



Inequality

The IFS Deaton Review

How families matter for understanding economic inequality

Cezar Santos
Michèle Tertilt

An IFS initiative funded by the Nuffield Foundation



ESRC
CENTRE FOR THE
MICROECONOMIC
ANALYSIS OF
PUBLIC POLICY

How families matter for understanding economic inequality

Cezar Santos (Inter-American Development Bank) and Michèle Tertilt (University of Mannheim)¹

Introduction

Income inequality in the United Kingdom is relatively high. The Gini coefficient of household disposable income was 0.366 in 2019, which is well above the European Union average and one of the highest in the OECD.² As Figure 1 shows, after a period of increasing inequality in the UK during the 1980s and 1990s, the Gini coefficient has stabilised at a relatively high level. Understanding what explains the high level of inequality is an important and much discussed topic. Here we zoom in on the connection between families and economic inequality, with a focus on the relationship between families and inequalities in income and consumption.

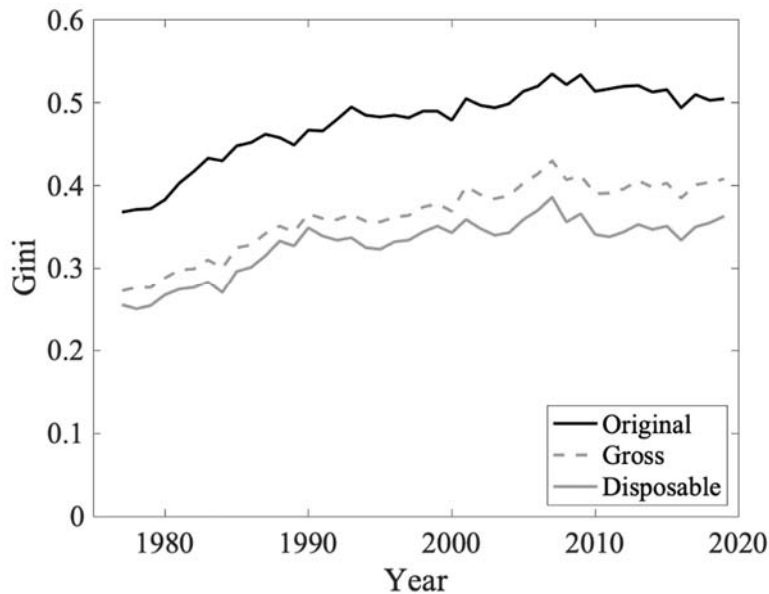
Family structure can, in principle, act as an amplifier or a mitigator of economic inequality – not only because family members can provide insurance to each other, but also because over time and across generations, much of the human capital investments for the next generation (and hence inequality in the future) are made within families. Thus, one would like to know what role UK family structure plays in dampening or amplifying economic inequality. And what about inequality within families – does it exist and is it quantitatively large? Can changes in family structure explain changes in inequality over time? And do cross-country differences in family structure play a role for understanding why inequality in the UK is particularly high?

Two channels stand out in our view on how the existence and structure of families shape economic inequality. First, families are relevant for inequality as they can amplify or mitigate inequality both within and across generations. In particular, a higher degree of assortative mating is related to higher inequality; whereas a more random matching across spouses mitigates inequality. So where people (i.e., future parents) meet and along what dimensions they match matters for the next generation. Similarly, parental investment in children can amplify existing inequalities. Second, a more subtle point that is mostly ignored in economic analysis is that inequality exists even within families. Thinking about within-family inequality requires us to depart from the concept of earnings inequality. Consumption inequality is more relevant when analysing inequality within families, but it is also much harder to measure. Earnings inequality within families, especially between husband and wife, is also interesting of course. But within-family earnings inequality is not very informative when thinking about welfare. Most households pool resources to some extent, so earnings inequality is a poor proxy of consumption inequality within families.

We will now zoom in on the two channels discussed above. First, we discuss the role of families as an amplifier of inequalities in the UK and then we elaborate on the importance of inequality within families.

¹ Financial support from the German Research Foundation (through the CRC-TR-224 project A3 and Leibniz prize TE966/2-1) is gratefully acknowledged. The views expressed in this commentary are those of the authors and do not necessarily represent those of the Inter-American Development Bank.

