

VIENNA INSTITUTE OF DEMOGRAPHY

Working Papers

1 / 2014

Bilal Barakat

Revisiting the History of Fertility Concentration and its Measurement



Vienna Institute of Demography
Austrian Academy of Sciences

Wohllebengasse 12-14
A-1040 Vienna - Austria

E-Mail: vid@oeaw.ac.at
Website: www.oeaw.ac.at/vid



Abstract

To date, the concentration of fertility has been studied almost exclusively through the lens of indicators derived from the Lorenz curve of completed parity. Arguably, this situation has resulted from a misapplication of criteria from the study of income inequality, without fully considering the distinctive characteristics. Here, alternative measures of concentration are justified on substantive grounds and applied to a wide range of empirical data from developing, transitional, and industrialised countries. In particular, the traditional scale-invariant concentration measures are complemented by translation-invariant and intermediate perspectives. The latter conveniently also connects the nonparametric indicator approach to the study of fertility concentration with established statistical models for count data. This connection provides the missing link for the kind of behavioural interpretation of fertility concentration dynamics that has been attempted in the past, but can be shown to have been ill-founded. Examining the phenomenon of fertility concentration through this tri-focal lens leads to a picture that is more coherent, more parsimonious, and more complete.

Keywords

Fertility, parity distribution, concentration, inequality

Authors

Bilal Barakat, Research Scholar, Wittgenstein Centre (IIASA, VID/ÖAW, WU)
Vienna Institute of Demography, Austria. Email: bilal.barakat@oeaw.ac.at

Revisiting the History of Fertility Concentration and Its Measurement

Bilal Barakat

1 Introduction

Completed cohort fertility is a core demographic phenomenon. As with many complex social phenomena, its average magnitude alone provides a rich field for varied research. Especially so if changes over time are considered, or differences between populations, generations, or specific sub-groups. Arguably, interest in cohort fertility differences between countries (Dorius, 2008) or socioeconomic groups has generally eclipsed interest in the patterns of interpersonal variation *within* more-or-less homogeneous groups.

Closer to present concerns is the relatively small body of research that specifically addresses questions of fertility «concentration», in particular its measurement (Goodwin & Vaupel, 1985), empirical patterns (Lutz, 1989), and its relationship to other aspects of demographic change (Lichter & Wooton, 2005; Spielauer, 2005). An open question is whether fertility concentration tends to increase as the overall level of fertility declines. The evidence to date is inconclusive (Shkolnikov *et al.*, 2007). Here, I argue that this literature has been held back by prematurely settling on a limited set of concentration measures that fail to capture the full range of useful conceptualisations of «concentration». This situation has resulted from a misapplication of criteria from research on income inequality, without fully considering the distinctive characteristics of birth parity distributions that call for a different treatment.

The contribution of the present study is to take steps to remedy this situation. Alternative measures of concentration are justified on substantive grounds. Applying them to empirical data leads to results that are not only qualitatively different in some cases, but also lead to more coherent and parsimonious accounts of the underlying dynamics, as well as to entirely novel observations. No new measures will be developed here, only the application to fertility of ones that have hitherto been neglected in this context. Indeed, there is little need for participating in «the process of rediscovering old implements left lying around in the inequality-analyst's toolshed [that] continues unabated, so that often several labels and descriptions exist for essentially the same concept» (Cowell, 2009, 180).

The presentation is structured as follows. First, in section 2, I review the existing literature on fertility concentration, including the methodological origins of indicators commonly employed in its study. The methodological debate is re-opened in sections 2.2 and 2.3 by justifying the utility of alternative concepts and measures of concentration. In particular, the reasons for their past dismissal are shown to be less than compelling. The practical difference to our understanding of fertility behaviour is illustrated in section 3 by applying this more diverse set of indicators to empirical parity data. The implications are discussed in the concluding section.

2 Concentration, Variation, Dispersion: Linkages and Distinctions

While there is seemingly a variety of fertility concentration measures commonly used in demographic research, this diversity is more apparent than real. Effectively, it is limited to measures that are derived from the Lorenz curve. This curve represents the cumulative distribution of some quantity of interest when individuals are ordered monotonically with respect to their shares (cf. Shkolnikov *et al.* (2007) for a recent presentation in the fertility context). The «concentration ratio» (CR) with respect to fertility, can be defined graphically in terms of the area between the Lorenz curve and the diagonal. Its intersection with the horizontal and vertical lines through 0.5 in turn determine the popular «havehalf» and «halfhave» measures. These capture the (smallest) share of women who account for half the births, respectively the share of births accounted for by the 50% of women with highest parity. This choice has been partly driven by how easy different measures are to comprehend and explain (Goodwin & Vaupel, 1985; Shkolnikov *et al.* , 2007). In addition to their close conceptual relationship as different one-dimensional summaries of the Lorenz curve, these three indicators tend to be highly correlated in empirical data.

It is useful in this context to recognise that much of the related literature from other fields is phrased in terms of «inequality» in the distribution of some good. This is obvious when noting that the CR is equivalent to the Gini coefficient, which is generally characterised as a measure of inequality. We may identify identify high inequality with high concentration, or equivalently with high variation in the individual shares.

2.1 Relative Concentration Measures, Scale Invariance, and Its Misinterpretation

Entire classes of concentration measures derived from alternative premises are generally disregarded in the fertility context. This is curiously and needlessly self-limiting. In the early days of the study of fertility concentration, researchers considered a much wider, but still incomplete, «smorgasbord» of indicators that all attempt to measure the level of concentration in a distribution, as well as criteria for choosing between them.

It is enlightening why the arguably most obvious statistical candidate for measuring variation in a parity distribution, namely its variance, was excluded from further consideration by Goodwin & Vaupel. One criterion that sharply narrows the range of concordant indicators, and that excludes the variance among others, is the insistence on multiplicative *scale invariance*. This means that the value of the indicator in question should remain the same if all observed outcomes were multiplied by a common factor. Formally, it requires of indicator I that:

$$I(\mathbf{X}) = I(c\mathbf{X}), \quad \forall c > 0. \quad (1)$$

Such an indicator is considered a «relative» measure¹.

In the following, I argue that both the justification of this invariance condition and the interpretation of its implications presented in the literature on fertility concentration must be reassessed in light of a careful consideration of what the criterion actually means.

In general, this condition is justified by noting that a reasonable measure of inequality should not depend on the arbitrary unit of measurement. When measuring income inequality, for example, the degree of inequality should be independent of whether we choose to measure money in dollars, euros, or cents. These units represent proportional scales, related to each other by some multiplier.

While commonsensical, this argument is less persuasive than it seems as a justification for requiring scale invariance in the sense of Eq. (1). With respect to measuring fertility concentration especially, the indispensability of scale invariance is more imagined than real.

To begin with, the naive interpretation conflates two conceptually distinct (Aczél

¹Note that usage has changed over time. While extremely well-established in the context of inequality indicators research, this use of «relative» is different from the usage by Lutz (1989), for example, as is the present use of «concentration ratio» for the Gini coefficient as applied to parity.

& Moszner, 1994) issues, namely «scale» as relating to the unit of measurement and as relating to real magnitude². Accordingly, we can distinguish between the requirement that an inequality measure should be robust to changes in the unit of measurement, and the requirement that it should be constant under uniform proportional change in the *real* values. The former condition, «unit-consistency» (Zheng, 2007), is a sensible requirement, but even with respect to income it has been argued that it does not require the strict condition of Eq. (1), merely that re-scaled measures under a change of units be *proportional* to each other. In other words, under unit-consistency, the inequality assessments are ordinally equivalent.

For rank comparisons, and even for a qualitative assessment of trends over time, ordinally equivalent measures give the same results. It is only when we move to comparisons of first differences that cardinal equivalence becomes crucial for the unit-consistency of statements such as: «The difference between countries A and B is larger/smaller than the difference between countries C and D», or for similar comparisons over time.

However, ordinal rather than cardinal equivalence is not a significant constraint, given that in fertility research the choice among scale invariant measures has been made without substantive consideration of their different (!) cardinal results (which are highlighted in decomposition analysis (Shorrocks, 1982; Kimhi, 2011), for example). In particular, the cardinal properties of the Gini coefficient are of questionable desirability in the case of income as well as for parity, in that it is more sensitive to transfers in the middle of the distribution (Atkinson, 1970). In other words, the Gini does not respect the Principle of Diminishing Transfers (Dahlby, 1987). Given the special status accorded to childlessness, it is questionable whether shifting a child from a mother at parity 2 to a childless women really reduces fertility concentration less than a transfer from a woman at parity 4 to one at parity 2, which will be the case with the CR if the modal parity is 3. In other words, cardinal equivalence under unit changes is not a compelling requirement as long as cardinal outcomes are incidental to the measure chosen rather than reflecting substantive considerations.

In any case, even (ordinal) unit-consistency is not relevant in the case of measuring the concentration of birth parity, because the latter possesses a *natural scale*. Most of us are not about to begin counting parity in half-children—nor in kilograms, for that matter.

