CSAE Working Paper WPS/2020-05

Divide to Conquer? Latent Preference Types and Country-level Heterogeneity

Yannick V. Markhof *

February 29, 2020

Abstract

Do country-level outcomes relate to latent preference types within the population? This paper uses six different preference measures obtained from a global sample of 80,000 individuals to construct a set of unique preference types. It then asks whether the prevalence of these types correlates with outcomes in the domains of income, entrepreneurship, conflict, democratization, and gender equality. To test this hypothesis, it draws on and contributes to three distinct streams of literature: Firstly, it draws on the psychological literature on personality traits and types and transfers this concept to economic preferences to explore the possibility of a preference typology. In a second step, it expands on previous work on global preference heterogeneity by adding a person-centered perspective. This accounts for complex relationships between different preference measures, lets them cluster into characteristic types and allows us to determine individuals' particular type membership. Exploiting a recently published, comprehensive dataset, the paper identifies four robust "pure" preference types. Furthermore, it discovers considerable global variation in their prevalence. To uncover this latent grouping structure of the data, the paper applies latent profile analysis, a type of mixture model related to unsupervised machine learning, and demonstrates its use for economic research. Therefore, this paper lastly makes a contribution on a methodological level. It provides a framework for a rigorous approach to typologies and highlights its applicability to a multitude of questions economists face. In this regard, it outlines how employing a similar approach to divide a population into latent groups can present a stepping-stone for new insights into development research.

*Email: yannick.markhof@gmail.com.
An earlier version of this paper was submitted as part of the MSc in Economics for Development at the University of Oxford and received the Luca D'Agliano Dissertation Prize. I am grateful for supervision and support by Sanjay Jain and helpful comments by Michael Koelle.

1 Introduction

The time is right to forge stronger links between personality psychology and economics.

– Ferguson, Heckman and Corr (2011: 207)

Are we frugal today so we can have more tomorrow? Do we take a risk to earn a high reward? Do we trust others? Return a favor? Care for them? Preferences sit at the heart of economic decision theory. And recent evidence provided by Falk et al. (2018) suggests that we are all different. We are different through our own set of distinct preference measures that govern our behavior. But are different preferences also associated with different development outcomes across countries?

Until recently, little was known about the exact interrelations between the various measures of preferences. This paper, investigates whether six distinct preference measures combine to create unique preference types in the population. Motivated by recent evidence suggesting a relationship between preferences and country-level outcomes, it sets out to determine the distinct preference type profiles of 76 countries and their correlation with development-relevant outcomes.

With the publication of the Global Preference Survey (GPS) and the associated paper analyzing preference heterogeneity, Falk et al. (2018) have laid the groundwork for an exploration of global variation in preferences. In a dataset containing representative samples of 80,000 individuals from 76 countries in total, they find considerable heterogeneity along six different preference measures, both between and within countries.[1] Furthermore, they find that differences in preference profiles between countries are associated with differences in outcomes in the domains of income, entrepreneurship, and conflict.

In analyzing these correlations, Falk et al. (2018) take what the psychological literature calls a "variable-centered" approach (Meeusen et al. 2018). Each of the six preference measures enters their model for itself and independently of the other measures. This allows them to measure the effect of each separately; however, it limits their ability to detect complex interactions between them.[2] The simplified correlation structures Falk et al. (2018) can detect might thus hide a more fine-grained picture. For example, such complex interactions have been found to play a crucial role for the relationship of personality traits and economic outcomes (Fréchette et

[1] These measures are: patience, risk-taking, positive and negative reciprocity, altruism, and trust.
[2] In fact, anything beyond bivariate interactions (certainly though if we go past three variables) presents a problem for interpretation.

al. 2017).

Economic development research thus faces the following challenge: A number of different preferences are crucial to decision making (e.g. Dohmen et al. 2009; Alan and Ertac 2015) and have been linked to development-relevant outcomes (Dohmen et al. 2018; Falk et al. 2018; Kremer, Rao and Schilbach 2019). However, considering the different preference measures in isolation might represent a substantial simplification: In contrast to the implicit conjecture of Falk et al. (2018), the global population might actually be heterogeneous in the way the different preference measures correlate amongst each other and with economic outcomes. Following Merz and Roesch (2011) who relate personality *types* to economic outcomes and a recent application explaining firm performance through different *types* of CEOs (Bandiera et al., 2017, forthcoming, JPE), we might thus want to think of preference heterogeneity in terms of entire preference *types* made up of unique combinations of the six distinct preference measures from Falk et al. (2018). This way, we can let preference measures relate to each other and economic outcomes in different ways across types and account for heterogenous effects in the global population. Such an approach that can account for complex interactions and gives us a holistic picture of a person's preference profile is called "person-centered" (Daljeet et al. 2017).

A person-centered approach shifts the focus from individual variables to the pattern they create together. It acknowledges that preference measures might not exist in isolation and combine in different ways between individuals. When Falk et al. (2018: 1647) thus talk about "distinct preference profiles", they mean several discrete preference measures. However, a person-centered approach does not take these measures as separate parts but as characteristics that jointly create a preference profile together. This has the benefit that we can detect recurring patterns in which the six preference measures combine. These latent subgroups made up of individuals with a similar pattern of preferences is what we will call "pure preference types".

Why are these types called "pure"? We extract them from a heterogenous, global sample. As such, they represent only the first layer of preference types, broad generalizations that approximate the finer-grained true type of a person. Identifying pure preference types thus gives us a global benchmark, a first approximation of what preference types might exist. From this starting point, we can explore the possibility of more nuanced underlying structures, "idiosyncratic" deviations from the pure types that are characteristic for a region or a certain income-level.

As the idea of a preference typology is new to economics, it is natural to turn to a discipline that has explored a similar concept for years: the psychological literature on personality traits and types (McCrae and John 1992; Merz and Roesch 2011;

Ferguson and Hull 2018). Most popularly, the so-called "five-factor model" outlines personality along five distinct traits. These traits represent the overarching structure to finer-grained subtraits and together span (most of) the domain of an individual's personality (Costa and McCrae 1992). Furthermore, they have been linked to a wide range of economic and social outcomes (e.g. Ozer and Benet-Martínez 2006; Mueller and Plug 2006; Soto 2019). How exactly these five traits might combine to create entire personality types has long been a contested topic (Gerlach et al. 2018). However, using the same unsupervised machine learning algorithm we will employ later, Merz and Roesch (2011), Ferguson and Hull (2018), and Gerlach et al. (2018) all recently identified robust personality types made up of the five traits.

The increasing popularity of machine learning has also induced new person-centered research in economics. Bandiera et al. (2017) establish a concept similar to ours for CEO behavior. They use unsupervised machine learning to identify two pure CEO behaviors made up of characteristic combinations of CEO activities. To transfer these concepts to economic preferences and explore the possibility of an equivalent preference typology, we need an econometric method that can detect recurring patterns of preference measures, gives us thorough guidance in identifying their number and quality, and is flexible enough to be applied to a wide range of other questions relevant for economists. Latent profile analysis (LPA) represents such a method and will be the tool of our choice in this paper (Daljeet et al. 2017).

LPA is a type of mixture model that is concerned with the identification of homogeneous clusters in a heterogenous population. It assumes the existence of an unobserved, "latent" variable that groups observations into distinct classes based on a number of indicators. In our case, it allows for the identification of distinct preference types (the classes) made up of combinations of the six preference measures (the indicators).

Employing LPA, I find four robust pure preference types with regional and development-dependent idiosyncrasies on a subordinate level. The first type is characterized by below average values across preference measures. Since this suggests a low level of engagement with others, I label this type *lone wolf*. The lone wolves represent the largest type globally (35%) and are also the most prevalent type in the largest number of countries (46 out of 76). They thus serve as our benchmark when analyzing correlations between the share of each type in a country and outcomes in the domains of income, entrepreneurship, conflict, gender equality, and democratization. The second type, which I label *retaliative*, constitutes the smallest share of the global population (14%). It is characterized by high levels of negative reciprocity along with low values of patience, positive reciprocity, and altruism and is associated with a higher prevalence of armed conflicts on a country-level. The third type I find is called *pro-social* for distinctly higher levels of positive reciprocity and altruism. It makes up 27% of the global population and constitutes a larger share in countries that are neither rich nor poor. Furthermore, higher shares of this type correlate with worse institutional quality. Lastly, I identify a *patient* type (24% globally) that is markedly more prevalent in high-income countries (36%). It is also this type that is most consistently associated with higher levels of development.

Taken together, this paper makes a contribution in three distinct ways. On a conceptual side, it transfers the psychological concept of personality traits and types to economic preferences. Therefore, it lays out a framework that opens up the possibility of a preference typology. Secondly, it expands on research on global preference heterogeneity conducted by Falk et al. (2018) and adds a person-centered perspective to it. By controlling for entire preference types, it accounts for complex interactions between preference measures and verifies that their actual prevalence on the intensive margin matters on top of simple country-averages of individual measures. Lastly, this paper also makes a contribution on a methodological level. By providing a detailed discussion of LPA and the associated empirical strategy, it extends the opportunities for a rigorous approach to typologies in economics. Furthermore, this paper outlines how its approach can be adapted to answer a number of questions relevant for (development) economists and thus pave the way for new insights into development research.