² OECD (2021), Income inequality (indicator), <https://doi.org/10.1787/459aa7f1-en> (accessed on 8 September 2021).

Figure 1. Gini coefficients for measures of original, gross and disposable income, UK

Note: This figure shows Gini coefficients for different measures of income. Original income includes all sources of income from employment, private pensions, investments and other non-government sources. The receipt of cash benefits is then added to original income to estimate gross income. Finally, direct taxes are subtracted from gross income to estimate disposable income.

Source: Office for National Statistics.

The role of families in propagating/mitigating economic inequality

Which families are formed or dissolved, their educational choices and their resulting income are all determined jointly. An interesting question is how families contribute to amplifying or mitigating inequality. Such contribution can happen in a static or dynamic sense. Statically, as individuals match to form families and make joint decisions, the resulting households can be more or less equal compared with individuals. Dynamically, current investment in children of different households can influence the inequality of future generations.

Inequality within a generation

Inequality across households affects family formation and dissolution (Kiernan, Crossman and Phimister, 2022). How about the other way around? First, marriage itself can have an impact on income inequality. Previous research has documented a 'marriage wage premium'. That is, upon marriage, individuals earn higher wages (see, e.g., Hill, 1979; Loh, 1996; Pilossoph and Wee, 2021). A popular explanation relies on intra-household considerations. With a spouse (usually the wife), one (usually the husband) can specialise in market work. Recent work with US data has shown that both husbands and wives seem to enjoy such a premium nowadays (McConnell and Valladares-Esteban, 2021). So, perhaps there is more to the story than only specialisation. Documenting the existence or not of such a premium in the UK and the underlying mechanisms seems to be an interesting research question.

In this section, instead, we focus on the role of assortative mating. Marriage patterns can have an effect on income inequality. Take the following simple example. Imagine that half of men earn £10 and the other half earn £30. Suppose the same is true for women. If low-earning women always marry high-earning men and high-earning women always marry low-earning men, the income of all households will be £40 and all the inequality at the individual level will disappear at the household level. Alternatively, if low-earning women marry low-earning men and vice versa, half of the households will make £20 and half will make £60. That is, the degree of assortative mating in society matters for income inequality. In this simple example, a more random matching between spouses mitigates individual income inequality.

For marital sorting to have any effect on any outcome, people must first get married. Changes in marriage rates and the fraction of single households over time also play a role in determining income inequality. Households composed of single individuals tend to be poorer than those of married couples. The first reason is mechanical: households with two spouses have more people to work and generate income. Secondly, poorer low-ability individuals may have a harder time on the marriage market and remain single for longer.

For the degree of marital sorting to have an impact on household income inequality, both spouses must work. If only one gender works in the market, the inequality at the household level will reflect the inequality at the individual level. In most countries, during the 20th century, the fraction of women engaging in market work increased considerably. Hence, changes in married female labour force participation interact with marriage patterns to influence income inequality.

Moreover, the gap between the rich and the poor obviously matters, and education is an important determinant of earnings. The fraction of people getting an education and the marriage rates of the more- and less-educated affect the distribution of income as well. Furthermore, the earnings college premium has changed substantially in several countries over the last few decades. Another important price is the gender wage gap that influences how much income wives can contribute to a married household.

To understand the forces behind how families relate to inequality, one must then jointly account for marriage patterns, female labour force participation, and education choices by both genders. Together, these factors shape the level of household income inequality. Greenwood et al. (2016) develop a structural model with such a unified framework and study, among other things, how changes in these variables have affected household income inequality in the US since 1960. They find that an increase in the college premium entices more high-ability people to go to college, which makes household incomes more disperse. Positive assortative mating also heightens inequality. For this latter effect to be operational, however, married women must work in the labour force. Hence, the rise in married female labour force participation also plays a role in generating household income inequality. Quantitatively, taken together, these forces explain a six-point increase in the Gini coefficient (out of the 12 points observed in the US data). Marriage patterns alone account for half of this rise. Such a systematic exercise for the UK seems not to exist. How do changes in educational attainment and the college wage premium interact with marriage patterns? How much does the rise in female labour force participation correlate with marriage and divorce and, ultimately, with income inequality? Nonetheless, some parts of the picture in the UK have been analysed previously.

