sociological science

What You Need to Know When Estimating Monthly Impact Functions: Comment on Hudde and Jacob, "There's More in the Data!"

Josef Brüderl,[a] Ansgar Hudde,[b] Marita Jacob[b]

a) LMU Munich; b) University of Cologne

Abstract: In life course research, it is common practice to analyze the effects of life events on outcomes. This is usually done by estimating "impact functions." To date, most studies have estimated yearly impact functions. However, Hudde and Jacob (2023) (hereafter H&J) pointed out that most panel data sets include information on the month of events. Consequently, they proposed exploiting this information by estimating monthly impact functions. In this adversarial collaboration, we address two issues regarding H&J's work. First, H&J did not provide sufficient guidance on how to estimate monthly impact functions. We will provide a step-by-step description of how to do so. Second, the procedure H&J proposed for smoothing monthly estimates produces confidence intervals (CIs) that are likely too narrow. This can lead to misleading conclusions. Therefore, we suggest using more appropriate bootstrapped CIs.

Keywords: dynamic treatment effects; event study; panel data; life course; happiness; motherhood

Reproducibility Package: Stata replication code is available on the Open Science Framework (OSF), https://osf.io/kx9ne/ (file: "Monthly Impact Functions-Replication File.zip"). The replication file includes the prepared pairfam data that we used for all of our analyses. If you would like to reproduce our data preparation (also included in the replication file), you can order the pairfam data at https://www.pairfam.de/en/data/data-access/.

Citation: Brüderl, Josef, Ansgar Hudde, Marita Jacob. 2025. "What You Need to Know When Estimating Monthly Impact Functions: Comment on Hudde and Jacob, "There's More in the Data!"" Sociological Science 12: 862-870.

Received: May 16, 2025

Accepted: August 31, 2025

Published: December 4, 2025

Editor(s): Arnout van de Rijt, Kristian B. Karlson

DOI: 10.15195/v12.a34

In life course research, it is common practice to analyze the (causal) effects of life events (e.g., marriage, divorce, parenthood, and unemployment) on outcomes (e.g., earnings, well-being, and health). When panel data are available, this is typically done by estimating so-called "impact functions" (Ludwig and Brüderl 2021). An impact function provides the time-varying effect of an event on an outcome (also called "dynamic treatment effect" in the event-study literature [see Miller 2023]). For example, one could study the effect of the birth of a first child on mothers' happiness in the year before the birth (anticipation effect) and several years after the birth. To date, most studies estimating impact functions have used yearly event time dummies included in a fixed-effect regression. This is because most panel data are collected annually.

Hudde and Jacob (2023) (hereafter H&J) rightly pointed out that many panel data contain more information about the timing of life events. Often not only the year but also the month of the event (and the month of the interview) is available, which allows for the estimation of more fine-grained monthly impact functions. Monthly impact functions provide more information than annual impact functions, and this is clearly an advantage.

However, the first author of this comment pointed out to H&J that their approach could be improved in two ways. He and H&J decided to join forces in an adversarial collaboration. The final result of this collaboration is described in this comment, which has been approved by all three authors.

First, if a researcher decides to estimate monthly impact functions, he/she will not find sufficient guidance in H&J on how to do so. In fact, H&J does not display a single monthly impact function (only predicted outcome curves). Next (and in the accompanying replication file), we provide a step-by-step description of how to estimate monthly impact functions.

Second, the smoothing procedure proposed by H&J for smoothing the monthly estimates produces confidence intervals (CIs) that are likely to be too narrow. The reason is that the smoothing algorithm used by H&J assumes that the observations are data. In fact, they are estimates (from a fixed-effect regression) that are themselves uncertain. The smoothing algorithm does not account for this additional uncertainty and underestimates the uncertainty of the smoothed impact function. The CI of the smoothed impact function is therefore likely to be too narrow, which can lead to misleading conclusions. Therefore, we suggest using more appropriate bootstrapped CIs.

Empirical Application: Effect of a First Birth on Mothers' Happiness

Let us illustrate these points with an empirical example. We examine the effect of the birth of a first child on mothers' happiness, one of the applications that H&J use to demonstrate the usefulness of monthly impact functions. H&J used the German Socio-Economic Panel (SOEP). Here, we use data from the German Family Panel (pairfam). We follow the design of Ludwig and Brüderl (2021), but in addition to estimating an annual impact function, we also estimate a more fine-grained monthly impact function.