The remainder of this paper proceeds as follows. Section 2 describes the data obtained through the GPS. Section 3 relates my approach to the existing literature and draws on the personality psychology literature to develop a framework of pure preference types. Section 4 introduces LPA, discusses its distinct advantage for our means and describes my empirical strategy. Section 5 investigates global preference type heterogeneity, their relationship with country-level outcomes, and verifies the robustness of my results. Section 6 concludes.

2 Data

In 2012, a group of researchers paired up with the private surveying institute Gallup to conduct the Global Preference Survey (GPS) in order to shed light on the distribution of preferences across the globe. In a recent paper, Falk et al. (2018) now publish first results of this survey along with the dataset.

The GPS features data on six different measures of preferences for more than 80,000 individuals from 76 countries. The countries were selected to constitute a representative selection of the global population and chosen to cover all geographic regions and development levels.[3] Furthermore, they provide the cultural, spatial, linguistic, historical, political and ecological variation desirable for our approach (Falk and Hermle 2018: appendix). The median sample size taken from each of the countries is 1,000 and, together with sampling weights they obtain from Gallup, is representative of the country's population.

For each individual, the GPS contains a standardized[4] score on patience, willingness to take risks, positive and negative reciprocity, trust, and altruism.[5] These were elicited in a standardized procedure and through an experimentally validated procedure detailed in Falk et al. (2016).

In Section 4, we link the prevalence of the different pure types to country-level outcomes. Outcomes were chosen to represent a broad set of indicators of economic development popularly used in the literature. Specifically, we investigate correlations with outcomes in the domains of per capita income, entrepreneurial activity, conflict, institutional quality and gender equality.

The next section discusses the conceptual framework for a preference typology in economics.

3 Conceptual Framework

This section is concerned with the motivation for a preference typology. It outlines how our research question fits into the economic literature on preferences and grouping exercises and highlights how the psychological literature on personality traits and types can help us to conceptualize a model that takes these findings to the economic (development) sphere.

Identifying latent preference types directly expands on the economic literature on preference measurement and their significance for economic outcomes (see for example Becker et al. 2012; Falk and Hermle 2018). Most notably, Falk et al. (2018) provide the basis for our research in that they find distinct country-level preference profiles that relate to country-level economic, social and political outcomes. Additionally, they find evidence of correlations between different preference measures

[3] In total, it constitutes 22 countries from Asia and the Pacific, 25 European countries, 15 from the Americas and 14 African countries of which 11 are sub-Saharan. This selection spans about 90% of the world population and income (Falk et al. 2018: 1651).

[4] i.e. with zero mean and a standard deviation of one.

[5] Patience represents time preferences and positive and negative reciprocity the propensity to return good actions toward one equivalently and seek revenge for bad actions directed against one, respectively.

and of these preferences with individual-level observables (such as gender or age) along with substantial geographic variation in preferences. Therefore, they conclude that the distribution of preferences is marked by substantive across-country as well as within-country variation.

The economic analysis of a comprehensive set of measures making up an individual's preference profile is still in its nascent stages though. Recently, this literature has received an important impetus through the publication of the GPS and the results of Falk et al. (2018). To advance on this topic, however, we can benefit from a similar stream of research in psychology on personality measurement that has been around for much longer. This literature can provide us with a broad structure that we can transfer to the economic sphere.

The most popular model of personality is the "five-factor model" (FFM) or "big five" (Costa and McCrae 1992). The FFM maps out personality along five stable traits across contexts: *openness to experience*, *conscientiousness*, *extraversion*, *agreeableness*, and *neuroticism*.[6] While these five traits have reached broad consensus in the literature, they are not meant to constitute an all-encompassing representation of personality (Ashton et al. 2004). In a similar way, studies have amended the FFM on a subsidiary level to capture a more nuanced picture of personality. Costa and McCrae (1992) augment the FFM by proposing six subcategories ("facets") underlying each of the five personality factors.[7] DeYoung et al. (2007) further suggest a number of aspects at the intermediate level between the big five and their facets.

A number of studies have found personality traits to correlate with a broad range of economic, social and political outcomes (e.g. Ozer and Benet-Martínez 2006; Mueller and Plug 2006; Soto 2019). Likewise, the classification of individuals into distinct types based on the big five traits has been on psychologists' agendas for a while (see e.g. Hofstee et al. 1992). However, it was not until recently that latent profile analysis allowed Gerlach et al. (2018) to identify four stable personality *types*.

To recap, we have seen how the psychological literature models personality along five distinct traits that represent the overarching level of finer, underlying structures. These traits are related to economic outcomes and can be combined to form entire personality types. How do these concepts transfer to economic preferences?

---

[6] Each of the five factors is measured on a continuum determined by a positive and negative pole. Openness reflects curiosity and creativity, conscientiousness reliability and goal-orientedness, extraversion reflects gregariousness and emotional expressiveness, agreeableness compassion and pro-sociality, and neuroticism emotional instability.

[7] For example, the facets underlying agreeableness are trust, altruism, straightforwardness, compliance, tendermindedness and modesty.

A comprehensive body of more recent literature examines the relationship between personality traits and preference measures. Both concepts are related in that they describe (relatively stable) intrinsic attributes of an individual that play an important role for behavior and decision making across a range of different contexts (Golsteyn and Schildberg-Hörisch 2017). However, they are markedly different in the way they are motivated and surveyed: Preferences are parameters of a utility function that agents typically seek to maximize in economics and usually inferred from the actions taken in experiments (ibid., Becker et al. 2012). On the other hand, personality traits are lexically derived and elicited through self-reported questionnaires (Golsteyn and Schildberg-Hörisch 2017). At the same time, these differences also mean that both concepts can complement each other (Ferguson, Heckman and Corr, 2011). Rustichini et al. (2016), for example, take an integrative approach to explaining economic behavior drawing on decision theory that is rooted in economics and personality theory with its origin in psychology. Becker et al. (2012) specifically examine the relationship of the same six preference measures included in the GPS to the big five. They conclude that while they find a number of correlations, the two concepts are largely complementary.[8] For us, this means that while the literature on personality can provide us with some guidance in identifying preference types, its benefits lie especially in the methodological and conceptual sphere.

In identifying such a latent grouping structure of preferences, our work resembles other studies conducting taxonomies and grouping exercises. Taxonomies have been a common exercise in the psychological (personality) literature for many years with latent variable models being a popular tool (e.g. Ferguson and Hull 2018; Merz and Roesch 2011). Through recent methodological advances and the advent of (unsupervised) machine learning, economists are now increasingly tapping into their potential. Mullainathan and Spiess (2017) highlight how these advances give economists new tools to answer old questions but also open up new possibilities for research. Su, Shi and Phillips (2016: 2216) motivate the "universally relevant" challenge to identify latent group structures in the data to account for heterogeneity across different subgroups in the population.

Such grouping exercises apply to a broad range of questions relevant to (development) economists: Vollmer et al. (2003) use a latent variable model to cluster countries into different "human development clubs". Others employ them for a multidimensional approach to measuring health (Landale et al. 2013) or poverty (Moisio

[8] For example, the strongest correlation they find is between the "social preferences" of altruism, trust and negative as well as positive reciprocity and the big five trait of agreeableness.

2004; DeWilde 2004), to cluster households according to their socioeconomic status (McParland et al. 2014), identify different types of family firms (Stanley et al. 2017) or model different types of risk factors for criminal offenders (Francis and Liu 2015).

Motivated by 1) evident global preference heterogeneity, 2) the psychological framework of the big five and its subtraits and 3) recent methodological advances opening up new, promising opportunities for economists employing grouping exercises, this paper sets out to identify "pure preference types" in a global sample.

The idea of latent "pure types" (or "ideal types") dates all the way back to Max Weber and underlies an extensive body of theory and research in the social sciences, including economics (Weber 1949). Along these lines, we propose a framework of pure preference types made up of distinct combinations of our six preference measures. Our approach is most closely related to Gerlach et al. (2018) who conduct a personality typology based on the big five and a very recent application of Bandiera et al. (2017) who identify pure types of CEOs and link them to real firm outcomes. To our knowledge, the unprecedented comprehensiveness of the GPS allows us to be the first ones to transfer this concept to economic preferences. Notwithstanding, our framework resembles the FFM conceptually in that it models a broader structure of types (traits) that further split up into ever finer subtypes (facets) (DeYoung et al. 2016). An important point to note here is that we make no claim on the real existence of these pure types, nor do we expect a perfect fit of individuals to any single one of them. They rather provide us with the degree of abstraction necessary to analyze latent preference types in a large sample of heterogeneous individuals and an approximation of the finer-grained (regional) idiosyncrasies.

The process of identifying the correct number of pure types to extract is described in the next section.

4 Methodology

This section introduces latent profile analysis and motivates its suitability for our research question. We highlight how it is employed in our specific context and propose a new strategy to study global preference heterogeneity that can be transferred to a wide range of other economic problems.

4.1 An introduction to latent profile analysis

Mixture models play an important role in several strands of the econometric literature (Henry et al. 2014: 123). As one subtype of mixture models, latent variable analysis is concerned with the identification of homogenous clusters in a heterogenous population, which makes it a vital tool across the social sciences (McCutcheon 1987). The assumption behind this approach is that there is a categorical latent (i.e. underlying/unobserved) variable that groups observations in the data into distinct classes.[9] It does so based on a number of characteristics of the individuals and identifies subgroups that share a similar vector of characteristics. If these characteristics are continuous, we are conducting a "latent profile analysis" (LPA). In our case, we have data on six preference measures of each individual which LPA uses to group them into different latent preference types. In doing so, we estimate two sets of parameters: a number of measurement parameters (the means, variances and covariances of the indicators in each class) and a structural parameter giving the proportions (i.e. sizes) of each class. We will come back to this in Section 4.3.