Scale invariance «proper», in the sense of Eq. (1) applied to real proportional change

²Although (Goodwin & Vaupel, 1985), for example, do seem to refer to magnitude proper.

(rather than a change in units) no longer enjoys unreserved support as an *axiom* of inequality measurement. Even when it comes to income inequality, it is not universally accepted as valid (Amiel & Cowell, 1992, 1999) or desirable (Sen, 1973), but recognised as encoding a particular conception of inequality and a value judgement; a judgement that is decidedly political, no less (Kolm, 1976a,b).

In the case of fertility concentration, the case for *requiring* real scale invariance is weaker still. If we recognise that it is a statement about what happens when *every individual instance* in a distribution is multiplied by a common factor, it is evident that this condition can tell us nothing whatsoever about real parity distributions, because parity distributions are never related in this way.

To begin with, individual parities are integer-valued, so to derive a valid parity distribution from another through proportional increase requires the factor c in Eq. (1) to be integer too. An immediate consequence is that if the means of two parity distributions are not such that their ratio (or its inverse) is integer, they *cannot* be related in a way such that scale invariance applies.

The least objection is that every parity distribution contains *some* individuals at close-to-maximal parity, and that by the above the smallest possible scale increase is to double every value, pushing the largest parities beyond the biologically feasible range. If high parities are rare, we might simply consider a truncated parity distribution only and be content for two parity distributions to be «approximately» proportional, of course, so this objection is not too damaging. In addition, however, incomes are rarely absolutely zero (especially after transfers), but there are (plenty of) absolutely childless women. It is at least questionable whether the inequality in the distribution of «the joy or burden of reproduction» (Lichter & Wooton, 2005) could reasonably be said to remain constant if all parents had twice, thrice, or x -times as many children, but the childless remained childless.

Moreover, it is an incoherent thought experiment that we should expect a fertility concentration measure to give the same answer in a counterfactual world where «everyone has x times as many children». Even when x is an integer, the implied counterfactual parity distribution is grotesquely implausible in terms of its structure. If every individual parity were doubled, for instance, this implies that in the resulting parity distribution there may be individuals at parities 0 and 2, but by construction there would be not a single individual at parity 1. Indeed, every second parity level would not be observed. For larger values of x the gaps would be even larger. Such is clearly not a parity distribution that we would ever expect to observe in a real population. In

other words: real-life parity distributions will—for all intents and purposes—*never* be multiplicatively related to each other.

In guiding our choice between different concentration measures, it would be myopic to categorically exclude non-scale-invariant alternatives only because they give the «wrong» answer under impossible circumstances.

The above shows that it is inappropriate to insist on scale invariance for measuring fertility concentration. In addition, however, even taking scale invariance as given, its limitations demonstrated above make it clear that the concept has been misapplied in the literature. One example is the question how fertility concentration changes as the overall fertility level declines. It has been claimed that scale invariance implies that an observed «general pattern of increasing concentration with declining fertility» (Lutz, 1989, 179) cannot be a statistical artefact (Lutz, 1988), and must therefore be treated as a substantive phenomenon in need of a behavioural explanation. This argument is not debated even by other researchers who question whether—empirically—such a negative relationship actually is confirmed in general (Shkolnikov *et al.* , 2007).

However, the notion that «if reduction in fertility were proportional it would not affect concentration» (Lutz, 1988, 25) is not warranted by scale invariance, as the latter’s definition shows. The argument actually fails, because the parity distributions referred to are not scalar multiples of each other at the individual level. The conditions under which scale invariance asserts constant concentration are never satisfied in real parity distributions, not even approximately. But if the premise of scale invariance is not met, its conclusion does not follow.³ On the contrary, far from warranting the conclusion that «this association between the level and the concentration of fertility [...] reflects real changes in the degree of heterogeneity in reproduction» (Lutz, 1988, 25), such an association can in fact be shown to appear as a purely mathematical effect under conditions of perfect homogeneity if stochastic variation is taken into account, as shown in section 2.3.

All the above is not to argue that scale invariant measures are necessarily inappropriate to the study of fertility concentration, merely that scale invariance is inappropri-

³The claim is not saved by modifying it to proportional fertility changes between groups of women (which would at least allow for fractional increases/decreases). Scale invariance relates to proportional change in *all* individual values, not in groups. Indeed, the claim fails even for the most artificial grouping for which it might conceivably be true, namely the two groups from the very definition of the havehalf measure. Suppose the parities are: $G_1 = \{1, 1, 2, 2\}$, $G_2 = \{2, 2, 2\}$. The havehalf measure now equals $\frac{3}{7}$. While it is *possible* for a proportional change in the group means, say by 1.5 to leave havehalf unchanged ($G'_1 = \{2, 2, 2, 3\}$, $G'_2 = \{3, 3, 3\}$, havehalf = $\frac{3}{7}$), this is by no means a necessary implication, as a straightforward counterexample demonstrates: $G''_1 = \{1, 1, 3, 4\}$, $G''_2 = \{2, 2, 5\}$, havehalf = $\frac{2}{7}$.

ate as a *sine qua non* criterion. In the empirical part, the Concentration Ratio (CR), will represent the class of relative, scale invariant, concentration measures. Other relative measures exist that arguably possess more attractive theoretical properties than CR, especially with respect to consistent decomposability (Shorrocks, 1980). But since the range of fertility concentration measures is already expanded in other directions below, the choice among relative measures conservatively follows the precedent for CR in the existing literature.

2.2 Translation Invariance, and Absolute and Intermediate Inequality Measures

Given the above, it appears worthwhile not to restrict attention to scale invariant concentration measures. A mathematically obvious alternative is to consider measures with *translation invariance*, in other words, invariance under constant *addition* of a fixed amount to all instances, so that:

$$I(\mathbf{X}) = I(\mathbf{X} + x\mathbf{1}), \quad x > 0. \quad (2)$$

In contradistinction to «relative» concentration measures, this condition defines «absolute» measures of inequality and concentration.

Of course, one such measure is simply the variance. Another is the equivalent of the Gini index derived from the *absolute* Lorenz curve. In effect, this «absolute Gini» is simply the standard relative Gini multiplied by the mean. However, for reasons explained further below, an alternative additively invariant concentration measure is preferred here, namely the Kolm index of inequality. This is defined as:

$$K_\alpha(\mathbf{X}) = \frac{1}{\alpha} \log \left(\frac{1}{n} \sum_{i=1}^n e^{\alpha(\lambda - X_i)} \right).$$

The parameter α measures the degree of sensitivity to inequality. As α increases, the Kolm index converges to the difference between the mean and minimum value. For parity the minimum is normally zero, so the index converges to the mean for high α . This justifies using a relatively low value of α . In the following analyses, it was set to $\alpha = 0.1$. The results are actually robust over a wide range of values.

With respect to income, the Kolm index may be loosely interpreted as measuring the (per capita) saving that would be possible under a uniform income distribution without loss of social welfare. In the absence of such a normative interpretation in

the context of parity, the Kolm index may simply serve as a descriptive statistical measure of concentration. While not as well-known as the Gini index or even its relative counterpart, the Atkinson index of inequality (Atkinson, 1970), it has on occasion been used in high-profile applied work (Atkinson & Brandolini, 2010). In fact, a case can be made for the Kolm index as the canonical choice of absolute inequality index (Blackorby & Donaldson, 1980), indeed it is the only such index that is decomposable, apart from the variance (Zheng, 2007).

Applying the same kind of scrutiny to translation invariance as to scale invariance, note that the former still requires x to be integer in Eq. (2). In addition, the absence of childlessness in the shifted distribution would be implausible. Nevertheless, a shifted parity distribution is at least conceivable as an approximation to one that might actually occur if childlessness is rare, because unlike scale invariance it does not imply a «gappy» distribution. In any case, the intention here is not to present translation invariance as an unreservedly superior criterion, merely as an equally reasonable alternative that is, by construction, outside the commonly considered pool of scale invariant ones.

In principle, scale invariant concentration measures on the one hand and translation invariant ones on the other may give different results. These differences pertain not just to the level of concentration, but also the magnitude of changes in concentration, and even to the direction of such changes. Such divergent results are not a deficiency. On the contrary: the point of considering alternative concentration measures is that they should provide independent rather than redundant information. The above discussion therefore serves to justify complementing CR with the genuinely different Kolm index, rather than with multiple other instances of scale invariant indices based on the Lorenz curve.

Perspectives for interpreting different kinds of concentration measures and for choosing between them (or for choosing a parameter value for α in the case of the Kolm and other families of indices) are largely predicated on the measurement of economic inequality. The «normative» approach to inequality measurement, for example, derives inequality indices from a Social Welfare Function (Chakravarty, 2009). Such a function assigns distributions across the population to scalar values that capture social utility or desirability. With the notable exception of China's One-Child-Policy, explicit social valuations of specific parity distributions are rare, even in policy environment that are generally pro-natalist or anti-natalist with respect to the mean fertility level. This approach is of limited interest in the case of measuring fertility concentration,

therefore. The same applies to the related interpretation of inequality measures from the perspective of risk-averse individuals choosing a distribution from behind a «veil of uncertainty» (Dahlby, 1987), since the notion of «riskiness» of a parity distribution likewise requires value comprehensive judgements that should not be a prerequisite for measuring the concentration of parity.