Chiappori et al. (2020a) present evidence on changes in assortative mating and the implications for household inequality by comparing the birth cohorts of 1945–54 and 1965–74 using data from the UK Labour Force Survey. They find that assortative mating increased among the most populous groups in society.³ In the actual data, the Gini coefficient increased from 0.445 for the older cohort to 0.460 for the younger cohort. The authors then perform two counterfactuals to disentangle the impact of marital sorting and changes in education on household income inequality. First, how would inequality change if the younger cohort had sorted just like the older cohort? Instead of increasing to 0.460, the Gini coefficient would have decreased to 0.425. So, the rise in assortative mating has contributed to an increase in household income inequality in the UK. However, changes in education patterns have contributed to mitigate this. In a second counterfactual, these authors find that the Gini coefficient would have increased even more (to 0.535) had the younger cohort's educational attainment been the same as for the older cohort. That is, the higher educational attainment of the younger cohort acts as an equaliser: with a higher fraction of the population having more education, income inequality declines.

The results described in the previous paragraph provide an important first look at how families have affected income inequality in the UK. There is still much to uncover, however. For instance, Chiappori et al. (2020a) take marital sorting and education as primitives in their analysis. But who marries whom and people's educational attainments are endogenous decisions. Can these decisions be affected by deep causes that themselves affect income inequality, such as changes in the college premium for example? How about the contributions of changes in female labour

³ There is stronger evidence of an increase in assortative mating in the US than in the UK; see Greenwood et al. (2014, 2015) and Chiappori et al. (2020b). For a comparison across the US, UK and a few other countries, see Eika, Mogstad and Zafar (2019).

force participation or the gender wage gap? According to data from the World Bank, female labour force participation increased in the UK from around 45% in the early 1980s to 59% in 2019. The Office of National Statistics reports that the gender pay gap for median gross hourly earnings decreased from 27.5% in 1997 to 15.5% in 2019. Applying a unified framework to understand the overall picture in the UK seems to be an important avenue for future research. With these tools, one can decompose the effects of individual channels (such as changes in female labour force participation, college and gender pay premia, marriage patterns, etc.) and their interactions on income inequality.

Inequality across generations

Parents can prepare children for their lives by supplying resources and support in their upbringing. Parents can invest time in teaching children how to behave and pay attention; they can also read to children and check their school work. They can also invest money in buying books, sending them to better schools, and so on. Throughout a child's life cycle, different inputs can become more or less important. The previous literature has also estimated complementarity of investments in various stages of a child's life cycle (e.g., Cunha, Heckman and Schennach, 2010). To assess how inequality among parents can be transmitted to the following generations, it is thus important to understand parental influence at different points in time. Kiernan et al. (2022) show that children of low-income families in the UK have worse outcomes in their early years. When children experience more frequent episodes of poverty, they are more likely to underachieve in vocabulary tests and have worse behaviour. They further find that parental time investment, not surprisingly, also matters. Children whose parents read to them more often are more likely to score well in vocabulary tests. But what about other inputs that parents use in their children's upbringing?

One interesting avenue of research is to study how parental investments change across different types of families. Such an analysis seems to be missing for the UK. How much time and money is devoted by parents of different socio-economic status (SES)? How are these investments affected by household composition and how does it affect children's outcomes? We can have a sense of the importance of these questions by looking at US data.

Blandin and Herrington (2022) study how family background affects how prepared children are for college using micro data for the US. They assemble data on expenditures on children's education and time investment on children made by families of different characteristics. Figure 2 reports annual expenditures on goods for 1973 and 2003.⁴ Married households in which at least one of the spouses has a college degree ('Married, college') always spend more on children. Interestingly, while other categories have kept their expenditures more or less constant over time, these high-resource families increased their annual expenditures the most. This statistic then suggests that resources are becoming increasingly more unequal across children. A similar conclusion arises when one looks at time investments (see Figure 3).⁵ Between 1965 and 1985, all types of families spent around four to five hours a week on childcare. In the more recent period (1993–2005), families headed by a single parent without a college degree spent on average about seven hours. Higher-resource families with two parents, of which at least one had a college degree, increase their average child time investment to almost 11 hours. Thus, the gap has widened substantially in the US. The childcare gap between more- and less-educated parents is also quite pronounced in the UK. Doepke et al. (2023) document that educated mothers and fathers spent about 20% more time on childcare in the UK in 2015 – a larger gap than in many other countries. But has it, as in the US, risen over time? And if so, to what extent has it contributed to rising inequality in the UK? More research to answer such questions would be desirable.


