Università Commerciale Luigi Bocconi
PhD School

PhD program in Economics and Finance
Cycle: 36
Disciplinary field: SECS-P/06

Essays in Labour Economics

Advisor: Simon Görlach
Co-advisor: Jérôme Adda

PhD Thesis by
Katarina Kuske
ID number: 3140898

Year 2026

Acknowledgements

First and foremost, I would like to thank Jérôme Adda and Simon Görlach for their thoughtful supervision and guidance throughout the PhD. They have been very generous with their time, and I walked away from every meeting with a fresh perspective.

I owe special thanks to Minji Bang and Hanna Wang, whose insights and encouragement were a great source of inspiration. Working with them has taught me a lot.

I am grateful to the many members of Bocconi's faculty who provided valuable feedback, and particularly to Alexia Delfino and Sarah Eichmeyer, who went above and beyond. I am also grateful for the support received by Angela Baldassarre, Patrizia Pellizzari, Marika Podelli, and Giulia Zenoni in dealing with all things administrative.

Some of the friends and colleagues who shared their support, advice, and company in celebrating the small wins throughout the years include Alberto, Amalia, Cristóbal, Diego, Enrico, Ines, Isabella, Jack, Julian, Maik, Maria-José, and Matteo. And I always appreciated the warm glow of spending time with Ricardito, Sofía, José Manuel, and Antonio. I was very lucky to find my Via Torricelli neighbours-turned-friends – Bini, Flavia, Lola, and Desi – whose dinners made the Covid time much more enjoyable.

A heartfelt thank you also to my parents, brother, and grandma for their emotional support and visits, which were cherished reasons to explore more of Milano and Italy. Finally, I am infinitely grateful to Iván for being a great partner both inside and outside of the office.

Thesis abstract

This thesis comprises three self-contained chapters. In these, I examine how non-wage factors and institutional frictions shape labour market outcomes and welfare at the individual and country level.

The first chapter investigates how co-parenting following divorce affects parents' labour market outcomes. Exploiting a Dutch reform that raised co-parenting uptake, I find that mothers' wages decline by 0.8% in an intention-to-treat framework post-reform, and that this effect rises to a substantial 10% wage loss for mothers who adopt co-parenting due to the reform (the complier population). This impact is driven by reduced geographical mobility, which constrains mothers from moving to access better-paid employment. Fathers' wages remain unchanged. The findings highlight an efficiency cost of co-parenting that falls disproportionately on mothers, thereby widening the gender wage gap.

The second chapter examines the role of non-wage factors in the labour market. Using a large-scale Dutch survey linked to administrative data, we find that workplace amenities – such as flexibility, workload, and social support –, like wages, strongly predict job satisfaction, search, absenteeism, and mobility. We document that desirable amenities are often bundled with high-paying jobs, which amplifies overall compensation inequality.

The final chapter turns to international capital markets and documents that foreign direct investment (FDI) inflows are procyclical while remittances are countercyclical with respect to recipient-country GDP growth. Estimating a model of both capital flows jointly, we evaluate how policies that lower barriers to these flows can enhance risk-sharing and improve welfare. The results show that reducing the cost of either enhances cross-border risk sharing, but the welfare gains from cheaper remittances dominate.

Taken together, the essays reveal how location constraints, worker sorting based on amenities, and cross-border capital flows influence the allocation of human and financial capital, with implications for gender inequality, job design, and macroeconomic risk sharing.

Contents

Chapter 1

Co-Parenting and Careers after Divorce

Abstract

I investigate how co-parenting affects parents' labour market outcomes after divorce, exploiting a custody reform in the Netherlands that encouraged co-parenting and increased its uptake by 7.6 percentage points among parents with young children. I find that mothers who divorce after the reform experience a 0.8% wage decline in an intention-to-treat framework relative to those divorcing before the reform, which implies an average wage loss of about 10% for compliers. This is driven by slower wage growth for mothers in the treatment group, who are less likely to move further away to access better-paid employment. The findings suggest that co-parenting ties both parents to a fixed location, reducing geographical mobility. Fathers' wages remain unaffected. The wage penalty is concentrated among mothers who were secondary earners during marriage and younger at the time of divorce. These patterns are consistent with couples placing greater weight on the primary earner's career when making location decisions, which makes the post-divorce location constraint under co-parenting bind more tightly for mothers than fathers, thereby widening the gender wage gap. My results indicate an efficiency cost of location constraints under co-parenting.

Keywords: Divorce, Parenthood, Family structure
JEL Codes: J12, J13

Results based on own calculations using non-public microdata from Statistics Netherlands (project 8999). The NFN data were collected by Utrecht University in collaboration with Statistics Netherlands (CBS) and were funded by grant 480-10-015 from the Medium Investments Fund of the Netherlands Organization for Scientific Research (NWO) and by Utrecht University.

1.1 Introduction

Parental divorce and separation represent a significant source of disruption in children's lives (Frimmel et al., 2024). Recognising this, many policies aim to safeguard children's well-being in the aftermath of parental separation. One key concern is that children tend to spend significantly less time with their fathers after divorce (Brown et al., 2025). The main reason for this is that, traditionally, physical custody was awarded to only one parent, usually the mother, even when legal custody was shared. More recently, joint physical custody (also termed co-parenting) – an arrangement in which children split their time equally between both parents after separation – has emerged as an increasingly prevalent alternative arrangement in many Western countries, reflecting judicial and legislative shifts toward equal parental responsibility. Co-parenting can indeed benefit children (Vowels et al., 2023; Nielsen, 2018; Spruijt and Duindam, 2009; Bauserman, 2002). However, its economic implications for parents themselves have been largely unexplored.

This paper asks how co-parenting affects parents' labour market trajectories. The effects are theoretically ambiguous: On the one hand, a more equal distribution of childcare responsibilities could free up time for mothers, who are typically the primary caregivers, thereby relaxing their time constraint and allowing them to increase their hours and earnings. On the other hand, the logistical demands of co-parenting, specifically the need for both parents to remain in close geographical proximity to their children, and thus also to their ex-partner, imposes a significant mobility constraint. This constraint could force individuals to remain in a local labour market that is not optimal for their career, potentially limiting their wage growth and career progression. For instance, a parent might be prevented from accepting a promotion that requires relocation, or be unable to pursue a better-paying job in another city or region, effectively shrinking their set of feasible job offers to a smaller area. The empirical question, therefore, is which of these opposing forces dominates.

To understand the causal effect of co-parenting on parents' labour market outcomes, I exploit a custody law reform in the Netherlands that provides an exogenous shock to the prevalence of its uptake. I find that for mothers, the reform leads to a persistent decline in hourly wages, while there is no increase in hours worked. The likely mechanism driving the negative wage effect is the mobility constraint, which prevents mothers from moving towards higher-paid employment opportunities when co-parenting. This finding reveals a previously unexamined economic friction created by non-market institutions and highlights a subtle but important tension between family and economic policy, where an institutional change designed to improve social welfare in one domain unintentionally generates a significant economic cost in another.

To illustrate the mechanism driving these results, I develop a theoretical framework that yields predictions on how a move from sole to joint physical custody affects mothers' and fathers' hourly wages, hours worked, and commuting time. The model treats post-divorce work decisions as a search problem with two different rules depending on the custody arrangement. If one parent keeps sole custody, each ex-partner is free to move anywhere in the country and simply takes the job that pays the highest wage. With joint physical custody, both parents must stay in the same location, which is assumed to be the location in which they lived prior to divorce, and can only accept jobs that are close enough to the location during marriage (or stay in their previous job). The model formalises the trade-off and constraints faced by a parent choosing a job: higher wages from a more distant location (up to a maximum distance under co-parenting) versus lower wages from a more accessible one. I assume that sole custody is only granted to mothers as in my setting, fathers receive sole custody in fewer than 5% of cases.

The model predicts that under co-parenting, mothers' wages will be lower than when they have sole physical custody, while fathers' wages will be unaffected. The intuition is simple: when married, couples typically choose their location based on the best opportunities for the main earner, frequently the father.[1] The family optimises its location to maximise household welfare, which can result in sub-optimal job matches for the secondary earner, who often foregoes better professional opportunities in other locations to accommodate their partner's career. When the marriage ends, a mother is more likely to find a better job match outside of the marital location, but the co-parenting arrangement prevents her from relocating. This creates an immediate friction: her job-search set is now a geographically constricted subset of her original, unconstrained set. My model also predicts that treated mothers will accept a longer commute and that the effect on their hours worked is theoretically ambiguous, as the positive effect of reduced childcare hours is offset by negative substitution effects of lower wages and longer commutes. The remaining part of the paper tests these predictions empirically.

The challenge in answering the question of how co-parenting affects divorced parents' labour market outcomes empirically lies in identifying an appropriate exogenous shift in co-parenting uptake. Purely correlational analyses risk misattributing outcomes based on selection into co-parenting as causal. For example, Augustijn (2023) shows that the positive effects of co-parenting on fathers' life satisfaction are driven by selection into co-parenting. Hence, it is crucial to rely on exogenous variation.

My empirical strategy exploits a nationwide custody law reform, passed in 2009, as a natural experiment. The setting is well-suited for a causal analysis as it provides exogenous variation in the probability of co-parenting uptake. The reform made a

[1] In a world with equal sharing of market work and home production within couples, the reform would be expected to affect mothers and fathers in equal proportion.

joint parenting plan mandatory and signalled a judicial preference for shared residence, thereby increasing the uptake of co-parenting among parents with young children. A key strength of this setting is that the reform was mild enough not to affect divorce rates or the characteristics of the divorcing population, alleviating concerns about selection bias.

My empirical analysis uses granular administrative data from Statistics Netherlands, which provides rich, longitudinal information on the employment histories of ~23,000 divorced individuals with young children. These data, which include precise firm location at the municipality level and 4-digit postcode, allow me to track not only labour market outcomes but also the geographical movements of parents across workplaces over time, which is essential for identifying the role of reduced mobility in shaping labour market outcomes under co-parenting arrangements.

I complement the administrative records with survey data covering a representative sample of parents who divorce just before or after the reform. This survey contains detailed information on post-divorce parenting arrangements, distance between ex-partners, commuting, and re-partnering outcomes. Using these survey data, I first show that the reform led to an increase in co-parenting rates of 7.6 percentage points among parents with young children, a 27% increase relative to the pre-reform share. In addition, I rely on the survey data to test if commuting and re-partnering outcomes change for parents who share physical custody of their children.

The main result of the paper is that the reform led to a persistent 0.8% decline in hourly wages among mothers with young children, in line with the theoretical predictions. Hours worked do not increase; they temporarily drop due to lower overtime before recovering to those in the control group. Fathers' labour market outcomes remain broadly unaffected. The local average treatment effect (LATE), computed by scaling the intention-to-treat (ITT) effects by the share of compliers, implies a 10% wage decrease for treated mothers with young children in co-parenting arrangements after the reform. The observed effects are not driven by business cycle fluctuations: Controlling for local unemployment rates at the time of divorce interacted with a time trend or only comparing mothers within the same month, region, and industry does not change the observed wage decline. A placebo check using divorcing mothers of older children, for whom co-parenting is an unlikely choice, shows no wage decrease.

The likely mechanism behind the negative wage effect is reduced geographical mobility. Co-parenting constrains the ability to move for higher-paying jobs as it introduces restrictions on moving the child without the other parent's consent. I show that the workplaces of treated parents remain in closer proximity to their ex-spouses' workplaces. Using data on employer locations, I quantify this friction, showing that treated parents

who take up co-parenting (compliers) work 13 kilometres closer to their ex-partners – compared to an average distance between parents of 27 kilometres. After the reform, ex-spouses are also more likely to be employed in the same region and even the same municipality, and additionally remain working in the same firm for longer after divorce if they worked together while married. This is consistent with co-parenting arrangements tying both parents to one location. The effects are concentrated among mothers who were secondary earners during their marriage and younger mothers – precisely those individuals who would have benefited most from moving for a better professional match.

The fact that the reform did not affect fathers' labour market outcomes can be rationalised through household behaviour during marriage. While married, both partners optimise as a couple, which can involve specialisation and choosing a residential location based on the career of the main earner, who is typically the father. This may result in worse job matches for mothers than fathers during marriage. Mothers may accept this trade-off, knowing that – at least before the reform – they are able to move to better professional opportunities together with their children should their marriage end in divorce.[2] Consequently, when the marriage ends, the father's career is less constrained by the marital location, while the mother, having been the secondary earner, is now stuck in a potentially sub-optimal location and faces a co-location friction.

I contribute to three strands of the literature: First, this paper forms part of the large literature on family formation or dissolution and labour supply (Calvo, 2022; Reynoso, 2019; Calvo et al., 2024; Goussé et al., 2017a,b; Wu and Pollard, 2000; Fernández and Wong, 2014). More specifically, I contribute to the insights on how post-separation legal accommodations affect labour supply decisions (Voena, 2015; Goussé and Leturcq, 2022; Halla, 2013, 2015; Nunley and Seals, 2011; Nguyen et al., 2018; Rasul, 2006; Weiss and Willis, 1985; Altindag et al., 2017; Vuri, 2018; Halla and Hölzl, 2007). While much of this literature sheds light on the effects of changes to divorce laws on still married couples, my contribution is to provide causal evidence on the labour market outcomes of divorcing parents. I show that legal institutions designed to promote child welfare can entail unintended economic costs.

Within this literature, there is scant evidence on how parents' labour market trajectories are affected by co-parenting, especially when it comes to causally identified estimates. The two most closely related papers are those by Halla (2013) and Bonnet et al. (2018). Bonnet et al. (2018) provide correlational evidence that mothers who choose co-parenting arrangements also work more hours. Thanks to the natural experiment provided by the Dutch reform, my results can instead be interpreted as causal estimates. Halla (2013) combines changes to laws regarding legal and physical custody at the

[2] If divorce risk is taken into consideration at all. Many individuals underestimate it for their own marriage (Berresheim and Koll, 2023).

level of U.S. states and investigates their effects on the fertility and divorce decisions of still married couples. In contrast, I cleanly isolate the effect of joint physical custody, which may well differ from that of legal custody, and document its effects on the labour market outcomes of divorcing couples. My setting allows me to provide causal evidence of the effect of co-parenting on wages and hours, and clearly illustrates the likely mechanism driving my results.

I also contribute to work on households' labour market and location decisions (Gemici, 2023; Braun et al., 2021), especially those of dual-earner couples (Guler et al., 2012; Foerster and Ulbricht, 2023; Alonzo, 2025; Ranosova, 2025; Rueda and Wilemme, 2025; Venator, 2024; Jayachandran et al., 2024; Foged, 2016; Costa and Kahn, 2000). These papers study how co-location frictions can involve trade-offs between partners' careers and how they limit the job search or migration decisions of one partner, often the woman, resulting in gender disparities in wages and employment. Divorced parents who share physical custody for their children are no longer part of the same household, but I demonstrate that they face similar location constraints, and that this version of co-location frictions also has important implications for gendered wage dynamics.

Finally, this paper also relates to the relevance of outside options for observing gendered effects of policies that are gender neutral on paper. Coviello et al. (2024) show that gender-neutral policies such as a minimum wage can result in unequal labour market outcomes for women and men when they face differences in outside options in the labour market.[3] I show that outside options also matter for creating gendered effects of policies in the domain of family law such as custody. In an on-the-job search and matching framework, my results can be interpreted as smaller offer acceptance sets for co-parents. The effect of co-parenting is a priori gender neutral, but results in gendered outcomes because mothers' current wages rank lower in their distribution of wage offers than the current wages of fathers relative to fathers' distribution of wage offers, i.e., mothers face a worse outside option, creating a differential incentive to move.

The remainder of the paper is organised as follows: Section 1.2 provides an overview of the Dutch institutional setting and details the 2009 reform. In Section 1.3, I develop a theoretical framework linking co-parenting to labour market outcomes. Section 1.4 describes data sources and sample construction. Section 1.5 presents the empirical strategy and discusses the identification assumptions. Section 1.6 reports main results and the evidence for the proposed mechanism. Section 1.7 discusses the broader implications of my results, and Section 1.8 concludes.

[3]Men and women may also react differently to the same outside option (Fredriksson et al., 2025).

1.2 Institutional Setting

1.2.1 2009 Reform

Evolution of Dutch divorce law From 1998, Dutch family law gave both parents joint legal custody after divorce. Thus, prior to 2009, Dutch family law had already favoured legal custody for both parents yet effectively granted sole physical residence only to one parent – typically the mother – unless both parties explicitly agreed to a shared arrangement. This de facto sole physical custody regime constrained fathers' access to children and limited judicial oversight over custodial schedules. In response to rising joint-residence advocacy by family sociologists and fathers' rights advocacy groups, the Dutch Parliament passed the Promotion of Continued Parenting and Careful Divorce Act (*Wet bevordering voortgezet ouderschap en zorgvuldige scheiding*, Stb. 2008, 500) in late November 2008 and it came into effect on 1 March 2009.

Importantly for my application, unlike in other countries that mandate a minimum period of separation before divorce, in the Netherlands there is no period of separation required before being able to file for divorce, neither before nor after the reform.

Reform The 2009 reform aimed to promote joint physical custody after divorce, for the sake of children's well-being. The reform amended the Dutch Civil Code, codifying the principle of "equal parenting" (*gelijkwaardig ouderschap*), meaning that children have the right to equal care from both parents, also once they are divorced.[4] It also states that from 1 March 2009 onward, parents filing for divorce (or the dissolution of a registered partnership) must include a signed parenting plan with their divorce petition. The reform reinforced that post-divorce, both parents remain jointly responsible for their children, not only financially but also in terms of parenting time. While the law stopped short of presuming a perfectly equal timeshare, the courts interpreted it as a strong encouragement of shared residence arrangements.[5] In other words, the expectation in 2009 became that, absent contraindications, children should spend a substantial amount of time with each parent. Sole physical custody now had to be justified, especially if the other parent requested a more equal arrangement. Moreover, the assumption of

[4]The law states (machine-translated) that "A child over whom the parents exercise joint parental authority, after dissolution of marriage (other than by death), dissolution of registered partnership (other than by death), or termination of cohabitation (where a § 252 annotation has been made), has the right to equal care and upbringing by both parents" (*Burgerlijk Wetboek Boek 1, artikel 247, lid 4*).

[5]One Family Court ruling shortly after the law came into effect explicitly noted that since 1 March 2009, Dutch children have the right to equal care from both parents. The ruling by the Dutch Supreme Court (Hoge Raad) on 21 May 2010 clarifies that post-reform, joint physical custody became the "prioriteitsmodel" (default preference), not an absolute rule. *Ruling ECLI:NL:HR:2010:BL7407.*

what is the child's best interest shifted the focus from stability of residence (maximised under sole physical custody) to continued equal parental engagement (maximised under co-parenting).

Throughout this paper, I only analyse individuals who experience divorces instead of also considering separations from cohabitation. The reason is that before 2023, unmarried fathers did not automatically receive legal custody of their children; the mother had to agree. Thus, claiming physical custody was more difficult for fathers without previous legal custody. To avoid the selection into joint legal custody of cohabiting couples, I only consider those couples who were married before separating.

Figure 1.1: Co-parenting share increased by 7.6 percentage points

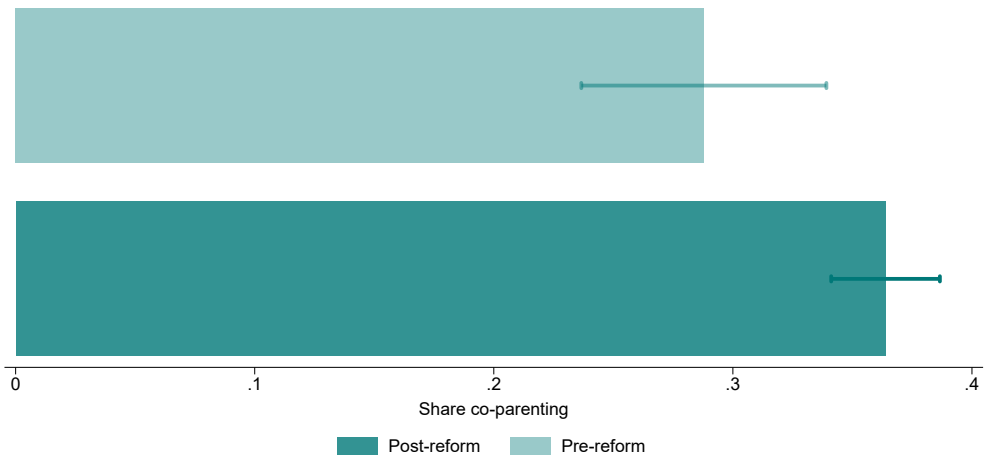

Note. Share of parents choosing co-parenting arrangements after divorce for parents separating between 2007 and 2012. The bars plot coefficients of a linear probability model in which co-parenting status is regressed on a reform indicator. Horizontal capped lines are 95% confidence intervals based on robust standard errors. Own calculations based on NFN data.

Parenting plan The reform also mandated that a parenting plan (*ouderschapsplan*) be submitted as part of the divorce petition of couples with minor children.[6] These parenting plans cover residence schedules (weekday/weekend split, holiday rotation), communication protocols (dispute resolution, information sharing), and financial support arrangements. The idea was to have parents plan in advance for how they would continue to share responsibilities, thereby reducing post-divorce conflicts. Crucially, in this plan, parents specify the living arrangement (either co-parenting or some form of visitation agreement). Importantly, the parenting plan also specifies the maximum mutual living distance (Dutch National Institute for Budget Information). This imposes

[6] A divorce petition involving children will not proceed without a parenting plan, unless the petitioner convinces the court that it was impossible to obtain one due to the other party's refusal *(Article 815 Rv)*.

restrictions on unilateral relocation: Once parents have agreed to joint physical custody, one parent cannot relocate with the child without the other parent's consent. The mandatory nature of the parenting plan also made the reform very salient (as shown in the Google searches for "parenting plan" in Figure C1.1) and effective: The share of divorcing parents who chose joint physical custody rose by 7.6 percentage points (Figure 1.1, Table A1.1).[7]

Distance between ex-spouses One of the requirements for making co-parenting work is that parents continue to live close to each other.[8] Indeed, survey data from the New Families in the Netherlands survey show that ex-spouses who co-parent live less than half as far apart as those who do not co-parent (Figure 1.2) and that this pattern persists over time (Figure C1.4). Of course, this pattern alone could either be due to parents who would choose to live closer to each other in any case deciding to co-parent, or be a consequence of co-parenting arrangements on the acceptable distance to one's former spouse. I provide evidence in support of the latter interpretation in Section 1.6.2.

Figure 1.2: Distance between divorced parents with minor children by co-parenting status

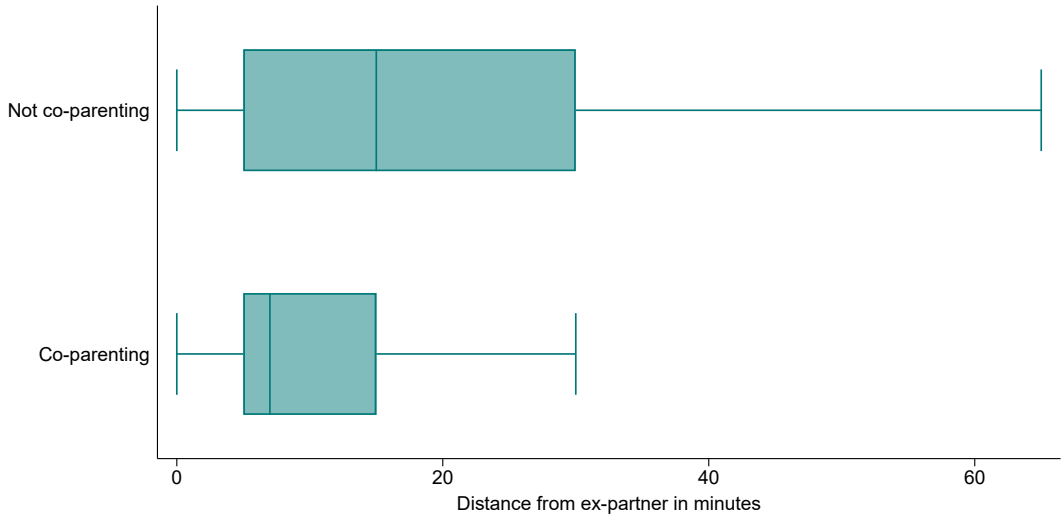

Note. This figure shows the travel distance between former spouses in minutes. The box indicates the 25 to 75 percentile, and the horizontal line inside the box denotes the median. Horizontal capped lines are 95% confidence intervals. Own calculations based on NFN data.

Childcare environment As I analyse the labour market behaviour of parents with young children, it is important to understand to what extent the childcare environment

[7] Poortman and van Gaalen (2017) also report a 7 percentage point increase in joint physical custody.
[8] They must also be on speaking terms and have little conflict.

may constrain or facilitate parents' employment. At the time of the custody reform, formal childcare in the Netherlands was cheap and relatively widely used. Thanks to a reform in 2005 that heavily subsidised formal childcare usage, take-up had risen. Parents' own contribution is dependent on their income and the number of children, but on average, they only paid for 19% of the childcare cost in 2008 (State Budget Archive). Thus, in the period under analysis, typical out-of-pocket costs to parents were very low by international standards.[9]

1.3 Theoretical framework

1.3.1 Intuition

After divorce, mothers work and earn more, compensating for some of the decrease in household earnings (Johnston et al., 2025). At the same time, adopting joint physical custody changes who spends how much time looking after the child and where the parents can live. The size of post-divorce labour market adjustment may thus differ by the prevailing custody regimes at the time of the reform. Figure 1.3 illustrates the key differences between the two custody regimes. There are three channels to consider when comparing the earnings of mothers choosing joint physical custody vs. sole physical custody:

(i) **Mobility constraint:** Joint physical custody arrangements put a stronger constraint on the distance between parents, restricting job search geography. In particular, those who were secondary earners during their marriage and who would have considered moving for higher wages under sole physical custody may now be locked into a smaller local labour market with lower best wage offers. At the same time, if mobility constraints bind, relatively lower wages could exert a negative substitution effect on hours.[10]

(ii) **Commuting costs:** Commuting across greater distances may be one adaptation to a tighter restriction on the location of residence in order to still be able to access employment opportunities further away. Doing so involves a trade-off between potentially higher hourly wages and time spent on paid work or leisure.

[9] Culturally, using formal childcare was also accepted in the Netherlands by the late 2000s, yet it is still common to have a grandparent day (*"opa- en omadag"*) of informally provided childcare. In terms of quality, survey responses for 2008 indicate that more than 90% of parents were satisfied with the childcare institution their child attended (SEO Economic Research).

[10] Sharing custody could also result in lower child-related expenses and a negative income effect on hours.

(iii) **Time constraint:** With childcare split more evenly under joint custody, the parent who was the main caregiver (typically the mother) could increase paid hours relative to the sole custody arrangement, as predicted by standard household models (Becker, 1985).

Figure 1.3: Stylised illustration of the two custody regimes

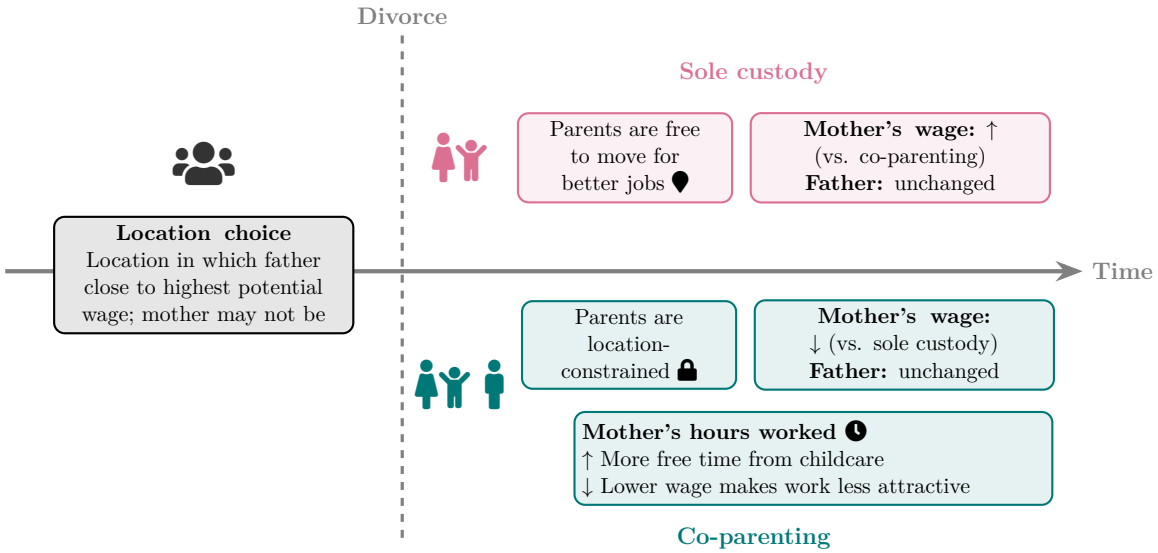

Note. Illustration of differences between sole and joint physical custody for mothers and fathers for the average Dutch family.

Which force dominates is an empirical question. If the mobility channel is paramount, we expect that implementing joint rather than sole physical custody arrangements would induce: (i) shorter distance moves by treated compared to non-treated parents; (ii) wage declines driven by parents who stay near their ex-spouse; (iii) negligible effects on wages and hours for individuals whose pre-divorce location already matched their job (mostly fathers); (iv) possibly longer commutes for co-parents, and (v) limited changes in hours worked.

1.3.2 Model setup

The aim of this theoretical framework is to derive predictions on the post-divorce labour market outcomes of moving from sole physical custody (SPC) by mothers to joint physical custody (JPC) by both parents. Section 1.6 tests these predictions empirically.

Labour market environment Ex-spouses $i \in \{m, f\}$ share a child, where m refers to mothers and f to fathers. There is a continuum (or discrete set) of local labour mar-

kets,[11] $x \in X$. Each labour market offers wage draws $w_i(x) \sim F_i(x_i)$ that differ across locations. A job is a pair of location and best wage offer in that location $(x_i, w(x_i))$. Individuals consider one job per location. If they receive multiple job offers in the same local labour market, they only consider the one paying the highest hourly wage.

Time constraint Parents spend their time on work, commuting, leisure, and childcare. There is an exogenous time amount of childcare parents must supply; the split between them varies by custody arrangement. Under regime $R \in \{\text{SPC}, \text{JPC}\}$, parent i supplies q_i^R. Children's total care need is fixed, such that $q_m^{\text{SPC}} + q_f^{\text{SPC}} = q_m^{\text{JPC}} + q_f^{\text{JPC}}$. I define σ_i^R as the share of total childcare that is provided by parent i under custody regime R: $\sigma_i^R = \frac{q_i^R}{q_m^R + q_f^R}$. Commuting time $k_i = d(x_i, x_{r,i})$ is a function of distance between the workplace x_i and the location of residence $x_{r,i}$. $d(x_i, x_{i,r})$ maps the distance between location of the job and location of residence into commuting time.[12] Individuals' time constraint is therefore given by

$$1 = l_i + h_i + k_i + q_i^R,$$

where l_i is leisure, h_i are market hours, k_i is commuting time, and q_i^R is child-care time.

Location constraint Under SPC each parent can pick any location of residence in the country; under JPC both parents must remain in the same location as during marriage $x_{r,i}$ and may only take jobs whose home-to-work distance is no more than a maximum distance $\bar{d}$. They face the following Euclidean distance constraint under JPC:

$$\|x_i - x_{r,i}\| \leq \bar{d}$$

This location constraint binds more strongly for mothers because of household optimisation during marriage: Mothers of young children are overwhelmingly employed part time, whereas the median father works full time. Hence, it is optimal for a representative couple to choose a location in which the father's wage is close to his best wage offer. The clearest example comes from movers: When couples move, they on average move to locations in which the husband sees an earnings increase, whereas his wife experiences no to negative earnings growth (Jayachandran et al., 2024). Once this location becomes the child's centre of life (school, friends, etc.) it becomes likely for the divorcing parents to remain close to that location. If parents cannot relocate unilaterally while maintaining JPC, fathers tend to be better matched to their jobs than mothers, in the sense that

[11] Assuming a continuum only simplifies derivatives; all results hold on a finite grid.

[12] Empirically, the mapping of distance to commuting time also depends on urbanity and the chosen means of transportation, but as this is not relevant for the model, it is omitted to simplify the mapping.

their wage is closer to the maximum wage in their job offer set after divorce. Mothers therefore have two job search sets: Under SPC, they search across all locations; under JPC, the only feasible locations are those within a circle of radius $\bar{d}$ around x_r.

Budget constraint Parents also face a period budget constraint:

$$c_i = w_i h_i - M_i^R,$$

where M_i^R is the amount that needs to be spent on the child. This is assumed to be exogenous and depends on the custody arrangement. For mothers, it is lower under JPC than SPC because the child is provided for by the father half of the time, i.e. $M_m^{\text{JPC}} < M_m^{\text{SPC}}$.[13] I abstract from spousal alimony as that is not affected by the choice of custody regime.

1.3.3 Optimisation problem

Parents maximise expected life-time utility (post-divorce):

$$\max_{\{h_i, x_i\}} U_i = \mathbf{E}\left[\sum_{t=t_{\text{div}}}^{T} \beta^t \left(\nu(w_i h_i - M_i^R) - \eta(h_i) + \phi(\sigma_i^R)\right)\right]$$

subject to the time and, under JPC, distance constraints. $\phi(\sigma_i^R)$ denotes the utility derived by parent i from spending share σ_i^R of total childcare time with the child.