By contrast, an interpretation that can—at a stretch—be applied to fertility concentration is the question of how to distribute an increment in the overall total in a fair manner. Fairness here is not meant in a Rawlsian sense of actively reducing inequality, but in a way that leaves inequality constant. Scale invariance implies that a windfall should be distributed proportionately. Translation invariance suggests that it should be distributed uniformly. This perspective makes it obvious why relative scale invariance and absolute translation invariance have also been characterised as «rightist» and «leftist» measures of inequality respectively. In a parity context, consider the hypothetical problem of assigning a group of refugee children to host families depending on the latter's size. For dramatic clarity suppose the host community is a small village and that the number of refugee children is similar in magnitude to the number of local children. The relative notion of fertility concentration would suggest that the childless take in no-one, and that a family with four children of its own take in four times as many refugee children as a family with a single child. One may question whether such a proportional distribution really should be considered not only more reasonable than distributing the extra children uniformly across households, but so much more so that the latter should not even be contemplated. This example illustrates the justification for adding an absolute, translation-invariant concentration measures at least for a complementary perspective.

The interpretation of invariance conditions as prescriptions for the fair distribution of incremental income (or incremental fertility) creates a natural definition of inequality indices that are neither strictly relative nor strictly absolute. Such indices correspond to distribution strategies for incremental income that is a compromise between distributing incremental income entirely proportionally or entirely uniformly. The classic «centrist» (Kolm, 1976a,b) respectively «intermediate» (Bossert & Pfingsten, 1990) inequality indices have since been shown to suffer from limitations, they are not unit-consistent, for example. For reasons that will become clear in a moment, I focus here not on the seemingly most active line of research for defining intermediate inequality indices, exemplified by del Río & Ruiz-Castillo (2000); del Río & Alonso-Villar (2010). Instead, it is worthwhile revisiting Krtscha's index (Krtscha,

1994; Subramanian, 2013).

In terms of the distribution of incremental income, Krtscha's index makes the intuitively plausible suggestion that inequality remains constant when half of each marginal incremental unit is distributed uniformly and half proportionately. In this sense, it falls exactly between the relative and absolute perspectives. Crucially, the proportional part should continuously track the changing relations induced by the distribution of previous increments. While neither the continuous limit of this operation nor the half-half splits for marginal increments make sense with respect to indivisible units such as births, this index is nevertheless attractive as the third, intermediate measure of fertility concentration to complement the relative and absolute perspectives provided by the CR and Kolm's index. The reason is that Krtscha's index derived from the above distribution rule turns out to be equal to the Variance-to-Mean ratio (VMR), also known as statistical dispersion. This immediately brings us conceptually closer to birth counts again, because the VMR is of course a key parameter of standard parametric models for count data. In particular, the relationship of a count distribution to the Poisson distribution is determined by whether the VMR, in other words, its «statistical dispersion», is less than, equal to, or greater than unity.

2.3 A Stochastic Perspective

This relationship with a parametric model is useful for at least two reasons.

First, recall that in general, neither proportional scaling nor a translation will generate a plausible parity distribution from another. If we wish to generate plausible counterfactual parity distributions that relate to each other through a level shift, in other words, to generate hypothetical parity distributions for given mean values, a parametric model is required.

Second, as noted already, there have been attempts to interpret changes in relative concentration measures, specifically the CR and havehalf indices, with a behavioural interpretation in terms of preferences and contraceptive use (Shkolnikov *et al.*, 2007; Lutz, 1989), for example. However, even uniform behaviour may result in stochastic results. Of four initially childless women who all desire exactly one child and act identically before and after the first birth, one may remain childless because she does not conceive, one may meet her target, another might end up with a second child after a contraceptive failure, and a fourth may have triplets.

These two reasons suggest an investigation of how the different concentration measures behave as known parametric distributions trace out a path of declining overall

mean. A recently revived generalisation of the Poisson function allows for both overdispersion (the only direction permitted by the negative Binomial distribution) and underdispersion, namely the Conway-Maxwell-Poisson distribution (Sellers *et al.* , 2012). It is without practical alternative for the first purpose mentioned above. Since it contains the regular Poisson distribution as a special case, it also serves the second purpose, by allowing us to study the concentration of outcomes that is induced in a perfectly homogeneous population by stochastic variation. It is at least an arguable proposition that the informal notion that «a parity concentration measure should be independent of the overall level of fertility» is captured coherently by identifying «the overall level of fertility» with the mean of the effectively stochastic distribution of parities. In particular, this corresponds with the standard interpretation of regression coefficients, by equating «mean independence» with «independence from the intercept».

It is clear analytically that stochastic mean independence differs from scale invariance. Doubling every value of a Poisson distribution does not lead to a Poisson distribution with twice the mean, after all. The variance will increase fourfold instead of doubling, and the result would not even be distributed according to an overdispersed generalised Poisson distribution due to the gaps.

We can simulate how different concentration measures, including the CR, relate to the mean in an independent, identically distributed (i.i.d.) sample of generalised Poisson counts, the standard model for count data, including births. Recall the assumptions underlying the Poisson distribution: such a distribution would arise if every women (in the relevant age range) experienced the same small annual probability of giving birth. This is clearly a population that is perfectly homogeneous in behaviour, and changing the underlying probability is conceptually a change in level. Figure 1 shows the outcome of this model for different values of the dispersion parameter.

First, as an aside, note that in addition to being closely related to CR by construction and therefore adding little information, the havehalf index suffers from a troublesome characteristic that has largely gone unnoticed. In particular, even in the simplest and smoothest case where parity is distributed as Poisson counts, it exhibits strong discretisation artefacts. Suppose, for example, that the mean of a Poisson-shaped parity distribution $\nu = 1$ declines smoothly from 3.75 to 3.25 children, then on to 2.75 and finally to 2.25. As is evident in Figure 1, that concentration hardly changes during the first step, then seems to undergo a sudden drop, but that this decrease stall agains during the third interval. This despite the fact that the underlying distributional changes are perfectly smooth. Matters are worse if the distribution is

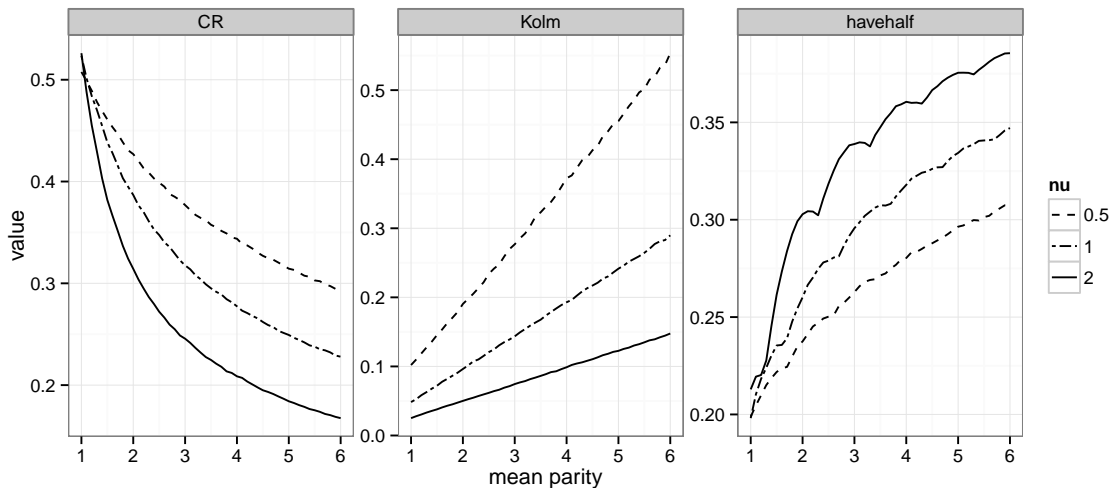


Figure 1: Concentration by mean parity for simulated parity distributions from a Conway-Maxwell-Poisson distribution, by value of Variance-to-Mean dispersion parameter ν . $\nu = 1$: Poisson, $\nu = 2$: overdispersed, $\nu = 0.5$: underdispersed.

overdispersed ($\nu = 2$), in which case there are even apparent trend reversals and concentration appears to alternately increase and decrease. This effect does not affect the CR, which in the range of interest (for fertility) has a smooth, almost inversely proportional relationship to the mean. The Kolm index is even *exactly* proportional to the variance in the special case of Poisson birth counts.⁴

From the results (Figure 1), it is evident that there is a negative relationship between the inverse of the CR and the mean in a Poisson model of fertility counts. This puts past empirical findings of a negative relationship between concentration and mean fertility in perspective. Even though the Poisson distribution is not necessarily an accurate model of reality, this means that when concentration is found to increase with declining fertility in empirical data, this observation is not *ipso facto* in need of an explanation. Occam’s razor would suggest to simply attribute this to the intrinsic concentration increase that goes hand in hand with a uniform decline in mean parity under standard distributional assumptions in a homogenous population, until proven otherwise. As a matter of fact, the mean-dependence of CR is *not* restricted to the Poisson case. In particular, it is not a consequence of the equi-dispersion assumption as such. The negative relationship between CR and mean parity continues to hold in the case of generalised Poisson distributions that exhibit overdispersion or underdispersion.

