INE PAN Working Paper Series

Paper number 51

Gaps in many dimensions: application to gender

Martyna Kobus, Marek Kapera, Esfandiar Maasoumi

Warsaw, 18.05.2023

wp@inepan.waw.pl

Gaps in many dimensions: application to gender

Martyna Kobus[*], Marek Kapera[†], Esfandiar Maasoumi[‡]

November 2022

Abstract

We treat the problem of measuring and estimating gaps between groups (e.g. gender, ethnicity) and counterfactual effects when the outcome is multivariate. This is an extension of the standard gender pay gap that has a univariate outcome, typically wage or income. We refer to such gaps as well-being gaps. We extend decomposition of the gap into counterfactual effects, pursuant to the multidimensionality of outcomes. However, in our setting, gaps can also be decomposed in a different form, namely, the one that distinguishes the dependence between outcomes from the impact of the marginal distributions. Greater concordance means, *ceteris paribus*, more inequality or divergence between groups. Combining two forms of decomposition gives the full decomposition of multivariate gaps. Overall, we provide a toolbox to perform multidimensional gap analysis based on known tools and newly proposed methods. We apply our methodology to study of the gender gap in wages and leisure. In recent years, there has been a growing body of research that has delved into the differences between genders when it comes to both paid and unpaid work. The utilization of time use surveys has greatly enhanced empirical research in this area. The data has revealed that on average, women tend to take on more unpaid care responsibilities than men. As a result, men earn more and enjoy more leisure time when the two are analyzed separately. We contribute to this literature by utilizing the American Time Use Survey to estimate and decompose gaps between genders in wages and leisure considered jointly.

Keywords: gender gaps; well-being; counterfactual distribution, decomposition

JEL classification: D30; I31; C02

1 Introduction

The interest in the gender gap, or racial gap, and similar questions of inequality derives from an interest in well-being of target populations and allocation and distributive questions in

[*]Institute of Economics, Polish Academy of Sciences, 00-330 Warsaw, Nowy Swiat 72, Poland, mkobus@inepan.waw.pl.

[†]Institute of Economics PAS, mkapera@inepan.waw.pl.

[‡]Emory University, Atlanta, GA 30322-2240, USA, esfandiar.maasoumi@emory.edu

general. The choice of wages or 'income' as a basis for measuring inequality, poverty, mobility, and the gaps thereof reflects convenience and primacy of income as an indicator of well-being. Health status, educational attainment, or social amenities provided by communities in which families and individuals reside, jointly determine their well-being. Well being is a latent concept, like happiness, and the challenges in its measurement have been widely acknowledged for many years. Recent Nobel prizes in economics (Angus Deaton for his efforts in these regards, and Amartya Sen for poineering contributions) reflect the centrality and urgency of metters herein. The multivariate analysis of well-being gaps is diminished when it is conducted in any one dimension, since the same dollar level of wages has vastly different values for a healthy or unhealthy individual, or one that has the investment and consumption benefits of high levels of educational attainment, living in a society that offers strong social nets. One-dimensional accounting is extremely subjective, as it places zero weights on all other sources of well-being and their interactions. Multivariate distributions of attributes must be contrasted for a more informed and realistic view of allocations and well-being.

This poses additional challenges to those that are increasingly being addressed in one dimension, income/wages, in the literature on counterfactual distributions. This new challenge comes from the dependence between dimensions of well-beinq, which makes multidimensional analysis irreducible to the analysis of unidimensional gaps. If multiple outcomes are analyzed separately, then an important piece of information is lost, namely, whether individuals who are worse (or better) on one dimension are the same individuals who are worse (better) on other dimensions. The presence of such association will lead to higher or lower estimations of differences between groups than can be inferred from multiple unidimensional analyses. For example, if women have worse wage and worse health outcome distributions than men, and additionally, women's wages and health are more dependent than men's (so that there is higher likelihood that a woman with bad health has also low income than it is for a man with bad health), then we should conclude that there is higher gap between men and women than we would conclude from just looking separately at the wage gap and health gap. Or, if women still have worse wage distribution but health comparisons are ambiguous, and in addition, men have higher correlation of wages and health for workers with bad health, it may happen that in the range of bad health outcomes, women actually have better joint distribution of wages and health than men. In such cases, a joint analysis can often resolve ambiguity arising from univariate analyses. The outcomes must be considered together to get the full picture of the well-being comparisons between the groups. As Blau (1998) notes: "First, it is important that a broad range of indicators of well-being be employed, encompassing not only labor market outcomes like wages and occupations, but also time available for leisure, the level of family income, the share of women who are single family heads (...) for forming a more complete picture of changes in women's well-being than may be obtained elsewhere in the literature." This broad view has for long been acknowledge in the analysis of well-being, inequality, poverty, but has not been extended to the analysis of gaps and counterfactual effects.

The aim of this paper is to propose a theory for identification and estimation of well-being gaps. We call them well-being gaps in line with well-being measurement literature that proxies well-being through the use of multiple outcomes, e.g., well-known indices such as OECD's Better Life Index or Human Development Index. These gaps can be differences between men and women, or any exogenuously defined groups. In accordance with the literature on counterfactual distributions and decomposition methods, we differentiate between distributional and quantile effects. In each case, we check which of the tools known in the literature for analyzing multidimensional distributions can be used to study these effects. We also offer new tools and new effects of substantive policy implications. We start with a statistically complete definition of the gap in the multidimensional framework. For quantile effects, this is additionally problematic since, in this setting, quantiles are sets. The quantile difference, which is very natural in the unidimensional setting, is no longer meaningful. We propose solutions to this problem. With each tool we use to analyze gaps we check whether it allows for the decomposition into structural and composition effects. Otherwise, this tool might be good for studying and comparing multidimensional distributions but not very useful for the analysis of type of comparisons which is of interest in the gap literature. Thus, some of the well-known tools are excluded as unhelpful.

There are two directions for decomposition in this setting. Apart from the traditional decomposition into structure and composition effects, we are also interested in decomposition by dimensions of well being. In particular, we are interested in what part of the overall gap between two groups comes from the gap in, say, wages, the gap in leisure and the gap in the dependence between wages and leisure in both groups. Are women with lower wages also the ones with the most free time, or vice versa? Is it the same for men or different, and if so, to what extent? Thus, we would ask what proportion of the wage-leisure gap between men and women is due to different patterns of dependence between wages and leisure in both groups?[1] If both forms of decomposition are possible, then we can be even more exact, and trace the sources of the gap to the dependence differences that are structural and that are compositional. In more detail, exact decomposition means that the gap can be decomposed into six additive elements, namely the structural effect of the first outcome, the structural effect of the second outcome and the structural effect of the dependence, the composition effect of the first outcome, the composition effect of the second outcome, and the composition effect of dependence. From this various other effects can be computed such as, for example, overall composition effect, or overall effect of the dependence. These decomposition effects

[1]Dependence also often has additional interpretation in various contexts, e.g., with income and health, dependence is the widely studied income gradient in health, or with several health indicators dependence is understood as comorbidity, namely, the joint occurrence of diseases.

are of great consequence for better informed policy decisions. We show when useful solutions exist and when they do not.

As mentioned, the new framework allows for new definitions of counterfactual effects. A counterfactual distribution results either from a change in the distribution of covariates (keeping women's wage structure fixed but with men's characteristics), or from a change in the relationship between covariates and outcome, that is, a change in the conditional distribution of outcome given a set of characteristics (keeping women's characteristics fixed but giving them men's wage structure). The latter effect is called the structural effect and the former is called the composition effect. In a multivariate setting it is possible to define a new counterfactual effect such that we give women access to marginal conditional distribution of men's outcomes, but we keep their preferences (substitution parameters) towards the two outcomes. This is possible, because in one of the solutions that we offer, quantile sets are modelled as a latent well-being and substitution weights can be recovered. This technique was developed by Maasoumi and Racine (2016) and we apply it to the analysis of gaps. Parameterization also helps to avoid the thorny problem of measuring the distance between quantile sets. As mentioned, the quantile difference is simple with a single outcome only. Already with two outcomes, there is a whole quantile curve associated with a given quantile, and it is not a priori known what is meant by quantile difference. We study some of the definitions from the mathematics literature in this respect, but they do not allow for decomposition, and hence are not useful. Maasoumi and Racine (2016) solution offers a way out and additionally addresses the heterogeneity challenge. The latter is an inherent problem of any aggregation problem, namely, typically used summary measures assume homogeneous evaluation functions across individuals and groups, which in the presence of substantial heterogeneity (i.e., individuals differ in substitution parameters) results in great arbitrariness of such measures. For example, to identify quantile effects, one typically has to assume rank invariance or similarity, an assumption that women endowed counterfactually with men's skill set will occupy the same ranks in counterfactual and observed setting. This is often unsupported empirically. The possibility that men and women may engage differently in skills substitutions is ruled out. This confirms the reservations of Heckman et al. (1997) as to the meaningfulness of quantile gaps. Applying Maasoumi and Racine (2016) methods to the analysis of quantile gaps, we allow for different weights and substitution values at different quantiles/groups andcheck whether heterogeneity is present.

All in all, for the analysis of multivariate distributional gaps, i.e. gaps in the values of distribution or density functions, we propose the Tsui (1995) inequality measure, which we show is decomposable by dimensions. Thus exact decomposition into dimensions and structural and composition effects is possible with this measure. Furthermore, the gap defined as a difference in log probability can also be decomposed in the exact manner utilizing the factorizing of the joint density into marginal densities and a copula density. Log probabilities are often used in computation to replace multiplication with addition of probabilities. Similarly to reporting chosen quantiles of the gender wage gap, we may report the full decomposition of the difference in log densities between men and women at chosen quantiles. However, we may also summarize this information using entropy measures. The average of log probability is known as entropy, and as we show, it is also decomposable in exact manner. In particular, entropy of a copula is a well-known concept in information theory, namely, mutual information. For the analysis of multivariate quantile effects, as already mentioned, we suggest parametrization of quantiles via a latent concept of well-being which is implemented using a CES utility function. The assumption is that individuals who belong to the same quantile set enjoy the same level of well-being/CES utility. With this assumption, at each quantile set, the parameters of the CES function can be estimated and the difference between well-being of men and women at a given quantile can be computed. To ensure some scale invariance, we use well-being shares rather than scores. The parameters of CES function will potentially differ between quantile levels reflecting heterogeneity in the aggregation of outcomes. A new counterfactual effect can be defined such that women are assigned men's conditional distribution of both outcomes, but they keep their substitution parameters. Less formally speaking, this is the counterfactual effect of women having access to men's structure of both outcomes, but retaining their own preferences towards these outcomes.

The estimation can be performed applying the DFL re-weighting method (DiNardo et al. 1996) to the non-parametric estimate of the multivariate conditional distribution. We apply the developed toolbox to the analysis of the American Time Use Survey (ATUS). The American Time Use Survey, also known as ATUS, is a yearly study that began in 2003. It utilizes data from the Current Population Survey to gather information on how individuals use their time over a 24-hour period, referred to as the "diary day". The survey allows for a detailed understanding of the activities that respondents engage in and the amount of time they spend on each. The survey also collects information on the respondent's education level, employment status, and earnings, as well as demographic information about the household. The survey has been collecting data from 2,190 households per month, or around 26,400 households per year, since December 2003. Initially, the sample size was 3,375 households per month, or about 40,500 households per year, but it was later decreased to cut costs. The households selected for the survey are divided into 12 strata, based on their race/ethnicity (e.g. Hispanic, Non-Hispanic black, Non-Hispanic non-black) and type of household (e.g. households with children under 6, households with children between 6 and 17, single adult households without children, households with two or more adults without children). We use years 2005-2021, because of some variables restrictions and study wage-leisure gaps.