Given the static nature of the set-up, the optimisation problem is equivalent to maximising flow utility subject to constraints:

$$u_i = \nu(c_i) - \eta(h_i) + \phi(\sigma_i^R),$$

where utility of consumption $\nu(c_i)$ and utility of child time $\phi(\sigma_i)$ are increasing and concave, and the disutility of working $\eta(h_i)$ is increasing and convex.

Taking the choice of custody arrangement as given, parents choose their optimal job, i.e. the utility-maximising location x_i^* with its wage offer $w_i^*(x_i^*)$ and their optimal hours h_i^*. Under JPC, the mobility constraint binds for mothers if and only if their optimal location x_m^* lies outside the radius $\bar{d}$ around x_r.

[13] I am implicitly assuming that child support payments under SPC are lower than the true cost of raising a child. Given that the level of child support payments is relatively low (about 250-300€ per month, depending on age of the child and income of the paying parent), this is empirically likely on average. The results do not depend on this assumption and hold equally when omitting the term M_i^R.

1.3.4 Effect of moving from sole to joint physical custody

What are the model predictions for hours, hourly wages, and commuting when switching from sole to joint physical custody? The model predicts wage decreases and commuting time increases for mothers who become location constrained. The theoretical prediction on mothers' hours is ambiguous. Fathers' hours are predicted to weakly decrease; there are no predicted changes to fathers wages or commuting time. Below, I state each of these propositions in turn. Detailed derivations are provided in Appendix C.

Hourly wages

Think of the wage at a location as a lottery ticket value $w(x)$. SPC lets a mother look at every ticket in the country and pick the highest-paying one. JPC removes any tickets lying beyond the distance $\bar{d}$. A simple "maximum over a subset" rule implies:

- If the best SPC ticket is outside the JPC circle (frequently the case for the lower-earning spouse), the maximum inside the circle must be strictly lower.
- If the best SPC ticket is already inside the circle (assumed for the main earner), the maximum is unchanged.

One can formally show that the wage weakly declines relative to SPC if the mobility constraint binds (refer to Appendix C for details).

Proposition 1. *The switch from SPC to JPC lowers mothers' expected hourly wages:*

$$w_m^{\text{JPC}} \leq w_m^{\text{SPC}}.$$

Proof. See Appendix D. □

Corollary 1 (No wage effect for fathers)**.** *If the father's SPC optimum already satisfies* $||x_f^{\text{SPC}} - x_r|| \leq \bar{d}$*, then the distance constraint is slack,* x_f^{SPC} *remains feasible after the reform and* $w_f^{\text{JPC}} = w_f^{\text{SPC}}$*.*

Commuting

The result on commuting follows trivially from the assumptions, but can also be shown formally.

Assumption A1 (Radial monotonicity of commuting time)**.** *For every pair of residence locations* $x_{r,1}, x_{r,2} \in X$*,*

$$\|x_{r,1} - x_i\| < \|x_{r,2} - x_i\| \implies d(x_{r,1}, x_i) < d(x_{r,2}, x_i).$$

Proposition 2. *Assume radial monotonicity of commuting time and recall that under JPC both parents reside in location x_r. Then for every realisation of wage draws*

$$d_m^{\text{JPC}} \geq d_m^{\text{SPC}} = 0, \qquad d_f^{\text{JPC}} = d_f^{\text{SPC}} = 0,$$

with $d_m^{\text{JPC}} > 0$ whenever the mother's JPC job lies at $x_m^{\text{JPC}} \neq x_r$ (strictly positive probability). Hence, under JPC, mothers commute for longer than under SPC. There is no change for fathers.

Proof. See Appendix D. □

Hours

The theoretical prediction for moving from SPC to JPC for mothers' working hours is ambiguous because it creates several counteracting effects for mothers. For fathers, the predicted effect is weakly negative.

Counteracting effects on hours First, fewer childcare hours translate into a relaxation of the time constraint: The switch to JPC reduces a mother's childcare obligations ($\Delta q_m = q_m^{\text{JPC}} - q_m^{\text{SPC}} < 0$). At the same time, mothers commute weakly more under JPC, reducing the hours available for work. But even if commuting increases, it is reasonable to assume that the reduction in childcare is typically larger, resulting in a net gain in disposable time. As a result, the marginal value of an hour of leisure (its shadow price λ_i) falls. This makes working relatively more attractive, creating an incentive to work more hours.

However, this positive effect on hours worked through a relaxation of the time constraint is counteracted by two forces:

(i) Mothers also experience a negative substitution effect due to lower wages. As established in Proposition 1, the mobility constraint under JPC forces mothers to accept a lower wage ($\Delta w_m = w_m^{\text{JPC}} - w_m^{\text{SPC}} \leq 0$). A lower wage makes working less rewarding compared to leisure. This substitution effect creates a clear incentive to work fewer hours. The strength of this effect depends on the elasticity of labour supply.

(ii) There is also an income effect regarding the cost of raising children when switching from SPC to JPC: For mothers, the monetary cost of children M_m decreases, which in turn lowers the marginal utility of income, and thus the incentive to work more hours. For fathers, the opposite is true.

In total, the effect of moving from SPC to JPC on hours worked is ambiguous for mothers. The ultimate change in a mother's working hours depends on which of these effects dominate. It is only positive if the relaxation of the childcare constraint outweighs (i) the additional commuting time *and* (ii) the negative substitution effect arising from lower wages *and* (iii) the income effect from lower child-related expenses.

For fathers, the predicted change in hours is also ambiguous. They do not experience a change in wages or commuting time, and thus there is no substitution effect for them. They should only experience a tightening of the time constraint due to increased childcare hours, leading to reduced market hours and/or leisure, and a potentially offsetting income effect through higher child-related expenditures that raises the incentive to work more hours.

Proposition 3. *The prediction on hours from moving from SPC to JPC is ambiguous but weakly positive for mothers when assuming that the time constraint effect dominates the substitution effect, and weakly negative for fathers (assuming the time constraint outweighs the positive income effect for fathers).*

Proof. See Appendix D. □

The theoretical predictions of the effect of joint physical custody on wages and hours worked are tested empirically in Section 1.6. I also provide suggestive evidence on commuting time.

1.4 Data

The main datasets used are administrative files from Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS). The labour market data consist of linked employer-employee records and provide information on the universe of all Dutch individuals' employment histories. They can be combined with data from the municipality registers that contain characteristics such as age, gender, and household composition. They can also be linked to the number and age of children as well as educational attainment.[14] On the employment side, I observe monthly hours and total earnings for all job spells. Based on these, I compute hourly wages. All employment histories can be linked at the firm level. Employer-side information includes sector, size, and, crucially for my proposed mechanism, the municipality and 4-digit postcode in which a firm is located. I use the municipality as a measure of firm location and compute the distance between

[14] Information on educational attainment is not available for all individuals.

firms based on these centroids.[15] I do not observe the municipality of residence. On the household side, I use the information on household composition to construct a panel of marriage spells and combine that with individuals' employment histories. Marriages include registered partnerships as these cannot be differentiated in the data. Legally, registered partnerships confer the same broad rights and responsibilities on partners as marriage, with only minimal differences between the two.[16] The reform also affects them equally.

I complement the administrative data with the New Families in the Netherlands survey (NFN), collected by Anne-Rigt Poortman (Utrecht University) and collaborators, which covers divorcing families before and after the reform. It contains detailed information on the parenting arrangement chosen, travel distance between ex-partners' households, commuting time to work, the time children spend with each parent, and demographic characteristics of parents and their children (Poortman et al., 2014; Poortman and Van Gaalen, 2019).

In the administrative data, I do not observe which custody regime is chosen. Hence, I use the NFN, which allows to observe the exact parenting regime chosen, to identify the share of compliers and for descriptive analyses with variables that are not available in the administrative records. My main analysis remains in an intention-to-treat framework.

1.4.1 Sample selection

In my sample, I include individuals who divorce between September 2007 and August 2010, i.e., in the 18 months before and after the reform. This window provides a large enough sample while still being relatively close together in terms of divorce date. Those divorcing between September 2007 and February 2009 form the control group, and those divorcing between March 2009 and August 2010 the treatment group. Extending it to 24 months before and after the reform does not change my results, see Section 1.6.1. For the main analysis, I restrict the sample to parents aged between 18 and 60 with the youngest child up to age ten at the time of divorce. There are some fathers for which the age ceiling is binding, i.e., they are older than 60 but have a child aged ten or younger. For mothers, it is not binding. I choose this age cut-off as to avoid capturing early retirees. I further restrict the sample to a balanced panel of those whose marriage

[15] There were 441 municipalities in the Netherlands on January 1st, 2009, so this measure provides detailed geographical information. The number of municipalities slightly shrinks over time as some are consolidated for administrative reasons. Using municipalities or 4-digit postcodes yields very similar results.

[16] In the administrative data, I observe the first of either 1) the actual date of divorce, 2) the date of separation when still-spouses cease to cohabit but have not yet finalised their divorce. Even when not yet legally divorced, the economically more relevant dimension is cohabitation status (also see Dasgupta et al. 2025).

lasted for at least 18 months, and who do not live with a new partner for at least 24 months after divorce. After applying these restrictions, the sample contains 40,000 divorces, 23,000 of which involve children under the age of ten. Conditioning on at least 18 months of marriage means that everyone married before the reform date, such that the marriage decision should not have been affected by custody considerations. The reason for excluding re-partnered individuals is that including them in the sample might alter labour supply incentives through income effects at the household level.[17] Re-partnering is an interesting and economically relevant outcome nonetheless, which I touch upon separately.

I choose the sample of parents with the youngest child aged ten and younger based on predictions of co-parenting using the NFN data. The group of compliers – those who did not adopt co-parenting before the reform but chose it after – is also highest among this group (Figure C1.2). Previous studies have also found joint physical custody to be most common when the children are of primary school age (Poortman and van Gaalen, 2017; Weston et al., 2011), specifically up to ten years old (Cancian and Meyer, 1998). Teenagers dislike the weekly moves that co-parenting entails. This is important because from the age of twelve, the child's preference regarding physical custody is strongly considered in the judge's decision.[18][19]

What are characteristics of parents who choose co-parenting? They tend to be well-educated (Figure C1.3) and enjoy a higher socio-economic status (Bakker and Mulder, 2013). Part of the reason for this are financial considerations: Co-parenting necessitates two bedrooms for the child and other duplicated costs, meaning that the total cost of supporting the child is generally higher than in a scenario in which the child lives only with one parent. Using NFN data, it can also be shown that parents who split market work more equally during marriage are more likely to opt for co-parenting arrangements after divorce.

1.4.2 Variable construction

Using the administrative data, I analyse how the reform affects two main sets of outcomes: labour market outcomes and geographical distance between former spouses. The labour market outcomes are standard: Employment is an indicator variable equal to 1 if an individual is employed in a given month and 0 otherwise; earnings are monthly and capture total compensation, including potential bonuses and allowances; hours worked

[17] The results are not sensitive to this and are unchanged when including re-partnered parents.

[18] Parents can also split their children between themselves. However, this option is generally not seen as in the best interest of the children, and therefore not commonly implemented by courts.

[19] The Australian Institute of Family Studies has found a very similar age range that makes co-parenting most likely (AIFS, 2011).

are monthly contractual hours plus overtime; and hourly wages are constructed as the ratio between earnings and hours. I exclude individuals employed for fewer than 10 hours per week. The data do not cover self-employed workers.

For geographical mobility, I use the following outcomes: Distance between former spouses, the probability of being employed in the same region or same municipality as the former spouse, and the probability of continuing to work in the same firm after divorce. The three measures are calculated based on the location of individuals' employers. For the distance measure, I compute the distance between the municipality centroid of the workplaces of both ex-spouses (in metres). The two measures of being employed in the same region or municipality as one's ex-spouse are indicator variables.[20]

The data contain firm location only at the statistical business unit (head office) level, not at the establishment (branch) level as the municipality or postcode information in the company register data refers to the seat of the administrative and legal unit. This means that firms with multiple establishments within one administrative unit cannot be attributed the correct municipality. Therefore, I limit the analysis of geographical mobility to individuals employed in firms with only one establishment.

1.5 Empirical Design

It is generally difficult to disentangle the causal effects of co-parenting from selection. To identify the causal impact on parents' labour market outcomes, I exploit the 2009 custody law reform as an exogenous shifter in the likelihood that parents choose co-parenting after divorce. The empirical strategy is a difference-in-differences (DiD) design: I compare parents of young children who divorced in the 18 months after the reform to those who divorced in the 18 months before the reform (see Figure 1.4 for a timeline).

The design is an intention-to-treat (ITT) design: The reform did not make co-parenting mandatory, but served more as a nudge. Nonetheless, it led to a 27% increase in its uptake. My empirical design uses the reform-induced variation as a source of identification. The DiD estimates provide reduced-form evidence on how co-parenting affects hours worked, wages, and earnings, as well as the distance between former spouses and the likelihood that parents remain in the same location.

[20] Throughout this paper, a region refers to the Dutch COROP regions, which are equivalent to the European NUTS3 level. There are 40 COROP regions in the Netherlands. At the time of the reform, there was no statistical unit for labour market regions. Such *arbeidsmarktregios* were introduced in 2012, and are slightly coarser than the COROP regions (there are 35 labour market regions). However, many *arbeidsmarktregios* were constructed to overlap with the COROP regions.

Figure 1.4: Timeline and definition of treatment and control groups.

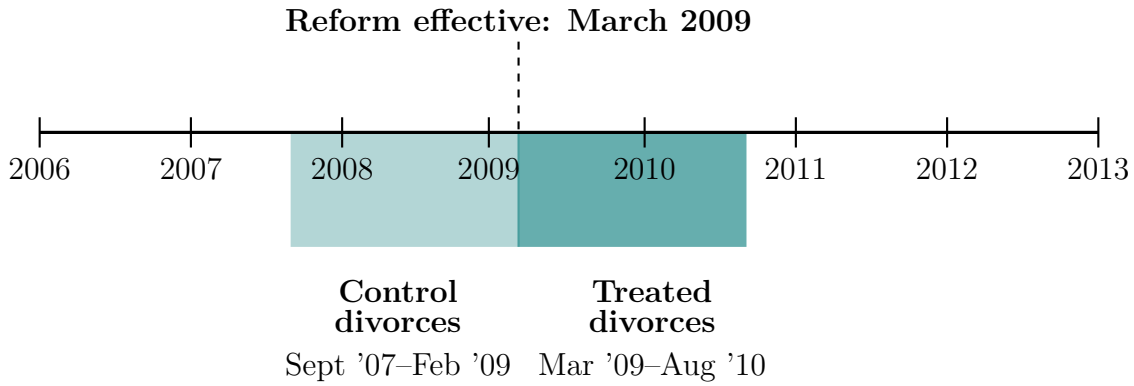

Note. The reform came into effect on 1 March 2009. The 18-month windows before and after define the control and treatment groups.

I estimate the following event-study specification:

$$Y_{itr} = \sum_{j \neq -1} \gamma_j \cdot \mathbf{1}[m_{it} = j] \cdot \mathbf{1}[\text{Treated}_i = 1] + \sum_{j \neq -1} \beta_j \cdot \mathbf{1}[m_{it} = j] + \rho_i + \delta_{tr} + \nu_{itr} \tag{1.1}$$

where Y_{itr} is the outcome of interest ((log) hours, wages, earnings, and later distance from ex-partner) for individual i in month-year t and region r. Treated_i is treatment status, equal to 1 if the individual divorced after the reform. γ_j is the coefficient of interest, comparing outcomes in the treated to those in the control group at the same time relative to divorce, m_{it}. To ensure that observed differences can be attributed to the reform and not to differential trends in local labour markets, I control for month-year×region fixed effects δ_{tr}. These capture macroeconomic and seasonal variation. These time fixed effects are crucial given the proximity of the reform to the Great Recession[21][22] and mean that the variation identifying the effects of the reform comes from comparing individuals divorcing in the same region, netting out region-specific time effects. I also include individual fixed effects ρ_i to account for time-invariant individual characteristics such as ability or preference. Standard errors are clustered at the individual level because treatment is assigned once per person and each individual appears in multiple periods. This allows the error terms to be arbitrarily serially correlated

[21] In the Netherlands, the Great Recession started in the fourth quarter of 2008 (CBS, 2009).

[22] As the primary focus lies on women's labour market outcomes, the Great Recession is somewhat less of a concern as it affected men's labour market outcomes more than women's. Moreover, if anything, women raise their labour supply during recession, cushioning the negative overall employment effects of the Great Recession (Doepke and Tertilt, 2016). The gendered differences were largely due to differential sorting into occupations and industries by women and men (Alon et al., 2022; Hoynes et al., 2012).

within each individual's labour market path (Bertrand et al., 2004). The identifying assumption is that – conditional on the set of fixed effects –, in the absence of the reform, labour market outcomes of parents in the treatment group would have followed the same trend as the labour market outcomes of those in the control group.

My preferred specification includes month-year×region fixed effects. The results are also robust to specifications with only month fixed effects and month-year×region×sector fixed effects. The latter strengthens the argument that the results are not due to the reform coinciding with the Great Recession as this specification compares individuals within the same sector and region, abstracting from macroeconomic conditions. This also rules out that potential effects are driven by changes in labour demand. All three specifications show very similar results. The similarity between the main specification and the month-year×region×sector fixed effects specification should assuage concerns about selection into industries, which may be endogenous. The sector is chosen at the two-digit level, a relatively broad definition to allow for sufficient variation within month-year×region×sector.

In addition to the difference-in-differences design on the sample of divorcing parents of young children, I conduct a placebo analysis comparing individuals divorcing before and after the reform with *older* children (aged 11 to 17). The parents of older vs. younger children should be more similar to each other than those without children, hence the two groups can be expected to be more comparable to each other on unobservables than other possible placebo groups. The difference lies in the fact that parents of teenagers are very unlikely to choose co-parenting arrangements. By demonstrating that the reform had no effect on parents of older children, this specification shows that the post-reform change can plausibly be attributed to the custody-law reform rather than to macroeconomic conditions.

1.5.1 Identification

No selection into treatment The treatment and control groups are balanced on observable pre-divorce demographic and labour market characteristics, including the number of children, hours worked, and secondary-earner status (see Tables 1.1 and 1.2).[23] The only meaningful differences are in earnings and wage levels, which rise over time and are thus higher among the treated group, as well as in age at divorce and the age of the youngest child at divorce, all of which also rises slightly because of rising ages at marriage. For those variables that are not balanced across the two samples, I show that the differences are due to general time trends (Figure C1.5). This is not a concern

[23]The samples are also very similar in terms of sectors of employment. Table A1.2 shows the top ten sectors of employment and their employment shares for mothers and fathers in my sample.

for my identification strategy as it does not require balance in levels, but in trends, for the parallel trends assumption to hold.

Table 1.1: Demographic characteristics of divorcees at the time of divorce

	Pre-Reform		Post-Reform		Post-Pre	*p*-value	
	Mean	Obs	Mean	Obs		unadj.	Holm
Age youngest child	5.40	15162	5.80	13117	0.40	0.000	0.000
Number of children	2.20	15162	2.20	13117	0.00	1.000	1.000
Age	37.20	15162	37.80	13117	0.60	0.000	0.000
Age ex-husband	40.00	15162	40.70	13117	0.70	0.000	0.000
High education	0.28	12103	0.29	10606	0.01	0.097	0.875
High education ex-husband	0.31	10423	0.33	9393	0.02	0.019	0.194

Note. Pre- and post-reform variable means and their difference. The last two columns report conventional (unadjusted) *p*-values and those adjusted for multiple hypothesis testing using the Holm correction. Variables are measured in the semester before divorce. High education refers to the share of individuals with higher vocational training or a university degree. Data on education levels is not available for all individuals.

Table 1.2: Labour market variables of divorcees around the 18 months before and after the 2009 custody reform

	Pre-Reform		Post-Reform		Post-Pre	*p*-value	
	Mean	Obs	Mean	Obs		unadj.	Holm
Employed	0.58	15162	0.60	13117	0.02	0.001	0.010
Earnings (monthly)	2010.00	8740	2132.00	7850	122.00	0.000	0.000
Hours (monthly)	103.00	8740	104.00	7850	1.00	0.072	0.860
Wage (hourly)	19.05	8740	19.90	7850	0.85	0.000	0.000
Part-time	0.87	8740	0.87	7850	0.00	1.000	1.000
Employed ex-husband	0.75	15162	0.75	13117	0.00	0.990	1.000
Earnings ex-husband	4274.00	11385	4403.00	9781	129.00	0.083	0.916
Hours ex-husband	164.00	11385	164.00	9781	0.00	0.978	1.000
Wage ex-husband	26.00	11385	26.80	9781	0.80	0.013	0.125
Part-time ex-husband	0.16	11387	0.16	9781	0.00	0.926	1.000
Secondary earner	0.72	13297	0.71	11600	-0.01	0.110	0.916
Equal earners	0.13	13297	0.13	11600	0.00	0.756	1.000
Main earner	0.15	13297	0.16	11600	0.01	0.068	0.860

Note. Pre- and post-reform variable means and their difference. The last two columns report conventional (unadjusted) *p*-values and those adjusted for multiple hypothesis testing using the Holm correction. Outcomes are measured one year before divorce. Earnings, hours, and wages are conditional on employment. Wages and earnings are denoted in 2015 prices.

In the NFN survey, the means of the relevant variables in the pre- and post-reform samples are similarly balanced (Table A1.3). One difference is that the pre-reform

sample is smaller than the post-reform one. To account for this and make the samples representative of the national population, sample weights are included. The NFN sample is also similar to the administrative sample in terms of demographic characteristics. One difference lies in how the labour market variables are measured: In the NFN, hours are weekly instead of monthly, but comparable when scaled up. While the NFN asks about paid work in general (including self-employment etc.), in the administrative records, I only observe employees.

No strategic timing of divorce A second condition for identification is that there is no break in the number of divorces around the date of the reform. One concern may be that the reform could have induced divorces that would otherwise not have taken place, prevented some individuals from initiating divorce or led them to anticipate or postpone it. Visual inspection of divorce rates of my sample does not support this (Figure 1.5) and neither does the Cattaneo-Janssen-Ma density test (Figure C1.6); there is no discontinuity in the number of divorces at the time of the reform or bunching right before. To further rule out an abrupt change in the number of divorces after the reform, I conduct a Chow test, allowing for heteroskedasticity. The Chow test does not reject the null hypothesis of no discontinuity at the time of the reform (Table A1.4). On top of providing supporting evidence for the identifying assumption, these results also show that the reform does not seem to have affected the number of divorces.

Figure 1.5: Monthly number of divorces

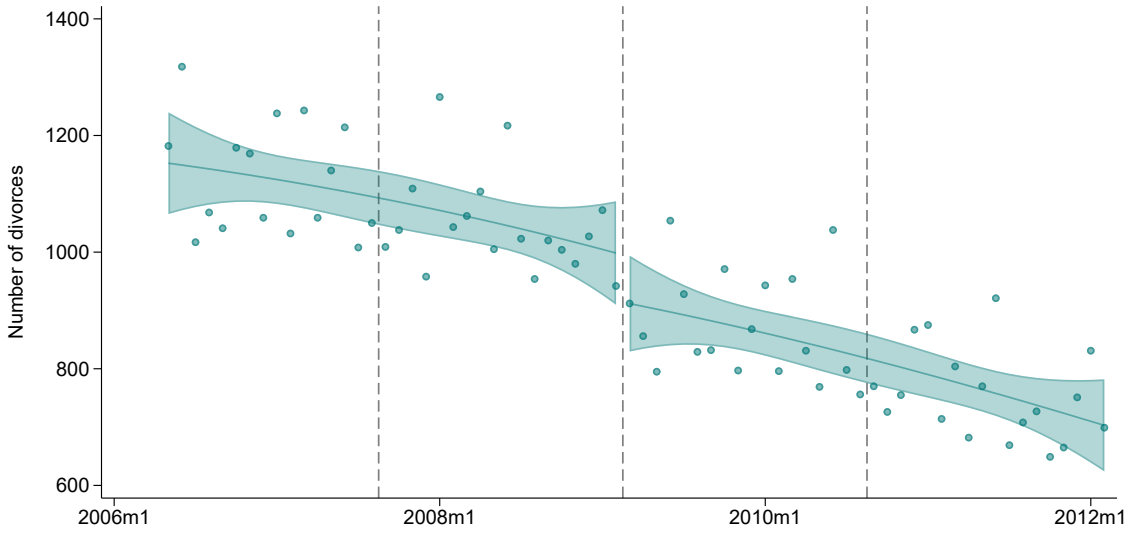

Note. Monthly number of divorces among parents with at least one child aged 10 or younger between September 2007 and August 2010. The central vertical line indicates March 2009, the time of the reform. The other two vertical lines indicate the pre- and post-reform horizons that define the control and treatment groups.

Parallel trends The assumption of parallel trends is supported by the lack of pre-trends, as shown in the event studies in Section 1.6.

1.6 Results

I start by showing that the reform did not meaningfully affect the extensive margin of labour supply. Mothers and fathers divorcing after the reform are no more or less likely to be employed than those who divorced before. Next, I test the predictions regarding wages and hours made in Propositions 1 and 3. I show that, as predicted, treated mothers experience a reduction in wages relative to those in the control group. The probable mechanism is reduced geographical mobility: Mothers divorcing after the reform are less likely to move to employers further away, but no less likely to change employers overall. They stay closer to their ex-partners, which is consistent with co-parenting arrangements requiring both parents to live close to each other. There is no clear increase in hours worked for mothers with young children after the reform.

1.6.1 Labour market outcomes

Employment Comparing mothers divorcing before and after the reform (specification in Equation 1.1), I find no significant differential effects of the reform on employment outcomes after divorce (Figure B1.1). Employment in the treated group is minimally lower, but not significantly so. In both the treatment and control group, the probability of being employed rises by 3 percentage points between one year before and one year after divorce (Figure B1.6). Some of this increase occurs in anticipation of the divorce. Overall, the change in employment status around divorce is relatively small – most mothers had already been employed before the divorce, reflecting high maternal labour force participation in the Netherlands – and, importantly, it does not differ significantly between treatment and control groups.

No pre-trends Figures 1.6 and 1.7 plot the estimated interaction coefficients $\hat{\gamma}_j$ of the main specification for monthly hours worked and hourly wages (both in logs), testing Propositions 1 and 3. The pre-divorce difference between those divorcing before and after the reform is very close to zero and insignificant, supporting the identifying assumption of parallel pre-trends. An F-test for joint significance of the pre-divorce coefficients is also insignificant (p-value = 0.32 for the wage result).

Hourly wages The key result is shown in Figure 1.6. Hourly wages decline for treated mothers, and these decreases are driven by wage gains in the control group, and a flat wage profile in the treatment group (Figure B1.4). Wages steadily decrease after divorce, from about 0.7% after the first year to 1.2% after two years. This result supports the theoretical prediction stated in Proposition 1. There is no effect on fathers' wages (Figure B1.10, Table A1.6). This, too, is in line with theoretical predictions (Corollary 1). The simple pre-post effects of the reform on labour market outcomes for mothers and fathers are summarised in Tables A1.5 and A1.6.

Figure 1.6: Hourly wages decrease persistently

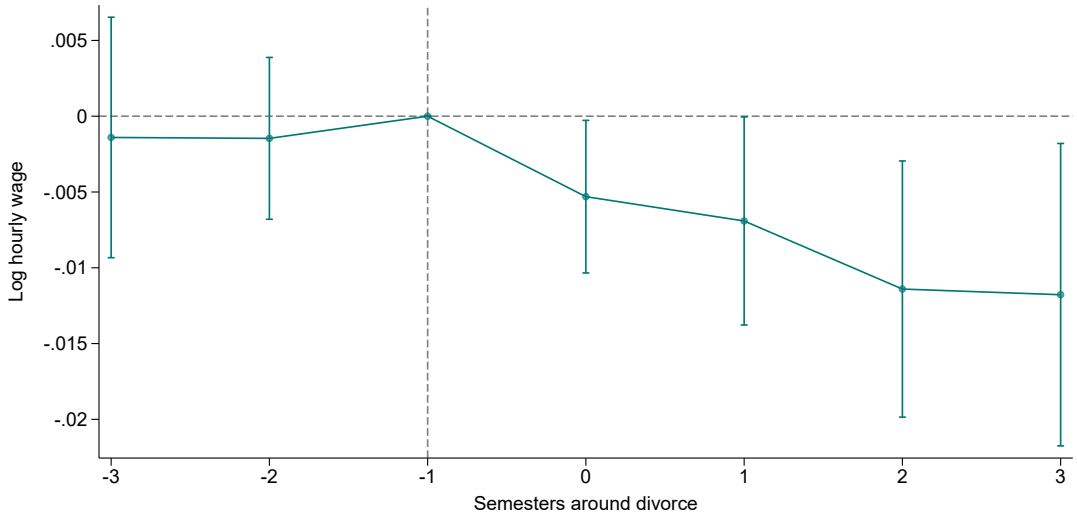

Note. This figure shows event study estimates of the effect of divorcing after the reform on mothers' hourly log wages. The sample includes mothers divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log wages between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

The observed wage decreases in the treatment group are not driven by business cycle factors: In the regressions, I use specifications controlling for month-year×region fixed effects as well as month-year×region×sector fixed effects, which capture possible regional variation in labour market opportunities over time and sectors. In Table A1.8, I additionally control for quarterly unemployment rates at the regional level, interacted with a linear time trend, and find an almost identical wage decrease as in the main specification.

The negative wage effect is indeed driven by those with young children. When regressing wages on the event time using the sample of mothers whose youngest child is older than 10 at the time of the divorce, the effect on wages is zero (Figure C1.7, Table A1.7). This

result affirms that the observed wage decrease is due to the increase in co-parenting, as parents of teenagers rarely co-parent.

Hours Mothers' hours worked did not increase significantly (Figure 1.7, Table A1.5).[24] This is in contrast to the idea that a relaxation of the time constraint due to a shift from largely sole physical custody by mothers to shared parenting outweighs the increased commuting time and negative substitution effect from lower wages. It is also in line with Proposition 3. In addition to negative substitution effects, income effects may also contribute: Expenses on children generally exceed what is covered by child support payments under sole physical custody. Such an income effect may further dampen the effect of a relaxation of the time constraint on hours. Finally, as laid out in Section 1.2.1, the availability and affordability of childcare do not appear to have been dramatic hindrances. Hence, the broadly null result on hours may be different in settings in which the lack of availability of quality childcare is a greater concern. The temporary dip in Figure 1.7 is due to a larger increase in hours in the control group, largely coming from overtime work.

Figure 1.7: There is no persistent change in mothers' hours worked

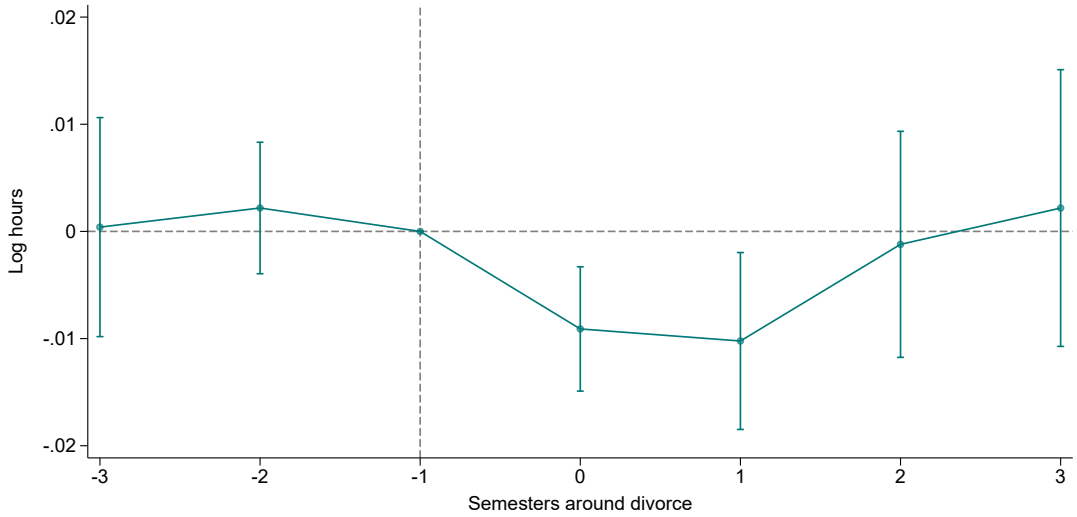

Note. This figure shows event study estimates of the effect of divorcing after the reform on mothers' monthly log hours. The sample includes mothers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log hours between treated and control individuals in a given semester relative to this reference period. The specification includes month-year × region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Finally, the NFN data provide suggestive evidence that those women who commit to

[24] When conditioning on the sample of continuously employed mothers, hours do increase by about 1%, with the increase concentrated in the last few semesters.

co-parenting consume more leisure and report less time pressure (Figure C1.10).[25] Unfortunately, the time use and leisure questions were not asked to individuals divorcing before the reform in the NFN. Therefore, the results should not be interpreted causally, but should instead be viewed as suggestive evidence. They could equally be explained by mothers with more leisure time selecting into co-parenting. Yet, they are consistent with being able to better concentrate on work during "off-duty" days and dedicate quality time to the child during "on-duty" days, whereas sole parents must juggle work and childcare contemporaneously.

Fathers' hours worked decrease slightly compared to before the reform, in line with the fact that those in co-parenting arrangements need to spend more time caring for their children (Figure B1.9, Table A1.6). Income effects from possibly spending more on their children may partly explain why the effects are small. Moreover, spousal alimony is not affected by the parenting arrangement, even if fathers wanted to work fewer hours (and thus earn less): Dutch law stipulates that the amount of alimony is not reduced if individuals decide to reduce their income voluntarily, e.g., by working fewer hours.