The German Family Panel is a nationwide longitudinal study of initially more than 12,000 randomly selected individuals from the birth cohorts of 1971–1973, 1981–1983, and 1991–1993. The pairfam started in 2008/2009 with approximately 1-h face-to-face interviews. Respondents were contacted annually in subsequent waves. A detailed description of the study can be found in Huinink et al. (2011). Release 14.2 is used for this analysis (Brüderl et al. 2024). We restrict the analyses to waves 1–11, which cover the observation period from 2008 to 2019, to avoid problems related to the onset of the pandemic during wave 12.

The estimation sample is constructed according to the recommendations of Ludwig and Brüderl (2021). Only never-treated women, that is, those who have never given birth before pairfam wave 1, are included. For first-time mothers, the observation window is censored at the second pregnancy/birth. The central variable for estimating the impact functions is the time an interview was conducted before/since the event occurred (event time). The event time is calculated based on the interview year and month as well as the birth year and month. Negative event time values indicate that the interview was conducted before the birth. Positive values indicate that the interview was conducted after the birth. As done by H&J, event time is truncated at 36 months before and after birth. Thus, in our study, event time can take on values $\{-36, \ldots, -1, 1, \ldots, 36\}$ (we exclude 0 to avoid confusion).

Note that the estimation sample also includes a control group (women who did not give birth during the pairfam observation period). This is advisable because it provides more power to estimate the age trajectory (which provides the counterfactual for the impact function). In total, the estimation sample consists of 2,977 women, of whom 507 gave birth to a first (biological) child. These women contributed 18,635 person-years.

The outcome variable, life satisfaction (happiness), is measured by the question: "Now I would like to ask about your general satisfaction with life. All in all, how satisfied are you with your life at the moment?" Responses are recorded on an 11-point scale ranging from "very dissatisfied" (0) to "very satisfied" (10).

Our estimand is the total causal effect of the birth of a first child on mothers' happiness. Therefore, we must control for potential confounders to identify the total causal effect.[1] Potential confounders are variables that are thought to influence both the treatment and the outcome, in our case first birth and life satisfaction. The most important confounder in our context is age. The observed age range is 15–47 years. Therefore, the models include age dummies for ages 16–47 (15 is the reference category). As further covariates, we include dummies for relationship status, dummies for subjective health in the past four weeks (even though these could also be considered as mediators rather than confounders), and a dummy for pregnancy (other than the pregnancy that resulted in the focal child).

Fixed-effect regression with cluster-robust standard errors is applied. The following model specification is used:

$$Y_{it} = \alpha_i + \gamma_a + \sum_k \beta_k D^k_{it} + \mathbf{X}'_{it}\delta + \varepsilon_{it},$$

where Y_{it} denotes the happiness of woman i in panel wave t, α_i is the person-specific fixed effects, γ_a is the age-specific fixed effects (age dummies, with one dummy excluded as the reference), $\mathbf{X}'_{it}$ represents the other controls, and ε_{it} is the error term. D^k_{it} is the event time dummies and thus the estimates of β_k provide the impact function.

An important modeling decision is the event time range of the impact function, that is, the range of k. H&J chose for the range -35 to 36, that is, only month -36 is in the reference group. They argue that the pre-event trend helps to identify to what degree the change around the event is indeed due to the event or rather a consequence of other unobserved circumstances or events. Ludwig and Brüderl (2021) advocate a more theoretically driven decision. They argue that only pre-event dummies that plausibly represent an anticipation process should be included (Miller [2023] provides an extensive discussion of these issues). For this methodical exercise, we follow Ludwig and Brüderl (2021) and limit the range of the impact function to -12 to 36 months (we assume that 12 months before birth is the maximum period during which women could anticipate the upcoming birth), that is, $k \in \{-12, \ldots, -1, 1, \ldots, 36\}$. The annual impact function is estimated by grouping the event time months.

Estimation is done with Stata 19. Below only graphical representations of the impact functions are shown: Basically, the estimates of β_k are plotted over the event time -12 to 36. A step function is best suited for comparing annual and monthly impact functions (Fig. 1). For smoothing, a coefficient plot is best (Fig. 2). Numerical results can be obtained by running the replication file.