Because family investments in the UK are unequally distributed, a natural follow-up question is what public policy can do to level the playing field. There is a sizeable literature by now asking this question for specific countries. For example, Lee and Seshadri (2019) find that early education subsidies can significantly reduce intergenerational persistence and lower inequality. Brotherhood and Delalibera (2020) analyse the aggregate and distributional impacts of public

⁴ Blandin and Herrington (2022) include a range of expenditures on children that plausibly contribute to human capital formation, including books, toys, games, computers, musical instruments, childcare, primary/secondary school tuition and tutoring.

⁵ The child time investment measure in Blandin and Herrington (2022) includes time caring for infants (under 5 years old) and older children (up to age 18), medical care for children, playing with children, supervising and helping with homework, reading and talking with children, and 'other' childcare.

school and college investment in Brazil and also find that earlier investments (in basic schooling rather than college) can substantially decrease inequality. How then is public investment in education allocated in the UK? Are schools able to level this playing field? How does the value-added of UK schools that serve mainly lower SES families compare with those that serve more affluent neighbourhoods? These are open questions that deserve further attention.

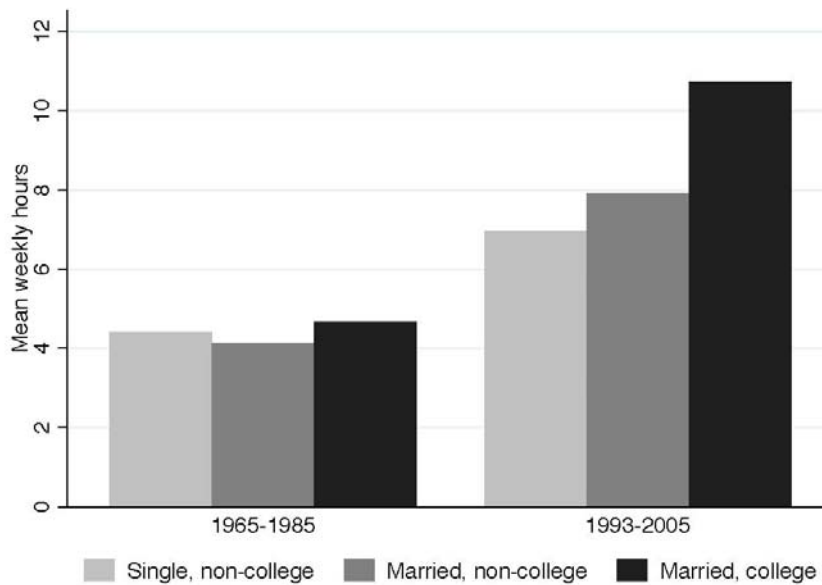
Figure 2. Trends in annual expenditures per child, by family type



Note: Figure shows mean expenditures per child by family type in the 1973 and 2003 waves of the CEX.

Source: Blandin and Herrington (2022).

Figure 3. Trends in child time investment per parent, by marital status and education



Note: Figure displays mean parent time spent with children per week by parent marital status and education from several waves of the American Heritage Time Use Survey.

Source: Blandin and Herrington (2022).

Inequality within families

A recent body of literature has also stressed the importance of inequality *within* families.⁶ At the end of the day, we care about inequality across people, not households. However, inequality is typically measured on the household level. Yet, it has become quite clear by now that not all household or family members necessarily consume the same amount, and hence disparities within families do exist.