Earnings Treated mothers experience a significant earnings decline of about 1% for the first year after divorce, which diminishes in later semesters. Fathers' earnings are not affected by the reform (Figure B1.11). The observed earnings decline for mothers is driven by both slightly lower hours in the immediate semester after divorce, mostly by the persistent wage decreases. It is important to note that mothers in both the treatment and control group work more hours and have higher earnings after divorce. However, the increases are smaller for those treated by the reform. The earnings decline is thus driven by larger earnings increases in the control group, rather than actual decreases of those in the treatment group (Figure B1.7).

Heterogeneity Survey results have shown that families with more highly-educated mothers are more likely to choose joint physical custody. This suggests stronger effects when restricting the sample to more highly-educated mothers.[26] They, in particular if they were secondary earners during their marriage, may also be able to realise greater wage gains when moving to employers outside of their region. For this reason, I also test if the wage effect is more strongly negative when considering specifically more

[25] Figure C1.10 shows the coefficient on co-parenting interacted with a gender dummy for different outcomes related to leisure time based on the following regression:

$$y_i = \beta_0 + \beta_1 \text{ co-parenting}_i + \beta_2 \text{ woman}_i + \beta_3 \text{ co-parenting}_i \cdot \text{woman}_i + \varepsilon_i$$

[26] High education is defined as having any university degree or higher vocational training, which in the Netherlands also covers professions that require a university degree in other countries, e.g. architecture.

highly-educated mothers who were the secondary earners during their marriage.[27] The secondary-earner component is potentially very important: In couples in which women are the main earner during marriage, the couple has likely optimised the location choice for her job, meaning that the scope for earning a higher wage by moving elsewhere is smaller than if the couple chose the optimal location for the (ex-)husband's job.

Indeed, the wage effect is larger for mothers who were secondary earners during their marriage and have above-median earnings potential: These are the women who would have seen the largest wage gains from moving to a job that is a better match, thereby realising their earnings potential. Figure 1.8 shows the coefficient on log hourly wages for different subsamples. Compared to the full sample, those who were secondary earners during marriage and those with high education see larger wage decreases from the reform. These are plausibly the mothers who could have earned higher wages by moving.

If the effect of lower wages for the women divorcing after the 2009 reform is mediated by lower geographical mobility, we can expect the negative wage effects to be driven by those women who did not move to other regions of the country. Indeed, I find that the negative effect on wages is larger for those mothers who stay in the same region as their ex-partner (Figure C1.8). Of course, this is only suggestive as the specification conditions on the outcome of moving.

If secondary earners in general – both mothers and fathers – are those negatively affected by the reform, the effect will show up in the gendered pattern documented because of the much higher share of secondary earners among mother than fathers (72% vs. 15%, see Table 1.2). Fathers who are secondary earners also experience lower wages after the reform, but the sample is too small to detect statistically significant effects.

Finally, it seems that the inability to relocate freely affects mothers' wages most negatively during the years during which wage growth is generally steepest: The estimated effects on wages are twice as large for mothers below the median age at divorce, which is 38 in my sample, than those above (Figure 1.8).

Robustness To address the potential concern that I cannot discern the exact date of filing for divorce, I run the same specification using a doughnut-hole design, leaving out the two months immediately before and after the reform date (i.e., omitting those individuals who divorce between January and April 2009). This doughnut design yields results that are slightly larger in magnitude but still very similar to those in my main specification (Table A1.8). Separating the treatment and control groups further in time

[27] More precisely, towards the end of their marriage: My measure of earnings share is based on the 12 to 18 months before divorce (excluding the year right before divorce to avoid capturing possible anticipation).

Figure 1.8: The wage effect is driven by below median age, highly educated, and secondary-earner mothers

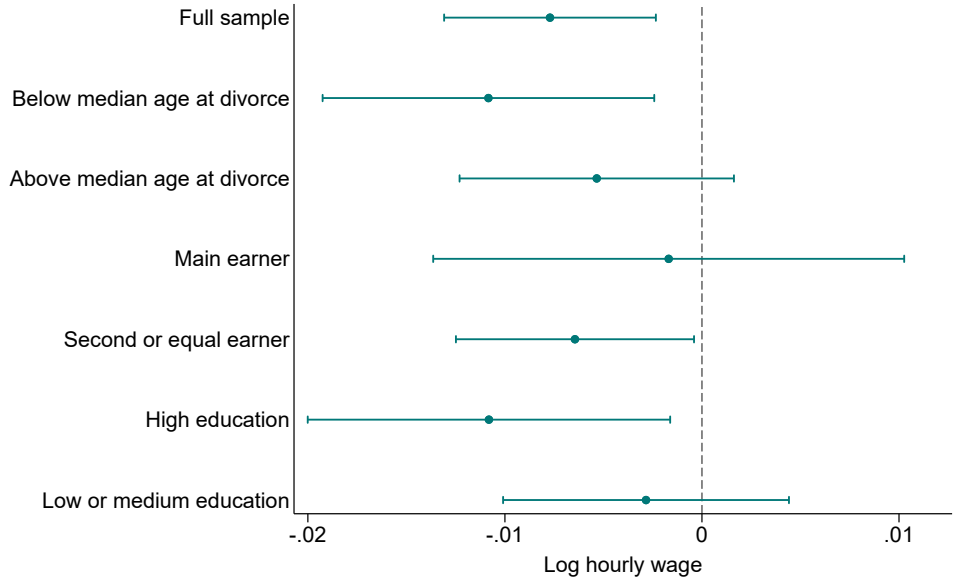

Note. This figure shows regression coefficients of a DiD regression of wages on the interaction between a post-reform and a treatment group indicator for different subgroups of mothers with young children. High education refers to university education or high vocational training. The median age at divorce of mothers in the sample is 38. Main earners are defined as earning at least 60% of household income. The sample includes mothers divorcing between September 2007 and August 2010. The specification includes month-year×region and individual fixed effects. Horizontal bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

also means that there is slightly less overlap in time to estimate the calendar month-year fixed effects. Therefore, the earliest and latest time fixed effects that could previously be identified based on observations of individuals from both the treatment and the control group are only identified from variation within either the control group (for the earliest) or the treatment group (for the latest).

I include all parents with young children in my main specification for hours and wages for the periods during which they are employed. This may raise concerns that the effects are driven by compositional changes. Yet, conditioning on the subsample of mothers who are continuously employed both before and after divorce does not change the results (Figure B1.12). The downside of this specification is that it excludes more than half of the number of observations.

The cut-off age of 10 years or younger was chosen to include likely compliers based on the NFN data. However, the results do not hinge on this specific age cut-off. Computing the main wage effect for each possible cut-off age (Figure C1.9) shows that the effects are highest for mothers of younger primary-school-age children and attenuate as the cut-off increasingly includes older children, but do not depend on the specific age cut-off chosen.

An alternative measure of the relative effects on wages by gender is to look at the wage difference between former spouses. This within-couple specification shows that the wages of mothers who were secondary earners are more negatively affected than those of fathers – mothers' wages relative to those of fathers decrease by 1.1-1.4 percentage points. The estimate is less precise than that of the main specification because of additional fixed effect for (ex-)spouse and the month×region of their job, but the results remain significant at the 10% level.

The month-year fixed effects should capture any labour market impacts that affect individuals observed in the same month. To additionally verify that the results are not driven by any remaining cross-sectional correlation among individuals observed in the same month – for example, shocks from nationwide macro news that are imperfectly absorbed by the month-year×region fixed effects – I also report a specification that clusters on both individual and month-year. This two-way clustering is conservative: It nests the individual-level correlation structure while additionally permitting arbitrary correlation across individuals that share the same month. The standard errors increase slightly given the relatively short panel, but remain highly statistically significant.

The fixed-effects specifications imply that the only potential threat to identification may arise from treated and control groups being differentially affected within month×region (×sector). To rule out that divorcing with children is worse for wages per se after March 2009, I look at mothers whose youngest child is between 11 and 17, which makes co-parenting very unlikely. Using these mothers as a placebo group, I find a very small and insignificant (positive) effect of the reform on these mothers' wages (Table A1.7). This result lends additional support to the claim that the observed effects are attributable to the reform.

1.6.2 Mechanism: Limited geographical mobility

Below, I provide evidence supporting the idea that the lack of post-divorce wage growth in the treated group is driven by limited geographical mobility. I show that the reform reduced the share of moves to municipalities that are different from the ex-spouses' by 1.5 percentage points, and that the average distance between ex-spouses shrinks by over one kilometre (13 kilometres for the compliers) two years post-divorce.[28] This reduced geographic mobility is the likely mechanism behind the relative wage decreases of mothers who were affected by the reform.

Using data on employer location, Figure 1.9 shows the results of the event study spe-

[28]The mean of individuals employed in the same municipality as their partner after divorce is 52.9% and the average post-divorce distance 27.5 kilometres.

Figure 1.9: Distance between ex-spouses' workplaces

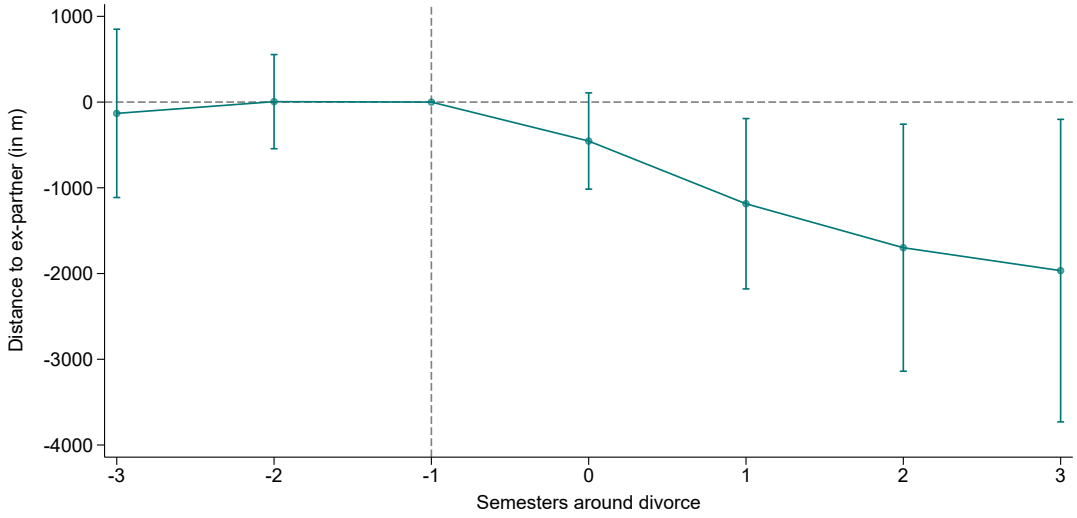

Note. This figure shows event study estimates of the effect of divorcing after the reform on the distance between workplaces of former spouses with young children. The sample includes individuals divorcing between September 2007 and August 2010 who are employed at single-establishment firms, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in the distance (in metres) between ex-spouses' workplaces for treated versus control individuals in a given semester, relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

cification with distance to ex-partner as an outcome. It shows that parents of young children divorcing after the reform stay working in greater proximity than before the reform (Figure 1.9). The effect is persistent and increases over time, consistent with gradual moves to better-paying jobs by mothers in the control group. It is also sizeable: Treated parents work one kilometre closer to each other than before the reform, out of a mean of 27.5km distance. In this sample, I exclude firms with multiple plants, for whom location information is based only on the head office and can thus be imprecise. Estimating the wage effect on the subsample of single-establishment firms yields the same estimate for the wage effect (Table A1.8, column 4).

Another measure of proximity between former partners is an indicator for employment in the same region or municipality. After the reform, ex-partners who are parents become more likely to have employers in the same municipality (Figure 1.10), in the same region (Table A1.9, column 3), and even in the same firm. These results are consistent with limited mobility towards better labour market opportunities as the mechanism behind a lack of wage growth for the treated mothers, and the finding that the negative wage effect is driven by those who do not move to employers outside of their previous region (Figure C1.8).

Figure 1.10: Probability of being employed in the same municipality as ex-spouse

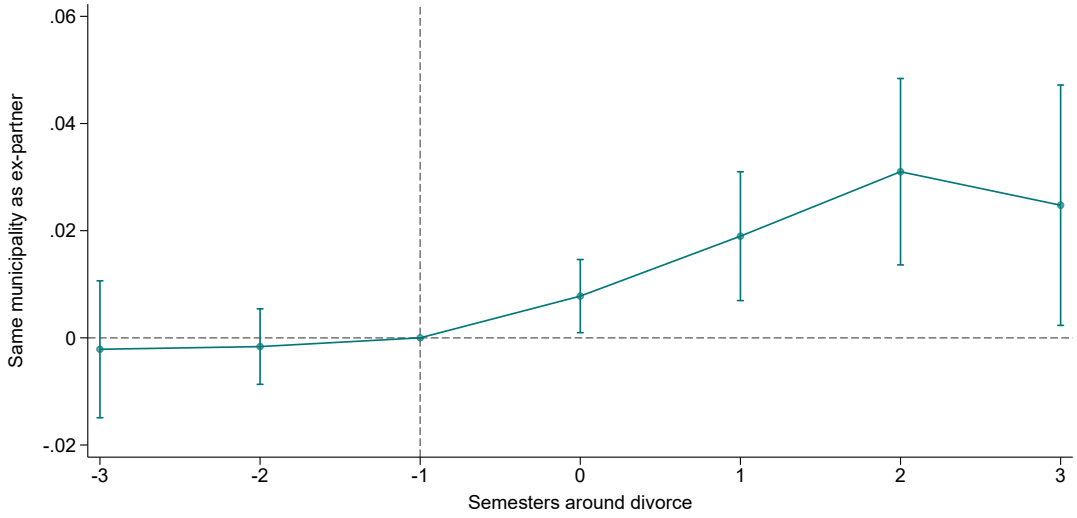

Note. This figure shows the probability of working in the same municipality as the former spouse. The sample includes individuals with young children divorcing between September 2007 and August 2010 who are employed at single-establishment firms. The bar heights represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform, relative to the semester before divorce. The specification includes month-year and individual fixed effects.

Do mothers who are further apart from their ex-partner after divorce see larger wage increases? While I cannot answer this causally, descriptive evidence points in this direction. Figure 1.11 shows mothers' log wages by 10-kilometre bins of post-divorce distance between ex-spouses workplaces. Mothers' wages are increasing in distance from their ex-partner up until 30-40km distance (there are relatively few mothers at distances above 50km from their ex-partner). The wage difference when moving from the <10km bin to the 30-40km one is about 9%. The relationship between distance and log wage does not change before or after the reform, so the observed wage decrease for treated mothers likely arises through fewer distance moves after the reform, as shown in Figure B1.5.

Given that the reform coincided with the Great Recession, one concern might be that decreased mobility could be driven by changes to house prices. To address this concern, I rely on the region×month-year fixed effects. In addition, I also use a specification controlling for yearly municipal house sale prices interacted with a linear trend, and find no change to my main results.

1.6.3 Ruling out alternative explanations

Income effects through alimony The reform increases fathers' bargaining power, which may be a concern for my identification strategy if they use this to negotiate lower

Figure 1.11: Mothers' log wages by distance from ex-partner

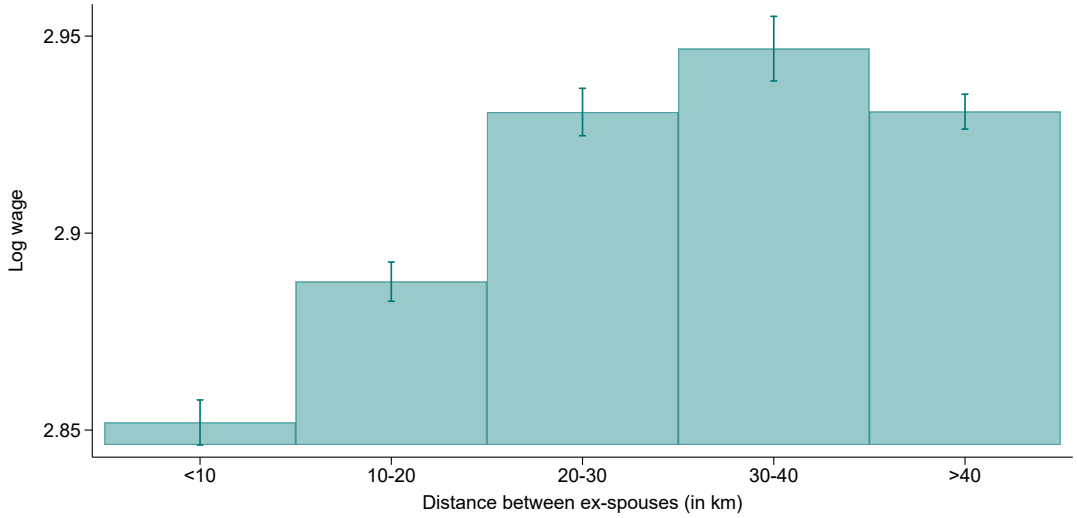

Note. This figure shows mothers' post-divorce log wages by distance from ex-partners' workplaces in bins of 10km.

spousal alimony payments, thereby differentially affecting mothers' incentives in the treatment and control groups. I do not observe alimony payments in the administrative data,[29] but can observe whether alimony payments are made and to whom in the NFN sample. Using that, I document that co-parenting status is uncorrelated with the probability of paying or receiving spousal alimony (Table A1.10), which does not suggest that fathers' newfound bargaining power translates into changes to alimony payments. Moreover, the overall share of couples in which alimony payments are made is relatively low, at less than 20%. Most important to ruling out differential income effects due to changes in alimony payments is the expected direction of income effects: A reduction in alimony would not be able to explain the pattern that wages and earnings *decrease* relative to the pre-reform divorcees. Instead, if alimony was reduced, income effects suggest that mothers should increase their earnings.

Lower productivity If co-parenting lowers mothers' productivity, this may be reflected in lower hourly wages. It seems highly unlikely that the time needed for coordination with the other parent more than outweighs the time that would have been spent on weekly childcare under sole physical custody. Co-parenting is not associated with increased time pressure that could make mothers less productive and thereby result in lower wages: On the contrary, leisure and time pressure results from NFN show *less* time pressure for co-parenting mothers (Figure C1.10).

[29] CBS data contain information on alimony payments, but only from 2011.

Search Effort/Intensity An alternative channel is a reduction in search intensity because of an increased cognitive load of co-parenting. If co-parenting mothers search less intensively instead of having a geographically restricted acceptance set, one may expect similar negative effects on wages. While mothers search less intensively in general (Bang and Wang, 2025), shared childcare should lighten the childcare burden and free up additional time for job search among parents with joint physical custody. However, lower search effort is unlikely to explain my results, as there is no differential effect on job separation rates comparing pre-/post-reform – the only difference lies in where they move. Finally, one would expect symmetric effects for mothers and fathers, which I do not find, as demonstrated in Figures 1.6 and B1.9.[30]

1.6.4 LATE: Scaling wage effect for compliers

The ITT estimate of a 0.8% wage decrease is relatively small. However, the share of compliers among parents with children aged up to ten is estimated to be 7.6 percentage points, so the likely effect for compliers is an order of magnitude larger. Therefore, I also provide estimates for the local average treatment effect for the compliers (LATE) by scaling the estimated ITT by the share of compliers, both for my main sample of parents with children ten and younger and the full sample of divorcing parents with minor children.

Using the NFN dataset that details post-divorce parenting arrangements, I estimate the share of compliers, i.e., the share of parents who chose co-parenting after the reform but would not have chosen it before. I regress the co-parenting status on treatment status. The results show a 6.7 percentage point increase among all parents of minor children, and a 7.6 percentage point increase among those parents whose youngest child is aged ten or younger (Table A1.1).

I estimate the LATE by scaling the estimated ITT by the estimated share of compliers:

$$\widehat{\text{LATE}} = \frac{\widehat{\text{ITT}}}{\widehat{\text{share of compliers}}} = \frac{0.00771}{0.076} \approx 0.101 \tag{1.2}$$

The estimated LATE is large: It suggests a 10% reduction in hourly wages for mothers who take up co-parenting due to the reform (Equation 1.2, Table 1.3). If considering not only those parents with children 10 and younger, but those with any minor children, the LATE decreases but still implies an economically meaningful 6.4% wage decrease. Given that co-parenting is much less likely for teenagers, the effect on parents with any

[30] Similarly, if having to frequently interact with the ex-partner causes psychological stress that negatively affects productivity, the expected effects would be the same for mothers and fathers.

Table 1.3: Estimated local average treatment effects of the reform for compliers

	Wage change	CI	
	relative to control	Delta method	Fieller's theorem
Youngest child under 10	-10.1%	[-0.1984, -0.0035]	[-0.3398, -0.0285]
Any minor children	-6.4%	[-0.1384, 0.0095]	[-0.2153, -0.0030]

Note. LATE estimates, combining estimation results of the effect of the reform on wages in Tables A1.5 with the share of compliers in Table A1.1.

minor children is still driven by those with younger children, among whom the compliers are concentrated.

LATE assumptions One potential issue with the LATE estimates is that there could be compositional changes. The presence of defiers – those parents who would have chosen co-parenting before the reform but decide against it after –, and thereby a violation of the monotonicity assumption, is unlikely. The balance checks described in Section 1.3.1 also do not suggest changes in the composition of the sample. Nonetheless, it is important to keep in mind that the compliers are specific types of parents: those with higher education who already shared market hours more equally than the average couples before divorce.[31] The effect may be different if co-parenting was assigned randomly, and is likely to be lower for individuals for whom wages depend less on good matches between their skills and the ones required for their employment.

1.6.5 Commute

One margin of labour market adaptation to not being able to relocate further away due to co-parenting could be to commute for longer. In the Netherlands, good public transport infrastructure makes it relatively convenient to commute over larger distances. However, even moderately long commutes may not be feasible for parents of young children, given the constraint to only leave for work after the start of school. To better understand if there is some adjustment on the commuting margin for co-parents, I use the NFN data. I compare the commuting duration of mothers and fathers by co-parenting status (Table 1.4). The regression displays the familiar result that women commute less than men (Le Barbanchon et al., 2021). Men who co-parent have a slightly shorter commute than those who do not. The interaction between co-parenting and women is large relative to the mean dependent variable and significant: Co-parenting mothers commute more. These results are consistent with mothers trying to compensate

[31] Bakker and Mulder (2013) also find that the couples who select into co-parenting are positively selected on earnings.

being location-bound with longer commutes, as predicted by Proposition 2. Nonetheless, based only on these data, I cannot rule out the alternative interpretation that only women with greater willingness to commute enter into co-parenting arrangements.

Table 1.4: Commuting duration by co-parenting status and gender

	Commuting time (in min)
Co-parenting	-3.816
	(1.277)
Woman	-11.23
	(1.126)
Co-parenting × Woman	3.982
	(1.745)
Mean dep. var.	24.68
Observations	2070

Note. The table shows coefficients for the regression of commuting duration (in minutes) on co-parenting status, controlling for five categories of educational attainment. The sample is that of parents who have been separated for up to three years. Data: NFN.

1.6.6 Re-partnering outcomes

A change in the presence of children in the household could also affect parents' re-partnering probabilities, whether in new cohabiting or married unions. Below, I investigate this descriptively using those parents who selected into co-parenting. As the time since separation matters strongly for the probability of living with a new partner, and respondents in the NFN were sampled at different times relative to their divorce, I control for the years since separation in the regression.

There are no large negative effects of co-parenting on the probability of re-partnering for mothers. However, fathers who co-parent are much less likely to move in with a new partner after divorce: 9 percentage points, which is more than a third lower than the re-partnering probability of fathers who are not co-parents (Table 1.5). Bakker and Mulder (2013) reach the same conclusion drawing on 35 in-depth qualitative interviews. On average, mothers are 6.6 percentage points less likely to live with a new partner than fathers. However, when co-parenting, both parents are equally (un)likely to live with a new partner after divorce. While co-parenting mothers appear disadvantaged relative to fathers in the labour market, when it comes to the dating market for cohabiting relationships, both are equally affected by the presence of children.

Table 1.5: Probability of living with a new partner by co-parenting status

	Women (1)	Men (2)
Co-parenting	-0.0153	-0.0926
	(0.0230)	(0.0233)
Years since separation	0.0320	0.0326
	(0.0076)	(0.0077)
Mean dep. var	0.275	0.341
Observations	3374	3369

Note. The table shows coefficients from regressions of the re-partnering probability on co-parenting status and parent gender, controlling for years since separation. Data: NFN.

1.7 Discussion

Should co-parenting become even more widely adopted as a post-divorce parenting arrangement? As an expansion would cover parents who currently do not choose co-parenting, and are thus different from the compliers for whom effects are estimated, this question does not have a straightforward answer. On the side of benefits to children's well-being, the quality of parental time matters for child development (Chan and Liu, 2025). Part of the positive developmental outcomes for children in co-parenting arrangements likely stems from the fact that the parents – including fathers – who choose it are positively selected. Therefore, extending co-parenting arrangements to the full population of divorced parents is likely to entail lower benefits for children. At the same time, the parents selecting into co-parenting are those that have the highest expected gains (in child welfare, time, etc.) net of costs. Thus, mandatory co-parenting may result in worse labour market outcomes for those parents who currently opt out of this arrangement, while yielding lower benefits to children than the ones previously estimated.

There are already small financial incentives for parents to adopt co-parenting. They include extending child tax credits and housing benefits to both co-parents.[32] Yet the

[32]In terms of child-related government benefits and allowances, there are slight advantages to co-parenting, especially for higher earners. The child benefit is paid by the Social Insurance Bank (SVB) to one of the parents, but they may then divide the money among themselves. The child-related budget (KGB) is paid by the Tax Authorities and is paid to the parent with whom the child is registered. It is income-dependent, so it is generally considered advantageous to register the child with the lower-earning parent (Uiteen Family Law). If there are two children or more, they can be registered with different parents, so that both receive some of the child-related budget. Moreover, there is a child tax credit for parents of minor children who are in work. This tax credit is generally granted to the parent with whom the child is registered, but in the case of co-parenting, both parents can apply. Finally,

fact that joint physical custody comes at a labour market cost to women may limit its uptake. At the same time, it also requires a certain income to be able to afford a bedroom for children in each parent's home etc., suggesting that negative labour market outcomes due to co-parenting will not hit the financially most constrained families. As governments worldwide promote co-parenting arrangements, alternative policy considerations to alleviate geographical labour market constraints could lie in facilitating options to work remotely, but it remains to be determined whether remote working may have adverse career effects itself.

Effectively, the child custody reform trades off child welfare considerations with mothers' career opportunities. Policy could consider spousal compensation for lower earnings in the joint location. Such transfers could function akin to alimony, and are similar to those envisioned by Reynoso (2019) to compensate the partner who specialised in home production for not being able to build human capital that is valued in the labour market. From an aggregate-economy perspective, joint physical custody limits the efficient spatial allocation of workers across the economy.

1.8 Conclusion

The labour market consequences of co-parenting are ex-ante ambiguous. It may alleviate time constraints for the primary caregiver, potentially expanding their capacity for market work. However, co-parenting also introduces a new constraint: Parents may face reduced geographical mobility, limiting career advancement opportunities that require relocation or long commutes. In this paper, I analyse the labour market outcomes of parents following the 2009 custody reform in the Netherlands, which encouraged and raised the share of divorced couples with young children who choose co-parenting.

Comparing parents who divorce before and after the reform in a difference-in-differences framework shows broadly no changes to fathers' labour market outcomes. Mothers, however, experience a decline in hourly wages of 0.8% compared to those who divorced before the reform. When scaled by the share of compliers, the magnitude of this wage effect increases to an economically meaningful -10% for mothers with young children. The wage decrease persists for the entire window of observation and grows over time. It is not driven by business cycle factors such as local unemployment conditions. I do not find strong evidence of a relaxation of time constraints for mothers that would translate into expanded market hours. These empirical findings are consistent with theoretical predictions.

rent allowance can also be claimed by both parents if they have a signed co-parenting agreement.

The negative effect on wages is likely rooted in lower geographical mobility. Co-parenting requires both parents to live in close geographical proximity for the arrangement to function. Indeed, after the reform, ex-spouses' workplaces are closer to each other and more likely to be in the same region and same municipality. This channel also helps explain why I find negative wage effects for mothers, but not fathers: While still married, the couple will have optimised at the household level, which for many entails choosing a location based on professional opportunities for the main earner. Given that this is the husband in a majority of couples, he will on average enjoy better professional opportunities in the chosen location than his wife. Once the marriage breaks down and ex-spouses optimise individually, mothers are more likely to face better employment opportunities elsewhere and have a financial incentive to move. Yet when co-parenting, the location of marriage becomes sticky, and the mothers who, with sole physical custody of the children, might have moved for higher-wage jobs are now tied stayers in a possibly sub-optimal location in terms of employment opportunities for themselves. The mobility constraint illustrates an important trade-off between family arrangements and spatial labour market frictions.

This paper estimates the partial equilibrium effects on parents who divorce right after the reform, so for whom the couple-location decision was made before the change to custody law. If individuals have rational expectations about divorce risk and take into account that they may end up in the divorced state, location decisions at the couple-level will give less weight to locations for the jointly optimising couple. Instead, they will attach more weight to locations in which either partner has decent career opportunities, ensuring a minimum satisfactory labour market option for both partners is attainable in the chosen location to insure against the possibility of divorce. This may lead to less efficient specialisation in marriage, but also less disparate labour market outcomes in case of its dissolution.

Regarding mothers' labour market outcomes after the reform, I conclude that the location constraint outweighs the benefits from a lighter care burden. These findings do not imply that co-parenting is not in the best interest of the child(ren) of divorcing parents or should otherwise be discouraged. Prior research finds upsides for child development and well-being, including that of parents. What my analysis does show, however, is that family policies targeting improved outcomes for one demographic (here, children) may have inadvertent economic consequences for another (here, parents, specifically mothers).

My findings also carry implications for policy design, particularly in the context of emerging work arrangements. Flexibility in work location may be a powerful tool for circumventing the geographical constraints of co-parenting, allowing mothers to access better

employment opportunities without compromising the post-divorce family arrangement. This insight is particularly salient in a post-pandemic world, where flexible work has become more widespread and offers a potential policy solution to the location constraint. Specifically, the ability to work remotely decouples a worker's physical location from their professional labour market, effectively freeing them from the co-location constraint.

Future research may exploit the reform or a similar setting to understand how co-parenting as a common post-divorce custody arrangement affects the labour supply and location decisions of intact couples in light of the changing outside option. There is already some very interesting work on the co-determination of marriage market and labour market outcomes through location choice (Alonzo, 2025; Foerster and Ulbricht, 2023), but these do not yet take into account post-divorce ties through shared children.

Bibliography

Alon, T., S. Coskun, M. Doepke, D. Koll, and M. Tertilt (2022): "From mancession to shecession: Women's employment in regular and pandemic recessions," *NBER Macroeconomics Annual*, 36, 83–151.

Alonzo, D. (2025): "Marrying Your Job: Matching and Mobility with Geographic Heterogeneity," *working paper*.

Altindag, D. T., J. Nunley, and A. Seals (2017): "Child-custody reform and the division of labor in the household," *Review of Economics of the Household*, 15, 833–856.

Augustijn, L. (2023): "Physical custody arrangements and fathers' post-separation well-being," *Journal of Family Studies*, 29, 2008–2024.

Bakker, W. and C. H. Mulder (2013): "Characteristics of post-separation families in the Netherlands: Shared residence versus resident mother arrangements," *GeoJournal*, 78, 851–866.

Bang, M. and H. Wang (2025): "Job Search and Mobility Over the Life-Cycle: Implications for the Child Penalty," .

Bauserman, R. (2002): "Child adjustment in joint-custody versus sole-custody arrangements: a meta-analytic review." *Journal of Family Psychology*, 16, 91.

Becker, G. S. (1985): "Human capital, effort, and the sexual division of labor," *Journal of Labor Economics*, 3, S33–S58.

Berresheim, U. and D. Koll (2023): "Staying together forever? Life-cycle effects of overoptimistic couples," *working paper*.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How much should we trust differences-in-differences estimates?" *The Quarterly Journal of Economics*, 119, 249–275.

Bonnet, C., B. Garbinti, and A. Solaz (2018): "Does part-time mothering help get a job? The role of shared custody in women's employment," *European Journal of Population*.

Braun, C., C. Nusbaum, and P. Rupert (2021): "Labor market dynamics and the migration behavior of married couples," *Review of Economic Dynamics*, 42, 239–263.

Brown, M., C. J. Flinn, and J. Mullins (2025): "Family Law Effects on Divorce, Fertility, and Child Investment," *Journal of Labor Economics*, 43, S351–S397.

Calvo, P. (2022): "The Effects of Institutional Gaps between Cohabitation and Marriage," *Working Paper*.

Calvo, P., I. Lindenlaub, and A. Reynoso (2024): "Marriage market and labour market sorting," *Review of Economic Studies*, 91, 3316–3361.

Cancian, M. and D. R. Meyer (1998): "Who gets custody?" *Demography*, 35, 147–157.

Chan, M. K. and K. Liu (2025): "Changing Families: Family Relationships, Parental Decisions, and Child Development," *Journal of Labor Economics*, 43, S399–S444.

Claessens, E. and D. Mortelmans (2025): "Joint Physical Custody in Europe: A Comparative Exploration," *European Journal of Population*, 41, 1–21.

Costa, D. L. and M. E. Kahn (2000): "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990," *The Quarterly Journal of Economics*, 115, 1287–1315.