Monthly Impact Functions Provide More Information

Impact functions following this design are plotted in Figure 1. The annual impact function (red curve) tells us that expectant mothers are 0.30 scale points happier (counterfactual: if the woman had aged one year without giving birth). The effect is statistically significant at the 5 percent level because the 95 percent CI does not include 0 (we use the 5 percent significance level throughout). This is the average effect over the 12 months before the birth. In the year following the birth, happiness is 0.57 scale points higher (the so-called "baby effect"). It drops to 0 in the second year and -0.10 in the third year.

The first part of the innovation of H&J is that they propose not to group the months but to estimate a monthly impact function (if this information is available in the data). This is the blue curve in Figure 1. As one can see, this provides much more information. The most important additional information is that both the anticipation effect and the baby effect are much more short lived. A significant anticipation effect can only be seen in the seven months prior to birth. Similarly, the baby effect is statistically significant only during the first six months after birth. After that the effects are (mostly) small and no longer significantly different from 0. Moreover, the annual estimate of 0.57 clearly masks the large baby effect immediately after birth (1.99 scale points). This is the main advantage of monthly impact functions: The annual estimates are a kind of average of the underlying monthly estimates.[2] As a result, they can hide much of the time variation of the causal effect that can be seen by using monthly estimates.

However, as shown in Figure 1, this increased information comes at a price: the monthly estimates are much more uncertain (compare the red and blue CIs). This is inevitable because they are necessarily based on fewer observations than the annual estimates. Therefore, more detail comes at the cost of higher volatility and more uncertainty. For example, the large baby effect right after birth is quite uncertain: the CI ranges from 0.95 to 3.03. Therefore, there is a trade-off, and each researcher must decide whether the increased detail gained by a monthly impact function is worth the increased uncertainty (this trade-off has also been discussed by H&J). This decision depends on the research question and may not always favor monthly impact functions. Researchers may be interested in getting a rough but more precise picture of the pattern of impact of a life event.

A final note on the identification assumptions for (monthly) impact functions: Their estimation requires an intricate identification assumption, namely that interview timing is not affected by the size of the treatment effect (the effect of the event of interest). If treatment effects are heterogeneous and respondents schedule their interviews according to the value of the treatment effect, the shape of the impact function may be biased. For example, if mothers with a bad birth experience delay their interview by a few months, the baby effect in the first few months will be overestimated, that is, the sudden change after birth will appear too high.[3] Thus, applied researchers should be cautious when interpreting sudden changes in monthly impact functions. Finally, it should be noted that selective interview timing