This result has implications in a regression setting. Here, I mean not the regression

⁴Since—to the best of my knowledge—this relationship has not yet been established in the statistical literature, an original proof is included in the appendix.

of inequality measures on covariates (Aaberge *et al.* , 2006), to answer question such as: «Is inequality higher among the educated?». Instead, what is at stake are questions such as: «How much inequality is accounted for by education?». The state-of-the-art approach to such questions is to regress the outcome on the covariates followed by decomposition (Wan & Zhou, 2005). However, the status of changes in the intercept is debated (Kimhi, 2011), at least with respect to income inequality. In other fields, such as education, inequality is uniquely identified with non-zero regression coefficients for background characteristics such as gender, parents' occupation or income, or migration status, independently of the intercept.

One conflict between the regression and decomposition perspectives pointed out by Giroux *et al.* is that at constant means for different groups (which may be defined by covariates), even with constant regression coefficients, the latter takes into account the (potentially changing) group sizes, in other words composition effects. A separate question, however, is how they react to changes in the coefficients, holding group sizes constant.

To make this point more drastically, consider groups of size one, in other words individual-level heterogeneity. The distribution becomes $\text{Parity}_i \sim \text{Poisson}(\lambda_i)$, with $\log \lambda_i = \beta_0 + \beta_i$ where the β_i vary over individuals. We know from Figure 1 that the CR declines with increasing Poisson mean, so it comes as no surprise that with a suitable increment of β_0 , it is trivial to construct examples where the spread (by whatever measure) of the β_i increases, yet where according to CR, the overall concentration of fertility is taken to be decreasing.

This flies in the face of the standard interpretation of such regression coefficients in other contexts. In some meaningful sense at least, inequality has not changed in the example of a Poisson with mean β_0 and with mean $2\beta_0$, so it is natural to consider the VMR as a measure that remains constant in this case at least for a complementary perspective. Such invariance with respect to the expected value only holds in the context of a given stochastic distribution model. This is not necessarily a constraint, because many applications require such a model in any case. The standard approach to regressions of completed parity requires an underlying parametric distributional specification (be it explicit or implicit) for the completed parity distribution. Simulations, too, require the ability to generate birth parities with a given level of concentration. Using a generalized count model, this is trivial when concentration is understood as dispersion. By contrast, when concentration is measured by the CR, havehalf or halhave indices, the generating mechanism is underdefined.

A key advantage of parametric descriptions of the parity distribution over non-parametric ones is that the latter are precisely that: descriptions, while the former are actually *models*. This means that the codify assumptions about fertility behaviour that can be interpreted in relation to theories. To give an example, the Poisson model is equivalent to the assumption that the probability of experiencing (another) birth does not depend on the parity already attained or on the waiting time since the last birth. We already know from qualitative studies and quantitative micro-level research that this is not an accurate description of fertility behaviour. But the correspondence between the parametric distribution and the theoretical behavioural model allows us to conclude this even for populations where no such in-depth research exists, purely from the aggregate observation that the Poisson model does not fit the parity distribution well. For example, consider the competing explanations for why China is a counterexample to the general pattern and did not experience increasing concentration as fertility declined. In light of the knowledge gained above that a homogeneous fertility decline is perfectly compatible with increasing concentration, it becomes clear that assuming the decline in China was particularly «egalitarian» (Lutz, 1989, 180) does less to explain the result than the assumption that it followed from strict parity control, which would indeed led to non-Poisson behaviour.

2.4 Summary

These arguments provide a case for adding statistical dispersion, in other words, the ratio of the variance to the mean, to the suite of concentration measures to be applied to the study of fertility concentration. Dispersion provides an additional perspective, with distinct advantages outlined above. The inclusion of dispersion also justifies the choice of the relatively complex Kolm index of absolute inequality over other absolute inequality indices, in particular the absolute Gini coefficient and the variance. As a matter of fact, all three absolute indices (Kolm, absolute Gini, variance) give similar results in the empirical analysis below. However, the latter two are simply the product of, respectively, the standard relative Gini coefficient or the dispersion with the mean. Accordingly, in graphs of how the concentration varies with the mean they are strictly redundant. The absolute Kolm index, by contrast, is not fully determined by either the relative Gini or the dispersion in combination with the mean, and is therefore preferred here, even though other theoretically attractive properties, such as its decomposability, are not exploited here.

The following example with data for the USA demonstrates that the CR, Kolm,

and dispersion really do provide distinct perspectives on empirical patterns of parity concentration. Together, they therefore provide a more comprehensive and nuanced picture.




Figure 2: Concentration of completed cohort parity distribution against mean completed fertility by birth cohort. Data: HFD, women aged 40–44.

There are two different ways of reading these graphs, depending on whether we treat the level of concentration as a function of time or of mean fertility. As a time series, we note that the three indicators give qualitatively different interpretations of trends in fertility concentration. In terms of CR, we observe the pattern: $CR(1918) > CR(1933) < CR(1945)$. By contrast, the Kolm index suggests the opposite time pattern, namely: $K(1918) < K(1933) > K(1945)$. Finally, the changes in dispersion suggest that there

was a continuous increase in concentration: $\text{VMR}(1918) > \text{VMR}(1933) > \text{VMR}(1945)$.

As a function of mean fertility, concentration as captured by the CR is generally declining with increasing fertility, while the Kolm index is increasing. Dispersion as measured by the VMR, finally, has an inconsistent relationship with mean fertility in this case. The historical US case will not be developed further here.

The above establishes that a triad of CR, Kolm index, and VMR dispersion each can be justified as reasonable measures of fertility concentration in their own right and, unlike the commonly deployed set of concentration measures, are fundamentally distinct by construction. As a result, together they allow for a more comprehensive and nuanced study of fertility concentration.

3 Empirical Patterns of Parity Dispersion

The relationship of dispersion to the mean level of fertility can be investigated whenever the latter exhibits some variation. Evidently, mean fertility levels vary across a number of dimensions. In particular, fertility can and does vary across countries, across cohorts and over time, across subpopulations, and not least over the life course. Accordingly, I attempt to briefly examine patterns of fertility dispersion along all these dimensions.

3.1 Data

The following analyses are based on data from three sources, each of which allows us to highlight a different angle of the phenomenon, not least by focusing on different parts of the spectrum of overall fertility.

The Human Fertility Database⁵ is focused on industrialised high-income countries, at least with respect to time series of completed parity. Cumulative fertility rates by birth order from the HFD were extracted for all countries for which these were available at the time of writing. The strength of these data for present purposes is the fact that they carefully account for exposure rates and mortality, and that they provide a consistent longitudinal perspective. The limitations (for present purposes) are, firstly, that, even for the earliest cohorts, only a relatively limited range of average fertility levels are represented, namely the low fertility end of the spectrum, and secondly, that

⁵Human Fertility Database. Max Planck Institute for Demographic Research (Germany) and Vienna Institute of Demography (Austria). Available at www.humanfertility.org (data downloaded on 2013-01-16)

high parities are aggregated at 5+.⁶

The Integrated Public Use Microdata Series (IPUMS), being based on census samples, is broader in coverage, but by necessity still limited to countries with the institutional capacity to conduct credible censuses, and to samples including the variable «Children Ever Born». The disadvantage of these data are that they are strictly cross-sectional snapshots of approximately completed fertility of women 40–44 alive at the time of census. As such, the figures differ slightly from the cohort fertility rates that would be calculated from complete birth counts and exposures by age, period and parity. Their advantages are, firstly, that they cover a much broader range of populations with vastly different levels of average fertility, and that their microdata format and large sample sizes allow disaggregation into socio-economic subgroups, for example.

Settings with very high overall fertility are largely limited to the poorest and least developed countries with poor availability or reliability of comprehensive parity data. Accordingly, the Demographic and Health Surveys (DHS) have established themselves as the main source of data for research and policy and also inform the present study. Aggregate tables of completed parity of women aged 40–44 were extracted from the DHS database.