Coviello, D., E. Deserranno, and N. Persico (2024): "Gender Disparities in the Welfare Effect of the Minimum Wage," *working paper*.

Dasgupta, K., A. C. Johnston, L. Kirkpatrick, M. N. Massenkoff, and A. Plum (2025): "Social and Health Outcomes around Divorce: Evidence from New Zealand," Working Paper 33873, National Bureau of Economic Research.

Doepke, M. and M. Tertilt (2016): "Families in macroeconomics," in *Handbook of macroeconomics*, Elsevier, vol. 2, 1789–1891.

Fernández, R. and J. C. Wong (2014): "Divorce risk, wages and working wives: A quantitative life-cycle analysis of female labour force participation," *The Economic Journal*, 124, 319–358.

Foerster, H. and R. Ulbricht (2023): "The Colocation Friction: Dual-Earner Job Search and Labor Market Outcomes," *working paper*.

Foged, M. (2016): "Family Migration and Relative Earnings Potentials," *Labour Economics*, 42, 87–100.

Fredriksson, P., D. Gülümser, and L. Hensvik (2025): "Outside Job Opportunities and the Gender Gap in Pay," Tech. rep., Rockwool Foundation Berlin (RF Berlin).

Frimmel, W., M. Halla, and R. Winter-Ebmer (2024): "How does parental divorce affect children's long-term outcomes?" *Journal of Public Economics*, 239, 105201.

Gemici, A. (2023): "Family Migration and Labor Market Outcomes," .

Goussé, M., N. Jacquemet, and J.-M. Robin (2017a): "Household labour supply and the marriage market in the UK, 1991-2008," *Labour Economics*, 46, 131–149.

——— (2017b): "Marriage, labor supply, and home production," *Econometrica*, 85, 1873–1919.

Goussé, M. and M. Leturcq (2022): "More or less unmarried. The impact of legal settings of cohabitation on labour market outcomes," *European Economic Review*, 149, 104259.

Guler, B., F. Guvenen, and G. L. Violante (2012): "Joint-search theory: New opportunities and new frictions," *Journal of Monetary Economics*, 59, 352–369.

Hakovirta, M., D. R. Meyer, M. Salin, E. Lindroos, and M. Haapanen (2023): "Joint physical custody of children in Europe," *Demographic Research*, 49, 479–492.

Halla, M. (2013): "The Effect of Joint Custody on Family Outcomes," *Journal of the European Economic Association*, 11, 278–315.

——— (2015): "Do joint custody laws improve family well-being?" *IZA World of Labor*.

Halla, M. and C. Hölzl (2007): *Bargaining at Divorce: The Allocation of Custody*, Department of Economics, University of Linz.

Hoynes, H., D. L. Miller, and J. Schaller (2012): "Who Suffers during Recessions?" *Journal of Economic Perspectives*, 26, 27–48.

Jayachandran, S., L. Nassal, M. J. Notowidigdo, M. Paul, H. Sarsons, and E. Sundberg (2024): "Moving to opportunity, together," Tech. rep., National Bureau of Economic Research.

Johnston, A. C., M. R. Jones, and N. G. Pope (2025): "Divorce, Family Arrangements, and Children's Adult Outcomes," Working Paper 33776, National Bureau of Economic Research.

Le Barbanchon, T., R. Rathelot, and A. Roulet (2021): "Gender differences in job search: Trading off commute against wage," *The Quarterly Journal of Economics*, 136, 381–426.

Meyer, D. R., M. Carlson, and M. M. U. Alam (2022): "Increases in shared custody after divorce in the United States," *Demographic Research*, 46, 1137–1162.

Nguyen, D. K., A. Tran Thi Van, and T. Phan (2018): "Child Custody and Family Labour Supply: Evidence from the United States," *Labour*, 32, 74–92.

Nielsen, L. (2018): "Joint versus sole physical custody: Children's outcomes independent of parent–child relationships, income, and conflict in 60 studies," *Journal of Divorce & Remarriage*, 59, 247–281.

Nunley, J. M. and R. A. Seals (2011): "Child-custody reform, marital investment in children, and the labor supply of married mothers," *Labour Economics*, 18, 14–24.

Poortman, A., T. Van der Lippe, and K. Boele-Woelki (2014): "Codebook of the survey New Families in the Netherlands (NFN). First wave." Tech. rep., Utrecht: Utrecht University.

Poortman, A. and R. Van Gaalen (2019): "New Families in the Netherlands (NFN): Wave 1," *DANS*.

Poortman, A.-R. and R. van Gaalen (2017): "Shared Residence After Separation: A Review and New Findings from the Netherlands," *Family Court Review*, 55, 531–544.

Ranosova, T. (2025): "Commute, Specialization and the Value of Marriage," .

Rasul, I. (2006): "The Economics of Child Custody," *Economica*, 73, 1–25.

Reynoso, A. (2019): "Marriage, marital investments, and divorce: Theory and evidence on policy non neutrality," *working paper*.

Rueda, V. and G. Wilemme (2025): "Career Paths with a Two-Body Problem: Colocation and Gendered Professional Crossroads," *Available at SSRN 5435755*.

Spruijt, E. and V. Duindam (2009): "Joint physical custody in the Netherlands and the well-being of children," *Journal of Divorce & Remarriage*, 51, 65–82.

Venator, J. (2024): "Dual-earner migration decisions, earnings, and unemployment insurance," *working paper*.

Voena, A. (2015): "Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples?" *American Economic Review*, 105, 2295–2332.

Vowels, L. M., C. L. Comolli, L. Bernardi, D. Chacón-Mendoza, and J. Darwiche (2023): "Systematic review and theoretical comparison of children's outcomes in post-separation living arrangements," *PLOS ONE*, 18, 1–23.

Vuri, D. (2018): "Joint custody law and mothers' labor market outcomes: Evidence from the USA," *Journal of Population Economics*, 31, 1203–1237.

Weiss, Y. and R. J. Willis (1985): "Children as Collective Goods and Divorce Settlements," *Journal of Labor Economics*, 3, 268–292.

Weston, R., L. Qu, M. Gray, J. De Maio, R. Kaspiew, L. Moloney, and K. Hand (2011): "Shared care time: An increasingly common arrangement?" *Family Matters*, 51–56.

Wu, Z. and M. S. Pollard (2000): "Economic circumstances and the stability of nonmarital cohabitation," *Journal of Family Issues*, 21, 303–328.

Appendix

Appendix A: Additional tables

Table A1.1: Increase in co-parenting

	Youngest child up to 10 (1)	All minor children (2)
Post-reform	0.076	0.067
	(0.0295)	(0.0218)
Constant	0.286	0.292
	(0.023)	(0.019)
Mean dep. var	0.345	0.327
Observations	2095	3555

Note. This table shows the coefficients of a regression of co-parenting status on a post-reform dummy. Standard errors in parentheses.

Table A1.2: Top-10 sector shares by gender

	Mothers — Pre-reform		Mothers — Post-reform		Fathers — Pre-reform		Fathers — Post-reform	
Rank	Sector	Share	Sector	Share	Sector	Share	Sector	Share
1	Health & social work	0.282	Health & social work	0.289	Other business activities	0.152	Other business activities	0.138
2	Other business activities	0.155	Other business activities	0.151	Wholesale trade	0.098	Wholesale trade	0.099
3	Retail trade	0.092	Retail trade	0.089	Construction	0.087	Construction	0.088
4	Education	0.079	Education	0.078	Public administration	0.080	Public administration	0.080
5	Public administration	0.077	Public administration	0.073	Land transport & pipelines	0.045	Health & social work	0.051
6	Wholesale trade	0.043	Wholesale trade	0.043	Health & social work	0.045	Land transport & pipelines	0.047
7	Financial intermediation	0.037	Financial intermediation	0.033	Financial intermediation	0.040	Financial intermediation	0.040
8	Hotels & restaurants	0.026	Hotels & restaurants	0.028	Retail trade	0.038	Retail trade	0.036
9	Recreation, culture & sport	0.017	Recreation, culture & sport	0.016	Education	0.035	Education	0.036
10	Other service activities	0.015	Construction	0.014	Computer & related activities	0.030	Computer & related activities	0.034

Note. Labels follow SBI-2003 (aligned with NACE Rev. 1.1) at the 2-digit division level. Shares are means of individual sector shares within each subgroup (pre/post, women/men), measured one year before divorce. Only the top 10 per subgroup are shown.

Table A1.3: NFN sample: Demographic and labour market characteristics of divorcees before and after the reform

	Pre-Reform		Post-Reform			*p*-values	
	Mean	Obs	Mean	Obs	Post-Pre	unadj.	Holm
MOTHERS							
Age at divorce	36.35	450	36.58	2324	0.22	0.285	1.000
High education	0.30	446	0.29	2304	-0.01	0.654	1.000
Worked before divorce	0.78	450	0.79	2321	0.00	0.813	1.000
Hours (weekly)	18.56	445	19.1	2304	0.53	0.269	1.000
FATHERS							
Age at divorce	39.1	450	39.65	2324	0.55	0.021	0.147
High education	0.33	442	0.34	2291	0.01	0.665	1.000
Worked before divorce	0.93	448	0.92	2322	-0.01	0.412	1.000
Hours (weekly)	38.23	446	36.5	2304	-1.73	0.001	0.006
CHILDREN							
Number of children	2.24	450	1.98	2324	-0.26	0.000	0.000
Age of youngest child	4.61	450	5.04	2324	0.44	0.000	0.001

Note. Pre- and post-reform means of demographic variables at the time of divorce and their difference. The last two columns report conventional (unadjusted) *p*-values and those adjusted for multiple hypothesis testing using the Holm correction. Employment status and hours worked refer to the year before divorce. High education refers to the share of individuals with higher vocational training or a university degree.

Table A1.4: Chow test for structural break in number of divorces after March 2009

	Number of divorces
Month	-7.151
	(0.667)
Constant	5172.9
	(399.3)
Chow Chi2	3.59
p-value	0.166
Mean dep. var	957.583
Observations	60

Note. Chow test for structural break in the number of divorces after March 2009 using robust standard errors. Standard errors in parentheses.

Table A1.5: Mothers' labour market outcomes - simple DiD

	Employment (1)	Log earnings (2)	Log hours (3)	Log wages (4)
Post	0.015	0.037	0.033	0.004
	(0.003)	(0.003)	(0.002)	(0.002)
Post × Treat	0.002	-0.014	-0.006	-0.008
	(0.004)	(0.004)	(0.003)	(0.003)
Month	✓			
Month×Region		✓	✓	✓
Observations	1,176,990	702,173	702,173	702,173
Individuals	28,279	22,140	22,140	22,140

Note. Specification includes individual and month-year (Column 1) or month-year×region (Columns 2–4) fixed effects. Estimated across 18 months before and 24 months after the reform. Standard errors are clustered at the individual level.

Table A1.6: Fathers' labour market outcomes - simple DiD

	Employment (1)	Log earnings (2)	Log hours (3)	Log wages (4)
Post	0.003	0.003	0.003	-0.000
	(0.002)	(0.002)	(0.001)	(0.002)
Post × Treat	-0.006	-0.003	-0.004	0.001
	(0.004)	(0.003)	(0.002)	(0.003)
Month	✓			
Month×Region		✓	✓	✓
Observations	1,086,480	770,165	770,165	770,165
Individuals	26,232	22,160	22,160	22,160

Note. Specification includes individual and month-year (Column 1) or month-year×region (Columns 2–4) fixed effects. Estimated across 18 months before and 24 months after the reform. Standard errors are clustered at the individual level.

Table A1.7: Wages - mothers of older children

	Young children			Older children		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treat	-0.008	-0.008	-0.006	0.001	0.000	0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Month	✓			✓		
Month×Region		✓			✓	
Month×Region×Sector			✓			✓
Observations	702,173	702,173	702,173	460,561	457,004	434,970
Individuals	22,140	22,140	22,140	13,360	13,350	13,180

Note. "Young children" refers to mothers whose youngest child is 10 years old or younger. "Older children" refers to mothers whose youngest child is aged 11-17. Standard errors are clustered at the individual level.

Table A1.8: Wages

	Baseline	Doughnut hole	Regional unemployment	Single-establishment firms
	(1)	(2)	(3)	(4)
Post × Treat	-0.008	-0.008	-0.007	-0.008
	(0.003)	(0.003)	(0.003)	(0.005)
Month×Region	✓	✓	✓	✓
Observations	702,173	665,261	702,173	186,623
Individuals	22,140	20,966	22,140	9,422

Note. Specifications include individual and month-year×region fixed effects. Standard errors are clustered at the individual level and shown in parentheses. Unemployment is computed at the quarter-region level using the Dutch Labour Force Survey.

Table A1.9: Ex-spouses move less to employers further away after the reform

	Distance (in metres) (1)	Same region (2)	Same municipality (3)
Post	823.51	-0.0079	-0.0087
	(208.37)	(0.0029)	(0.0026)
Post × Treat	-990.59	0.0092	0.0115
	(370.01)	(0.0050)	(0.0044)
Mean dep. var	26105.38	0.5437	0.2713
Observations	457168	457168	457168
Individuals	23616	23616	23616

Note. Specifications include individual and month-year fixed effects. Estimated across 18 months before and 24 months after the reform. The sample is restricted to single-establishment firms. Distance in metres. Standard errors are clustered at the individual level.

Table A1.10: Probability of receiving spousal alimony by co-parenting status and gender

	Alimony to ex-wife (1)	Alimony to ex-husband (2)	No alimony payments (3)
Co-parenting	-0.002	0.001	0.001
	(0.012)	(0.002)	(0.012)
Constant	0.131	0.006	0.863
	(0.007)	(0.001)	(0.007)
Mean dep. var.	0.180	0.008	0.813
Observations	5294	5294	5294

Note. The table shows coefficients from regressions of the probability of receiving alimony on co-parenting status by parents' gender. The three mean dependent variables do not add up to one perfectly because of rounding. Data: NFN.

Appendix B: Additional event studies

Figure B1.1: The reform does not positively affect mothers' post-divorce employment

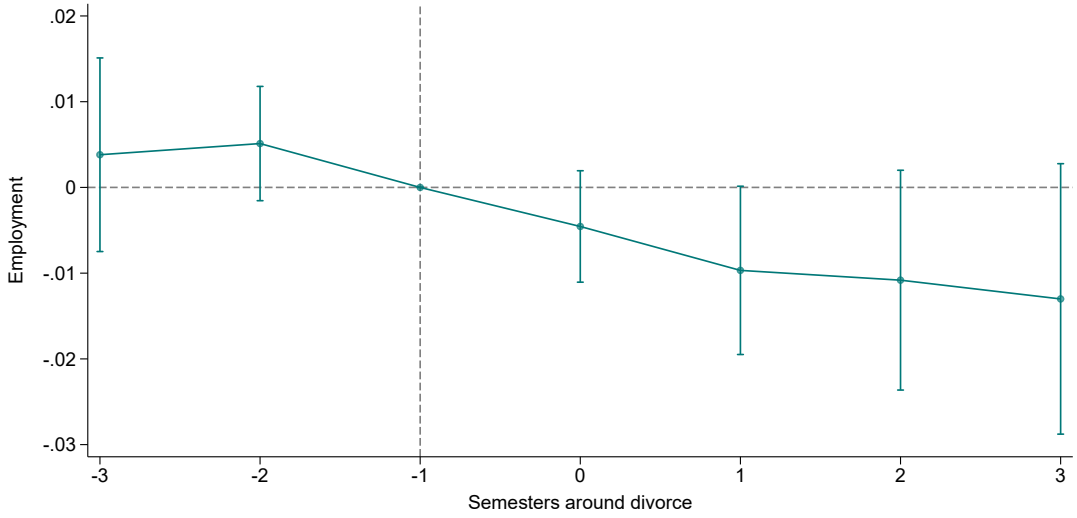

Note. This figure shows event study estimates of the effect of divorcing after the reform on mothers' probability of being employed. The sample includes mothers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in employment rates between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.2: Women's earnings decrease by about 1%

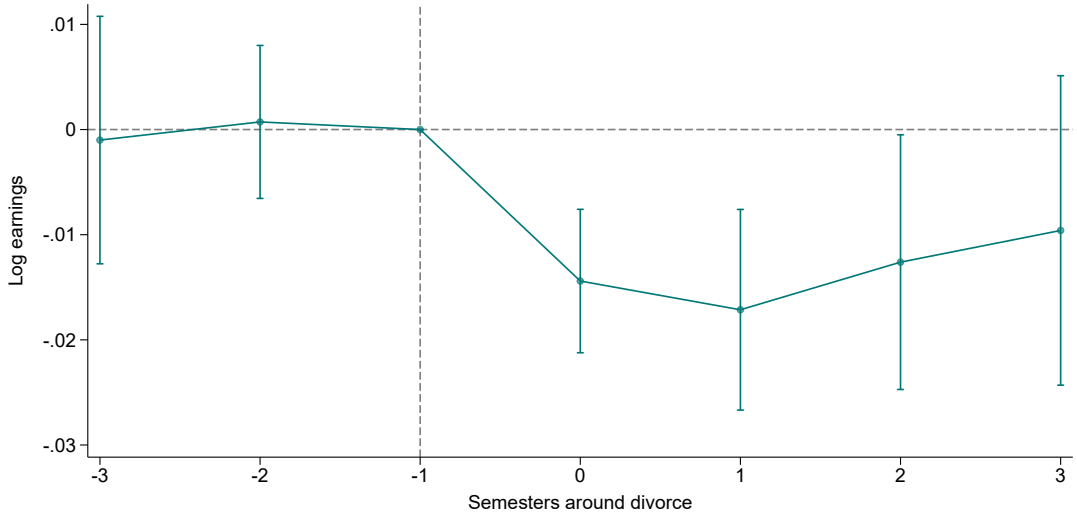

Note. This figure shows event study estimates of the effect of divorcing after the reform on mothers' monthly log earnings. The sample includes mothers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log earnings between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.3: Mothers' hours by treatment status

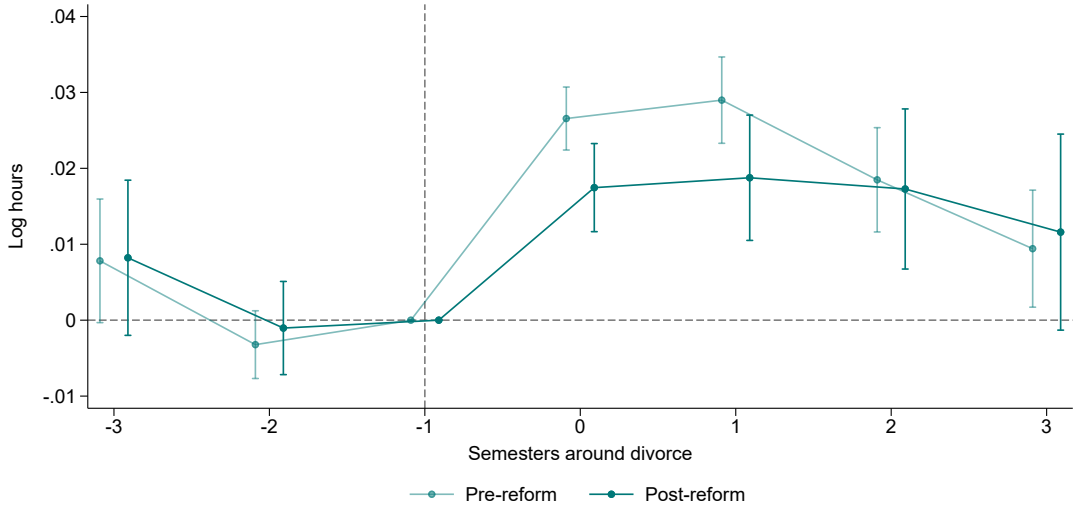

Note. This figure shows event study estimates of the changes in log hours relative to the time of divorce for mothers with young children divorcing before and after the reform. The pre-reform sample includes mothers divorcing between September 2007 and February 2009; the post-reform sample those divorcing between March 2009 and August 2010. The reference category is the semester immediately before divorce (semester -1). The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.4: Mothers' wages by treatment status

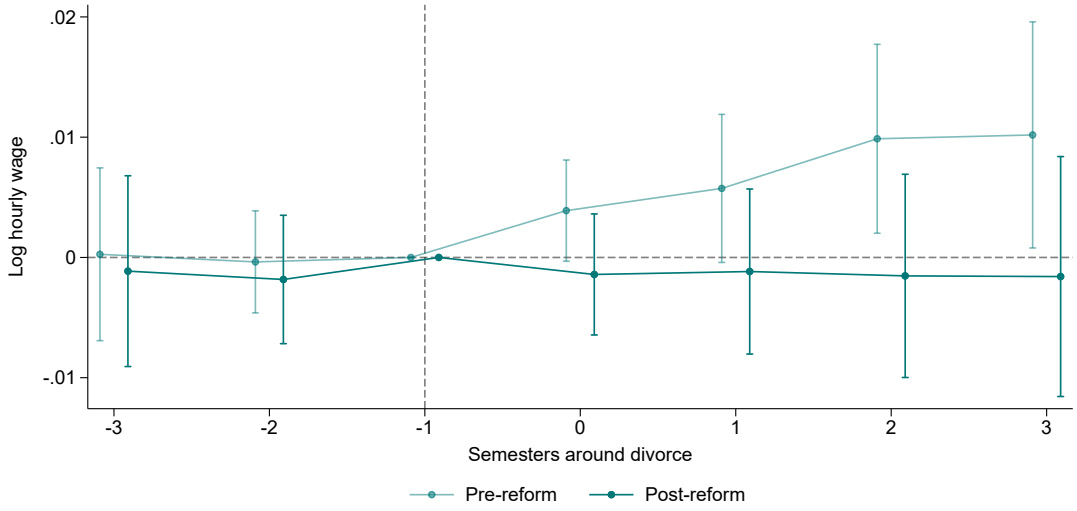

Note. This figure shows event study estimates of the changes in log wages relative to the time of divorce for mothers with young children divorcing before and after the reform. The pre-reform sample includes mothers divorcing between September 2007 and February 2009; the post-reform sample those divorcing between March 2009 and August 2010. The reference category is the semester immediately before divorce (semester -1). The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.5: Distance between ex-spouses' workplaces by treatment status

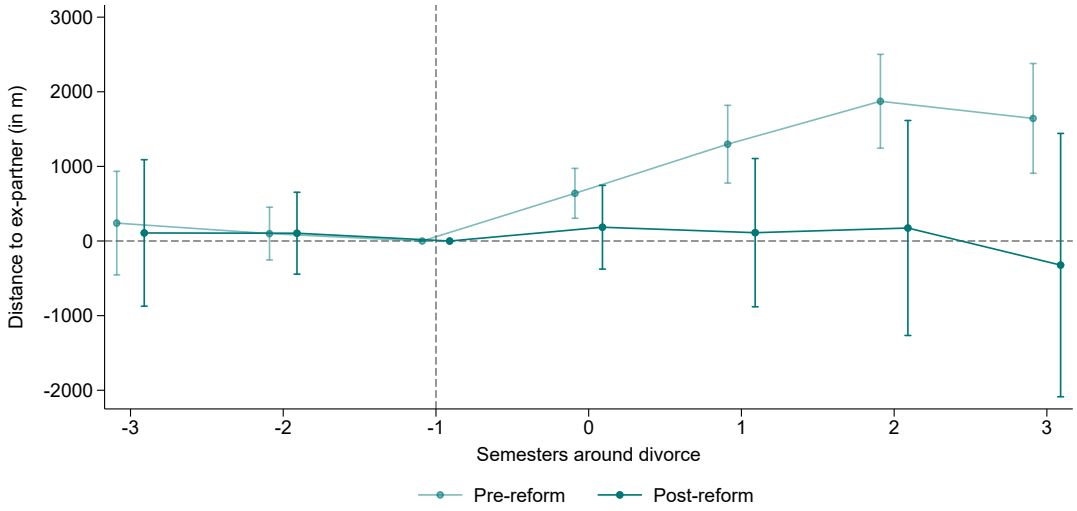

Note. This figure shows event study estimates of the changes in distance between ex-spouses' workplaces (in metres) relative to the time of divorce for parents with young children divorcing before and after the reform. The pre-reform sample includes couples divorcing between September 2007 and February 2009; the post-reform sample those divorcing between March 2009 and August 2010. The reference category is the semester immediately before divorce (semester -1). The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.6: Mothers' employment by treatment status

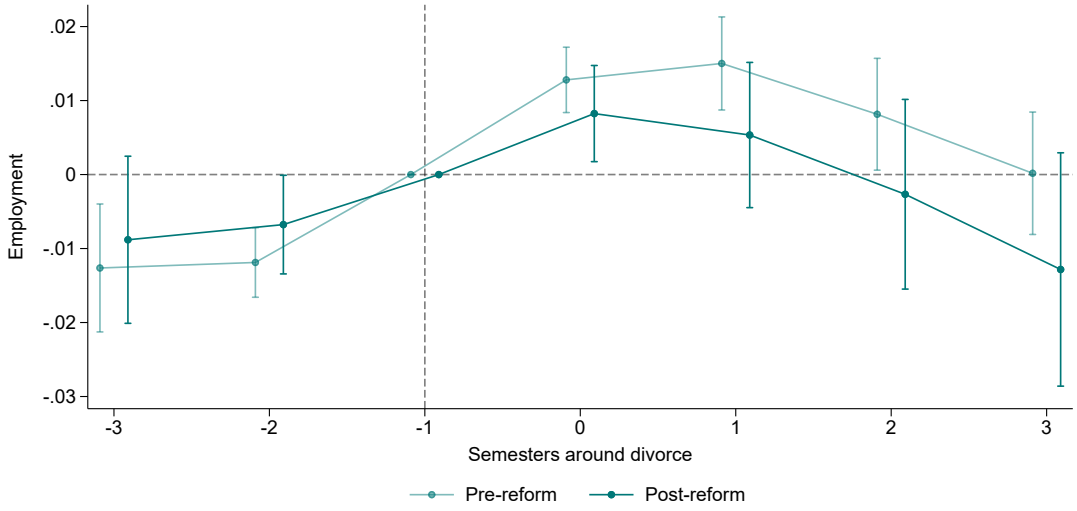

Note. This figure shows event study estimates of the changes in employment probability relative to the time of divorce for mothers with young children divorcing before and after the reform. The pre-reform sample includes mothers divorcing between September 2007 and February 2009; the post-reform sample those divorcing between March 2009 and August 2010. The reference category is the semester immediately before divorce (semester -1). The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.7: Mothers' earnings by treatment status

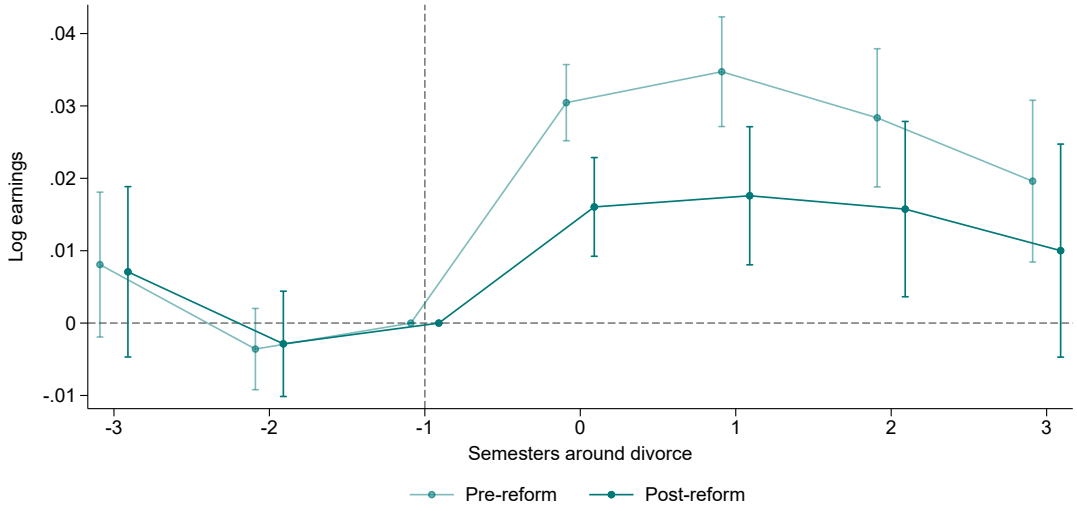

Note. This figure shows event study estimates of the changes in log earnings relative to the time of divorce for mothers with young children divorcing before and after the reform. The pre-reform sample includes mothers divorcing between September 2007 and February 2009; the post-reform sample those divorcing between March 2009 and August 2010. The reference category is the semester immediately before divorce (semester -1). The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.8: Fathers' employment

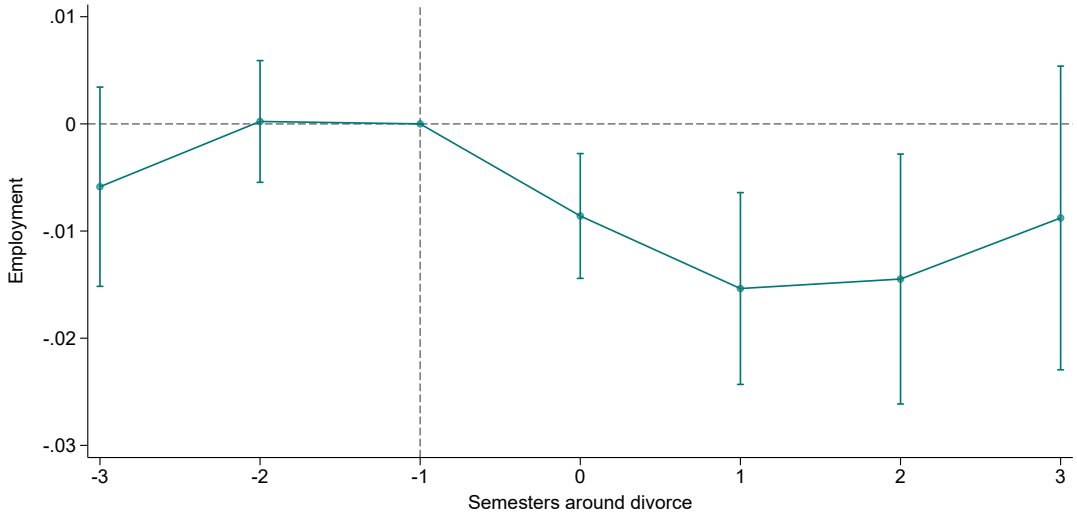

Note. This figure shows event study estimates of the effect of divorcing after the reform on fathers' probability of being employed. The sample includes fathers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in employment rates between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.9: Fathers' hours

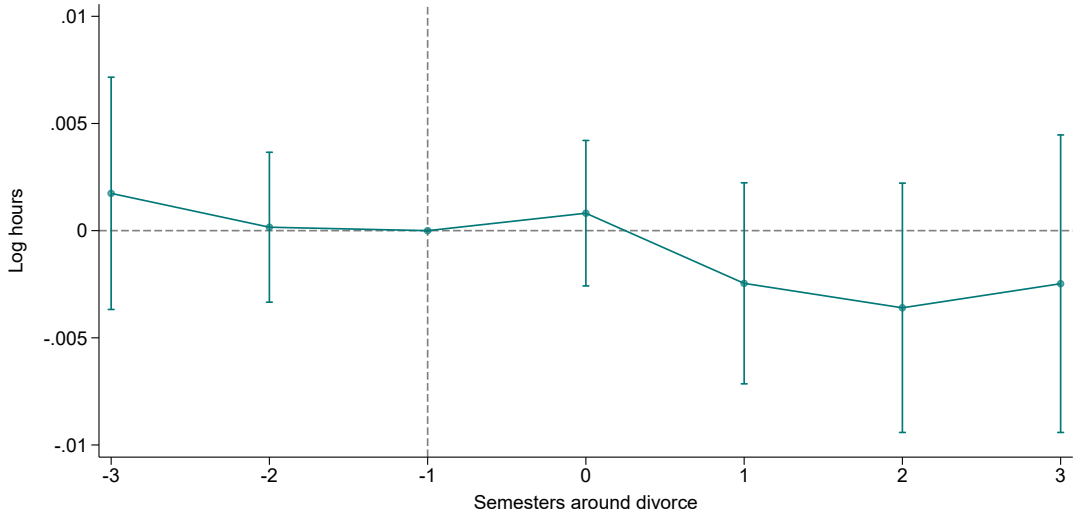

Note. This figure shows event study estimates of the effect of divorcing after the reform on fathers' monthly log hours. The sample includes fathers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log hours between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.10: Fathers' wages

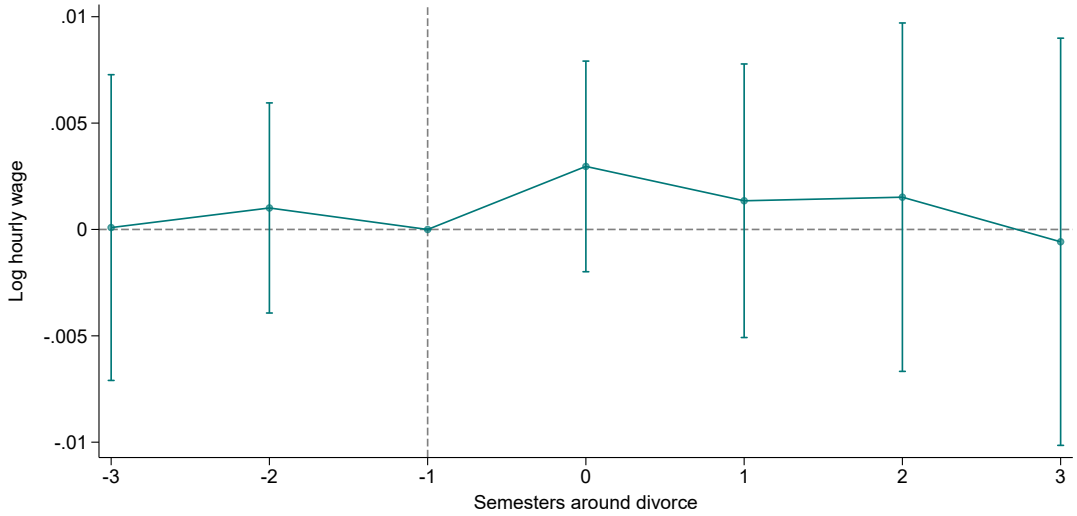

Note. This figure shows event study estimates of the effect of divorcing after the reform on fathers' hourly log wages. The sample includes fathers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log wages between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.11: Fathers' earnings

