THESE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par l' Université Toulouse 1 Capitole
Discipline : Sciences Economiques

Présentée et soutenue par

Héctor PIFARRÉ i AROLAS

Le 19 juin 2015

Titre :

Essays in Health and Demographic Economics

Ecole doctorale : Toulouse School of Economics

Unité de recherche : LERNA – TSE

Directeurs de Thèse : Franck Portier and Hippolyte d'Albis

JURY

Franck PORTIER, professeur, Université Toulouse 1 Capitole
Patrick FEVE, professeur, Université Toulouse 1 Capitole
Hippolyte d'ALBIS, professeur, Université Paris I Panthéon-Sorbonne
Ronald LEE, professeur, Berkeley, University of California
Emmanuel THIBAULT, professeur, Université de Perpignan

Essays in health and demographic economics

Héctor Pifarré i Arolas

We made it! Dedicat a l'Adeline, el Guillem i la Sonia. Sempre m'heu recolzat d'una manera o altra.

I would like to thank: Franck Portier for helping me get started and my supervisor Hippolyte d'Albis for his support throughout the process; Laura Pellisé, Natàlia Pascual and Helena M. Hernández for being so welcoming during my stay at the Center for Research in Health and Economics (CRES); Ronald Lee and Ryan Edwards for their hospitality and insightful discussions during my visit to UC Berkeley; Marc Saez for his generosity with his time and his excellent comments; my classmates in Toulouse: Joe, Margaret, Kartika, Daniel, Jorge, Hoch, Ananya, Leo, Nils, Sinem, Zeynep for their friendship and the fun times; Guillem Lopez for being an academic father to me; Adeline for her patience, her support and simply for being there; Sonia for making sure that it happened.

Abstract

This dissertation consists of four essays on health and demographic economics. In the first chapter, I explore the implementation of the theory of equality of opportunity (EOp) developed by Roemer (1998) to health in a joint work with Guillem Lopez and Frederic Udina. According to the theory of EOp, individuals' outcomes depend both on effort and circumstances and, since individuals are only reponsible for their effort, inequalities rooted in differences in circumstances ought to be compensated. A common impediment to the achievement of EOp applications with given resources constraints is that it is unlikely that public policies can fully compensate for existing unfair inequalities. This scenario is particularly relevant in the case of health policy, where public spending coexists with a large private spending component. We argue that if social justice is not attainable, social deliberation should not only focus on choosing the circumstances that ought to be compensated but also reflect on which groups suffering unfair inequalities should be prioritized. Roemer's proposal is to compensate individuals who are the worst off. However, this Ralwsian approach ignores the capacity to benefit from the recipients of public help, a shortcoming that becomes especially problematic in applications to health. Instead, we propose a framework in which society decides on the weight to place on the potentially conflicting objectives of helping individuals (or groups) who would benefit the most or individuals who are the worst off.

The second chapter examines the impact of income-related reporting heterogeneity on the measurement of health inequality. While most studies of health inequality rely on self-reported measures of health, recent research has studied the possibility that part of the existing differences in self-reported health could be due to systematic differences in reporting across socioeconomic groups. The concern is that part of the existing inequalities may not be founded on differences in the "true" health status of individuals. In particular, some studies have concluded that reliance on self-reported health might have resulted in an overstatement on the degree of health inequality of some countries.

I study the income-related reporting heterogeneity hypothesis in the 2006 wave of the Catalan Survey of Health. My hypothesis is that the finding of higher income groups overstating their "true" health is related to a failure of accounting for health states for people suffering from similar conditions. Better living conditions and life styles may result in health gains that could explain part of the "excess health" reported by higher income groups. My results indicate that including life style variables in the analysis explains at least part of the reporting differences across income groups that were previously attributed to reporting heterogeneity. In the second part of this chapter I examine the contribution to existing income-related inequality of the factors involved in self-reported health, accounting for life-style differences. I find that the main contributor to health inequality is the disproportionate concentration of the prevalence the reported conditions in lower income groups. However, the second largest contributor is the difference in educational attainment. This finding could be attributed to reporting heterogeneity but an alternative explanation is that, similar to the case of income, reporting differences might in fact be capturing underlying differences in objective health statuses.

The third chapter, joint with Hippolyte d'Albis and Loesse Jacques Esso, studies the trends in mortality convergence across developed countries from 1960 to 2008. While the epidemiological transition has provided a theory behind the expectation of convergence in mortality patterns, our results reject the convergence hypothesis for a sample of industrialized countries. We study the disparities across the mortality distributions of the countries and our sample and find no evidence of convergence towards a common mortality distribution. After a short period of convergence in the 1960s, countries have diverged for the rest of the period, with a pronounced increase in the divergence speed at the end of the 1980s. We show part of the observed increase in differences, especially at the end of the 1980s, is explained by the trends within former Socialist countries (Eastern countries) in Europe. Western countries, or countries with no Socialist past, have remarkably stable cross country differences in mortality patterns. Eastern countries, on the other hand, demonstrate convergence up until the 1990s, when they abruptly diverge. Our analysis also reports that, for both Western and Eastern countries, the relationship between life expectancy and the variance of the mortality distribution has radically altered over the period. While Western countries have transitioned from a strong negative correlation between life expectancy and variance to no association, Eastern countries have experienced the opposite evolution. Taken together, our results lend support to the theories that postulate the importance of structural conditions that lead to multiple steady-state mortality distributions.

The fourth and final chapter of this dissertation examines the relationship between unemployment and fertility. I offer a possible explanation for the apparent contradiction between the empirical work that finds a negative relationship between unemployment and fertility and the theoretical work that emphasizes the lower opportunity cost of childbearing while unemployed. I reconcile these perspectives by distinguishing two forms of unemployment. The first form is structural unemployment while the second is cyclical unemployment, a less permanent component of unemployment that is linked to the economic cycle. I study both effects over the life cycle using cohort data on a panel of developed countries. I find that while structural unemployment has an unambiguous negative effect on fertility, reactions to cyclical unemployment depend on the age at which it is experienced. Cyclical unemployment has no effect at younger ages and a negative effect at the end of the fertile life, but it is associated with an increase in fertility in the middle periods. In the final part of chapter 4, I show how these findings are consistent with a narrative that emphasizes the role of career considerations in fertility choices. I am able to obtain results that are coherent with my empirical findings in a model of fertility timing once I account for career considerations. At younger ages, delaying fertility has a larger impact on future employability due to the decreasing returns to labor market experience.

Contents

List of Figures

List of Tables

CHAPTER 1

Limited budgets and equality of opportunity in health (with Guillem Lopez and Frederic Udina)

1.1. Introduction

While Roemer's theory of equality of opportunity (Roemer, 1998) has gained ground in numerous applications, an unexplored area has been its application to too-real situations where social justice is not attainable due to resource constraints. We argue that when it is not possible to achieve full equality of opportunity (EOp), in addition to identifying unfair sources of inequality, it becomes crucial to decide which groups to target with public compensation policies. A reason why equality of opportunity might not be attainable is because EOp policies are likely affected by several factors that disregard the principles of EOp. In the case of health care, untargeted public provision, such as universal health care systems, are a primary example of a public policy that does not necessarily contribute to the achievement of EOp. We show how through the existing inequalities created by these factors the funds to target EOp can be insufficient and we offer a framework to guide public choice in these situations.

In our scenario an EOp planner must make the following decision: given a number of groups of individuals facing unfair inequalities and a dearth of resources to compensate all individuals, which group should be targeted first? Roemer suggests a Rawlsian approach: compensate the groups that are worst off. We argue that this choice is not inherent to the ethical criterion of EOp. Crucially, a critique to his stance is that the resulting allocation does not account for the effectiveness of the policy in the reduction of *unfair inequalities*. Our work shows that accepting Roemer's formulation of the principle of EOp does not require implementation according to Rawlsian principles. We contribute to the theory of EOp by first showing that a separation of the ethical criterion from the method chosen to allocate the resources is possible while maintaining the core of the EOp criterion. We propose an allocation method whereby the social decision is broadened vis a vis selecting targeted groups under an explicit trade-off between efficiency and equity[1].

[1] Equality of opportunities has been studied according to different definitions by a number of works. For a comparison between Roemer's and Van der gaer's approach see Ooghe, Schokkaert and Van de gaer (2007).

To contextualize our argument we briefly review Roemer's proposal. In his book (Roemer, 1998), Roemer formalizes external circumstances and effort, individual responsibility, and public compensation, in a way that recovers the concept of individual responsibility. According to his formalization, individuals' outcomes are attributed to both effort and circumstances. Roemer proposes effort as the only legitimate source of differences, implying that inequalities rooted in circumstances should be compensated. Possibly, the major strength of Roemer's work is providing a framework for society to consider inequality, while leaving the responsibility for most of the normative content for society. Many authors have taken advantage of the flexibility of the framework and have applied Roemer's strategies for EOp in different fields. These include education (Peragine and Serlenga 2007 and 2009, Bratti 2008, Blanco and Villar 2010), healthcare (Williams and Cookson 2000, Fleuerbaey and Shockhaert 2009) and development aid (Llavador and Roemer 2001).

Returning to Roemer's original contribution, the normative choice open to society is the identification of tolerable sources of inequality; in Roemer's terminology, this is the distinction between circumstances and effort. The framework offers a very large degree of flexibility, since it is possible to encompass, within Roemer's theoretical construction, ethical criteria that may be considered to be on opposite ends of conceptions of justice. We recover utilitarianism when society judges individual outcomes as determined solely by effort and embrace Rawlsian ideals when outcomes are judged as dependent only on circumstances.

In this chapter, we regard Roemer's formalization of EOp as two separate contributions. First, it defines an ethical criterion based on an intepretation of equality of opportunity: any two individuals exerting the same effort should attain the same level of health. Second, it offers an allocation method such that a given budget meets EOp. In a context of scarce resources where social justice according to EOp might not be achievable, the allocation method allows a stage for identifying which groups to compensate, opening an important dimension for social deliberation. We propose a different allocation that accounts for the capacity to benefit of public compensation recipients. The rationale behind our proposal is that, in the application of EOp, society might also be concerned with the degree to which groups are expected to improve their health when given public help. We favour an allocation method based on a tradeoff between compensating the individuals who are worst off and those who benefit the most. Hence, our proposal effectively extends society's choice to cover all aspects of the application of the theory of EOp.

Our work is related to the literature on the so-called bankruptcy problems, which proposes solutions to the allocation of a divisible good[2] among agents when the total amount of the good is insufficient to cover all their demands. Although the problem in the bankruptcy solution literature (for an extensive survey of the literature, see Thompson 2003) resembles ours to a large extent, our formulation diverges in that we follow Roemer's formalization of an EOp ethical criterion. In a similar way, our particular proposal to implement EOp is related to that discussed by Herrero and Villar (1994). The authors present a method for the allocation of a public budget to different objectives when the available funds are insufficient to completely satisfy all the objectives. In their work, the authors study the properties of a number of sharing rules inspired by both the bankruptcy and axiomatic negotiations literatures. An important distinction between our work and the rules suggested in the bankruptcy literature (and in Herrero and Villar 1994) is that we introduce concerns for recipients' capacities to benefit from transfers.

The chapter is organized as follows. In the first section, we formally review the framework for EOp presented in Roemer (1998). The second section is devoted to justifying our allocation strategy and our priority setting method. Preliminary, we impose further assumptions on the health state functions and effort decisions. We continue by dividing between impediments to achieving a fair health status distribution from budgetary restrictions and from complexities of the compensation policy. We tackle them separately for

[2] The literature tends to focus on divisible goods; nevertheless, for an exploration of the problem under indivisible goods see Herrero and Martinez (2008)

clarity. First, in order to convey the basic intuition, we use a simplified case to present how budgetary limitations might preclude the social planner from achieving EOp in health regardless of the allocation method. We then illustrate how simple policies can only attain EOp in very specific settings. Since these situations are not the focus of our work, we offer a modified social objective which allows us to ignore the concerns regarding the otherwise complexity of the policies. In this modified setting, we present the derivation of our allocation method and discuss how the choice of different metrics allows us to cover the full spectrum of social decisions regarding which groups to intervene.

1.2. The formalization of EOp

In this section we briefly introduce basic elements of the theory of EOp. The problem, in short, is the following. Consider a set of individuals who can achieve a certain health status (or any other relevant dimension of welfare). Suppose the health status is a function of the amount of health care consumed by individuals, the effort they exert (individually), and their circumstances. Our goal is to decide on the distribution of a given public budget across individuals to achieve a health status distribution which meets the ethical criteria of EOp.

The basis of EOp's normative structure lies precisely in the distinction between circumstances and effort, which we elaborate in this section. Define *circumstances* to be that which society judges beyond the responsibility of the individual. *Effort* is constituted by all the actions that society judges as within the responsibility of the individual. It follows from the characterization of circumstances, that we may create a classification of individuals by types. A *type* is a subset of individuals who share the same circumstances relevant to the attainment of health. We can imagine this to be anything from their genetical predisposition to illnesses, education or their capacity to benefit from treatment, etc.

We denote $\top = \{1, \ldots, T\}$ as the set of T types in which we divide population. The relationship between resources, effort and health is given by the health status function. The *health status* of an individual is a function of the effort and resources allocated, indexed by the type.

We denote $u^t(x,e)$ as the health status function for type t, where x are the resources and e is the effort[3]. As stated before, society must choose an allocation of health care spending (potentially) dependent on effort and type. While we know that the aim of the policy is to attain EOp, let it be for now any rule which satisfies the following definition:

DEFINITION 1. (Policy) A policy is a T-tuple of functions which specify, for each type, the resources devoted as a function of effort. We denote it $\phi = \left(\phi^1, \ldots, \phi^T\right)$ and call each function ϕ^t an allocation rule. Then, $\phi^t(e)$ is the amount of resources a type t individual receives if she exerts effort e.

It is then reasonable to believe, in turn, that the effort exerted by individuals is dependent on the policy. Given a policy, the individuals of a given type t generate a distribution of effort, given by a cumulative probability function $F^t_{\phi^t}$; $F^t_{\phi^t}(e)$ is the cumulated probability up to and including e. In some cases we may assume that effort has a discrete probability distribution, in others we assume a continuous distribution with a convex support, say an interval (infinite or not) with a density function $f^t_{\phi^t}$.

We now specify the notation for the budget constraint. Let ω be the per capita disposable resources. The amount of (per capita) resources assigned to type t is

$$\omega^t = \sum_e \phi^t(e) P(e) \text{ or } \omega^t = \int_R \phi^t(e)\, dF^t_{\phi^t}(e)$$

depending on whether effort is a discrete or a continuous variable. Denote α^t the proportion of type t in the population. Then the global constraint is

3 Indexing by type is equivalent to including another variable or vector of variables which includes type characteristics, i.e., $u(C,x,e)$, where C is the vector of individual characteristics that constitute a type (see Roemer, 2002 for a complete exposition of this notation).

$$\omega = \sum_t \alpha^t \omega^t. \tag{1.2.1}$$

Let $\rho^t = \omega^t / \omega$ be the per capita share of the resource for type t. Then $\sum_t \rho^t = 1$.

A key contribution by Roemer is in formalizing a method to justly compare the effort exerted by different types. Given that the effort distribution is influenced by the circumstances, how can we *fairly* compare the amount of effort between different types? A solution is to draw a distinction between the *level* and the *degree* of effort as formalized in this definition.

DEFINITION 2. (Effort level, effort degree and indirect advantage) For $\pi \in (0,1)$, let $e^t(\pi, \phi^t)$ be the *level of effort* exerted by an individual of the type t in the π^{th} quantile of effort of the type . We call π a *degree of effort*. These levels and degrees are characterized by the equations

$$\pi = \int_0^{e^t(\pi,\phi^t)} dF^t_{\phi^t}, \qquad t \in \top$$

in the case of continuous effort distribution, and similarly for the discrete case.

The *indirect health status function* gives the health status of an individual of type t receiving the resources determined by the policy ϕ and exerting the π degree of effort of the type distribution of effort, and is defined by

$$v^t\left(\pi, \phi^t\right) = u^t\left(\phi^t\left(e^t\left(\pi, \phi^t\right)\right), e^t\left(\pi, \phi^t\right)\right)$$

The main idea of the concept of equality of opportunity is that we should be comparing individuals according to their degree of effort and not their level, which is influenced by their circumstances. We formalize this ethical criterion in the following definition.

DEFINITION 3. (Social criteria and strong social justice) The ethical criterion of *equality of opportunity* (EOp) states that given any two individuals, independent of type, exerting the same *degree* of effort (individuals in the same position in their respective type distributions), they must achieve the same level of advantage. That is,

$$\forall i, j \in \top,\ \forall \pi \in (0,1)\,, \quad v^i\left(\pi, \phi^i\right) = v^j\left(\pi, \phi^j\right) \tag{1.2.2}$$

By *strong social justice* according to a given criterion we indicate, the state in which the advantages of all individuals satisfy the requirements of the chosen ethical criterion, in this case, EOp.

Having defined the fair distribution of health status, the remaining question is how to design policies that achieve that social objective. Roemer proceeds in the following manner. In his presentation, he refers to EOp as in definition 3 (1.2.2). However, when implementating EOP — i.e. the choice of policy — Roemer defines the resulting policy from his proposed methodology as an EOp policy regardless of whether it actually achieves the EOp in definition 3. In a sense, one could argue that Roemer has made EOp a criterion subsidiary to the allocation method, as what is actually treated as EOp is the outcome of his rule for choosing the policy.

