised in communications technologies, although the absolute share of communication technologies on the industry of Finland is low. Indices for absolute specialisation include for instance, the Hirschman-Herfindahl Index (Herfindahl, 1950 or Hirschman, 1964) or the Ogive Index (Tress, 1938). Indices for relative specialisation are for instance the Krugman Specialization Index (Krugman, 1991) or the Relative Gini Index (Hoover, 1936). Palan finds that results differ widely depending on the measure used, and indices on absolute specialisation cannot be compared with indices on relative specialisation. Moreover, she concludes that ”Due to the quite limited availability of consistent input-output data over a long time horizon, the application of more sophisticated measures of specialization is hard to accomplish in empirical studies.”

Ariel Burstein and Jonathan Vogel (2016) analyse the consequences of international trade on the skill premium. In their model, data from 60 countries, 76 merchandise sectors (i.e. agriculture, mining and manufacturing) and 81 service sectors is aggregated between the years 2005-2007 where feasible. A skilled worker is defined as someone who has completed at least a tertiary degree. Burstein and Vogel find that international trade shapes the skill premium in two ways. First, ”Trade reallocates factors towards a country’s comparative advantage sectors, increasing the skill premium in countries with a comparative advantage in skill-intensive sectors and decreasing it elsewhere.” Second, ”Trade reallocates factors towards skill-intensive producers within sectors, increasing the skill premium in all countries.”

2.4 Evolutionary Biology

Specialisation has been explored not only in the context of economics, but also in evolutionary biology. Fenster et al. (2004) illustrate that pollinators of different function groups (e.g. long-tongued flies or nectar-collecting bees), which differ in specialisation, vary greatly in their effectiveness as pollinators for different plant species, which in turn is a driver for floral evolution. Fenster et al. found that ”approximately 75% (209/278) of flowering plant species exhibit specialisation onto functional groups". It can also be concluded that moths, for example, vary in their effectiveness as pollinators, but that there seems to exists an upper limit in the degree of specialisation. Moreover, Fenster et al. describe that specialisation in pollination ecology may also involve, apart from function groups, the time of day that flowers open or where the pollen is placed.

Predator-prey dynamics have also previously been explored with prey specialisation being identified as a result, meaning that predators can be specialised in hunting specific types of prey. Rana et. al. (2002) describe, "Generally, predator–prey associations are thought to be determined more by characteristics of the habitat and/or the phenology, size or abundance of prey or natural enemies than by chemical or other intrinsic characteristics of prey." In Rana et al.'s study, an experiment with ladybirds and and two types of prey (pea aphids and black bean aphids) was conducted. Performance was measured and specialisation was developed over five generations of ladybirds, divided into different groups, by controlling the type of prey available. They found that specialisation on one aphid resulted in a trade-off in performance on the other aphid. Additionally, switching from pea to black bean aphid had a greater effect on performance, along with an increase in mortality during development, than the reverse.

3 Analysis

3.1 Setup

In this paper we will model the endogenous variable *time allocation*, chosen with respect to the objective variable utility. We explore how the dynamics of the choice depends on the market structure we assume for the model. All our models can be described in terms of the chart below, with variations between the models explained by differences in the functions.

How much an agent produces depends on their productivity and time allocation, through the production function. Productivity can be thought of as the maximum number of that good which can be produced in one unit of time. Furthermore, productivity in a task improves with time allocated to that task, determined by a skill acquisition function.

The link between production and consumption is determined by the structure of the market which the agents navigate. The market structure is captured in a trade function. Finally, the utility function serves to rank the agents consumption alternatives.

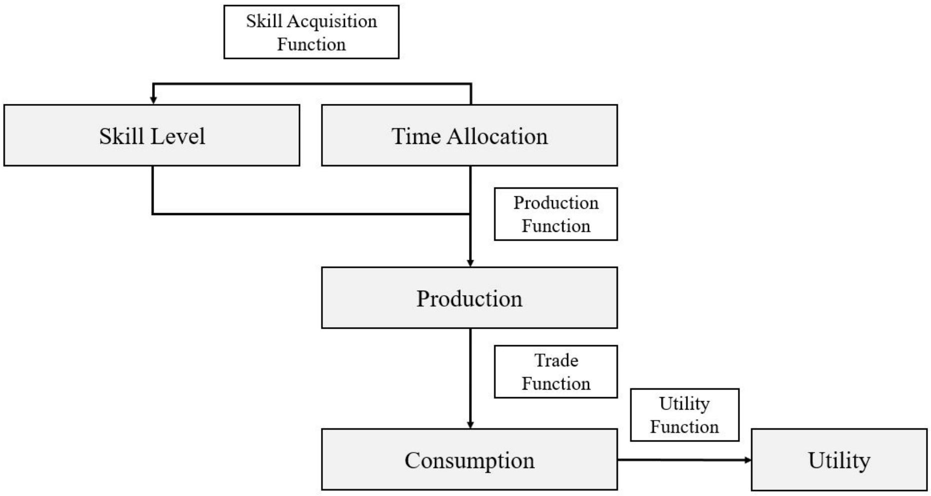

Figure 1: High-level overview of this paper's models

The specifics of the functions are defined according to assumptions made regarding the nature of the market. Each of the three models uses a different set of assumptions. We have chosen to explore the sets of assumptions which are associated with clear conclusions that could contribute to the theoretical understanding in this field.

We use the variable α to denote time allocation. It is defined such that $\frac{1+\alpha}{2}$ of one's time is dedicated to producing apples and $\frac{1-\alpha}{2}$ of one's time is dedicated to producing bananas. In other words, someone with an α of -1 is a specialist in bananas, $\alpha = 0$ means that one is a generalist and $\alpha = 1$ means that one is a specialist in apples. The term *strategy* is used to refer to choice of α.

3.1.1 General Considerations

Before exploring the assumptions made in the functions of the model, and variations thereof, we describe some general considerations necessary to define in order to model the problem.

We have decided to model this problem in continuous time, rather than discrete time. In other words, production, trade, consumption, and skill improvement take place at every infinitesimally short time instance.

We can view the endogenous variable of time allocation as either continuous or discrete. Viewing it as a continuous variable implies that any allocation, from allocating all one's time to producing one good to allocating it all to the second, are feasible options. A discrete approach to time allocation divides the variable into two strategies: generalists who split their time between both goods and specialists who invest all their time in one good. We assume that half of specialists will specialise in one good and the other half in the other. We start by viewing time allocation as a continuous variable, and then move on to describe it as discrete.

In our models, we model only two goods: apples and bananas. For a more general model which can apply to a set of more than two goods we can consider one good to be the set of all other goods in the market. Note that the goods of the model are rivalrous and that the results cannot be used with regard to non-rivalrous goods.

A limitation of the models is that they do not account for innovation that stems from broad knowledge in disparate fields. Nor do they account for the boredom one may experience as a consequence of repeating the same task.

3.1.2 Skill Acquisition Function

Recall that productivity is defined as one's maximum capacity for production of a good per time period. We use M_A to denote productivity in apples and M_B to denote productivity in bananas. All our models view productivity as a function of total time invested into improving production in that good. We have conceived

Under learning by doing, the same process that results in production also leads to an improvement in productivity. We can view this as dynamic skill acquisition because the productivity is changing dynamically at the same time as production and consumption occur.

With learning by practice, productivity is a function of time spent practicing. Practice is a process that occurs prior to any production takes place. We assume that the time allocation between different goods during practice is the same as the time allocation will be during production.

Regardless of the learning framework used, we must define some properties of the learning curve, the function which describes productivity as a function of time invested into the activity that improves productivity. We have chosen to conceive of the learning curve as linear starting from zero. In reality, the learning curve is almost certainly non-linear and can be argued to be either convex (potentially exponential), concave (potentially approaching some asymptotic productivity) or a combination such as the sigmoid function. If a convex learning curve had been used, this would have benefited the specialists due to increasing returns from time investments. Similarly, a concave learning curve would benefit generalists because of diminishing returns to time investments.