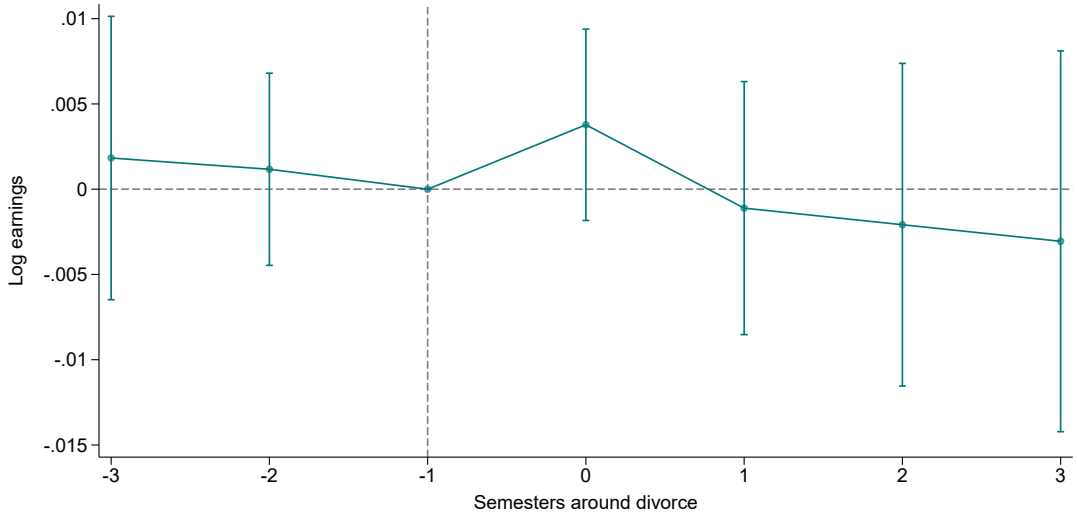

Note. This figure shows event study estimates of the effect of divorcing after the reform on fathers' monthly log earnings. The sample includes fathers with young children divorcing between September 2007 and August 2010, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log earnings between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure B1.12: Continuously employed mothers

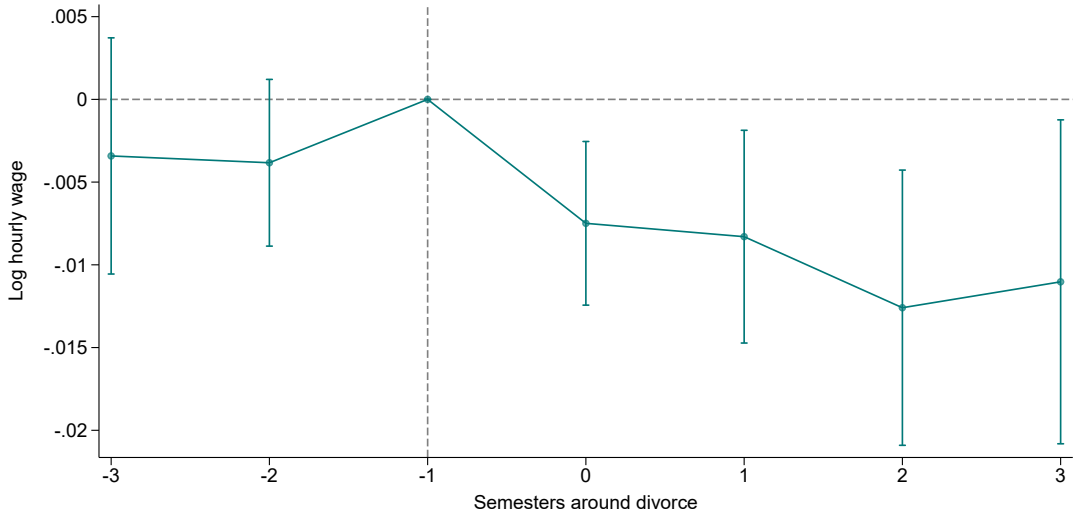

Note. This figure shows event study estimates of the effect of divorcing after the reform on mothers' hourly log wages. The sample includes mothers with young children divorcing between September 2007 and August 2010 who are continuously employed during the event time around divorce, and the coefficients represent interactions between post-divorce semester dummies and a treatment indicator for those divorcing after the joint custody reform. The omitted category is the semester immediately before divorce (semester -1). Each point estimate reflects the difference in log earnings between treated and control individuals in a given semester relative to this reference period. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Appendix C: Additional figures

Figure C1.1: Searches for "parenting plan" (*ouderschapsplan*) spike in March 2009 and remain elevated thereafter.

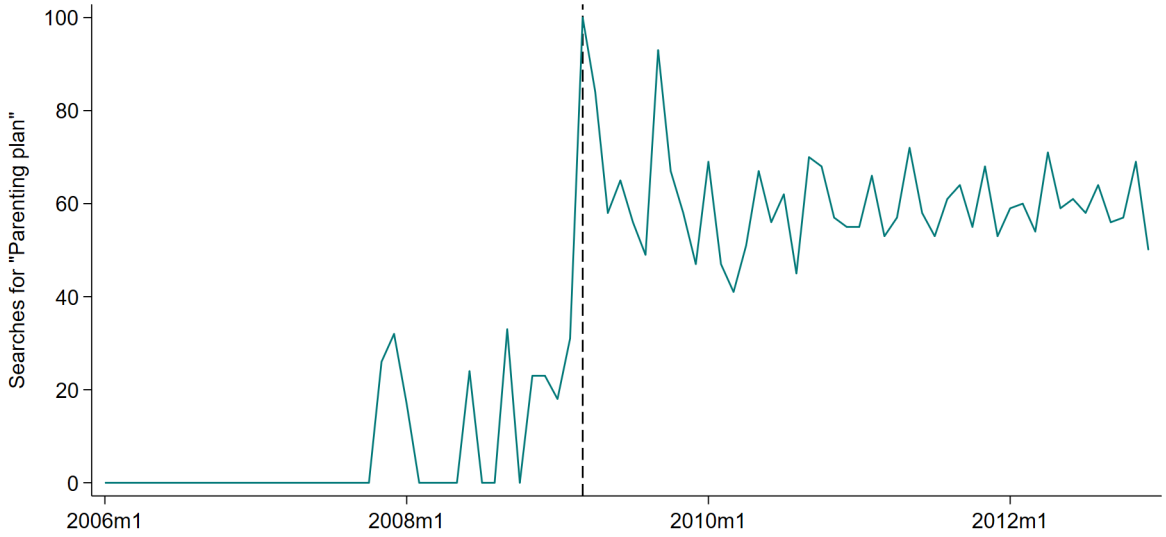

Note. Google trends data for the Dutch term for "parenting plan" between 2006 and 2012. Google trends data are reported as a value relative to the highest search volume, which receives a value of 100 (here, in March 2009). (Data: Google Trends, Netherlands; search term 'ouderschapsplan').

Figure C1.2: Compliers are mostly parents with children aged up to 10

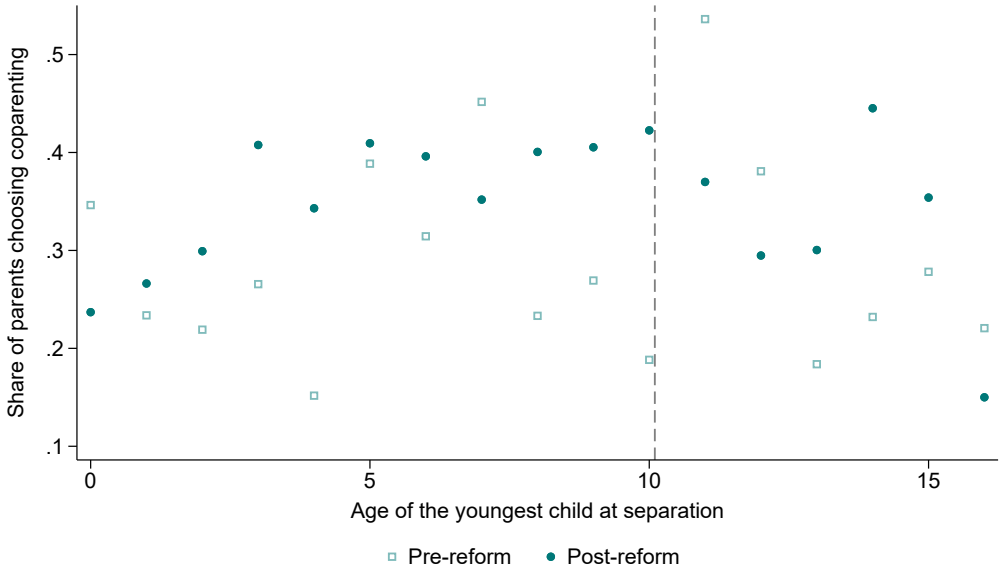

Note. This figure shows the share of parents who choose a co-parenting arrangement after divorce by the age of the youngest child and whether they were affected by the reform (Data: NFN).

Figure C1.3: Post-divorce child custody arrangements in the Netherlands by education of the mother, 2010.

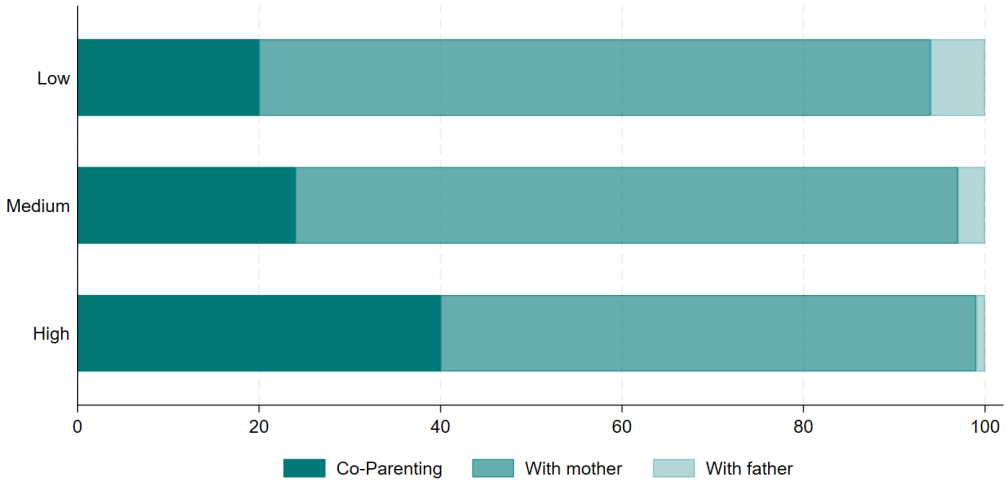

Note. This figure shows the share of mothers who choose co-parenting arrangements by educational attainment (Data: CBS). Co-parenting is most common among highly-educated mothers.

Figure C1.4: Distance between divorced parents with minor children by co-parenting status

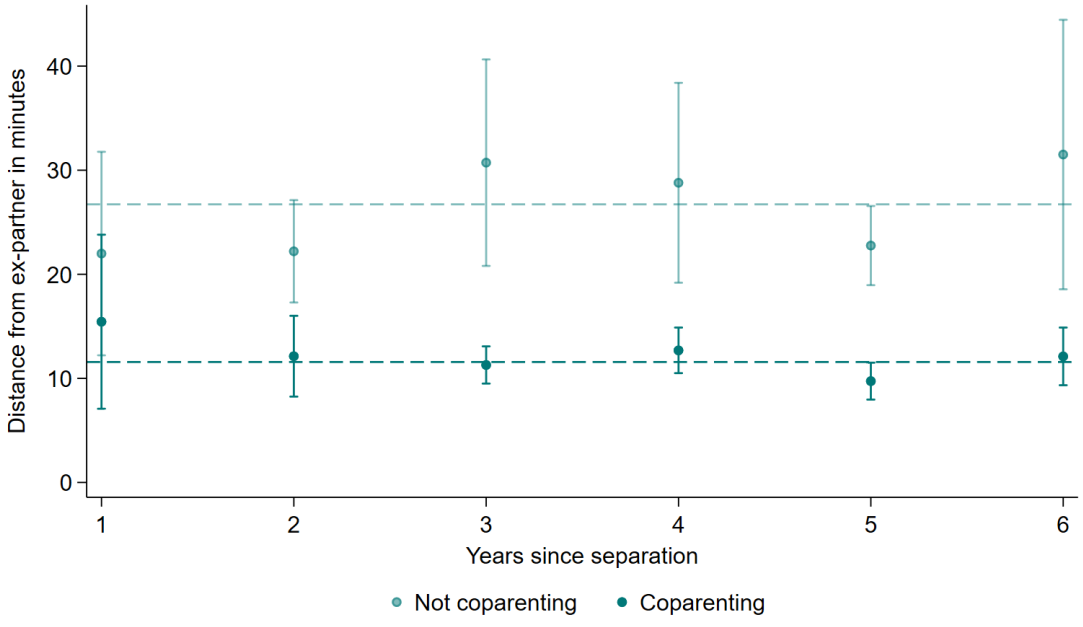

Note. This figure shows the travel distance between former spouses in minutes by year of separation. Markers plot coefficients of a linear model in which distance between ex-partners is regressed on years since separation by co-parenting status. The horizontal lines indicate the mean distances by co-parenting status. Vertical capped lines are 95% confidence intervals based on robust standard errors. Own calculations based on NFN data.

Figure C1.5: Treatment and control groups are balanced in trends

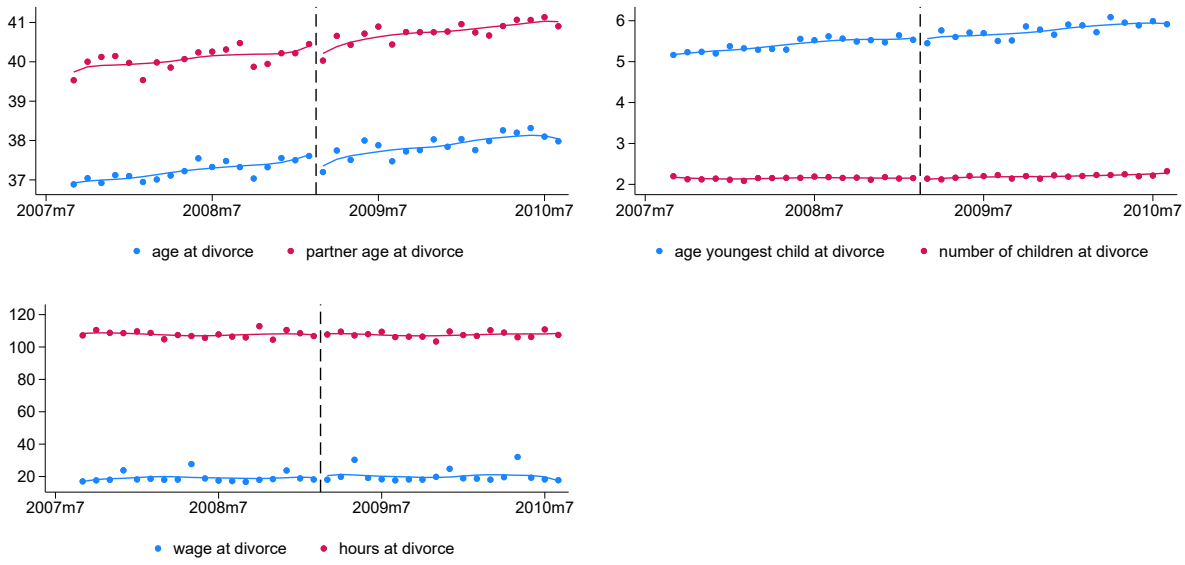

Note. This figure shows time trends of the variables for which the treatment and control samples are not balanced in levels.

Figure C1.6: There is no bunching right before the reform

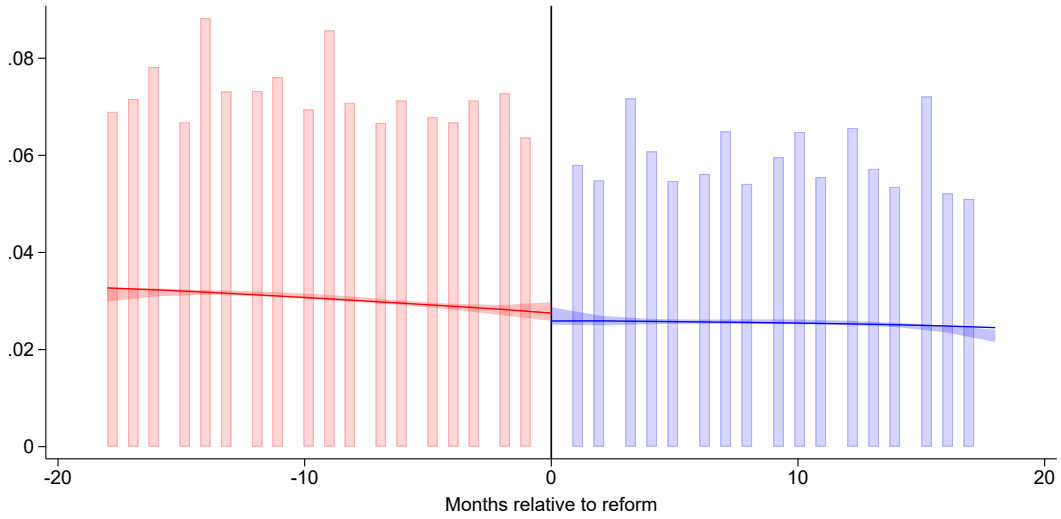

Note. This figure shows the density of divorces by month in the 18 months before and after the reform, and the 95% confidence bands based on the Cattaneo-Janssen-Ma density test.

Figure C1.7: The wage effect is driven by mothers with young children

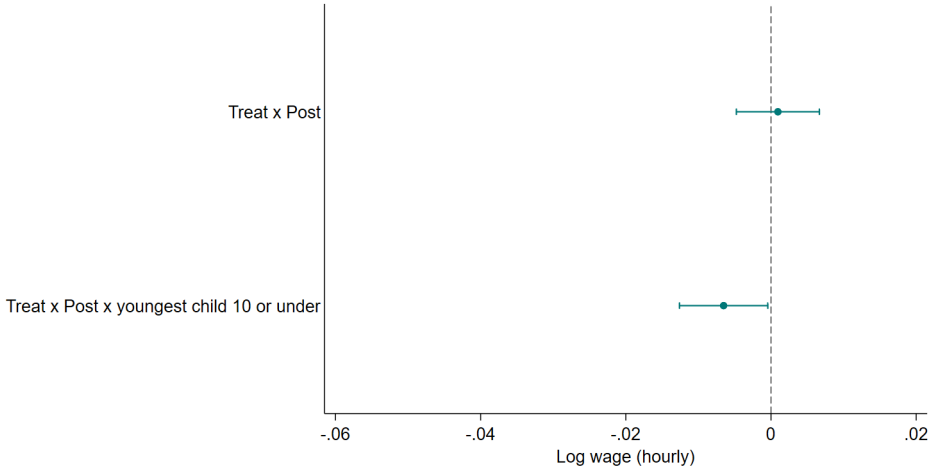

Note. This figure shows regression coefficients of a DiD regression of wages on a triple interaction between a post-reform indicator, a treatment group indicator, and an indicator for the age of the youngest child. The sample includes mothers divorcing between September 2007 and August 2010. The specification includes month-year×region and individual fixed effects. Horizontal bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure C1.8: Wage decreases are driven by those who stay in the same region as their partner: triple difference

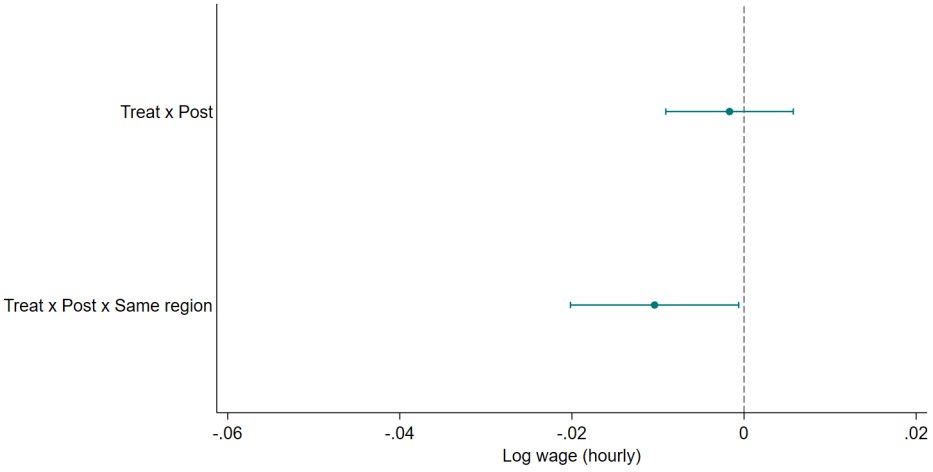

Note. This figure shows regression coefficients of a DiD regression of wages on a triple interaction between a post-reform indicator, a treatment group indicator, and an indicator for whether the ex-spouses work in the same region. The sample includes mothers with young children divorcing between September 2007 and August 2010. The specification includes month-year×region and individual fixed effects. Horizontal bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure C1.9: Wage effect by cut-off age

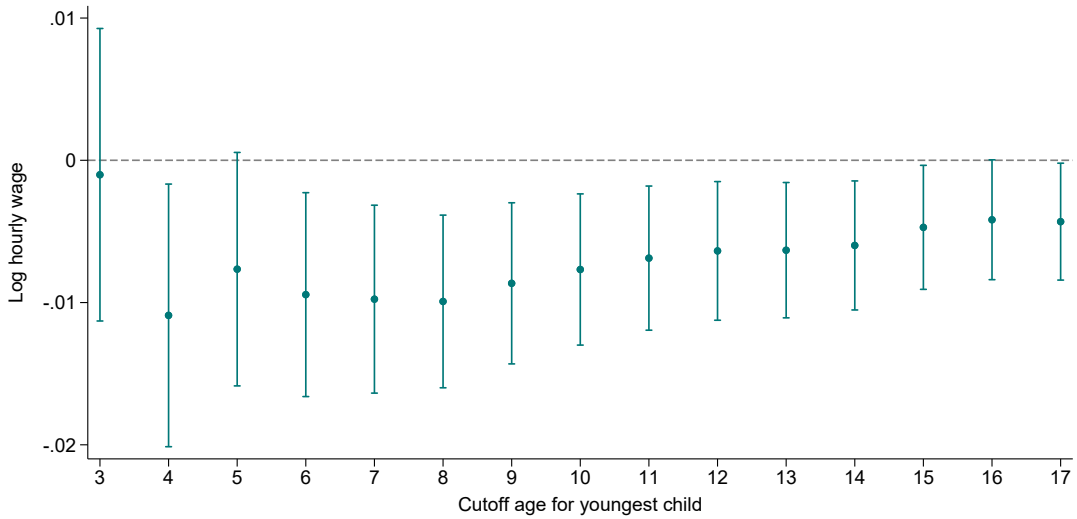

Note. This figure shows the effect of the reform on mothers' wages by cut-off age for inclusion in the sample. The specification includes month-year×region and individual fixed effects. Vertical bars indicate 95% confidence intervals, with standard errors clustered at the individual level.

Figure C1.10: Mothers who co-parent report less time pressure

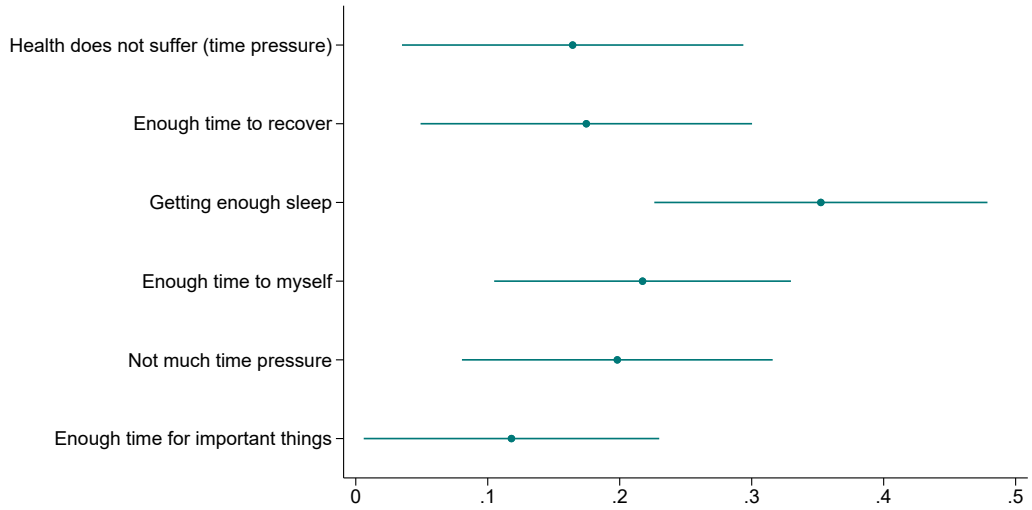

Note. This figure shows coefficients of a DiD regression of different outcomes on the interaction between indicators for co-parenting and being a woman (Data: NFN).

Appendix D: Detailed derivations

Setup

Having substituted the period budget constraint, parents maximise expected life-time utility (post-divorce):

$$\max_{\{h_i,x_i\}} U_i = \mathbf{E}\left[\sum_{t=t_{div}}^{T} \beta^t \Big(\nu(w_i h_i - M_i^R) - \eta(h_i) + \phi(\sigma_i)\Big)\right]$$

subject to the time and possible distance constraints as well as commuting:

$$\begin{aligned} l_i + h_i + k_i + q_i^R &= 1 && \text{time constraint} \\ \|x_i - x_{r,i}\| &\leq \bar{d} && \text{distance constraint (under JPC only)} \\ k_i &= d(x_i, x_{r,i}) && \text{commuting, where } d'(x_i, x_{r,i}) > 0 \end{aligned} \tag{1.3}$$

In addition to the environment outlined in Section 1.3, I make the assumption that the utility function is moderately concave in consumption: $-\nu''(c) \cdot c < \nu'(c)$. This bound on the degree of concavity of the utility function implies that the function is concave, but not sharply curved, i.e. that marginal utility does not decline rapidly. A common example is CRRA utility with relative-risk aversion parameter $\gamma < 1$.[2] This assumption is used in the derivation of the result on hours as it allows to determine the sign of the derivative of hours worked with respect to wages.

Solving

Lagrangian

$$\max_{\{x_i,h_i\}} \mathcal{L} = \nu(w_i(x_i)h_i - M_i^R) - \eta(h_i) + \phi(\sigma_i) + \lambda_i(1 - l_i - h_i - d(x_i) - q_i^R) + \mu_i(\bar{d}^2(x_i, x_{r,i}) - \|x_i - x_{r,i}\|^2)$$

FOCs

$$\frac{\partial \mathcal{L}}{\partial h_i} : \nu'(w_i(x_i)h_i - M_i^R) \cdot w_i(x_i) = \eta'(h_i) + \lambda_i \tag{FOC 1}$$

[2] Connecting to expected utility theory, this is equivalent to saying that the utility function exhibits bounded absolute risk aversion, with curvature constrained by $A(x) < \frac{1}{c}$. The utility function $u(\cdot)$ is strictly increasing and concave, with bounded absolute risk aversion:

$$A(x) := -\frac{\nu''(c)}{\nu'(c)} < \frac{1}{c} \quad \forall\, c > 0.$$

Intra-temporal Euler equation: marginal utility of consumption of working one additional hour = marginal cost of working an additional hour in terms of disutility of working and the opportunity cost of leisure. Importantly, there is also an income effect of moving from SPC to JPC: For mothers, M_m decreases, decreasing the marginal utility of income (the left-hand side of the first order condition). Thus, from the right-hand side, mothers' hours decrease. The opposite holds for fathers.

$$\frac{\partial \mathcal{L}}{\partial x_i} : \nu'(w_i(x_i)h_i - M_i^R) \cdot w_i'(x_i) \cdot h_i = \lambda_i \cdot d'(x_i, x_{r,i}) + 2\mu_i(x_i - x_{r,i}) \qquad \text{(FOC 2)}$$

Parents relocate for higher wages until the expected wage gradient benefit is exactly offset by (i) the commuting cost ($\lambda_i \cdot d'(x_i, x_{r,i})$) and/or, under joint custody, (ii) the distance constraint bites. Under SPC, the last term drops out as $\mu_i = 0$.

KKT

$$\mu_i(\bar{d}^2(x_i) - \|x_i - x_{r,i}\|^2) = 0$$

Under SPC, $\mu_i = 0$; under JPC, $\mu_i > 0$.

Comparative Statics

What happens to hours, wages, and commuting when switching from SPC to JPC? Wages should decrease for mothers who become location constrained and they should commute more. The theoretical prediction on hours is ambiguous. No effects are expected for fathers. Below, I derive each of these predictions in turn.

Hourly wages

Think of the wage at a location as a lottery ticket value $w(x)$. SPC lets a mother look at every ticket in the country and pick the highest-paying one. JPC removes any tickets lying beyond the distance $\bar{d}$. A simple "maximum over a subset" rule implies:

- If the best SPC ticket is outside the JPC circle (empirically common for Dutch mothers), the maximum inside the circle must be strictly lower.
- If the best SPC ticket is already inside the circle (assumed for fathers), the maximum is unchanged.

Below, I formally show that the wage weakly declines relative to SPC if the mobility constraint binds:

Proposition 1. *The switch from SPC to JPC lowers mothers' expected wages:*

$$w_m^{\text{JPC}} \leq w_m^{\text{SPC}}.$$

Proof. Let the mother's unconstrained optimum be x_m^{SPC}. The feasible set of locations under JPC is $\chi^{\text{JPC}} := \{x : \|x_i - x_{r,i}\| \leq \bar{d}\}$. Then

$$w_i^{\text{JPC}} = \max_{x_i \in \chi_{\bar{d}}^{\text{JPC}}} w_i(x_i) \leq w_i^{\text{SPC}} = \max_{x_i \in \chi^{\text{SPC}}} w_i(x_i)$$

because utility is increasing in wages ($\nu_w > 0$) and

$$\chi_{\bar{d}}^{\text{JPC}} \in \chi^{\text{SPC}}$$

i.e. $x_{\bar{d}}^{\text{JPC}}$ could have been chosen under SPC, but was not. By revealed preferences, $w_{\bar{d}}^{\text{JPC}} \leq w_m^{\text{SPC}}$.

The relation holds with strict inequality whenever

$$w_i^{\text{SPC}} \notin \chi_{\bar{q}}^{\text{JPC}}.$$

□

Corollary 1 (No wage effect for fathers)**.** *If the father's SPC optimum already satisfies* $||x_f^{SPC} - x_{r,i}^{SPC}|| \leq \bar{d}$*, then* $\mu_f = 0$*,* x_f^{SPC} *remains feasible after the reform and* $w_f^{JPC} = w_f^{SPC}$*.*

Commuting

This follows rather trivially from the assumptions, but can also be shown formally:

Assumption 1 (Radial monotonicity of commuting time)**.** *For every pair of locations* $x_{r,1}, x_{r,2} \in \mathcal{X}$*,*

$$\|x_{r,1} - x_i\| < \|x_{r,2} - x_i\| \implies d(x_{r,1}, x_i) < d(x_{r,2}, x_i).$$

Proposition 2. *Assume radial monotonicity of commuting time and recall that under JPC both parents reside in location* x_f*. Then for every realisation of wage draws*

$$d_m^{\text{JPC}} \geq d_m^{\text{SPC}} = 0, \qquad d_f^{\text{JPC}} = d_f^{\text{SPC}} = 0,$$

with $d_m^{\text{JPC}} > 0$ *whenever the mother's JPC job lies at* $x_m^{\text{JPC}} \neq x_f$ *(strictly positive probability). Hence, under JPC, mothers commute for longer than under SPC. There is no*

change for fathers.

$$d_m^{JPC} > d_m^{SPC}$$

and

$$d_f^{JPC} \approx d_f^{SPC}.$$

Proof. Under SPC the mother can relocate to her chosen job, hence $d_m^{\text{SPC}} = 0$. Under JPC the residence is fixed at x_f, so her accepted job x_m^{JPC} satisfies $0 < \|x_m^{\text{JPC}} - x_r\| \leq \bar{d}$ unless the wage-maximising job is co-located with the home. Assumption 1 then implies $d(x_m^{\text{JPC}}) \geq d(x_r) = 0$; strict inequality holds if $x_m^{\text{JPC}} \neq x_r$. Fathers' residence and job coincide in both regimes, so their commute is unchanged. □

Hours

Proposition 3. *The prediction on hours from moving from SPC to JPC is ambiguous but weakly positive for mothers and weakly negative for fathers.*

Proof. To obtain the derivative of hours with respect to the shadow value of leisure ($\frac{\partial h_i}{\partial \lambda_i}$), first note that h is implicitly a C^1 function of (λ, w) by the implicit function theorem because $F_h = \nu''(wh)w^2 - \eta''(h) < 0$ as ν is concave and η convex. I rewrite the first-order condition with respect to hours (equation FOC 1) as

$$F(h_i, w_i, \lambda_i) = \nu'(w_i h_i - M_i^R)w_i - \eta'(h_i) - \lambda_i = 0 \tag{1.4}$$

By the implicit function theorem:

$$\frac{\partial h_i}{\partial \lambda_i} = -\frac{F_\lambda}{F_h} < 0$$

as $F_\lambda = -1$ and $F_h = \nu''(w_i h_i - M_i^R) \cdot w_i^2 - \eta''(h_i) < 0$ given concavity of ν and convexity of η ($\nu'' < 0, \eta'' > 0$).