While one may accept his proposed method, which is presented in the remaining of the section, it does not follow directly from the theory of EOp. In other words, a society that accepts the normative criterion of definition 3 does not need to support the ethical choices inherent to his proposed method. This is where our contribution lies. While we accept EOp as a criterion, we propose a different methodology to guide the choice of policies. When we refer to EOp, we are considering the criterion in definition 3. This allows for the possibility that EOp is not actually achieved, something that is otherwise ruled out by the nature of Roemer's procedure. If we proceed according to this separation of *ethical criterion* and *method,* we can still discuss Roemer's rule for a choice of policy as a particular method. Recall that we seek a distribution of resources that leads to the achievement of EOp. Ideally, we aim at meeting this criterion for every percentile of effort, whereby all types accomplish the same level of advantage. The method proposed by Roemer is the maximization of the minimum advantage among types for every π of effort.

$$\max_{\phi} \min_{t \in \mathrm{T}} v^t \left(\pi, \phi^t\right)$$

Note that this allocation method is Rawlsian in its conception. It centers the attention on the type that is worst off, for every π of effort. This is the essence of Roemer's implementation and what follows are technical considerations. Given limitations on the sophistication of the implementation policy, the solution to the program for a given quantile might not correspond with the solution that equalizes the advantage among other quantiles. Therefore, we encounter the problem of potentially obtaining as many policies as quantiles. The proposed solution consists of assigning a weight equivalent to the population weight for every quantile and solving for this modified problem. Effectively, it concedes the same importance to every quantile, so in this sense, it becomes utilitarianist across quantiles.

$$\max_{\phi} \int_0^1 \min_{t \in \mathrm{T}} v^t \left(\pi, \phi^t\right) \, d\pi \tag{1.2.3}$$

The remaining part of the chapter is devoted to constructing and justifying our proposed allocation method.

1.3. Preliminaries

In this section we want to explore why is it that EOp might not be achieved and why this is particularly relevant in applications to health. We begin with the justification of the relevance of our proposal. Suppose first that EOp cannot be achieved, disregarding the reason. The main concern that motivates our contribution is that the Rawlsian approach to the choice of policy proposed in Roemer's original contribution is particularly ill-suited for health applications. This is because it excludes any consideration over the capacity to benefit from health care of the individuals to be treated among those facing unfair inequalities. In the extreme case where the type that is the worst off cannot improve with medical treatment, using *1.2.3* would still lead to allocating the entirety of the budget to this type. This is a criticism made to all the methods based on Ralwsian inspiration; interestingly, this precise point was raised by Harsanyi (1975) in his critique to Rawls' maxmin. The following example, part of Harsanyi's text, illustrates our critique.

As a first example, consider a society consisting of one doctor and two patients, where both patients are critically ill with pneumonia. Their only chance for recovery is via antibiotic treatment, but the amount of treatment available is sufficient for treatment of only one patient. Of these two patients, individual A is an otherwise healthy person (apart from his present attack of pneumonia). On the other hand, individual B is a terminal cancer victim whose life would be prolonged by merely several months, given treatment of the antibiotic. Which patient should be given the antibiotic? According to the difference principle, it should be given to the cancer victim, who is obviously the less fortunate of the two patients.

It is important to note that while one might reject *1.2.3* , this does not invalidate EOp as an ethical principle since *1.2.3* is not the unique program that implements EOp. Recognizing this fact, we propose abandoning the pure Ralwsian program and to instead design policies based on a more flexible framework. While society might be concerned with the well-being of types that are the worst off, the capacity to benefit should be a key component behind the adoption of a given policy. Our contribution is to allow society to choose the particular weight given to both potentially conflicting objectives in the program that obtains the EOp implementing policy. However, before we describe our proposal in detail, we clarify why we think it is most reasonable to assume that EOp cannot be achieved in the majority of applications.

In several contexts and most certainly in the case of health care policy, individuals have a starting level of advantage (health endowments or health gained by income, for example) that is not distributed following ethical considerations. The origin of these initial differences can be traced to a variety of factors. In the case of health these might be behavioral or biological reasons, such health related habits or genetics. Without disregarding the importance of such factors, in this chapter, we emphasize the role of health care spending in creating initial differences in health status. Regardless of the amount devoted within the health budget to the pursuit of EOp, there is bound to be a large fraction of total spending allocated without regard to the ethical principle of EOp. One might think of high income elasticity private spending on expensive healthy food, free time for exercise, health education and lower opportunity costs for prevention for certain groups; or public spending, perhaps with a universalist flavor, that is allocated on different premises. Regardless of the source of these initial differences in health status accross types, it is entirely possible and even reasonable to think that it might not be feasible to fully compensate those unfair inequalities with the budget allocated to EOp policies. It is in such context when the choice of a EOp implementing program becomes relevant. As an aside, another reason why strong social justice (according to EOp) might not be a feasible state is technical in nature. As explained in the previous section, there could be a conflict between EOp achieving policies. When the complexity of the policy is limited, this is more likely to become an issue. We illustrate this point later on for the case of policies that only depend on the type. In this section, we explore separately both reasons why a EOp distribution might not be feasible separately before we turn to our proposal.

1.3.1. Assumptions. Throughout the remainder of the chapter, we assume the following regarding effort decisions, advantage function, and behavior of individuals.

(1) Achievement functions $u^t(x,e)$ are defined for all nonnegative values of their arguments, unbounded for any fixed positive value of any of its arguments, and twice differentiable with continuity, i.e. C^2 functions. First partial derivatives are strictly positive and second derivatives with respect to the same argument twice are strictly negative.

(2) Both resources and effort are necessary and sufficient for obtaining a positive achievement: $u^t(x,e) = 0$ if and only if $x = 0$ or $e = 0$.

(3) The assignment policy is determined by Roemer's method, so it is the optimal solution of (*1.2.3*) when all resources are allocated. Furthermore, we restrict to constant policies.

(4) Denote e^* the individual choice of effort of any given type. Then, effort depends positively on the amount of resources allocated to this type: $\forall x$, $\frac{de^*}{dx} > 0$.

We clarify here the extent to which our assumptions are restrictive. Our first assumption is mainly technical, as it states that the health status function is well-behaved. However, we also assume that there is no upper bound to how much health can ameliorate with health care. That is, we allow for different capacities to benefit, but we assume that even though some individuals may have an arbitrarily small capacity to translate health care spending into a better health status, it is always possible to improve their health. The second assumption states that health status requires a positive amount of resources, but we do not require the resources to be fully determined by the social planner. Resources may originate from other public interventions (not aimed at achieving EOp), privately by individuals ("other resources"), or as allocated by the social planner. We do impose a restrictive assumption further along by requiring "other resources" to be unaffected by the social planner to obtain EOp. This simplifying assumption precludes important considerations such as crowding-out of private resources.

We restrict to constant policies in our exercise. While our message would remain mostly unchanged by allowing for more complicated policies, we choose to restrict our analysis to compensation policies that only depend on the type for two main reasons. First, this simplifies the presentation of our proposal and it presents our results under the realistic assumption that the policy space is much restricted. This could be because of the feasibility constraints related to its implementation or because of some political economy reasons. Finally, our most restrictive assumption states that allocating more resources to a type or group of individuals does not imply that the group will exert less effort. For the purposes of our exercise, we assume that the group of individuals will exert more effort. This is important for our implementation method since we rely on the possibility of achieving social justice. If agents diminish their efforts in a way that fully offsets the advantage gains garnered from allocation of further resources, our method would no longer result in a well-defined problem. Our claim, however, is that our implementation method can be applied even if there is some "crowding out" as long as there is an arbitrarily large public budget such that social justice can be attained. That is, that achieving EOp is feasible. We think that it would be far-fetched to assume otherwise, since it would imply the rather extreme vision that individuals would simply undo the public compensating effort by reducing their effort.

1.3.2. Insufficient budget. We present the context of an insufficient budget. In order to focus solely on this issue, we present a limit case by ruling out the possibility of conflicts in the policies required by each quantile by assuming a single level of effort for each type (although not necessarily the same level across types). That is, the distribution of effort is characterized by a single effort for each type with the immediate implication that EOp is achieved by completely equalizing advantages across types. This is the simplest scenario where the availability of public funds represents a limitation in the achievement of EOp. We want to emphasize that we use this simplistic scenario strictly for presentation purposes, as our proposal does not rely on single effort by type distributions. In fact, this is an assumption we abandon later on. Having established the aim of our exercise, suppose now that all the resources available for individuals are publicly provided and distributed according to the ethical objective of EOp. Then, as we formalize in proposition 1, strong social justice is always achieved in our context.

PROPOSITION 4. *Let each type have a unique nonzero effort level. Then, under assumptions 1, 2, 3 and 4, strong social justice is achieved.*

PROOF. Under a single effort level for a given type, functions v^t are now simply

$$v^t\left(\phi^t\right) = u^t\left(\phi^t, e^t\left(\phi^t\right)\right)$$

where $e^t(\phi^t)$ denotes the effort level applied by all individuals of type t as a response to receiving ϕ^t resources. By assumption 4, v^t as a function of its single argument is monotonic and strictly increasing with continuity. It can be shown that, if two types have different advantages, it is possible to reassign resources to reduce the difference. To reassign: choose i, j such that $v^i(\phi^i)$ is the minimum among all types and $v^j(\phi^j)$ is the maximum among types. Since functions v^t are continuous and increasing, there exists $\delta > 0$ such that

$$v^i\left(\phi^i\right) < v^i\left(\phi^i + \delta\alpha^i\right) < v^j\left(\phi^j - \delta\alpha^j\right) < v^j\left(\phi^j\right).$$

Then we can adjust ϕ^i by $\phi^i + \delta\alpha^i$ and ϕ^j by $\phi^j - \delta\alpha^j$ to obtain a new policy that still satisfies the global constraint (1.2.1). It may be that several types share the same minimum given by $v^j\left(\phi^j\right)$. To achieve a policy that improves Roemer's criterion, reassignment may require repeating the application. □

What poses a threat to the achievement of EOp is the existence of initial differences, not the availability of budgetary resources per se. After all, if types did not have initial levels of health, it would always be possible to equalize them at the bottom by assigning zero advantage for all of them. Proposition 1 formalizes the idea that in the absence of initial levels of health, EOp is achieved regardless of the budget.

Assume now that the types have some initial health status independent of the policy selected. We consider the case where this is due to some arbitrary initial assignment of resources that is given without necessarily respecting any justice requirement. It can be interpreted as public funds distributed in a previous stage according to some other criterion (for instance, utilitarism), as privately-provided resources (for instance, provisions by the family) or a combination of both. Regardless of the origin, we assume that all types hold some initial amount of resources (with at least one amount strictly positive) and that this does not depend on posterior public resources allocated. Then, in our context, it is strictly a matter of the the size of the public budget whether EOp can be reached or not.

PROPOSITION 5. *Let each type have a unique nonzero effort level. Then, under assumptions 1 and 2 and a given nonnegative initial assigment* x_0^t, $t \in \mathrm{T}$ *(where at least one assignment is positive, and at least one is zero), there are values* ω_S *and* ω_L *for the amount of resources to be assigned such that the following hold.*

(i) For $\omega < \omega_S$*, equality of achievement among types is unattainable.*

(ii) For $\omega \geq \omega_L$*, strong social justice is achieved.*

(iii) If strong social justice has been achieved (all types have equal advantage) any marginal increment of resources is distributed such that

$$\forall i, j \in \mathrm{T},\ \frac{\partial u^i}{\partial x}\frac{d\phi^i}{d\omega} + \frac{\partial u^i}{\partial e}\frac{de^i}{d\phi^i}\frac{d\phi^i}{d\omega} = \frac{\partial u^j}{\partial x}\frac{d\phi^j}{d\omega} + \frac{\partial u^j}{\partial e}\frac{de^j}{d\phi^j}\frac{d\phi^j}{d\omega}$$

.

PROOF. If (i) were false, letting ω go to zero will give us equality of advantages among $u^t\left(x_0^t + 0, e^t\right)$ with some x_0^t positive and some zero; this is impossible given assumption 2. To prove (ii), apply proposition 4 with advantage functions $u_*^t(x,e) = u\left(x_0^t + x, e\right) - u\left(x_0^t, e\right)$. We have a budget ω_* with equality of achievements that we may assume to be the minimal one. Then, take $\omega_L = \omega_* + \sum_t x_0^t$. (iii) follows by differentiating with respect to

ω (using the chain rule) the equality $u^i\left(\phi^i, e^i\left(\phi^i\right)\right) = u^j\left(\phi^j, e^j\left(\phi^j\right)\right)$ (see proof of prop. 4) where now $\phi^t = \phi^t\left(\omega\right)$. □

We extract a number of valuable insights from proposition 2. First, the amount public funds available matters for the equality of achievements. Second, once a state of strong social justice is reached, further resources are allocated so that every type obtains a share. The size of the portion depends on the ability of types of transforming resources into achievement. This last point is insightful and we build our modification around it. Intuitively, if all types receive some share of an hypothetical marginal increment of resources, it should mean that we are *as close as possible* to social justice. Otherwise, the additional funds would be channeled to those types experiencing unfair inequalities. Having characterized the issue of the scarcity of resources to the extent that is needed for our exercise, we now turn to the problem posed by limitations on the complexity of the policy.

1.3.3. Simplicity of the policy. In the presentation of his method, Roemer addresses the possibility of conflicts in the policies necessary to obtain strong social justice for every given quantile. In our context, this is the second main reason why strong social justice might not be achieved. In this part we exclude initial differences in resources and focus on illustrating the problems posed by restricting to the particular case of constant policies. Our motivation is that in a context of limitations on the availability of information to design compensation schemes, we believe there is a gain in designing a method that works under the simplest possible policies. We find that very stringent requirements are needed to be able to attain the ethical objective if we restrict ourselves to constant policies when dealing with multiple efforts. Through Example 6, we show how, under very particular homotheticity properties between both types and quantiles of effort, strong social justice is attained. We provide the example as an illustration of the degree to which special circumstances are necessary to completely fulfill the desired ethical criterion.

EXAMPLE 6. Assume:

(i) There are just two types $t = A, B$ each with a continuum of effort levels, and achievement functions of the form $u^t(\varphi, e) = \lambda^t \varphi^{\alpha^t} e^{1-\alpha^t}$.

(ii) There is a policy φ that assigns a constant amount of resources $\varphi^A + \varphi^B = 1$ (total amount of resources per capita is normalized to 1) and achieves strong social justice (equality of indirect achievements), that is $\forall \pi \in (0,1)$, $v^A(\pi, \varphi^A) = v^B(\pi, \varphi^B)$.

Then, the frontier of advantages is homothetic on the quantiles. If e^t_π denotes the π quantile in the effort distribution of type t,

$$\forall \phi \in (0,1), \quad \forall \pi, \pi' \in (0,1), \quad \frac{u^A(\phi, e^A_\pi)}{u^B(1-\phi, e^B_\pi)} = \frac{u^A(\phi, e^A_{\pi'})}{u^B(1-\phi, e^B_{\pi'})} \tag{1.3.1}$$

Furthermore, this assumption forces the effort quantiles by $e^A_\pi = c_1 \left(e^B_\pi\right)^{c_2}$ for some constants c_1, c_2.

PROOF. Simple substitution of u^t into the social justice condition 1.3.1 gives

$$\left(e^B_\pi\right)^{1-\alpha^B} = \frac{\lambda^A \left(\varphi^A\right)^{\alpha^A}}{\lambda^B \left(\varphi^B\right)^{\alpha^B}} \left(e^A_\pi\right)^{1-\alpha^A}$$

and this gives both results. □

We do not wish to restrict our contribution to such special cases nor relax our restrictions on policies, so our strategy is to develop a weaker definition of social justice as an objective. In particular, we relax the requirements on the definition of social justice by only accepting differences that arise from the simplicity of the policy.

1.4. The method

In application, the feasible level of complexity of policies may become an issue. We tackle this problem by proposing a weakened version of strong social justice which is enough for our objective. Namely, we consider a new social objective in which only inequalities derived from the simplicity of the policy are tolerated. To that effect, recall that proposition 2 establishes that all types receive a positive amount any marginal increase in available public funds once social justice is met. We build our new definition of social justice around this notion by defining a state of *weak social justice* in which all types receive a share of any marginal increase in the public budget. In the particular situation where proposition 2 is established, this coincides with the achievement of strong social justice. This is no longer the case when given multiple levels of effort. However, it is still possible to find an amount of resources (total budget) such that by using Roemer's allocation method, all types obtain a portion of a marginal increase of the public funds. By choosing Roemer's implementation, we relax our definition of social justice in the same way that Roemer did in his contribution. However, as emphasized earlier, we depart from Roemer in allowing only for differences in advantage arising from the simplicity our policy. We maintain that given exogenous initial differences across types, our relaxed version of social justice might not be achieved. The formalization of this explanation is given in the following definition.

DEFINITION 7. (Sufficient budget and weak social justice): Assume the policy is decided by solving Roemer's program. A sufficient budget is such that, for any budget in excess of the sufficient budget, all types would receive a strictly larger assignment under the decided policy associated with the larger budget. Formally, a total budget ω is *sufficient,* if, for all types t and for all $\bar{\omega} > \omega$, $\phi^t_{\bar{\omega}} > \phi^t_{\omega}$, where ϕ_{ω} denotes the assignment under the optimal policy in assumption 3. Denote as *weak social justice* the state in which the total budget is at least *sufficient*.