Under learning by doing we define the time derivative of productivity as $\dot{M}_A = \frac{1+\alpha}{2}$ and $\dot{M}_B = \frac{1-\alpha}{2}$. Under learning by practice, productivity is written as $M_A = k * \frac{1+\alpha}{2}$ and $M_B = k * \frac{1-\alpha}{2}$ where k is an arbitrary constant.

3.1.3 Production Function

The production function is defined such that every agent faces a linear production possibilities frontier (PPF) at every period. A specialist will produce at a point where the PPF intersects with an axis. An alternative approach, which is not explored at depth in this paper, would involve a concave PPF, indicating diminishing marginal product during each time period. However, a concave PPF would unilaterally benefit generalists. It is therefore possible that the models we set up underestimate the viability of the generalist strategy. It is not the goal of this paper, however, to establish whether it is generally preferable to be a specialist or a generalist, but to explore which circumstances influence the decision.

3.1.4 Trade Function

The trade function shows significant variations between all three models. To define the trade function we must first decide whether trade takes place through an intermediary or directly between agents. Then, if trade is carried out directly between agents we must decide how agents are matched and the rules by which they bargain.

We model trade as taking place through an intermediary in the first model, and directly between agents after that. When trade occurs through an intermediary, we assume a stable relative price of one. We also assume that transaction costs are present, resulting in an agent receiving less than they give up. Conceptually, we may view this transaction cost as a fee paid to the intermediary in exchange for using the exchange.

If trade is instead exercised through pairwise matching, dynamics at the population level may be derived, because the success of a strategy will depend on the strategies employed by other agents in the economy. The matching process can be either random or non-random. We have chosen to construct random matching models as a non-random matching model would be either too complex to derive clear conclusions from or it would end up with no meaningful differences from the intermediary model. Random matching poses a disadvantage to specialists as they can never expect to be matched with a complementary specialist more than half of the time, assuming the specialist population is distributed equally between both goods.

Furthermore, the models in which agents trade directly with one another require rules for bargaining by which the transaction is determined after the agents have been matched. One option, similar to the intermediary model, is to use a fixed relative price of one. Here, the pair simply trades one unit of one good for one unit of the other as long as both agents derive benefit from the transaction. The problem with such a set-up is that it leaves mutually beneficial trades on the table. That is, the final allocation of goods may not be Pareto efficient. A justification of such a set-up, though inefficient, could be that agents feel taken advantage of if they give up more of their good than their counterpart does, causing them to rather take the loss in potential utility.

Alternatively, there could be some bargaining solution which always allows a Pareto efficient allocation of goods to be reached, without either party ending up worse off. One such algorithm is the following: one agent in the pair is randomly designated the role of *offeror*, who then decides which trades should be made so as to maximise their own utility without compromising the utility of their counterpart. The resulting allocation of goods from such a transaction will always be Pareto efficient, and it can be considered fair because the designation of who becomes offeror is random.

To summarise the variations of the trade function, we can conceive the market as either taking place through some intermediary, or directly between agents. If agents trade directly, they may be matched randomly or non-randomly (we produce no model based on non-random matching). When two agents have been matched, the transactions they make with one another may either be constrained by a fixed relative price or it can be determined so as to guarantee a Pareto efficient allocation of goods.

3.1.5 Utility Function

We have chosen to use the Cobb-Douglas utility function $U_t = \sqrt{C_{At} * C_{Bt}}$. Diminishing marginal utility is exhibited in the consumption of one good if consumption in the other good is remained constant. If consumption of both goods increases by an equal percent amount, overall utility increases by the same percentage (constant utility to scale).

An unfortunate specialist who is unable to trade their goods in a random matching model will gain no utility from consuming their goods. Therefore, the model allows for a storage mechanism, wherein goods can be stored for one period if consuming them would yield no utility. Goods cannot be stored for longer than one time instance.

3.2 Intermediary Model

The intermediary model consists of two main sections. In the first section, we model lifetime utility as a function of α, assuming α is held constant throughout a lifetime. In the second section the assumption of constant α is relaxed, with the aim to see how utility changes as over time when strategy is altered. The following assumptions are constant throughout the model:

- The skill acquisition function follows a learning by doing framework. The learning curve is linear.
- At every time instance, trade takes place through an intermediary. Transaction costs, potentially viewed as commission charged by the intermediary, is represented by τ. τ is defined such that an agent who gives up X of one good in a transaction receives $(1-\tau) * X$ of the other good.

We start this analysis off assuming that the time allocation chosen at the beginning of one's lifetime cannot be altered. A lifetime lasts for T units time.

Symmetry between specialisation of either good lets us simplify the analysis by observing only the segment of $0 \leq \alpha \leq 1$, recognising that any features observed in this segment will also be present in the domain $-1 \leq \alpha \leq 0$. Mathematical analysis demonstrates that, for any positive transaction costs, an agent's degree of specialisation must exceed a threshold value for trade to become beneficial. We refer to this level of specialisation as the *point of marginal trade*. For this reason, in order to model utility as a function of α, we model the function separately both with and without trade.

The left picture below shows the graphs of these two functions, derived in the appendix. The red curve assumes autarky, meaning that production equals consumption. We call this *the autarky curve*. The blue curve assumes that the agents trade quantities to maximise utility from consumption. We call this *the trade curve*. To derive lifetime utility as a function of α, we combine the autarky curve for levels of α below the point of marginal trade with the trade curve. After all, the trade curve is only defined for the points where trading increases utility. Combining the two curves yields the curve on the right. We will refer to this curve as the USG (Utility of Specialists and Generalists) curve. Both graphs below are modeled using $\tau = 60\%$

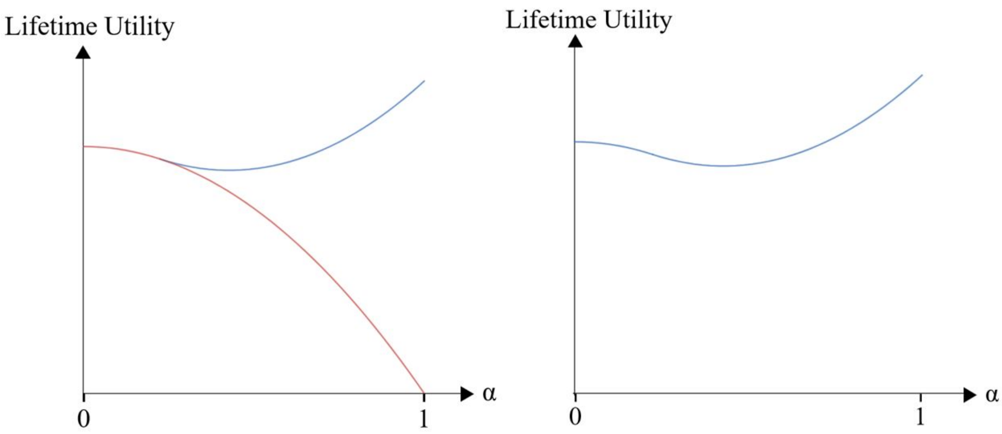

Figure 2: Autarky curve and trade curve (left), and combined USG-curve (right)

A few interesting observations come out of a mathematical analysis of this model. First, comparative statics show that decreased transaction costs shift the entire trade curve upward and moves the point of marginal trade to lower levels of α. To illustrate this effect, we show below three separate trade curves at different levels of τ.