To assess inequality within families, one must naturally go beyond income inequality. Many single-earner families exist, but clearly it would not make sense to declare all such families as highly unequal simply because one person has high earnings, while the other has zero earnings. Both spouses as well as their children will obviously consume out of the single income. At the same time, just because all family members consume out of the same income does not mean that there is no inequality within such a family. Several dimensions of inequality within families come to mind. First, not all children within a family receive the same resources, often preference is given to male children. Second, access to resources often also differs across adult household members. For example, Lise and Yamada (2019) document in Japanese data that women, on average, have less private consumption than men and that there is a large dispersion in consumption shares across households. In addition to unequal consumption, gender gaps in leisure and different measures of psychological well-being have been documented.⁷ One could take the stand that these issues are best left to psychologists and perhaps sociologists, but we believe that thinking about these questions is important for economists as well. Welfare concepts in economics are defined for individuals, not families. Yet, in empirical analyses, these concepts are often applied as if a household (or family) was an individual. While this is done for practical reasons, such a choice is not innocuous. A recent small body of literature on inequalities within families by and large finds that consumption inequality is understated if inequality within families is ignored.

So what are the reasons to believe that family members do not share resources equally? Often it is appealed to altruism and insurance to argue that within-family inequality must be negligible. However, there are several reasons why this does not have to be the case. First, even though perfect insurance implies *constant* consumption shares over time, it does not imply *equal* shares. Consider a family consisting of one male m and one female f , who receive earnings each period of w_{mt} and w_{ft} , respectively. For simplicity, assume there is no storage and no capital markets, so that all income needs to be consumed each period. The full insurance problem is

$$\max_{c_{mt}, c_{ft}} \sum_t \lambda_m u(c_{mt}) + (1 - \lambda_m) u(c_{ft}) \quad (1)$$

$$c_{mt} + c_{ft} \leq w_{mt} + w_{ft} \quad \forall t. \quad (2)$$

The solution is

$$\frac{u'(c_{mt})}{u'(c_{ft})} = \frac{1 - \lambda_m}{\lambda_m}$$

for all t . So clearly, even though consumption shares are constant over time, there is no reason for them to be equal. Assuming log utility, we find $c_{mt} = \lambda_m (w_{mt} + w_{ft})$ and $c_{ft} = (1 - \lambda_m) (w_{mt} + w_{ft})$. Thus, a key question is how λ_m is determined. In a full equilibrium model with endogenous marriage, λ_m would adjust to clear the marriage market. Thus, λ_m would be determined by marriage market conditions such as the sex ratio of singles, but outside options (such as how attractive it is to remain single, which may differ by gender) would also matter. As long as there is

⁶ See Chiappori and Meghir (2015) for an excellent survey.

⁷ For example, Oreffice and Quintana-Domeque (2021) document that women's mental health was worse than men's on average in the UK during the COVID-19 pandemic and that the gender gap in housework and childcare increased during this time period.

a gender wage gap, for example, it seems quite plausible that $\lambda_m > 0.5$ as the outside option of remaining single is better for men than women.

So far we have argued that even a perfect insurance world allows for consumption inequality across spouses. Yet, in reality, perfect insurance will not always hold for several reasons. The main reason for perfect insurance to break down in couples is limited commitment (Mazzocco 2007). As long as either partner can leave the union unilaterally (which is possible today in most countries), partners will only stay in the union if at any point in time staying is better than leaving. Thus, there will be a participation constraint each period, leading to the possibility of λ_m changing over time. One example could be that one spouse gets a big pay raise and then feels entitled to purchase a new motorcycle or another large private consumption item. Another example would be that one spouse falls in love with someone else, but in the end the couple stays together, yet the spouse that wanted to stay together has to compensate the other in form of more private consumption. Such behaviours can be explained in a limited commitment world. Similarly, private information might be relevant within families as well. Private information about income is perhaps somewhat less plausible, especially if spouses file jointly for taxes, but private information about preferences seems quite plausible. Changes in the preferences (perhaps due to the arrival of a child or because of changes in health conditions) would then not be fully insured and would lead to changes in relative consumption (Doepke and Tertilt, 2016). Thus, both limited commitment and private information can cause inequality within families and possibly changes therein over time.