One of the main findings of Falk et al. (2018) are correlations between the preference measures. However, their approach to analyzing preference heterogeneity and its association with country-level outcomes is a *variable*-centered one. As such, they model preference heterogeneity under the implicit assumption that relationships among preferences and their correlations with the outcomes are homogenous in the global population (Daljeet et al. 2017).

However, given the complex relationships Fréchette et al. (2017) or Ferguson and Hull (2018) find among different personality traits, it is conceivable that this approach masks more nuanced relations between preference measures. Namely, while variable-centered approaches can account for some heterogeneity in effects through interaction terms, these can quickly get extremely convoluted, impinge on statistical power, and are hard to interpret (Merz and Roesch 2011). They are thus typically limited to bivariate interactions. We therefore opt for a *person*-centered approach that can account for different relationships across different subgroups in the data (Pastor et al. 2007; Daljeet et al. 2017).[10] As a result, we might be able to identify distinct patterns and add a more comprehensive picture of preference heterogeneity to Falk et al.'s results.[11]

While being conceptually similar to (k-means) cluster analysis (Landale et al. 2013), LPA has a number of distinct advantages. It features a model-based, stochas-

[9] I use the terms "class", "profile" or "type" interchangeably when talking about the different clusters identified.

[10] However, its application is by no means limited to individuals. As we have seen already, "person" might as well refer to households, firms or entire countries.

[11] An important point to stress, though, is that person- and variable-centered approaches are not rival approaches but rather complementary. As such they are often combined for thorough investigation of the matter (Masyn 2013).

tic approach thus modelling the relationship between the latent variable and its indicators as probabilistic. Hence, it does not rely on some (arbitrarily chosen) distance measure and can account for uncertainty in class assignment (Moisio 2004). Among other improvements, it also allows us to select our model specification based on objective goodness-of-fit and information criteria and thus presents a more rigorous approach to a person-centered analysis (Stanley et al. 2017). Among the different person-centered methodologies, LPA is therefore "arguably the most flexible" and adequate for "the widest array of research questions" (Meyer and Morin 2016: 597).

4.2 Model selection

Su, Shi and Phillips (2016) identify model selection and the assignment of individuals to classes (or pure preference types in our case) as the two major challenges of person-centered approaches. Concerning model selection, we have two simultaneous decisions to make. First, we have to identify the correct number of pure preference types to extract from the data. The second decision is concerned with the within-class variance-covariance structure of the indicators and its relation across classes. Specifically, the flexibility of LPA adds an additional layer of complexity to our model selection in that it requires us to specify a covariance matrix (Masyn 2013). This is an issue often overlooked or neglected in many practical approaches and entails great implications for the estimated parameters (Masyn 2013, Pastor et al. 2007).[12] In order to present a rigorous approach to LPA, we therefore consider the matter in detail next.[13]

The covariance matrix includes the variances of the indicators along the main diagonal for each class as well as their covariances. For the variances, we can either set them to be equal across classes or varying (i.e. class-specific). The covariances can either be constrained to be zero, equal across classes or freely estimated for each class individually. This gives rise to five different specifications of the covariance matrix, $\mathbf{\Sigma}_k$, which I label A to E from the most parsimonious model constraining variances to be equal and covariances to be zero to the most complex model that allows for variances and covariances to be freely estimated across classes k.[14] Below, I show an example of the covariance matrix for the most complex model (i.e. model

[12] For example, the covariance matrix is often specified without much consideration either for simplicity (Marsh et al. 2009) or convenience (Geiser et al. 2014).

[13] The notation in this and the next subsection will follow Pastor et al. (2007).

[14] Specifically, the options are A (equal variances, covariances fixed to 0), B (var=equal, cov=equal), C (var=varying, cov=0), D (var=varying, cov=equal), and E (var=varying, cov=varying).

E) and r different indicators.

$$\Sigma_k = \begin{bmatrix} \sigma_{1k}^2 & & & \\ \sigma_{21k} & \sigma_{2k}^2 & & \\ \vdots & \vdots & \ddots & \\ \sigma_{r1k} & \sigma_{r2k} & \cdots & \sigma_{rk}^2 \end{bmatrix} \tag{1}$$

As would be expected for an issue that is so complex yet so central to mixture modelling, there is an extensive body of literature on how to pick between different numbers of classes and specifications of the covariance matrix. The first lead is obtained by looking at several information criteria. Although there is a large number of possible criteria here, the ones most commonly used in the literature are the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC) and the adjusted BIC (aBIC) (Nylund et al. 2007). These measures have in common that they are based on the (maximum) value of the likelihood function while taking the sample size into consideration differently and applying different penalties for the number of estimated parameters. As such, they provide us with a score that can be compared across different models with lower values representing better model fit.[15] When comparing whether a model with one more class fits significantly better than one with one less class but the same $\mathbf{\Sigma}_k$ specification, we can consult statistical tests of which the Vuong-Lo-Mendell-Rubin likelihood ratio test (VMLR-LRT) and the bootstrap likelihood ratio test (BLRT) are the most common. Simulation studies by Nylund et al. (2007) have found the BIC and the BLRT to consistently outperform the other measures which is why we will rely on them for model selection.

A second way to judge model fit is by the average accuracy with which we can assign individuals to their most likely class (Masyn 2013, Pastor et al. 2007). Average highest posterior probabilities of 0.7 or higher indicate good separation between classes and high classification accuracy (Nagin 2005).

Since none of the above criteria can give a definitive answer to our model selection problem and they sometimes do not coincide, the literature lastly also stresses the importance of assessing the interpretability and meaningfulness of the different solutions when selecting a model (Ram and Grimm 2009, Hu and Bentler 1998). For example, a class size containing less than 5% of the population raises questions about the relevance of the extracted profile and the same goes for several non-differentiated classes (Hipp and Bauer 2006, Pastor et al. 2007). We will come back to this in

[15] Their actual value has no meaning though, so they are only useful when comparing two or more models.

more detail in Section 5.1.

4.3 Estimation

In the next subsection, we go through our estimation strategy step-by-step and discuss common pitfalls as we go along.

To estimate our model using a maximum likelihood approach, we use the specialist software *Mplus* that is tailored to latent variable modelling.[16] To deal with missing values in the data, we use full information maximum likelihood (FIML). This has the advantage that we can make full use of all information available to us (Masyn 2013). Because there is no closed-form solution to the ML estimates, we rely on the Expectation Maximization (EM) Algorithm, an iterative approach to finding the maximum likelihood (Muthén and Shedden 1999). A common issue with the EM algorithm is failure to converge to a global maximum since the log-likelihood in our case will likely have multiple (local) maxima. Because it is crucial for a reliable solution to make sure that we arrive at a global maximum, we randomize over the starting values for the EM algorithm. Replicating the maximum likelihood value for different sets of starting values should make us confident that we have arrived at a global maximum. In line with the recommendation of Muthén and Muthén (2017) I use 100 random sets of start values and 20 final stage optimizations by default. In the case of convergence issues, I increase this number to up to 1000 random start values and 250 final stage optimizations which is double what is recommended for a thorough investigation of such issues.

We then estimate the following LPA model with K pure preference types where $\mathbf{y}_i$ is the vector containing our six preference measures for individual i and $\boldsymbol{\phi}$ is a vector of the structural and measurement model parameters we estimate.

$$f(\mathbf{y}_i|\boldsymbol{\phi}) = \sum_{k=1}^{K} \alpha_k f_k(\mathbf{y}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \tag{2}$$

This shows that the multivariate distribution of the six preference measures is a mixture of K pure type specific multivariate density functions[17] with mean vector μ_k and covariance matrix $\boldsymbol{\Sigma}_k$, each weighted by the proportion of individuals belonging to the respective type α_k.

[16] Many of the functions are also available by combining different packages in *R* and *Stata 15* now also provides some basic coverage of finite mixture models for the first time.
[17] We usually assume these to be multivariate normal (Masyn 2013).

Once we have obtained the structural and measurement parameters described in Section 4.1, we can calculate *posterior probabilities* that give us the estimated membership probability of each individual for each pure preference type:

$$\boldsymbol{\pi}_{k|\mathbf{y}_i} = \frac{\boldsymbol{\alpha}_k f_k(\mathbf{y}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{k=1}^{K} \boldsymbol{\alpha}_k f_k(\mathbf{y}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)} \tag{3}$$

In the next subsection, we will see how we can use these posterior probabilities to our advantage by accounting for imperfect fit of individuals to the pure types.

4.4 Grouping individuals

We have already touched upon a strength of LPA in comparison to other clustering approaches, which is the ability to calculate posterior probabilities to account for classification uncertainty. Nonetheless, an important question remains: How do we assign individuals to pure types?