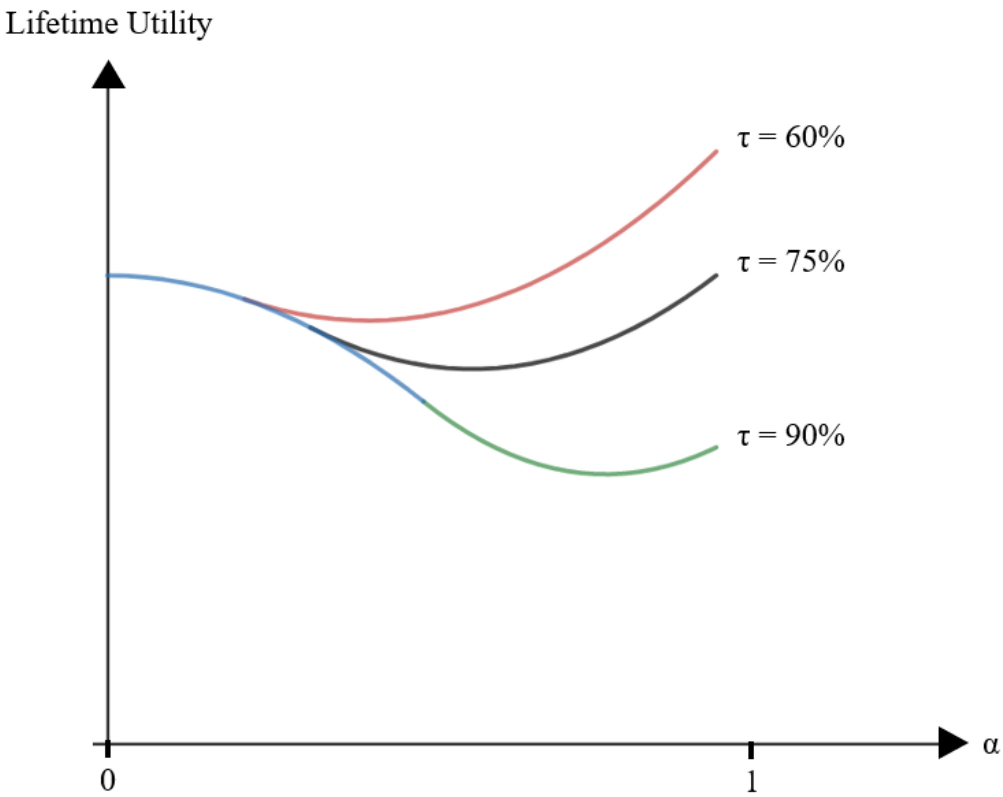

Figure 3: USG-curve for different levels of transaction costs

Second, an agent's lifespan has no effect on the optimal choice for time allocation. Mathematically, every point on the USG-curve is scaled by the square of lifespan, so the shape of the curve remains the same, regardless of lifespan.

In the appendix, we prove that the autarky curve is strictly decreasing as α rises in this interval. This implies that without access to trading, the best strategy is to be a generalist. This fact, along with the convex shape of the trade curve and the continuous nature of the USG-curve as a whole, we find that utility can only be maximised for values $\alpha = 0$, $\alpha = 1$ and, by symmetry, $\alpha = -1$. The conclusion reached is therefore that only a pure generalist or a pure specialist can maximise utility. Which one is superior is determined by transaction costs. The two strategies are equally viable at $\tau = 75\%$. For transaction costs below this level, specialists enjoy greater utility.

We use this finding to justify viewing α as a discrete variable from here on out. We reduce the variable to two potential strategies: a generalist strategy ($\alpha = 0$) and a specialist strategy ($\alpha = 1$ or $\alpha = -1$ with even distribution). We also relax the assumption that α is fixed at its initial value.

The second part of the intermediary model will explore the consequences to utility when an agent who has previously been a generalist becomes a specialist, in the presence of transaction costs. The transition to specialising entails an instantaneous drop in utility per unit time. Intuitively, utility drops because the total production of goods is only marginally higher than it was before the transition. The presence of transaction costs means therefore that consumption must be strictly lower just after than it was just before the transition.

In the appendix, we prove that the rate at which utility increases is higher for a specialist than for a generalist if transaction costs are lower than 75%. The model therefore demonstrates that in the range of $0 < \tau < 75\%$, a transition from generalist to specialist entails a short-term cost and a potential long-term benefit. The picture below illustrates this effect for transaction costs of 50%.

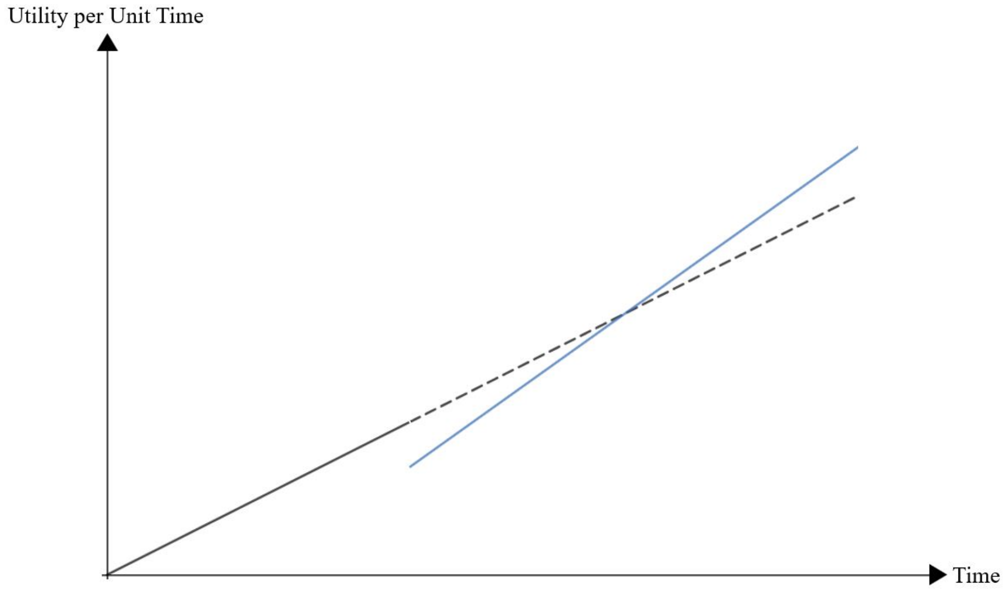

Figure 4: Development of utility after transitioning from generalist to specialist

The initial black line shows the development in utility while the agent acts as a generalist. The blue line shows the development in utility after transitioning to a specialist. This can be compared to the dotted black line which shows the utility that would have been, if the agent had remained a generalist. The predicted effect can be observed here, utility falling initially to later catch up.

3.3 Random Matching Models

The chief purpose of these models is to analyse how the utility associated with either strategy depends on the strategies used by others in the economy. To this end, the models we set up were based on the following assumptions:

- Time allocation is a discrete variable which can be described in terms of two strategies: a generalist strategy and a specialist strategy. Generalists divide their time equally between both goods. Specialists allocate all their time to one good. Half of all specialists produce apples and the other half produce bananas.
- Productivity is developed during an initial practice period of fixed length during which nothing is produced. The time allocation during this practice is the same as it will be during production for that agent.
- The mechanism by which a transaction is determined is varied between two random

matching models. In the first model, the relative price is fixed at one and goods will be traded under this condition until at least one party no longer wants to trade. In the second model, a random agent in each pair is assigned the role of offeror. This agent can suggest any trade, and the other agent will accept as long as their utility does not become worse off. The offeror maximises their own utility under this condition.

- If an agent is left with only one type of good after their option to trade, they will get no utility from consuming it. They will then choose to save these goods for the next period. Goods cannot be saved for more than one period.
- No transaction costs are present in these models.

For reasons outlined in the appendix, a specialist produces twice the total quantity of goods than generalists do. The non-dynamic nature of skill acquisition in these models makes expected utility constant between time periods. Thus, instead of lifetime utility on the vertical axis, it is labeled expected utility per time instance. We model this variable as a function of the portion specialists in the population, a variable labeled S in the appendix.

3.3.1 Relative Price Fixed

We will first analyse the random matching model in which the relative price is fixed to one. Under these circumstances, a generalist will never want to trade, so their expected payoff per time period is independent of the distribution of strategies in the population. A specialist in an economy of only generalists will never have an opportunity to trade and will therefore experience zero utility. As the number of specialists in the economy rises, the opportunity for a specialist to trade will increase.