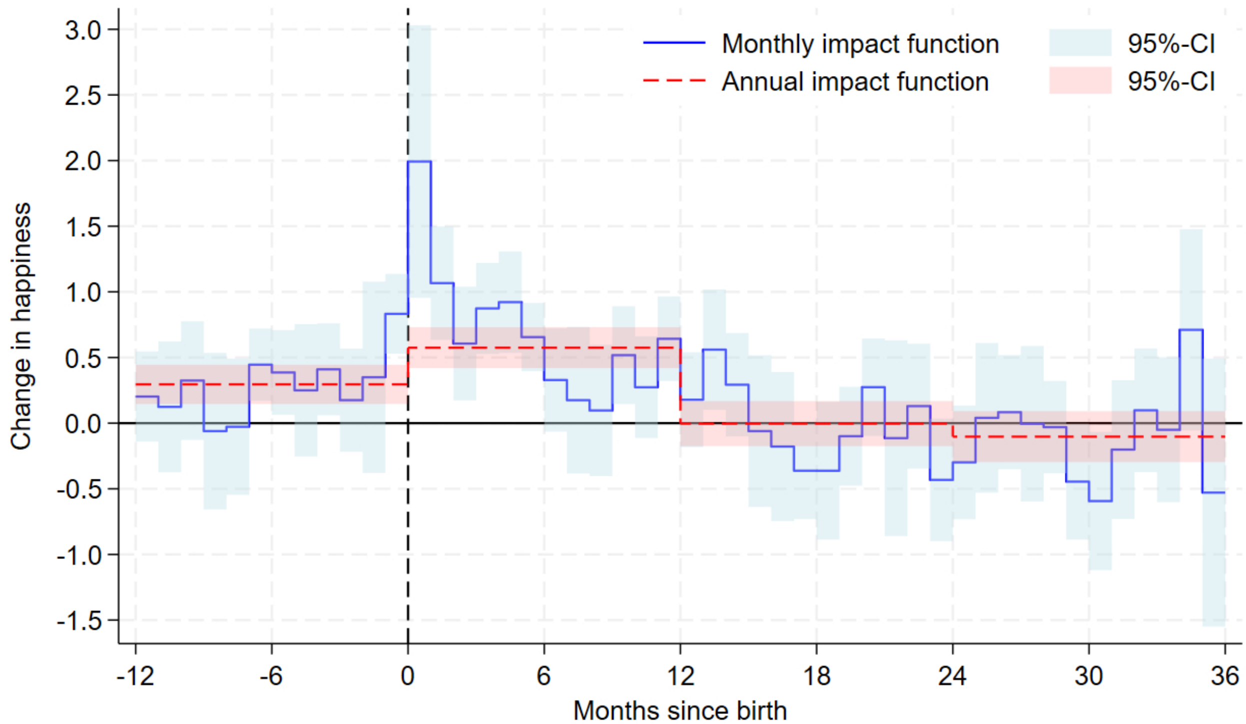

Figure 1: Effect of a first birth on mothers' happiness: comparing annual and monthly impact functions. Note: Results of two fixed-effect regressions. Coefficients of the event time dummies and their CIs are shown as step functions. Controls are age dummies, relationship status dummies, subjective health dummies, and a dummy for pregnancy (other than focal child). The dashed line indicates the birth. Source: pairfam release 14.2, own calculations.

would also affect annual estimates. Thus, this identification assumption is also needed when estimating annual impact functions. H&J discussed this identification problem under the label "selective interview timing" and suggest an exploratory test for this (see their online supplement).

Smoothing Monthly Impact Functions

Monthly estimates are often quite volatile (unless they are based on a very large data set). Therefore, H&J suggest smoothing the monthly impact function. They suggest using weighted local linear smoothing (lpoly in Stata; polynomial degree 1 and bandwidth 5) and presenting this smoothed impact function including its confidence band.[4] However, there is a drawback to this strategy: smoothing algorithms assume that the observations to be smoothed are data. In the case of a monthly impact function, however, the observations are estimates (from a fixed-effect regression). And these estimates are themselves uncertain. The smoothing algorithm does not take this additional uncertainty into account and may thus underestimate the uncertainty of the smoothed impact function. The CI of the smoothed impact function is therefore likely too narrow and this can lead to misleading conclusions.

To arrive at more appropriate CIs for the smoothed impact function, we recommend bootstrapping the procedure (a textbook description of bootstrapping can be found in Cameron and Trivedi [2022]: chap. 12): (1) generate multiple samples by resampling from the current sample (resample individuals with replacement, keeping all person-years for each selected individual to preserve the panel structure). (2) The monthly fixed-effect impact function is estimated for each bootstrap replication. Impact function estimates at each event time are saved. (3) The impact function estimates are smoothed for each replication. The smoothing algorithm provides an estimate of the smoothed impact function at each event time, which is also saved. (4) The standard deviation of these estimates across the bootstrap replications provides the bootstrap standard error of the smoothed impact function at each event time. (5) Using the bootstrap standard error, the normal-based 95 percent CI is computed.

We implement this bootstrap procedure with Stata's bootstrap command (9,999 replications; for more details, see the replication file). We find that in fact the CIs obtained by the bootstrapping procedure are, on average, 23 percent larger than those provided by lpoly (averaged over all 48 smoothed impact function estimates). The average width of the bootstrapped CIs is 0.426 and the average width of the lpoly CIs is 0.347. Therefore, basing one's interpretation on the lpoly CIs would overestimate the precision of the smoothed impact function. Therefore, we strongly recommend using the bootstrapped CIs.

Nevertheless, the bootstrapped CIs (average 0.426) are much narrower than the monthly estimates (average 0.961). The bootstrapped CIs are closer to the yearly estimates' CIs (average 0.335). This greater similarity to the yearly estimates is plausible given that the yearly estimates are based on all observations within 12 months, whereas the smoothed estimates are based on a weighted average of 10 months.

Figure 2 shows the results. The monthly estimates are plotted as dots, and their CIs are shown as spikes. The red line shows the smoothed impact function. The red area shows the confidence band obtained using the bootstrapping procedure described above. As can be seen, the smoothed impact function accurately captures the general pattern of the monthly impact function. An obvious advantage of the smoothed impact function is that it averages the somewhat volatile monthly estimates. One might argue that a disadvantage is its tendency to over smooth. For instance, the smoothed impact function indicates that anticipation effects become significant nine months before birth. However, the monthly impact function indicates that it is seven months. Thus, there is a trade-off between bias and variance (Fan and Gijbels 1996). Unsmoothed monthly estimates have no bias, but they carry a high risk of modeling random variation rather than true association. The more one smooths the data, the more precise the estimation becomes; however, the risk of estimating a biased curve also increases. Users must decide what is more important for their research question.