But is such *within-family inequality* economically relevant? Is it quantitatively large for the UK specifically and has it changed over time? While first studies exist suggesting that within family-inequality is meaningful and sizeable (Lise and Seitz, 2011; Chiappori and Meghir, 2015; Lise and Yamada, 2019; Obermeier, 2019, 2021; Lechene, Pendakur and Wolf, 2021), there are important measurement and data collection issues that are the main reason why more progress in answering these questions has not been made. To assess whether inequality exists within a family, one needs to know how much each family member consumes. Relevant dimensions would include consumption goods, but also leisure. Let us start with leisure as it is relatively easy to measure. Many modern time use surveys include detailed questions on time use of each individual family member, so that a measure of leisure for each member can be calculated in principle. In practice, however, it is not always easy to assign categories to the three broad categories of market work, home production and leisure. For example, is cooking in the category of leisure or home production? Similarly, is time with children leisure or home production? The answers to such questions will likely differ across individuals and possibly by gender. There is a conceptually clean way of distinguishing the two though. A person who considers cooking a chore would be happy to delegate this (at no cost) to someone else. However, a person who considers cooking a leisure activity would not. Thus, in principle, one could add questions of this sort to time use surveys to assess whether each individual considers a particular time use category as leisure or home production. In fact, some time use surveys include questions on how much an individual enjoys a particular activity, which could possibly be used to make such a distinction. It would not fully resolve the issue though as many parents consider childcare a chore, yet even at zero cost would not outsource all of it, as they consider parental time an important input for child development. In this case, perhaps one could ask whether the respondent would prefer the partner to spend more time on childcare.

Now let's move to consumption, which is even trickier to measure on a per person level than leisure. For starters, much of consumption within a family is a public good, such as a house that is consumed by the entire family together, similarly cleaning services, a TV or a car.⁸ Second, how does one attribute private goods to individual family members? Empirically, the concept of 'assignable goods' is typically used. But most goods are unfortunately not assignable. This leaves much research with using data on clothing (where several surveys ask separately about male, female and children's clothing) to try to estimate consumption shares. Clearly, many strong assumptions are needed to extrapolate from clothing shares to consumption shares. If spouses evaluate different types of consumption goods differently (it seems quite plausible that women care more about clothing than men), then such procedures can easily lead to flawed results. So what is the alternative? One could ask survey respondents more explicitly about their own

⁸ Of course, a TV is not perfectly public if family members want to watch different channels, and similarly a car is not perfectly public if people want to drive to different places, but many goods consumed in the family have a large public goods component. See Salcedo, Schoellmann and Tertilt (2012) for a discussion.

consumption in individual categories. So far there are only very few surveys that attempt this (two notable exceptions are the Longitudinal Internet studies for the Social Sciences in the Netherlands and the Japanese Panel Survey of Consumers). Even though it is far from trivial to design such questions well, this should not prevent us from trying. The Dutch and Japanese data have been used successfully already, so even though they are far from perfect, there is much to learn from these surveys. Also, it is likely that the questions could be further fine-tuned and improved. For example, rather than asking about the entire consumption basket to be assigned to individuals, one could ask about whether the family had incurred large expenses lately that benefited one person only (such as consumer electronics, a motorcycle, a restaurant dinner, an individual vacation) and then ask to specify the amount and family member.

Yet, even without collecting better data, one can use theory together with data on variables relevant for bargaining weights to try to estimate within-family inequality. Some obvious variables that come to mind include the sex ratio, the gender wage gap, the generosity of the welfare state (affecting, in particular, the income of single mothers) and even the progressivity of the tax system.⁹ Have these factors changed over time in the UK? And if so, have they affected within-family inequality in the UK over time or in comparison with other countries? Lise and Seitz (2011), for example, argue that the large increase in inter-household inequality over the 1970–2000 time period in the UK was partly compensated by a sizeable decrease in intra-household inequality so that total inequality increased by far less than typical measurements show. Can similar statements be made for the last 20 years? As discussed above, the gender wage gap (for median gross hourly earnings) was almost halved between 1997 and 2019. Theory would predict that within-family inequality decreased in response. Is there any evidence that this actually happened? And does it mean that overall inequality has actually declined over the last 20 years? If so, this would be contrary to what one would conclude from simply looking at the standard Gini coefficient, which has remained quite flat over the same time period.