How is λ_i affected? Taking the total derivative of the time constraint

$$-dl_i - dh_i - dd(x_i, x_{r,i}) - dq_i^R = 0$$

$$dh_i = -dl_i - dd(x_i, x_{r,i}) - \Delta q_i \,, \tag{1.5}$$

where $\Delta q_i = q_i^{\text{JPC}} - q_i^{\text{SPC}}$. On average, $\Delta q_i < 0$ for mothers and $\Delta q_i > 0$ for fathers.

Totally differentiating the first-order condition with respect to hours (equation FOC 1)

yields

$$\begin{aligned}
&\nu''(w_i(x_i)h_i - M_i^R) \cdot w_i(x_i)^2 \cdot dh + \nu''(w_i(x_i)h_i - M_i^R) \cdot h_i \cdot dw + \nu'(w_i(x_i)h_i - M_i^R) \cdot dw \\
&= \eta''(h_i) \cdot dh + d\lambda_i \\
&\left[\nu''(w_i(x_i)h_i - M_i^R) \cdot w_i(x_i)^2 - \eta''(h_i)\right] dh + \left[\nu''(w_i(x_i)h_i - M_i^R) \cdot h_i + \nu'(w_i(x_i)h_i - M_i^R)\right] dw \\
&= d\lambda_i
\end{aligned} \tag{1.6}$$

Substituting equation 1.5, I obtain

$$d\lambda_i = \underbrace{\left[\nu'' \cdot w^2 - \eta''\right]}_{<0} \cdot (-dl_i - dd(x_i, x_{r,i}) - \Delta q_i) + \left[\nu'' \cdot h_i + \nu'\right] dw \tag{1.7}$$

Wages and commuting: effect of mobility restriction

Let x_i^{SPC} be an unconstrained maximiser of expected utility. The feasible set under JPC is:

$$\chi_i^{\text{JPC}} = \left\{x : ||x_i - x_{r,i}^{\text{SPC}}|| \leq \bar{d}\right\}$$

Given that $\chi_i^{\text{JPC}} \in \chi_i^{\text{SPC}}$:

$$\begin{aligned}
w_i^{\text{JPC}} &\leq w_i^{\text{SPC}} \\
d_i^{\text{JPC}} &\geq d_i^{\text{SPC}}
\end{aligned}$$

Empirically, $\chi_i^{\text{SPC}} \notin \chi_i^{\text{JPC}}$ more likely for mothers than fathers.

Define

$$\begin{aligned}
\Delta w_i &= w_i^{\text{JPC}} - w_i^{\text{SPC}} \leq 0 \\
\Delta d_i &= d_i^{\text{SPC}} - d_i^{\text{JPC}} \geq 0
\end{aligned}$$

Again, using the implicit definition of h from the first-order condition with respect to hours (equation 1.4) and linearising F_h around the SPC point $(\lambda^{\text{SPC}}, w^{\text{SPC}})$ using the implicit function theorem yields the two partial derivatives:

$$\begin{aligned}
\frac{\partial h}{\partial \lambda} &= -\frac{F_\lambda}{F_h} = -\frac{-1}{\underbrace{\nu''(wh - M^R)w^2 - \eta''(h)}_{<0}} < 0 \\
\frac{\partial h}{\partial w} &= -\frac{F_w}{F_h} = -\frac{\nu''(wh - M^R)hw + \nu'(wh - M^R)}{\underbrace{\nu''(wh - M^R)w^2 - \eta''(h)}_{<0}} \geq 0 \quad \text{if } -\nu''(wh - M^R)wh \leq \nu'(hw - M^R)
\end{aligned} \tag{1.8}$$

The second result holds with inequality as long as the assumption about bounded absolute risk aversion is satisfied. This is the case for standard classes of utility such as CRRA. Under log utility, the numerator is zero. I also assume that earnings always

cover the cost of children under any custody arrangement, i.e. that $wh \geq M_i^R$.

First-order Taylor expansion for small values of $\Delta\lambda, \Delta w$:

$$\begin{aligned}\Delta h &\equiv h\Big(\lambda^{\text{JPC}}, w^{\text{JPC}}\Big) - h\Big(\lambda^{\text{SPC}}, w^{\text{SPC}}\Big) \\ &\approx \underbrace{\frac{\partial h}{\partial \lambda}}_{<0} \cdot \Delta\lambda + \underbrace{\frac{\partial h}{\partial w}}_{\geq 0} \cdot \Delta w + \mathcal{O}\Big((\Delta\lambda, \Delta w)^2\Big)\end{aligned} \tag{1.9}$$

Depending on the size of the income effect of the custody regime (ΔM_i), the change in hours will be attenuated relative to a scenario in which the choice of custody regime has no income effect. This assumes that the substitution effect dominates the income effect. If the income effect were larger than the substitution effect, the sign on the change in hours could reverse.

Moving from SPC for mothers to JPC implies

$$\begin{aligned}\Delta\lambda_m &< 0 \qquad \text{for mothers} \\ \Delta\lambda_f &> 0 \qquad \text{for fathers}\end{aligned} \tag{1.10}$$

On the wage side

$$\begin{aligned}\Delta w_m &< 0 \qquad \text{for mothers} \\ \Delta w_f &\approx 0 \qquad \text{for fathers}\end{aligned} \tag{1.11}$$

Finally, putting everything together:

$$\begin{aligned}\Delta h_m &\geq 0 \qquad \text{for mothers} \\ \Delta h_f &\leq 0 \qquad \text{for fathers}\end{aligned} \tag{1.12}$$

□

Appendix E: Co-parenting in other countries

In 2021, about 20.7% of separated parents in Europe had some form of joint physical custody, with the highest rates in Northern European countries, and the lowest in Eastern and Southern Europe (Claessens and Mortelmans, 2025; Hakovirta et al., 2023).[3] Canadian data document a rise in joint physical custody from about 10% in 2006 to 31% in 2018/19 (Department of Justice Canada, 2021). In France, about a quarter of divorce cases involving children resulted in co-parenting arrangements. In the United States, joint physical custody is estimated to have covered 34% of children after divorce between 2010 and 2014. This rise has been attributed to changing norms and policies favouring co-parenting (Meyer et al., 2022). Since 2018, several U.S. states have introduced the legal presumption of splitting parenting time 50/50 unless it is not in the child's best interest, and other U.S. states have debated the adoption of similar legislation in recent years. In short, co-parenting has become a widespread post-divorce parenting arrangement in many Western jurisdictions over the past two decades, though the pace and extent of its adoption vary.

In Asia and the Middle East, joint physical custody is much less prevalent. For example, Japan does not legally recognise joint custody after divorce (one parent receives sole custody by law), and culturally the concept of equal co-parenting post-divorce remains uncommon. However, even in such jurisdictions there is growing debate about fathers' access to their children and reforms to facilitate more contact with non-resident parents.

In most countries other than the Netherlands, the increase in shared physical custody occurred gradually, without a clear policy reform. Legal regimes differ in whether they mandate, presume, or simply allow such arrangements. Coupled with its administrative data records, the reform for a legal preference for co-parenting makes the Netherlands a good country to study its labour market effects.[4]

The idea that co-parenting is the ideal post-divorce parenting arrangement has been codified in international policy, most notably in the Council of Europe's 2015 landmark statement on "Equality and shared parental responsibility: the role of fathers."[5] This resolution urges all 47 member states to "introduce into their laws the principle of shared residence following a separation," with exceptions only for cases of abuse or violence. It emphasises that equal involvement of both parents is beneficial for children's development and gender equality, and calls on governments to facilitate shared residence (for example, through enforcing contact orders, encouraging mediation, and adjusting social benefits to accommodate dual-household families).

[3] In Sweden, the share with equal parenting is as high as 42.5%.
[4] Similar reforms were passed in Australia and Belgium, both in 2006.
[5] Council of Europe Parliamentary Assembly Resolution 2079.

Appendix F: Data construction

First, I identify all divorces between 2006 and 2013. Then, I keep all observations in that time window related to those ever divorcing during it in the dataset. I proceed by creating a monthly relationship panel with partner ID.

For all individuals in the sample, I collect their employment history between 2006-2013, at the monthly level. If individuals hold multiple jobs in the same month, I keep the highest-paying one. If the jobs are paid equally, I choose the one with more hours worked.

In a final step, I combine the relationship and employment panels and keep only individuals aged 18 to 60 divorcing in heterosexual marriages.

Chapter 2

What Makes a Good Job? Evidence from a Large-Scale Survey of Workplace Amenities

This chapter is co-authored with Minji Bang (University of Cambridge) and Hanna Wang (Universitat Autònoma de Barcelona and Barcelona School of Economics).

Abstract

By linking large-scale survey data on workplace amenities, representing 2.4% of the Dutch workforce, to administrative records, we provide new evidence on how amenities shape labor market outcomes. From a rich set of survey questions, we construct ten amenity indices, covering dimensions such as work from home, support from colleagues and supervisors, task complexity and intensity, and learning opportunities. We document how amenities vary with wages and by worker characteristics, sectors, and occupations. Variance decompositions show that firms account for the largest share of variation for most amenities, exceeding the share explained by occupations and individual characteristics. Amenities strongly predict job satisfaction, job search, worker retention, and absenteeism, even when controlling for standard worker and job attributes. Finally, when treating these outcomes as proxies for workers' job value, we consistently find that amenities amplify compensation inequality beyond what is captured by wages alone.

Keywords: Non-wage amenities, Working conditions
JEL Codes: J32, J81

2.1 Introduction

For most workers, the value of a job depends not only on pay but also on a broad set of non-wage amenities. These encompass features shaped by the nature of work, workplace culture, and the broader work environment, including physical and mental demands, relationships with colleagues, and opportunities for learning and development. Such amenities can have important implications for workers' utility, job mobility and welfare inequality.

Empirical research has long been constrained by data limitations. Most existing studies observe only a narrow set of job attributes or infer amenity values indirectly, leaving open questions about which aspects of job quality matter most and how they covary with wages across the labour market. While the theory of compensating differentials predicts that workers accept lower wages in exchange for better working conditions (Rosen, 1986), empirical evidence is mixed, with many studies finding that favourable working conditions are in fact positively correlated with wages (Mas and Pallais, 2017). A central challenge is unobserved heterogeneity: more productive workers may sort into both higher-paying and higher-quality jobs, or better-managed firms may jointly offer higher pay and superior amenities.

This paper overcomes these challenges using large-scale, representative survey data from the Netherlands that provide exceptionally rich measures of workplace conditions across numerous dimensions. By linking these data to administrative records and controlling for a wide range of worker characteristics, including individual fixed effects, we offer new evidence on how workplace amenities relate to wages and other labour market outcomes.

Our dataset comprises over 300,000 observations spanning multiple years and covers roughly 2.4% of the Dutch workforce. The sample is drawn randomly every year from employed individuals and includes more than 11,000 workers observed in multiple waves. In addition to self-reported outcomes such as job satisfaction and job search behaviour, the data are linked to detailed administrative records on workers and firms, offering a comprehensive view of how amenities are distributed across the labour market and how they shape worker mobility and welfare.

Drawing on over 70 survey items from the Dutch National Working Conditions Survey (NEA), we construct ten amenity indices capturing complexity and work intensity, autonomy, quality of colleagues and supervisors, innovation and learning opportunities, emotional demands, irregular hours, health risks, job security, and remote work. To derive these measures, we apply hierarchical clustering based on pairwise correlations and extract latent factors using Item Response Theory (IRT), which groups related items

and filters out measurement noise.[1]

Our analysis reveals substantial heterogeneity in amenity provision across workers and jobs. We document striking patterns across worker characteristics: men enjoy greater autonomy but face more irregular hours and health risk, while women report having more supportive colleagues and higher job security yet experience greater emotional demands and lower learning opportunities. Education creates clear gradients in amenities, with the most educated workers enjoying substantially more opportunities to work from home and facing lower health risks, though they also encounter the highest emotional demands, likely reflecting occupational sorting.

At the firm level, larger and higher-wage-premium employers provide better amenities across most dimensions, though they also impose greater work intensity and emotional demands. Decomposing amenities into firm, occupation, and time effects, as well as individual-level characteristics, we find that firms account for the largest share of amenity variation, followed by workers' occupations.

Higher-wage positions typically offer more favourable amenities, with notable exceptions: such jobs are often characterized by greater emotional demands, heavier workloads, and more irregular hours. These patterns are robust. Wage-amenity correlations remain largely unchanged when examined individually or jointly in regression analysis, and they persist after controlling for a range of individual and job characteristics.

We further consider the influence of work amenities on workers' utility and mobility. We examine a broad set of outcomes capturing both subjective assessments and revealed preferences: workers' self-reported job satisfaction, willingness to stay in their current job, job search activity, absenteeism, and the probability of leaving the current job and transitioning to another job (derived using administrative employment histories). These outcomes are important not only for worker welfare but also for firm profitability.[2]

Our results show that amenities have broadly similar effects across all outcomes. Beyond wages, job security, emotional demands, supervisory quality, and opportunities for innovation and learning have the largest impact across all outcomes. For example, a one-standard-deviation improvement in learning and innovation opportunities reduces job search by 2.1 percentage points – an effect comparable in magnitude to a one-standard-deviation wage increase. Most results are robust to controlling for individual

[1] We employ graded-response IRT with marginal maximum likelihood/EM estimation (Samejima, 1969; Bock and Aitkin, 1981).

[2] These outcomes are important for both worker welfare and firm profitability. Job satisfaction strongly predicts worker retention (Clark, 2001; Böckerman and Ilmakunnas, 2009), which matters because replacing workers is costly. Absenteeism lowers team productivity (Adhvaryu et al., 2024), and hiring costs have been estimated at 10-17 weeks of wages in Switzerland (Blatter et al., 2012) to a full year in Denmark (Bertheau et al., 2022), with these costs increasing for more specialized skills (Jäger and Heining, 2022).

and job characteristics.

Finally, we show that amenities substantially widen compensation inequality beyond wage differences. Interpreting the above outcomes as indicators of job value, we compare the top and bottom of predicted values based on wages alone versus wages and amenities jointly. Across all outcomes, amenities consistently amplify compensation dispersion, indicating that workplace amenities reinforce rather than offset inequality arising from wages alone. This is most pronounced for job satisfaction and the other self-reported measures. Examining the overall variance, the standard deviation of predicted values based on wages alone is only 10-45% of the standard deviation when amenities are included, further confirming that amenities substantially amplify compensation dispersion.

Our analysis advances research on the measurement of workplace amenities and their roles in job mobility and inequality. Our main contribution lies in directly analysing multi-dimensional workplace conditions, moving beyond revealed-preference approaches that infer composite amenity values. Hall and Mueller (2018) estimate amenity valuations from data on job applications, reservation wages and accepted offers, revealing substantial heterogeneity in non-wage preferences. A recent line of work develops and estimates dynamic search models to recover job values from mobility patterns (Sorkin, 2018; Lentz et al., 2023; Morchio and Moser, Forthcoming; Lamadon et al., 2024). The key advantage of our analysis is that it allows us to open the black box and disentangle the contributions of specific amenities.

Recent work has made important progress in analysing observed amenities but often features limited amenity measures. For example, Wiswall and Zafar (2018) show that pre-labour market stated preferences predict actual job attributes such as job security, gender ratio and bonuses. Mas and Pallais (2017) embed discrete choice experiments in call center applications to measure valuations of schedule flexibility and remote work. Using Brazilian administrative data and collective bargaining agreements, Corradini et al. (2025) demonstrate that unions can improve workplace amenities without reducing wages or employment. Koenig and Annelli (2025) develop a bunching method to measure willingness to pay for amenities and apply it to Covid infection risks.

A handful of papers study a broader set of amenities but rely on selected or smaller samples. Bonhomme and Jolivet (2009) estimate a job search model using reported satisfaction with various job amenities from the European Community Household Panel. Using stated-preference experiments in the American Working Conditions Survey, Maestas et al. (2023) find that a range of workplace conditions often exacerbate rather than reduce wage inequality. Analysing online employee reviews, Sockin (2022) shows that higher-paying firms offer better amenities, amplifying compensation dispersion by 50-

65%.

The most closely related paper to ours is Humlum et al. (2025a), who embed survey responses on a number of amenities and reservation wages from a large sample of job movers in a search model. Our approach complements and extends this literature by leveraging representative survey data for both movers and stayers combined with administrative records. Our findings align with Sockin (2022) and Maestas et al. (2023), who similarly find that higher-paying positions offer better amenities, exacerbating measured inequality. In contrast, Humlum et al. (2025a) find the opposite. Studies focusing on gender inequality also find that women's lower wages are partially offset by better workplace amenities (Le Barbanchon et al., 2021; Morchio and Moser, Forthcoming).

The remainder of this paper is structured as follows: Section 2.2 introduces our data and methodology for constructing amenity measures. Section 2.3 examines how workplace amenities are distributed across workers and firms. Section 2.3.2 investigates the relationship between wages and amenities. Section 2.4 explores how amenities affect labour market outcomes and inequality, and Section 2.5 concludes.

2.2 Data and Construction of Amenity Measures

We use various data sources that can be linked through unique person IDs spanning the years 2007 to 2021. The main data source is the Dutch National Working Conditions Survey (NEA), a large cross-sectional survey of employees conducted annually since 2005.[3] We link the NEA to administrative records covering the universe of Dutch workers. The administrative data provide comprehensive employment histories, including monthly hours and wages for all job spells, as well as worker demographics (age, gender, education, household composition) and a range of firm-level characteristics such as sector, size, and productivity. This merged dataset allows us to observe both self-reported workplace amenities and actual labour market outcomes, including job-to-job transitions and transitions to non-employment conditioned on both worker and firm characteristics.

Survey Data on Workplace Amenities The NEA provides exceptionally rich measures of workplace conditions.[4] Between 20,000 and 60,000 employed individuals respond to the survey annually, yielding over 300,000 observations across our sample

[3] We begin our sample in 2007, as many of the amenity-related questions used in our analysis were not included in earlier waves of the survey.

[4] Documentation on sampling, fieldwork, and questionnaire design is provided by CBS/TNO (Statistics Netherlands (CBS), 2024).

period.[5] The survey includes a comprehensive battery of questions covering multiple dimensions of job quality, with responses recorded on 3 to 7 point Likert scales.[6] These questions capture various aspects of the work environment, including the nature of tasks (e.g. complexity), workplace relationships (quality of colleagues and supervisors), work arrangements (remote work options).

We restrict the sample to respondents aged 22-64 who are employed for at least 10 hours per week and earn at least the minimum wage.[7] We further limit the sample to workers employed in firms surveyed at least twice, which allows us to control for firm fixed effects. This restriction excludes some smaller firms and reduces the sample by approximately 47,000 observations, leaving a final sample of 308,415. Most respondents are surveyed only once, but over 11,000 workers appear in multiple survey waves. This panel dimension allows us to control for individual time-invariant characteristics that may drive sorting into jobs with particular amenity bundles.

The NEA sample is balanced by gender, with women making up 50.5% of the respondents (Table B2.1). Slightly more than a quarter (26%) of respondents are aged 22–34. The remaining respondents are split roughly equally between the 35–49 and 50–64 age groups.[8] Close to half of the sample is highly educated (university or highest vocational training). 39% have a medium level of educational attainment (medium-level vocational training), and 12.6% are in the lowest education group (only schooling or lowest vocational training). Half of the respondents are married or cohabiting with at least one child. An additional 29% are married or cohabiting without children.

The outcomes that we consider as different dimensions of job value are described in Table B2.2. The first five outcomes are based on the NEA survey, and the last two are calculated from administrative employment histories. Job satisfaction is measured on a 1-5 Likert scale and averages 3.9 (where 4 means satisfied).[9] About half of the respondents report no absences in the prior year; across all workers, the average attendance rate is 95.5% of scheduled work time. Three quarters of respondents are not looking for another job at the time of the survey, and 69% state they would like to remain in their current jobs for the next five years. When considering actual transitions, 93% of

[5]The differences in sample size are due to a gradual expansion of the survey, not widely varying response rates across years.

[6]Most responses are in terms of frequencies ("Always", "Often", "Sometimes" etc.). Examples of the selected raw variables and answer options can be found in Table A2.1.

[7]In 2014, the relevant age group was expanded from up to 64 to up to 75 years of age. To keep the early and later years comparable, we decided to exclude the workers above age 64. They make up fewer than 5% of the post-2014 sample.

[8]Younger workers are oversampled by design due to lower response rates; provided survey weights correct any remaining imbalances.

[9]The job satisfaction measure ("To what extent are you, all things considered, satisfied with your job?") should capture utility from the job as a whole, including both wage and amenities. This should also be clear to respondents as another question specifically refers to satisfaction with the workplace.

employees remain in their current firm for at least 12 months after the survey. Of those that do not stay, 4.9% transition to another firm, and the remaining 1.7% transition into non-employment.

Construction of Amenity Measures From the more than 600 variables broadly related to working conditions in the NEA, we select 70 amenity-related questions based on data availability across survey years and respondents. We exclude variables that reflect individual-specific circumstances or perceived match quality and promotion prospects, focusing instead on amenities that characterize the job and workplace itself. Table A2.2 provides a complete list of the included questions and their response scales.

We first cluster individual questions into amenity groups based on pairwise rank-based correlation.[10] The key idea is that items answered similarly across individuals likely capture aspects of the same underlying workplace characteristic. We then construct latent amenity factors within each cluster using item response theory (IRT). Factors built from multiple subitems are inherently more robust to measurement error than relying on a single survey response for each amenity. This approach allows us to include variables not available in every survey year, as IRT can estimate factor scores using all available information for each respondent without requiring complete responses.[11] The ten final amenity measures are work complexity/intensity, autonomy, quality of colleagues, quality of supervisors, innovation and learning opportunities, emotional demands, irregular hours, health risks, job security, and remote work.

Finally, we validate the internal consistency of the constructed factors using McDonald's ω and the proportion of variance explained within each cluster (Table A2.3).

A more in-depth description of how we construct our amenity measures is detailed in Appendix D.

2.3 Workplace Amenities Across Workers, Jobs, and Firms

In this section, we examine how the different amenity dimensions vary systematically with worker and job characteristics. Documenting these patterns serves two purposes. First, it provides insights into the extent of sorting in the labour market and inequality in access to amenities. Second, it offers a validation check on our measurement strategy:

[10] We define pairwise distance as one minus the correlation coefficient in the clustering analysis.

[11] For robustness, we alternatively use PCA, vary the number of clusters, and make different selections of raw variables. The resulting amenity measures remain very similar, and, importantly, result in similar coefficient estimates when regressed on job satisfaction and other outcomes.

if the constructed amenity indices align with well-established correlations (e.g., by occupation, education, or industry), this increases confidence that our measures capture meaningful aspects of job quality rather than noise.[12]

Worker Characteristics Our results show substantial heterogeneity in access to workplace amenities across worker characteristics. Figure 2.1a illustrates gender differences. Men report greater autonomy and a higher ability to work from home, but also face higher health risks and more irregular hours. Women, in contrast, enjoy greater job security and more supportive colleagues, yet experience more emotionally demanding work environments. These gender gaps are sizable, often exceeding 0.1 standard deviations, while differences in supervisor quality or opportunities for learning and innovation are comparatively small.[13]

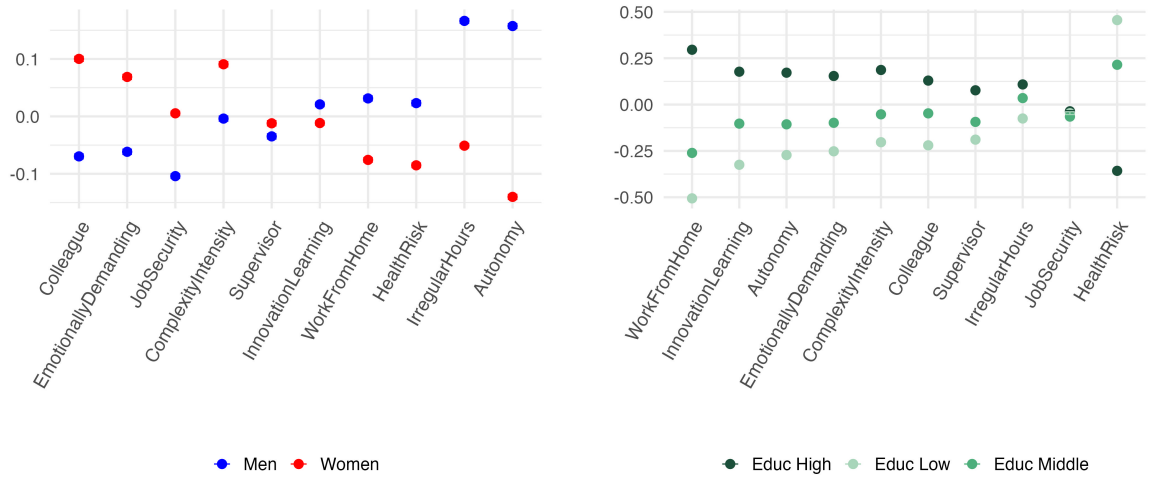

(a) Amenities by gender

(b) Amenities by education

Figure 2.1: Average amenity levels by gender and education.

Note. (a) Average amenity level (standardised) by gender of respondent. Estimates in blue refer to men and those in red to women. Amenity measures are standardised with mean zero and standard deviation one. (b) Average amenity level (standardised) by educational attainment of respondent. High education is any university education or the highest vocational training, middle refers to the intermediate and lower tracks of vocational training, and low makes up educational attainment below professional training. Amenity measures are standardised with mean zero and standard deviation one.

We also find pronounced gradients in workplace amenities across education levels (Figure 2.1b). With the exception of job security, nearly all amenities display a clear and monotonic relationship with education. The differences between the highest and lowest

[12]We further compare our worker-reported measures with manager-reported amenities from a companion survey, which can be linked at the firm level for a subset of our sample. The results (available upon request) show that worker and manager assessments are closely aligned, providing additional validation of our measures.

[13]For comparison, the gender wage gap in our sample is 0.39 standard deviations.

education groups range from roughly 0.3 to 1.0 standard deviations—substantially larger than the corresponding gender gaps. Respondents with higher education levels mostly report better working conditions: they are more likely to be able to work from home, face lower health risks, and enjoy more opportunities to innovate and greater autonomy. At the same time, they are also more exposed to emotionally demanding and cognitively intense tasks, suggesting that higher-skilled jobs combine both greater rewards and higher pressures.

Grouping workers by individual productivity, measured via fixed effects in an AKM framework, yields patterns similar to those by education. The most productive workers are most likely to work from home, while the least productive face the highest health risks (Figure C2.1). The systematic patterns in amenities across worker characteristics reflect occupational and firm sorting. We next examine amenity variation along these job-level dimensions directly.

Occupations, Sectors and Firms Job tasks are likely an important determinant of workplace amenities. To illustrate this, Figure C2.2 compares amenity levels across the ten largest occupational groups. The patterns align with common expectations: managers and professionals rank highest in opportunities for innovation and learning, whereas plant and machine operators and assemblers report the lowest levels. Workplace amenities also differ markedly across sectors, as shown in Figure C2.3 for the ten largest sectors. Health risks are lowest in retail and trade but highest in mining. The ability to work from home shows an almost inverse pattern—retail and trade, followed by transport, exhibit the highest prevalence of remote work, while construction, utilities, and mining rank at the bottom.

Beyond sector and occupation, amenities additionally vary along firm characteristics such as size and productivity. Figure 2.2 shows the distribution of amenities by firm size and illustrates that large firms offer many favourable amenities such as work from home, autonomy, low health risks and innovation and learning. However, jobs in large firms are also more complex, emotionally demanding and have lower job security and higher incidence of irregular hours. Interestingly, the quality of supervisors and colleagues shows a U-shaped relationship in firm size: It is high both for small and large firms, and lower for those in between.

We also group firms by their AKM productivity and show average amenity levels for each in Figure C2.4. The correlations with firm productivity align closely with the ones for worker productivity. Firms with the highest productivity are most likely to facilitate work from home, and expose their employees to the lowest health risks. Moreover, the quality of colleagues and supervisors rises steadily with firm productivity, in contrast

Figure 2.2: Amenities by firm size

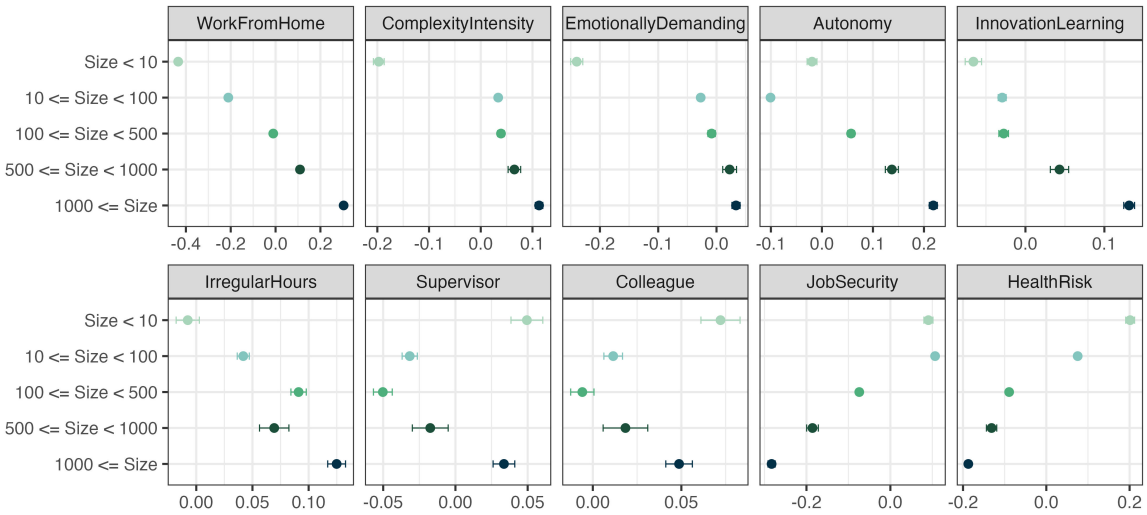

Note. Average amenity level (standardised) by firm size. Amenity measures are standardised with mean zero and standard deviation one.

to the non-linear relationship observed with firm size.

Besides worker and job characteristics, we also document the time trends for four amenities in Figure C2.5. The steep drop in job security until 2014 and its subsequent recovery track the employment rate and reflect the Great Recession. The ability to work from home sees a substantial jump upwards in 2021 due Covid-19. Workplace health risks decreased, while reported emotional demands rose for much of the time period we observe.

2.3.1 Explaining Variation in Amenities

Using a variance–covariance decomposition, we quantify the share of total variation in amenities attributable to key factors. Specifically, we decompose amenities into firm (j), occupation (o), time (year) (t) and individual characteristics (X_{it}). Individual characteristics include age, education, gender, and household composition. We estimate the following:

$$\begin{aligned} A_{iojt} &= \beta X_{it} + Firm_j + Occ_o + Time_t + Resid_{iojt} \\ \underbrace{Var(A)}_{\text{variance of amenity}} &= \underbrace{Cov(\beta X, A)}_{individual} + \underbrace{Cov(Firm, A)}_{firm} + \underbrace{Cov(Occ, A)}_{occupation} \\ &\quad + \underbrace{Cov(Time, A)}_{time} + \underbrace{Cov(Resid, A)}_{residual} \end{aligned} \tag{2.1}$$

Figure 2.3 presents the results of the variance decomposition described in Equation (2.1). Overall, between 18% and 53% of the variation in amenities can be accounted for by observable factors. Firms emerge as the dominant determinant, accounting for more than half of the explained variation for most amenities and around 15% of the total variation. This highlights the role that firm- specific factors including management practices, organisational culture, and human resource policies play in determining workers' experiences on the job.

Occupations explain smaller but still significant shares of variation. They are of comparable or greater importance than firms for certain dimensions, particularly health risks, work-from-home opportunities, and autonomy, which are closely tied to the nature of job tasks rather than to firm-level practices. Individual covariates and time effects account for only a small fraction of overall variation, once firm and occupational factors are controlled for. This suggests that neither discrimination based on observable worker characteristics nor systematic subjective reporting biases contribute substantially to explaining heterogeneity in amenities. Together, these results emphasize that the workplace itself, more than the worker, is the primary source of differences in job amenities.