The utility achieved by a specialist depends on who they are matched with and whether either party has saved goods. Expected utility therefore becomes a weighted average of the utility associated with each scenario. The probability of ending up in each one depends on the distribution of strategies in the population. Mathematical details are derived in the appendix.

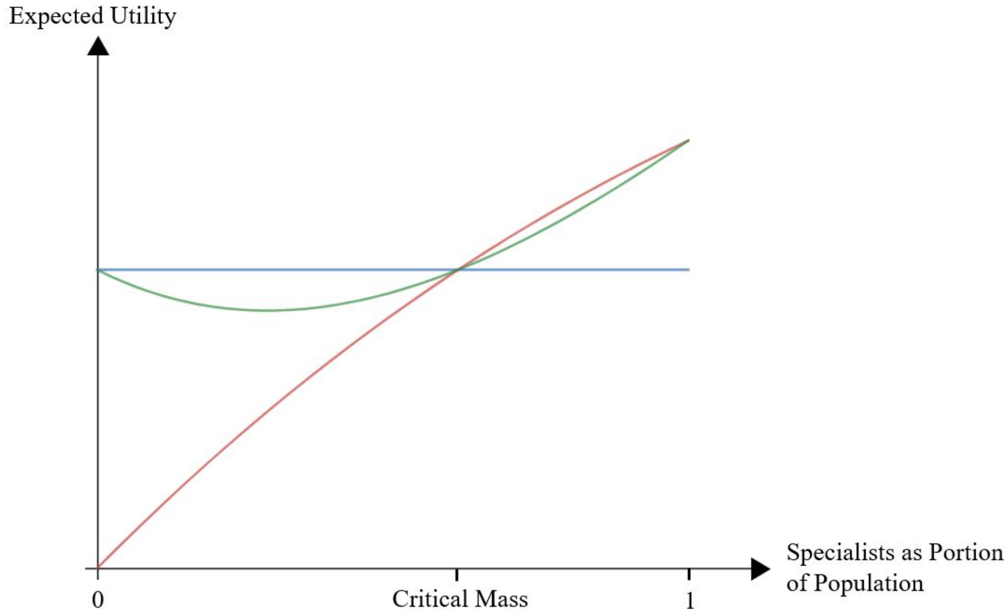

Figure 5: Utility of each strategy for different population distributions of strategies

Inserting the probabilities and their associated utilities into a graphing calculator yields the graph above. The blue line shows the utility of the generalist (the generalist curve), the red curve shows the utility of specialist (the specialist curve), and the green curve shows the weighted average representing the expected utility in the whole economy.

The green line is maximised where the whole population consists of specialists. In other words, this is the socially optimum state. We notice, however, that when the portion of specialists in the population is low, it is beneficial to be a generalist.

We can do a basic analysis based on evolutionary game theory by adopting the simple principle of replicator dynamics that the group with higher expected utility will replicate more rapidly. For levels of specialisation below a critical mass, generalists will proliferate. Above this critical mass, specialists will proliferate. The two population distributions which correspond to stable steady states are $S = 0$ and $S = 1$. Based on population expected utility, $S = 1$ is clearly preferred. However, an economy that at some point ends up below the critical mass will likely find itself stuck with a generalised economy in the long run.

It is worth noting that while expected utility is higher for the specialist at $S = 1$, the variance of expected utility is lower for the generalist. Because the problem is modeled in continuous time, each time instance independent of the previous, this variance is less significant than it would have been in discrete time. However, in a more realistic setting, having reliable access to all goods would be valued by the agents.

3.3.2 Pareto Efficient Trading

We will now analyse a similar problem but with the key difference that the relative price at which transactions take place is not fixed at one. Instead, the rules by which the transaction is determined are such that a randomly assigned offeror in each pair is free to decide the transaction that will maximise their own utility under the condition that their counterpart is not left worse off than they would be if no trade was made. This algorithm guarantees a Pareto efficient allocation of goods. It can be considered fair because the offeror is assigned randomly.

Under this rule, unlike the previous rule, transactions will take place between generalists and specialists. If the generalist becomes offeror, they will take all of the specialist's goods (recall that specialists gain no utility from consuming only what they produce). If the specialist becomes offeror, they will trade much of the good they specialise in for some of the other good. Because generalists will still not trade with other generalists (trade is not beneficial if the two parties have equal *marginal rates of substitution*), they will now benefit from higher levels of S. The result of these two effects is that the left part of the specialist curve and the right side of the generalist curve are both increased from before.

By assigning appropriate weights and their associated utilities we retrieve the graph displayed above. The colour codes are identical to before. We can observe that the total expected social benefit (the green curve) strictly increases as more of the population specialises. The socially optimal outcome is therefore, as in the previous model, one where every agent specialises.

The interesting difference from the previous model is that at any distribution of strategies in the population, the expected payoff of the generalist strategy exceeds that of the specialist, implying that the generalist strategy is evolutionarily stable. The situation that emerges from this set up is similar to the well-known game of the prisoners' dilemma. What characterises a prisoners' dilemma is that the Pareto efficient outcome that occurs when everyone cooperates requires each party to disregard their own best interest. For this reason, the Nash equilibrium (here also the evolutionary stable strategy) is associated with a sub-optimal outcome.

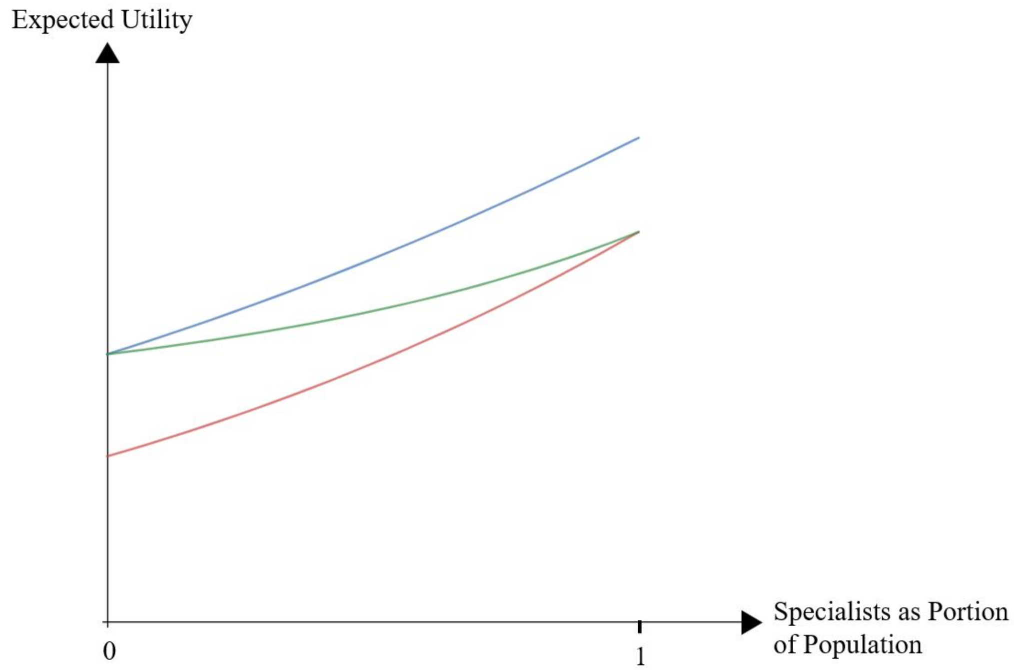

Figure 6: Utility of each strategy for different population distributions of strategies

4 Conclusion

The three models presented in the analysis all describe challenges associated with specialising, both on an individual and on a societal level.

In the intermediary model, the transition from generalist to specialist entails an initial sacrifice of utility. A rational agent would be willing to make this sacrifice if the long-term benefit outweighs this initial cost. However, because people frequently choose short-term gratification over long-term benefits, it is not obvious that everyone would choose to specialise when they recognise that this is the advantageous option. Therefore, the conclusion drawn from this model is that an agent acting in their own short-term interest may remain a generalist to the detriment of their own long-term interest.