The definition captures the intuition that if every type receives a share of the additional resource pie, it is not possible to move closer to the social justice state. Therefore, a total budget is sufficient if it compensates for initial inequalities among types. We now prove that such a budget exists under fairly general conditions.

PROPOSITION 8. *Under assumptions 1 , 2 , 3 and 4 and a compact support effort distribution for each type (or finite number of values if effort levels are discrete), there exists a (finite) sufficient budget.*

PROOF. With these assumptions, functions $v^t(\pi, \phi^t)$ are continuous and non decreasing in both arguments, not bounded on the second argument but bounded for $\pi \in [0,1]$. Function $v(\pi, \Phi) = \min_t v^t(\pi, \phi^t)$ also has these properties. Function $R(\omega) = \max_\Phi \int_0^1 v(\pi, \Phi)\, d\pi$ is increasing and not bounded as ω grows.

If $\Phi = \Phi_\omega$ is the optimal policy for a budget ω, a type t has $\phi^t = 0$ if and only if $v^t(\pi, \phi^t_\omega) > v(\pi, \Phi_\omega)$ for all values of π except possible for some isolated ones. This may occur either because type t has some initial assignment x^t_0 or because the distribution of effort does not contain the zero effort. The only case where ϕ^t_ω does not increase when ω grows is when $\phi^t_\omega = 0$.

Now we proceed by induction on the number of types T. For $T = 1$ the proposition is clearly true. For T, choose a budget ω_0 large enough to be a sufficient budget for types $t = 1, \ldots, T-1$ while also satisfying $v(\pi, \Phi_{\omega_0}) > v^T(\pi, 0)$ for all π in some interval of positive length. This is also a sufficient budget for all types. □

Having proved existence, we select a particular *sufficient* budget. Recall, we seek to approach *weak social justice* using our selected policy. However, this new and more lax requirement of social justice could be achieved by a number of budgets. Our natural choice of budget, which in turn characterizes the advantages in a particular *weak social justice* state, is that which requires the least amount of resources.

DEFINITION 9. (Minimum sufficient budget): Given initially allocated resources $x^t \geq 0$, whereby some are positive, the infimum of all socially acceptable budgets is defined as the *minimum sufficient* budget and the corresponding optimal policy is defined as the *minimum sufficient policy* (note that this is also a sufficient budget).

How do we choose a policy when the budget is short of a sufficient budget? Roemer's method remains valid; however, it presumes priority for compensating types that are the worst off, a judgement that is not implied by EOp. We propose a method that allows society to decide which types to compensate by taking a stand on the tradeoff between the effectiveness of the public intervention and helping the worst off. We denote the choice in allocation resulting from society's stand on the effeciency-equity tradeoff in a socially-selected policy. Formally, we suggest that, given an *insufficient budget*, we can obtain our policy from the minimization of the distance to the state of *weak social justice* generated by the *minimum sufficient policy*. What follows is a formal exposition of our method .

Let ω_0 be the minimum sufficient budget for a given set of initially allocated resources $\{x^t\}$, let ϕ_0 be the corresponding minimum sufficient policy , and let $v^t(\pi, \phi_0)$ the associated indirect achievement functions. We denote by $\mathbf{v}(\pi, \phi_0)$ the vector of all types' achievements for degree π. For a smaller budget ω we say that the *socially-selected* policy is the solution for the program

$$\min_{\phi} \int_0^1 \|\mathbf{v}(\pi, \phi_0) - \mathbf{v}(\pi, \phi)\|_p \, d\pi \tag{1.4.1}$$

where $\|.\|_p$ denotes a norm (see below) that measures the distance to the minimum sufficient policy and $\mathbf{v}(\pi, \phi)$ is the vector of achievements under policy ϕ.

It is clear that the choice of distance has profound implications on the resulting policy. Consider a p-norm; then, 1.4.1 becomes

$$\min_{\phi} \int_0^1 \left(\sum_{t=1}^{t} \left| v^t(\pi, \phi_0) - v^t\left(\pi, \phi^t\right) \right|^p \right)^{\frac{1}{p}} d\pi \tag{1.4.2}$$

It is desirable to formulate the problem in this fashion as every metric comprised in the p-norm corresponds to a choice of weights in the trade off between effectiveness of public resources and compensating those who are worse off. For instance, for $p = 1$, 1.4.2 becomes the minimization of the sum of differences.

$$\min_{\phi} \int_{0}^{1} \left(\sum_{t=1}^{t} \left| v^{t}\left(\pi, \phi_{0}\right) - v^{t}\left(\pi, \phi^{t}\right) \right| \right) d\pi$$

The choice of this metric entails a utilitarian policy or simply, targeting *the type that could benefit the most*. At the margin, resources are allocated to the type with a larger partial derivative of the advantage function with respect to resources. Since in general, there is no guarantee that an interior solution is achieved, this could lead to abandoning types with lower abilities to transform resources into advantage.

On the other extreme, when $p \to \infty$, 1.4.2 is then:

$$\min_{\phi} \int_{0}^{1} \left(\max_{t} \left| v^{t}\left(\pi, \phi_{0}\right) - v^{t}\left(\pi, \phi^{t}\right) \right| \right) d\pi$$

The solution allocates resources to types farther away from the minimum sufficient policy advantages, disregarding any consideration regarding the effectiveness of the public funds .We acknowledge, however, a non-desired consequence that follows from our definition of weak social justice. We cannot prove that in general, when limited to simple constant policies, our sufficient policy leads to the equality of achievements for every quantile across types[4].

1.4.1. Discussion. We have assumed in this chapter the standard case of a continuous achievement function; this might be considered a limitation of our proposal as many applications require considering categorical data such as health states. This issue has been studied in recent work that has explored extensions of the EOp framework categorical data (Herrero and Villar 2012). Our methodology extends to these situations as well;

[4] Through our example 1 we have shown that in general, when limited to simple constant policies, there is no guarantee that our sufficient policy leads to the equality of achievements for every quantile across types.There can be quantiles for given types that result in higher levels of achievement than others, leading to the counterintuitive result that a Ralwsian policy would prioritize their compensations. This is entirely a result based on the limitations we have imposed to our policy and its actual relevance depends on how different advantages are in weak social justice.

when dealing with categorical data, our method requires minimizing the distance to the proportion, for each category and type, that is in accordance to the weakened definition of EOp.

On a more general note, a critique leveled against EOp is that it does not include the traditional tradeoff between efficiency and equity. Our work includes considerations regarding efficiency in the application of EOp, but only related to the attainment of social justice. Hence we acknowledge the possibility that pursuing the optimal policies according to our method could be inefficient from a broader standpoint. Nevertheless, we conceive the framework of EOp as providing the guidelines for the design of the policies of a public authority in charge exclusively of the attainment of equity. That is not to say that society, at a broader level, might still be allocating its resources considering the traditional balance between efficiency and equity. In the case of health care, for instance, the relevant ministry might devote part of its budget to reduce avoidable mortality and some other fraction to equitable access. Our proposal is addressed exclusively at the latter objective.

1.5. Conclusion

We have argued that full capacity of decision over the normative content of the theory of equality of opportunity formulated by Roemer (1998) requires transfering the decision power over its allocation method. In order to apply the theory it is not only relevant for society to establish the legitimacy of inequalities; when the scarcity of resources results in the impossibility of attaining equality of opportunity, it is also crucial to decide which types to compensate first. In this work, we outlined the conditions under which this concern matters, namely when public funds are insufficient to compensate for initial differences across types. In this context, Roemer's original proposal advocates for a Rawlsian approach prioritizing redistribution towards groups that are the worst off. Our critique is that this approach neglects consideration of targeted individuals' capacities to benefit from transfers, something that is crucial in applications to health care.

We then present a new allocation method based on the intuitive idea of the minimization to an unattainable objective. Through the choice of different metrics, our methodology allows inclusion, with varying importance, of the weight attached to the potentially conflicting objectives of compensating the types that are worst of and types that can benefit the most. The main contribution of this chapter is, therefore, to extend the framework of EOp to allow societal control of the application of equality of opportunities.

CHAPTER 2

Reporting heterogeneity or better health? Revising the income-related reporting heterogeneity hypothesis

2.1. Introduction

Researchers and policymakers often rely on self-reported health measures to assess health inequality. Part of the popularity of self reported health indicators is due to the fact that they are relatively inexpensive to collect and, despite the debate over their accuracy (Baker, Stabile and Deri, 2004), have been shown to be a reasonably good predictor of health utilization (van Doorslaer, Koolman and Jones, 2004) and mortality (Idler and Benyamini, 1997 and 1999, van Doorslaer and Gerdtham, 2003, and Jürges, 2008). Even when reliable information on objective health measures is available, such as the conditions suffered by the individual, methods that map objective health variables onto overall health status incorporate are often based on revealed preferences (Dolan, 2000).

A concern surrounding self reported health measures is that they are not based exclusively on underlying "true" health but rather incorporate reporting heterogeneity across individuals. In this work, I investigate whether there are indeed systematic differences in self-reported health across socioeconomic groups. My main finding is that at least part of the observed differences in reporting across socieconomic groups can be attributable to unaccounted differences in "true" health. Understanding this type of difference is particularly relevant because of the policy and research focus on the socioeconomic gradient of health. Arguably, health inequality evaluations should only consider differences in the "true" health status of individuals and hence we must account for reporting heterogeneity by socieconomic group.

This is not a new question; there exists an extensive literature that has investigated the validity of self-reported measures of health and, in particular, its effect on health inequality evaluations. There exist a plethora of works that evaluate the accuracy and utility of self-reported measures of health. While it is generally accepted that self-reported health measures are correlated with objective indicators of health, the debate over their accuracy has resulted in some recent works turning to vignettes as an alternative methodology to finding objective measures of health. This method has been deemed particularly appropriate in the context of developing countries (Bago d'Uva, van Doorslaer, Lindeboom and O'Donell, 2008), but has also been used in studies in the United States (Dowd and Todd, 2011) and in European countries (Bago d'Uva, O'Donnell and van Doorslaer, 2008).

Finally, the strand of the literature most closely related to this work has focused on studying individual reported health[1] and finds reporting heterogeneity using surveys that contain both self-reported health variables and "objective" measures of health, such as self-reported incidence of health conditions. It is important to note that these "objective" measures of health have been found to be somehow inaccurate (Baker, Stabile and Deri, 2004) and subject to income related reporting heterogeneity (Johnston, Propper and Shields, 2009). However, self-reported incidence is considered to be relatively more reliable in developed countries, given the improved likelihood of diagnosis. Interestingly, the results of this literature are mixed with regards to the direction of the self-reporting bias. Sen (2002) compares life expectancy and self reported morbidity across two Indian states and the United States. Surprisingly, reported morbidity is positively associated with life expectancy; this is something that he labels as a "health paradox". Dowd and Zajacova (2010) use education instead to represent socioeconomic status and also find evidence of heterogeneity in the association between objective health (as measured by biomarkers) and self reported health measures. Their findings suggest that reporting heterogeneity leads to an understatement of the health socioeconomic gradient. On the other hand, van Doorslaer and Humphries (2000) find that higher income groups report relatively smaller declines in health as their health status worsens, which implies that researchers could be overstating socioeconomic health inequalities. Etilé and Milcent (2006) also find evidence of understatement of health declines by higher income groups.

This work is closest to these last papers in that it combines measures of subjective and objective health to assess reporting heterogeneity. I study the reporting heterogeneity

[1] Reporting heterogeneity has also been studied in other dimensions. For example, (Jürges, 2008) finds evidence of cross-country heterogeneity.

hypothesis using data from the Catalan Survey of Health (ESCA, in Catalan) survey that covers individual reports on health, socioeconomic variables, information on conditions suffered, other measures of health (such as the body mass index), and on individual's life styles. My hypothesis is that part of the finding that high income groups overstate their objective health is at least partially explained by the failure to properly account for differences in health states by people suffering from similar conditions. Indeed, higher income groups may in fact have better health despite suffering from the same conditions. Better care, living conditions and life styles might result in health gains from palliative care which could explain the apparent over reporting observed by income.

To assess this possibility, I test if reporting heterogeneity is robust to the introduction of variables that reflect better self care such as smoking and exercising (also self reported). In some of my specifications the terms capturing reporting heterogeneity by income group disappear, lending support to my hypothesis that differences in reporting across income groups might in fact reflect true differences in health. I then decompose the observed socioeconomic inequality in health into five components: i) the different prevalence of conditions across socioeconomic groups and other objective health measures, ii) differences across occupational statuses, iii) age a gender iv) the effect of education groups, and v) a final part based on differences in life styles. The results of my exercise indicates that, while differences in objective health measures are the largest source of inequality, education variables are also responsible for a substantial amount of health inequalities.

2.2. Data

I use data from the ESCA from the year 2006[2]. The ESCA covers variables on socioeconomic status such as education, occupation, income, self reported health measures, a wide array of conditions, life style variables and health service utilization. Table A1 (in

[2] The year 2006 is particularly appropriate for our exercise because unemployment is relatively low compared to more recent waves that have unemployment levels well above the 20% mark. Given the evidence on the relationship between unemployment and health (Roelfs, Shor, Davidson and Schwartz, 2011), this could distort my results.

the appendix) is a full list of the variables used in this study. It is noteworthy to clarify that, in the survey, questions regarding conditions are asked in a very general manner that does not allow the researcher to distinguish the severity with which the condition is suffered by the respondent.

For instance, the below question is asked in the survey:

- Could you tell us if you suffer or have suffered from any of the chronic conditions that I will proceed to read? (In the original Catalan: Ens podria dir si pateix o ha patit algun dels trastorns crònics que ara li llegiré?)

The variable for income is categorical, which is problematic for the computation of the concentration index, which is the measure I use for the decomposition of health inequality (Chen and Roy, 2009). To address this, I create a continuous variable as follows. First, I estimate a model that explains income group affiliation as a function of the rest of socioeconomic variables. Then, I assign continuous income as the expected income resulting from multiplying the probability of belonging to each of two adjacent groups by the midle point of the income bracket they represent.

2.3. Methods

2.3.1. Testing the hypothesis. In order to test my hypothesis that apparent reporting heterogeneity by income might be due to higher income groups having better objective health, I estimate two different models. The first model explains self-reported health as a function of income, other socioeconomic variables, incidence of conditions and the interaction between income and objective health variables. Reporting heterogeneity is captured by both the coefficients on income and on the interaction between income and objective health variables. While both terms capture reporting heterogeneinty by income group, their interpretation is different. A significant and positive main effect of income implies that, for the same reported objective health, higher income groups report better health. The interaction signals that income groups react in terms of reporting to variation in objective health.

In the second model I introduce variables of self care. In particular, I consider smoking status and the frequency of physical activity. If the terms signaling self reporting heterogeneity were simply capturing differences in objective health within people suffering the same conditions, I expect them to no longer appear significant. Given the number of conditions in the survey, estimating the interactions between all the conditions and income would result in a very large number of parameters. Instead, I construct health indices with a principle components analysis (PCA)[3] on the twenty four conditions of the survey. PCA helps reduce the high dimensionality of the problem by exploiting the correlation between the different conditions, which can be considered "redundant" information. Instead of including the set of X conditions, which are highly correlated, I use the first five PC that are by construction orthogonal to each other.

Table 1 reports the correlations between conditions and the first five principle components (PC) and the variance captured by each component. The first five components explain roughly 40% of the variance. In my analysis, I report the results for models that consider from one to five components. The first PC is the easiest to interpret since it is positively correlated with all the conditions; it creates a divide between respondents who are healthy and unhealthy. The remaining components are more difficult to interpret as they contain both positive and negative associations with the different conditions. For example, the second PC is positively correlated with prostrate conditions and negatively correlated with arthrosis and migraines, a combination of conditions most likely to be found in men. The third PC is heavily correlated with respiratory conditions, such as asthma and chronic pulmonary conditions. In the case of the forth component the dominant (and positively correlated) conditions are "ulcera", prostate, cancer and skin conditions. The fifth and last PC is strongly correlated with reporting having suffered a stroke.