The paper is incomplete when it comes to empirical application of the proposed methods. However, a few things that emerge are to note. Both the gaps measured by inequality measure and entropy are decreasing, for entropy the decrease happened in the first years of 2000s and then the gap remains stable. There is a more moderate decrease in the leisure gap alone than in the wage gap, but both gaps are decreasing and as such contribute to the decrease of the overall gap. The dependence gap, on the other hand, is increasing. It is also negative, meaning that women exhibit less dependence between wages and leisure than men and this difference is becoming greater in the last years. In each year, the structural effects for all dimensions and dependence are greater than composition effects. In fact, for wages composition effects start to become negative, meaning that women have surpassed men when it comes to the distribution of observable characteristics. Maasoumi and Wang (2019) also find this analysing CPS data and wages only. The results reveal substantial heterogeneity in the substitution parameters with varying quantiles. Willingness to substitute is higher at lower quantiles of joint wage-leisure distribution. Women have lower elasticity of substitution between wage and leisure than men almost uniformaly, that is, at each quantile. However, these differences were greater in 2005 than in 2021.

The paper is organized as follows. In Section 2 we review the related literature. In Section 3 we provide a theory of multidimensional distributional gaps, followed by decomposition possibilities. In Section 4 we treat multidimensional quantile gaps. Section 5 contains empirical application. In Section 6 we offer discussion on further work and conclude.

2 Relation to the literature

This paper builds a bridge between the literature on well-being measurement, which now widely adopts the view that well-being is inherently a multidimensional concept (Stiglitz et al. 2009) and the econometrics literature on group inequalities. In more detail, the paper relates to several big areas of economics research, namely, empirical evidence on gender gaps (Blau and Kahn 1997), wellbeing research (Deaton 2015), estimation of counterfactual distributions and policy effects (Chernozhukov et al. 2013) and decomposition methods (Fortin et al. 2011).

The distribution of wages in the US has evolved differently for men and women over several decades. Using conventional gap measures as the mean and quantiles, economists have shown that the gender gap has decreased over time, especially in the 1980s and early 1990s. This trend has slowed down since the mid-1990s. Women are catching up with men (Blau and Kahn 1997, Goldin 2014). The quantile gaps, however, have evolved differently over the past decades (Albrecht et al. 2003). The measurement of the gap is filled with problems of heterogeneity and selection. In the presence of significant heterogeneity, all gap measures are biased. Rank invariance assumption, namely, that agents occupy a given quantile in both factual and counterfactual setting, does not hold. Men and women, when endowed with each other's skill set or market returns, may engage differently in skills substitutions, work and job decisions. Gap measures are biased also when there is selection into employment.

That is, if non-working men and women systematically differ from working men and women, measures of the gap would be biased. For example, if only high-earning women enter labor market we may observe convergence in the gap with no actual progress. There are only few attempts in the gender gap literature that account for selection (Blau and Kahn 2006, Mulligan and Rubinstein 2008). Recently, Maasoumi and Wang (2019) examine the gap using General Entropy measures which overcomes the problem of heterogeneity. These are anonymous aggregative measures of the gap. They also use the novel copula approach developed by Arellano and Bonhomme (2017) to model selection. Once selection is accounted for, the gap converges slower than without selection and in fact the trend reverses in some parts of the distribution between mid-1990s and the most recent recession. In the great recession, there was a marked decline in the gap among low-skilled workers, perhaps due to a relative deterioration in wages of the low-skilled males. Convergence is much smaller amongst the least educated or black women, especially during the recent years.

Economists have for long acknowledged that improvements in well-being are not confined to economic growth (Sen 1973). Governments (for example in UK, France, Canada, Japan) and international organizations have been increasingly responding to economists' calls to go beyond GDP in measuring nations' progress. Measuring well-being is now a prominent item on the agenda of many statistical offices. There is a wide spread recognition that well-being measurement is critical for informing policy making about areas that matter to people. Specific measures taken in the last decade include the construction of the OECD's Better Life Index, the report by Stiglitz-Sen-Fitoussi commission on the measurement of economic performance and social progress, the program initiated by the British Office for National Statistics on measuring nation's well-being or the 2011 resolution of the United Nations General Assembly (No. 65/309) which explicitly "invites Member States to purse elaboration of additional measures that better capture the importance of the pursuit of happiness and well-being" than GDP-based indicators.

Gender gap is an important source of well-being inequalities in the population. On one hand, women live longer than men on average and they are often more educated. On the other hand, they report a lower health status, have worse job prospects and fewer professional networks. Furthermore, various gaps are interrelated. In the context of our empirical application, gender gaps in time allocation produce gender inequalities in work experience and earnings. Research has shown (Aguiar and Hurst 2007, Aliaga 2006, Fisher and Robinson 2011, Gimenez-Nadal and Sevilla 2012, Gimenez-Nadal and Molina 2020) that women tend to spend more time on household tasks, such as cooking, cleaning, and caring for others, compared to men. These patterns of specialization, with women devoting more of their time to household production, may be detrimental for their labour market opportunities. In a recent article, (Campaña et al. 2023) write: "The overall picture is that the bulk of unpaid work is still performed by women. The gender difference in unpaid work is larger than the gender difference in labour market work, reflected in the allocation of leisure time by gender, indicating that men spend slightly more time on leisure than do women. (...) the results show that the gender gap in both paid and unpaid work has decreased in most countries, leading to a more egalitarian gender distribution of leisure time. However, these gender gaps are still significant (...)." It thus seems natural to think about well-being differences and not just wage differences when comparing groups (men and women, black and white, and other groups). This justifies the types of problems tackled in this paper. This broad view of the gap requires developing the theory and estimation procedures for measuring it. This problem has not been considered so far in econometrics literature, which constitutes the pioneering nature of the paper.

Since the gaps considered in the paper are multidimensional we briefly describe the state of art and issues present in multidimensional aggregation. Aggregation is a key challenge in the multi-outcome setting. The classic problem of aggregation over agents (as in any inequality analysis) is compounded by an inevitable aggregation over outcome indicators. Aggreation over individuals is subject to Arrow's impossibility, and aggregation over outcomes requires empirically valid weights and substitution values. The latter are not readily available since many outcomes may not even be priced through markets. Due to inevitable subjectivity of aggregation rules (associated to ANY scalar measure), researchers have examined uniform (partial) ranking i.e. stochastic dominance criteria. In the multidimensioned context, this approach was first illucidated by Atkinson and Bourguignon (1982). In the SD setting, ranking of outcomes is over classes of welfare functions and outcome aggregators. The conditions for rankings are, however, restrictive. The problem of heterogeneity is significant i.e. groups with different needs cannot be represented with identical utility functions. On the other hand, certain plausible assumptions about the trade-offs between incomes and other variables, such as health or other "needs" (and if groups can be ranked with respect to "needs"), allow partial comparability (Atkinson and Bourguignon 1987). However, Maasoumi's and Racine's (2016) example based on Indonesian data suggests that heterogeneity is so high that the above assumptions seem generally implausible. Consequently, they do not find many findings of statistically significant multidimensional rankings.[2]

As a solution they propose the following procedure. They follow the two step approach of Maasoumi (1986) in which outcome aggregation is a transparent first step, producing an aggregate "index" of well being, followed by measures based on the distribution of the aggregate index. Next, they offer a nonparametric data-dependent method to recover substitution and weights values required for an aggregator index.[3] Such index based evaluation functions (ag-

[2]This has major implications for all summary measures of inequality, poverty, or equivalence scales, which all make assumptions about relatively homogeneous evaluation functions across individuals and even groups. Robust analysis may have to avoid summary measures as the assumptions they employ seem to be empirically unsupported.

[3]Another data-dependent method is given in Maasoumi and Tong (2015) who find estimates of a happiness

gregators) provide for complete multidimensional ranking of welfare states. Their solution stems from the insight that non-parametrically estimated joint distribution of outcomes contains information in the purest form, free from arbitrary assumptions. Non-parametrically estimated quantile sets are iso-quants of well-being states. They choose a particular aggregator, namely, generalized entropy measure that has the advantage that it minimizes the divergence between the joint distribution and the distribution of its constituent dimensions. The CES aggregation is a member of this class. Mapping the non-parametric joint distribution, possibly conditional on other characteristics, to the chosen aggregator functions provides estimates for weights and substitution parameters. In what follows, we use their approach to study multidimensional quantile gaps.

3 Distribution effects

The basic framework with which we work follows Chernozhukov et al. (2013) modified to the case of two outcomes.[4] Let us suppose that we have two populations, of men (denoted 0) and women (denoted 1). For each population $k = 0, 1$, there is a random d_x-vector X_k of covariates and a random outcome d_y-vector Y_k, where for now we assume $d_y = 2$. We denote the support of X_k by $\mathcal{X}_k \subseteq \mathbb{R}^{d_x}$ and denote the region of interest for outcomes Y_k by $\mathcal{Y}_k \subseteq \mathbb{R}^2$. The conditional distribution functions $F_{Y_0|X_0}(y|x)$ and $F_{Y_1|X_1}(y|x)$ describe the stochastic assignment of wages to workers with characteristics x, for men and women, respectively. Let $f_{0|0}$ and $f_{1|1}$ represent the observed distribution function of wages for men and women, and let $f_{0|1}$ represent the counterfactual distribution of wages that would have prevailed for women had they faced the men's wage schedule $F_{Y_0|X_0}$. Further, $f^1_{0|0}, f^1_{1|1}, f^1_{0|1}$ and $f^2_{0|0}, f^2_{1|1}, f^2_{0|1}$ denote, respectively, the first and the second marginal of the three distributions of interest.

Definition 1. *The counterfactual distribution $F_{0|1}$ is defined as*

$$F_{0|1}(y) := \int_{\mathcal{X}_1} F_{Y_0|X_0}(y|x) dF_{X_1}(x), \; y \in \mathcal{Y}_0 \tag{1}$$

Such integral is well-defined if the support condition holds if $\mathcal{X}_0$, the support of characteristics for men, includes $\mathcal{X}_1$, the support of women's charactertistics, namely, $\mathcal{X}_1 \subseteq \mathcal{X}_0$. This distribution is called counterfactual because it does not arise as a distribution from any observable population. Rather, it is constructed by integrating the conditional distribution of both outcomes for men with respect to the distribution of characteristics for women.

The definition of the distributional gap is simply $F_{1|1} - F_{0|0}$. This difference can be decomposed in the spirit of Oaxaca (1973) and Blinder (1973), which as already mentioned, we call "decomposition by sources".

index in China and show that incorporating substitution among attributes, and taking into consideration group heterogeneity are very important in multidimensional analysis of well-being.

[4]The extension to several outcomes is straightforward in this setting and we comment on it in Section 6.

Definition 2. *The decomposition by sources of the distributional gap is the following*

$$F_{1|1} - F_{0|0} = \underbrace{\left(F_{1|1} - F_{0|1}\right)}_{\text{Structural effect}} + \underbrace{\left(F_{0|1} - F_{0|0}\right)}_{\text{Composition effect}} . \tag{2}$$

The difference $F_{1|1} - F_{0|0}$ is the overall effect i.e. the gap, and the two effects on the right-hand side of the equation are called counterfactual effects (CE). The first term in brackets is due to differences in the wage structure and is called *the structural effect*. The second term is *the composition effect* that arises due to differences in characteristics between men and women. These counterfactual effects are well defined econometric objects and are widely used in empirical analysis. For example, the structural effect is a measure of gender discrimination. It is important to note that these effects do not necessarily have a causal interpretation without additional conditions. In particular, the assumption of conditional exogeneity, selection on observables or unconfoundedness is often invoked to interpret CE as casual effects. This can be separated from the measurement methods applied separately as explained in Chernozhukov et al. (2013).

The definition of the multivariate distributional gap shows consistency in a sense that it relates naturally to the univariate gaps.

Proposition 1. $F^1_{0|1} = F_{Y_1<0|1}$ *and* $F_{0|1>2} = F_{Y_2<0|1}$

Proof.

$$\begin{aligned} F_{0|1>1}(y_1) &= \int_{\mathrm{supp} Y_{0|0>2}} f_{0|1}(y_1, y_2)\mathrm{d}y_2 = \int_{\mathrm{supp} Y_{0|0>2}} \int_{\mathrm{supp} X_1} F_{Y_0|X_0}(y_1, y_2|x)\mathrm{d}x\mathrm{d}y_2 = \\ &= \int_{\mathrm{supp} X_1} \int_{\mathrm{supp} Y_{0|0>2}} F_{Y_0|X_0}(y_1, y_2|x)\mathrm{d}y_2\mathrm{d}x = \int_{\mathrm{supp} X_1} F_{Y_{01}|X_0}(y_1|x)\mathrm{d}x = \\ &= F_{Y_1<0|1} \end{aligned}$$

□

Proposition 1 states that the cumulative counterfactual distribution of Y_1 (and of Y_2) is the marginal cdf of the joint counterfactual distribution of Y_1, Y_2. That is, the construction of counterfactual marginal distribution can be recovered from counterfactual joint distribution by taking its marginals.