Figure 2.3: Variance Decomposition Results

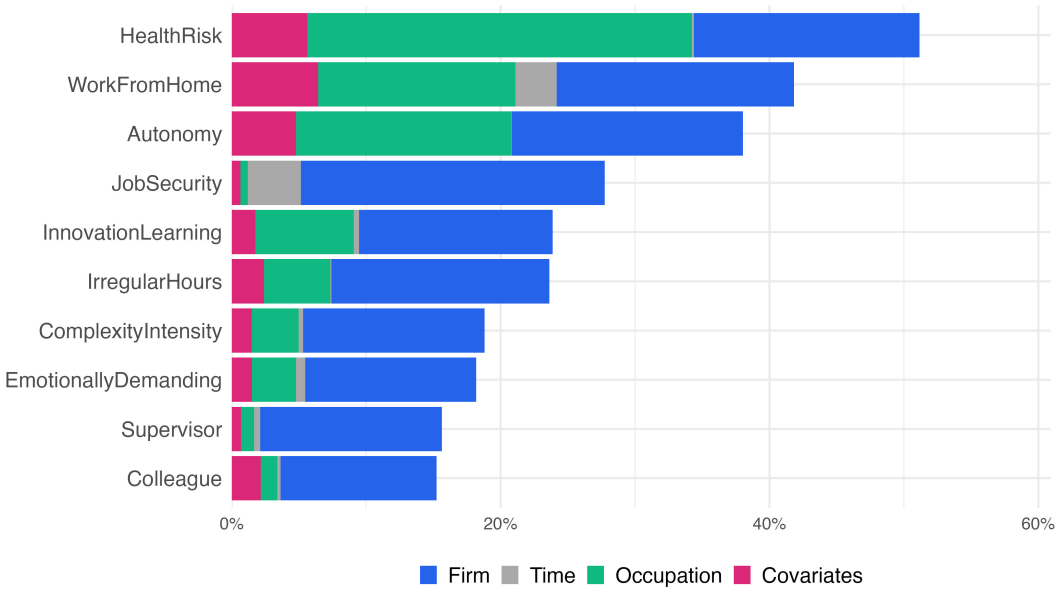

Note. Decomposition of amenity variation into firm, time (year), occupation, and individual covariates. The x-axis measures the total share of explained variation, and colors indicate the share explained by each of the variables.

We further show that two-digit sectors account for substantially less variation in amenities than firms, and also less than occupations or individual characteristics (Figure C2.6). In the main decomposition, we use occupation at the 2-digit ISCO code, as earlier survey

waves lack finer occupational detail. For the subsample from 2010 onward, we replicate the analysis using three- and four-digit ISCO codes to obtain a more granular occupational breakdown. The share of variation explained by occupations increases slightly, but the central result remains unchanged: for most amenities, firm-level factors are the dominant source of heterogeneity.

2.3.2 Wages and Amenities

Examining pairwise correlations between wages and individual amenities (shown in Figure 2.4), we find that higher wage are mostly positively correlated with better workplace amenities. For example, higher-wage respondents also report a greater ease of working from home, more autonomy, better supervisors, lower health risk. Two exceptions are that they are also faced with more emotionally demanding and more complex and intense jobs. The correlation of wage with colleagues and job security is close to zero. This pattern mostly aligns with amenity differences by education level from Figure 2.1b.

These findings cast doubt on the existence of systematic compensating wage differentials. If firms offering poor working conditions systematically compensated workers with higher pay, we would expect to see negative correlations between wages and desirable amenities. Instead, the positive associations we observe suggest that high-wage jobs tend to bundle together both pecuniary and non-pecuniary advantages.

The results from our regressions presented in Table B2.3, where standardised wages are regressed on standardised amenities, confirm these raw correlations. The coefficients in column (1), which controls only for year fixed effects, are generally positive and significant for desirable amenities like the ability to work from home (0.26 s.d.) and autonomy (0.18 s.d.), and negative for undesirable ones like health risk (-0.17 s.d.). These magnitudes are a strong reflection of the pairwise correlations, suggesting the positive association is not merely driven by time trends.

Even after introducing the full set of controls – including year, firm, occupation, and demographic controls including individual fixed effects – the persistence of these positive coefficients is notable. In Column (5) of Table B2.3, the coefficients for nearly all amenities persist with the same sign as in the raw correlation, albeit with significantly attenuated magnitudes, especially for high-variance amenities like working from home (0.024 s.d.). This persistence, even after isolating within-person, within-firm, and within-occupation variation, suggests a robust connection between amenity quality and wages.

Beyond wage-amenity relationships, Figure 2.4 also reveals intuitive correlations among amenities themselves. For instance, working from home is associated with lower health

risks, higher work intensity correlates with greater emotional demands, and better supervisors tend to work alongside better colleagues. These patterns suggest that certain amenities naturally cluster together, reflecting either technological complementarities (e.g., remote work reducing physical hazards) or workplace cultures that bundle multiple job quality dimensions.

Figure 2.4: Correlation between Wage and Amenity Measures

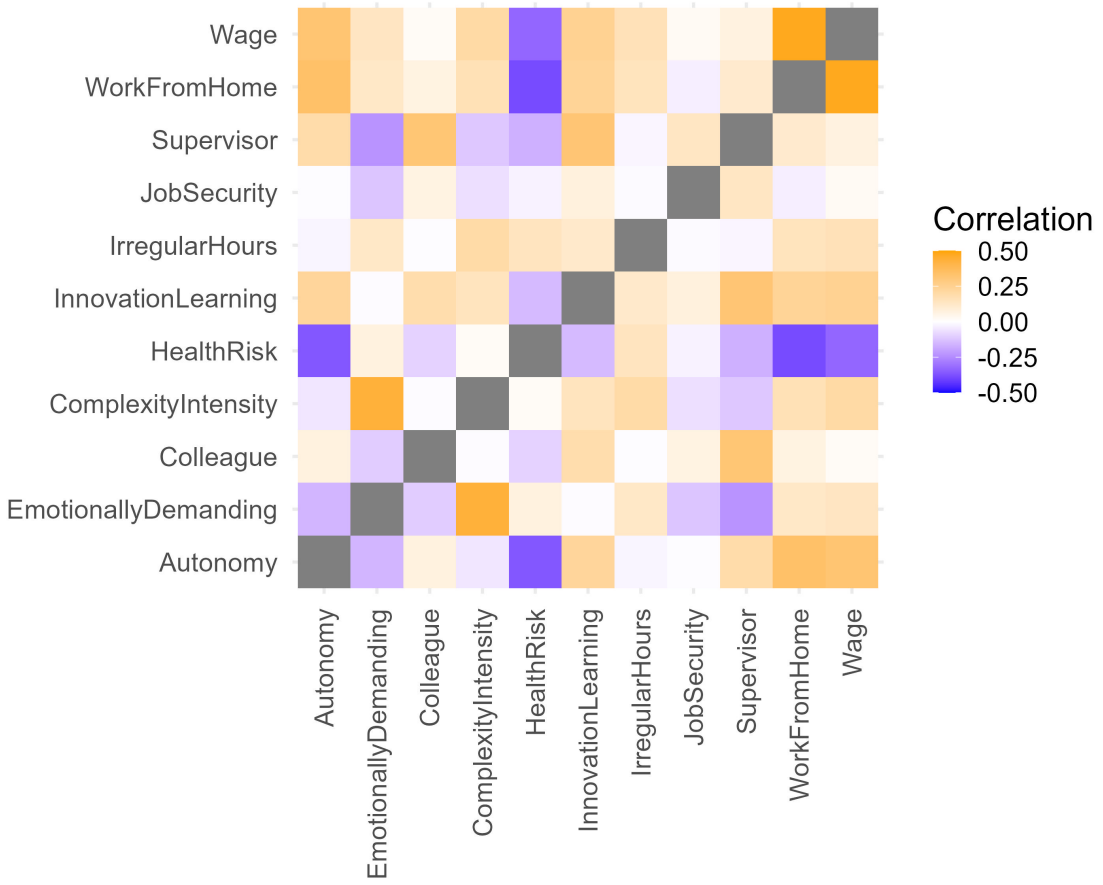

Note. Correlation matrix heatmap showing pairwise correlations between wages and all amenity clusters.

2.4 Labour Market Outcomes and Inequality

2.4.1 Job Value, Mobility and Absenteeism

In this section, we examine how amenities and wages relate to seven important labour market outcomes. These outcomes are relevant both for workers, through their impact on job value and mobility, and for firms, as they affect absenteeism and retention.

We use workers' self-reported job satisfaction as a key proxy for job value. A large

empirical literature shows that job satisfaction predicts job search behaviour and actual quits (Freeman, 1978; Clark et al., 1998; Clark, 2001; Lévy-Garboua et al., 2007; Böckerman and Ilmakunnas, 2009; Card et al., 2012). In addition to job satisfaction, we consider six further outcomes: (i) stated intention to remain with the current employer for the next five years (if it were up to the worker), (ii) current on-the-job search, (iii) whether the worker had any absences in the past year, (iv) absenteeism rate (share of scheduled hours missed), (v) staying with the same employer for the next year, and (vi) making a direct job-to-job transition. The first five outcomes are self-reported, while the last two are computed from administrative employment histories.

Flipping the sign of absenteeism and job search so that higher values indicate greater job value, Figure 2.5 shows that the seven outcomes are positively correlated with each other, but only moderately so overall. The two realized mobility measures are extremely highly correlated; consequently, we focus on one of them, retention in the next 12 months, when presenting results. Similarly, the two absenteeism variables are highly correlated, so we focus on the binary indicator ("Not Absent 12 Mo") in our main text. Among the self-reported measures, satisfaction, job search, and intention to stay are fairly strongly correlated with one another, whereas absenteeism shows weaker associations with the other outcomes.

Although no single measure perfectly captures job value, each provides a different and valuable perspective on worker welfare and labour market behaviour. Self-reported job satisfaction directly reflects workers' overall assessment of their jobs and strongly predicts subsequent mobility (Freeman, 1978; Clark, 2001). Stated intentions to stay and job search behaviour provide forward-looking indicators of workers' attachment to their current positions (Böckerman and Ilmakunnas, 2009). Absenteeism reflects both job quality and working conditions, with higher absence rates linked to poor workplace amenities and lower worker wellbeing (Adhvaryu et al., 2024). Finally, actual job mobility represents revealed preferences that account for the full set of constraints workers face in the labour market.

These outcomes matter not only for worker welfare but also for firm profitability. Absenteeism directly lowers team productivity (Adhvaryu et al., 2024), and replacing workers is costly, with estimates ranging from 10-17 weeks of wages in Switzerland (Blatter et al., 2012) to a full year in Denmark (Bertheau et al., 2022). Moreover, worker satisfaction has been causally linked to higher productivity in both experimental (Oswald et al., 2015) and matched survey-register data (Böckerman and Ilmakunnas, 2012), with broader firm-level benefits including improved stock performance (Edmans, 2011). Meta-analytic evidence across hundreds of studies confirms the positive relationship between employee wellbeing and firm performance (Krekel et al., 2019). Together, these

multiple measures provide a comprehensive picture of how amenities shape job value and labour market dynamics.

Figure 2.5: Correlation between outcome variables

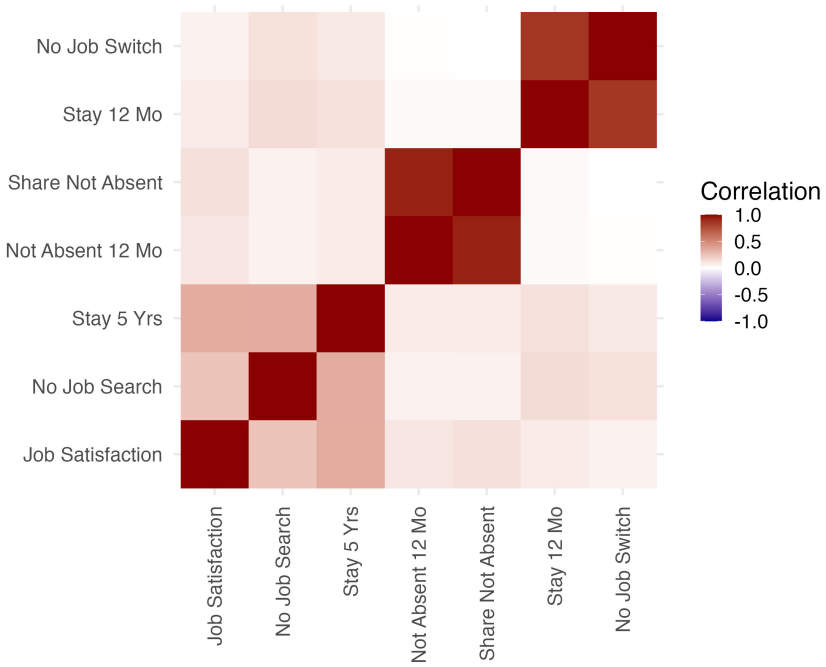

Note. Left-to-right: job satisfaction, no job search in the last 12 months, intention to stay for five more years, no absenteeism in the past 12 months, share of working hours not absent in the past 12 months, stayed for the following 12 months, no job-to-job transition in the following 12 months. The first five outcomes are self-reported; the last two are computed based on employment histories.

To identify the influence of wages and amenities on outcomes, we specify a framework in which the value that workers derive from a job is additive in wages and amenities.[14] We regress our outcomes of interest on amenities and wages:

$$Y_{ijt} = \beta_0 + \beta_w wage_{ijt} + \beta_1 a_{1ijt} + ... + \beta_{10} a_{10ijt} + \gamma_x X_{it} + \gamma_z Z_{jt} + \gamma_t + \varepsilon_{ijt} \qquad (2.2)$$

where $a_{1ijt}, ..., a_{10ijt}$ refer to the ten standardised amenity measures, X_{it} are worker-level characteristics such as age, gender, education, and family status, and Z_{jt} are job-level characteristics, e.g., occupation, sector, size. For comparability across regressors, wages are also scaled according to standard deviations. Standard errors are clustered at the firm level.

Our baseline specification includes workers' demographic characteristics and year fixed

[14]This linear additive specification could be relaxed to allow for interactions between amenities or nonlinear relationships. We adopt the simpler specification because it facilitates interpretation of individual amenity effects and, as discussed in Section 2.4.2, enables a cleaner decomposition of compensation inequality into wage and amenity components.

effects. For job satisfaction, we standardise the outcome (originally measured on a 1-5 Likert scale), while for binary outcomes we estimate linear probability models. This setup allows us to compare outcomes across similar workers net of time trends while preserving the variation in amenities that stems from occupation and firm choices, which Section 2.3 showed to be key sources of heterogeneity. To demonstrate robustness, we estimate specifications that add occupation fixed effects, firm fixed effects, and individual fixed effects (for workers observed multiple times). These progressively isolate within-occupation, within-firm, and within-worker variation, showing that our main findings remain largely unchanged.

The coefficient plots in Figure 2.6 show that amenities and wages significantly affect all outcomes and that the results exhibit strong similarities. Job satisfaction and intention to stay show nearly identical patterns, perhaps because these two outcomes most directly target perceived job value. Job search and actual retention show similar patterns as well. The biggest difference across outcomes lies in job security: its coefficient is larger for job search (6.4 p.p.) and realized retention (2.2 p.p.) than for stated intention to stay (0.6 p.p.). This likely reflects the fact that job security captures involuntary separation risk. Workers may be laid off or search preemptively when their positions are at risk. The smaller effect on intention to stay suggests that job security matters less for purely voluntary job switches.

A few amenities stand out for their pronounced effects on job value. Having a good supervisor strongly increases almost all outcomes. A one-standard-deviation increase in supervisory quality reduces job search by 4.3 percentage points and raises the probability of realized retention by 0.6 percentage points. In contrast, emotionally demanding work has large adverse effects on satisfaction, absenteeism, job search, and intention to stay. The effect is slightly more muted for realized retention. A one-standard-deviation increase in emotionally demanding work raises job search by 4.1 percentage points and absenteeism by 8.3 percentage points. Opportunities for learning and training consistently have a substantial positive impact across both subjective and behavioural measures: a one-standard-deviation increase raises job satisfaction by 0.14 standard deviations and increases the probability of actually staying with the employer for the next 12 months by 0.6 percentage points.

Wage coefficients are positive across all outcomes but are often smaller in magnitude than those for supervisor quality, emotional demands, and learning opportunities. The main exception is realized retention, where the wage effect (1.7 percentage points per standard deviation) exceeds that of all amenities except job security, suggesting that financial considerations play a particularly important role in actual job mobility de-

Figure 2.6: Regressions Results of amenities on job value and mobility

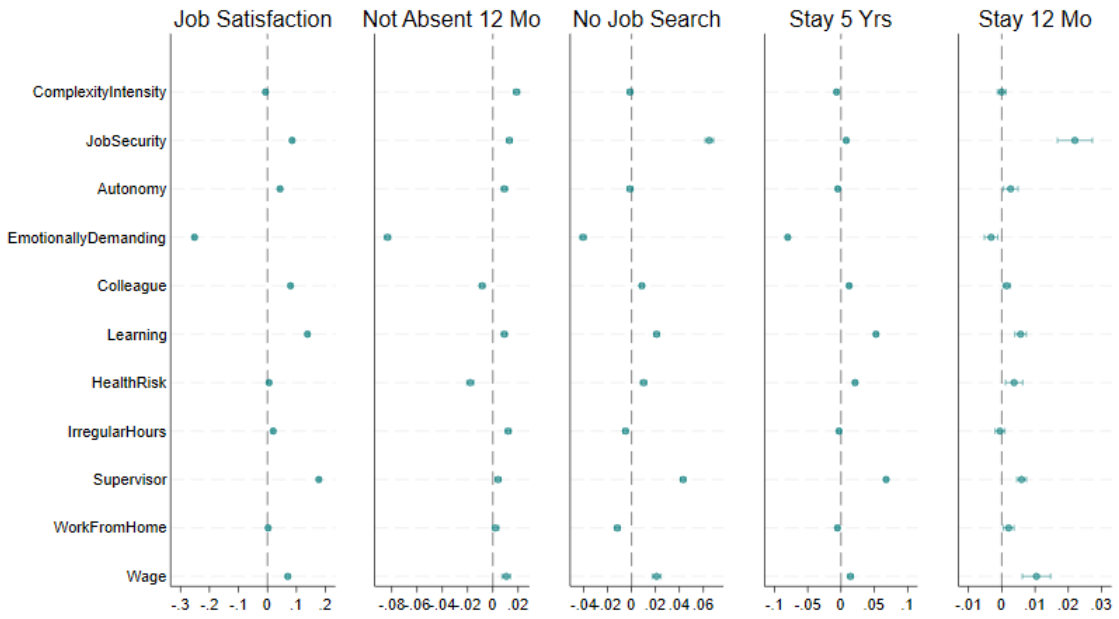

Note. All amenities, including wages, and the job satisfaction measure are standardised with mean zero and standard deviation one. Absenteeism in the past 12 months, intention to stay for five more years, no job search, and staying for the following 12 months are binary outcomes estimated using linear probability models. All regression specifications control for demographic characteristics (age, gender, education, family status) and year.

cisions.[15] One possible explanation is that wages are immediately observable and salient to workers when evaluating a new job offer, making them more influential in realized mobility decisions.

Finally, our results are remarkably robust to rich controls. The estimated effects of amenities and wages on outcomes remain largely unchanged after accounting for year fixed effects, demographic characteristics, firm fixed effects, occupation fixed effects, and individual fixed effects (Figure 2.7).[16] For example, the job security coefficient on job satisfaction ranges only from 0.077 to 0.087 across all five specifications. Results are also consistent when using firm-level means rather than individual responses: job security, emotional demands, and supervisor quality remain the strongest predictors of firm-level outcomes, and higher firm wage fixed effects predict better retention (Figure C2.7, Table B2.11).[17]

[15]Similar results hold for job-to-job transitions (Appendix Table B2.7).

[16]Full regression results for all outcomes and specifications are in Appendix Tables B2.4-B2.10.

[17]Firm-level amenities average responses from firms with at least three observations, netting out demographic characteristics and time effects.

Figure 2.7: Comparison between alternative specifications

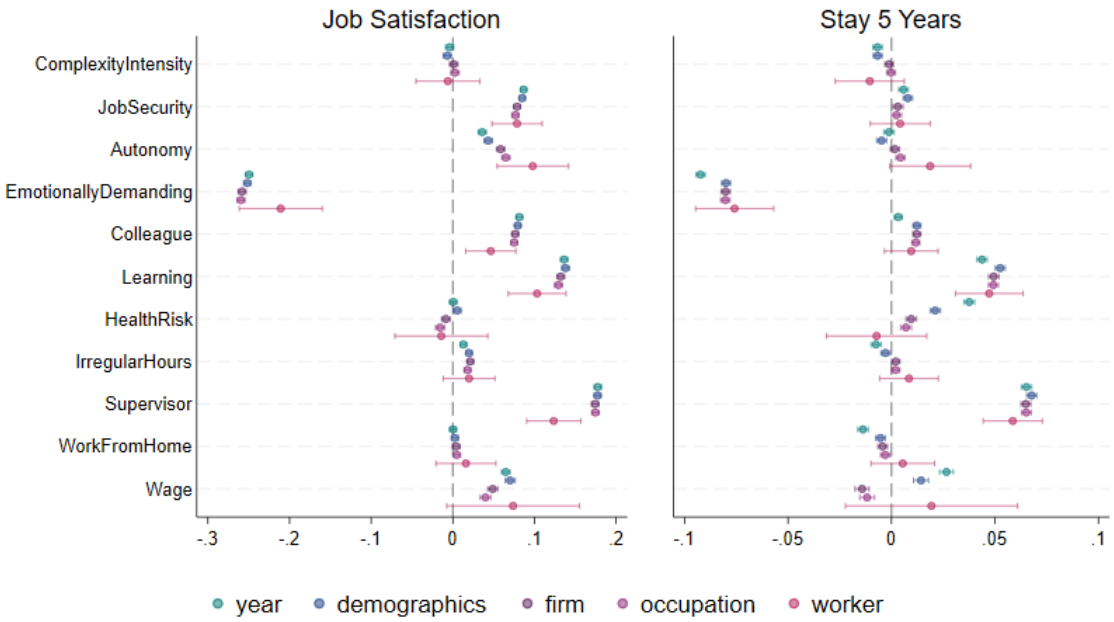

Note. All amenities, including wages, and the job satisfaction measure are standardised with mean zero and standard deviation one. Intention to stay for five more years is a binary outcome.

2.4.2 Compensation Inequality

Non-wage amenities are a potentially important source of compensation inequality among workers. While the theory of compensating differentials predicts that wage adjustments should offset differences in working conditions, search frictions and heterogeneity across workers and firms may instead generate positive correlations between wage and amenity values, thereby amplifying inequality. As seen in Figure 2.4, favourable amenities tend to be positively associated with wages, with some exceptions. The net effect on inequality depends on how much each amenity contributes to overall job value relative to wages.

To assess whether amenities increase compensation inequality beyond what is captured by wages alone, we examine predicted values for five outcomes that serve as proxies for job value. Specifically, we estimate job value based on (i) wages only and (ii) wages and amenities jointly, and then compare the gap between the top and bottom of the distribution for both specifications. We measure this gap as the difference between the 80th and 20th percentiles (namely, "80-20" gap), which provides an intuitive measure of inequality by capturing how predicted job value differs between high and low valued positions. In addition, we compute the correlation between predicted values based solely on wages and those based solely on amenities. Unlike the 80–20 gap, this correlation measure captures co-movement across the entire distribution and is unit-free.

The predicted values ($\hat{Y}_W$, $\hat{Y}_A$, $\hat{Y}_{W+A}$) are derived from Equation 2.2 and calculated as follows:

$$\begin{aligned} \hat{Y}_{W,ijt} &= \hat{\beta}_0 + \hat{\beta}_w wage_{ijt} \\ \hat{Y}_{A+W,ijt} &= \hat{\beta}_0 + \hat{\beta}_w wage_{ijt} + \hat{\beta}_1 a_{1ijt} + ... + \hat{\beta}_{10} a_{10ijt} \end{aligned} \tag{2.3}$$

The results, displayed in Table 2.1, show that the 80-20 gap is smaller when using wage-only predictions (first row) than when amenities are also considered (second row). This holds across all outcomes. For example, in the case of job satisfaction, the 80–20 gap increases from 0.104 to 0.760 standard deviations when amenities are included. Similarly, when predicting retention, the gap between the top and bottom quintiles is 1.5 percentage points based on wages alone, but 4.9 percentage points when amenities are incorporated. Examining the overall variance, the standard deviation of predicted values based on wages alone is only 10-34% of the standard deviation when amenities are included, further confirming that amenities substantially amplify compensation dispersion. Wages play the largest role relative to amenities for the dispersion in actual retention. For the other outcomes, the standard deviation in wage predictions make up less than 20% of the standard deviation in predictions based on wage and amenities.

Table 2.1: Compensation inequality

	Job Sat	Not Absent	No Search	Stay 5 Yrs	Stay 12Mo
$\Delta_{80,20}\hat{Y}_W$	0.104	0.016	0.031	0.021	0.015
$\Delta_{80,20}\hat{Y}_{W+A}$	0.760	0.145	0.185	0.229	0.049
$SD(\hat{Y}_W)$	0.070	0.011	0.021	0.014	0.010
$SD(\hat{Y}_{W+A})$	0.472	0.089	0.111	0.143	0.029

Note. Differences between the top and bottom quintiles in predicted job value based on wages only (first row) or wages and amenities jointly (second row). Variances of the predicted job values based on wages only (third row) or wages and amenities jointly (fourth row). Outcomes for columns left to right are: job satisfaction, not absent in the last 12 months, no search for new job in the last 12 months, intention to stay in job for next five years, and actually staying for next 12 months.

Our findings for the contribution of amenities in total compensation dispersion are somewhat larger but roughly comparable to the literature. Hall and Mueller (2018) estimate that the standard deviation of wages is around 63% of that of wage and non-wage component jointly. Taber and Vejlin (2020) finds that the overall log wage variance accounts for about 45% of the variation in utility. Lastly, Sockin (2022) shows that considering amenities leads to a 6 percent increase in the variance of job satisfaction at the firm level. On the contrary, Humlum et al. (2025b) find that amenities lower compensation inequality, such that variance in wages overstates the variance of utility

by 29%.

2.5 Conclusion

Workplace amenities are central to both worker welfare and employee retention. Drawing on a large-scale Dutch survey of employment conditions linked to administrative records, we construct ten multidimensional measures of job amenities to provide a comprehensive portrait of the non-wage aspects of work. We document substantial heterogeneity in amenity provision and show that firms, more than occupations or sectors, are the primary drivers of these differences, highlighting their role in shaping labour market inequality beyond wage-setting.

Contrary to the simple logic of compensating differentials, we find that high-wage jobs tend to combine higher pay with better amenities. These amenities are strong predictors of job satisfaction, absenteeism, willingness to stay, job search behaviour, and realized mobility. In particular, emotional demands, job security, supervisory quality, and opportunities for learning and innovation stand out as the most influential dimensions.

Finally, incorporating amenities alongside wages reveals greater dispersion in predicted job value, implying that overall welfare inequality across workers exceeds what is captured by income differences alone. Our findings imply that improvements to workplace amenities could significantly enhance worker welfare and, if not disproportionately targeted to the highest-wage workers, also lower inequality.

Bibliography

Adhvaryu, A., J.-F. Gauthier, A. Nyshadham, and J. Tamayo (2024): "Absenteeism, Productivity, and Relational Contracts Inside the Firm," *Journal of the European Economic Association*, 22, 1628–1677.

Bertheau, A., P. Cahuc, S. Jäger, and R. Vejlin (2022): "Turnover costs: Evidence from unexpected worker separations," Working paper.

Blatter, M., S. Muehlemann, and S. Schenker (2012): "The costs of hiring skilled workers," *European Economic Review*, 56, 20–35.

Bock, R. D. and M. Aitkin (1981): "Marginal Maximum Likelihood Estimation of Item Parameters: Application of an EM Algorithm," *Psychometrika*, 46, 443–459.

Böckerman, P. and P. Ilmakunnas (2009): "Job Disamenities, Job Satisfaction, Quit Intentions, and Actual Separations: Putting the Pieces Together," *Industrial Relations*, 48, 73–96.

——— (2012): "The Job Satisfaction–Productivity Nexus: A Study Using Matched Survey and Register Data," *ILR Review*, 65, 244–262.

Bonhomme, S. and G. Jolivet (2009): "The Pervasive Absence of Compensating Differentials," *Journal of Applied Econometrics*, 24, 763–795.

Card, D., A. Mas, E. Moretti, and E. Saez (2012): "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *American Economic Review*, 102, 2981–3003.

Clark, A. E. (2001): "What Really Matters in a Job? Hedonic Measurement Using Quit Data," *Labour Economics*, 8, 223–242.

Clark, A. E., Y. Georgellis, and P. Sanfey (1998): "Job Satisfaction, Wage Changes, and Quits: Evidence from Germany," in *Research in Labor Economics*, ed. by S. W. Polachek, JAI Press, vol. 17, 95–121.

Corradini, V., L. Lagos, and G. Sharma (2025): "Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs," *The Quarterly Journal of Economics*, published online July 9, 2025.

Edmans, A. (2011): "Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices," *Journal of Financial Economics*, 101, 621–640.

Freeman, R. B. (1978): "Job Satisfaction as an Economic Variable," *American Economic Review*, 68, 135–141.

Hall, R. E. and A. I. Mueller (2018): "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values," *Journal of Political Economy*, 126, 1594–1637.

Humlum, A., M. Rasmussen, and E. K. Rose (2025a): "Firm Premia and Match Effects in Pay vs. Amenities," Working Paper 33884, National Bureau of Economic Research, bFI Working Paper No. 2025-67.

——— (2025b): "Firm Premia and Match Effects in Pay vs. Amenities," Working Paper 33884, National Bureau of Economic Research, also circulated as BFI Working Paper No. 2025-67.

Jäger, S. and J. Heining (2022): "How substitutable are workers? Evidence from worker deaths," Tech. rep., National Bureau of Economic Research.

Koenig, F. and M. Annelli (2025): "Willingness to Pay for Workplace Amenities," *Working Paper*.

Krekel, C., G. Ward, and J.-E. De Neve (2019): "Employee Wellbeing, Productivity, and Firm Performance," CEP Discussion Paper 1605, Centre for Economic Performance, London School of Economics, London, meta-analysis of 339 studies.

Lamadon, T., J. Lise, C. Meghir, and J.-M. Robin (2024): "Labor Market Matching, Wages, and Amenities," *Working Paper*.

Le Barbanchon, T., R. Rathelot, and A. Roulet (2021): "Gender Differences in Job Search: Trading off Commute against Wage," *The Quarterly Journal of Economics*, 136, 381–426.

Lentz, R., S. Piyapromdee, and J.-M. Robin (2023): "The Anatomy of Sorting—Evidence From Danish Data," *Econometrica*, 91, 2409–2455.

Lévy-Garboua, L., C. Montmarquette, and V. Simonnet (2007): "Job Satisfaction and Quits," *Labour Economics*, 14, 251–268.

Maestas, N., K. J. Mullen, D. Powell, T. von Wachter, and J. B. Wenger (2023): "The Value of Working Conditions in the United States and Implications for the Structure of Wages," *American Economic Review*, 113, 2007–2047.

Mas, A. and A. Pallais (2017): "Valuing Alternative Work Arrangements," *American Economic Review*, 107, 3722–3759.

McDonald, R. P. (1999): *Test Theory: A Unified Treatment*, Mahwah, NJ: Lawrence Erlbaum Associates.

Morchio, I. and C. Moser (Forthcoming): "The Gender Pay Gap: Micro Sources and Macro Consequences," *American Economic Review*.

Murtagh, F. and P. Legendre (2014): "Ward's Hierarchical Agglomerative Clustering Method: Which Algorithms Implement Ward's Criterion?" *Journal of Classification*, 31, 274–295.

Oswald, A. J., E. Proto, and D. Sgroi (2015): "Happiness and Productivity," *Journal of Labor Economics*, 33, 789–822.

Rosen, S. (1986): "The theory of equalizing differences," *Handbook of Labor Economics*, 1, 641–692.

Samejima, F. (1969): "Estimation of Latent Ability Using a Response Pattern of Graded Scores," *Psychometrika Monograph Supplement*.

Sockin, J. (2022): "Show Me the Amenity: Are Higher-Paying Firms Better All Around?" CESifo Working Paper No. 9842.

Sorkin, I. (2018): "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.

Statistics Netherlands (CBS) (2024): "National Working Conditions Survey (NEA): Methodology," .

Taber, C. and R. Vejlin (2020): "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market," *Econometrica*, 88, 1031–1069.

Ward Jr., J. H. (1963): "Hierarchical Grouping to Optimize an Objective Function," *Journal of the American Statistical Association*, 58, 236–244.

Wiswall, M. and B. Zafar (2018): "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 133, 457–507.

Appendix

Appendix A: Constructing Summary Measures from Individual Amenity Variables

The survey questions in the NEA collect information about different aspects of the work environment using multiple measures. In our case, we have 70 different amenity-related variables measured on various scales (3-point, 5-point, and 7-point Likert scales, plus some continuous measures).[2] While each question captures something meaningful, analysing 70 separate variables would be unwieldy and would likely miss the underlying patterns that connect related workplace features.

Rather than arbitrarily deciding which amenities belong together, we let the data reveal natural groupings based on how workers actually respond to different questions. The key insight is that if two amenity questions tend to be answered similarly by the same workers, they likely measure aspects of the same underlying workplace characteristic.

A2.0.1 Identifying Natural Clusters of Related Variables

We calculate correlations between every pair of amenity variables using rank-based correlation methods that are appropriate for each combination of variables type: Likert-scale variables, mixed pairs (one Likert scale, one continuous variable), and continuous variables. This approach ensures we capture relationships accurately regardless of the measurement scale used for each variable.