3.2 Results

3.2.1 Across Cohorts

The evolution over time of completed parity concentration according to the different measures for the HFD countries with appropriate time series is displayed in Figure 3. The CR perspective is well-known, of course. The alternative perspectives in terms of dispersion and absolute inequality lead to a number of significant differences in interpretation. While dispersion agrees with the CR regarding the uniform decline in concentration for the pre-WWII birth cohorts, absolute concentration in the USA actually increased up to about 1930 (unfortunately it is the only country with data extending that far back), before declining. All three approaches agree that a universal and highly uniformly paced decline in fertility concentration occurred over cohorts from that time onward up to the mid-late 1940s, i.e., up to those cohorts who came of age

⁶As the aggregation occurs at the level of cumulative fertility rates (CCFR), the calculated share at parity 4 is also affected, corresponding to the difference between CCFR4 and CCFR at exactly 5 (rather than 5+). Three different synthetic assumptions about the true spread of the reported CCFR5+ over parities 5-10 were tested, namely a uniform distribution, and a linear or exponential decline. All presented analyses are based on the exponential model, but, unless noted otherwise, the conclusions are qualitatively robust in the sense of not being sensitive to the choice of imputation.

in the mid-1960s. Comparing the pre-WWII and post-WWII birth cohorts, it is clear that the CR suggests an increasing cross-country variation in fertility concentration, whereas in terms of the Kolm index, there was more of a convergence and in terms of dispersion there was little overall change.

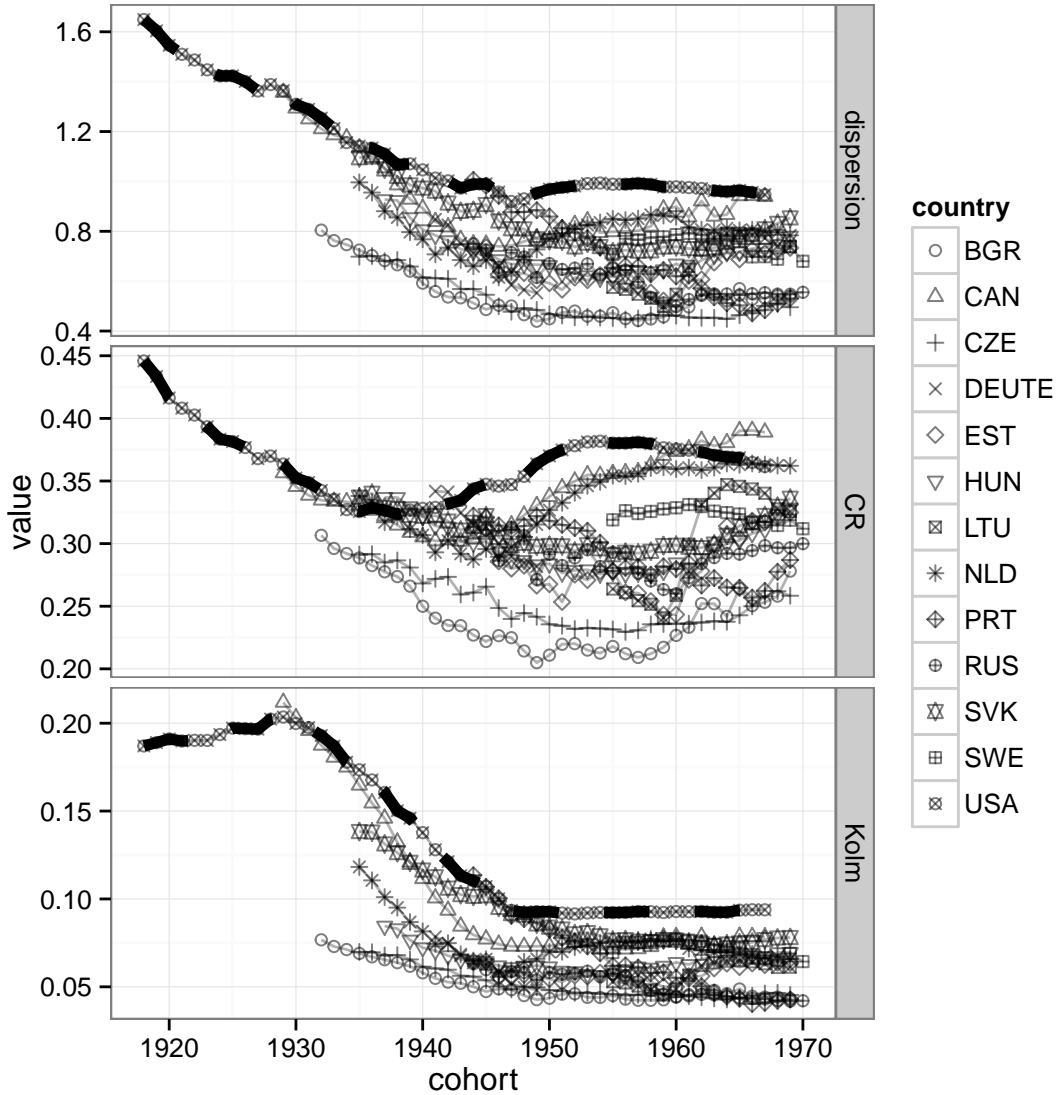


Figure 3: Concentration of completed cohort parity distribution by birth cohort. USA highlighted with dashed overlay. Data: HFD, women aged 40–44.

The apparent increase in CR in the USA during the middle decades of the century is therefore in no need of a behavioural explanation in terms of differential preferences or otherwise. The trajectory of dispersion shows that the latter declined steadily and uniformly as everywhere else, and then remained essentially constant during the post-WWII period. Moreover, it settled at a value of almost exactly unity. In other words,

the pattern is quite compatible even with a model of homogeneous secular shifts in fertility. Both in the case of the USA and more generally, it seems that the dynamics of concentration over time can be described much more parsimoniously from the intermediate and absolute perspectives of dispersion and the Kolm index than from the standard relative CR perspective.

3.2.2 Cross-country Comparisons

The development of fertility concentration over time is only half the story. Of similar, if not greater, interest is the question whether it systematically varies with the overall level of fertility. Conceptually, we would like to treat the level and concentration of the parity distribution as conceptually independent parameters. As shown in Section 2.3, the CR and related relative concentration measures actually fall short in this regard, because they are in fact mean-dependent. It is crucial, therefore, to also consider the dispersion.

Continuing first with the HFD data, Figure 4a displays concentration against mean parity. The smoothed trend should not be taken at face value. Part of the curvature is evidently generated cross-sectionally, because not all countries cover the entire range of mean fertility. While there is much greater variety in the longitudinal structure, it is also evident that the smoothed trend does not contradict this structure. There is some evidence that below replacement fertility, lower fertility is associated with greater, but still underdispersed, concentration. Above replacement level, the direction of the relationship appears reversed.

The HFD only provides a view of the low end of the fertility spectrum. Pushing towards higher mean fertility levels, we turn to examining dispersion by mean parity for the countries from the IPUMS data (Fig. 4b). At above replacement level, this confirms the patterns tentatively suggested by the HFD data, namely a strong positive relationship between dispersion and cohort fertility level: all observations with a cohort fertility below 2 are underdispersed, the situation is mixed at cohort fertility between 2 and 3, and all observations where the cohort fertility exceeds 3 are overdispersed. There is, however, some evidence that the direction of the relationship might change again at very high fertility levels.

This suspicion is confirmed by DHS data (Fig. 4c). While the few observations in the range of mean parity between 2 and 4 support the above positive relationship, the bulk of observations at the high end clearly show a trend reversal. This more comprehensive evidence validates the tentative conclusion of a «reversal of the clear

negative association between concentration and the level of fertility» (Lutz, 1989, 181) at low fertility.

Mentally placing the HFD, IPUMS, and DHS plots of concentration against mean parity alongside each other yields a «slanted s-shaped», or wavelike pattern overall. At least cross-sectionally, there is a pronounced dip in dispersion around parity 2 and a pronounced peak around 4 to 5. Mindful of the dangers of «reading history sideways» (Thornton, 2001), by itself this does not support a dynamic, longitudinal interpretation in terms of a wavelike «path» of concentration during the course of the fertility transition. However, evidence for the robustness of this pattern is provided by the fact that, as shown in the following, it can be seen to hold across sources of variation in fertility other than cross-country differences.