The two models based on random matching share many similarities. In both models, we described the expected utility associated with each strategy and total societal payoff as a function of the distribution of strategies in the population. Both were set up in a way so that the optimal social outcome is achieved in a specialised economy.

The first of the two models required a critical mass of the agents in the economy to be made up of specialist in order for specialising to become the advantageous strategy. In other words, generalists win in a generalised economy and specialists win in a specialised economy. Under this model, an economy which begins in a generalised state will not naturally specialise. If a collective effort of retraining is undertaken to achieve a more specialised economy, it will only be successful if enough agents participate in the effort. A government that believes the model describes its economy should subsidise specialisation until it believes with confidence that the critical mass has been surpassed. Once this state has been reached, the government can step back and if the judgement was correct, the economy will stabilise in a specialised state.

The second random matching model displays a prisoners' dilemma. Here, as the economy becomes increasingly specialised, total utility increases. In other words, it is in the collective interest that the economy be as specialised as possible. The game theoretical analysis shows, however, that the collective interest is in direct opposition to the individual interest. Regardless the distribution of strategies in the population, it is never beneficial for a single agent to specialise. Because of the misalignment between the interest of the individual and that of society, a government that believes its economy is in a situation like the one described in the model should consider permanently subsidising specialisation.

With these models, we present a hypothesis of a situation that we call *the generalist trap*, where specialisation on a societal level is the desired outcome, but barriers of various forms prevent this from being achieved. As our models show, the generalist trap can vary in nature. It can originate from a failure to adopt a long term perspective or it can arise from the game theoretical implications of the market. Moreover, the effect can be temporary or permanent, depending on whether the specialised economy constitutes a stable steady state.

A discussion about potential tests of this work is warranted. An economic model must by definition make simplifying assumptions. The key to a good model is to make these simplifications in a way that the model can still provide useful insights and predictions relating to empirical phenomena. Therefore, to know whether our models are useful, they should be tested using empirical data.

Ideally, we would like to suggest a range of econometric tests which could falsify our hypothesis of the generalist trap. It is a difficult hypothesis to test for a number of reasons. We will discuss the issue for each model separately.

The intermediary model predicts how the utility of agent that specialises will develop over time. Utility is not possible to measure perfectly, but self-reported happiness or income can potentially be used as a proxy. If a data set of self-reported happiness or income over time of generalists who specialise can be found and compared with generalists who make no change, then perhaps the hypothesis can be tested. However, generating such a specific data set is idealistic at best and impossible at worst.

To test the random matching models we would need data sets showing how the degree of specialisation within an activity tends to change over time. If an economy starts becoming increasingly specialised, but soon enough returns to where it started, this may lend credence to either model. If the specialisation process does not reverse but accelerates after some point, this could lend credence to the first random matching model. A significant drawback with these tests is that there is no outcome in which the models could clearly be falsified. Moreover, we cannot isolate the variable of other agents' strategies from effects related to technology or habit.

5 Discussion

We will discuss below a number of limitations of the models. The limitations generally stem from simplifying assumptions necessary to derive clear results.

The trade function determines the degree to which the distribution of strategies used in the economy affects the viability of any one strategy. In the intermediary model, there was no such effect. In the random matching models, on the other hand, the effect was strong. Reality most likely resembles something between these two options. The viability of specialising is certainly influenced by the choices of others in the economy, but the interactions in the economy cannot realistically be considered random. Rather, specialists would be expected to gravitate toward finding specialists in the other good.

This framework also ignores the possibility of agents to form organisations to utilise economies of scale. As previously mentioned, economies of scale, in combination with a larger market size, are potentially drivers for specialisation (Krugman et al., 2015).

Another limitation is that we have chosen not to model the option to exchange goods for money, which in turn can be used to acquire goods for consumption. This decision was made to take a unique approach to the problem of specialisation. If goods can be traded for money then the conclusion that individuals will invest their time where they enjoy the highest hourly wage, adjusted for enjoyment and earnings development, is relatively evident. Though this makes our model less realistic, it lets us explore aspects of the problem previously unexplored.

6 Appendix

This section contains all mathematical proofs required to support the analysis. The proofs are organised by model and ordered as in the analysis.

6.1 Intermediary Model - Calculations

Subscript A denotes apples and subscript B denotes bananas. At each time instance, an agent maximises their utility from consumption to determine how much to trade:

$$\max U_t = \sqrt{C_{At} * C_{Bt}} \tag{1}$$

Suppose the agent chooses to trade apples for bananas. If X is the number of apples sacrificed then $C_{At} = P_{At} - X$ and $C_{Bt} = P_{Bt} + (1 - \tau) * X$. We can now rewrite the maximisation problem as:

$$\max_X U_t = \sqrt{(P_{At} - X) * (P_{Bt} + (1 - \tau) * X)} \tag{2}$$

Differentiating with respect to X gives:

$$\frac{dU_t}{dX} = \frac{P_{At} * (1 - \tau) - P_{Bt} - 2 * (1 - \tau) * X}{2 * \sqrt{P_{At} - X} * \sqrt{P_{Bt} + X * (1 - \tau)}} = 0 \tag{3}$$

Solving for Xgives:

$$X = \frac{P_{At} * (1 - \tau) - P_{Bt}}{2 * (1 - \tau)} \tag{4}$$

We use the second order condition to ensure that we are, in fact, maximising utility with respect to X:

$$\frac{d^2U_t}{dX^2} = -\frac{(P_{At} * (1 - \tau) + P_{Bt})^2}{4 * ((P_{At} - X) * (P_{Bt} + X * (1 - \tau)))^{3/2}} < 0 \tag{5}$$

The second order condition holds. X must be a positive number. The denominator is positive for relevant values of τ so the numerator must be positive, thus:

$$P_{At} > P_{Bt}/(1 - \tau) \tag{6}$$

This is a necessary condition to be fulfilled for the agent to trade apples for bananas.

Inserting $X = P_{At} - C_{At}$ into equation 4 and solving for C_{At} gives:

$$C_{At} = \frac{P_{At} * (1 - \tau) + P_{Bt}}{2 * (1 - \tau)} \tag{7}$$

Inserting $X = (C_{Bt} - P_{Bt})/(1-\tau)$ into equation 4 and solving for C_{Bt} gives:

$$C_{Bt} = \frac{P_{At} * (1-\tau) + P_{Bt}}{2} \tag{8}$$

Inserting equations 7 and 8 into the utility function gives:

$$U_t = \frac{P_{At} * (1-\tau) + P_{Bt}}{2 * \sqrt{(1-\tau)}} \tag{9}$$

We can find the derivative of this utility with respect to transaction costs with the quotient rule:

$$\frac{dU_t}{d\tau} = \frac{P_{Bt} - P_{At} * (1-\tau)}{4 * (1-\tau)^{3/2}} \tag{10}$$

Which is negative if the necessary condition for trade $P_{At} > P_{Bt}/(1-\tau)$ holds. Thus the utility of an agent who trades goes down with as transaction costs rise.