[3] While PCA analysis is not recommended on categorical data, the case of binary dummy variables is special Gower, 1966)

TABLE 1. Principle components

	PC1	*PC2*	*PC3*	*PC4*	*PC5*
Hyptertension	0.48	0.21	-0.31	-0.15	-0.20
Diabitis	0.36	0.30	-0.26	-0.23	-0.02
Heart	0.46	0.27	-0.16	-0.12	0.25
Asthrosis	0.57	-0.19	0.01	0.01	-0.30
Asthma	0.25	0.31	0.69	-0.24	-0.07
Chronic pulmonar	0.37	0.37	0.61	-0.12	-0.10
Ulcer	0.31	0.02	0.07	0.45	-0.16
Prostate	0.25	0.47	-0.07	0.46	-0.12
Cholesterol	0.40	0.06	-0.34	-0.24	-0.31
Catacracts	0.49	0.23	-0.15	0.16	0.10
Chronic skin	0.25	-0.06	0.11	0.32	-0.05
Constipation	0.46	-0.25	0.02	0.10	0.05
Depression	0.53	-0.31	0.05	0.00	-0.05
Stroke	0.32	0.23	-0.07	-0.13	0.53
Migraine	0.37	-0.46	0.16	-0.04	-0.21
Circulatory	0.59	-0.17	-0.03	-0.13	0.04
Cancer	0.27	0.02	0.00	0.43	0.26
Osteoporosis	0.43	-0.12	-0.06	-0.02	-0.03
Anemia	0.34	-0.28	0.21	0.07	0.35
Thyroids	0.27	-0.27	-0.02	-0.18	0.35
	variance	*variance*	*variance*	*variance*	*variance*
	16.2	6.8	6.4	5.1	5

2.3.2. Decomposing inequality. Next we turn to breaking down existing health inequality according to contributing factors of inequality. My analysis follows the methodology first presented by (Wagstaff and van Doorslaer, 2003). Intuitively, this consists of an Oaxaca-Binder (Blinder, 1973 and Oaxaca, 1973) style of decomposition for a measure of inequality instead of using the mean the variable of interest . In particular, the focus of this method is the concentration index (Wagstaff, Paci and van Doorslaer, 1991). The concentration index is a measure of inequality widely used in health economics that measures the degree of income based inequality. Its construction is very similar to that of the Gini coefficient, but while the concentration index similarly ranks population according to income, it plots it against the proportion of health. Formally, the concentration index is defined from a concentration curve $L(s)$ that plots the cummulative proportion of the population ranked by income, starting with the lowest group, against the commulative proportion of health. A concentration curve with slope 1 indicates perfect equality between income groups, a value of -1 signals that health is concentrated purely in the lowest income group and 1 that it is accumulated by the highest income group.

Following Wagstaff and van Doorslaer (2003), one can rewrite the concentration index as a weighted sum of the concentration indexes of the different explanatory variables of self reported health. The result is as follows. Let the following be a linear model of health as a function of its k determinants. In my case, this is objective health, socioeconomic and life style variables:

$$h_i = \alpha + \sum_k \beta_k x_{k_i} + \varepsilon_i$$

Then, it is possible to rewrite the concentration index C as:

$$C = \sum_k \left(\frac{\beta_k \bar{x}_k}{\mu} \right) C_k + \frac{GC_\varepsilon}{\mu}$$

where μ is the mean of the concentration index, $\bar{x}$ the mean of x_k , and C_k the concentration index of variable x_k. The term GC_ε is the generalized concentration index of the residuals. According to this decomposition, it is possible to assess the contribution of each variable to the overall income related inequality in health according to three terms: i) its mean, ii) its effect on health, measured by the health status mode, and iii) how it is distributed across income groups, which is captured by its concentration index. I perform this decomposition on the model resulting from my test of the health paradox and report the relative importance of the different components of health inequality in Catalunya in 2006.

2.4. Results

2.4.1. Testing the hypothesis. Table 2 displays the results of testing the socioeconomic reporting heterogeneity. The results are presented as follows. First, I report the estimates of the coefficients of the model interacting the health indexes with income. Then, I estimate the same model with additional variables on objective health and self care. While not completely ruling out reporting heterogeneity in all the models, my results lend support to the hypothesis that part of the differences across income groups in reporting is due to failure of properly accounting for true health differences. In all the models the main effect of income loses significance after introducing the additional variables. However, the interaction between the health indexes and income shows mixed results across specifications. The first PC loses significance in models one through five, but the fourth and fifth PCs are significant.

TABLE 2. Testing the hypothesis

	Model 1		*Model 2*		*Model 3*		*Model 4*		*Model 5*	
	A	*B*	*A*	*B*	*A*	*B*	*A*	*B*	*A*	*B*
PC1*income	0.001 *	0.001 *	0.001 *	0.001	0.001 *	0.001	0.001	0.001	0.001 *	0.001
PC2*income	-	-	0.000	0.000	0.000	-0.001	0.000	-0.001	0.000	-0.001
PC3*income	-	-	-	-	-0.001	0.000	-0.001	0.000	-0.001	0.000
PC4*income	-	-	-	-	-	-	-0.001 *	-0.002 *	-0.001 *	-0.002 *
PC5*income	-	-	-	-	-	-	-	-	-0.001	-0.001 *
BMI	-0.991 *	-1.045 *	-1.038 *	-1.091 *	-1.101 *	-1.134 *	-1.066 *	-1.107 *	-1.024 *	-1.058 *
Smoking		1.392 *		1.371 *	-	1.349 *	-	1.353 *	-	1.307 *
Exercise		-3.566 *		-3.576 *	-	-3.566 *	-	-3.568 *	-	-3.614 *
Hospital	-1.964 *		-2.013 *	-1.662 *	-1.985 *	-1.659 *	-1.984 *	-1.647 *	-2.045 *	-1.708 *
PC1	-8.930 *	-8.773 *	-8.665 *	-8.408 *	-8.503 *	-8.276 *	-8.398 *	-8.079 *	-8.448 *	-8.041 *
PC2	-	-	0.848 *	1.105 *	0.987 *	1.205 *	1.139 *	1.444 *	1.186 *	1.597 *
PC3	-	-	-	-	-0.005	-0.039	0.016	-0.025	0.035	-0.028
PC4	-	-	-	-	-	-	1.539 *	1.808 *	1.520 *	1.780 *
PC5	-	-	-	-	-	-	-	-	1.083 *	1.551 *
Income	0.002 *	0.001	0.002 *	0.001	0.002 *	0.001	0.002 *	0.001	0.002 *	0.001
M35-44	-2.372 *	-2.272 *	-2.347 *	-2.251 *	-2.406 *	-2.301 *	-2.443 *	-2.316 *	-2.382 *	-2.247 *
M45-54	-3.921 *	-4.182 *	-3.987 *	-4.253 *	-4.146 *	-4.364 *	-4.224 *	-4.405 *	-4.097 *	-4.256 *
M55-64	-4.067 *	-4.059 *	-4.382 *	-4.317 *	-4.702 *	-4.533 *	-4.808 *	-4.584 *	-4.704 *	-4.472 *
M65-74	-2.484 *	-1.957	-3.141 *	-2.589 *	-3.461 *	-2.835 *	-3.555 *	-2.830 *	-3.500 *	-2.852 *
M75+	-5.849 *	-5.095 *	-6.783 *	-5.972 *	-7.066 *	-6.219 *	-7.243 *	-6.273 *	-7.320 *	-6.382 *
F15-34	-0.276	-0.513	-0.072	-0.324	-0.045	-0.306	0.003	-0.265	-0.037	-0.312
F35-44	-2.674 *	-3.042 *	-2.337 *	-2.722 *	-2.328 *	-2.713 *	-2.287 *	-2.689 *	-2.307 *	-2.713 *
F45-54	-4.716 *	-4.632 *	-4.348 *	-4.303 *	-4.422 *	-4.354 *	-4.370 *	-4.316 *	-4.337 *	-4.253 *
F55-64	-5.179 *	-5.931 *	-4.925 *	-5.762 *	-5.142 *	-5.925 *	-5.128 *	-5.944 *	-4.964 *	-5.739 *
F65-74	-6.591 *	-6.433 *	-6.656 *	-6.647 *	-6.974 *	-6.883 *	-6.907 *	-6.872 *	-6.835 *	-6.827 *
F75+	-9.995 *	-8.818 *	-10.210 *	-9.188 *	-10.650 *	-9.529 *	-10.670 *	-9.624 *	-10.810 *	-9.800 *
Unemployed	-2.078 *	-2.894 *	-2.046 *	-2.839 *	-2.057 *	-2.850 *	-1.971 *	-2.821 *	-1.972 *	-2.854 *
Disabled	-13.320 *	-13.060 *	-13.270 *	-12.960 *	-13.280 *	-12.970 *	-13.280 *	-12.960 *	-13.340 *	-12.880 *
Retired	-2.382 *	-2.887 *	-2.322 *	-2.771 *	-2.332 *	-2.785 *	-2.370 *	-2.797 *	-2.312 *	-2.668 *
Home maker	-0.233	-0.801	-0.198	-0.723	-0.225	-0.734	-0.071	-0.590	-0.053	-0.569
Student	2.162 *	1.160	2.166 *	1.153	2.105 *	1.130	2.240 *	1.231	2.062 *	0.986
Other	-3.970	-1.727	-3.975	-1.676	-4.076	-1.638	-3.777	-1.407	-3.868	-1.511
Private ins.	0.825 *	1.027 *	0.847 *	1.045 *	0.832 *	1.030 *	0.808 *	1.010 *	0.814 *	1.023 *
Primary ed.	3.656 *	3.595 *	3.727 *	3.672 *	3.725 *	3.677 *	3.718 *	3.667 *	3.682 *	3.612 *
Secondary ed.	5.177 *	5.213 *	5.226 *	5.275 *	5.215 *	5.280 *	5.194 *	5.267 *	5.126 *	5.173 *
Tertiary ed.	6.136 *	5.810 *	6.143 *	5.820 *	6.138 *	5.829 *	6.105 *	5.815 *	6.047 *	5.738 *
Others ed.	1.147	0.807	1.213	0.903	1.166	0.881	1.130	0.853	0.861	0.504

* $p < 0.05$

A shortcoming of current measures of objective health, such as the incidence of certain conditions, is that they also contain information on the severity of the conditions and its impact in individual well-being. Hence, evaluations based on those variables cannot allow for reported severity levels to be a reflection of true health as opposed to an illusion stemmed from reporting heterogeneity.

By adding some additional information on the individuals' self care I have been able to show that at least part of the effects previously identified as reporting heterogeneity are in fact differences in "true" health states. A possible explanation for this result is that the diminished severity could be due to the palliative effect of self care. In the next section I discard the interaction terms, estimate a model of health with the full array of conditions available in the ESCA and discuss more generally the effect of the different explanatory variables.

2.4.2. The decomposition of the contributions to health inequality. Table 3 shows the results of estimating self reported health on a model without interactions between socioeconomic status and health. This model, supported by the results in the previous section, has the advantage in that it allows for separately estimating the effect of the different conditions on health. The signs of the effects are mostly significant and in the expected direction. Most conditions negatively affect reported health, with the exception of strokes and intestinal problems that are not significant. In the case of strokes, this is most likely an issue of statistical power due to its very low prevalence (around 2% of the population). Age negatively affects self-reported health for both males and females, with older ages progressively reporting lower health statuses. The coefficient on education is positive and significant for all education groups except the residual category of "others". Interestingly, there does not appear to be significant differences between the effect of secondary and tertiary education. Both variables of the body mass index and hospital stays result in lower self reported health as expected and non smokers and individuals who report more physical activity report to be healthier.

TABLE 3. Self reported health model with all the conditions

	Coefficient	*Std.error*		*Coefficient*	*Std. error*
Intercept	76.72	(0.74) *	M35-44	-2.19	(0.56) *
Smoking	-1.16	(0.31) *	M45-54	-4.13	(0.63) *
Exercise	3.54	(0.42) *	M55-64	-4.03	(0.76) *
Hospital	-1.59	(0.26) *	M65-74	-2.07	(1.08)
BMI	-1.00	(0.29) *	M75+	-5.49	(1.23) *
Hypertension	-1.88	(0.39) *	F15-34	-0.27	(0.47)
Diabetes	-3.45	(0.62) *	F35-44	-2.64	(0.59) *
Heart	-2.70	(0.58) *	F45-54	-4.28	(0.63) *
Arthrosis	-5.27	(0.32) *	F55-64	-5.45	(0.73) *
Asthma	-2.26	(0.63) *	F65-74	-6.43	(0.89) *
Chronic pulmonar	-2.41	(0.64) *	F75+	-9.54	(0.93) *
Ulcer	-3.08	(0.60) *	Unemployed	-2.92	(0.69) *
Prostrate	-3.96	(0.94) *	Disabled	-12.53	(0.92) *
Cholesterol	-1.37	(0.42) *	Retired	-2.69	(0.71) *
Cataracts	-1.75	(0.59) *	Home maker	-1.20	(0.55) *
Chronic skin	-1.96	(0.55) *	Student	0.81	(0.58)
Constipation	-0.60	(0.50)	Other	-2.03	(4.59)
Depression	-7.29	(0.40) *	Primary ed.	3.56	(0.48) *
Stroke	-1.70	(1.18)	Secondary ed.	5.04	(0.55) *
Migraine	-2.50	(0.37) *	Tertiary ed.	5.61	(0.61) *
Circulatory	-3.44	(0.40) *	Others ed.	-0.91	(7.60)
Cancer	-3.75	(0.84) *	Private ins.	1.03	(0.30) *
Osteoporosis	-4.79	(0.63) *			
Anemia	-2.90	(0.56) *			
Thyroids	-1.83	(0.67) *			
Adjusted R^2	0.38				

* $p < 0.05$

Table 4 shows the different components of the decomposition of the concentration index and Table 5 has a summary of the overall contributions of the different categories to income related health inequality. The largest contributor to health inequality is the income gradient in the prevalence of the conditions considered. In particular, depression, artrosis and circulatory problems stand out as the largest contributors within the group. The three conditions are relatively highly prevalent and largely concentrated among lower income groups. Furthermore, depression is among the largest contributors to overall inequality, second only to university degrees. Differences in educational achievements are the second group in contributions to inequality. The uneven distribution of university degrees across income groups is a strong positive contribution to inequality.

TABLE 4. Contribution to the concentration index

	CI	*Contrib.*	*Cont./total**		*Cont./total*	*Contrib.*	*Cont./total*
Smoking	0.0531	-0.0003	-2.6	M35-44	0.2552	-0.0007	-6.2
Exercise	0.0050	0.0002	1.9	M45-54	0.2886	-0.0013	-11.1
Hospital	-0.0768	0.0002	1.9	M55-64	0.4420	-0.0016	-13.1
BMI	0.0289	-0.0002	-1.8	M65-74	0.4646	-0.0007	-5.7
Hypertension	0.0332	-0.0002	-1.6	M75+	0.4381	-0.0016	-13.3
Diabetes	-0.0139	0.0000	0.4	F15-34	-0.2412	0.0001	1.2
Heart	0.0147	0.0000	-0.4	F35-44	-0.1447	0.0005	3.9
Arthrosis	-0.0251	0.0009	7.3	F45-54	-0.1893	0.0009	7.3
Asthma	-0.0778	0.0002	1.3	F55-64	-0.2069	0.0010	8.6
Chronic pulm.	-0.0423	0.0001	0.8	F65-74	-0.1764	0.0009	7.6
Ulcer	0.0255	-0.0001	-0.6	F75+	-0.1368	0.0013	10.9
Prostrate	0.4031	-0.0011	-8.9	Unemployed	-0.6329	0.0011	9.3
Cholesterol	0.0461	-0.0001	-1.1	Disabled	-0.1885	0.0011	9.4
Cataracts	0.0028	0.0000	-0.1	Retired	0.2942	-0.0019	-15.9
Chronic skin	-0.0040	0.0000	0.1	Home maker	-0.6591	0.0014	11.9
Constipation	-0.1012	0.0001	0.7	Student	-0.9129	-0.0006	-5.4
Depression	-0.1277	0.0025	20.9	Other	-0.2465	0.0000	0.1
Stroke	-0.0186	0.0000	0.1	Primary ed.	-0.1700	-0.0036	-30.4
Migraine	-0.1305	0.0009	7.3	Secondary ed.	0.0488	0.0010	8.1
Circulatory	-0.0966	0.0010	8.1	Tertiary ed.	0.4804	0.0058	48.3
Cancer	0.0229	0.0000	-0.3	Others ed.	-0.3259	0.0000	0.0
Osteoporosis	-0.1084	0.0004	3.6	Private ins.	0.2298	0.0011	9.3
Anemia	-0.1528	0.0005	3.8				
Thyroids	-0.1461	0.0002	1.5				

* Contribution over total is in percentual terms.

TABLE 5. Contributions to the concentration index by groups of variables

	Contrib./total
Conditions and obj. health variables	43.0
Gender and age	-10.0
Occupation	9.4
Education	25.9
Self care and private insurance	8.6

Interestingly, primary education contributes to the reduction of inequality. This is due to the fact that, while having a positive effect on self reported health, it is heavily concentrated in lower income groups. Disparities in reported health across occupation groups are also a source of inequality. The largest and most straightforward contribution is that of incapacitated people, mainly due to the size of its effect on health. Unemployment also has a significant contribution to overall inequality, in this case a heavy component is its large concentration index. Interestingly, among the life style variables, smoking is responsible for a reduction in health inequality. This is due to the higer prevalence of smoking among males, who are also among the higher income groups. Another important contributor to inequality are private insurances. Despite having a modest effect on self reported health, their concentration index is relatively high and about two thirds of the individuals report having a private insurance. Finally, the age-gender interaction contributes negatively to inequality, mostly due to the combination of positive concentration indexes and negative contributions to health of males.

2.5. Discussion

The main conclusion of this work is that failing to account for differences in "true" objective health may be partially responsible for the income-related heterogeneity in self-assessed health that has been found in previous studies. While my results do not completely rule out the presence of heterogeneity in reporting, they do suggest that at least part of the differences come from "true" health differences derived from better self care which is just one possible source of differences in "objective" health across income groups. Other works that have used a wide array of measures of objective health, such as Etilé and Milcent (2006) might have failed to account but relevant circumstances affecting the overall "true" health of individuals that otherwise share similar conditions. Here, by "true" health I mean measures of well-being that are not affected by reporting differences and could include considerations beyond the purely physical condition of the individual.