Final remarks on COVID-19, families and inequality

There are several reasons to believe that inequality increased during the COVID-19 pandemic for reasons related to families. School closures led to a major childcare challenge for most families with children, with women taking on the majority of the additional childcare responsibilities. This has been documented widely for both the UK and other countries. Thinking through carefully what this means for inequality in consumption and welfare, and quantifying the effect, would be very interesting.

Specifically, both of the channels laid out in this paper likely interact with the ongoing COVID-19 pandemic. First, the school closures heavily affected children's ability to learn, with the largest impact on the most disadvantaged families. This should, in principle, increase across-family inequality in the long run. Several recent papers use calibrated models to quantify the impact of the pandemic-related education disruptions in the US and find sizeable effects on future inequality (Fuchs-Schündeln et al., 2020; Agostinelli et al., 2022). As far as we know, no such analysis exists for the UK yet.

Within-family inequality was likely affected as well. It has been widely documented that inequality in the gendered division of childcare rose during the COVID-19 pandemic in the UK (Sevilla and Smith, 2020; Andrew et al., 2020). Changes in the division of childcare generally should have implications for inequality in the labour market. Hence, one would have expected mothers to reduce hours worked more than fathers. Interestingly, while most countries experienced a large 'shecession' driven partly by mothers reducing hours worked, there appears to be no sizeable gender gap in employment and hours reductions in the UK (Alon et al., 2022; Hupkau and Petrongolo, 2020). Whether this means that there were no changes in within-family inequality is an open question. Oreffice and Quintana-Domeque (2020) document an increase in gender gaps in mental health during the pandemic in the UK. It seems plausible that the increase in the childcare burden without a compensating decrease in work hours led to a sizeable reduction in leisure for working mothers, hence widening inequality in well-being across genders in the UK. These remain interesting angles to pursue in future research.

⁹ How the progressivity of the tax system matters for within-family inequality is nicely laid out in Obermeier (2019).

References

- Agostinelli, F., Doepke, M., Sorrenti, G., and Zilibotti, F. (2022), 'When the Great Equalizer Shuts Down: Schools, Peers, and Parents in Pandemic Times', *Journal of Public Economics*, 206, 104574.
- Alon, T., Coskun, S., Doepke, M., Koll, D., and Tertilt, M. (2022), 'From Mancession to Shecession: Women's Employment in Regular and Pandemic Recessions', in M. S. Eichenbaum and E. Hurst (eds), *NBER Macroeconomics Annual 2021*, Volume 36, Chapter 2, Chicago, IL: University of Chicago Press, 83–151.
- Andrew, A., Cattan, S., Costa Dias, M., Farquharson, C., Kraftman, L., Krutikova, S., Phimister, A., and Sevilla, A. (2020), 'The Gendered Division of Paid and Domestic Work under Lockdown', IZA Discussion Paper 13500.
- Blandin, A., and Herrington, C. (2022), 'Family Heterogeneity, Human Capital Investment, and College Attainment', *American Economic Journal: Macroeconomics*, 14 (4), 438–78.
- Brotherhood, L., and Delalibera, B. R. (2020), 'Minding the Gap between Schools and Universities', *Journal of Economic Dynamics and Control*, 120, 104010.
- Chiappori, P.-A., Costa-Dias, M., Crossman, S., and Meghir, C. (2020a), 'Changes in Assortative Matching and Inequality in Income: Evidence for the UK', *Fiscal Studies*, 41, 39–63.
- Chiappori, P.-A., Costa-Dias, M., Crossman, S., and Meghir, C. (2020b), 'Changes in Assortative Matching: Theory and Evidence for the US', Institute for Fiscal Studies (IFS) Working Paper W20/10.
- Chiappori, P.-A., and Meghir, C. (2015), 'Intrahousehold Inequality', in A. B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*, Volume 2, Chapter 16, Amsterdam: Elsevier, 1369–418.
- Cunha, F., Heckman, J. J., and Schennach, S. M. (2010), 'Estimating the Technology of Cognitive and Noncognitive Skill Formation', *Econometrica*, 78, 883–931.
- Doepke, M., Hannusch, A., Kindermann, F., and Tertilt, M. (2023), 'The Economics of Fertility: A New Era', in S. Lundberg and A. Voena (eds), *Handbook of Family Economics*, Chapter 4, Amsterdam: North-Holland.
- Doepke, M., and Tertilt, M. (2016), 'Families in Macroeconomics', in J. B. Taylor and H. Uhlig (eds), *Handbook of Macroeconomics*, Volume 2, Chapter 23, Amsterdam: North-Holland, 1789–891.
- Eika, L., Mogstad, M., and Zafar, B. (2019), 'Educational Assortative Mating and Household Income Inequality', *Journal of Political Economy*, 127, 2795–835.
- Fuchs-Schündeln, N., Krueger, D., Ludwig, A., and Popova, I. (2020), 'The Long-Term Distributional and Welfare Effects of Covid-19 School Closures', National Bureau of Economic Research (NBER) Working Paper 27773.
- Greenwood, J., Guner, N., Kocharkov, G., and Santos, C. (2014), 'Marry Your Like: Assortative Mating and Income Inequality', *American Economic Review*, 104 (5), 348–53.
- Greenwood, J., Guner, N., Kocharkov, G., and Santos, C. (2015), 'Corrigendum to Marry Your Like: Assortative Mating and Income Inequality', Mimeo, <https://www.jeremygreenwood.net/papers/ggksPandPcorrigendum.pdf>.
- Greenwood, J., Guner, N., Kocharkov, G., and Santos, C. (2016), 'Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation', *American Economic Journal: Macroeconomics*, 8 (1), 1–41.
- Hill, M. S. (1979), 'The Wage Effects of Marital Status and Children', *Journal of Human Resources*, 7, 579–94.