Recall that production is constrained by $P_{At} = \frac{1+\alpha}{2} * M_{At}$ and $P_{Bt} = \frac{1-\alpha}{2} * M_{Bt}$. Because both M_{At} and M_{Bt} begin at zero and have time derivatives $\dot{M}_{At} = \frac{1+\alpha}{2}$ and $\dot{M}_{Bt} = \frac{1-\alpha}{2}$, then $M_{At} = t * \frac{1+\alpha}{2}$ and $M_{Bt} = t * \frac{1-\alpha}{2}$. Therefore, we conclude that production of each good can be written as the following function of time and α:

$$P_{At} = t * \left(\frac{1+\alpha}{2}\right)^2 \tag{11}$$

and likewise:

$$P_{Bt} = t * \left(\frac{1-\alpha}{2}\right)^2 \tag{12}$$

When these are inserted into the utility function the following expression is obtained for utility in a time instance:

$$U_t = \frac{t * (1+\alpha)^2 * (1-\tau) + t * (1-\alpha)^2}{8 * \sqrt{1-\tau}} \tag{13}$$

We differentiate utility with respect to time to obtain the rate at which utility increases from one period to the next:

$$\frac{dU_t}{dt} = \frac{(1+\alpha)^2 * (1-\tau) + (1-\alpha)^2}{8 * \sqrt{1-\tau}} \tag{14}$$

Note that this rate of change does not depend on the time variable t thus utility increases linearly over time. Total lifetime utility can therefore be calculated as the area of a triangle with base T and height U_T.

$$U = \int_0^T U_t = \frac{T * U_T}{2} = T^2 * \frac{(1+\alpha)^2 * (1-\tau) + (1-\alpha)^2}{16 * \sqrt{1-\tau}} \tag{15}$$

Now we compute the derivative of lifetime utility with respect to α to explore which level of specialisation maximises lifetime utility.

$$\frac{dU}{d\alpha} = T^2 * \frac{(1+\alpha)*(1-\tau)+(\alpha-1)}{8*\sqrt{1-\tau}} = 0 \tag{16}$$

The first order condition holds true for $\alpha = \frac{\tau}{2-\tau}$. We check the second order condition to know whether this is a maximum or a minimum.

$$\frac{d^2U}{d\alpha^2} = \frac{2-\tau}{8*\sqrt{1-\tau}} \tag{17}$$

The second derivative is positive for all values of $\tau < 1$. This implies that the point $\alpha = \frac{\tau}{2-\tau}$ is a minimum. Furthermore, we learn that the trade curve is strictly convex, implying that the trade curve can only be maximised at the edges of its domain. We now consider further the lower bound of the trade curve's domain. Previously we have defined this as $P_{At} > P_{Bt}/(1-\tau)$. Inserting $P_{At} = t * \left(\frac{1+\alpha}{2}\right)^2$ and $P_{Bt} = t * \left(\frac{1-\alpha}{2}\right)^2$ into this inequality we find it can be written as

$$\left(\frac{4}{\tau} - 2\right) * \alpha - \alpha^2 > 1 \tag{18}$$

To find the lower bound we transform this inequality into an equality and solve it using the quadratic formula. This yields:

$$\alpha = \frac{2-\tau \pm 2*\sqrt{1-\tau}}{\tau} \tag{19}$$

The greater root is greater than 1 and is therefore rejected. Thus we confirm that the lower bound of the trade curve's domain is given by:

$$\alpha = \frac{2-\tau - 2*\sqrt{1-\tau}}{\tau} \tag{20}$$

Thus the domain of the trade curve is given by:

$$\frac{2-\tau - 2*\sqrt{1-\tau}}{\tau} < \alpha \leq 1 \tag{21}$$

As previously stated, on the trade curve lifetime utility may only be maximised at the edges of its domain, that is $\alpha = \frac{2-\tau-2*\sqrt{1-\tau}}{\tau}$ or $\alpha = 1$. After having examined the trade curve, we will now consider the shape of the autarky curve more closely. The points on this curve are characterised by the fact that consumption of each good equals production, thus utility at a time instance is given by:

$$U_t = \sqrt{P_{At} * P_{Bt}} = t * \frac{(1+\alpha)*(1-\alpha)}{4} \tag{22}$$

We differentiate with respect to time to find the rate at which utility increases.

$$\frac{dU}{dt} = \frac{(1+\alpha)*(1-\alpha)}{4} \tag{23}$$

Just as with the trade curve, the time derivative of the autarky curve is not dependent on time. Thus utility increases linearly whether or not one trades. We find total lifetime utility similarly to how we did for the trade curve, by calculating the area of a triangle.

$$U = \int_0^T U_t = \frac{T*U_t}{2} = T^2 * \frac{(1+\alpha)*(1-\alpha)}{8} \tag{24}$$

We now differentiate with respect to α to see where the autarky curve is maximised.

$$\frac{dU}{d\alpha} = -\frac{T^2*\alpha}{4} \tag{25}$$

The first order condition holds only when $\alpha = 0$. We explore if the second order condition holds:

$$\frac{d^2U}{d\alpha^2} = -\frac{T^2}{4} \tag{26}$$

Which is negative for all positive values of T. We then conclude that if trade is unavailable, lifetime utility is maximised at $\alpha = 0$. Let us now consider the full USG-curve, the combination of the relevant parts of the trade curve with the relevant parts of the autarky curve. Transactions will only take place when they increase an agent's utility. Therefore, the USG-curve, its domain being $0 \leq \alpha \leq 1$, will consist of the trade curve in the previously defined domain of $\frac{2-\tau-2*\sqrt{1-\tau}}{\tau} < \alpha \leq 1$ and of the autarky curve in the remaining domain of $0 \leq \alpha \leq \frac{2-\tau-2*\sqrt{1-\tau}}{\tau}$. To ensure that the USG-curve does not jump at the point of marginal trade, we define the equation that gives the marginal benefit of the first sliver of a unit traded as a function of α. We do this by evaluating the derivative in equation 3 at $X = 0$ and rewriting the production in terms of t and α.

$$\frac{dU_t}{dX}|(X=0) = t * \frac{\left(\frac{1+\alpha}{2}\right)^2 * (1-\tau) - \left(\frac{1-\alpha}{2}\right)^2}{2*\sqrt{\left(\frac{1+\alpha}{2}\right)^2 * \left(\frac{1-\alpha}{2}\right)^2}} \tag{27}$$

This expression, as expected, equals zero at the point of marginal trade. At this point, therefore, the utility evaluated on the trade curve or the autarky curve must be equal. This means that the point of marginal trade cannot be a global maximum of the USG-curve, because the value on the autarky curve must be higher at $\alpha = 0$. Thus we conclude that the only points that could potentially correspond to the global maximum of the USG curve are $\alpha = 0$ or $\alpha = 1$ and $\alpha = -1$. The latter two are by symmetry associated with the same utility level.

Next will be presented the calculations associated with a transition from generalist to specialist. Suppose a generalist who decides to specialise has a productivity of $2k$ in each good. The instance before they make the transition they produce and consume k units of each good and therefore earn k utility. Equation 4 tells us that regardless of transaction costs, specialists will want to give up half of their produced goods in a transaction. This means that in the first instance after the transition, the agent will consume k of the good they specialise in and $(1 - \tau) * k$ of the other good. The utility they will enjoy is therefore $\sqrt{1 - \tau} * k$. In other words, the transition will instantly reduce their utility per unit time to $\sqrt{1 - \tau}$ of what they enjoyed before.

Next, we shall examine the rate of increase in utility. According to equation 23, a generalist's utility increases by 1/4 units per time period. For specialist we find the corresponding number in equation 14, corresponding to $\sqrt{1 - \tau}/2$. For all numbers of τ less than 75% this growth rate exceeds that for the specialist.

We now calculate how long it will take for the agent to catch up to their counterfactual generalist utility level (that which they would have enjoyed had they remained generalists). $4k$ units of time has passed from the beginning when they decide to switch strategies. We label the time beyond this point t. If one continues as a generalist their utility t time periods after $4k$ will be:

$$U_t = k + \frac{t}{4} \tag{28}$$

After becoming a specialist, utility will follow this trajectory:

$$U_t = k * \sqrt{1 - \tau} + t * \sqrt{1 - \tau}/2 \tag{29}$$

To determine the point where these values for the utility are equal, we set the expressions equal to one another and solve for t. This gives:

$$t = \frac{4 * k * (1 - \sqrt{1 - \tau})}{2 * \sqrt{1 - \tau} - 1} \tag{30}$$

This is the time it will take beyond that point to reach the same level of utility per time period as one would have enjoyed, had one remained a generalist. To find the time it takes until total lifetime utility is higher than it would have been without a change in strategy, we double this time.

$$t = \frac{8 * k * (1 - \sqrt{1-\tau})}{2 * \sqrt{1-\tau} - 1} \tag{31}$$

We see that this time increases with k. To know how it changes with τ we use the quotient rule to take the first derivative with respect to τ.

$$\frac{dt}{d\tau} = \frac{4 * k}{(2 * \sqrt{1-\tau} - 1)^2 * \sqrt{1-\tau}} \tag{32}$$

The derivative is positive in the relevant range $0 < \tau < 75\%$. This means that the time it takes until utility has caught up is greater the longer one has already lived as a generalist and the higher transaction costs are.