The second part of this work finds that differences in the prevalence objective conditions are the largest contributor to income-related health inequality. This is because of the large impact of prevalence objective conditions on self-assessed health they have and the fact that concentration index of a majority of conditions, with notable exceptions like prostrate conditions, is negate, indicating a disproportionate concentration among lower socioeconomic grups. This is particularly true for the case of the three largest negative contributors to well being: depressions, circulatory problems, and artrosis. Despite its lower contribution to income-related health inequality, educational differences are responsible for a relatively large fraction of it, ahead of occupational based differences. Although this is not a possibility that I have explored in the present work, it may very well be the case that education is, similarly to income, capturing the effect of actual differences in objective health. Given the findings of the literature that related this variable to self-reported heterogeneity (Dowd and Zajacova, 2010 and Bago d'Uva, O'Donnell and van Doorslaer, 2008), this is a promising direction for future research.

CHAPTER 3

Persistent Differences in Mortality Patterns Across Industrialized Countries (with Hippolyte d'Albis and Loesse Jacques Esso)

3.1. Introduction

Whether there is global convergence in well-being across countries remains in the realm of scientific debate. The hypothesis of convergence has been studied for a variety of dimensions of well-being, with mixed results. While economic convergence is not a general reality (Barro and Sala i Martin, 2003), recent research has emphasized the existence of global demographic convergence (Wilson, 2001) — in life-expectancy and fertility — and its importance for reductions in the inequality in living conditions (Becker, Philipson and Soares, 2005). We investigate the evolution of mortality patterns for a large group of industrialized countries through an analysis of their ages-at-death distributions. The ages-at-death distribution is given by the number of deaths at every given age in the period life table.

Our research advances the current understanding of mortality convergence by, first, formally testing the implications of the theory on the epidemiological transition, then second, uncovering trends in mortality patterns beyond the evolution of the mean, i.e. life expectancy, and variance of the considered distribution. Our chosen indicator of divergence, the Kullback-Leibler divergence (KLD), provides a comprehensive measure of the overall differences between distributions.

Studies of convergence classify differences between countries using two broad categories: first there is consideration for the unequal positions of countries in the stages of their development; second are structural differences, dissimilarities that persist even should countries become equally developed (Barro and Sala i Martin, 2003). Thus far, expectations for convergence in mortality have been based on a catching-up process between countries in different stages of development. In particular, the hypothesis of convergence in mortality patterns has been a natural corollary to the theory of epidemiological transition (Omran, 1971), whereby countries lagging in their transition paths experience relatively faster gains in life expectancy and acatch-up with countries in the later stages (Wilmoth, 2001 and Vallin and Meslé, 2004). This is a consequence of the first stage in development whereby reductions of death rates arise from reductions in infectious diseases, a phenomenon with greater relative impact among infants. However, once the the reduction of infant mortality and relatively easily preventable deaths at younger adult ages have been realized, further gains in life expectancy are due to gains at older ages, which are increasingly costly and occur at much slower rates.

In this work, we test and reject the convergence hypothesis for industrialized countries in the period 1960-2008. However, we acknowledge that the lack of convergence for our whole sample does not necessarily imply that there are not subgroups of countries converging. In fact, the concept of convergence among subgroups or clubs of countries has already received some attention in the mortality convergence literature (Bloom and Canning, 2007). The basic theoretical difference is that proponents of the hypothesis of convergence club postulate that there might exist not one but several long term ages-at-death distributions. To address this possibility, we divide our sample in two groups, Eastern and Western countries, based on former pertinence to a common political history. We find that the trajectories of the two groups of countries are remarkably different. These results point out that the overall divergence trend is partly driven by trends in the differences between the two subgroups.

Our findings on the lack of convergence are coherent with recent findings that highlight the relatively high variability of mortality at young adult ages across countries and its contribution to international differences in mortality patterns (Edwards and Tuljapurkar, 2005, and Fillespie and Trotter, 2014). While the reduction of mortality at earlier life stages described by the theory of the epidemiological transition has contributed to convergence, differences in young adult mortality have acted as a countervailing

3.2. Materials and methods

Our object of study, mortality patterns, is extracted from the period life-tables available in the Human Mortality Database. Following the work by Edwards and Tuljapurkar (Edwards and Tuljapurkar, 2005), we evaluate the dissimilarities between the ages-at-death distributions of all our countries with the Kullback-Leibler divergence (KLD), a measure of the overall dissimilarities between distributions. An advantage of using this index of dissimilarity is that we are no longer restricted to the study of the mortality distribution through its first two moments — life-expectancy and variance — which is particularly relevant given the non-normality of ages-at-death distributions. We describe any group of countries converging in mortality if the dissimilarities across their ages-at-death distributions diminish. Since the KLD is a measure of pairwise differences, we compute each sum of KLDs between individual countries and the period's average distribution and study its evolution. As we are only concerned with the trend (and not levels), this is formally equivalent to computing the pairwise sum of differences across countries.

The divergence of the age-at-death distribution for country P from that of country Q is given by the expression:

$$KLD(Q \parallel P) = \sum_{\alpha=0}^{110} \ln\left(\frac{P(\alpha)}{Q(\alpha)}\right) P(\alpha)$$

where $P(\alpha)$,$Q(\alpha)$ are the probability masses in each age group $\alpha = 0, ..., 110$ .The higher the value of $KLD(Q \parallel P)$, the larger the differences between the two distributions. In our exercise, the average age-at-death distribution is computed as the unweighted average of ages-at-death: for every age group α, the arithmetic mean was computed. For a given sample of size N, where i denotes a country, our measure of dispersion can be written as:

$$\sum_{i=1}^{N} KLD\left(Q \parallel P_i\right) \; with \; Q\left(\alpha\right) = \frac{1}{N}\sum_{i=1}^{N} P_i\left(\alpha\right)$$

Instead of focusing on the particular value of the KLD, in our exercise we set 1960 as the base year.

Motivated by concerns of structural similarities and dissimilarities, we turn to the well-known literature on the effect of political and economic transitions of former Soviet countries, a process which has exerted a major influence in the form of a mortality shock for said countries (Leon, Saburova and Tomkins, 2007, and Meslé, 2004). In order to perform an exploration of the importance of convergence clubs, we split the sample in two large groups: Western and Eastern countries (Table 1). Countries within Eastern Europe that experienced communism belong to our Eastern group of countries. In the next section, we first present a graphical analysis of the mean and variance for both samples then comment on the evolution of the KLD in the period 1960-2008.

3.3. Results

3.3.1. Mean-variance analysis. Before we turn to KLD trends we first consider the traditional mean and variance study of the mortality distribution. We provide a graphical analysis of the mean and variance of the ages-at-death distributions over time that also uncovers interesting features of the epidemological transition. A remarkable feature of the data is that the organization of Western countries in 1960 is mostly along a line that orders countries in the space of high life expectancy — low inequality and low life expectancy — high inequality. This correlation reflects the reality of countries in different stages of their epidemiological transitions. After 50 years, the picture that emerges is drastically different. Although there have been common trends among the majority of countries, i.e. the generalized reduction of variance and the increase in life expectancy, the resulting distribution of countries encompasses a heterogenous landscape. The previous correlation is less apparent, with countries with new profiles emerging from the distribution.

FIGURE 1. Variance and life expectancy profiles of Western countries

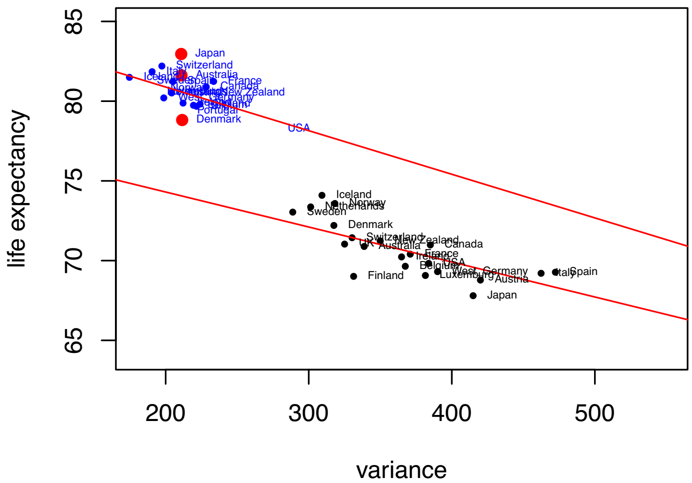

Figures 1 and 2 plot the samples of Western and Eastern countries in the mean-variance space for the years 1960 and 2008. In blue are the variance and life expectancy profile for the year 2008 and in black for the year 1960.

FIGURE 2. Variance and life expectancy profiles of Eastern countries

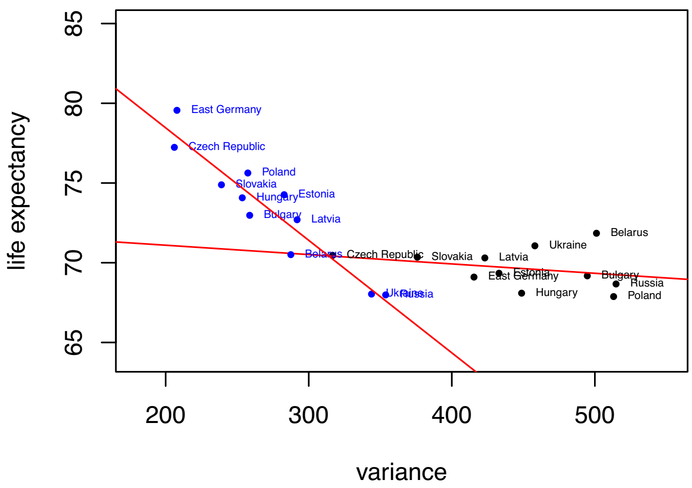

For instance, we observe countries with similar variances (Denmark, Australia, Japan) but with large differences in life expectancy. Performing a similar exercise for Eastern countries yields an interesting contrast. In the latter case, the historical evolution is reversed, with countries evolving from many different profiles to the single dimension ordering as seen amongst Western countries in the 1960's. The mean-variance profiles we report are coherent with previous work (Robine, 2001, and Vaupal, Zhang and van Raalte, 2011) that relates the relationship between life expectancy and variance to the different stages of the epidemiological transition.

In the next section we provide an evaluation of convergence based on the KLD. The main strength of the KLD is that it encompasses the whole differences in distributions, providing a clearer picture than analysis based on only a collection of moments. This is particularly relevant for the analysis of the mortality distributions as infant mortality breaks the normality of the ages-at-death distribution. The ages-at-death distributions of the earlier periods contained a considerably larger number of infant deaths than the contemporary distributions, making mean-variance comparisons a less accurate evaluation.

3.3.2. KLD trends. Our results indicate a clear pattern of divergence in our sample of industrialized countries, as can be seen in Figure 3. The sum of KLD divergences to the mean ages-at-death distribution is normalized with reference to the starting value in 1960. The KLD for the whole sample is in blue, the Eastern countries in green and Western countries are in red. The KLD diminishes from the 1960's until the 1970's, then rises over the 1980's. In fact, the KLD more than doubles by the end of the 1980's, a period that coincides with the dissolution of the Soviet Union, before returning to the positive trend similar to the earlier part of the 1980's. The trends reported remain very similar for both genders, lending support to the robustness of our findings. Figure 4 depicts the KLD for the whole sample is in blue, the female-only in brown and male-only in black.

FIGURE 3. Disparities across mortality distributions

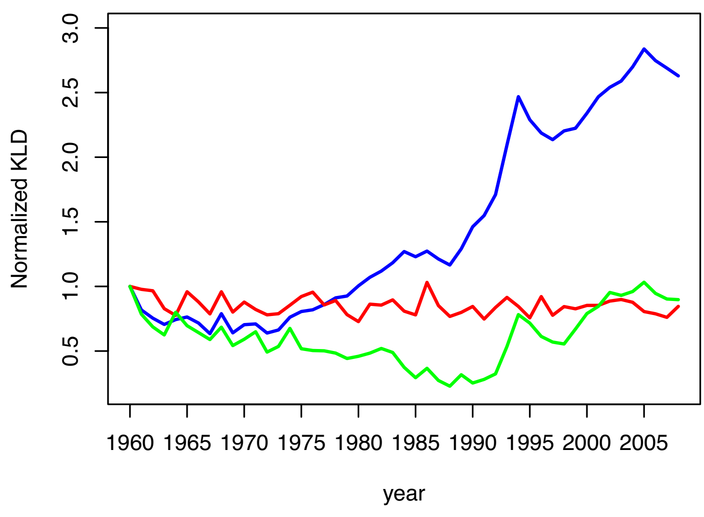

It is of note that the increase in male disparities is of a larger magnitude than that of the female distribution. While by 2008 the female's sample KLD is 1.76 larger that in 1960, the male's sample KLD is over 3 times larger.

FIGURE 4. Disparities across mortality distributions by gender

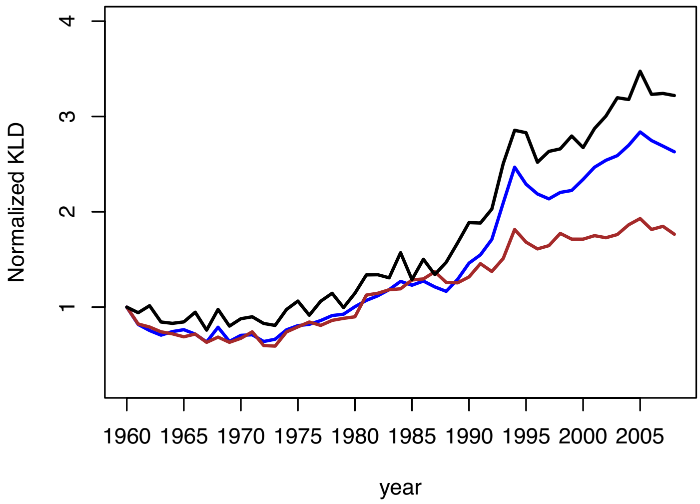

The relevance of taking into account the existence of clubs of countries becomes apparent when observing the differences in the trends between the two groups of countries. In the sample of Western countries we find a relatively flat profile of convergence for mortality patterns (Figure 3); that is, no reduction in the overall differences in mortality patterns. Our conclusion of no convergence is mainly based on the comparison with the strength of the other trends we report. Given the volatility of the KLD a statistical analysis of the trend is very sensitive to the period considered. The sample of Eastern countries (Figure 3) reveals, as expected, a break immediately following the dissolution of the Soviet Union. The first 30 years are marked by convergence, whereas in the late 1980's the trend is completely reversed and we observe a large increase in the differences, followed by sustained increases in the KLD.

3.4. Discussion

Our main finding suggests that, while the reduction in differences due to the catching up process of countries at earlier stages of development is an important catalyst towards convergence (Wilson, 2001), in the later stages structural differences between countries have emerged and halted the convergence process, with young adult mortality differences playing an important role (Gillespie, Trotter and Tuljapurkar, 2014). Based on previous research that provided evidence of a mortality shock in former Soviet countries, we investigate the trajectories of two distinct groups of countries, Eastern and Western countries. Our observations on Eastern countries, in line with previous studies (McMichael, Mckee, Shkolnikov and Valkonen, 2004), point out that structural sources of differences may be strong enough to set back the realized gains in life expectancy. We also find evidence of the lack of a robust convergence process even amongst Western countries.

While our results uncover group specific trends, our classification of countries does not identify convergence clubs. There exists a period of convergence for Eastern countries prior to the 1990s but the recent history of development shows that Eastern countries are no longer approaching a common distribution. Given the robustness of the lack of overall convergence and the crucial role that group-specific dynamics play, further research is needed to uncover the determinants of club membership and pertinence.

The discussion on the determinants of club membership can be divided into two areas of study depending on whether the focus is on the differences between developed and developing countries or the differences amongst industrialized countries. When focusing on the former, the literature has highlighted the existence of mortality traps (Vallin and Meslé, 2004). That is, there is evidence that there might exist two types or clubs of countries, ones with low life expectancy and low growth of life expectancy and ones with high life expectancy and high growth of life expectancy. In order to transition from one club to another, a certain threshold of life expectancy must be reached. However, the variables that define a club still remain relatively unexplored. A first step in the search for these structural conditions can be to draw from the rich literature on the historical path of mortality transitions for current high life expectancy countries (Cutler, Deaton and Lleras-Muney, 2006). When considering industrialized countries, likely suspects are the variables that influence young adult mortality. Given that a large part of the disparities across countries \begin{tabular}{llll} \multicolumn{3}{c}{ } \\ \multicolumn{2}{c}{\bf{Western countries}} & \multicolumn{1}{c}{\bf{Eastern countries}} \\ \hline Australia & Luxemburg & Belarus \\ Austria & Netherlands & Bulgary \\ Belgium & New Zeland & Czech Republic \\ Canada & Norway & East Germany \\ Denmark & Portugal & Estonia \\ Finland & Sweden & Hungary \\ France & United Kingdom & Latvia \\ Iceland & US & Poland \\ Ireland & West Germany & Russia \\ Italy & ~ & Slovakia \\ Japan & ~ & Ukraine \\ \hline \end{tabular}in mortality patterns can be attributed to young adult mortality (Gillespie, Trotter and Tuljapurkar, 2014), it is reasonable to believe that the variables influencing mortality for those age groups might play an important role in determining the formation of convergence clubs.