To sum up, the distributional effects, definitions and decomposition by sources are the same as in the univariate setting. However, they are rarely reported as such, but ultimately the decomposition results are quantified using some distributional statistic such as, for example, an inequality measure. We will now show how this extends to the multivariate setting. In particular, we will inspect how the multivariate gap can be analyzed using well-known tools for dealing with multivariate distributions and examine the key role of dependence between outcomes (i.e. "decomposition by outcomes"). In what follows, we use F to denote joint cumulative distribution function, and F_1, F_2 to denote its marginals.

3.1 Multivariate inequality measures

Although distribution effects can be easily defined, the decomposition results are ultimately quantified using some distributional statistic ν such as, for example, an inequality measure. In our setting these will be multivariate inequality measures, that is, measures that measure inequality in multiple outcomes and not just income. The definition of the gap using distributional statistic is simply $\nu(f_{1|1}) - \nu(f_{0|0})$ and its decomposition is the following,

Definition 3. *The decomposition of the gap into structural and composition effects using distributional statistic is the following*

$$\nu(f_{1|1}) - \nu(f_{0|0}) = \big(\nu(f_{1|1}) - \nu(f_{0|1})\big) + \big(\nu(f_{0|1}) - \nu(f_{0|0})\big)\,. \tag{3}$$

Now ν can be any multivariate inequality measure. However, because our goal is to also be able to distinguish the impact of outcomes' dependence on the overall gap, we will use the measure that can be decomposed in such a way that dependence can be isolated from the impact of marginal distributions. This way of decomposing the measure is known as attribute decomposition and was introduced into inequality measurement literature in a paper by Abul Naga and Geoffard (2006).[5]

Definition 4. *Let I denote the multivariate inequality measure and I_1, I_2 denote univariate inequality measures. I is attribute decomposable if and only if*

$$I(F) = g_1(I_1(F^1) + g_2(I_2(F^2)) + g_3(\kappa(F)),$$

where g_1, g_2, g_3 are increasing functions and κ is a measure of association between Y_1, Y_2 with joint cdf F.

We will consider the following measure, which (slightly abusing notation) we also denote I.

$$I(F) = 1 - \left(\int \frac{y_1^{r_1} y_2^{r_2}}{\bar{y}_1^{r_1} \bar{y}_2^{r_2}} \mathrm{d}F(y_1, y_2)\right)^{\frac{1}{r_1+r_2}}, \tag{4}$$

where $\bar{y}_2, \bar{y}_2$ denotes the mean of Y_1, Y_2, respectively, and r_1, r_2 are dimensional weights. This measure is a simple multidimensional generalization of the Atkinson (1970) measure derived by Tsui (1995). Therefore, $r_q = 1 - \epsilon_q$, where ϵ_q is the degrees of inequality aversion in outcome q. As ϵ_q rises, more weight is attached to inequality-reducing transfers at the bottom of the distribution than at the top. In his seminal work, Atkinson (1970) arbitrarily set the inequality aversion parameter equal to 1, 1.5 or 2. Subsequently, empirical research has tried to infer plausible values of ϵ_q from tax schedules (Aristei and Perugini, 2016; Gouveia and Strauss, 1994; Young, 1990). These estimates also range between 1 and 2 depending on the country and time period.

[5]Although attribute decomposition is a desirable property of multivariate measures, it is technically difficult and there are not many general results concerning this form of decomposition; see e.g. Kobus (2012).

Abul Naga and Geoffard (2006) provide the decomposition of Tsui (1995) measure into attributes and a residual term that captures the effect of attributes' dependence in a discrete setting. Here we reformulate this result in a continuous setting and then apply linearization to obtain additive decomposition.

Proposition 2. *I is attribute decomposable as follows*

$$I(F) = \frac{r_1}{r_1+r_2} I_1 + \frac{r_2}{r_1+r_2} I_2 + \frac{1}{r_1+r_2}(1-\kappa) + R, \tag{5}$$

where

$$\begin{array}{rl} 1 - \left(\int \frac{y_1^{r_1}}{\bar{y}_1^{r_1}} \mathrm{d}F_1(y_1) \right)^{\frac{1}{r_1}} & \text{unidimensional index } I_1 \\ 1 - \left(\int \frac{y_2^{r_2}}{\bar{y}_2^{r_2}} \mathrm{d}F_2(y_2) \right)^{\frac{1}{r_2}} & \text{unidimensional index } I_2 \\ \frac{\int y_1^{r_1} y_2^{r_2} \mathrm{d}F(y_1,y_2)}{\int y_1^{r_1} \mathrm{d}F_1(y_1) \int y_2^{r_2} \mathrm{d}F_2(y_2)} & \text{measure of association } \kappa \\ R & \text{approximation residual.} \end{array}$$

Proof. Please see the Appendix. □

Proposition 2 states that Tsui (1995) measure is attribute decomposable and provides the exact decomposition into I_1, I_2, Atkinson's (1970) unidimensional indices of inequality, and a measure of association κ. The precise decomposition[6] is not additive, but multiplicative

$$I(F) = 1 - (1-I_1)^{\frac{r_1}{r_1+r_2}} (1-I_2)^{\frac{r_2}{r_1+r_2}} (\kappa)^{\frac{1}{r_1+r_2}}.$$

Applying linear decomposition around the point of perfect equality (i.e. $I_1 = I_2 = 0, \kappa = 1$) gives (5) and R is a residual from this approximation.

Let us now shed more light on κ. Abul Naga and Geoffard (2006) show that it increases following correlation-increasing switch. This is a transformation that puts probability mass on the diagonal, hence increases correlation between two variables. A reasonable measure of association should react positively to this transfer and so does κ. Furthermore, it can be seen easily that $\kappa = 1 + \frac{\mathrm{Cov}\left(Y_1^{r_1}, Y_2^{r_2}\right)}{\mathbb{E}Y_1^{r_1}\mathbb{E}Y_2^{r_2}}$. That is, κ is a function of covariance normalized by $\mathbb{E}Y_1^{r_1}\mathbb{E}Y_2^{r_2}$. Therefore, $\mathrm{Cov}\left(Y_1^{r_1}, Y_2^{r_2}\right) = 0 \iff \kappa = 1$. For uncorrelated variables, κ equals 1. Therefore, in a bidimensional model, $\kappa = 1$ together with unidimensional inequalities equal to zero (i.e. $I_1 = I_2 = 0$) represent the case of perfectly equal distribution. Furthermore, for positive variables, we also have that $\mathrm{Cov}\left(Y_1^{r_1}, Y_2^{r_2}\right) < 0 \iff \kappa < 1$ (negatively dependent variables) and $\mathrm{Cov}\left(Y_1^{r_1}, Y_2^{r_2}\right) > 0 \iff \kappa > 1$ (positively dependent variables).

With this decomposition, we can now study the role of dependence in the overall gap. Coming back to our original setting with observed and counterfactual distribution, we can further decompose (3).

[6] Please refer to the Appendix.

Proposition 3. *The gap $I(f_{1|1}) - I(f_{0|0})$ can be decomposed as follows*

$$
\begin{aligned}
I(f_{1|1}) - I(f_{0|0}) = & \\
\frac{r_1}{r_1+r_2}\Big(I_1(f^1_{1|1}) - I_1(f^1_{0|1})\Big) & \quad \text{structural effect for outcome 1} \\
+\frac{r_2}{r_1+r_2}\Big(I_2(f^2_{1|1}) - I_2(f^2_{0|1})\Big) & \quad \text{structural effect for outcome 2} \\
+\frac{1}{r_1+r_2}\big(\kappa(f_{1|1}) - \kappa(f_{0|1})\big) & \quad \text{structural effect for dependence} \\
+\frac{r_1}{r_1+r_2}\Big(I_1(f^1_{0|1}) - I_1(f^1_{0|0})\Big) & \quad \text{composition effect for outcome 1} \\
+\frac{r_2}{r_1+r_2}\Big(I_2(f^2_{0|1}) - I_2(f^2_{0|0})\Big) & \quad \text{composition effect for outcome 2} \\
+\frac{1}{r_1+r_2}\big(\kappa(f_{0|1}) - \kappa(f_{0|0})\big) & \quad \text{composition effect for dependence} \\
+R & \quad \text{residual from linear approximation}
\end{aligned}
$$

Given Proposition 3 we can compute the overall effect of dependence and also specific effects (i.e. structural and composition) in the gap. For example, $\frac{\left(\kappa(f_{1|1})-\kappa(f_{0|1})\right)}{I(f_{1|1})-I(f_{0|0})}$ is the proportion of the gap attributed to the structural effect of dependence; $\frac{\left(\kappa(f_{0|1})-\kappa(f_{0|0})\right)}{I(f_{1|1})-I(f_{0|0})}$ is the proportion of the gap attributed to the composition effect of dependence; finally, $\frac{\left(\kappa(f_{1|1})-\kappa(f_{0|1})\right)+\left(\kappa(f_{0|1})-\kappa(f_{0|0})\right)}{I(f_{1|1})-I(f_{0|0})}$ is the proportion of the gap that is attributed overall to dependence. Similar effects can be computed for marginal distributions.

To sum up, using well-known multivariate measure of inequality such as (4), the gap in multiple outcomes can be easily analyzed. Furthermore, full decomposition that separates dependence from marginal distributions is possible. However, as noted by DiNardo et al. (1996), although using inequality measures for the measurement of the gap is very common, there is a drawback. Such measures provide little information on the different shapes of two distributions as they are defined for a given distribution. This issue is better addressed by divergence measures that take directly into account distributional differences.

3.2 Divergence measures

The most well-known divergence measure is the so called the Kullback-Leibler difference between two densities f and $\tilde{f}$

$$KL(f, \tilde{f}) = \int \left[f(y) - \tilde{f}(y)\right] \ln \frac{f(y)}{\tilde{f}(y)} \mathrm{d}y$$

Please note that in this definition y can be multivariate. This measure, however, is not decomposable.

Proposition 4. *KL divergence measure is not decomposable into structural and composition effects.*

Proof. We have

$$\left(\int \left[f_{1|1}(y) - f_{0|1}(y)\right] \ln \frac{f_{1|1}(y)}{f_{0|1}(y)} \mathrm{d}y\right) + \left(\int \left[f_{0|1}(y) - f_{0|0}(y)\right] \ln \frac{f_{0|1}(y)}{f_{0|0}(y)} \mathrm{d}y\right) =$$

$$\int f_{1|1}(y) \ln \frac{f_{1|1}(y)}{f_{0|1}(y)} - f_{0|1}(y) \ln \frac{f_{1|1}(y)}{f_{0|1}(y)} \mathrm{d}y +$$

$$\int f_{0|1}(y) \ln \frac{f_{0|1}(y)}{f_{0|0}(y)} - f_{0|0}(y) \ln \frac{f_{0|1}(y)}{f_{0|0}(y)} \mathrm{d}y \neq$$

$$\int f_{1|1}(y) \ln \frac{f_{1|1}(y)}{f_{0|0}(y)} - f_{0|0}(y) \ln \frac{f_{1|1}(y)}{f_{0|0}(y)} \mathrm{d}y = KL(f_{1|1}, f_{0|0}).$$

□

Kullback-Leibler is a member of *k-class entropy distance measures*

$$I_k(f_{1|1}, f_{0|0}) = \frac{1}{k-1} \left(\int \frac{f_{1|1}^k}{f_{0|0}^k} \mathrm{d}f_{0|0} - 1\right),$$

such that $\lim_{k\to\infty} I_k = I_1$, where I_1 is KL. It is evident that the same reasoning as in the proof of Proposition 4 will apply to the whole class. In fact, below we give a general result concerning decomposition of such measures. Let us first define what is the general form of decomposition for such measures.