6.2 Random Matching - Calculations

First we derive the specialist and generalist curves for the random matching model where the relative price is fixed to one. Using the definition of productivity in a learning by practice framework and the productivity function, we see that generalists produce $k/4$ units of each good in a time instance and specialists produce k units of the good they specialise in, twice the total units produced of a generalist. The utility achieved by a specialist is determined by the strategy used by their counterpart and whether either party has goods saved from the previous instance. We organise the outcomes in a tree diagram, specifying probabilities and utilities associated with each outcome. S denotes the portion of the population made up by specialists.

Expected utility can be written as a weighted average of the potential realisations of utility, yielding the following formula:

$$\frac{S}{2} * \left(\frac{S}{2} * 0.5k + (1 - \frac{S}{2}) * \left(\frac{S}{2} * \sqrt{0.75}k + (1 - \frac{S}{2}) * k\right)\right) \tag{33}$$

We insert the formula into a graphing calculator to obtain the graph displayed in the analysis. Now we compute the formulas related to the specialist and generalist curves in the second random matching model, where transactions outcomes are Pareto efficient. Suppose agent 1 becomes the offeror of the transaction, producing P_{A1} apples and P_{B1} bananas. Agent 2, the non-offeror, produces P_{A2} apples and P_{B2} bananas. The superscript on utility the role in the transaction. O denotes offeror and N denotes non-offeror. Superscript E denotes expected utility. Agent 1 faces the following constrained maximisation problem.

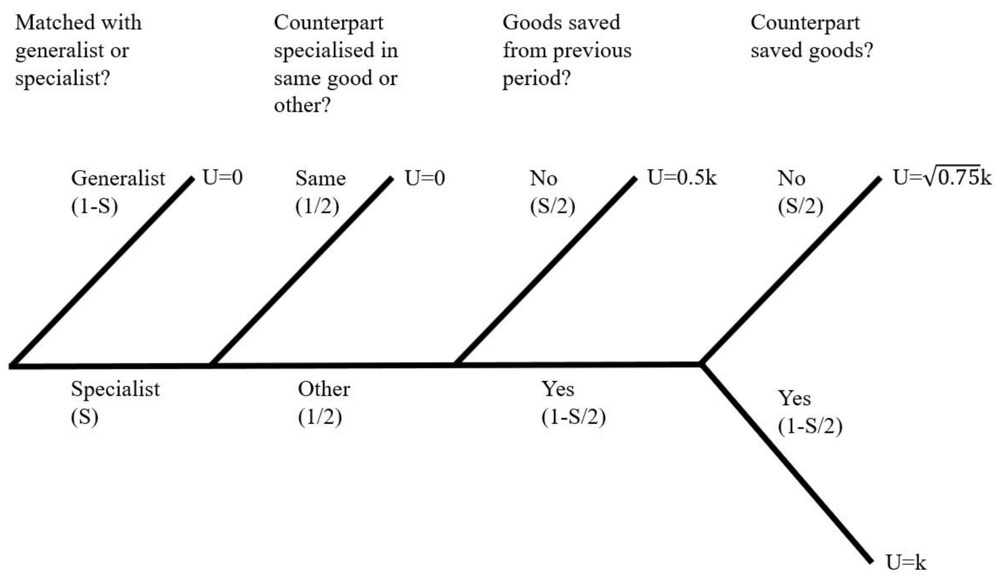

Figure 7: Tree diagram showing potential outcomes for a specialist

$$\max_{T_A,T_B} \quad U_1^O = \sqrt{(P_{A1}+T_A)*(P_{B1}+T_B)}$$
$$\text{subject to} \quad U_2^N = \sqrt{(P_{A2}-T_A)*(P_{B1}-T_B)} \geq \sqrt{(P_{A2})*(P_{B2})}$$

Where T_A and T_B are agent 1's net gain of either good in the transaction. Because utility strictly increases with consumption, the condition will be binding. By isolating T_B we find can rewrite the condition as:

$$T_B = P_{B2} - \frac{P_{A2}*P_{B2}}{P_{A2}-T_A} \tag{34}$$

By replacing T_B in the maximisation problem with this expression, the optimisation becomes unconstrained and with T_A as the only variable to choose.

$$\max_{T_A} \quad U_1^O = \sqrt{(P_{A1}+T_A)*(P_{B1}+P_{B2}-\frac{P_{A2}*P_{B2}}{P_{A2}-T_A})}$$

Because the square root is a strictly increasing function, maximising the expression without the square root will give the same results as it will with. We find the first order condition and isolate T_A to find:

$$T_A = P_{A2} - \sqrt{\frac{P_{A1}*P_{A2}*P_{B2}+(P_{A2})^2*P_{B2}}{P_{B1}+P_{B2}}} \tag{35}$$

We verify that this utility is maximised by seeing if the second order condition holds:

$$\frac{d^2U_1^O}{dT_A^2} = -\frac{P_{A2} * P_{B2}}{(P_{A2} - T_A)^2} * (P_{A1} + P_{A2}) \tag{36}$$

The second order condition holds unless one of the following holds true: $P_{A2} = 0$ or $P_{B2} = 0$. We notice also that the traded quantities are undefined if $P_{A2} = 0$ or $P_{B2} = 0$. In other words, the outcome is ambiguous if the non-offering party is a specialist. We consider this case intuitively instead. If the non-offeror is a specialist, they receive no utility from consuming what they produce, so the offeror will suggest a trade that gives them all of the non-offeror's goods (unless both parties are specialists in the same good). If an infinitesimally small number is added to the denominators of the fractions of each trade function, the utility retried reflects this intuitive finding. In summary, we can write utility for an offeror in the following way:

$$U_1^O = \sqrt{(P_{A1} + T_A) * (P_{B1} + T_B)} \tag{37}$$

Where:

$$T_B = \lim_{q \to 0} \quad P_{B2} - \frac{P_{A2} * P_{B2}}{P_{A2} - T_A + q} \tag{38}$$

$$T_A = \lim_{q \to 0} \quad P_{A2} - \sqrt{\frac{P_{A1} * P_{A2} * P_{B2} + (P_{A2})^2 * P_{B2}}{P_{B1} + P_{B2} + q}} \tag{39}$$

Now suppose agent 1 is the non-offeror and agent 2 is the offeror. In this case, agent 1's utility cannot be lower than it would be because that would break the rules of transaction. It also cannot be higher, because that would leave room for agent 2 to further raise their own utility. Utility for the agent 1 as a non-offeror, therefore, is simply their utility without trade.

$$U_1^N = \sqrt{P_{A1} * P_{B1}} \tag{40}$$

Agent 1 has equal chance to become offeror as they have to become non-offeror. Therefore, expected utility is simply the average of these two utility levels.

$$U_1^E = \frac{\sqrt{(P_{A1} + T_A) * (P_{B1} + T_B)} + \sqrt{P_{A1} * P_{B1}}}{2} \tag{41}$$