3.5. Appendix

TABLE 1. List of countries

Western countries		**Eastern countries**
Australia	Luxemburg	Belarus
Austria	Netherlands	Bulgary
Belgium	New Zeland	Czech Republic
Canada	Norway	East Germany
Denmark	Portugal	Estonia
Finland	Sweden	Hungary
France	United Kingdom	Latvia
Iceland	US	Poland
Ireland	West Germany	Russia
Italy		Slovakia
Japan		Ukraine

CHAPTER 4

A cohort perspective of the effect of unemployment on fertility

4.1. Introduction

Two broad perspectives feature in the extensive literature on the relationship between economic conditions and fertility. The first studies fertility and the economic cycle as dynamic phenomena where shifts in fertility affect economic conditions and vice versa (Lee, 1973). However, as evidence on fertility-based economic cycles has become more tenuous over time (Poterba, 2001 and Abel, 2001), a second perspective, that of focusing more on the effects of economic conditions on fertility, has emerged.

This second perspective has gained momentum in part because of the recent recession from 2008 to 2009[1]. The recession has resulted in a renewed interest in the effects of economic conditions on fertility because it has coincided with a period of low fertility rates in developed countries (Kohler, Villari and Ortega, 2002). Importantly, it is now well-understood that sustained low fertility rates pose a serious threat to the sustainability of transfer schemes such as the pension system. Developed countries have experienced a downward trend in fertility rates in the last few decades. Some researchers claim that part of this downward trend can be attributed to "delayed children". That is, individuals still choose to have children but choose to have them later in their fertility lives. According to this strain of thought, families that were temporarily postponing fertility and were due to start having the "delayed children" were yet again halted by the onset of the recession (Goldstein, Kreyenfeld, Jasilionene, 2013) in 2008.

Among the variety of channels through which economic conditions affect fertility, current reseach has highlighted the importance of unemployment (Sobotka, Skirbekk and Philipov, 2011). In particular, a few key research questions have become of interest. The first is, what are the differences between the permanent and temporary effects of unemployment? It is possible that part of the current depression in fertility corresponds to families opting to postpone childbearing. The second question is: what are the effects

[1] These are dates for the United States according to the National Bureau of Economic Research. Exact periods vary by country.

of unemployment at different childbearing ages? In a context of a rise in mothers' ages-at-birth it becomes important to distinguish between the intensity of the effects by age and accurately predict the effects of the economic crisis. Finally, what are the separate effects of the cycle or variations in unemployment and the effects of the structural levels of unemployment on fertility?

In this chapter, I empirically address these three questions through the lens of cohort fertility. I study 12 cohorts with birth years from 1957 to 1969 for eight developed countries for which I have complete histories of fertility and unemployment. The main advantage in using a cohort approach is that it allows capture of temporary and permanent effects of unemployment at each age and can clearly separate the effects of unemployment trends on fertility or evaluate postponement beyond specific birth orders.

I estimate the effects of unemployment on fertility using a panel data Bayesian approach. The Bayesian approach, used here with weak or uninformative priors, can provide gains in efficiency, which is particularly relevant given the relatively small size of the panel. My findings suggest distinct effects of both the levels and the cyclical variations of unemployment. Throughout this chapter, I define the level of unemployment as the average level of unemployment during a period whereas the cycle is captured by the positive and negative yearly variations within a given time span. I elaborate on these two concepts in the presentation of my empirical strategy. While higher levels of unemployment negatively affect fertility, the effect of cyclical variations depends on the age at which they are experienced. Generally speaking, at younger ages there is no effect while for later periods we observe countercyclical fertility. That is, under certain conditions cyclical increases of unemployment can lead to raises in fertility rates. Interestingly, I find evidence of permanent and temporary responses associated to both level and cyclical variations in unemployment. This is particularly relevant given the puzzling and seemingly contradictory perspectives in the literature on fertility and unemployment as my approach suggests that both perspectives capture true effects, and indeed are not contradictory when appropriately accounting for level and cyclical effects.

The chapter is organized as follows. First, I discuss the background literature within which this work is situated and the places where I make contributions. Then, I present the data with my operationalization of the research questions and the Bayesian estimation method. The next section contains the empirical results, which I discuss with a theoretical model that rationalizes the heterogeneity by age of the effects of the level and variation of unemployment. The last section concludes.

4.2. Literature review

Most studies conclude that worsening economic conditions lead to lower fertility (for a recent comprehensive review of the empirical literature, see Sobotka, Skirbekk and Philipov, 2011). For instance, Goldstein, Kreyenfeld, Jasilionene and Karaman Örsal (2013) study the effect of unemployment on fertility in the context of the current recession. However, earlier research has shown (Butz and Ward, 1979) that in a context of high levels of female labor participation we can encounter countercyclical fertility. This empirical result is aligned with the conclusion derived from theoretical microeconomic models that identify recessions as periods where the opportunity cost of having children for working women diminish (see Becker, 1981 and Heckman, 2014). It is commonly accepted that this is the driving force behind the observed negative relationship between female labor participation and fertility at the country level in earlier papers in the literature.

More recent work, however, has claimed that the earlier negative relationship between participation levels and fertility may have changed from the late 1980's onwards (Ann and Mira, 2002)[2]. Adsera (2004), for instance, finds that high unemployment levels are associated with lower fertility. Since countries with high levels of participation today coincide with countries with lower levels of unemployment, the effect of unemployment

[2] There is still controversy over whether the relationship between participation and fertility has changed sign. For instance, Kögel (2004) does not find evidence of this change using time series models).

levels could be enough to reverse the earlier negative relationship. A possible pathway for this finding is presented in Ann and Mira (2001). The authors find support for the hypothesis that high unemployment leads to delays in family formation and marriage; through this mechanism, unemployment thereby lowers fertility. Thus, a tension exists in the current literature. Higher levels of unemployment can decrease the opportunity costs of having children, thereby increasing fertility rates. Yet, the patterns in existing countries today seems to correlate high unemployment levels with lower fertility rates.

In this chapter, I propose to reconcile the tension behind these two findings by studying the differential effects of the unemployment cycle and unemployment level. On the one hand, temporary increases in unemployment, what I term unemployment cycle, can increase fertility through lowering the opportunity cost of children via foregone wages. On the other hand, high unemployment levels can drive down fertility because, among other reasons, high unemployment increases the opportunity costs of building a professional career. The final relationship between unemployment and fertility is then the sum of these two conflicting forces. I consider both effects and as such include both in my models. Unemployment cycle and level effects are both statistically significant, with the overall direction of the total effect of unemployment depending on the relative strength of each effect. Without considering the breakdown of the unemployment effect into both cycle and level effects, as I propose here, we would have more difficulty identifying precisely how and in what directions unemployment can affect fertility rates.

A relatively unexplored topic has been the distinction between the permanent and temporary effects of unemployment (Sobotka, Skirbekk and Philipov, 2011). Neels, Theunyck and Wood (2013) are among the few papers that study the topic. The authors finds evidence of postponement in a microeconometric model that relates increases in the probability of occurrence of a first birth at later ages with the level of unemployment at younger ages. In another work, Goldstein, Kreyenfeld, Jasilionene and Karaman Örsal (2013) disentangle the effect of postponement from a pure reduction of fertility by performing separate analyses on fertility rates and temporally-adjusted fertility rates. Their findings point toestfd the existence of both effects. The cohort perspective I take in this work is particularly well-suited to investigate the persistence of the effects of unemployment. Because I track the fertility and unemployment histories of the cohorts throughout their entire fertile lives, I can include measures of both contemporary and past unemployment in my models. In addition, unlike Neels, Theunynck and Wood (2013), I am not limited to first births. My findings also identify both postponement and permanent effects of unemployment on fertility; interestingly, I find that there is substantial heterogeneity by age group.

4.3. Data and methods

4.3.1. Data. I use data on age-specific fertility rates by cohort from various countries, available in the Human Fertility Database. I obtain the complete fertility history for twelve cohorts from eight developed countries, born from 1957 to 1969. This includes information on fertility rates from 1969 to 2009. The series on unemployment rates is from the online repository of the International Labor Organization (ILO) for the period 1969-2009. My main control variables are female participation rate and net migration flow, which I obtain from the ILO and the Organisation for Economic Cooperation and Development (OECD) statistical repository respectively, also for the period 1969-2009.[3] A complete list of the countries used in my analysis is presented in Table 1 in the appendix.

4.3.2. Descriptive statistics and methods. There are several important advantages to studying the effects of unemployment on fertility at a cohort level. First, given knowledge

[3] Some countries have missing years for the female labor participation variable. In the majority of cases, the missing values are concentrated in the early periods and is usually a single missing year. In some cases, however, up to 5 years are missing. To deal with this, I have imputed the values assuming a linear trend between the last and the first data values before and after the missing years. I have chosen a linear imputation given how steady and linear the participation trend is at a cohort level. Furthermore, the imputation is less important when considering the fact that I use cohort average levels of participation by taking the mean of the participation level from ages 12 to 40.

on the entire fertility history of the cohort, there is no uncertainty over the future possible effects of contemporary and past unemployment. This is not the case, for instance, in Goldstein, Kreyenfeld, Jasilionene and Karaman Örsal (2013). Second, I am able to observe the lifelong effects of unemployment at different ages, which allows for more accurate capture of the effects of postponement. Finally, the cohort level perspective makes it possible to control for trends and unobserved effects at both country and cohort levels[4].

My key dependent variable is cohort age-specific fertility rate. For a given cohort, I have information on the number of births per 1000 women for all years between ages 12 and 40. For tractability, I divide this variable into five age groups: 12 to 20, 21 to 25, 26 to 30, 31 to 35, and 36 to 40 years old. I sum the age-specific cohort fertility rates of years included for each age group. While I have enough data to narrow down the age groups, this would create a problem of dimensionality in the estimation. Since I include the unemployment history in every model, more age groups would quickly increase the number of parameters to estimate.

FIGURE 1. Cohort completed fertility

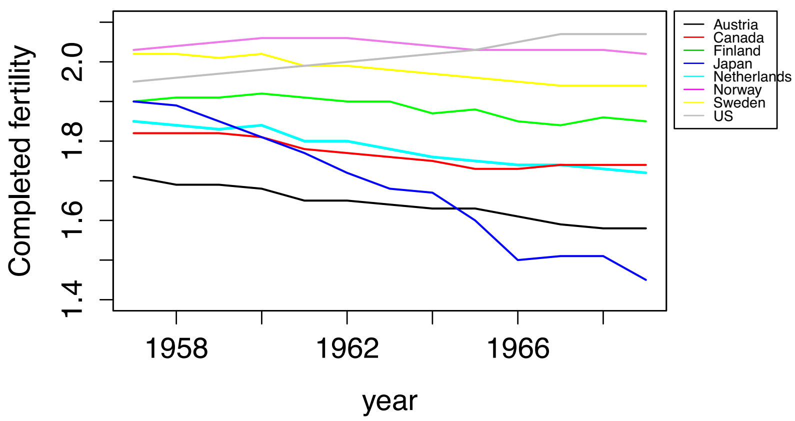

A key issue in the estimation is how to properly account for fertility trends. In the period considered, most countries are experiencing both general delays and reductions

[4]As a consequence of this reasoning, incomplete fertility histories of younger cohorts are not a part of this analysis.

in fertility. Figure 1 plots the cohort completed fertility for the 12 cohorts. The figure shows a general decrease in completed fertility across cohorts for most countries. The two extreme cases are Japan, with a very significant drop, and the United States, which has a slight increase over the period. I also observe delays for all countries in the timing of fertility during this period. This can be seen more clearly in Figure 2, which displays cummulated fertility as a proportion of completed fertility of the cohort from ages 12 to 40 for the United States and Japan in 1957 and 1969.

FIGURE 2. Proportion of births by age 12-40

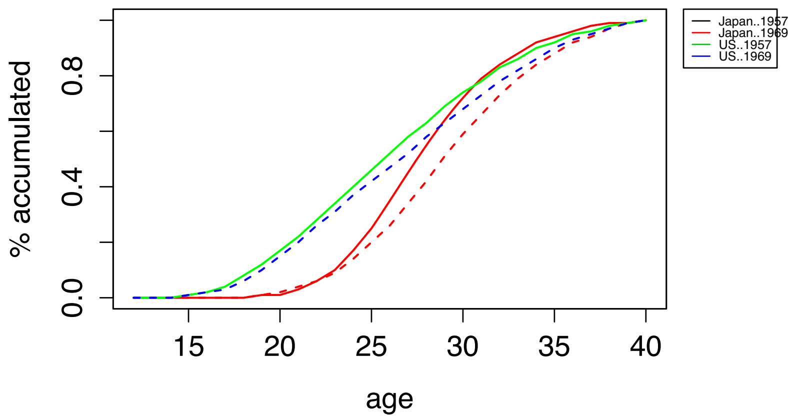

I will address three main questions in this chapter. These three questions require empirical measures for the concepts of unemployment cycle, unemployment level, postponement and foregone fertility. First, do variations in unemployment (cycle) and levels of unemployment (mean) have separate impacts on fertility? I characterize the level of unemployment as the mean of the yearly unemployment for a given childbearing age group. For example, for the group that includes ages 25 to 30, the fertility level is captured by the mean unemployment experienced from ages 24 to 29. On the other hand, the cycle is meant to represent the variation in unemployment. I measure unemployment cycle through the sum of the increases or decreases of unemployment in a given time span. For instance, if in any given group of 5 years unemployment increased between any two years twice and decreased three times, I separately keep track of the sum of the two increases and the three decreases. Do families postpone fertility when facing adverse economic conditions? (and hence attenuate the contemporary negative effect of unemployment and fertility by increasing fertility at later ages?) I introduce variables on the history of both unemployment levels and cycles on my modeling of cohort fertility. My specification includes a complete record of the unemployment levels and both contemporary and lagged measures of unemployment cycles. The objective is to be able to understand the life-long effects of unemployment.

4.3.3. Model and estimation. I organize my panel so that countries are treated as individuals and the different cohorts correspond to observations of a given country in different moments in time. The structure of my models reflect the life cycle perspective of cohort fertility. I model each age group separately including, when relevant, information on their completely history of unemployment. For instance, for the first age group, ages between 12 and 20, I estimate the age group cummulative fertility $k \in (1,...,8)$, cohort $c \in (1,...12)$, and age group $a = 20$ in the following way:

$$CF_{k,c,20} = \alpha_{k,20} + t_{c,20} + \beta_{20}u_{k,20,c} + \gamma_{20}p_{k,20,c} + \theta_{20}n_{k.20,c} + \delta_{20}Z + e_{k,20,c}$$

My dependent variable here is the cummulated fertility of an age group, $CF_{k,c,20}$. The terms $\alpha_{k,20}$ are age-specific fixed effects and $t_{c,20}$ is a nonparametric country and age group specific trend. In this case, time is measured in terms of cohorts. The term $u_{k,20,g}$ is the mean unemployment between ages 12 and 19[5]. The terms $p_{k,j,g}$ and $n_{k,j,g}$ capture the positive and negative variations in unemployment respectively. The vector Z includes controls controls at a cohort level; I include the cohort mean level of female labor participation and mean level of immigration[6].

[5]The one year lag represents pregnancy.

[6]This is simply the average of both variables across the 28 years I follow the cohort.

The model for the second age group (ages 21 to 25), includes information on the previous level of unemployment both in terms of cycle and level. The model is the following:

$$CF_{k,25,c} = \alpha_{k,25} + t_{c,25} + \beta_{20} u_{k,20,c} + \beta_{25} u_{k,25,c} + \gamma_{20} p_{k,20,c} + \gamma_{25} p_{k,25,c}$$

$$+\theta_{20} n_{k.20,c} + \theta_{25} n_{k.25,c} + \delta_{25} Z + e_{k,25,c}$$

More generally, the model for age groups beyond the first age group can be written in the following way. For a country $k \in (1,...,8)$, age group $a \in (2,...,5)$ and cohort $c \in (1,...12)$, the sum of the fertility each year of the age interval is estimated from the following model:

$$CF_{k,a,c} = \alpha_{k,a} + t_{c,a} + \sum_{i=1}^{a} \beta_i u_{k,i,c} + \sum_{j=a-2}^{a-1} \gamma_j p_{k,j,c} + \sum_{l=a-2}^{a-1} \theta_j n_{k.l,c} + \delta_a Z + e_{k,a,c}$$

I estimate the model with Bayesian panel data methods. In particular, I use the integrated nested Laplace approximation (INLA) algorithm recently developed by Rue, Martino and Chopin (2009). The INLA algorithm estimates a latent Gaussian model, in my case with weak priors. I test the robustness of the results to several specifications of the underlying distribution of the priors and my main results hold.

4.4. Results

4.4.1. Description of the results. Table 1 through 5 depict results from the main models for the five age groups under study. Standard deviations of the estimated parameters are provided in parentheses next to each coefficient, while the three columns to the right are upper bound, median and lower bound for the credible intervals of the estimation.

This is a concept can be somewhat paralleled in interpretation to the significance level in a frequentist approach. The credible intervals probabilistically identify where the estimated parameters lie. We are most interested in effects that do not include 0 within their credible intervals.