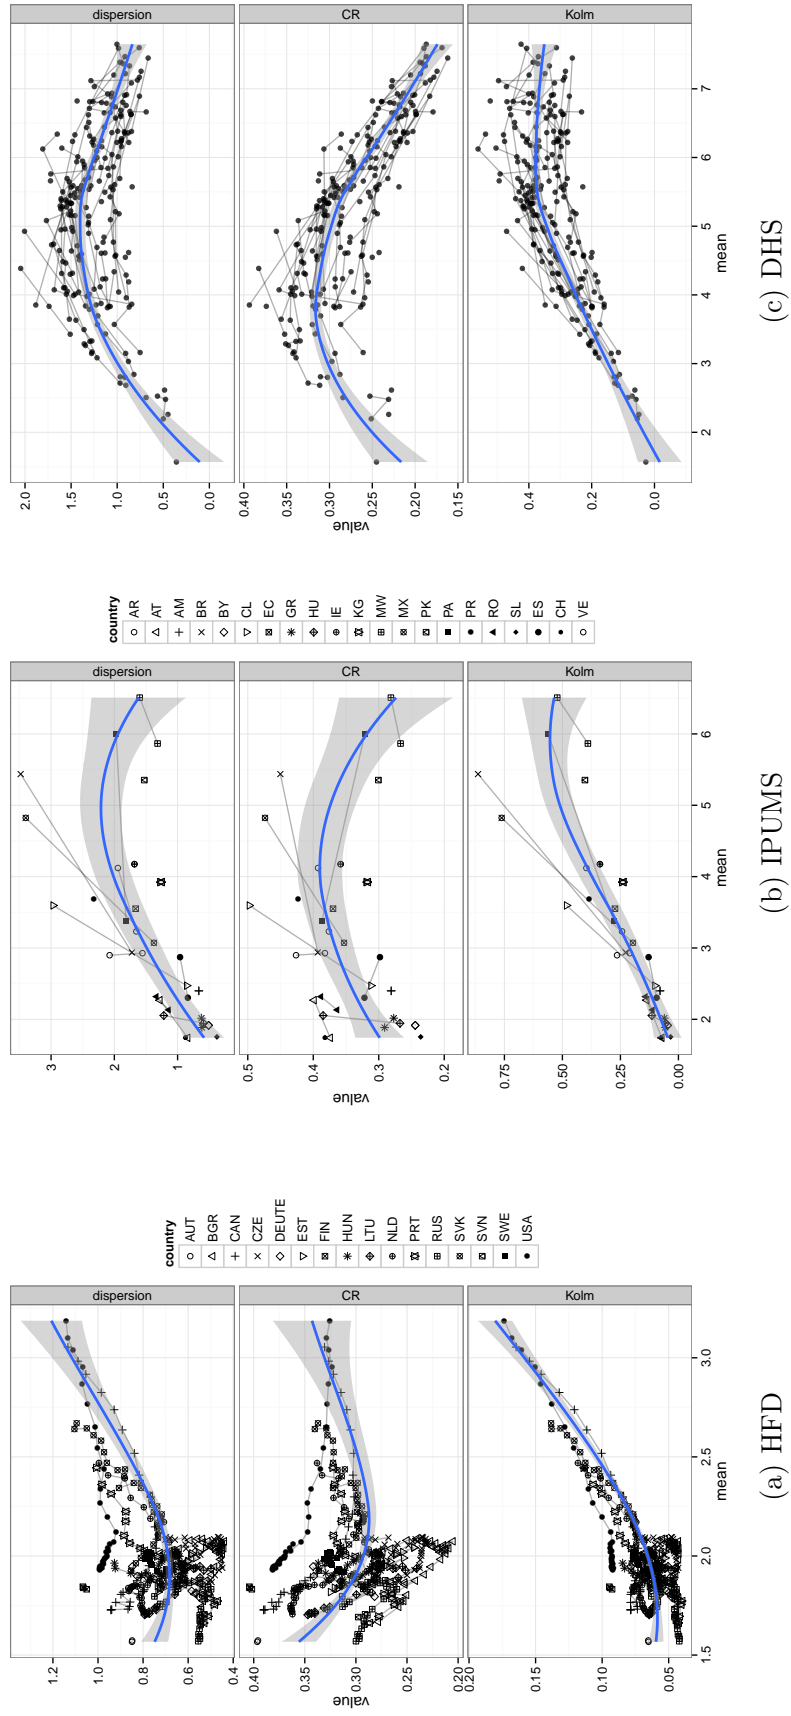


Figure 4: Concentration by mean completed parity. Solid line: weighted regression smoother. Data: as labelled, women aged 40–44

3.2.3 Across Subpopulations

Both the average parity and its concentration vary not only across countries and over national cohorts, but also between socioeconomic subgroups. To investigate this dimension, Figure 5 plots the different concentration measures against mean completed parity for women aged 40–44 in subgroups defined by educational attainment (4 levels) and occupation status (11 categories). The analysis includes all countries for which the necessary data are available in the IPUMS repository, specifically 14 countries from 3 continents as labelled in the figure.

These results demonstrate why the search for a universal linear relationship between the level and concentration of fertility has been inconclusive. For one, there is a non-linear U-shaped or J-shaped relationship that makes the possibility of finding a linear relationship highly dependent on the specific range of overall fertility examined. Moreover, we see that a quantitative assessment of the relationship depends on the choice of concentration measure. For example, it could reasonably be claimed that in Latin America and the Caribbean, at low fertility levels, the relationship is flat (dispersion), negative (CR), or positive (Kolm).

However, it is also evident that viewed in a different way, the pattern is robust to the choice of concentration measure, even though these are not ones that are intrinsically related (as the CR, havehalf, and halfhalf measures are), but are derived from *distinct* notions of distributional concentration. In particular, in all regions and by any measure, the relationship between level and concentration of fertility is: a) essentially convex, b) with a trend break around roughly parity 2, that is c) sharpest in Europe (making the curve almost piecewise linear), and smoothest in Latin America and the Caribbean, with Asia in between. In other words, the relationship between concentration and mean fertility is more negative / less positive at low fertility and more positive or neutral at higher overall fertility.

Applying a suite of conceptually distinct concentration measures has already paid off in demonstrating that this pattern is robust. The CR and Kolm index only allow us to note the existence of this pattern, but offer no deeper theoretical implications. The benefit of considering statistical dispersion specifically arises from the fact that it can actually generate a hypothesis to explain the location of the trend break, because it maps to a behavioural model through the Poisson distribution. This is discussed in more detail in Section 3.3 below.




Figure 5: Dispersion and mean completed fertility, disaggregated by education, occupation. Solid line: weighted regression smoother. Data: IPUMS, women aged 40–44.

3.2.4 Over the Life Course

The need to take seriously the pronounced dip in dispersion of parity around 1.8 or so in European countries that is evident in the above cross-section of socio-economic subgroups is underlined further by the fact that this phenomenon can be shown to also occur using an entirely different source of variation. In particular, instead of comparing subgroups, Figure 6 displays the development of fertility dispersion over mean parity as a single cohort ages.

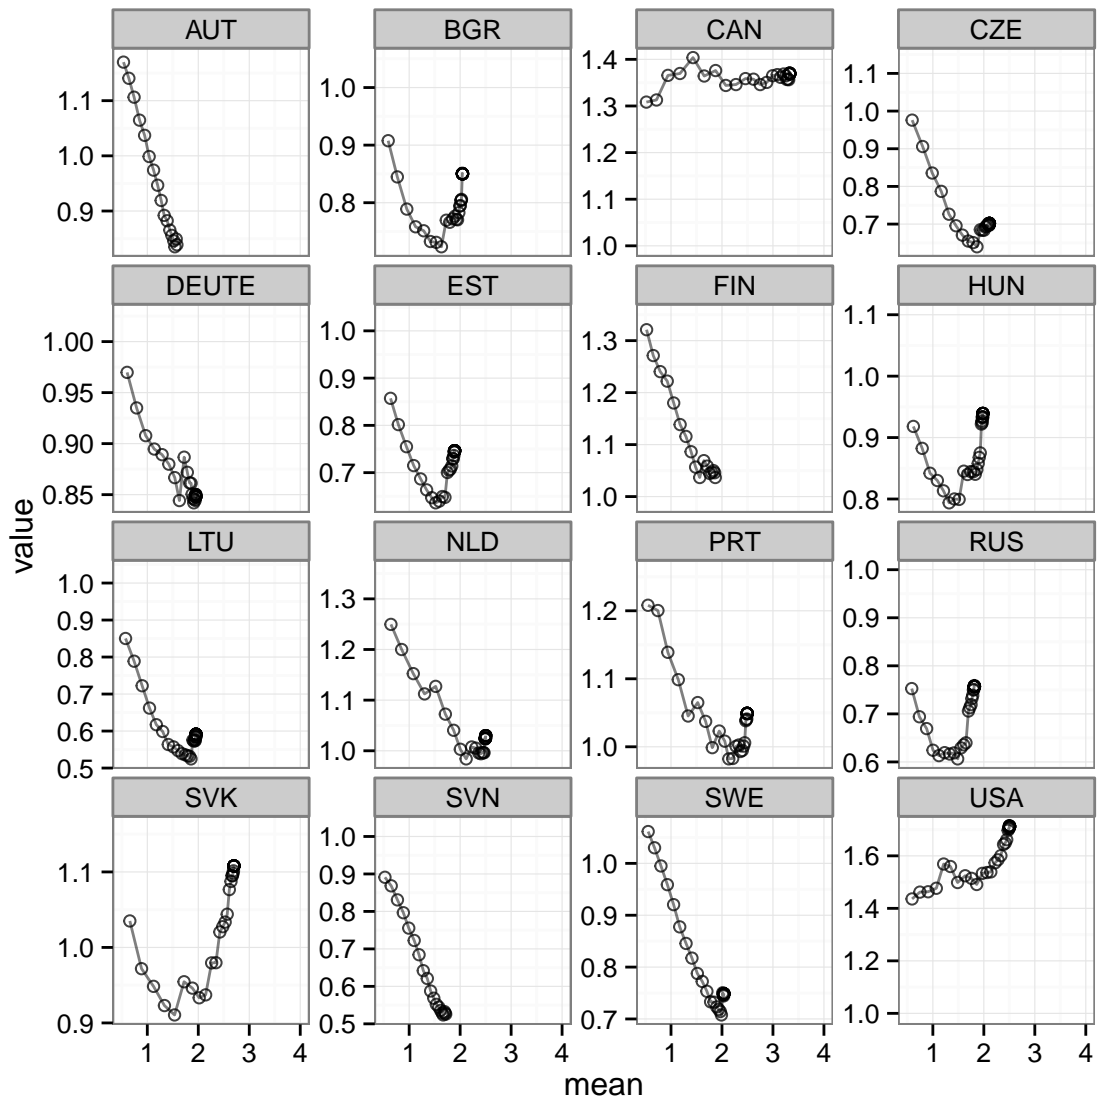


Figure 6: Dispersion by mean parity over increasing age for the earliest available cohort in each country's HFD parity data.