Definition 5. *The decomposition of the gap into structural and composition effects using divergence statistics is the following*

$$\rho(f_{1|1}, f_{0|0}) = \rho(f_{1|1}, f_{0|1}) + \rho(f_{0|1}, f_{0|0}). \tag{6}$$

Please note that (6) is different from (3): ν is a functional of one distributions, whereas ρ is a functional of two distributions. Consequently, while the decomposition using ν obtains trivially, the decomposition using ρ holds only under a restrictive condition as Proposition 5 shows.

Proposition 5. *ρ is decomposable in the form* (6) *if and only if* $\rho(f_{1|1}, f_{0|0}) = g(f_{1|1}) - g(f_{0|0})$*, where g is some function.*

Proof. (6) is known as Sincov's functional equation (Marshall and Olkin 1979, p. 223) and as such has the solution described in Proposition 5 without any further assumptions. □

Additionally, if the decomposition takes a more specific but still very general form

$$\psi(f_{1|1} - f_{0|0}) = \psi(f_{1|1} - f_{0|1}) + \psi(f_{0|1} - f_{0|0}), \tag{7}$$

then ψ is a linear function (of then one variable).

Proposition 6. *ψ is decomposable in the form (7) if and only if $\psi(t) = at$, where a is some constant.*

Proof. Let $u = f_{1|1} - f_{0|1}$ and $v = f_{0|1} - f_{0|0}$, then $\psi(u) + \psi(t) = \psi(u + v)$, which is the well-known Cauchy functional equation that has solution $\psi(t) = at$ for some constant a. $\square$

We can see from Propositions 5 and 6 that the restrictions on the form of decomposition using two distributions are very restrictive. In order for the divergence measure to be decomposable into structural and composition effects it needs to be a simple difference of distributions evaluated by g, and in case the decomposition is specifically the decomposition of the difference in distributions it needs to be a linear function. The widely used entropy divergence measures such as members of k-class are not of these forms.

3.3 The exact decomposition of the distribution

We showed in Section 3.1. that a concrete family of measures (4) can be decomposed in a way that distinguishes the role of marginal distributions from the role of their association. Clearly, not every family of functions can be decomposed in this way. In this section, we study such decompositions more generally, namely, working directly with the decomposition of the distribution itself.

To this end, we utilize a famous result by Sklar (1959) who asserts that there exists a copula function such that the joint cdf can be written as $F(y_1, y_2) = C(F_1(y_1), F_2(y_2))$. $C : [0, 1]^2 \mapsto [0, 1]$ is the so called copula function. It can be regarded too as the joint cdf of two variables u and v that have standard uniform distribution. That is, $u = F_1(y_1)$ and $v = F_2(y_2)$ so that (u, v) can always be found. Copula density c is thus the density of the uniform random vector, $c(u, v) = \frac{\partial^2}{\partial u \partial v} C(u, v)$. Given Sklar's result, we get $f(y_1, y_2) = c(f_1, f_2) f_1(y_1) f_2(y_2)$. Thus we can write the following comparison of our distributions of interest

$$\frac{f_{1|1}}{f_{0|0}} = \frac{c_{1|1}}{c_{0|0}} \frac{f^1_{1|1}}{f^1_{0|0}} \frac{f^2_{1|1}}{f^2_{0|0}}. \tag{8}$$

Using logarithm, we are back to the additive gap, which now takes the form $\ln f_{1|1} - \ln f_{0|0}$. This gap can of course be decomposed.

Definition 6. *The decomposition of the gap into structural and composition effects using log density is the following*

$$\ln f_{1|1} - \ln f_{0|0} = \left(\ln f_{1|1} - \ln f_{0|1}\right) + \left(\ln f_{0|1} - \ln f_{0|0}\right). \tag{9}$$

The use of log density instead of density means representing probabilities on a logarithmic scale rather than on a standard unit interval. Log probabilities are used in many applications, such as natural language processing and machine learning, because they can make calculations involving very small numbers more stable and efficient. In particular, the logarithm of a probability can be used to transform the product of many probability values (such as the likelihood of a sequence of events) into the sum of their logarithms. This can be more computationally stable because the logarithm function "compresses" large numbers and "stretches out" small numbers, making it less likely to encounter underflow or overflow errors. The use of log probabilities is very common e.g., in statistical physics, where it is quite natural to switch back and forth between quantities proportional to log probabilities (energy, entropy, enthalpy, free energy) and quantities proportional to probability (number of microstates, partition function, density of states).

Altogether, we have the following result concerning the gap in log probabilities.

Proposition 7. *The gap* $\ln f_{1|1} - \ln f_{0|0}$ *can be decomposed as follows*

$$\begin{aligned}
&\ln f_{1|1} - \ln f_{0|0} = \\
&= \left(\ln f^1_{1|1} - \ln f^1_{0|1}\right) && \text{structural effect for outcome 1} \\
&+ \left(\ln f^2_{1|1} - \ln f^2_{0|1}\right) && \text{structural effect for outcome 1} \\
&+ \left(\ln c_{1|1} - \ln c_{0|1}\right) && \text{structural effect for dependence} \\
&+ \left(\ln f^1_{0|1} - \ln f^1_{0|0}\right) && \text{composition effect for outcome 1} \\
&+ \left(\ln f^2_{0|1} - \ln f^2_{0|0}\right) && \text{composition effect for outcome 2} \\
&+ \left(\ln c_{0|1} - \ln c_{0|0}\right) && \text{composition effect for dependence}
\end{aligned}$$

Proof. Coming back to (8) and taking log, we can see that the gap can be decomposed by attributes, namely

$$\ln f_{1|1} - \ln f_{0|0} = \left(\ln f^1_{1|1} - \ln f^1_{0|0}\right) + \left(\ln f^2_{1|1} - \ln f^2_{0|0}\right) + \left(\ln c_{1|1} - \ln c_{0|0}\right).$$

The rest follows from (9). □

Thus utilizing results from copula theory applied to densities, the exact decomposition using both forms of decomposition into counterfactual effects and into marginals and dependence, can be obtained. Applying this result to the data, the decomposition of the probability mass concentrated at some interesting points (e.g. at quartiles or at the median) of the distribution can be reported. Alternatively, we could present a summary statistics by taking the average of log probabilities. This is known as famous Shannon's (1948) entropy which for continuous distributions is called differential entropy

$$S(f) = \int \ln(f(x,y))f(x,y)\mathrm{d}x\mathrm{d}y.$$

We can also write entropy of marginal distributions f^1, f^2 and of copula density c thus arriving at the following representation (Calsaverini and Vicente 2009)

$$S(f) = S(f_1) + S(f_2) + S(c), \tag{10}$$

where

$$S(c) = \int_{[0,1]^2} f(y_1, y_2) \ln\left(\frac{f(y_1, y_2)}{f_1(y_1) f_2(y_2)}\right) \mathrm{d}y_1 \mathrm{d}y_2 \tag{11}$$

is the so-called *copula entropy* known also as *mutual information.* This is a very well-known concept in information theory. Mutual information is a measure of the amount of information that one random variable contains about another. It quantifies the amount of reduction in uncertainty about one variable given knowledge of the other variable. In other words, it measures the amount of dependence between two random variables, hence is can be expressed in the language of copulas. The mutual information is always non-negative and it is equal to zero if and only if the two variables are independent. Using entropy we can define the gap $S\left(f_{1|1}\right) - S\left(f_{0|0}\right)$.

Proposition 8. *The gap $S\left(f_{1|1}\right) - S\left(f_{0|0}\right)$ can be decomposed as follows*

$$\begin{aligned}
&S\left(f_{1|1}\right) - S\left(f_{0|0}\right) = \\
&= S\left(f^1_{1|1}\right) - S\left(f^1_{0|1}\right) && \text{structural effect for outcome 1} \\
&+ S\left(f^2_{1|1}\right) - S\left(f^2_{0|1}\right) && \text{structural effect for outcome 1} \\
&+ S\left(c_{1|1}\right) - S\left(c_{0|1}\right) && \text{structural effect for dependence} \\
&+ S\left(f^1_{0|1}\right) - S\left(f^1_{0|0}\right) && \text{composition effect for outcome 1} \\
&+ S\left(f^2_{0|1}\right) - S\left(f^2_{0|0}\right) && \text{composition effect for outcome 2} \\
&+ S\left(c_{0|1}\right) - S\left(c_{0|0}\right) && \text{composition effect for dependence}
\end{aligned}$$

Proof. It follows from (10). □

Let us give an example for Gaussian distributions, that is, let us assume that all distributions, i.e., men's, women's and counterfactual distribution, are normal distributions with their respective set of parameters. As it is known a bivariate normal distribution can be constructed by plugging in marginal Gaussian distributions $\Phi(x) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2} \mathrm{d}u$ into the Gaussian copula defined as

$$C(u, v) = \Phi_\rho\left(\Phi^{-1}(u), \Phi^{-1}(v)\right),$$

where $\Phi_\rho(x, y) = \int_{-\infty}^{x} \int_{-\infty}^{y} \frac{1}{\sqrt{4\pi^2(1-\rho^2)}} e^{-\frac{u^2+v^2-2uv\rho}{2(1-\rho^2)}} \mathrm{d}u\mathrm{d}v$ and ρ is a dependence parameter. The expressions for the logarithm of densities are the following:

$$\log f(x) = -\frac{1}{2}\log(2\pi\sigma^2) - \frac{(x-\mu)^2}{2\sigma^2}$$

for the marginal distributions,

$$\log f(x,y) = -\frac{1}{2}\log(2\pi) - \log(\sigma_x) - \log(\sigma_y) - \frac{1}{2}\log(1-\rho^2) - \frac{1}{2(1-\rho^2)}\left(\frac{(x-\mu_x)^2}{\sigma_x^2} + \frac{(y-\mu_y)^2}{\sigma_y^2} - \frac{2\rho(x-\mu_x)(y-\mu_y)}{\sigma_x\sigma_y}\right)$$

for bivariate distributions, and finally

$$\log c(u,v) = -\frac{1}{2}\log(2\pi) - \frac{1}{2}\log(1-\rho^2) - \frac{1}{2(1-\rho^2)}[\tilde{\Phi}^{-1}(u)^2 - 2\rho\tilde{\Phi}^{-1}(u)\tilde{\Phi}^{-1}(v) + \tilde{\Phi}^{-1}(v)^2]$$

for copula density, where $\tilde{\Phi}^{-1}$ is the inverse of the standard normal CDF. Using these expressions we can perform analytically decomposition prescribed in Proposition 7.

We can do the same for the entropy decomposition. The expressions for the entropy are the following:

$$S(f) = \frac{1}{2}\log(2\pi e\sigma^2)$$

for the entropy of the marginal normal distribution,

$$S(f) = \frac{1}{2}\log(2\pi e\sigma_1^2\sigma_2^2(1-\rho^2))$$

for the entropy of the joint normal distribution, and

$$S(c) = \frac{1}{2}\log(1-\rho^2)$$

for the copula entropy or mutual information. Using these expressions we can perform decomposition prescribed in Proposition 8. One thing to note: in our case, since both marginals and copula are Gaussian, ρ is measured by the standard linear correlation coefficient. Then, $S(c)$ gives minimal mutual information contained in the dependence between variables. If marginals are non-Gaussian, then linear correlation estimate will underestimate dependence parameter ρ and information excess will appear: estimated mutual information will be higher than $\frac{1}{2}\log(1-\rho^2)$. Thus this easily computable estimate can serve as a detection tool for non-Gaussian dependencies in the data and together with some marginal-invariant dependence measure such as e.g. Kendall's tau, can be used to identify the proper copula to represent these dependencies. This method of finding the best fit distribution and copula, namely, via entropy decomposition and information content matching, is equivalent to the usual maximum likelihood methods (Fernandez 2008).