In many instances, modal assignment based on the posterior probabilities is used (i.e. we assign an individual to the type it has the highest probability of being a member of). On second thought, however, this approach has a clear shortcoming: Not only do we discard the information given by the posterior probability in further analyses by deterministically assigning individuals to classes, we also treat this membership as known from that point on for each individual. A quick example lets us clearly see the inaccuracy in doing so: Assume a model with two pure types. One individual's highest posterior probability is 0.90 and for pure type 1, representing a very confident assignment and good fit of type 1 for this individual. Now consider another person that also has her highest probability for type 1 but this time only with probability 0.60. With modal assignment, we would assign both individuals to type 1 and not differentiate any further between them from this point on although the first individual clearly represents type 1 better than the second. This problem is known as "classify-analyze" in the literature and leads to biased results when using class membership for further analysis (Bray et al. 2015).

Our approach circumvents this problem and exploits the posterior probabilities for subsequent analysis. Similar to Pastor et al. (2007), we calculate the size of each pure type in the sample by adding up posterior probabilities. This accounts for imperfect fit of individuals to the different types as would be expected from the model of *pure* types proposed in Section 3. Moving to the country level, we can calculate the share of each pure type in each country in the same way. However, in order to claim that these shares are an adequate representation of each pure type's prevalence in a country, we need to make sure that the sample we have from each country is representative of its whole population. Therefore, we calculate weight-adjusted shares of each type for each country with individual weights obtained from Falk et al. (2018). By calculating the prevalence of each pure type in each country like this, we circumvent problems arising when deterministically assigning individuals to types. Furthermore, this allows us to analyze both the between- and within-country heterogeneity in preferences Falk et al. (2018) identify.

4.5 Country-level outcomes

In a last step, we look at whether the prevalence of a preference type correlates with country-level outcomes. We should be very clear from the start that we are not claiming a causal relationship here (nor can we). In the spirit of Falk et al. (2018) the goal here is thus not to maximize explanatory power but to provide a first exploration of possible correlations that will need to be further analyzed in future research.

Our approach expands on Falk et al.'s (2018) work by hypothesizing that it is not the simple *mean* of *individual* preferences relating to country outcomes, but the actual *prevalence* of preference *types* made up of *several* preference measures. As such, our approach here is somewhat of a hybrid between variable- and person-centered analysis as suggested by Masyn (2013) for example. Equation (4) formalizes our regression model

$$y_{jc} = \gamma + \boldsymbol{\tau}'_{\neg s}\boldsymbol{\beta} + \mathbf{X}'\boldsymbol{\delta} + \varepsilon_{jc} \tag{4}$$

where y is outcome j in country c, $\boldsymbol{\tau}_{\neg s}$ is the vector containing the shares of the different pure preference types for the country leaving share s out because shares add to 1 in a country (making them perfectly collinear) and $\mathbf{X}$ is a vector of country controls.

To sum up, we have seen how a person-centered approach can bring new value to the analysis of (global) preference heterogeneity and that LPA is the tool of choice here. We have highlighted several challenges that come with it and how our methodology resolves them. Lastly, we have outlined an identification strategy that is consistent with our theoretical framework of Section 3 in that it takes it takes into account classification uncertainty and how this approach constitutes an interesting addition to Falk et al.'s. The next section will take the framework to the data and use the methodology outlined in this chapter to identify pure preference types and country-level heterogeneity in a global sample.

Table 1: Model Fit Across Specifications

Model	#classes	BIC	BLRT	Posterior Probabilities Type 1	Type 2	Type 3	Type 4	Type 5
A	5	1293520	0.00	0.74	0.78	0.80	0.90	0.78
B	4	1283533	0.00	0.81	0.77	0.92	0.92	
C	4	1289380	0.00	0.81	0.81	0.81	0.80	
D	4	1275744	0.00	0.64	0.76	0.75	0.94	
E	4	1270503	0.00	0.69	0.78	0.76	0.94	

5 Results

This section presents the results of a latent profile analysis of the Global Preference Survey following the approach outlined in the previous section. I first discuss the results of the model selection procedure and introduce the pure preference types identified. I then analyze their global dispersion and correlation with country-level outcomes and verify the robustness of my results.

5.1 Preference types

As outlined in Section 4, we begin by fitting five different models to the data, increasing the number of classes step-by-step and carefully taking note of the various fit indices for each model specification.[18] [19]

As a first observation, our indicators in Table 1 consistently favor more complex models and a higher number of extracted classes.[20] This does not come as a surprise considering the large size of our sample: The typical "large" sample size for papers assessing the performance of the different fit indices lies between 1000 (Nylund et al. 2007) and 2000 (Tofighi and Enders 2008), a size our dataset exceeds by orders of magnitude. Consequently, Meyer and Morin (2016) point out that this might result in the fit indices suggesting additional classes without reaching a recommended solution in very large samples (see also Masyn 2013). This often leads to a well-known trade-off between fit and interpretability in the literature (Mullainathan and Spiess 2017; Chang et al. 2009; Marsh et al. 2009). For example, Gerlach et al. (2018) find 13 "optimal" classes in their study, however, they argue that this is

[18] As introduced in Section 4, I denote these model A to E depending on their covariance matrix with a number indicating the number of extracted classes following, e.g. model E4.

[19] For reasons of limited space, I only include a slimmed-down table containing just the best fitting model for each specification and the most meaningful fit indices as described in Section 4 here.

[20] As explained in Section 4.2, lower BIC values generally represent better fit and significant p-values of the BLRT a better fit than the same model with one less extracted class.

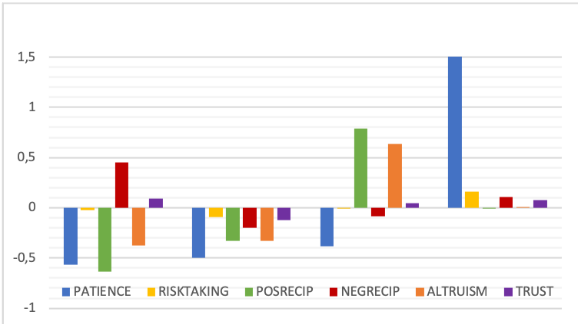

(A) Extracted Types E4

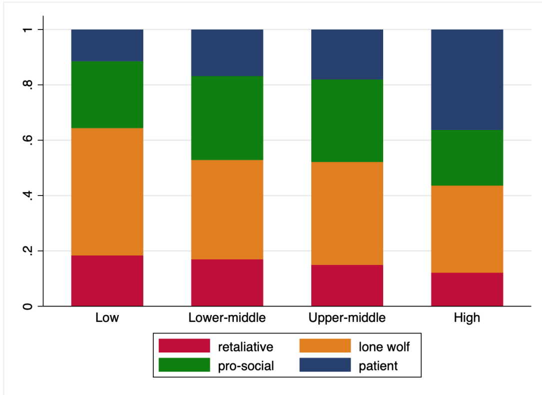

(B) Type by Income Level

FIGURE 1: Figure 1 shows the types we find for model E4 in Panel (A) and their prevalence by income level in Panel (B).

due to a number of spurious classes adding little value while severely affecting the interpretability of the solution and settle for a 4-class model. Similarly, Bandiera et al. (2017) defend their choice of an easily-interpretable model of two pure CEO types compared to the 11-class solution fit-indices suggest in their case.

Consistent with our framework, we need to keep in mind that we are aiming to extract a small number of archetypes from a very large and heterogeneous sample and that we leave room for deviations from these pure types on a subordinate level. Fortunately, however, going by the methodology outlined in Section 4 does indeed point to such a preferred solution: By ruling out classes with fewer than 5% of total observations and stopping when the model is not well identified anymore[21], we arrive at 5 suggested classes for model A and 4 optimal classes for models B to E. Additionally, the BIC, AIC and aBIC all suggest that the best fit is for model E. Furthermore, we can verify that model E4 fits significantly better than model E3 through the BLRT and that it fits better than the second-best model (D4) by means of a scaling-corrected Satorra-Bentler chi-square difference test (Satorra and Bentler, 2010). Inspecting the average highest posterior probability for each class, lastly, suggests good fit and separation of the classes. However, it should be noted that one class narrowly misses out on the desirable threshold of 0.7 with a value of 0.69. Taken together, the data strongly favors E4 as the model of choice for our subsequent analysis.

Figure 1.A plots the four resulting profiles we find with the corresponding in-class means for each preference measure. In labelling the types, we will follow

[21] As expected, the EM algorithm encounters convergence issues for more complex models past a 4-class solution (see also Masyn (2013) and Tofighi and Enders (2008).

the suggested approach by Meyer and Morin (2016) and identify the characteristic preference measure(s) for each of them. The first pure type we can identify has notably higher levels of negative reciprocity along with low levels of patience, positive reciprocity, and altruism. I therefore label this type "retaliative". It is also the smallest type claiming a share of about 14% of the sample. The second type is characterized by below average values across all indicators and accounts for the largest share in the population (35%). I label this type "lone wolf" for its seemingly low engagement with others. High levels of positive reciprocity and altruism stand out for type three (27%) which is why I label this type "pro-social". Lastly, I name type 4 "patient" for its very noticeable spike in the corresponding measure.[22]

In general, the types we identify reflect the correlations between individual measures at the country-level Falk et al. (2018) report in Table IV, most notably between pro-social traits. However, an interesting observation concerning "retaliative" types is that, as opposed to the raw correlation Falk et al. (2018) find at the country level, negative reciprocity and patience are *negatively* correlated, providing a good example of a more complex interaction not detectable by a variable-centered approach. Additionally, our analysis does not rely on averaging scores for the different preference measures at the country level and subsequently correlating them but builds from correlations at the individual level.[23] This is because the pure types we extract are based on the global sample in its most disaggregated form, i.e. at the individual level, and hence captures recurring correlation patterns when comparing over 80,000 individuals and not just 76 countries. Consequently, our analysis uncovers a much more nuanced picture than what Falk et al. (2018) (can) find when correlating only pairs of preference measures at the individual level and not leaving room for multivariate and complex correlations (Online Appendix Table 12) or when aggregating at the country level at a similar loss of information (Table IV).