TABLE 1. Model 1: Age group 12 to 20

		Credible interval		
variables		**0.025**	**0.5**	**0.975**
Unemployment cycle				
Positive				
12-19	0.0249* (0.0076)	0.0098	0.0248	0.0399
Negative				
12-19	0.0001 (0.00090)	-0.0017	0.0019	0.0001
Unemployment level				
12-19	-0.0007 (0.0029)	-0.0064	-0.0007	0.0050
Controls				
Immigration	0.0091* (0.0037)	0.0018	0.0091	0.0163
Participation	-0.7339* (0.0734)	-0.8780	-0.7340	-0.5897

In the different models I observe that the effect of unemployment varies greatly by age, both in terms of cycle and level and with regard to postponement and reductions in fertility. I later discuss how these results can be interpreted as suggestive evidence of heterogeneity across families in the fertility response to unemployment. The first two models correspond to the 12-20 year (Table 1) and 21-25 year (Table 2) age groups. We can consider these groups as younger periods during which a fraction of the population is still acquiring formal education. In the first age group fertility is heavily concentrated in the later years of the interval. We note expected effects of both types of control variables on the age groups: immigration increases fertility at younger ages whereas participation diminishes it. Interestingly, for ages 12 to 20 only the unemployment cycle matters and not the unemployment level: increases in unemployment have a positive (countercyclical) influence on fertility.

TABLE 2. Model 2: Age group 21 to 25

		Credible interval		
variables		**0.025**	**0.5**	**0.975**
Unemployment cycle				
Positive				
12-19	-0.0230 (0.0284)	-0.0784	-0.0232	0.0332
20-24	0.0023 (0.0122)	-0.0215	0.0022	0.0265
Negative				
12-19	0.0004 (0.0021)	-0.0038	0.0004	0.0046
20-24	0.0039* (0.0013)	0.0013	0.0039	0.0065
Unemployment level				
12-19	-0.0145* (0.0071)	-0.0284	-0.0145	-0.0005
20-24	-0.0092* (0.0018)	-0.0127	-0.0092	-0.0057
Controls				
Immigration	0.0194* (0.0077)	0.0041	0.0194	0.0346
Participation	-0.7224* (0.1572)	-1.0339	-0.7214	-0.4166

This last result does not carry on to the second model for fertility between ages 21 and 25. The overall effect of unemployment in this model appears to come from the negative influence of unemployment levels. I find that both contemporary and lagged unemployment levels have a negative impact on fertility. That is, experiencing high levels of unemployment at an earlier age diminishes fertility between 21 and 25 years old. The results for the second age group are consistent with earlier findings (for instance Adsera (2004) that link these effects to the career choices of families.

TABLE 3. Model 3: Age group 26 to 30

		Credible interval		
variables		**0.025**	**0.5**	**0.975**
Unemployment cycle				
Positive				
20-24	0.08520* (0.0153)	0.0548	0.0852	0.1151
24-29	0.0508* (0.0195)	0.0123	0.0508	0.0891
Negative				
20-24	0.0004 (0.0021)	-0.0038	0.0004	0.0046
25-29	0.0123* (0.0016)	0.0091	0.0123	0.0155
Unemployment level				
12-19	-0.0476* (0.0085)	-0.0641	-0.0476	-0.0306
20-24	-0.0154* (0.0037)	-0.0226	-0.0154	-0.0081
25-29	-0.0238* (0.0030)	-0.0297	-0.0238	-0.0178
Controls				
Immigration	-0.0083* (0.0116)	-0.0310	-0.0084	0.0145
Participation	1.0858* (0.2386)	0.6112	1.0881	1.5483

Model 3 encompasses the 26-30 year old age group (Table 3). Within this age group, the effects of the controls are reversed, which is coherent with the idea that higher levels of female labor participation are associated with a general postponement in childbearing. In this model I find effects for both the cycle and level of unemployment. Unemployment levels for the earlier years of the cohort (both ages 12 to 20 and 21 to 25) and the contemporary level of unemployment have a negative effect on fertility. In addition, the impact of the cycle appears to be quite complex. First, contemporary increases in unemployment have a countercyclical effect and increase fertility. At the same time, reductions in unemployment also have a positive effect on fertility. Furthermore, increases in unemployment that took place between ages 21 and 25 also increase fertility for this age group indicating there may have been some postponement of children at earlier ages.

TABLE 4. Model 4: Age group 31 to 35

		Credible interval		
variables		**0.025**	**0.5**	**0.975**
Unemployment cycle				
Positive				
24-29	0.0482* (0.0124)	0.0240	0.0482	0.0727
30-34	0.0599* (0.0161)	0.0283	0.0598	0.0916
Negative				
25-29	-0.0019 (0.0015)	-0.0049	-0.0019	0.0010
30-24	0.0033* (0.0016)	0.0002	0.0033	0.0064
Unemployment level				
12-19	0.0148* (0.0046)	0.0058	0.0147	0.0238
20-24	0.0015 (0.0021)	-0.0026	0.0015	0.0057
25-29	-0.0013 (0.0026)	-0.0064	-0.0014	0.0037
30-34	-0.0023 (0.0020)	-0.0061	0.0016	0.0016
Controls				
Immigration	-0.0178* (0.0078)	-0.0332	-0.0178	-0.0025
Participation	0.2745* (0.1256)	0.0251	0.2751	0.5205

The complexity of the effects of the variation of unemployment persists in the model for the age group between 31 and 35 (Table 4). In this model the previous effects of variation also coexist. Increases in unemployment during ages 26 to 30 seem to be associated with postponement since they positively affect fertility for the current age group. In addition, both contemporary reductions and increases of fertility have a positive effect on unemployment. However, the main difference with Model 3 is that fertility is unaffected by current and all levels of unemployment with the exception of the unemployment level between ages 12 and 20. This last effect can be interpreted as further evidence of postponement.

TABLE 5. Model 5: Age group 36 to 40

variables		Credible interval 0.025	0.5	0.975
Unemployment cycle				
Positive				
30-34	- 0.0107 (0.0090)	-0.0283	-0.0107	0.0068
35-39	-0.0147* (0.0102)	-0.0346	-0.0147	0.0053
Negative				
30-34	-0.0019 (0.0009)	-0.0012	0.0007	0.0026
35-39	0.0033 (0.0014)	-0.0023	0.0005	0.0033
Unemployment level				
12-19	0.0075* (0.0031)	0.0015	0.0075	0.0135
20-24	0.0021 (0.0011)	-0.0002	0.0021	0.0043
25-29	0.0009 (0.0017)	-0.0024	0.0009	0.0042
30-34	-0.0009 (0.0013)	-0.0035	-0.0009	0.0017
35-39	0.0013 (0.0021)	-0.0029	0.0013	0.0056
Controls				
Immigration	0.0179* (0.0048)	0.0084	0.0179	0.0273
Participation	0.3816* (0.0840)	0.2171	0.3813	0.5474

The final model (Table 5) explains the effects of unemployment on fertility between 36 and 40 years old. The effects of unemployment are quite different from the earlier periods. First, only positive contemporary variations of unemployment have an unambiguous impact: they reduce fertility (albeit in a more generous credible interval). In terms of levels, the only impact comes from the unemployment level at ages 12 to 20, hinting at a postponement effect.

4.4.2. Discussion of the results. Estimating models on the whole fertility history of cohorts provides the distinct advantage of being able to interpret the results as a reflection of lifelong choices. The caveat is that since the data is in the aggregate, a variety of responses to unemployment coexist in my results. In general terms, the results are coherent with the narratives that emphasize the importance of career considerations in the timing of fertility (Adsera, 2004). Between ages 21 and 25, fertility is negatively affected by high levels of unemployment and there are no observable increases of fertility associated to positive variations in unemployment. This seems to indicate lessened opportunity cost due unemployment is not a strong enough force to compensate for the potential negative effects of interrupting early on labor market involvement. Between ages 26 and 35 I detect countercyclical effects of unemployment, as well as mixed results for the effects of unemployment levels. Finally, unemployment in the older age group seems to be strictly related to lower fertility, particularly through the effect of positive unemployment variations.

It is noteworthy that there exists some seemingly contradictory effects in the results. For example, for some age groups, increases in unemployment simultaneously lead to higher fertility in the period in which they occur (countercyclical fertility) and to increases in the fertility at later stage in life (postponement). I also observe a positive effect of reductions and increases in unemployment for a same age group. These situations are suggestive of heterogeneous family responses to similar conditions. This could be explained through the uneven impact of unemployment across families or because families with different characteristics might be responding with different strategies to the same situation. In any case, these situations indicate that further microeconometric analysis in the spirit of Neels, Theunynck and Wood (2013) are needed. In the next section I explore the career effects of unemployment with a formal model of fertility timing.

4.5. Discussion

In this section, I rationalize my empirical findings with a model of fertility delays inspired in the work by Radjan (1999) and Adsera (2004). I propose an explanation for the heterogeneity across ages of the effects of unemployment based on the career and experience effects on unemployment. This exercise is conducted to verify the feasibility of observing the previous section's empirical results in a formal model with a reasonable set of assumptions on the career effects of fertility. In the model, labor market career considerations are behind the timing of fertility.

Although I follow a similar model structure as in Radjan (1999), the mechanism through which unemployment affects fertility is different. Radjan models childbearing as an irreversible decision in the context of income uncertainty. The main finding of the model is that as income levels soar, fertility postponement becomes more attractive. The reason is that the cost of having a child is consumption foregone, which due to the concavity of the utility function is higher when facing lower levels of income. Adsera (2004) modifies Radjan's model to incorporate the effect of childbearing on unemployment. In a sense, Adsera's work endogenizes income uncertainty to a certain extent through an unemployment probability function that depends on fertility decisions. The main differences between my model and Adsera's are two. While Adsera models fertility in a two period model I include an additional period to allow for heterogeneity across ages of the effects of unemployment. My empirical findings, along with those existing in the literature, strongly support the hypothesis that there is significant variability in the fertility response to unemployment across ages. In addition, I distinguish between the effect of the general unemployment level and an increase in unemployment. In my model this is captured by directly considering a general parameter, the unemployment level, and studying the decision of a family in a situation of unemployment.

In my simplified setting, I consider a model where a family (or female) lives 3 periods and must to decide whether and when to have a maximum of one single child in either of the first two periods. It is a model concerned exclusively with fertility timing. The utility of a family in any given period is given by their consumption and the utility derived from the child:

$$U(c_t) + \phi \kappa_t$$

where κ_t is an indicator function equal to one if there a child is present in the household and ϕ is the utility brought by him/her.

In each period, prior to the family making a fertility decision, the family learns about their employment status, which depends upon their previous participation in the labor market and on whether a child was born in the previous period. The probability of employment is given by the function:

$$e_t\left(\sum_{1}^{t-1} l_i; \rho_t; \kappa_{t-1}\right)$$

where ρ_t is the aggregate level of unemployment , κ_{t-1} captures if a child was born the previous period and l_i the amount of time spent in labor market activities in previous periods. Previous attachment to the market, represented by the sum of labor market time over the life time, increases the probability of employment: $e(.)_l > 0$. However, it has diminishing returns: $e(.)_{ll} > 0$, a standard assumption in the literature of wage determination. I also assume that previous experience is more valuable when there is a high level of unemployment: $e(.)_{l\rho} > 0$.

This last assumption is crucial for some of the results in the model. Recent work in the empirical literature on wage determination (Oreopoulos, von Wachter and Heisz, 2012) has found that workers entering the market during recessions are more affected by the adverse labor market conditions than those with more years of experience. In my model, the fact that past experience is more valuable in terms of employabiltiy in bad times has the direct implication that structural unemployment diminishes fertility.

The time allocation of the familiy is described as follows. Conditional on employment, their time allocation is given by:

$$T \geq l_t + m_t$$

In each period, families have T amount of time; m_t represents time allocated towards childcare. The decision to have a child has two direct implications in the model. First, having a child restricts the labor market experience and hence future employability. Furthermore, following Adsera, having a child imposes a direct penalty on employment in the next period: $e(.)_{\kappa} < 0$. Regardless of the employment status, the family is guaranteed a minimum income y.[7] Hence, the budget constraint of the family is given by:

$$wl_t + y \geq c_t$$

The model can be solved by backwards induction and a solution is given by the decision to give birth the first period, in the second or to forego childbearing altogether.

4.5.1. The effect of the general level of unemployment. We first consider the effect of an increase on the general level of unemployment, which corresponds to the unemployment level effects of the previous section's empirical exercise. Instead of solving the model with a particular set of parameters, I assume a given history and observe how an increase in the level of unemployment affects the fertility decision. In particular, I consider the case where the family is childless, employed by period 2 and has had employment since period 1. This scenario is interesting because the value of further employment is the lowest of all possible histories. As it is thus the least likely case where we might observe a delay in fertility, it is a "hard case" for observing consistent results with my empirical findings in the previous section. Suppose that ρ (the unemployment level) increases, then it follows from the assumptions on the employment function that foregoing fertility becomes more attractive. To see why, consider the indirect utility of having a child:

$$V^k = U(.)_{|\kappa_2=1} + \phi\kappa_2 + e(.)_{|\kappa_2=1}\left(U(.)_{|l_3>0} + \phi\kappa_3\right) + \left(1 - e(.)_{|\kappa_2=1}\right)\left(U(.)_{|l_3=0} + \phi\kappa_3\right)$$

and the utility of not having a child:

[7]This can be interpreted, for instance, as the unemployment subsidy.

$$V^{nk} = U(.)_{|\kappa_2=0} + e(.)_{|\kappa_2=0} U(.)_{|l_2>0} + \left(1 - e(.)_{|\kappa_2=0}\right) U(.)_{|l_3=0}$$

where $U(.)_{|\kappa_t=1}$ is the utility of consumption in the first period when the family has a child and the rest of terms have a similar interpretation. The difference between the two indirect utilities consists in three terms. The first term is the utility derived from the child. The second term is the additional probability of employment evaluated at the additional utility gained from employment. The final term is the loss of income in the current period. The family simply chooses max $\{V^k, V^{nk}\}$. Denote $\triangle_3 e^{k_2=0}_{k_2=1}$ the difference in the future probability of employment between a familiy with and without a child, $\triangle_2 U^{k_2=0}_{k_2=1}$ the difference in the utility derived from consumption from having a child, and $\triangle_3 U^{l_3=0}_{l_3>0}$ the difference in utility for the future period between an employed family and an unemployed family

The family decides against having a child if the following is true:

$$\Omega = \triangle_3 e^{k_2=0}_{k_2=1} \triangle_3 U^{l_3=0}_{l_3>0} + \triangle_2 U^{k_2=0}_{k_2=1} - 2\psi > 0$$

Given the assumption that a greater past attachment to the labor market is more important in times of higher general levels of unemployment, $e(.)_{l\rho} > 0$, this implies that $\Omega_\rho > 0$ since $\frac{d\triangle_3 e^{k_2=0}_{k_2=1}}{d\rho} > 0$. That is, the option to forego fertility, in the case of the third period, or more generally to delay fertility becomes more attractive as the general level of unemployment rises.

4.5.2. The differential effect of unemployment at younger ages. Another finding of the model is that unemployment has countercyclical effects that can occur in the middle stage of a fertility life for a given family. In the model, this mirrors the decision of a family that is currently unemployed. The idea behind countercyclical effects is that individuals facing an unemployment spell may face an incentive to have children as there is no longer an opportunity cost in terms of foregone salaries. However, whether this incentive is a dominant force depends on the family's previous labor experience and also on their age. In the model, this is because at older ages the probability of having a positive labor history is larger. To illustrate the forces at play, I turn to the effect of unemployment in periods 1 and 2 of the model separately.

As with the previous exercice, I study a limit case. Consider first the case of a family that is unemployed in the first period and knows that choosing to not have a child in the first period makes it a dominant strategy to have a child in the second period regardless of their employment status. In this scenario the problem becomes strictly a matter of timing.

In this case, the family delays childbearing if the value of early labor experience is high enough. To see why, consider lifelong utility of the two options. Recall that the presence of a child is represented by the indicator function κ_t, hence only consider $\{\kappa_1 = 0, \kappa_2 = 0, \kappa_3 = 0\}$ for V^{nk}.

$$V^k$$

$$U(\cdot)_{|l_1=0}+\phi\kappa_1+e_2(\cdot)_{|k_1=1}U(\cdot)_{l_2>0}+\phi\kappa_1 \qquad U(\cdot)_{|l_1=0}+\phi\kappa_1+(1-e_2(\cdot)_{|k_1=1})U(\cdot)_{l_2=0}+\phi\kappa_1$$

$$e_3(\cdot)_{l_2>0}U(\cdot)_{l_3>0}+\phi\kappa_3 \qquad (1-e_{t+2}(\cdot)_{|l_{t+1}>0})U(\cdot)_{l_2=0}+\phi\kappa_3 \qquad e_3(\cdot)_{|l_2=0}U(\cdot)_{|l_3>0}+\phi\kappa_3 \qquad (1-e_3(\cdot)_{|l_2=0})U(\cdot)_{|l_3=0}+\phi\kappa_3$$

The comparison between the two options has three main components. The first term is the additional utility provided by having the child in period one . The second term is the higher expected wages in the second period if a child is born in the first period, but at a lower probability of employment in the second period.

$$\Phi = \left[e_2\left(.\right)_{|k_2=0} U\left(.\right)_{|l_2=T-m}\right] - \left[e_2\left(.\right)_{|k_2=1} U\left(.\right)_{|l_2=T}\right]$$

The third term is the effect on employability in the third period, evaluated at the utility gains from employment, of the two fertility timing options. Having the child while unemployed is a strategy that emphasizes the lower opportunity cost of childbearing while unemployed. It implies lowering the opportunity of employment in the second period but freeing up time in case employment is found in that period. On the other hand, the decision against having a child when unemployed prioritizes early involvement in the labor market by increasing the chances of employment in the second period and lowering them the third. In terms of the model, the comparison is between these two terms:

$$\Psi = \left[e_2\left(.\right)_{|k_2=0} e_3\left(.\right)_{|k_3=1,l_3=T-m} U\left(.\right)_{|l_3>0}\right] - \left[e_2\left(.\right)_{|k_2=1} e_3\left(.\right)_{|k_3=1,l_3=T} U\left(.\right)_{|l_3>0}\right]$$

Which strategy prevails in this comparison depends on the concavity of the employment probability function. If the returns to experience have sufficiently high diminishing returns, then early involvement in the labor market is the dominant strategy. To see why, suppose that both childbearing and working part-time incur the same penalty in terms of employability, and that this penalty takes form of a reduction in accumulated experience: $e\left(\sum l - p\right)$ where p denotes the penalty. Then, without considering general unemployment, the comparison becomes:

$$\Psi = \left[e_2\left(\sum l\right)_{|k_2=0} e_3\left(\sum l - 2p\right)_{|k_3=1,l_3=T-m} U\left(.\right)_{|l_3>0}\right]$$

$$- \left[e_2\left(\sum l - p\right)_{|k_2=1} e_3\left(\sum l\right)_{|k_3=1,l_3=T} U\left(.\right)_{|l_3>0}\right]$$

If $e()$ is concave *enough* then, $\frac{e_2(\sum l)_{|k_2=0}}{e_2(\sum l-p)_{|k_2=1}} > \frac{e_3(\sum l-2p)_{|k_3=1,l_3=T-m}}{e_3(\sum l)_{|k_3=1,l_3=T}}$ and the family decides to postpone childbearing. The intuition is that the concavity of $e(.)$ makes facing employability penalties more attractive once the family has accumulated labor experience.