3.4 The case of ordinal outcomes

Some of the methods presented so far are best suited for cardinal ratio-scale variables. On the other hand, dimensions of well-being are often ordinal indicators, such as for example, self-reported health status, or educational attainment, or socioeconomic status, or occupational class, and many others. Then the well-known scaling problems arise (see e.g. Allison and

Foster 2004). Inequality measures such as (4) cannot be used as there is no natural scale that could be assigned to y_1, y_2. Neither would quantile parametrization work and there would also be problems with factorization of joint probability mass function into marginal probability and copula, because discrete copulas are not unique (Kobus and Kurek 2018). Taking into account these problems, we turn to the literature on inequality measurement for multivariate ordinal data to provide tools that can be further utilized to measure the gap. The literature is scarce at the moment, but Kobus and Kurek (2023) offer some solutions. They propose dominance ordering and measures that are sensitive to increased spread and bipolarity of the multidimensional ordinal variable, which are two properties necessary for the measurement of inequality in multidimensional ordinal data. Their proposed dominance relation states that a given distribution is 'more unequal than the other if it has greater partial sums of its cumulative distribution function below the multidimensional median and greater partial sums of its survival function above the multidimensional median. They develop a simple inequality measure based on this relation and show that this measure fulfills desired properties:

$$\mathcal{P}_{a,b}(\mathbb{P}, \mathbb{m}) = \frac{\Sigma_{\mathbb{i} \prec \mathbb{m}} a_{\mathbb{i}} \mathbb{P}(\mathbb{i}) + \Sigma_{\mathbb{i} \succeq \mathbb{m}} b_{\mathbb{i}} \bar{\mathbb{P}}(\mathbb{i})}{C}$$

where $a_{\mathbb{i}} = a\Pi_{j=1}^{k}(n_j - m_j + i_j)$, $b_{\mathbb{i}} = b\Pi_{j=1}^{k}(n_j + m_j - i_j)$, $C = \frac{a}{2}\Sigma_{\mathbb{i} \prec \mathbb{m}} \Pi_{j=1}^{k}(n_j - m_j + i_j) + \frac{b}{2}(\Sigma_{\mathbb{m} \preceq \mathbb{i} \prec \mathbb{n}} \Pi_{j=1}^{k}(n_j + m_j - i_j) - 1)$. $\mathbb{P}, \bar{\mathbb{P}}$ denote the values of, respectively, cumulative distribution function and survival function. Parameters a, b allow for differential treatment of inequality below and above the median, i.e., if $a > b$, then more weight is attached to inequality in the lower tail of the distribution and the reverse holds for $a < b$. Weights $a_{\mathbb{i}}$ attached to categories below the median are increasing and above the median weights $b_{\mathbb{i}}$ are decreasing to ensure consistency with the proposed dominance ordering. Both weights and the denominator C have been chosen to ensure that the index is normalized between [0,1].[7] As an additive measure $\mathcal{P}_{a,b}$ allows for the decomposition by sources:

$$\mathcal{P}_{a,b}(\mathbb{P}_{1|1}) - \mathcal{P}_{a,b}(\mathbb{P}_{0|0}) = \left(\mathcal{P}_{a,b}(\mathbb{P}_{1|1}) - \mathcal{P}_{a,b}(\mathbb{P}_{0|1})\right) + \left(\mathcal{P}_{a,b}(\mathbb{P}_{0|1}) - \mathcal{P}_{a,b}(\mathbb{P}_{0|0})\right).$$

The decomposition is not hundred percent formally correct, because we omitted the dependence of $\mathcal{P}$ on $\mathbb{m}$, the multidimensional median, and this would potentially be problematic if the median differ between observed and counterfactual distributions. However, Kobus and Kurek (2023) develop methods to precisely deal with this problem and that allow for comparisons of distributions with potentially different medians. Thus, measuring gaps in multivariate ordinal variables can be performed by utilizing their results rather than choosing arbitrary cardinalization of ordinal indicators. The shortcoming is that there is lack of attribute decomposability; this remains an open problem for ordinal measures.

[7] Parameters $a_{\mathbb{i}}$ and $b_{\mathbb{i}}$ are easy to interpret; they are a multiplication of dimensional weights that would hold on each dimension if there was only a single dimension.

4 Quantile effects

As already mentioned, gaps are typically reported using some distributional statistic or as differences between quantiles. We will now turn to the latter and analyze quantile gaps in a multivariate setting.

First we need to define counterfactual effects in terms of related quantiles. In a unidimensional setting, with distribution treatment effects, we say that the probability that women earn x income is higher by $\mathrm{d}\tau$ had they had the distribution of men's characteristics. With quantile treatment effects, we say that with probability τ women earn x more income had they had the distribution of men's characteristics. The definition of the counterfactual quantile is the same for both univariate and multivariate settings.

Definition 7. *The counterfactual quantile is defined as*

$$Q_{0|1}(\tau) = F^{-1}_{Y_0|X_1}(\tau) = \{y \in \overline{\mathcal{Y}_0} : F_{Y_0|X_1}(y) < \tau\},\ \tau \in (0,1) \tag{12}$$

The counterfactual effect is then $Q_{1|1}(\tau) - Q_{0|1}(\tau)$, but given that in $\mathbb{R}^2$ quantiles are not numbers but sets, this difference is in fact the difference between two sets and as such is not meaningful. With quantiles, when going from unidimensional to multidimensional setting there are problems already at the level of the very definition of the gap. We propose two solutions to this problem. The first solution is to find a reasonable representative of the quantile set that also has other properties e.g. decomposition. The second solution is the parametrisation of quantile sets. This is the approach developed by Maasoumi and Racine (2016) to study multidimensional well-being, which we utilize for multidimensional gaps.

4.1 Distance between quantile isoquants

In the mathematics literature there are definitions of the difference between two subsets of a metric space. The most well-known is the Hausdorff distance metric which is the greatest of all the distances from a point in one set to the closest point in the other set. It is a proper metric and as such it will not be decomposable into structural and composition effect. The triangle inequality which it fulfils as a metric requires that the distance between one quantile curve and the other is not necessarily equal but it can be less than the distance that takes into account also the counterfactual quantile curve. Thus, we propose the following simple alternative, namely, the representative of a given quantile isoquant will be the vector of means.: the mean value of the first and the mean value of the second outcome for a given quantile. We call a gap defined using those means as *the mean difference quantile gap*. For a given τ, let us define intermediate notions such as $Q_1(y_2) = y_1$ and $Q_2(y_1) = y_2$, which denote the quantile as a function of one dimension that for a given value of this dimension returns the value of the second dimension.[8] Now, the the gap can be defined as follows.

[8]Here we implicitely assume that the quantile curve is a bijection. If it is not, then it can be defined using infimum how the value of $y_1(y_2)$ is chosen, but we will not delve deeper into those technical details.

Definition 8. *The mean difference quantile gap is defined as*

$$\begin{aligned} \Delta_1(Q_{Y_{1|1}}, Q_{Y_{0|0}})(\tau) &= \frac{\int Q_{1,Y_{1|1}}(y_2) - Q_{1,Y_{0|0}}(y_2)\mathrm{d}y_2}{y_2^{MAX}(\tau) - y_2^{MIN}(\tau)} \\ \Delta_2(Q_{Y_{1|1}}, Q_{Y_{0|0}})(\tau) &= \frac{\int Q_{2,Y_{1|1}}(y_1) - Q_{2,Y_{0|0}}(y_1)\mathrm{d}y_1}{y_1^{MAX}(\tau) - y_1^{MIN}(\tau)} \end{aligned} \tag{13}$$

Although the definition may seem technical because of the difficulty of having quantiles as sets, it is a very simple concept. For a given value of τ, we summarize the whole quantile set by the mean of the first and the mean of the second dimension. For example, let $\tau = 0.2$. The quantile isoquant for this value of τ is summarized by mean life expectancy and mean income, if these are our outcomes.[9]

Proposition 9. *The mean difference quantile gap is decomposable into structural and composition effects.*

Proof. We have

$$\begin{aligned} &\Delta_1(Q_{Y_{1|1}}, Q_{Y_{0|0}})(\tau) = \\ &\frac{\int Q_{1,Y_{1|1}}(y_2) - Q_{1,Y_{0|0}}(y_2)\mathrm{d}y_2}{y_2^{MAX}(\tau) - y_2^{MIN}(\tau)} = \\ &\frac{\int Q_{1,Y_{1|1}}(y_2) - Q_{1,Y_{0|1}}(y_2) + Q_{1,Y_{0|1}}(y_2) - Q_{1,Y_{0|0}}(y_2)\mathrm{d}y_2}{y_2^{MAX}(\tau) - y_2^{MIN}(\tau)} = \\ &\frac{\int Q_{1,Y_{1|1}}(y_2) - Q_{1,Y_{0|1}}(y_2)\mathrm{d}y_2}{y_2^{MAX}(\tau) - y_2^{MIN}(\tau)} + \frac{\int Q_{1,Y_{0|1}}(y_2) - Q_{1,Y_{0|0}}(y_2)\mathrm{d}y_2}{y_2^{MAX}(\tau) - y_2^{MIN}(\tau)} = \\ &= \Delta_1(Q_{Y_{1|1}}, Q_{Y_{0|1}})(\tau) + \Delta_1(Q_{Y_{0|1}}, Q_{Y_{0|0}})(\tau) \end{aligned}$$

and the same for $\Delta_2(Q_{Y_{1|1}}, Q_{Y_{0|0}})(\tau)$. □

Please note that we use the same denominator $y_2^{MAX}(\tau) - y_2^{MIN}(\tau)$ for all three curves, namely, for $Q_{Y_{1|1}}(\tau), Q_{Y_{0|0}}(\tau)$ and $Q_{Y_{0|1}}(\tau)$. In practice, with real world data, the range of y_2 for these three distributions may be different. In this case, we choose the maximum $y_2^{MAX}(\tau) - y_2^{MIN}(\tau)$. Furthermore, please note that $y_2^{MAX}(\tau) - y_2^{MIN}(\tau)$ changes with each τ. It could be easier to just choose one global interval $y_2^{MAX} - y_2^{MIN}$ that is the same for every τ, however, if these values vary substantially, we may not reflect the true mean difference for a given τ, therefore we choose to have $y_2^{MAX} - y_2^{MIN}$ as a function of τ.

Although we found the definition of the distance between two quantile sets that can be decomposed into structural and composition effect, it is not possible to fully separate out dependence. There is still dependence in the vector of means that summarize each quantile set, however, we neglected large part of dependence by averaging out the quantile sets. This section is more to show that because in our framework are more complicated objects than in

[9] Given the data, $y^M AX_1(y^M AX_2)$ always exist (and similarly for MIN). In theory, they may not exist, however, the assumption of the bounded support of the cdf ensures that they do.

the unidimensional framework, there is no easy way to redefine their difference and preserve the properties that are necessary for the exercise that we have.

4.2 Parametrization of quantile isoquants

Quantile sets can be aggregated by an aggregation function that expresses the latent concept of well-being. The fundamental problem is then to obtain credible weights of such aggregation functions. We follow Maasoumi and Racine (2016) who fit a desired parametric form to non-parametric estimates of such iso-quants and discover the unknown substitution parameters. Multidimensional quantiles are equi-probable surfaces that correspond to the values of aggregation functions, parameters of which we seek to discover. Their proposed aggregation functions were first derived in Maasoumi (1986) as ideal aggregators in the sense of minimising the divergence between its distribution and the distribution of its constituent components. Without loss of generality, for two outcomes, the CES aggregation function is given by

$$S(y_1, y_2) = A(\alpha y_1^{-\beta} + (1-\alpha)y_2^{-\beta})^{\frac{-1}{\beta}}$$

where $A > 0$ and $0 < \alpha < 1$. The partial derivatives are as follows,

$$S_{y_1} = \frac{\partial S(y_1, y_2)}{\partial y_1} = A\alpha(\alpha y_1^{-\beta} + (1-\alpha)y_2^{-\beta})^{(\frac{-1}{\beta})-1} y_1^{-\beta-1}$$

$$S_{y_2} = \frac{\partial S(y_1, y_2)}{\partial y_2} = A(1-\alpha)(\alpha y_1^{-\beta} + (1-\alpha)y_2^{-\beta})^{(\frac{-1}{\beta})-1} y_2^{-\beta-1}$$

Along an iso-well-being quantile we have $\Delta S = 0$ (i.e. $S_{y_1}\partial y_1 + S_{y_2}\partial y_2 = 0$), therefore

$$-\frac{S_{y_1}}{S_{y_2}} = \frac{\partial y_2}{\partial y_1} = \frac{\alpha}{\alpha - 1}\left(\frac{y_2}{y_1}\right)^{\beta+1} \tag{14}$$

We exploit the fact that, for $y = (y_1, y_2)$, conditional on x, we can obtain estimates of $\frac{\partial y_2}{\partial y_1}$ directly from the estimated quantile $\widehat{Q}_{Y_j|X_k}(\tau)$ (i.e. for a given value of τ we can compute $\frac{\partial y_2}{\partial y_1}$ since the level of multidimensional well-being is constant). The estimates of α and β are then obtained via (nonlinear) regression of our non-parametrically estimated $\frac{\partial y_2}{\partial y_1}$ on $\frac{y_2}{y_1}$ using (14). The values of α and β may vary with the quantile. This means that different dimensions of well-being contribute differently to well-being at each level of well-being. This accounts for heterogeneity in preferences.