5.2 Global dispersion of preference types

What types are more prevalent in which countries? Figure 1.B depicts the share of the respective types by income level. While no single type clearly dominated the other types globally, the *lone wolf* type takes up 46% in low income countries. Another remarkable pattern is the increase in the *patient* type with income level.

[22] Section 5.4 will discuss the reasons for the apparent greater distinctiveness of patience and positive reciprocity compared to the other measures in detail.

[23] Individual-level correlations are reported in Table 12 of the Online Appendix to Falk et al. (2018).

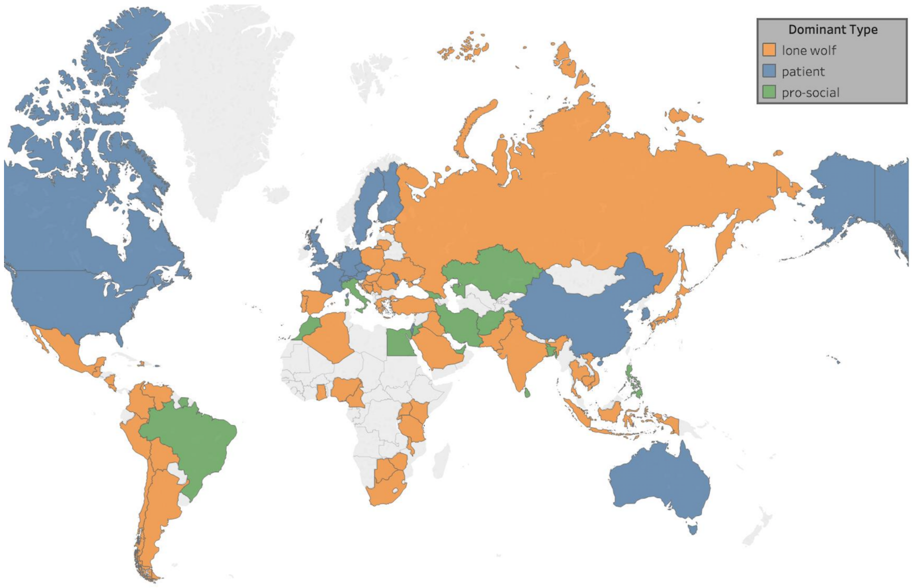

FIGURE 2: Dominant Types by Country

In high income countries, this type takes up more than double of what it does in upper-middle income countries and lies almost 13 percentage points above the global average. The *pro-social* type, on the other hand, seems most prevalent in countries that are neither among the poorest nor richest. While there is little dispersion in the *retaliative* type, its share is gradually decreasing with income level.

Figure 2 dedicates itself to another interesting question: What is the dominant preference type for a country (i.e. the type with the largest share)? Consistent with their shares in the global sample, the *lone wolf* type constitutes the dominant type in 46 out of our 76 countries, while no country has a majority of *retaliative* types. The *patient* and *pro-social* types are the most prevalent in 16 and 14 of the countries, respectively. In order for a preference type to be a country's dominant one, it consistently takes a share between 41% (*lone-wolf*-dominated countries) and 47% (*patient*-dominated) of the population on average.

Since we calculate shares by adding up posterior probabilities, it might matter whether a type amounts to a large share because we add up a lot of small shares (i.e. a lot of individuals but with bad fit to the class respectively) or if a type's prevalence is marked by a few very well-fitting individuals. To get an idea for what countries our pure types provide a better or worse fit, I also calculate the average highest posterior probability for each class for every country separately. However, the results show only small deviations from the global classification accuracy if we split up the countries by income-level or world region suggesting that our model fits equally well across the sample.[24]

Lastly, we might think that it is not the prevalence of a certain type that matters but the degree of within-country preference type fractionalization. To test this hypothesis, I calculate the Herfindahl index that was initially developed to measure market competition (e.g. Amiti and Konings 2007) but has also been used to quantify political competition (Kosec et al. 2018), exporter concentration (Ludema and Mayda 2013) or ethnic, linguistic, and religious fractionalization (Alesina et al. 2003) for each country. Yet, this reveals no evident patterns or significant correlations between the degree of preference type fractionalization and income-levels or any of the outcome measures we will take a closer look at in the following subsection.

5.3 Country-level outcomes

In this subsection, we will explore whether the respective prevalence of our pure preference types correlates with country-level outcomes. In order to facilitate comparability with the variable-centered approach in Falk et al. (2018), we will focus on a similar set of outcome variables the literature has emphasized as potentially related to preferences. Additionally, we add the domains of gender (in)equality recently studied by Falk and Hermle (2018) and a measure of democratization. Taken together, these variables span a wide range of development-relevant outcomes and help us to assess whether preferences are relevant beyond their separate, individual effects. In particular, we analyze whether there is a relationship of entire preference *types* with country-level outcomes and account for the intensive margin, i.e. a type's share of the total country population, as opposed to the simple average value of single preferences in Falk et al. (2018). For this analysis, we need to pick one type as a benchmark as mentioned in Section 4.5. We choose that to be the lone wolf type for two reasons: Firstly, it is representative of the most individuals worldwide and the biggest type in the largest number of countries. Additionally, it does not exhibit any single characteristic spikes of the six preference measures but rather a quite uniform profile. This makes it a good choice for our benchmark.[25]

[24] However, this method has the obvious drawback that it cannot account for a bad fit of all types simultaneously in a country as the posterior probabilities will always add up to 1.

[25] Coefficients can hence be interpreted in the following way: They present the effect of a 10 percentage point increase in the respective preference type's share while simultaneously decreasing the share of the lone wolf type by the same amount. However, as these types are "pure" and therefore only represent a generalization of the actual idiosyncratic types in a country and the mechanisms behind the correlations we find likely complex, we will refrain from giving the coefficients a causal interpretation.

Table 2: Country-Level Outcomes

	log [GDP p/c]		Patent applic. p/c		Armed conflicts		Gender Inequality Index		Polity score	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Patient	0.85***	0.49***	1.28***	0.98**	0.51***	0.28	-0.099***	-0.050***	0.80	-0.32
	(0.20)	(0.15)	(0.36)	(0.38)	(0.17)	(0.18)	(0.021)	(0.018)	(0.78)	(0.71)
Retaliative	0.17	-0.053	0.25	0.53	0.71**	0.73**	-0.027	-0.019	-0.73	-0.75
	(0.38)	(0.29)	(0.67)	(0.64)	(0.35)	(0.34)	(0.046)	(0.034)	(1.67)	(1.29)
Pro-Social	0.24	-0.007	0.71*	0.67	0.26	0.20	-0.028	-0.006	-1.31	-1.75**
	(0.24)	(0.20)	(0.40)	(0.45)	(0.21)	(0.21)	(0.027)	(0.019)	(0.98)	(0.80)
Constant	5.42***	7.28***	1.04	0.52	-1.11	-1.40	0.69***	0.47**	7.85	16.20**
	(1.58)	(1.88)	(2.76)	(3.30)	(1.36)	(2.07)	(0.17)	(0.20)	(6.15)	(7.38)
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	76	73	64	61	76	73	75	72	75	73
Adjusted R^2	0.45	0.67	0.30	0.37	0.11	0.21	0.41	0.65	0.15	0.43

Notes: OLS estimates, robust standard errors in parentheses. Controls include the distance to the Equator, average temperature, average precipitation, the share of the population living in (sub) tropical zones, terrain ruggedness, average distance to the nearest waterway, and an island dummy (s. Falk et al. 2018).*** p<0.01, ** p<0.05, * p<0.1

Consistent with the income-related pattern we saw in the last subsection, the *patient* type seems to be the one most closely associated with desirable outcomes across development indicators. For three out of our five outcome variables, this relationship also remains robust to the inclusion of a number of geographic and climatic controls.[26] As such, the results in Table 2 substantiate the strong relationship of patience with income levels (Falk et al. 2018) and other factors relevant for development such as the accumulation of physical and human capital and productivity (Dohmen et al. 2018). Furthermore, our results suggest that the prevalence of *patient* type individuals in a country also correlates with entrepreneurial activity measured as patent applications per capita and is inversely related to the UNDP Gender Inequality Index (GII). The latter adds a potential relationship between gender equality and differences in preference types to Falk and Hermle (2018) who find the GII and gender specific preferences to be strongly correlated.

Another interesting relationship is that of the retaliative type with the number of armed conflicts obtained from the UCDP/PRIO dataset. In accordance with this type's high value of negative reciprocity, along with little patience and below-average values of the pro-social traits (altruism and positive reciprocity), a larger share of retaliative types is associated with higher conflict prevalence. This confirms the role Falk et al. (2018) find for negative reciprocity, however, it adds the additional twist of low patience associated with the *retaliative* type, a relationship Falk et al. (2018) cannot detect in the raw data.

[26] These are the same used in Falk et al. (2018).