The main difference in evaluating the decision of an unemployed family in their second period is that the costs do not include such lifelong career impacts. In this case, the family balances the lower possibilities of employment in the third period against the utility derived from having a child. Furthermore, a larger potential pool of experience lowers the value of postponing (in this case, foregoing) fertility. For the sake of illustration, consider the tradeoff faced by a family that is currently unemployed in period two, but was employed in period one:

$$V^k = U(.)_{|l_2=0} + \phi\kappa_2 + e_2(.)_{|k_2=1} U(.)_{|l_3>0} + \phi\kappa_3 + \left(1 - e_2(.)_{|k_2=1}\right) U(.)_{|l_3=0} + \phi\kappa_3$$

$$V^{nk} = U(.)_{|l_2=0} + e_2(.)_{|k_2=0} U(.)_{|l_3>0} + \left(1 - e_2(.)_{|k_2=0}\right) U(.)_{|l_3=0}$$

The family will decide to forego childbearing if the utility of the additional expectation of employment is sufficiently large. Denote $\triangle_3 e^{k_2=0}_{k_2=1}$ the difference in the future probability of employment between a family with and without a child, and $\triangle_3 U^{l_3=0}_{l_3>0}$ the difference in utility for the future period between an employed family and an unemployed family. The family decides foregoing fertility if the following is true:

$$\Gamma = \triangle_3 e^{k_2=0}_{k_2=1} \triangle_3 U^{l_3=0}_{l_3>0} - 2\psi > 0$$

The penalty for childbearing is diminished here, compared to an unemployed family in period one, for two reasons. The first is that, given a previous history of unemployment, the effect of the childbearing penalty is smaller given the concavity of $e(.)$. The second is that the penalty on employment affects only the third and final period, without compounding for additional periods. Naturally, in models with more periods, the lifelong impact on the career of the family is greater the earlier the family has the child. Taken together, the results in this model indicate that, under plausible assumptions, some of the observed empirical relationships between fertility and unemployment are coherent with an explanation based on the career effects of fertility. My theoretical is a contribution to the existing literature on the effects of unemployment on fertility; in particular, I extend existing models to capture the differences in the impact of unemployment across ages.

4.6. Conclusion

In this chapter, I link the complete histories of both fertility and unemployment for 12 cohorts in eight developed countries to study the temporary and permanent impacts of unemployment on fertility. I can reconcile the seemingly contradictory results of earlier related literature that have separately highlighted the importance of counterclycical fertility and structural unemployment by decomposing the effects of unemployment into both cycle and level effects and separately estimating the two. Indeed, my results show that, in order to correctly assess the impact of unemployment, it is important to separately study the effects of unemployment levels and unemployment cycles across age groups. Estimation results on the panel data demonstrate a high variability of the fertility responses across ages. I find that the unemployment level has a negative effect that fades at older ages. The unemployment variation (cycle) triggers both delays and a countercyclical response between ages 25 and 35, however countercyclical effects seem to be the dominant contemporary effect.

In the second part of the chapter, I explored whether my empirical observations are coherent with a formal model of fertility timing with career considerations. In general terms, it is possible to obtain my main results from a model of fertility timing where labor participation has long lasting career effects. In the model, high unemployment levels are associated with fertility delays when labor experience is more important in times of job scarcity. Also, the costs associated with temporary interruptions of labor market participation are greater at younger ages. The age specific effects of unemployment variations can be explained through the lifelong impact of an early interruption of the career.

A natural next step is to bring the life cycle perspective of this work to microeconometric models. While I find a solid set of results regarding the lifelong effects of unemployment that can be interpreted in terms of career choices, some of these results are indicative of heterogeneous responses. For instance, between ages 25 and 35 I observe delays in fertility and a countercyclical behavior, which is coherent with different strategies across families. A microeconometric approach would also benefit from some of the specification choices in my work. First, I have showed how unemployment has effects well beyond the year it takes place, hence it would be necessary to take a life cycle perspective in the analysis. In addition, an insight from my work is the necessity to distinguish the unemployment cycle and level.

Bibliography

[1] A.B. Abel. 2001. Will Bequests Attenuate the predicted meltdown in stock prices when baby boomers retire?. *Review of Economics and Statistics*, 83(4):589-595, 2001.

[2] A. Adserà. Changing fertility rates in developed countries. The impact of labor market institutions. *Journal of Population Economics*, 17(1):17-43, 2004.

[3] A. Adserà. Vanishing children: from high unemployment to low fertility in developed countries. *American Economic Review*, 95(2):189-193, 2005.

[4] N. Ahn, N. and P. Mira. Job Bust, Baby Bust? Evidence from Spain. *Journal of Population Economics*, 14(3):505-521, 2001.

[5] N. Ahn and P. Mira. 2002. A Note on the changing relationship between fertility and female employment rates in developed countries. *Journal of Population Economics* 15(4):667-682, 2002.

[6] T. Bago d'Uva, O. O'Donnell and E. van Doorslaer. Differential health reporting by education level and its impact on the measurement of health inequalities among older Europeans. *International Journal of Epidiemology*, 37(6):1375-1383, 2008.

[7] T. Bago d'Uva, E. van Doorslaer, M. Lindeboom and O. O'Donell. Does reporting heterogeneity bias the measurement of health disparities?. *Health Economics*, 17(3):351-375, 2008.

[8] M. Baker, M. Stabile, and C. Deri. What do self-reported, objective, measures of health measure?. *Journal of Human Resources*, 39(4):1067-1093, 2004.

[9] R.J. Barro and X. Sala-i-Martin. Economic growth. 2nd edition, Cambridge: MIT Press, 2003.

[10] G.S. Becker, T.J. Philipson, R.R. Soares. The quantity and quality of life and the evolution of world inequality. *American Economic Review*, 95(1):277-291, 2005.

[11] G. Becker. A treatise on the family.Cambridge, MA: Harvard University Press, 1981.

[12] A. Blinder. Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 8(4):436-55, 1973.

[13] D.E. Bloom and D. Canning. Mortality traps and the dynamics of health transitions. *Procedures of the National Academy of Sciences*, 104:16044-16049, 2007.

[14] W. Butz, and M. Ward, The Emergence of countercyclical US fertility. *The American Economic Review*, 69(3):318-328, 1979.

[15] A. Calo-Blanco and A. Villar. Education, Utilitarianism, and Equality of Opportunity. Fundación BBVA - Documentos de Trabajo No. 13, 2009.

[16] C. Calsamiglia. Decentralizing Equality Of Opportunity. *International Economic Review*, 50(1):273-290, 2009.

[17] D. Checchi , V. Peragine and L. Serlenga . Income inequality ans opportunity inequality in Europe. *Rivista di Politica Economica*, 98: 266-293, 2009.

[18] Z. Chen and K. Roy. Calculating concentration index with repetitive values of indicators of economic welfare. *Journal of Health Economics*, 28(1):169-175, 2009.

[19] D. Cutler, A. Deaton and A. Lleras-Muney. The determinants of mortality. *Journal of Economic Perspectives*, 20(3):97-120, 2006.

[20] P. Dolan. The measurement of health-related quality of life for use in resource allocation decisions in health care. In A. J. Culyer and J. P. Newhouse, editors, *Handbook of Health Economics*, edition 1, volume1, chapter 32, pages 1723-1760. Elsevier, 2000.

[21] E. van Doorslaer and K. Humphries. Income-related health inequality in Canada. *Social Science and Medicine*, 50(5):663-671, 2000.

[22] E. van Doorslaer and U.G. Gerdtham. Does inequality in self-assessed health predict inequality in survival by income? Evidence from Swedish data. *Social Science and Medicine*, 57(9):1621-1629, 2003.

[23] E. van Doorslaer, X. Koolman and A.M. Jones. Explaining income-related inequalities in doctor utilization in Europe. *Health Economics,* 13(7):629-647, 2004.

[24] J.B. Dowd and A. Zajacova. Does self-rated health mean the same thing across socioeconomic groups? Evidence from biomarker data. *Annals of Epidemiology*, 20(10)-743-749, 2010.

[25] J.B. Dowd and M. Todd. Does Self-reported Health Bias the Measurement of Health Inequalities in U.S. Adults? Evidence Using Anchoring Vignettes From the Health and Retirement Study. *Journal of Gerontology: Social Sciences*, 66(4):478-489, 2011.

[26] R.D. Edwards and S. Tuljapurkar. Inequality in life spans and a new perspective on mortality convergence across industrialized countries. *Population and Devlopment Review*, 31(4):645-674, 2005.

[27] F. Etilé and C. Milcent. Income-related reporting heterogeneity in self-assessed health: evidence from France. Health Economics, 15(9):965-981, 2006.

[28] M. Fleuerbaey and E. Shockhaert. Unfair inequalities in health and health care. *Journal of Health Economics*, 28(1):73-90, 2009.

[29] D.I.S Gillespie, M.V. Trotter and S. Tuljapurkar. Divergence in age patterns of mortality change drives international divergence in lifespan inequality. *Demography*, 51(3):1-15, 2014.

[30] J.R. Goldstein, T. Sobotka and A. Jasilioniene. The end of “lowest-low” fertility?. *Population and Development Review*, 35(4):663-699, 2009.

[31] J.R. Goldstein, M. Kreyenfeld, A. Jasilionene, and K. Örsal, D. Fertility reactions to the "Great Recession" in Europe: recent evidence from order-specific data. *Demographic Research*, 29:85-104, 2013.

[32] J.C. Gower. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, 53(3):325-338, 1966.

[33] J. Heckman .2014. Introduction to a theory of the allocation of time by Gary Becker. IZA Discussion Papers 8424, Institute for the Study of Labor (IZA), 2014.

[34] C. Herrero, R. Martínez. Egalitarian Rules in Claims Problems with Indivisible Goods. *Social Choice and Welfare*, 30: 603-617, 2008.

[35] C. Herrero, A. Villar. (1994) La assistencia sanitaria como un problema de asignacion en presencia de objectivos inalcanzables. In: G. Lopez-Casasnovas, editor. *Analisis economico de la sanidad*, chapter 4, pages 145-164, Servei catala de salut, 1994.

[36] C. Herrero, A. Villar. An Equity-based Proposal for the Evaluation of Health States. Fundación BBVA - Documentos de Trabajo No. 15, 2012.

[37] E.L. Idler and Y. Benyamini. Self-Rated Health and Mortality: A Review of Twenty-Seven Community Studies. *Journal of Health and Social Behavior*, 38(1);21-37, 1997.

[38] E.L. Idler and Y. Benyamini. Community studies reporting association between self-rated health and mortality: additional studies, 1995 to 1998. *Research on Aging*, 21:392-401, 1999.

[39] H. Jurges. Self-assessed health, reference levels and mortality. *Applied Economics*, 40(5):569-582, 2008.

[40] D.W. Johnston, C. Propper and M.A. Shields. Comparing subjective and objective measures of health : evidence from hypertension for the income/health gradient. *Journal of Health Economics*, 28(3):540-552, 2009.

[41] T. Kögel. Did the association between fertility and female employment within OECD countries really change its sign?. *Journal of Population Economics*, 17(1):45-65, 2004

[42] H.-peter Kohler, F. Billari and J. Ortega the emergence of lowest-low fertility in Europe during the 1990s. *Population and Development Review*, 28(4): 641-680, 2002.

[43] R. Lee. Population in preindustrial England: an econometric analysis. *Quarterly Journal of Economics*, 87(4):581-607, 1973.

[44] D. Leon, S. Saburova and S. Tomkins. Hazardous alcohol drinking and premature mortality in Russian: a population based case-control study. *Lancet*, 369(9578):2001-2009, 2007.

[45] H. Llavador, J.E.Roemer. An equal-opportunity approach to the allocation of international aid. *Journal of Development Economics*, 64(1):147-171, 2001.

[46] A.J. McMichael, M. Mckee, V. Shkolnikov and T. Valkonen. Mortality trends and setbacks: global convergence or divergence?. *Lancet*, 362(9415):1155-1159, 2004.

[47] F. Meslé. Mortality in Central and Eastern Europe, *Demographic Research Special*, 2:45-70, 2004.

[48] K. Neels, Z. Theunynck, and J. Wood. Economic recession and first births in Europe: recession-induced postponement and recuperation of fertility in 14 European countries between 1970 and 2005. *International Journal of Public Health*, 58(1):43-55, 2013.

[49] R. Oaxaca. Male-female wage differentials in urban labour markets. *International Economic Review*, 14(3):693-709, 1973.

[50] A.R. Omran. The epidemiologic transition: a theory of the epidiomology of population change. *Milbank Memorial Fund Quarterly*, 49:509-538, 1971.

[51] E. Ooghe, E. Schokkaert, D. Van de gaer. Equality of Opportunity versus Equality of Opportunity Sets. *Social Choice and Welfare*, 28(2):209-230, 2007.

[52] P. Oreopoulos, T. von Wachter and A. Heisz.The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics*, 4(1):1-29, 2012.

[53] V. Peragine, L. Serlenga. Higher education and equality of opportunity in Italy. In: J. Bishop and D. Zheng D, editors. *Research in Economic Inequality*, vol 16, pages 67-97, 2008.

[54] J.M. Poterba. Demographic Structure and Asset Returns. *Review of Economics and Statistics*, 83(4):565-584, 2001.

[55] P. Ranjan . Fertility behaviour under income uncertainty.*European journal of population = Revue européenne de démographie*, 15(1):25-43, 1999.

[56] J.M. Robine. Redefining the stages of the epidemiological transition by a study of the disperion of life spans: The case of France. *Population: An English Selection*, 13:173-94, 2001.

[57] D.J. Roelfs, E. Shor, K.W. Davidson and J.E. Schwartz. Losing life and livelihood: a systematic review and meta-analysis of unemployment and all-cause mortality. *Social Science and Medicine*, 72(6):840-854, 2011.

[58] J.E. Roemer. Equality of opportunity: A progress report. *Social Choice and Welfare*, 19(2):455-471, 2002.

[59] J.E. Roemer. Equality of opportunity. Harvard University Press, Cambridge, MA, 2008.

[60] H. Rue, S. Martino and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society*, 71:319-392, 2009.

[61] A. Sen. Health: perception versus observation. *BMJ*, 324(7342):860-861, 2002.

[62] T. Sobotka, V. Skirbekk and D. Philipov. Economic Recession and Fertility in the Developed World. *Population and Development Review*, 37(2):267-306, 2011.

[63] W. Thomson. Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: a survey. *Mathematical Social Sciences* 45(3):249-297, 2003.

[64] J. Vallin and F. Meslé. Convergences and divergences in mortality. *Demographic Research Special*, 2:11-44, 2004.

[65] J.W. Vaupel, Z. Zhang, A. van Raalte. Life expectancy and disparity: an international comparison of life table data. *BMJ Open 1*: e000128, 2011.

[66] A. Wagstaff, P. Paci and E. van Doorslaer. On the measurement of inequalities in health. *Social Science and Medicine*, 33(5):545-557, 1991.

[67] A. Wagstaff and E. van Doorslaer. On decomposing the causes of health sector inequalities with an application to malnutrition inequalities in Vietnam. *Journal of Econometrics*, 112(1):207-223, 2003.

[68] A. Williams, R. Cookson. Equity in Health. In: A. Culyer, J.P. Newhouse, editors. *Handbook of Health Economics*, edition 1, volume 1, chapter 35, pages 1863-1910. Elsevier, 2000.

[69] J.R. Wilmoth. Is the pace of Japanese mortality decline converging toward international trend?. *Population and Development Review*, 24(3):593-600.

[70] C. Wilson. On the scale of global demographic convergence 1950-2000. *Population and Development Review*, 27(1):155-171, 2001.