Coming back to counterfactual effects, the S parametrization is related to the distribution via quantile sets $Q_{0|1}(\tau) = \{y \in \mathbb{R}^2 : f_{0|1}(y) = \tau\}$ and $Q_{1|1}(\tau) = \{y' \in \mathbb{R}^2 : f_{1|1}(y') = \tau\}$. We use the following notation: $y = \left(y_1^{0|1}(\tau), y_2^{0|}(\tau)\right)$ and $y' = \left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)$. For the S function we put $S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)$ and $S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)$. That is, $S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)$ is the value of function S that is associated with distribution $f_{0|1}$ for a given value τ of the quantile. It is important to highlight that S is ssociated with $f_{0|1}$ (by putting $S_{0|1}$), because $S_{1|1}$ will have a different set of parameters (substitution weights), so it will be a different function than $S_{0|1}$.

We can now define the gap and the counterfactual effects using S parametrization. It might be more intuitive to talk about shares than raw S scores. In what follows we assume the following definition of the gap.

Definition 9. *The well-being gap using S parametrization is defined as*

$$\frac{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{0|0}\left(y_1^{0|0}(\tau), y_2^{0|0}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)}. \tag{15}$$

Such definition of the gap for wages, namely the difference in average wages for men and women divided by the average wages for women, was used by Ñopo (2008), who believes that this better corresponds to the concept of the gender wage gap than simple difference. Definition 9 is the well-being gap between men and women defined as the difference in well-being between men and women expressed as a share of women's well-being at every point of the distribution of well-being.

Definition 10. *The well-being gap using S parametrization is decomposable into structural and composition effects*

$$\begin{aligned} &\frac{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{0|0}\left(y_1^{0|0}(\tau), y_2^{0|0}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)} = \\ &\frac{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)} + \\ &\frac{S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right) - S_{0|0}\left(y_1^{0|0}(\tau), y_2^{0|0}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)}. \end{aligned} \tag{16}$$

That is, at any given point τ, the first difference in (16) is the difference in the well-being of women and counterfactual women who are given men's conditional distribution of well-being dimensions. This is the structural effect in well-being. The second difference is the difference in well-being between counterfactual women and men who differ by their characteristics but have access to the same conditional distribution of well-being dimensions. This is the composition effect in well-being. At each point of the distribution we compare whether men's, women's, or counterfactual women's share of women's well-being is greater. Note that both effects are a difference of two numbers, because all S functions are defined exactly so that they are constant along the isoquant determined by τ in their respective distributions.

With S parametrization we can construct new counterfactual effects by further decomposing the structural effect.

Definition 11. *Let us consider the structural effect in* (16). *It can be decomposed it as follows*

$$\begin{aligned}
&\frac{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)} = \\
&\frac{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{1|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)} + \\
&\frac{S_{1|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right) - S_{0|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)}{S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right)}
\end{aligned} \tag{17}$$

We notice the presence of a new term, namely, $S_{1|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)$. This is the parametrization of men's conditional distribution of well-being dimensions, but using women's preferences i.e., substitution parameters α, β from $S^{1|1}$. These newly created counterfactual women keep their preferences towards two outcomes butas a counterfactual exercise they are given access to men's structure of these outcomes. Thus, the first difference in (17) is the comparison of well-being of women with these hypothetical women and reflects the change in well-being that corresponds to changing the conditional distribution of outcomes, but using women's evaluation function of these outcomes. The second difference in (17) concerns the change of substitution parameters and thus the impact of this isolated change only. Please note that $S_{1|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)$ is not constant along the iso-quant determined by τ, because the parameters are not suited for it to be constant. Therefore, $S_{1|1}\left(y_1^{1|1}(\tau), y_2^{1|1}(\tau)\right) - S_{1|1}\left(y_1^{0|1}(\tau), y_2^{0|1}(\tau)\right)$ is a vector of numbers, where the first expression is constant, but a second expression changes. For the purposes of the ease of exposition, we did not want to highlight this while introducing the new counterfactual effect, but in applications one should take the average of these differences over data points related to a given τ.

Distributions S are unidimensional objects, thus such defined well-being gaps can be studied in the same way as wage gaps are studied. One can report them at some quantiles, or compare their distributions using standard tools such as stochastic dominance or unidimensional inequality measures. These well-being shares can act as an independent variable in a regression framework and permit detailed decomposition.

5 Empirical application

We apply the above described tools to analyse wage-leisure gap in the US using American Time Use Survey data. Wage is used typically in gap analysis, but as to leisure, according to ATUS, fathers, on average, have about three hours more leisure time per week than mothers. The disparity in leisure time between mothers and fathers has been a consistent trend for at least the past decade. This gap is particularly evident when it comes to time spent in front of the television or using other forms of media. According to the American Time Use Survey (ATUS), fathers spend an average of 2.8 hours more each week than mothers watching TV or using other media. Additionally, fathers tend to spend more time engaging in sports or exercising, while mothers spend more of their leisure time participating in social activities, such as attending or hosting parties.

The ATUS not only measures how individuals spend their time, but also asks about their feelings and attitudes towards different activities. The survey found that mothers tend to find their leisure time to be more meaningful than fathers. 63% of mothers rated their leisure activities as "very meaningful," compared to only 52% of fathers. Despite this, mothers also reported feeling more exhausted and stressed during their leisure time. This discrepancy in mothers' leisure experiences may be related to the way they experience their time. Mothers' free time is often interrupted, making it difficult for them to fully relax. Additionally, mothers tend to spend more time multitasking, with the majority of this additional time spent on household and child care responsibilities.

The variables used are: log of weekly wages; log of weekly leisure time (in minutes), where leisure is understood in its narrow definition as time spent on activities described in category 12th of ATUS, namely, socializing, relaxing and leisure[10]; sex where 1 denotes male and 0 female; years of education; number of children; marital status, where 1 means married or stable partnership; log of minutes of housework (weekly); race, where 1 means white, 2 means black and 3 means other race; residence, where 1 means metropolitan area and 0 means non-metropolitan area and age reported in the following intervals: $<$25, 26-35, 36-45, 46-55, 56-65, 66+. Tables 1 and 2 contain summary statistics for these variables.

Table 1: Summary statistics 1: years 2005 and 2021

	2005				2021			
	Mean	Median	Min	Max	Mean	Median	Min	Max
Log wage (weekly) (k$)	10.92	11.00	00.00	12.57	11.33	11.45	1.09	12.57
Log leisure (weekly)	5.24	5.34	1.38	7.20	5.29	5.41	0.00	7.11
Sex	0.458	0.00	0.00	1.00	0.475	0.00	0.00	1.00
Education (years)	13.73	14.00	8.00	20.00	14.61	14.00	8.00	20.00
No. of children	0.99	0.00	0.00	3.00	0.73	0.00	0.00	3.00
Marital status	0.58	1.00	0.00	1.00	0.56	1.00	0.00	1.00
Log of housework (weekly)	4.97	4.60	2.30	4.82	4.96	4.49	2.70	6.76

Note: Data come from American Time Use Survey, years 2005 and 2021. The description of the variables is in the text above.

Figure 1 contains gaps as measured by the difference in inequality measure I. From this

[10]https://www.bls.gov/tus/lexicons/lexiconnoex2005.pdf.

Table 2: Summary statistics 2: years 2005 and 2021

Age	2005	2011	Race	2005	2021	Residence	2005	2021
>25	0.12	0.09	White	0.83	0.79	Metropolitan	0.81	0.87
26-35	0.22	0.23	Black	0.11	0.12	Non-metropolitan	0.19	0.13
36-45	0.28	0.23	Other	0.05	0.08			
46-55	0.22	0.19						
56-65	0.10	0.17						
>66	0.02	0.07						

Note: Data come from American Time Use Survey, years 2005 and 2021. The description of the variables is in the text above.

we can a general decreasing pattern of the gap. Female's inequality is higher than men, but these differences are smaller over time. Decomposing this trend, we observe the fastest decrease of the gap for the wage dimension, and for the leisure dimension, after initial decrease in the first decade of 2000s, the gap remains fairly stable. Altogether, the two dimensions contribute to a decrease in the overall gap (first picture), but a more moderate decrease than when considering only wage dimension (second picture). The dependence gap is also negative, meaning that there is less dependence of wage and leisure among women than men. The dependence gap is increasing (in absolute value).

A more detailed decomposition of the gap is provided in Table 3. Here we have an exact decomposition into all seven components, together with the residual from linear approximation, computed so far for years 2005-2015. In general we note that in each year case structural effects are greater than composition effects. This is in line with the finding of Maasoumi and Wang (2019) for wages. This is true for wages, leisure and also for dependence. We also note that starting in 2008, composition effect of wages is negative, meaning that it would hurt women more if their distribution of characteristics was replaced by men's. This again is in line with the finding of Maasoumi and Wang (2019) for wages, and it indicates that women have surpassed men in terms of the distribution of observable characteristics. The residual is not more than 1.2%. The composition effects for leisure and dependence require further inspection.

The evolution of the entropy gap (Figure 2) confirms some of the behaviour of the inequality gap. The entropy gap between men and women has decreased but more so in the first decade and remains fairly constant since then. Unidimensional entropies exhibit similar pattern. Dependence gap is not much different from zero which is different than the dependence gap measured by the difference in κ's (Figure 1). The detailed entropy decomposition results have not been yet computed.

As to the quantile parametrization, we report weight and substitution parameters for

Figure 1: Inequality measure gaps

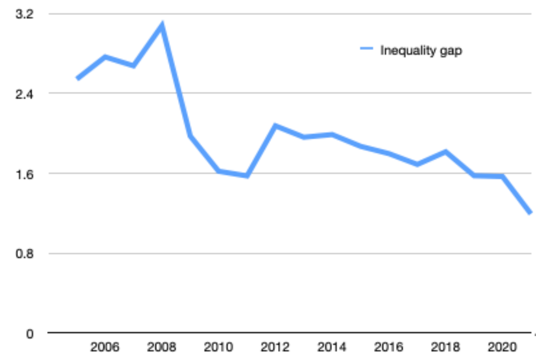

(a) Total inequality

(b) Wage gap

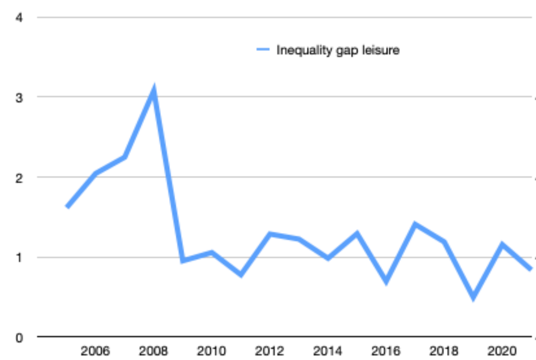

(c) Leisure gap

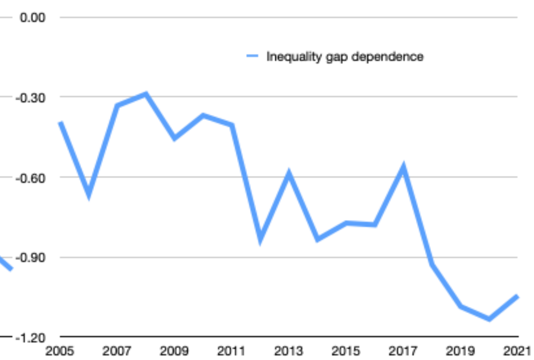

(d) Dependence gap

Table 3: Decomposition of inequality gap (total inequality gap =1)

	S wages	S leisure	S dependence	C wages	C leisure	C dependence	Residual
2005	0.384	0.368	0.128	0.153	-0.050	0.025	0.011
2006	0.380	0.364	0.130	0.158	-0.051	0.027	0.009
2007	0.383	0.381	0.135	0.064	0.038	-0.011	-0.007
2008	0.535	0.250	0.061	-0.126	0.249	0.032	0.002
2009	0.515	0.200	0.131	-0.120	0.231	0.032	0.002
2010	0.433	0.123	0.345	-0.004	0.202	-0.117	-0.008
2011	0.834	0.271	0.423	-0.333	-0.023	-0.165	0.006
2012	0.326	0.405	0.374	-0.050	-0.095	0.026	-0.012
2013	0.508	0.193	0.353	-0.021	0.118	-0.156	-0.003
2014	0.490	0.220	0.410	-0.091	0.007	-0.023	-0.003
2015	0.481	0.454	0.551	-0.227	-0.108	-0.138	0.012

Note: Data come from American Time Use Survey, years 2005-2021. S stands for “structural” effect and C for “composition“ effect.