In the institutional domain, an increase in the *pro-social* type relative to the benchmark lone wolf type is negatively associated with an index of democratization obtained from the Polity IV Project. This might point to a "norms-based" explanation of institutional development (Greif 1994; Greif and Kingston 2011): Given sufficient stability of the types over time, a higher share of pro-social types might historically be associated with a collectivist organization of society in which reciprocal relations substitute for "rules-based" institutions.[27]

In sum, I find four pure preference types with varying prevalence across the globe and income-levels. Furthermore, I find that their respective shares in a country significantly relate to country-level outcomes. Mostly, my results confirm the correlations Falk et al. (2018) find for preference measures individually. However, this paper adds that these relationships stay relevant when controlling for entire preference *types* and that their actual prevalence on the intensive margin matters on top of simple country-averages of individual measures.

This is a relevant addition to the findings in Falk et al. (2018) as they only accounted for the effect of each preference measure in isolation. Given the recurring pattern of entire preference types I find and the complexity of human behavior, it would have been conceivable that their results (and the emergence of types) are not driven by a single preference measure, but a characteristic combination of several of them. Along the same lines, a valid concern with the results in Falk et al. (2018) was that the same preference measure could have had very different implications for country-level outcomes depending on the other preference measures with which it occurs. As an example, *retaliative* types, *lone wolves*, and *pro social* types all display low levels of patience. Each comes with a distinct combination of the other preference measures though which creates the respective type and leads to diverging correlations with our outcome indicators for each type. To illustrate the relevance of this further, consider the possibility that we would have found another “high-patience profile? with a different pattern of the remaining outcome indicators. Arguably, the two resulting profiles could have related to development outcomes in very different ways if the complex interactions between preference measures, and not simply patience in isolation, would have driven the correlations we observe.[28]

[27] In another specification (not reported), I also test for heterogenous effects of the different preference types by income level for each outcome measure. The results suggest that the relationships we find in Table 2 are stable across income levels.

[28] For example, Falk et al. (2018) find both patience and risk taking to matter for GDP p/c, a relationship that is dominated by the effect of patience though. However, prior to our analysis of complex interactions between indicators, it was not evident that patience actually “outperforms? trust and risk taking (Falk et al. 2018: 1684). Concretely, the observed correlation they solely

Lastly, my results confirm that it is not just the simple country-average of one or more preference measures that might matter for development-relevant outcomes but also the actual frequency with which they occur within a country vis-à-vis other combinations of preference measures. We consider the robustness of these findings in the next subsection.

5.4 Robustness checks

A first concern might be that latent profile analysis is not a part of the standard toolkit of economists to date. I therefore verify the robustness of the extracted preference types using k-means clustering which has the drawbacks compared to LPA described in Section 4.1 but has been more widely applied in economics so far. Considering the notably different results obtained through similar robustness checks in other papers (e.g. Stanley et al., 2017), the results I obtain are remarkably similar: There is a profile exhibiting a notable spike in patience (the *patient* type), one with all below average values (the *lone wolves*) and one with characteristically high values of altruism and positive reciprocity (the *pro-social* type). Additionally, their respective sizes roughly compare to the ones found in our preferred specification with LPA: the *pro-social* type occupies the same share (26.6%) while the *patient* and *lone wolf* types are slightly less prevalent (18.5% and 26.6%). The last profile obtained through the k-means algorithm corresponds to the *retaliative* type in that it has a distinctly higher mean for negative reciprocity paired with low patience, however, it also comes with above average values in the willingness to take risks and positive reciprocity. Additionally, this profile is much bigger than what we find (28.3% vs. 14.4%). This confirms our previous suspicion that the *retaliative* profile is the least well-identified.

Next, a number of questions might relate to the shape of the extracted profiles directly. Specifically, patience and positive reciprocity stand out in a number of profiles, while trust and risk taking do not display such characteristic deviations from the mean in any of the profiles. This suggests that people with particularly distinct preferences for these measures are somewhat of a "rare breed" and do not share a common pattern that would warrant a pure type of its own.[29] Additionally, we can

assign to patience could have hidden a more nuanced picture where high patience in combination with average levels of trust and risk taking (as observed for our *patient* type) correlates differently with GDP p/c than if it is paired in a type with very low levels of trust and a high propensity to take risks, for example. The same argument, of course, goes for all other outcomes analyzed in Falk et al. (2018).

[29] It would thus be logical that solutions with a larger number of classes or more fine-grained

find a hint on the source of the particularly large spikes of patience and positive reciprocity by looking at their distribution in the raw data: Both are distributed with markedly higher kurtosis and skewness globally. This means that our profile solution will naturally pick up these spikes even with few extracted classes whereas they are less pronounced for the remaining measures.

A related query might concern the distinctiveness of the different pure types. Our labelling and interpretation crucially rest on each type having one or more "characteristic" preference measures. I therefore confirm that the means of these measures exhibit satisfactory separation from the other types.[30] Along the same lines, it might not be sufficient to have a notable spike in the characteristic measure when the in-class variance of the measure remains high. Comparing the variances of each type's characteristic measure(s) to their variance in the whole sample, we verify that grouping similar individuals together does indeed lead to a notable reduction in variance.[31] Taken together, these findings confirm that our labelling really bases on characteristic deviations from the other profiles and we extract distinct types.

Next, we check the sensitivity of our results to the model selection we made in Section 5.1. Figure 3 shows the profile solution for the four other specifications of the covariance matrix, each for a 4-class solution (which was the maximum we could identify for all models). While Masyn (2013) highlights the crucial role model selection plays for the solution obtained, three out of four types, *patient*, *lone wolf* and *pro-social*, are extracted consistently across specifications. The exception is the *retaliative* type which displays some deviations across models but was also our smallest and least accurately identified type. Reassuringly, we furthermore see that similarly complex models yield similar solutions with the model specification most similar to E4 (i.e. D4) yielding almost identical profiles. Subsequently, I also confirm that our results in Sections 5.2 and 5.3 are robust to using the prevalence of each type obtained through model D4.[32]

In Section 3 we developed a framework of *pure* types representing overarching generalizations of regional or development-dependent deviations from them. Therefore, I check the concept behind our framework by performing a latent profile anal-

approaches on a regional level for instance would then also pick up more characteristic variation among these indicators. Indeed, this is exactly what we find.

[30] For the patient profile, for example, we would check that the high mean for patience is indeed characteristic for this type when compared to its value in other types using the adapted formula for Cohen's d given in Masyn (2013).

[31] For all but the patient profile the preference measures actually have their smallest variances in the profiles they are characteristic for, a reassuring finding.

[32] For reasons of limited space, I do not include a new table here.

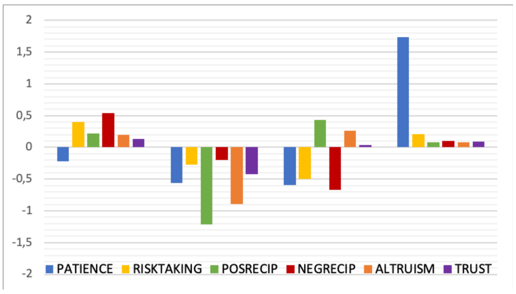

(A) Model A4

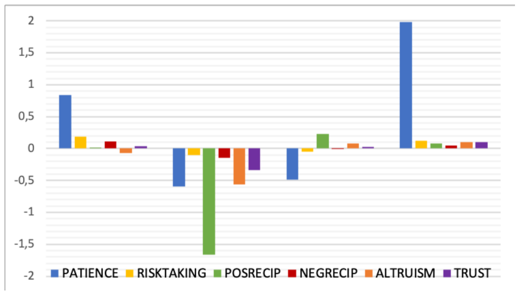

(B) Model B4

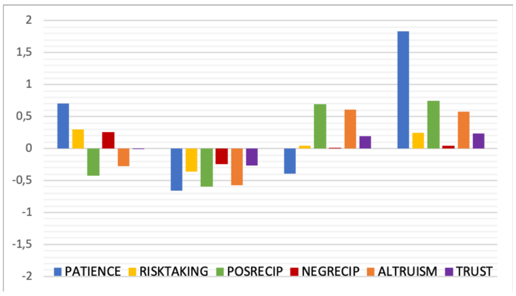

(C) Model C4

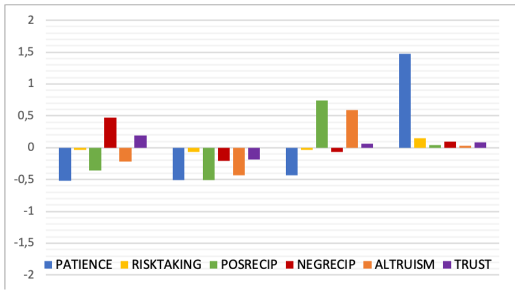

(D) Model D4

FIGURE 3: Profiles obtained from four alternative model specifications A-D. A is the specification of the covariance matrix most different to our model and D the one most similar.