Hupkau, C., and Petrongolo, B. (2020), 'Work, Care and Gender during the COVID-19 Crisis', *Fiscal Studies*, 41, 623–51.

Kiernan, K., Crossman, S., and Phimister, A. (2022), 'Families and Inequalities', IFS Deaton Review of Inequalities.

Lechene, V., Pendakur, K., and Wolf, A. (2021), 'OLS Estimation of the Intra-Household Distribution of Expenditure', CEPR Discussion Paper 15955.

Lee, S. Y. (T.), and Seshadri, A. (2019), 'On the Intergenerational Transmission of Economic Status', *Journal of Political Economy*, 127, 855–921.

Lise, J., and Seitz, S. (2011), 'Consumption Inequality and Intra-Household Allocations', *Review of Economic Studies*, 78, 328–55.

Lise, J., and Yamada, K. (2019), 'Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use', *Review of Economic Studies*, 86, 2184–219.

Loh, E. S. (1996), 'Productivity Differences and the Marriage Wage Premium for White Males', *Journal of Human Resources*, 31, 566–89.

Mazzocco, M. (2007), 'Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment', *Review of Economic Studies*, 74, 857–95.

McConnell, B., and Valladares-Esteban, A. (2021), 'On the Marriage Wage Premium', Unpublished Manuscript, University of Southampton.

Obermeier, T. (2019), 'The Marriage Market, Inequality and the Progressivity of the Income Tax', Collaborative Research Center (CRC) TR 224, Discussion Paper 109.

Obermeier, T. (2021), 'Individual Welfare Analysis: What's the Role of Intra-Family Preference Heterogeneity?', Unpublished Manuscript, Institute for Fiscal Studies.

Oreffice, S., and Quintana-Domeque, C. (2020), 'Gender Inequality in COVID-19 Times: Evidence from UK Prolific Participants', *Journal of Demographic Economics*, 87, 261–87.

Pilossoph, L., and Wee, S. L. (2021), 'Household Search and the Marital Wage Premium', *American Economic Journal: Macroeconomics*, 13 (4), 55–109.

Salcedo, A., Schoellmann, T., and Tertilt, M. (2012), 'Families as Roommates: Changes in U.S. Household Size from 1850 to 2000', *Quantitative Economics*, 3, 133–75.

Sevilla, A., and Smith, S. (2020), 'Baby Steps: The Gender Division of Childcare during the COVID-19 Pandemic', *Oxford Review of Economic Policy*, 36 (S1), S169–86.