Another advantage of smoothed impact functions is that they allow for easier comparison across studies. For example, Figure 2 shows that the impact function pattern found in this article using pairfam data is similar to the pattern found by H&J using SOEP data. This similarity is more apparent when comparing smoothed impact functions. The detail of the monthly impact functions may obscure it.

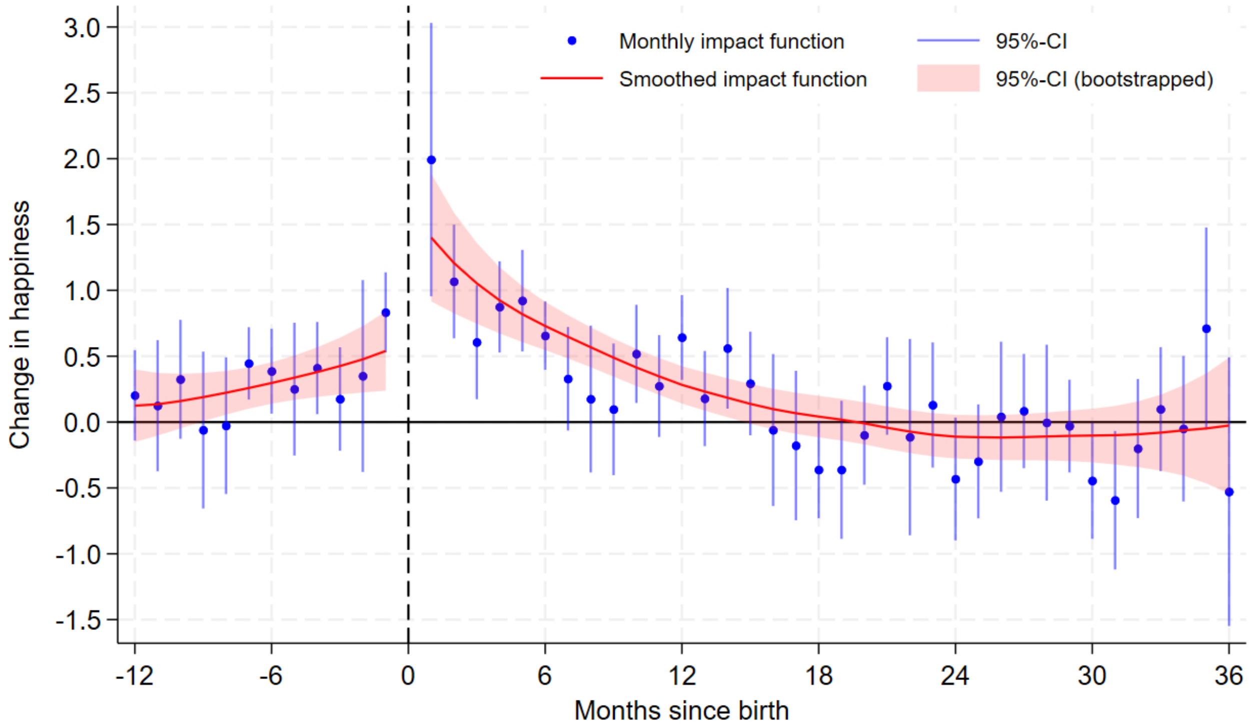

Figure 2: Effect of a first birth on mothers' happiness: comparing monthly and smoothed impact functions. Note: The monthly impact function shows the results of a fixed-effect regression (coefficients of the monthly event time dummies and their CIs). Controls are age dummies, relationship status dummies, subjective health dummies, and a dummy for pregnancy (other than focal child). The smoothed impact function applies weighted local linear smoothing (lpoly in Stata) to the monthly estimates (polynomial degree 1 and bandwidth 5). The confidence band is obtained by bootstrapping (bootstrap in Stata; 9,999 replications). The dashed line indicates the birth. Source: pairfam release 14.2, own calculations.

Conclusion

Based on our discussion, we have two recommendations. First, if your data include information on the month of the interview and the event, estimate and plot a monthly impact function. If you are not interested in the details of the monthly impact function, do this at least as a robustness check. Second, to better understand the general pattern of the monthly impact function, add a smoothed impact function with a bootstrapped confidence band to the plot. Decide whether you prioritize information reduction or detail.