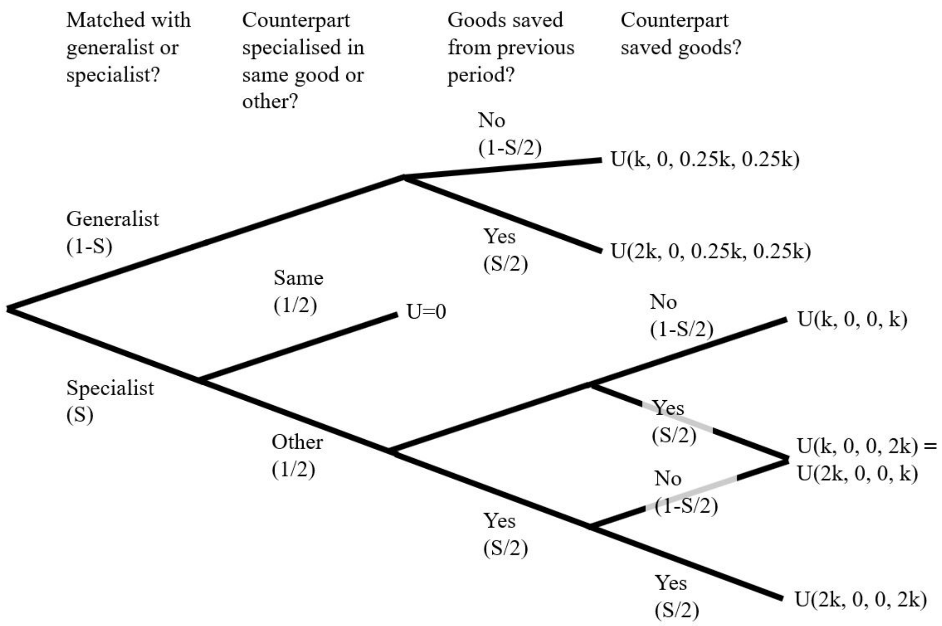

Figure 8: Tree diagram showing potential outcomes for a specialist

Where T_A and T_B are defined according to equations 38 and 39. Because the traded quantities are functions of what either party produces, expected utility can be calculated as a function of only both parties' produced quantities. We write this function as $U(P_{A1}, P_{B1}, P_{A2}, P_{B2})$. Like in the first random-matching model, we consider every possible outcome with its associated probability and utility according to a tree diagram. We assume that no goods will be exchanged when two specialists in the same good are matched, because the offeror would not gain any utility from taking the non-offeror's goods. Below we have mapped the possible outcomes for a specialist:

The expected utility for a specialist as a function of S can therefore be graphed according to the formula below:

$$
\begin{aligned}
U_S^E = (1-S) * ((1-S/2) * U(k, 0, 0.25k, 0.25k) + (S/2) * U(2k, 0, 0.25k, 0.25k)) \\
+(S/2) * ((1-S/2)((1-S/2) * U(k, 0, 0, k) + (S/2) * U(k, 0, 0, 2k)) \\
+(S/2)((1-S/2) * U(2k, 0, 0, k) + (S/2) * U(2k, 0, 0, 2k))
\end{aligned}
\tag{42}
$$

We must also determine the equivalent formula for generalists. To this end, we draw the corresponding tree diagram:

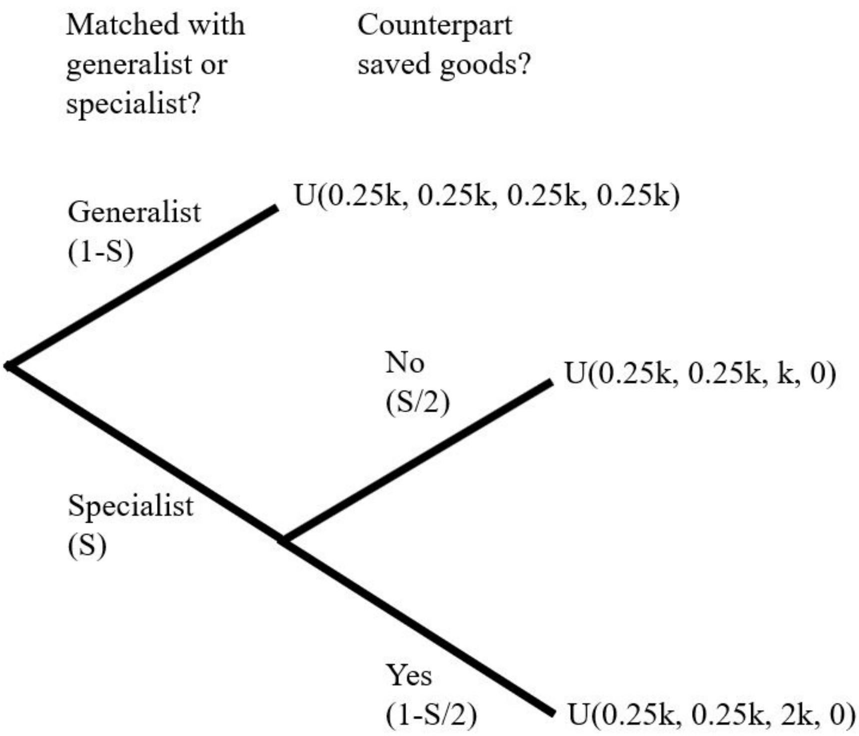

Figure 9: Tree diagram showing potential outcomes for a generalist

The following formula then gives the generalist curve in the analysis.

$$\begin{aligned} U_G^E &= (1-S) * U(0.25k, 0.25k, 0.25k, 0.25k) \\ &+S * ((S/2) * U(0.25k, 0.25k, k, 0) + (1-S/2) * U(0.25k, 0.25k, 2k, 0)) \end{aligned} \tag{43}$$

7 References

Acemoglu, Daron. (2009). Human Capital and Economic Growth. Introduction to Modern Economic Growth. *Princeton: Princeton University Press,* 359-383.

Fallick, Bruce. (1993). The Industrial Mobility of Displaced Workers. *Journal of Labor Economics, 11*, 302-323.

Fenster B. Charles, Armbruster W. Scott , Wilson Paul, Dudash R. Michele, Thomson D. James. (2004). Pollination Syndromes and Floral Specialization. *Annual Review of Ecology, Evolution, and Systematics, 35:1*, 375-403.

Helpman, Elhanan. (1981). International Trade in the Presence of Product Differentiation, Economies of Scale and Monopolistic Competition: A Chamberlin-Heckscher-Ohlin Approach. *Journal of International Economics, 11:3*, 305-340.

Henrich, J., Boyd, R. (2008). Division of Labor, Economic Specialization, and the Evolution of Social Stratification. *Current Anthropology, 49:4*, 715-724.

Herfindahl, O. C. (1950). Concentration in the Steel Industry. Ph. D. thesis, *Columbia University*.

Hirschman, A.O. (1964). The Paternity of an Index. *The American Economic Review, 54*, 761-762.

Hoover, E. M. (1936). The Measurement of Industrial Localization. *The Review of Economics and Statistics, 18:4*, 162-171.

Jones, W., Ronald. (1971). A Three-Factor Model in Theory, Trade and History, in Trade, Balance of Payments and Growth, ed. J. N. Bhagwati, R. W. Jones, R. A. Mundell. *North-Holland Publishing*, 3-21.

Kilbridge, M., Wester, L. (1966). An Economic Model for the Division of Labor. *Management Science, 12:6*, 255-269.

Krugman, P. (1991). Geography and Trade. *MIT Press*.

Krugman, P., Obstfeld, M., Melitz J. M. (2015). International Trade - Theory and Policy. *Harlow: Pearson Education Limited*.

Ohlin, G., Bertil. (1933). Interregional and International Trade. *Cambridge: Harvard University Press*.

Palan, Nicole. (2010). Measurement of Specialization - The Choice of Indices. *FIW - Research Center International Economics, Vienna, 62*.

Rana, J.S., Dixon, A.F.G. and Jarošík, V. (2002). Costs and Benefits of Prey Specialization in a Generalist Insect Predator. *Journal of Animal Ecology, 71,* 15-22.

Redding, Stephen. (2001). Specialization Dynamics. *Journal of International Economics, 58:2*, 299-334.

Ricardo, David. (1911). The Principles of Political Economy & Taxation. *London : New York :J.M. Dent; E.P. Dutton*.

Rosen, Sherwin. (1983). Specialization and Human Capital. *Journal of Labor Economics, 1:1*, 43-49.

Smith, Adam. (1776). An Inquiry into the Nature and Causes of the Wealth of Nations. *London: W. Strahan; and T. Cadell*.

Tress, R. C. (1938), Unemployment and the Diversification of Industry. *The Manchester School, 9*, 140-152.