Figure 2: Entropy measure gaps

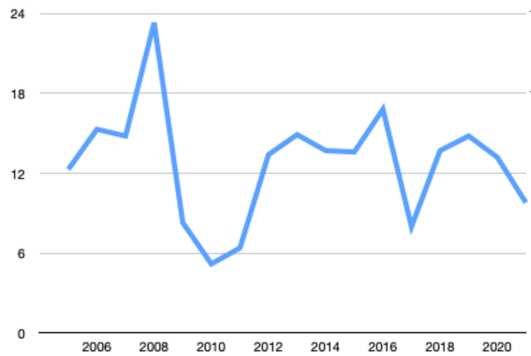

(a) Total entropy gap

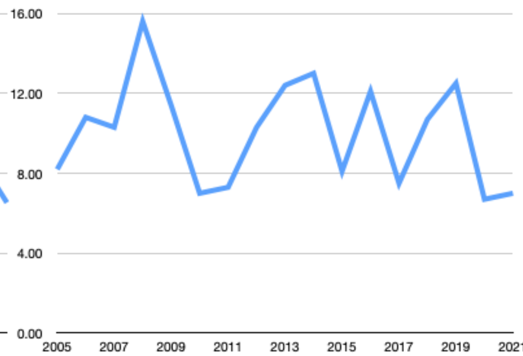

(b) Entropy gap wages

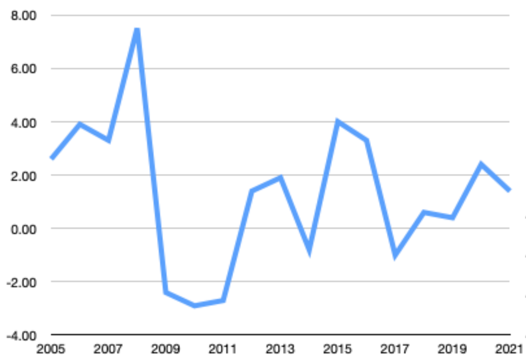

(c) Entropy gap leisure

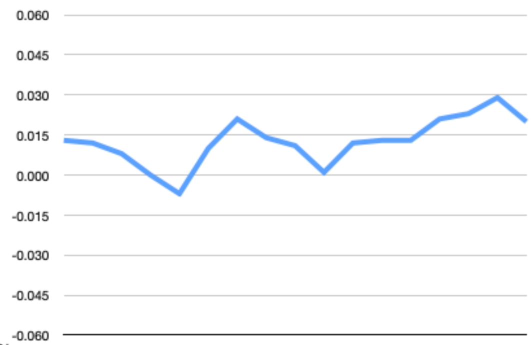

(d) Entropy gap dependence

years 2005 and 2021 (Table 4). We draw the following observations: (i) There is substantial heterogeneity with varying quantiles, which confirms reservations that applying homogeneuos evaluation functions may mask real trade-offs. (ii) Weight parameters are similar for men and women in both years. (iii) Both women and men have higher substitution parameters in 2021 than in 2005 at each quantile. (iv) Women have higher substitution parameters than men at each quantile in both years, except for women at the top decile in 2021. Please note that higher substitution parameters mean lower elasticity of substitution. Goods are then complements rather than substitutes. The difference between men and women in this respect should be contrasted with the analysis of what happens to home production (housework, childcare etc.) at different quantiles, which ATUS data enable. This is left for further inspection.

Not surprisingly, substitution dominates at lower quantiles. Women have lower levels of substitution but these differences are narrowing, in fact, at each quantile differences in substitution parameters between men and women are lower in 2021 than in 2005.

Table 4: Parameters' values for different quantiles: years 2005 and 2021

2005				
τ	α_M	β_M	α_F	β_F
0.2	0.84	6.00	0.85	6.78
0.3	0.86	7.54	0.87	8.83
0.4	0.88	9.16	0.89	10.57
0.5	0.89	10.72	0.90	12.77
0.6	0.90	12.65	0.91	14.98
0.7	0.91	14.88	0.92	17.85
0.7	0.93	18.67	0.93	21.13
0.9	0.94	24.71	0.94	25.51
2021				
τ	α_M	β_M	α_F	β_F
0.2	0.82	4.93	0.84	5.81
0.3	0.85	6.72	0.86	7.66
0.4	0.87	8.15	0.88	9.41
0.5	0.88	9.81	0.89	10.93
0.6	0.89	11.50	0.90	12.62
0.7	0.90	13.34	0.91	14.63
0.7	0.92	16.06	0.92	16.98
0.9	0.94	23.52	0.93	21.66

Note: Data come from American Time Use Survey, years 2005-2021. Subscripts M, F denote, respectively, male and female.

6 Conclusions

In this paper we tackle the problem of measuring gaps and counterfactual effects, but when the outcome is multidimensional. This concerns for example estimation of well-known gender gaps, but in the case when we compare not only men's and women's wages, but also for example their life expectancy, or any other outcome of interest. Therefore, we call these effects well-being gaps.

The multidimensional framework requires new definitions of well-known counterfactual effects. For distribution effects, the definitions are the same, but new counterfactual effects can be proposed given the multidimensional nature of outcomes. The problem that appears in a multidimensional context that is absent from a unidimensional framework concerns quantile effects. In our context, quantiles are sets, and there appears a problem of how to define the difference between two sets. Another issue that appears only in the multidimensional framework is outcomes' dependence and its role in the differences between considered groups.

As to possible extensions of our framework, let us discuss two extensions that naturally come to mind. First is the extension that allows for comparisons of more than two groups. For example, with racial gaps, it could be the comparison of Whites, Blacks and Hispanic, and other groups. This extension is straightforward and requires only new notation. Let the populations be labeled $k \in \mathcal{K}$ with $|K|$ the number of populations, and for each population k, there is a random d_x-vector X_k of covariates and a random outcome d_y-vector Y_k. We denote the support of X_k by $\mathcal{X}_k \subseteq \mathbb{R}^{d_x}$ and denote the region of interest for Y_j by $\mathcal{Y}_j \subseteq \mathbb{R}^{d_y}$. The counterfactual distribution and quantile functions are constructed by integrating the conditional distribution $F_{Y_j|X_j}$ in population j with respect to the covariate distribution F_{X_k} in population k, namely

$$F_{Y<j|k}(y) := \int_{\mathcal{X}_k} F_{Y_j|X_j}(y|x)\mathrm{d}F_{X_k}(x),\ y \in \mathcal{Y}_j$$

We define a counterfactual effect as the result of a shift from one counterfactual distribution $F_{Y<l|m}$ to another $F_{Y<j|k}$ for some $j, l \in \mathcal{J}$ and $m, k \in \mathcal{K}$

$$\Delta(y) = F_{Y<j|k}(y) - F_{Y<l|m}(y).$$

The rest of the analysis follows using such defined objects.

Another extension is an extension to more than two outcomes. Such model can be written theoretically by taking Y_l, a member of $\mathcal{Y}_l \subseteq \mathbb{R}^l$, where $l > 2$. Multivariate inequality measures are defined for such distributions and their attribute decomposition will have the following form

$$I(F) = \frac{r_1}{\sum_{i=1}^{l} r_i} I_1(F_1) + \ldots + \frac{r_l}{\sum_{i=1}^{l} r_i} I_l(F_l) + \frac{1}{\sum_{i=1}^{l} r_i}\kappa^l + R$$

where I_j are still unidimensional Atkinson indices and

$$\kappa^l = 1 - \frac{\int y_1^{r_1} \ldots y_l^{r_l} \mathrm{d}F(y_1, \ldots, y_l)}{\int y_1^{r_1}\mathrm{d}F_1(y_1) \ldots \int y_l^{r_l}\mathrm{d}F_l(y_l)}.$$

The copula density can also be defined for random vectors of dimensionality more than two. Both the logarithm decomposition and entropy decomposition can be computed. However, as to the quantile effects and parametrization, the generalization includes conditions on higher order derivatives in order for parameters to be recovered. This may become uninterpretable for many dimensions, so there seems to be a practical (not theoretical) upper limit on the number of dimensions to be used, but so is the case in practical usages of, e.g., stochastic dominance.

Finally, the methods presented here allow for aggregate decomposition. A detailed decomposition using I or entropy, or S parametrization or any other distributional statistic could be obtained via Shapley decomposition based on Shorrocks (2013). According to

Shapley value decomposition, the total contribution of a characteristic/factor is defined as the expected marginal decrease of the considered distributional statistic caused by the elimination of this factor, when the expectation is taken over all possible combinations of the elimination of this factor. The decomposition has the added value of being path independent.[11] If not Shorrocks decomposition which uses a particular notion of decomposition, namely Shapley value, then the problem of detailed decomposition in many dimensions remains open for future research.[12]

Acknowledgement

This research was funded by National Science Centre, Poland, grant number 2017/26/M/HS4/00905.

References

Aaberge R, Brandolini A. 2015. Multidimensional Poverty and Inequality, in: A. B. Atkinson, F. Bourguignon (eds.), Handbook of Income Distribution 2:141-216.

Abul Naga, R. H. and P.-Y. Geoffard. 2006. Decomposition of bivariate inequality indices by attributes. Economics Letters 90 (3):362–367.

Aguiar, M., Hurst, E. 2005. Consumption versus expenditure. Journal of Political Economy, 113, 919–948.

Albrecht J, Bjorklund A, Vroman S. 2003. In there a glass ceiling in Sweden?, Journal of Labor Economics 21: 145-177.

Aliaga, C. (2006). "How is the time of women and men distributed in Europe?" Statistics in Focus, Population and Social Conditions 4/2006. EUROSTAT.

Allison R A, Foster J E. 2004. Measuring health inequality using qualitative data, Journal of Health Economics 23(3):505-524.

Arellano M, Bonhomme S. 2017. Quantile Selection Models: with an application to understanding changes in wage inequality, Econometrica 85: 1-28.

Aristei D, Perugini C. 2016. Inequality aversion in post-communist countries in the years of the crisis, Post-Communist Economies 28 (4), pp. 436–448.

Atkinson A, Bourguignon F. 1982. The comparison of multi-dimensioned distributions of economic status, Review of Economic Studies 49: 183-201.

Atkinson A, Bourguignon F. 1987. Income distribution and differences in needs. In: Feiwel G. (Ed.), Arrow and the Foundations of the Theory of Economic Policy, Macmillan.

[11]For example, Devicienti (2010) describes how to use Shorrocks decomposition in regression-based wage decomposition. Although, this approach is well-known in inequality measurement literature, one should note that it has not become popular in the empirical literature on gender gaps.

[12]But also for aggregate decomposition, there are alternative approaches than presented here, for example, through the use of simultaneous equation models. We, however, wanted to avoid a parametric structure.

Blau F. 1998. Trends in the well-being of American women in 1970-1995, Journal of Economic Literature 36: 112-165.

Blau F D, Kahn L M. 1997. Swimming Upstream: Trends in the Gender Wage Differential in the 1980s, Journal of Labor Economics 15: 1-42.

Blau F D, Kahn L M. 2006. The U.S. Gender Pay Gap in the 1990s: slowing convergence, Industrial and Labor Relations Review 60: 45-66.