The initial increase is mathematically inevitable, since the starting point for disper-

sion at the start of the reproductive career of a cohort is zero, as constant zero parity to begin with implies zero variance. Evidently, a widespread pattern is that, following the initial upshot, dispersion decreases consistently as mean parity approaches values of 1.5 to 2, at which point it begins to increase. Since even completed fertility at age 40 does not increase much beyond this level for the countries in question, it is not possible to see whether the positive relationship continues for higher fertility levels, although the example of Slovakia suggests it might well do so. In other words, for many European countries, the U-shaped relationship with changes in mean fertility over age of a single cohort correspond closely to the purely cross-sectional relationship across sub-groups shown above.

Again, this is consistent with an explanation in terms of targetted behaviour. Dispersion decreases as those who have already reached their target parity stop their childbearing, and the ‘late-comers’ catch up, and increases again once those with a target of two or one have met it and the heterogeneity among the high-parity individuals begins to show as they keep progressing.

What is evident is that the USA and Canada follow an entirely different pattern, namely one of essentially monotonically increasing dispersion. The extent to which this reflects differences in the volume of immigration cannot be answered based on the present data. This also raises the question whether we should necessarily expect the finding that expected family size strictly diverges over age between different groups in the US (Hayford, 2009) to generalise to European countries.

3.3 Discussion

A priori there is no reason to expect the relationship between the level and dispersion of fertility to be consistent across different sources of fertility variation. Nevertheless, a number of general patterns can be observed across the above analyses. A general pattern seen above is that high levels of mean fertility tend to be associated with overdispersion, and low fertility with underdispersion. In terms of gradients, and taking the full range of mean fertility into account, then as mean parity declines, there is a tendency for dispersion to follow a «slanted s-shaped» pattern. In stylised terms, it first increases from equidispersion to overdispersion, then declines into underdispersion, then increases back up to equidispersion.

Overdispersion at high levels of overall fertility is, perhaps, unsurprising. One cause of overdispersion is (unobserved) heterogeneity, and it seems plausible that sources of heterogeneity with respect to fertility, be they fecundity, partnership status, for

example, will have more opportunity to exert their cumulative influence the closer the fertility experience moves to the frontier of natural fertility. Some women experience circumstances that would allow them to have ten children while others do not, but anyhow this does not matter if all women choose to only have one or two. In terms of the dynamics, the first stage of fertility decline is likely to be limited to an avant-garde, providing yet another source of heterogeneity.

The observation that low levels of mean fertility are associated with underdispersion can be explained by any number of (mutually non-exclusive) explanations. Statistically speaking, underdispersion could arise as a consequence of any of the following: a positive inter-personal correlation in terms of the child count («if others have more/fewer children, so will I»), a parity progression rate that is negatively related to the parity already achieved («the more children I already have, the less likely I am to have more»), or a parity progression rate that is positively (!) correlated with the waiting time since the last birth («the longer it's been since the last birth, the more likely I am to have another child»). The two latter mechanisms both diminish the occurrence of «runaway» parity, where some women tend towards extremely high birth counts, while others are «stuck» at low levels, a constellation that would elevate the variance, and hence the dispersion, observed. While highly stylized, both of the latter two models would be compatible with a qualitative explanation in terms of intentional or targetted behaviour. Since an average fertility of around or below replacement level is likely to only be possible as a result of such deliberate behaviour, the association of underdispersion with low mean fertility could be explained in those terms.

The theory proposed by Lutz (1989) is, in a nutshell, that «birth control in certain segments of society leads first to increased concentration, but, once contraception becomes almost universal and highly efficient, concentration tends to decline because of an approximation of actual family size to desired family size», the latter being assumed to be more even. While it was shown in this study that the original analysis in terms of scale invariant concentration that gave rise to this theory did not by itself warrant a behavioural explanation, the argument does fit well with the dispersion results presented here that do. The pattern of dispersion does indicate an initial increase in heterogeneity at overall fertility levels typical of the initial diffusion of modern contraception, and a decidedly non-stochastic, i.e., targeted, pattern at a level where actual family size and desired family are close.

Returning to the results by socioeconomic subgroup (Section 3.2.3), one possible interpretation of the results is, therefore, that in Europe, subgroups with close to aver-

age parity 2 reach that level by actively aiming for a two-child target, and that groups with lower fertility are experiencing diverse kinds of obstacles that create additional heterogeneity. By contrast, in Latin America, sub-replacement fertility at all levels appears to be distributed with equi-dispersion, i.e., in a way consistent with roughly a Poisson generating process. An interpretation of this might be that there are differences between socio-economic groups, but little differences within them, with respect to factors leading to lower fertility, and that these factors rather than specific targets drive fertility outcomes.

We may therefore speculate about whether the present result suggests that the gap between desired and realised fertility among lowest-low fertility subgroups is greater in Europe than in Latin America. The assessment of «ideal», «desired», or «intended» fertility and the existence of a «gap» in relation to the parity actually realised is a highly nuanced matter that cannot be pursued further here. However, the fact alone that the consideration of dispersion can generate testable hypotheses on such a gap purely from data on realised parity demonstrates the potential value of this approach. In contrast to basing such a conclusion on the CR or havehalf measures, deriving it from the pattern of dispersion can actually be theoretically justified.

4 Conclusion

Examining the phenomenon of fertility concentration through the tri-focal lens of a relative inequality measure, an absolute inequality measure, and statistical dispersion leads to a picture that is more coherent, more parsimonious, and more complete. It is more coherent because conclusions about behavioural change can be based on a measure that has a theoretically well-defined relationship with standard stochastic models of fertility behaviour. It is more parsimonious because in a given empirical context, changes over time can be accounted for in terms of the measure that exhibits the clearest trends and breakpoints; and as we have seen this is not always the standard multiplicatively-invariant measure based on the relative Lorenz curve. Finally, it is more complete because, unlike an indicator set that merely consists of different variations on the multiplicatively-invariant theme, the above set captures strictly more phenomena than any single one of its components. This demonstrably includes patterns that have hitherto gone unnoticed, such as a minimum in parity dispersion close to two children in many European (sub)populations. This confirms that «considering different inequality views can reverse conclusions but may also improve the robustness of some results, and thus give more convincing arguments for policy recommendations»

(Kelly & Bresson, 2007, 43). If a single fertility concentration measure had to be chosen, statistical dispersion would have much to recommend it. It enjoys the double justification that: a) in a rigorous sense it represents an intermediate compromise between relative and absolute views of inequality invariance, neither of which is strictly applicable to parity distributions in their pure axiomatic form, and b) in addition to being an «intermediate» inequality index in its own right, it is mean-independent and a standard parameter in the context of established statistical models for count data.

The methodological broadening of the study of fertility concentration can and should be carried further. It is not clear, for example, that the concept of *inequality* is even the most relevant concept. In the study of income distribution, it is now generally accepted that «the notion of polarization is distinctly different from that of inequality» (Esteban *et al.*, 2005, 4), and has its own measures (see also Wolfson (1994); Esteban & Ray (1994); Duclos *et al.* (2004)). One reasonable interpretation of «rising fertility concentration» is that «people are becoming more likely to *either* stay childless *or* have several children if they do». This intuition would be captured by the notion of polarisation rather than that of inequality, which also addresses the question of modality of the parity distribution (Sullivan, 2005) more directly than concentration measures do, paving the way for further research in this direction.