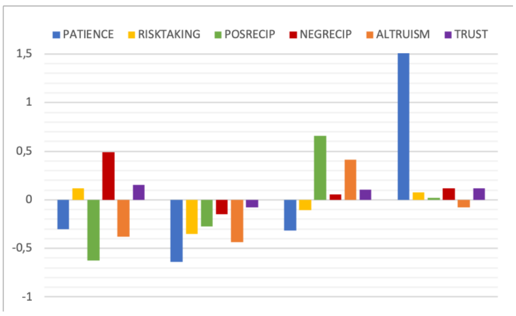

(A) High Income

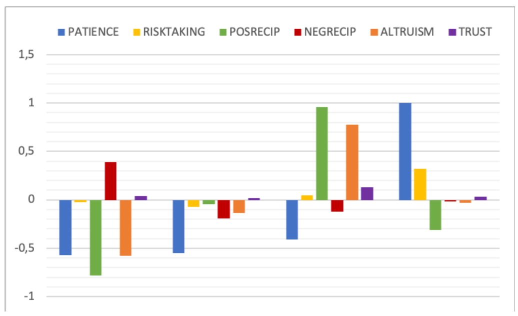

(B) Upper-Middle Income

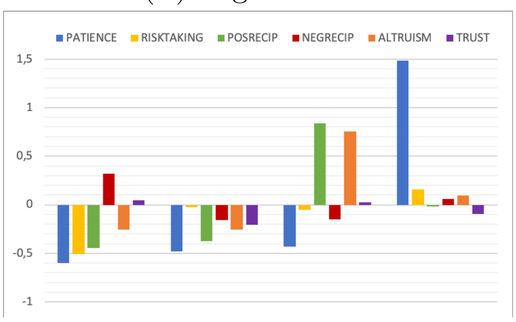

(C) Lower-Middle Income

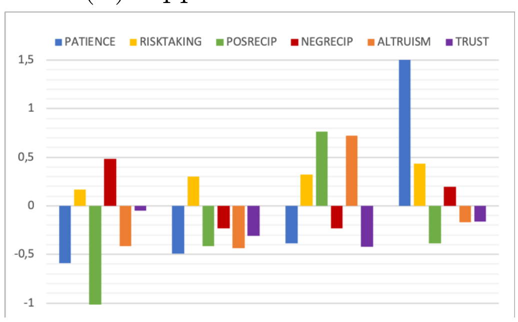

(D) Low Income

FIGURE 4: Subtypes if we perform our analysis separately by income level.

ysis only on subsets of the data depending on income-level.[33] Figure 4 allows us to make two observations. Firstly, the evidence neatly aligns with our framework set out in Section 3. Our four pure types are clearly identifiable across income levels with minor deviations in their shape. Secondly, the deviations we see are mostly associated with more distinctiveness on trust and risk-taking which corroborates our conjecture that moving to a more nuanced analysis also allows for more separation in these measures.

6 Discussion and Conclusion

This paper takes a person-centered approach to analyzing preference heterogeneity in a global sample. It exploits the unprecedented comprehensiveness of the Global Preference Survey, a new data collection of six different preference measures in 76 countries presented in Falk et al. (2018), to explore how different preference measures combine to create unique preference types. Based on the psychological literature on personality measurement and motivated by the recent examples of Gerlach et al. (2018) and Bandiera et al. (2017) establishing typologies as a so-far underutilized but promising field of research, I develop a framework of "pure preference types". In order to identify these clusters of preferences in the population, I use latent profile analysis, a form of mixture model that allows us to take into account an imperfect fit of individuals to the pure types.

The paper identifies four different pure preference types in the population: a *patient* type (24% of the global population), a *retaliative* type (14%) characterized by particularly high values on negative reciprocity, a *pro-social* type (27%) exhibiting high levels of positive reciprocity and altruism, and a type we name *lone wolf* (35%) for their low engagement with others marked by below-average values on all six measures. Furthermore, it finds global heterogeneity in the prevalence of these types within different countries across income levels. The paper then hypothesizes that the prevalence of these types relates to country-level outcomes. To test this, it analyzes correlational patterns between the prevalence of the different preference types and outcomes in the domains of income, entrepreneurship, conflict, gender equality, and democratization. Relative to the lone wolves which we take as benchmark, I find that countries with a larger share of *patient* types tend to exhibit higher levels of economic and social development, while the *retaliative* type is associated with higher prevalence of conflicts and larger shares of *pro-social* types negatively correlate with

[33] I conduct the same analysis with similar results for all seven world regions.

democratization.[34] These results confirm the ones found by Falk et al. (2018), however, they expand on them in that they suggest that it is not the simple mean of individual preferences, but the actual prevalence of a type made up of several preference measures that might account for the patterns they find.

The main contribution of this paper is therefore threefold. The first contribution is methodological in nature. This paper proposes a new approach to modelling preference heterogeneity by taking a person-centered perspective. In doing so, it demonstrates a research design easily transferable to other questions that could pave the way for a new stream of research leading to novel insights into a number of issues relevant in the (economic) development literature and beyond. For example, following its approach could allow for a rigorous, multidimensional analysis of different development trajectories or political economy types of questions. An intriguing expansion here might be to combine this approach with a regression discontinuity design to explore whether threshold levels exist beyond which there is a distinct change in the effect of the latent grouping structures found. Additionally, taking a person-centered approach offers a change in perspective that could yield fresh insights when re-examining the findings of many variable-centered studies (Daljeet et al. 2017).

Secondly, this paper also makes a contribution on the conceptual side. I transfer long-established concepts and insights from the psychological literature on personality and taxonomies to the economic literature. To my knowledge, this allows this paper to be the first one to open up the possibility of a preference typology.

Thirdly, the paper expands on the findings of Falk et al. (2018) and sheds additional light on global preference heterogeneity and its implication for country-level outcomes.

There are some obvious limitations to this study that open up a number of opportunities for future research. Most fundamentally, we have to acknowledge that this study advances to so far uncharted territory with a preference typology. Therefore, its approach should be taken as a first exploration of the topic and its results as a rough sketch which necessitate a thorough investigation into their robustness in future research. This could also include a more fine-grained analysis that can pick up preference type idiosyncrasies by focusing on a specific group of countries for example. Furthermore, while we identify correlational structures between preference types and country-level outcomes, establishing causality and analyzing the underlying mechanisms is beyond the scope of this paper. Establishing panel data

[34] However, these come under the caveat that we do not establish causality here.

on global preference heterogeneity might be a first step towards this and would also allow for a latent transition analysis that could investigate whether and how the composition and shares of our preference types change over time. Additionally, the historic emergence of distinct pure preference types and unique country profiles is an important aspect to further pursue. Becker et al. (2017) suggest that differences between preference profiles of countries have emerged over time. They posit that, departing from one universal preference profile for our very early ancestors, migration has led to populations constantly breaking up into sub-populations, subjecting them to different historic experiences and splitting the genetic pool which resulted in diverging preference profiles according to the time that elapsed since two populations split. Consistent with the framework set out in Chapter 3, this would bear the interesting question whether the pure preference types we find are a contemporary phenomenon that traces back to a single preference type in the very distant past. Lastly, the possibility to acquire microlevel data on the individuals included in the GPS (e.g. their gender) from the Gallup World Polls offers an exciting opportunity for future research into the determinants of preference type membership and its implications for microlevel outcomes.

References

Alan, S., & Ertac, S. (2015). Patience, Self-control and the Demand for Commitment: Evidence from a Large-scale Field Experiment. *Journal of Economic Behavior & Organization, 115*, 111-22.

Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., & Wacziarg, R. (2003). Fractionalization. *Journal of Economic Growth, 8*, 155-94.

Amiti, M., & Konings, J. (2007). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *The American Economic Review, 97*(5), 1611-1638.

Ashton, M. C., Lee, K., Perugini, M., Szarota, P., de Vries, R. E., Di Blas, L., Boies, K. & De Raad, B. (2004). A six-factor structure of personality-descriptive adjectives. *Journal of Personality and Social Psychology, 86*(2), 356-366.

Bandiera, O., Hansen, S., Prat, A., & Sadun, R. (2017). CEO Behavior and Firm Performance. *Harvard Business School Working Paper*, No. 17-083. (Forthcoming, Journal of Political Economy.)

Becker, A., Enke, B., & Falk, A. (2018). Ancient Origins of the Global Variation in Economic Preferences. *NBER Working Paper*, No. 24291, National Bureau of Economic Research, Cambridge.

Becker, A., Deckers, T., Dohmen, T., Falk, A., & Kosse, F. (2012). The Relationship Between Economic Preferences and Psychological Personality Measures. *Annual Review of Economics, 4*(1), 453-478.

Bray, B. C., Lanza, S. T., & Tan, X. (2015). Eliminating Bias in Classify-Analyze Approaches for Latent Class Analysis. *Structural Equation Modeling, 22*(1), 1-11.

DeWilde, C. (2004). The Multidimensional Measurement of Poverty in Belgium and Britain: A Categorical Approach. *Social Indicators Research, 68*(3), 331-369.

Chang, J., Boyd-Graber, J., Gerrish, J., Wang, C., & Blei, D. M. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. *NIPS'09 Proceedings of the 22nd International Conference on Neural Information Processing Systems*, 288-296.

Geiser, C., Okun M., & Grano, C. (2014). Who is Motivated to Volunteer? A Latent Profile Analysis Linking Volunteer Motivation to Frequency of Volunteering. *Psychological Test and Assessment Modeling, 56*(1), 3-24.

Costa, P. T., & McCrae, R. R. (1992). *Revised NEO personality inventory, NEO PI-R, and NEO five-factor inventory, NEO-FFI*, (1. print. ed.). Odessa, Fla: Psycholog. Assessment Resources.

Daljeet, K. N., Bremner, N. L., Giammarco, E. A., Meyer, J. P., & Paunonen,

S. V. (2017). Taking a Person-Centered Approach to Personality: A Latent-Profile Analysis of the HEXACO Model of Personality. *Journal of Research in Personality, 70*, 241-251.