Blau F D, Kahn L M. 2017. The Gender Wage Gap: Extent, Trends, and Explanations, Journal of Economic Literature 55(3): 789-865.

Blinder, A. S., "Wage Discrimination: Reduced Form and Structural Estimates." The Journal of Human Resources 8, 436–55, 1973.

Calsaverini R S, Vicente R. 2009. An information-theoretic approach to statistical dependence: Copula information, Europhysics Letters 88(6):68003

Campaña, J.C., Gimenez-Nadal, J.I. Velilla, J. 2023. Measuring Gender Gaps in Time Allocation in Europe. Soc Indic Res 165, 519–553. https://doi.org/10.1007/s11205-022-03026-0

Chernozhukow V, Fernández-Val I, Melly B. 2013. Inference on counterfactual distributions, Econometrica 81(6): 2205-2268.

Chernozhukov V, Imbens G, Newey W K. 2007. Instrumental variable estimation of non-separable models, Journal of Econometrics 139(1): 4-14.

Chernozhukov V, Hansen C. 2005. An IV model of quantile treatment effects, Econometrica 73(1): 245-261.

Chesher A. 2005. Nonparametric identification under discrete variation, Econometrica 73(5): 1525-1550.

Deaton A. 2015. The Great Escape: Health, Wealth, and the Origins of Inequality, Princeton University Press.

Devicienti F. 2010. Shapley-value decompositions of changes in wage distributions: a note. J Econ Inequal 8: 35–45.

DiNardo J, Fortín N, Lemieux T. 1996. Labor market institutions and the distribution of wages, 1973–1992: a Semi-parametric approach, Econometrica 64(5): 1001–44.

Fernandez V. 2008. Copula-based measures of dependence structure in assets returns, Physica A: Statistical Mechanics and its Applications 387(14): 3615-3628.

Fisher, K., Robinson, J. P. (2011). Daily life in 23 countries. Social Indicators Research, 101, 295–304.

Fortin N, Lemieux T, Firpo S. 2011. Decomposition Methods in Economics, Handbook of Labor Economics 4: 1-102.

Frandsen B, Lefgren L J. 2017. Testing rank similarity, Review of Economics and Statistics (forthcoming).

Giménez-Nadal, J., Molina, J. The gender gap in time allocation. IZA World of Labor 2022: 497 doi: 10.15185/izawol.497

Giménez-Nadal, J. I., Sevilla, A. 2012. Trends in time allocation: A cross-country analysis. European Economic Review, 56, 1338–1359.

Goldin C. 2014. A Grand Gender Convergence: Its Last Chapter, American Economic Review 104: 1091-1119.

Gouveia M, Strauss R P. 1994. Effective federal individual income tax functions: an exploratory empirical analysis. National Tax Journal 47 (2): 317–338.

Heckman J J, Smith J, Clements N. 1997. Making the most out of programme evaluations and social experiments: accounting for heterogeneity in program impacts, Review of Economic Studies 64: 487-535.

Kobus, M., 2012. Attribute decomposition of multidimensional inequality indices. Econ. Lett. 117, 189–191.

Kobus M, Kurek R. 2018. Copula-based measurement of interdependence for discrete distributions, Journal of Mathematical Economics 79: 27-39.

Kobus, M., Kurek R. 2023. Measuring inequality in the joint distribution of socioeconomic status and health, Economics Letters 226: 110093.

Maasoumi E. 1986. The Measurement and Decomposition of Multi-dimensional Inequality, Econometrica, 54(4):991-997.

Maasoumi E. 1989. Composite Indices of Income and Other Developmental Indicators: A General Approach, Research on Economic Inequality 1: 269-286.

Maasoumi E. 1989. Continuously Distributed Attributes and Measures of Multivariate Inequality, Journal of Econometrics 42(1): 131-144.

Maasoumi E, Wang L. 2019. The Gender Gap between Earnings Distributions, Journal of Political Economy 127(5): 2438–2504.

Maasoumi E, Racine J. 2016. A solution to aggregation and an application to multidimensional well-being frontiers, Journal of Econometrics 191: 374-383.

Maasoumi E, Tong X 2015. Weights and Substitution Degree in Multidimensional Wellbeing in China, Journal of Economic Studies 42(1): 4-19.

Maasoumi E, Zandvakili S 1990. Generalized Entropy Measures of Mobility for Different Sexes and Income Levels, Journal of Econometrics 43(1-2): 121-133.

Maasoumi E, Zandvakili S. 1989. Mobility Profiles and Time Aggregates of Individual Incomes, Research on Economic Inequality 1: 195-218.

Maasoumi E, Zandvakili S. 1968. A Class of Generalized Measures of Mobility with Applications, Economics Letters 22 (1): 97-102.

Maasoumi E, Nickelsberg G. 1988. Multivariate Measures of Well-being and an Analysis of Inequality in the Michigan Data, Journal of Business and Economic Statistics 6(3): 327-334.

Marshall A W, Olkin I. 1979. Inequalities: theory of majorization and its applications. Academic Press, San Diego.

Mulligan C B, Rubinstein Y. 2008. Selection, investment and women's relative wages over

time, Quarterly Journal of Economics 123: 1061-1110.
Nelsen R B. 2006. An Introduction to Copulas. 2nd edition, New York: Springer-Verlag.
Oaxaca, R. 1973. Male-female wage differentials in Urban Labor markets, International Economic Review 14(3).
Sen A K. 1973. On economic inequality. Oxford, UK: Clarendon Press.
Shannon, Claude E. 1948. A Mathematical Theory of Communication. Bell System Technical Journal. 27 (3): 379–423.
Shorrocks, A. F. (2013). "Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value". Journal of Economic Inequality 11 (1): 99.
Sklar A. 1959. Fonctions de repartition an dimensions et leurs marges, Publ. Inst. Statist. Univ. Paris, 8:229–231.
Stiglitz, J. E., Sen, A., Fitoussi, J. P. Mismeasuring Our Lives: Why GDP Doesn't Add Up. New York: The New Press, 2010.
Tsui, K.-Y. (1995). Multidimensional generalizations of the relative and absolute inequality indices: the Atkinson-Kolm-Sen approach. Journal of Economic Theory 67(1):251-265.
Young H P. 1990. Progressive Taxation and Equal Sacrifice. American Economic Review 80 (1): 253–266.

Appendix

Proof of Proposition 2. We derive and proof the attribute decomposability of $I(F)$ as defined in (4). We focus on the case of two outcome dimensions, but this restriction can be easily relaxed given the results in Abul Naga and Geoffard (2006). Let $W(F) = \int U(y_1, y_2)\mathrm{d}F$ with utility function U denote the social welfare function with standard properties of being concave and increasing. δ is defined as $W(F) = W(\delta(F)\bar{y})$, where $\bar{y}$ is a vector of dimensions' means $\bar{y}_1, \bar{y}_2$. That is, δ is a proportion of each outcome that if given to every individual in the population results in the same level of welfare as the original distribution F. Now $I(F) = 1 - \delta(F)$, where $\delta(F) \in [0, 1]$ is a measure of *equality*. This is the procedure that produces inequality measures that are normatively significant, in a sense that welfare W goes down whenever the associated inequality measure I goes up. In the same way as with I, we can consider attribute decomposability of δ.

Lemma 1. *$\delta(F)$ is attribute decomposable as follows:*

$$\ln \delta(F) = \frac{r_1}{r_1 + r_2} \ln \gamma_1(F_1) + \frac{r_2}{r_1 + r_2} \ln \gamma_2(F_2) + \frac{1}{r_1 + r_2} \ln \kappa(F), \tag{18}$$

where

Proof. Let δ be the proportion of $\bar{y}$ that is necessary to achieve the same level of welfare if all attributes were distributed equally. Formally, let $w_0 = \int U(\delta\bar{y}_1, \delta\bar{y}_2)\mathrm{d}F$ denote the welfare level associated with $\bar{y}$. Second, let ρ_1 be the proportion of $\bar{y}_1$ that is necessary to attain w_0, if the first attribute was equally distributed and the distribution of the second attribute remained as is. Formally, $w_0 = \int U(\rho_1 \bar{y}_1, \bar{y}_2)\mathrm{d}F$. Third, let γ_1 be the proportion of $\bar{y}_1$ that is necessary to attain w_0, if both attributed were distributed equally. Formally, $w_0 = \int U(\gamma_1 \bar{y}_1, \rho_2 \bar{y}_2)\mathrm{d}F$.

For the utility function that embeds Tsui (1995) inequality measure $U(y_1, y_2) = -y_1^{r_1} y_2^{r_2}$ with $r_1, r_2 < 0$ to ensure increasingness and sensitivity to increases in association between variables, it follows that

$$w_0 = \int (\delta \bar{y}_1)^{r_1} (\delta \bar{y}_2)^{r_2} \mathrm{d}F = \int (\gamma_1 \bar{y}_1)^{r_1} (\rho_2 \bar{y}_2)^{r_2} \mathrm{d}F.$$

After modification, we get $\delta^{r_1+r_2} = (\gamma_1)^{r_1} (\rho_2)^{r_2}$. Going further,

$$\ln(\delta) = \frac{r_1}{r_1+r_2} \ln(\gamma_1) + \frac{r_2}{r_1+r_2} \ln(\gamma_2) + \frac{1}{r_1+r_2} \ln\left(\rho_2/\gamma_2\right)^{r_2}, \tag{19}$$

which is the desired decomposition with $\kappa := (\rho_2/\gamma_2)^{r_2}$. We now need to derive γ_1, γ_2 and κ.

Note that $w_0 = \int (\gamma_1 \bar{y}_1)^{r_1} (\rho_2 \bar{y}_2)^{r_2} \mathrm{d}F = \int (\bar{y}_1)^{r_1} (\rho_2 \bar{y}_2)^{r_2} \mathrm{d}F$. Solving for γ_1 yields:

$$\gamma_1 = \left(\int \frac{y_1^{r_1}}{\bar{y}_1^{r_1}} \mathrm{d}F_1(y_1 \right)^{\frac{1}{r_1}}.$$

Similarly, for γ_2 we get:

$$\gamma_2 = \left(\int \frac{y_2^{r_2}}{\bar{y}_2^{r_2}} \mathrm{d}F_2(y_2 \right)^{\frac{1}{r_2}}.$$

Furthermore, we use $w_0 = \int (\bar{y}_1)^{r_1} (\rho_2 \bar{y}_2)^{r_2} \mathrm{d}F = \int (\bar{y}_1)^{r_1} (\bar{y}_2)^{r_2} \mathrm{d}F$ to obtain

$$\rho_2 = \left(\frac{\int y_1^{r_1} y_2^{r_2} \mathrm{d}F}{\int y_1^{r_1} \bar{y}_2{}^{r_2} \mathrm{d}F} \right)^{\frac{1}{r_2}}.$$

Finally, substituting the expressions for γ_2 and ρ_2 into $\kappa := (\rho_2/\gamma_2)^{r_2}$ we get:

$$\kappa = \frac{\int y_1^{r_1} y_2^{r_2} \mathrm{d}F(y_1, y_2)}{\int y_1^{r_1} \mathrm{d}F_1(y_1) \int y_2^{r_2} \mathrm{d}F_2(y_2)}.$$

□

Linear Approximation. Collecting terms and reversing the log-linearization of δ, we obtain the attribute decomposition of I:

$$I = 1 - (\gamma_1)^{\frac{r_1}{r_1+r_2}} (\gamma_2)^{\frac{r_2}{r_1+r_2}} (\kappa)^{\frac{1}{r_1+r_2}}. \tag{20}$$

Applying a linear approximation around the point of perfect equality (i.e. $\gamma_1 = \gamma_2 = \kappa = 1$), we get the linear decomposition.

$$\begin{aligned} I(F) &= \tfrac{r_1}{r_1+r_2}(1-\gamma_1) + \tfrac{r_2}{r_1+r_2}(1-\gamma_2) + \tfrac{1}{r_1+r_2}(1-\kappa) + R, \\ &= \tfrac{r_1}{r_1+r_2} I_1 + \tfrac{r_2}{r_1+r_2} I_2 + \tfrac{1}{r_1+r_2}(1-\kappa) + R. \end{aligned} \tag{21}$$