References

- Aaberge, Rolf, Bjerve, Steinar, & Doksum, Kjell. 2006. Modeling inequality and spread in multiple regression. *Lecture Notes-Monograph Series*, 120–130.
- Aczél, János, & Moszner, Zenon. 1994. New results on scale and size arguments justifying invariance properties of empirical indices and laws. *Mathematical Social Sciences*, **28**(1), 3–33.
- Amiel, Yoram, & Cowell, Frank. 1992. Measurement of income inequality: Experimental test by questionnaire. *Journal of Public Economics*, **47**(1), 3–26.
- Amiel, Yoram, & Cowell, Frank. 1999. *Thinking about inequality: Personal judgment and income distributions*. Cambridge University Press.
- Atkinson, Anthony B. 1970. On the measurement of inequality. *Journal of economic theory*, **2**(3), 244–263.

- Atkinson, Anthony B, & Brandolini, Andrea. 2010. On analyzing the world distribution of income. *The World Bank Economic Review*, **24**(1), 1–37.
- Blackorby, Charles, & Donaldson, David. 1980. A theoretical treatment of indices of absolute inequality. *International Economic Review*, **21**(1), 107–136.
- Bossert, Walter, & Pfingsten, Andreas. 1990. Intermediate inequality: concepts, indices, and welfare implications. *Mathematical Social Sciences*, **19**(2), 117–134.
- Chakravarty, Satya R. 2009. *Inequality, polarization and poverty: Advances in distributional analysis*. Vol. 6. Springer.
- Cowell, Frank. 2009. *Measuring inequality*. Oxford University Press.
- Dahlby, Bev G. 1987. Interpreting inequality measures in a Harsanyi framework. *Theory and Decision*, **22**(3), 187–202.
- del Río, Coral, & Ruiz-Castillo, Javier. 2000. Intermediate inequality and welfare. *Social Choice and Welfare*, **17**(2), 223–239.
- del Ró, Coral, & Alonso-Villar, Olga. 2010. New unit-consistent intermediate inequality indices. *Economic Theory*, **42**(3), 505–521.
- Dorius, Shawn F. 2008. Global demographic convergence? A reconsideration of changing intercountry inequality in fertility. *Population and Development Review*, **34**(3), 519–537.
- Duclos, Jean-Yves, Esteban, Joan, & Ray, Debraj. 2004. Polarization: concepts, measurement, estimation. *Econometrica*, **72**(6), 1737–1772.
- Esteban, Joan, & Ray, Debraj. 1994. On the measurement of polarization. *Econometrica: Journal of the Econometric Society*, 819–851.
- Esteban, Joan, Ray, Debraj, *et al.* . 2005. *A comparison of polarization measures*. Working Paper 310. Barcelona Graduate School of Economics.
- Giroux, Sarah C, Eloundou-Enyegue, Parfait M, & Lichter, Daniel T. 2008. Reproductive Inequality in Sub-Saharan Africa: Differentials versus Concentration. *Studies in Family Planning*, **39**(3), 187–198.

- Goodwin, Dianne G, & Vaupel, James W. 1985. *Curves and have-statistics for ecological analysis of diversity. Part III: Comparison of measures of diversity*. Tech. rept. Working Paper WP-85-91, International Institute for Applied System Analysis, Laxenburg, Austria.
- Hayford, Sarah R. 2009. The evolution of fertility expectations over the life course. *Demography*, **46**(4), 765–783.
- Kelly, LABAR, & Bresson, Florent. 2007. " *Leftist*," " *Rightist*" and *Intermediate Decompositions of Poverty Variations with an Application to China from 1990 to 2003*. Tech. rept. CERDI.
- Kimhi, Ayal. 2011. Comment: On the Interpretation (and Misinterpretation) of Inequality Decompositions by Income Sources. *World Development*, **39**(10), 1888–1890.
- Kolm, Serge-Christophe. 1976a. Unequal inequalities. I. *Journal of Economic Theory*, **12**(3), 416–442.
- Kolm, Serge-Christophe. 1976b. Unequal inequalities. II. *Journal of Economic Theory*, **13**(1), 82–111.
- Krtscha, Manfred. 1994. A new compromise measure of inequality. *Pages 111–119 of: Models and measurement of welfare and inequality*. Springer.
- Lichter, Daniel T, & Wooton, Jillian. 2005. The concentration of reproduction in low-fertility societies: The case of the United States. *The New Population Problem: Why Families in Developed Countries Are Shrinking and What It Means*. Mahwah, NJ, London: Lawrence Erlbaum Associates, 213–224.
- Lutz, W. 1988. *Comparative analysis of completed parity distributions: a global WFS perspective*. Working Paper W P-88-09. International Institute for Applied Systems Analysis (IIASA). also published as: Population bulletin of the United Nations 1989 Nr. 28:25–57.
- Lutz, W. 1989. *Distributional aspects of human fertility: A global comparative study*. London: Academic Press.
- Sellers, Kimberly F., Borle, Sharad, & Shmueli, Galit. 2012. The COM-Poisson model for count data: a survey of methods and applications. *Applied Stochastic Models in Business and Industry*, **28**(2), 104–116.

- Sen, Amartya. 1973. *On economic inequality*. Oxford University Press.
- Shkolnikov, Vladimir M., Andreev, Evgueni M., Houle, René, & Vaupel, James W. 2007. The Concentration of Reproduction in Cohorts of Women in Europe and the United States. *Population and Development Review*, **33**(1), 67–100.
- Shorrocks, Anthony F. 1980. The class of additively decomposable inequality measures. *Econometrica: Journal of the Econometric Society*, 613–625.
- Shorrocks, Anthony F. 1982. Inequality decomposition by factor components. *Econometrica: Journal of the Econometric Society*, 193–211.
- Spielauer, Martin. 2005. Concentration of reproduction in Austria: General trends and differentials by educational attainment and urban-rural setting. *Vienna Yearbook of Population Research*, 171–195.
- Subramanian, Sreenivasan. 2013. *Assessing Inequality in the Presence of Growth: An Expository Essay*. Presented at the ‘Inclusive Growth in Africa: Measurement, Causes, and Consequences’ conference held in Helsinki, Finland, 20–21 September 2013.
- Sullivan, Rachel. 2005. The age pattern of first-birth rates among U.S. women: The bimodal 1990s. *Demography*, **42**(2), 259–273.
- Thornton, Arland. 2001. The Developmental Paradigm, Reading History Sideways, and Family Change. *Demography*, **38**(4), 449–465.
- Wan, Guanghua, & Zhou, Zhangyue. 2005. Income Inequality in Rural China: Regression-based Decomposition Using Household Data. *Review of Development Economics*, **9**(1), 107–120.
- Wolfson, Michael C. 1994. When inequalities diverge. *The American Economic Review*, **84**(2), 353–358.
- Zheng, Buhong. 2007. Unit-Consistent Decomposable Inequality Measures. *Economica*, **74**(293), 97–111.

Working Papers

Buber-Ennsner, Isabella, *Attrition in the Austrian Generations and Gender Survey*, VID Working Paper 10/2013.

De Rose, Alessandra and Maria Rita Testa, *Climate Change and Reproductive Intentions in Europe*, VID Working Paper 09/2013.

Di Giulio, Paola, Thomas Fent, Dimiter Philipov, Jana Vobecká and Maria Winkler-Dworak, *State of the Art: A Family-Related Foresight Approach*, VID Working Paper 08/2013.

Sander, Nikola, Guy J. Abel and Fernando Riosmena, *The Future of International Migration: Developing Expert-Based Assumptions for Global Population Projections*, VID Working Paper 07/2013.

Caselli, Graziella, Sven Drefahl, Marc Luy and Christian Wegner-Siegmundt, *Future Mortality in Low-Mortality Countries*, VID Working Paper 06/2013.

Basten, Stuart, Tomáš Sobotka and Kryštof Zeman, *Future Fertility in Low Fertility Countries*, VID Working Paper 05/2013.

Sharygin, Ethan, *The Carbon Cost of an Educated Future: A Consumer Lifestyle Approach*, VID Working Paper 04/2013.

Winkler-Dworak, Maria and Heiner Kaden, *The Longevity of Academicians: Evidence from the Saxonian Academy of Sciences and Humanities in Leipzig*, VID Working Paper 03/2013.

Feichtinger, Gustav, Alexia Prskawetz, Andrea Seidl, Christa Simon and Stefan Wrzaczek, *Do Egalitarian Societies Boost Fertility?*, VID Working Paper 02/2013.

Muttarak, Raya, *Is it (dis)Advantageous to Have Mixed Parentage? Exploring Education & Work Characteristics of Children of Interethnic Unions in Britain?*, VID Working Paper 01/2013.

Testa, Maria Rita and Stuart Basten, *Have Lifetime Fertility Intentions Declined During the "Great Recession"?*, VID Working Paper 09/2012.

Buber, Isabella, Ralina Panova, and Jürgen Dorbritz, *Fertility Intentions of Highly Educated Men and Women and the Rush Hour of Life*, VID Working Paper 08/2012.

Testa, Maria Rita, Laura Cavalli, and Alessandro Rosina, *The Decision of Whether to Have a Child: Does Couple Disagreement Matter?*, VID Working Paper 07/2012.