McParland, D., Gormley, I. C., McCormick, T. H., Clark, S. J., Kabudula, C. W., & Collinson, M. A. (2014). Clustering South African Households Based on Their Asset Status Using Latent Variable Models. *The Annals of Applied Statistics, 8*(2), 747-776.

Nagin, D. (2005). *Group-Based Modeling of Development.* Cambridge: Harvard University Press.

DeYoung, C. G., Carey, B. E., Krueger, R. F., & Ross, S. R. (2016). Ten Aspects of the Big Five in the Personality Inventory for DSM?5. *Personality Disorders, 7*(2), 113-123.

DeYoung, C. G., Quilty, L. C., & Peterson, J. B. (2007). Between Facets and Domains. *Journal of Personality and Social Psychology, 93*(5), 880-896.

Dohmen, T., Enke, B., Falk, A., Huffman, D., & Sunde, U. (2018). Patience and Comparative Development. *Working Paper*.

Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2009). Homo Reciprocans: Survey Evidence on Behavioural Outcomes. *The Economic Journal, 119*(536), 592-612.

Falk, A., & Hermle, J. (2018). Relationship of Gender Differences in Preferences to Economic Development and Gender Equality. *Science, 362*(6412), 307.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics, 133*(4), 1645-1692.

Ferguson, E., Heckman, J. J., & Corr, P. (2011). Personality and Economics: Overview and Proposed Framework. *Personality and Individual Differences, 51*(3), 201-209.

Ferguson, S. L., & Hull, D. M. (2018). Personality Profiles: Using Latent Profile Analysis to Model Personality Typologies. *Personality and Individual Differences, 122*, 177-183.

Francis, B., & Liu, J. (2015). Modelling Escalation in Crime Seriousness: A Latent Variable Approach. *Metron, 73*(2), 277-297.

Fréchette, G. R., Schotter, A., & Trevino, I. (2017). Personality, Information Acquisition, and Choice under Uncertainty: An Experimental Study. *Economic Inquiry, 55*(3), 1468-1488.

Gerlach, M., Farb, B., Revelle, W., & Nunes Amaral, L. A. (2018). A Robust Data-Driven Approach Identifies Four Personality Types Across Four Large Data Sets. *Nature Human Behaviour, 2*(10), 735-742.

Golsteyn, B., & Schildberg-Hörisch, H. (2017). Challenges in Research on Preferences and Personality Traits: Measurement, Stability, and Inference. *Journal of Economic Psychology, 60*, 1-6.

Greif, A. (1994). Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies. *Journal of Political Economy, 102*(5), 912-950.

Greif, A., & Kingston, C. (2011). Institutions: Rules or Equilibria? In N. Schofield, & G. Caballero (Eds.), *Political economy of institutions, democracy and voting* (pp. 13-43) Springer.

Henry, M., Kitamura, Y., & Salanié, B. (2014). Partial Identification of Finite Mixtures in Econometric Models. *Quantitative Economics, 5*(1), 123-144.

Hipp, J. R., & Bauer, D. J. (2006). Local Solutions in the Estimation of Growth Mixture Models. *Psychological Methods, 11*(1), 36-53.

Hofstee, W. K. B., de Raad, B., & Goldberg, L. R. (1992). Integration of the Big Five and Circumplex Approaches to Trait Structure. *Journal of Personality and Social Psychology, 63*(1), 146-163.

Hu, L., & Bentler, P. M. (1998). Fit Indices in Covariance Structure Modeling: Sensitivity to Underparameterized Model Misspecification. *Psychological Methods, 3*(4), 424-453.

Kosec, K., Haider, H., Spielman, D. J., & Zaidi, F. (2018). Political Competition and Rural Welfare: Evidence from Pakistan. *Oxford Economic Papers, 70*(4), 1036-1061.

Kremer, M., Rao, G., & Schilbach, F. (2019). Behavioral Development Economics. In E. Cartwright (Ed.), *Handbook of Behavioral Economics* (3rd ed., pp. 345-458). London; New York: Routledge.

Landale, N., Lanza, S., Hillemeier, M., & Oropesa, R. S. (2013). Health and Development Among Mexican, Black and White Preschool Children : An Integrative Approach Using Latent Class Analysis. *Demographic Research, 28*, 1302.

Liangjun Su, Zhentao Shi, & Peter C. B. Phillips. (2016). Identifying Latent Structures in Panel Data. *Econometrica, 84*(6), 2215-2264.

Ludema, R. D., & Mayda, A. M. (2013). Do Terms-of-Trade Effects Matter for Trade Agreements? Theory and Evidence from WTO Countries. *The Quarterly Journal of Economics, 128*(4), 1837-1893.

Marsh, H. W., Lüdtke, O., Trautwein, U., & Morin, A. J. S. (2009). Classical Latent Profile Analysis of Academic Self-Concept Dimensions: Synergy of Person- and Variable-Centered Approaches to Theoretical Models of Self-Concept. *Structural Equation Modeling: A Multidisciplinary Journal, 16*(2), 191-225.

Masyn, K. E. (2013). Latent Class Analysis and Finite Mixture Modelling. In T. D. Little (Ed.), *The Oxford Handbook of Quantitative Methods* (2nd ed., pp. 551-611). Oxford, U.K.: Oxford University Press.

McCrae, R. R., & John, O. P. (1992). An Introduction to the Five-Factor Model and its Applications. *Journal of Personality, 60*(2), 175-215.

McCutcheon, A. L. (2015). *Latent Class Analysis* (2., revised ed. ed.). London [u.a.]: Sage.

Meeusen, C., Meuleman, B., Abts, K., & Bergh, R. (2018). Comparing a Variable-Centered and a Person-Centered Approach to the Structure of Prejudice. *Social Psychological and Personality Science, 9*(6), 645-655.

Merz, E. L., & Roesch, S. C. (2011). A Latent Profile Analysis of the Five Factor Model of Personality: Modeling Trait Interactions. *Personality and Individual Differences, 51*(8), 915-919.

Meyer, J. P., & Morin, A. J. S. (2016). A Person?Centered Approach to Commitment Research: Theory, Research, and Methodology. *Journal of Organizational Behavior, 37*(4), 584-612.

Moisio, P. (2004). A Latent Class Application to the Multidimensional Measurement of Poverty. *Quality & Quantity, 38*(6), 703-717.

Mueller, G., & Plug, E. (2006). Estimating the Effect of Personality on Male and Female Earnings. *Industrial and Labor Relations Review, 60*(1), 3-22.

Muthén, B., & Shedden, K. (1999). Finite Mixture Modeling with Mixture Outcomes Using the EM Algorithm. *Biometrics, 55*(2), 463-469.

Muthén, L. K., & Muthén, B. O. (2017). *Mplus User's Guide* (8th ed.). Los Angeles, CA: Muthén & Muthén.

Nylund, K. L., Asparouhov, T., & Muthén, B. O. (2007). Deciding on the Number of Classes in Latent Class Analysis and Growth Mixture Modeling: A Monte Carlo Simulation Study. *Structural Equation Modeling: A Multidisciplinary Journal, 14*(4), 535-569.

Ozer, D. J., & Benet-Martínez, V. (2006). Personality and the Prediction of Consequential Outcomes. *Annual Review of Psychology, 57*(1), 401-421.

Pastor, D. A., Barron, K. E., Miller, B. J., & Davis, S. L. (2007). A Latent Profile Analysis of College Students? Achievement Goal Orientation. *Contemporary Educational Psychology, 32*(1), 8-47.

Ram, N., & Grimm, K. J. (2009). Methods and Measures: Growth Mixture Modeling: A Method for Identifying Differences in Longitudinal Change Among Unobserved Groups. *International Journal of Behavioral Development, 33*(6), 565-576.

Rustichini, A., DeYoung, C. G., Anderson, J. E., & Burks, S. V. (2016). Toward the Integration of Personality Theory and Decision Theory in Explaining Economic Behavior: An Experimental Investigation. *Journal of Behavioral and Experimental Economics, 64*, 122-137.

Satorra, A., & Bentler, P. M. (2010). Ensuring Positiveness of the Scaled Difference Chi-Square Test Statistic. *Psychometrika, 75*(2), 243-248.

Vollmer, S., Holzmann, H., Ketterer, F., Klasen, S., & Canning, D. (2013). The Emergence of Three Human Development Clubs. *PLoS One, 8*(3), 1-7.

Mullainathan, S., & Spiess, J. (2017). Machine learning: An Applied Econometric Approach. *The Journal of Economic Perspectives, 31*(2), 87-106.

Soto, C. J. (2019). How Replicable are Links Between Personality Traits and Consequential Life Outcomes? The Life Outcomes of Personality Replication Project. *Psychological Science, 30*(5), 711-727.

Stanley, L., Kellermanns, F. W., & Zellweger, T. M. (2017). Latent Profile Analysis: Understanding Family Firm Profiles. *Family Business Review, 30*(1), 84-102.

Tofighi, D., & Enders, C. K. (2007). Identifying the Correct Number of Classes in Growth Mixture Models. In G. R. Hancock (Ed.), *Advances in Latent Variable Mixture Models* (pp. 317-341). Charlotte, NC: Information Age Publishing.

Weber, M. (1949). In Shils E. A., Finch H. A. (Eds.), Max Weber on the Methodology of the Social Sciences, The Free Press of Glencoe.