Notes

1 When the research interest is not causal identification but the description of trajectories, researchers may want to factor out certain variables, such as changes due to general aging. Control variables are handled equally, whether in the yearly or (smoothed) monthly

approach (H&J did not discuss this topic; the application paper by Hudde (2024) includes control variables).

2 Due to "negative weighting bias" (see Ludwig and Brüderl 2021), the annual estimate will not be the arithmetic mean of the monthly estimates. Instead, it will be a weighted average, with later months weighted less.

3 In fact, this may be the case in our empirical application. There were only six interviews in the month of the first birth. In the next month, there were 36 interviews, and in the third month, there were 49 interviews.

4 An alternative approach is estimating a parametric impact function (see the replication file). However, the drawback of a parametric approach is that it often yields inaccurate predictions, particularly at the edges of the event time range (Ludwig and Brüderl 2021). Another alternative would be to impose splines (Miller 2023).

References

Brüderl, Josef, Sonja Drobnič, Karsten Hank, Franz J. Neyer, Sabine Walper, Christof Wolf, Philipp Alt, Irina Bauer, Simon Böhm, Elisabeth Borschel, Christiane Bozoyan, Pablo Christmann, Rüdiger Edinger, Felicitas Eigenbrodt, Madison Garrett, Svenja Geissler, Tita Gonzalez Avilés, Nicolai Gröpler, Tobias Gummer, Kristin Hajek, Michel Herzig, Renate Lorenz, Katharina Lutz, Timo Peter, Richard Preetz, Julia Reim, Barbara Sawatzki, Claudia Schmiedeberg, Philipp Schütze, Nina Schumann, Carolin Thönnissen, Katharina Timmermann, and Martin Wetzel. 2024. "The German Family Panel (pairfam)." GESIS Data Archive, Cologne. ZA5678 Data file Version 14.2.0. https://doi.org/10.4232/pairfam.5678.14.2.0

Cameron, A. Colin and Pravin K. Trivedi. 2022. *Microeconometrics Using Stata*. 2nd ed. College Station, TX: Stata Press.

Fan, Jianqing and Irène Gijbels. 1996. *Local Polynomial Modelling and Its Applications*. New York: Routledge. https://doi.org/10.1201/9780203748725

Hudde, Ansgar. 2024. "Do They Think That Joy and Misery Are Temporary? Comparing Trajectories of Current and Predicted Life Satisfaction across Life Events." *European Societies* 26(4):1121–36. https://doi.org/10.1080/14616696.2023.2289653

Hudde, Ansgar and Marita Jacob. 2023. "There's More in the Data! Using Month-Specific Information to Estimate Changes Before and After Major Life Events." *Sociological Science* 10:830–56. https://doi.org/10.15195/v10.a29

Huinink, Johannes, Josef Brüderl, Bernhard Nauck, Sabine Walper, Laura Castiglioni, and Michael Feldhaus. 2011. "Panel Analysis of Intimate Relationships and Family Dynamics (pairfam): Conceptual Framework and Design." *Journal of Family Research* 23(1):77–101. https://doi.org/10.20377/jfr-235

Ludwig, Volker and Josef Brüderl. 2021. "What You Need to Know When Estimating Impact Functions with Panel Data for Demographic Research." *Comparative Population Studies* 46:453–86. https://doi.org/10.12765/CPoS-2021-16

Miller, Douglas L. 2023. "An Introductory Guide to Event Study Models." *Journal of Economic Perspectives* 37(2):203–30. https://doi.org/10.1257/jep.37.2.203

Acknowledgments: We thank Katrin Auspurg for her helpful comments. This article uses data from the German Family Panel pairfam, coordinated by Josef Brüderl, Sonja Drobnič, Karsten Hank, Johannes Huinink, Bernhard Nauck, Franz J. Neyer, and Sabine Walper. From 2004 to 2022, pairfam was funded as a priority program and a long-term project by the German Research Foundation (DFG).

Josef Brüderl: Department of Sociology, LMU Munich. E-mail: bruederl@lmu.de
Ansgar Hudde: Department of Sociology and Social Psychology, University of Cologne. E-mail: hudde@wiso.uni-koeln.de
Marita Jacob: Department of Sociology and Social Psychology, University of Cologne. E-mail: marita.jacob@uni-koeln.de