We consider the pairwise rank-based correlations between all amenity variables and rely on block-diagonal patterns to indicate natural groupings among related variables. Using these correlations, we group variables that tend to move together by treating correlation as similarity—variables with high correlations are "close" to each other, while those with low correlations are "distant." We then use hierarchical clustering (specifically Ward's method) to identify natural groupings.[3]

The clustering process works like building a family tree in reverse: we start with each variable as its own group, then progressively merge the most similar groups until we have a manageable number of distinct clusters. We determine the optimal number of clusters using the elbow method—looking for the point where adding more clusters provides diminishing returns in explaining the data structure. The elbow method suggests 10

[2]We have more variables that we did not include in our main analysis. We dropped some of them as they are available only in a few recent years.

[3]See Ward Jr. (1963); Murtagh and Legendre (2014) for details of this method.

clusters as optimal.

A2.0.2 Creating Summary Measures Using Item Response Theory

Once we have identified clusters of related amenity variables, we create a single summary score for each cluster using Item Response Theory (IRT). This approach is superior to simple averaging because it accounts for different response scales (a 3-point scale response and a 7-point scale response are weighted appropriately), considers item quality (questions that better distinguish between workers receive more weight), handles missing data (workers who did not answer every question in a cluster can still receive a score based on their available responses), and preserves meaningful variation (the resulting scores maintain their natural scale, so clusters with more variation across workers have larger standard deviations).

An important advantage of this approach is its handling of incomplete responses. Many workplace surveys are administered in alternating years or with rotating question batteries, meaning no single respondent answers all questions. Our IRT-based approach naturally accommodates this by using all available information for each respondent, rather than requiring complete responses or resorting to imputation methods that might introduce bias.

Table A2.2 presents the complete composition of each cluster, showing how the 70 individual variables are grouped into 10 coherent domains. Each cluster represents a distinct dimension of workplace amenities as perceived by workers themselves.

The individual items contribute to the overall amenity measure depending on their individual weights. These are shown in Figure A2.1.

A2.0.3 Validation of Summary Measures

To ensure that our clustering approach produces meaningful and reliable summary measures, we conduct several validation checks:

Internal Consistency

We assess the internal consistency of each cluster using McDonald's ω (McDonald, 1999), which measures how well the items in each cluster collectively measure a single underlying construct. All clusters bar two (JobSecurity and IrregularHours) exceed the

Table A2.1: Example questions and answer options for raw variables

Does your work require intensive thinking?
[Never/Sometimes/Often/Always]
Can you decide for yourself how to do your work?
[Yes, regularly/Yes, sometimes/No]
Do you do work where you have to exert a lot of force?
[Yes, regularly/Yes, sometimes/No]
At my work, employees are encouraged to think about ways to do work better
[Never/Sometimes/Often/Always]
Does your supervisor stimulate development of your knowledge and skills?
[Yes, to a limited extent/Yes, to a large extent/No]
Did you experience physical violence by customers?
[No, never/Yes, once or a few times/Yes, often/Yes, very often]
I need to work very fast.
[Never/Sometimes/Often/Always]
Does your work bring you into emotionally difficult situations?
[Never/Sometimes/Often/Always]

Note. Example questions from the NEA between 2005 and 2023 including answer options.

Figure A2.1: Contribution of subitems to each amenity factor

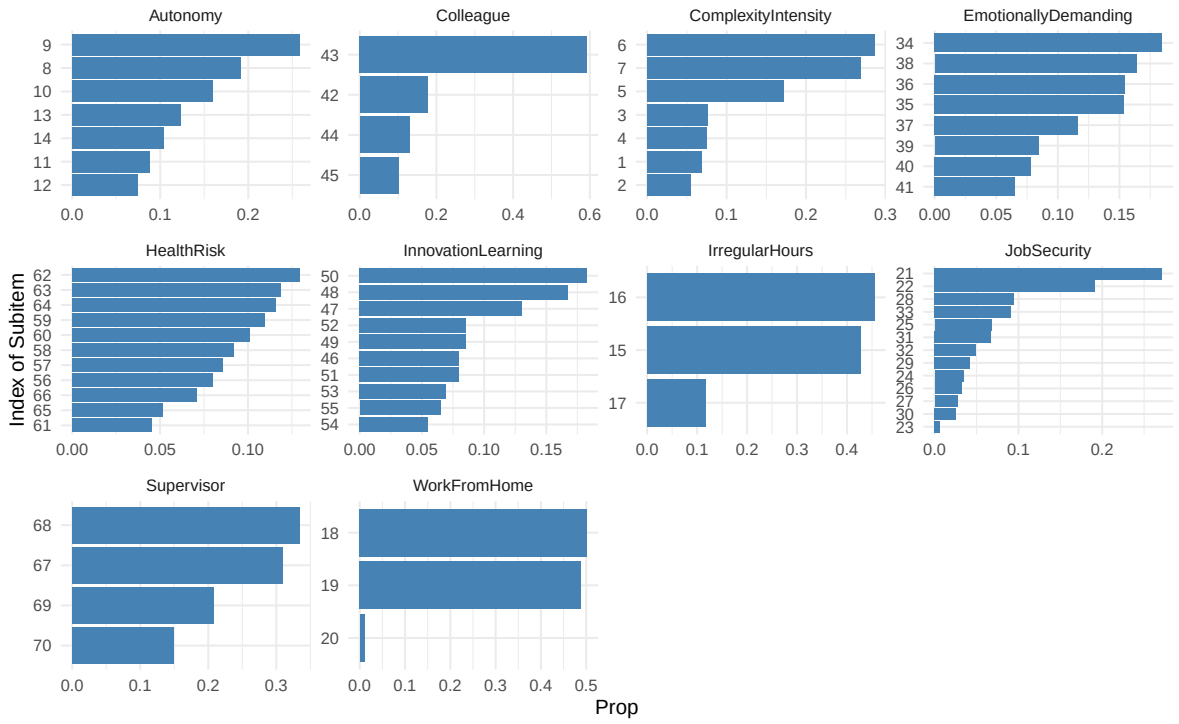

Note. Relative weights of subitems in each amenity factor. Each subitem is indexed by the subitem number presented in Table A2.2.

conventional threshold of 0.70, indicating acceptable to excellent internal consistency (Table A2.3).

Table A2.2: Composition of amenity clusters

Amenity	**Subitems**
Complexity & Work Intensity	Does your work require extensive thinking? (1) Does your work require you to keep your mind on it? (2) Does your work require a lot of attention from you? (3) Do you often receive so much information in a working day that you have trouble processing it quickly enough? (4) I need to work very fast. (5) I need to do a lot of work. (6) I need to work extra hard. (7)
Autonomy	Can you decide how to perform your work? (8) Can you determine the order of work? (9) Can you control your own pace? (10) In your work, do you have to come up with your own solutions? (11) Can you take leave whenever you want? (12) Can you decide where you work? (13) Can you determine for yourself what times you work? (14)
Irregular Hours	Do you ever work in the evenings (7pm-12am) or at night (12am-6am)? (15) Do you ever work on the weekend? (16) Do you regularly or sometimes work overtime? (17)
Work From Home	Do you ever telework (via a connection to your company network)? (18) Do you telework at least half a day per week? (19) I usually or always work from home. (20)
Job Security	Are you at risk of losing your job? (21) Are you worried about keeping your job? (22) Firm-level firing rate. (23) What is your position in the workforce? (24) Has there been a major reorganisation (25) / takeover by another organisation (26) / acquisition of another organisation (27) / downsizing with forced layoffs (28) / downsizing without forced layoffs (29) / merger with another company (30) / outsourcing of support services (31) / automation of business activities (32) / relocation of the organisation (33)?
Emotionally Demanding	I feel emotionally drained by my work. (34) At the end of the working day, I feel empty. (35) I feel tired when I get up in the morning and face my work. (36) It takes a lot out of me to work with people all day long. (37) I feel completely exhausted by my work. (38) Does your work bring you into emotionally difficult situations? (39) Is your work emotionally demanding? (40) Do you get emotionally involved in your work? (41)
Colleague	My colleagues take a personal interest in me. (42) My colleagues are friendly. (43) My colleagues are good at their work. (44) My colleagues help to get the work done. (45)
Innovation & Learning	Does your job require creativity? (46) In my work, I get time to develop new ideas. (47) In my work, I make a clear contribution to improving my company's products/services. (48) At my work, employees are encouraged to think about ways to do their work better. (49) In my work, I make a clear contribution to developing new products/services. (50) Does your job require you to learn new things? (51) Do you learn a lot or little from the tasks you perform for your job? (52) Do you learn a lot or little from people at work, such as colleagues, managers, and customers? (53) Does your supervisor stimulate development of your knowledge and skills? (54) Is your work varied? (55)
Health Risk	Do you need to do hazardous work? (56) Do you NOT need to do hazardous work? (57) Do you work with water or aqueous solutions? (58) Do you get substances on your skin during your work? (59) Do you inhale substances while working? (60) Are you in contact with potentially infectious persons, animals or material? (61) Do you do work that requires a lot of force? (62) Does your work involve the use of any tool that causes vibration or shaking? (63) Do you do work in an uncomfortable work position? (64) Do you do work that requires repetitive movements? (65) Is there so much noise in your workplace that you have to talk loudly? (66)
Supervisor	My supervisor is attentive to the well-being of employees. (67) My supervisor pays attention to what I say. (68) My supervisor helps get the work done. (69) My supervisor can get people to work well together. (70)

Note. Cluster composition and descriptions of individual subitems. Aggregation of detailed workplace characteristics into groups of amenities. Each question or statement refers to a survey item.

Proportion of Variance Explained by First Factor

We evaluate the unidimensionality of each cluster by examining the proportion of variance explained by the first factor (PV1). This metric indicates how much of the total variation in the cluster's items is captured by a single underlying factor. Values closer to 1.0 suggest that items within a cluster are primarily measuring one underlying construct, while lower values indicate that additional dimensions may be present. Only JobSecurity – comprised of 13 items – has a low PV1, with 0.1 All other PV1 values range from 0.35 to 0.7, with most clusters showing values, indicating that a substantial portion of the variance in each cluster can be attributed to a single common factor.

Table A2.3: Internal validation of amenity clusters

Cluster	McDonald's ω	PV1	S.D.	# Subitems
Autonomy	0.82	0.41	0.88	6
EmotionallyDemanding	0.86	0.44	0.94	8
Colleague	0.85	0.58	0.88	4
ComplexityIntensity	0.81	0.39	0.93	6
HealthRisk	0.86	0.37	0.87	11
InnovationLearning	0.85	0.36	0.81	7
IrregularHours	0.60	0.35	0.78	3
JobSecurity	0.53	0.10	0.70	13
Supervisor	0.87	0.63	0.91	4
WorkFromHome	0.87	0.70	0.80	3

Note. McDonald's ω measures how well the items in each amenities cluster collectively measure the underlying amenity. PV1 refers to the proportion of the variance explained by the first factor.

The variation in standard deviations across domains is particularly informative. Domains with larger standard deviations (such as ComplexityIntensity, SD = 0.93, and EmotionallyDemanding, SD = 0.94) indicate areas where workers experience more heterogeneous conditions, while domains with smaller standard deviations (such as IrregularHours, SD = 0.78, and Job Security, SD = 0.70) suggest more uniform experiences across workers.

The distribution of scores for each domain reveals the shape and range of variation in worker experiences across different amenity dimensions. Some domains show approximately normal distributions, while others exhibit skewness or discrete patterns that reflect the underlying structure of workplace experiences. The discrete or sparse appearance of certain distributions arises from the nature of the underlying survey items—domains based on binary variables or fewer items naturally produce more limited ranges of possible factor scores.[4]

[4]For example, the IrregularHours domain consists of only three variables with only two to three answer options each, which constrains the possible combinations of responses and results in discrete

factor score values. Similarly, domains with fewer items (like WorkFromHome with only three items) have fewer possible response patterns compared to domains with many items and varied response scales (like InnovationLearning with 10 items). This discrete structure is methodologically appropriate and reflects the actual response patterns in the survey data.

Appendix B: Additional tables

Table B2.1: Summary statistics of the NEA sample

Variable	Mean	Std	N
Gender: Women	0.505	0.500	308415
Gender: Men	0.495	0.500	308415
Education: High	0.484	0.500	308415
Education: Middle	0.390	0.488	308415
Education: Low	0.126	0.332	308415
Age group: 22–34	0.263	0.440	308415
Age group: 35–49	0.384	0.486	308415
Age group: 50–64	0.353	0.478	308415
Family status: Married with child	0.504	0.500	308415
Family status: Married w/o child	0.290	0.454	308415
Family status: Single	0.147	0.354	308415
Family status: Single parent	0.047	0.211	308415
Firm size < 10	0.058	0.223	307208
Firm size 10-99	0.367	0.482	307208
Firm size 100-500	0.273	0.445	307208
Firm size 500-999	0.078	0.269	307208
Firm size ≥ 1000	0.224	0.417	307208
Hours worked (monthly)	142.68	36.17	308415

Table B2.2: Summary statistics of the outcomes considered

Variable	Mean	Std	N
Job satisfaction	3.868	0.829	308,415
Not absent past 12 months	0.493	0.500	307,921
Share not absent	0.955	0.137	303,729
Job search	0.755	0.430	308,415
Intention to stay next 5 years	0.690	0.462	308,415
Stay next 12 months	0.934	0.248	308,415
Job-to-job transition next 12 months	0.049	0.215	308,415

Note. Job satisfaction is measured on a 5-point Likert scale. Share not absent is the share of scheduled working hours worked (i.e., not absent) during the previous 12 months. All other outcomes are binary indicators.

Table B2.3: Regression results: Wages (standardised)

	(1)	(2)	(3)	(4)	(5)
ComplexityIntensity	0.116	0.090	0.087	0.073	0.001
	(0.003)	(0.003)	(0.002)	(0.002)	(0.005)
JobSecurity	-0.021	-0.007	0.022	0.017	-0.001
	(0.007)	(0.006)	(0.002)	(0.002)	(0.005)
Autonomy	0.176	0.088	0.082	0.070	0.019
	(0.005)	(0.004)	(0.004)	(0.003)	(0.007)
EmotionallyDemanding	0.033	0.023	0.024	0.015	0.009
	(0.003)	(0.003)	(0.002)	(0.002)	(0.006)
Colleague	-0.018	0.009	0.005	0.003	-0.004
	(0.002)	(0.002)	(0.001)	(0.001)	(0.004)
InnovationLearning	0.084	0.071	0.047	0.024	-0.004
	(0.003)	(0.003)	(0.002)	(0.002)	(0.005)
HealthRisk	-0.166	-0.101	-0.108	-0.103	-0.014
	(0.006)	(0.005)	(0.003)	(0.002)	(0.008)
IrregularHours	0.117	0.095	0.099	0.076	0.010
	(0.003)	(0.003)	(0.002)	(0.002)	(0.005)
Supervisor	-0.011	-0.001	-0.006	-0.006	0.005
	(0.003)	(0.002)	(0.002)	(0.001)	(0.005)
WorkFromHome	0.258	0.157	0.122	0.096	0.024
	(0.004)	(0.003)	(0.003)	(0.002)	(0.006)
Year FE	✓	✓	✓	✓	✓
Demographics		✓	✓	✓	✓
Firm FE			✓	✓	✓
Occupation FE				✓	✓
Individual FE					✓
Observations	308,415	308,415	308,415	308,415	22,331
R^2	0.297	0.490	0.658	0.693	0.974

Note. Regression of standardised wages on standardised amenities. Columns (1)–(5) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE; +Individual FE).

Table B2.4: Job satisfaction

	(1)	(2)	(3)	(4)	(5)
ComplexityIntensity	-0.004	-0.007	0.001	0.003	-0.006
	(0.003)	(0.003)	(0.003)	(0.003)	(0.020)
JobSecurity	0.087	0.085	0.079	0.077	0.079
	(0.002)	(0.002)	(0.002)	(0.002)	(0.016)
Autonomy	0.036	0.043	0.059	0.065	0.098
	(0.003)	(0.003)	(0.003)	(0.003)	(0.022)
EmotionallyDemanding	-0.250	-0.252	-0.258	-0.260	-0.211
	(0.002)	(0.002)	(0.003)	(0.003)	(0.026)
Colleague	0.081	0.080	0.076	0.075	0.046
	(0.002)	(0.002)	(0.002)	(0.002)	(0.016)
InnovationLearning	0.136	0.138	0.132	0.129	0.103
	(0.002)	(0.002)	(0.002)	(0.002)	(0.018)
HealthRisk	0.000	0.005	-0.009	-0.016	-0.014
	(0.002)	(0.002)	(0.003)	(0.003)	(0.029)
IrregularHours	0.013	0.020	0.021	0.018	0.020
	(0.002)	(0.002)	(0.002)	(0.002)	(0.016)
Supervisor	0.177	0.177	0.174	0.175	0.124
	(0.002)	(0.002)	(0.002)	(0.002)	(0.017)
WorkFromHome	0.000	0.002	0.004	0.005	0.016
	(0.002)	(0.002)	(0.002)	(0.002)	(0.019)
Wage	0.065	0.070	0.049	0.040	0.074
	(0.003)	(0.003)	(0.003)	(0.003)	(0.042)
Year FE	✓	✓	✓	✓	✓
Demographics		✓	✓	✓	✓
Firm FE			✓	✓	✓
Occupation FE				✓	✓
Individual FE					✓
Observations	308,415	308,415	308,415	308,415	22,331
R^2	0.220	0.223	0.322	0.324	0.732

Note. Regression of job satisfaction on standardised amenities and wages. Columns (1)–(5) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE; +Individual FE).

Table B2.5: Stay 5 years

	(1)	(2)	(3)	(4)	(5)
ComplexityIntensity	-0.007	-0.007	-0.001	-0.000	-0.010
	(0.001)	(0.001)	(0.001)	(0.001)	(0.009)
JobSecurity	0.006	0.008	0.003	0.002	0.004
	(0.001)	(0.001)	(0.001)	(0.001)	(0.007)
Autonomy	-0.001	-0.005	0.002	0.004	0.019
	(0.001)	(0.001)	(0.001)	(0.001)	(0.010)
EmotionallyDemanding	-0.092	-0.080	-0.080	-0.080	-0.076
	(0.001)	(0.001)	(0.001)	(0.001)	(0.010)
Colleague	0.003	0.012	0.012	0.012	0.010
	(0.001)	(0.001)	(0.001)	(0.001)	(0.007)
InnovationLearning	0.044	0.053	0.049	0.049	0.047
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
HealthRisk	0.038	0.021	0.010	0.007	-0.007
	(0.001)	(0.001)	(0.001)	(0.001)	(0.012)
IrregularHours	-0.007	-0.003	0.002	0.002	0.008
	(0.001)	(0.001)	(0.001)	(0.001)	(0.007)
Supervisor	0.065	0.068	0.065	0.065	0.059
	(0.001)	(0.001)	(0.001)	(0.001)	(0.007)
WorkFromHome	-0.014	-0.005	-0.004	-0.003	0.005
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
Wage	0.027	0.014	-0.014	-0.012	0.019
	(0.002)	(0.002)	(0.002)	(0.002)	(0.021)
Year FE	✓	✓	✓	✓	✓
Demographics		✓	✓	✓	✓
Firm FE			✓	✓	✓
Occupation FE				✓	✓
Individual FE					✓
Observations	308,415	308,415	308,415	308,415	22,331
R^2	0.101	0.132	0.255	0.257	0.742

Note. Regression of stay 5 years, a binary indicator, on standardised amenities and wages. Columns (1)–(5) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE; +Individual FE).

Table B2.6: Stay 12 months

	(1)	(2)	(3)	(4)
ComplexityIntensity	0.000	0.000	0.002	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
JobSecurity	0.022	0.022	0.016	0.017
	(0.003)	(0.003)	(0.001)	(0.001)
Autonomy	0.004	0.003	0.003	0.003
	(0.001)	(0.001)	(0.001)	(0.001)
EmotionallyDemanding	-0.005	-0.003	-0.005	-0.005
	(0.001)	(0.001)	(0.001)	(0.001)
Colleague	-0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.000)	(0.001)
InnovationLearning	0.003	0.006	0.004	0.004
	(0.001)	(0.001)	(0.001)	(0.001)
HealthRisk	0.006	0.004	0.003	0.004
	(0.001)	(0.001)	(0.001)	(0.001)
IrregularHours	-0.003	-0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Supervisor	0.005	0.006	0.007	0.007
	(0.001)	(0.001)	(0.001)	(0.001)
WorkFromHome	0.001	0.002	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Wage	0.017	0.010	0.002	0.003
	(0.002)	(0.002)	(0.001)	(0.001)
Year FE	✓	✓	✓	✓
Demographics		✓	✓	✓
Firm FE			✓	✓
Occupation FE				✓
Observations	308,415	308,415	308,415	308,415
R^2	0.016	0.024	0.315	0.316

Note. Regression of stay 12 months, a binary indicator, on standardised amenities and wages. Columns (1)–(4) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE).

Table B2.7: No job-to-job transition in the past 12 months

	(1)	(2)	(3)	(4)
ComplexityIntensity	-0.000	-0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.000)
JobSecurity	0.012	0.012	0.007	0.007
	(0.002)	(0.002)	(0.001)	(0.001)
Autonomy	0.003	0.003	0.002	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
EmotionallyDemanding	-0.002	-0.000	-0.002	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
Colleague	-0.001	0.001	0.001	0.001
	(0.000)	(0.000)	(0.000)	(0.000)
InnovationLearning	0.001	0.003	0.001	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
HealthRisk	0.005	0.002	0.002	0.002
	(0.001)	(0.001)	(0.001)	(0.000)
IrregularHours	-0.002	-0.001	0.000	0.000
	(0.001)	(0.001)	(0.000)	(0.000)
Supervisor	0.003	0.004	0.005	0.005
	(0.001)	(0.001)	(0.000)	(0.000)
WorkFromHome	-0.001	0.001	-0.000	-0.000
	(0.001)	(0.001)	(0.000)	(0.001)
Wage	0.013	0.007	0.002	0.002
	(0.002)	(0.002)	(0.001)	(0.001)
Year FE	✓	✓	✓	✓
Demographics		✓	✓	✓
Firm FE			✓	✓
Occupation FE				✓
Observations	308,415	308,415	308,415	308,415
R^2	0.009	0.016	0.290	0.290

Note. Regression of no job-to-job transition in the past 12 months, a binary indicator, on standardised amenities and wages. Columns (1)–(4) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE).

Table B2.8: No job search

	(1)	(2)	(3)	(4)
ComplexityIntensity	-0.002	-0.001	0.000	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
JobSecurity	0.064	0.065	0.059	0.058
	(0.002)	(0.002)	(0.001)	(0.001)
Autonomy	-0.002	-0.001	0.001	0.003
	(0.001)	(0.001)	(0.001)	(0.001)
EmotionallyDemanding	-0.048	-0.041	-0.041	-0.041
	(0.001)	(0.001)	(0.001)	(0.001)
Colleague	0.002	0.009	0.009	0.008
	(0.001)	(0.001)	(0.001)	(0.001)
InnovationLearning	0.015	0.021	0.018	0.017
	(0.001)	(0.001)	(0.001)	(0.001)
HealthRisk	0.021	0.010	0.002	-0.003
	(0.001)	(0.001)	(0.001)	(0.001)
IrregularHours	-0.010	-0.005	-0.002	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
Supervisor	0.042	0.043	0.041	0.042
	(0.001)	(0.001)	(0.001)	(0.001)
WorkFromHome	-0.020	-0.012	-0.011	-0.010
	(0.001)	(0.001)	(0.001)	(0.001)
Wage	0.026	0.021	0.004	0.006
	(0.002)	(0.002)	(0.001)	(0.002)
Year FE	✓	✓	✓	✓
Demographics		✓	✓	✓
Firm FE			✓	✓
Occupation FE				✓
Observations	308,415	308,415	308,415	308,415
R^2	0.070	0.089	0.208	0.210

Note. Regression of not searching for a new job, a binary indicator, on standardised amenities and wages. Columns (1)–(4) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE).

Table B2.9: No absenteeism (past 12 months)

	(1)	(2)	(3)	(4)	(5)
ComplexityIntensity	0.016	0.019	0.014	0.013	0.013
	(0.001)	(0.001)	(0.001)	(0.001)	(0.010)
JobSecurity	0.010	0.013	0.011	0.011	0.008
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
Autonomy	0.013	0.009	0.010	0.010	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.011)
EmotionallyDemanding	-0.088	-0.083	-0.079	-0.079	-0.057
	(0.001)	(0.001)	(0.001)	(0.001)	(0.011)
Colleague	-0.013	-0.009	-0.007	-0.007	0.003
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
InnovationLearning	0.007	0.009	0.011	0.011	0.013
	(0.001)	(0.001)	(0.001)	(0.001)	(0.009)
HealthRisk	-0.011	-0.018	-0.028	-0.027	-0.006
	(0.001)	(0.001)	(0.002)	(0.002)	(0.014)
IrregularHours	0.013	0.012	0.007	0.007	0.004
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
Supervisor	0.004	0.004	0.004	0.003	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.008)
WorkFromHome	0.000	0.002	0.009	0.009	0.030
	(0.001)	(0.001)	(0.001)	(0.001)	(0.010)
Wage	0.023	0.011	0.027	0.023	0.014
	(0.001)	(0.002)	(0.002)	(0.002)	(0.026)
Year FE	✓	✓	✓	✓	✓
Demographics		✓	✓	✓	✓
Firm FE			✓	✓	✓
Occupation FE				✓	✓
Individual FE					✓
Observations	307,921	307,921	307,879	307,879	22,246
R^2	0.039	0.046	0.167	0.168	0.715

Note. Regression of no absenteeism in past 12 months, a binary indicator, on standardised amenities and wages. Columns (1)–(5) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE; +Individual FE).

Table B2.10: Share of working time not absent

	(1)	(2)	(3)	(4)	(5)
ComplexityIntensity	0.330	0.366	0.282	0.265	0.467
	(0.037)	(0.036)	(0.037)	(0.037)	(0.337)
JobSecurity	0.679	0.706	0.760	0.766	0.580
	(0.034)	(0.033)	(0.039)	(0.039)	(0.264)
Autonomy	0.329	0.257	0.260	0.241	0.392
	(0.034)	(0.034)	(0.038)	(0.040)	(0.348)
EmotionallyDemanding	-2.036	-2.117	-2.098	-2.104	-2.132
	(0.047)	(0.048)	(0.051)	(0.051)	(0.407)
Colleague	0.097	0.024	0.020	0.019	0.162
	(0.029)	(0.029)	(0.031)	(0.031)	(0.252)
InnovationLearning	0.255	0.122	0.125	0.131	-0.196
	(0.031)	(0.030)	(0.033)	(0.034)	(0.289)
HealthRisk	-0.656	-0.576	-0.577	-0.513	0.239
	(0.039)	(0.036)	(0.041)	(0.044)	(0.433)
IrregularHours	0.573	0.392	0.393	0.409	0.642
	(0.030)	(0.030)	(0.031)	(0.031)	(0.241)
Supervisor	0.062	0.027	0.036	0.033	0.032
	(0.033)	(0.033)	(0.036)	(0.036)	(0.262)
WorkFromHome	0.047	-0.065	0.011	0.003	0.215
	(0.035)	(0.034)	(0.037)	(0.037)	(0.286)
Wage	0.305	0.424	0.571	0.562	1.478
	(0.036)	(0.040)	(0.043)	(0.046)	(0.666)
Year FE	✓	✓	✓	✓	✓
Demographics		✓	✓	✓	✓
Firm FE			✓	✓	✓
Occupation FE				✓	✓
Individual FE					✓
Observations	303,729	303,729	303,383	303,383	21,643
R^2	0.036	0.046	0.139	0.139	0.643

Note. Regression of the share of working days not absent in the past 12 months on standardised amenities and wages. Columns (1)–(5) add controls/FEs cumulatively (Year FE; +Demographics; +Firm FE; +Occupation FE; +Individual FE).

Table B2.11: Firm-level regressions of outcomes on amenities (one specification)

	Job satisfaction	Not absent 12 months	Share not absent	No job search	Stay 5 years	Stay 12 months	Firm wage FE
ComplexityIntensity	-0.012	0.019	0.333	0.001	-0.024	-0.011	0.000
	(0.007)	(0.002)	(0.049)	(0.002)	(0.002)	(0.001)	(0.007)
JobSecurity	0.103	0.007	0.362	0.043	0.000	0.001	-0.007
	(0.007)	(0.002)	(0.044)	(0.001)	(0.002)	(0.001)	(0.008)
Autonomy	0.017	-0.001	0.125	-0.003	0.014	0.022	-0.023
	(0.007)	(0.002)	(0.050)	(0.002)	(0.002)	(0.001)	(0.007)
EmotionallyDemanding	-0.249	-0.049	-1.058	-0.024	-0.008	0.003	0.015
	(0.008)	(0.002)	(0.052)	(0.002)	(0.002)	(0.001)	(0.008)
Colleague	0.086	-0.007	0.001	0.002	0.003	0.002	-0.005
	(0.007)	(0.002)	(0.046)	(0.002)	(0.002)	(0.001)	(0.007)
InnovationLearning	0.155	0.000	0.043	0.014	0.024	0.000	0.167
	(0.007)	(0.002)	(0.050)	(0.002)	(0.002)	(0.001)	(0.008)
HealthRisk	0.015	-0.002	-0.412	0.007	0.021	-0.002	0.015
	(0.008)	(0.002)	(0.051)	(0.002)	(0.002)	(0.001)	(0.008)
IrregularHours	0.028	0.010	0.122	-0.007	0.021	-0.002	0.015
	(0.007)	(0.002)	(0.045)	(0.001)	(0.002)	(0.001)	(0.008)
Supervisor	0.198	0.006	0.015	0.029	-0.002	-0.003	-0.097
	(0.008)	(0.002)	(0.052)	(0.002)	(0.002)	(0.001)	(0.007)
WorkFromHome	0.012	-0.000	-0.119	-0.011	0.043	0.002	0.255
	(0.008)	(0.002)	(0.054)	(0.002)	(0.002)	(0.001)	(0.009)
Firm wage FE	0.090	-0.019	-0.096	0.027	-0.007	-0.003	
	(0.007)	(0.002)	(0.046)	(0.002)	(0.002)	(0.001)	
Year FE	✓	✓	✓	✓	✓	✓	✓
Observations	18,228	18,228	18,225	18,228	18,228	18,228	18,228
R^2	0.264	0.050	0.044	0.141	0.178	0.040	0.136

Note. All specifications contain year FE.

Appendix C: Additional figures

Figure C2.1: Amenities by worker productivity

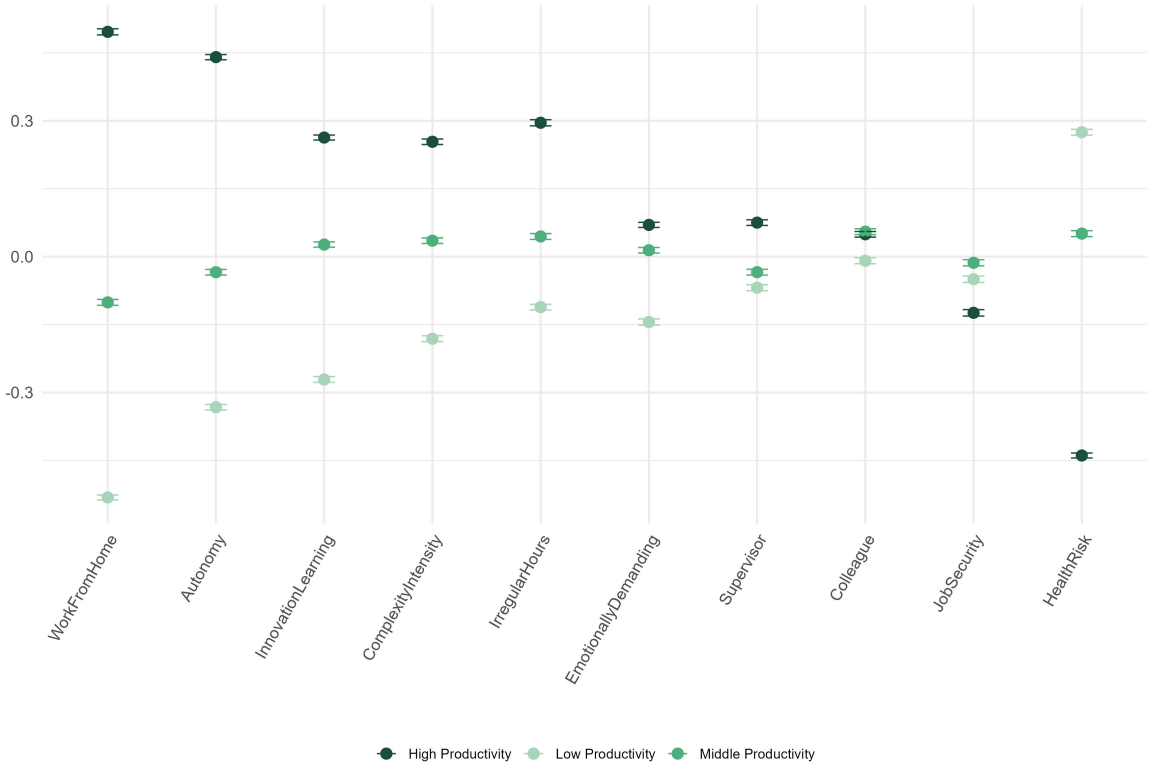

Note. Average amenity level (standardised) by productivity of respondent. Individual productivity is measured in worker AKM fixed effects and low/medium/high are terciles of these. Amenity measures are standardised with mean zero.

Figure C2.2: Amenities by broad occupational group

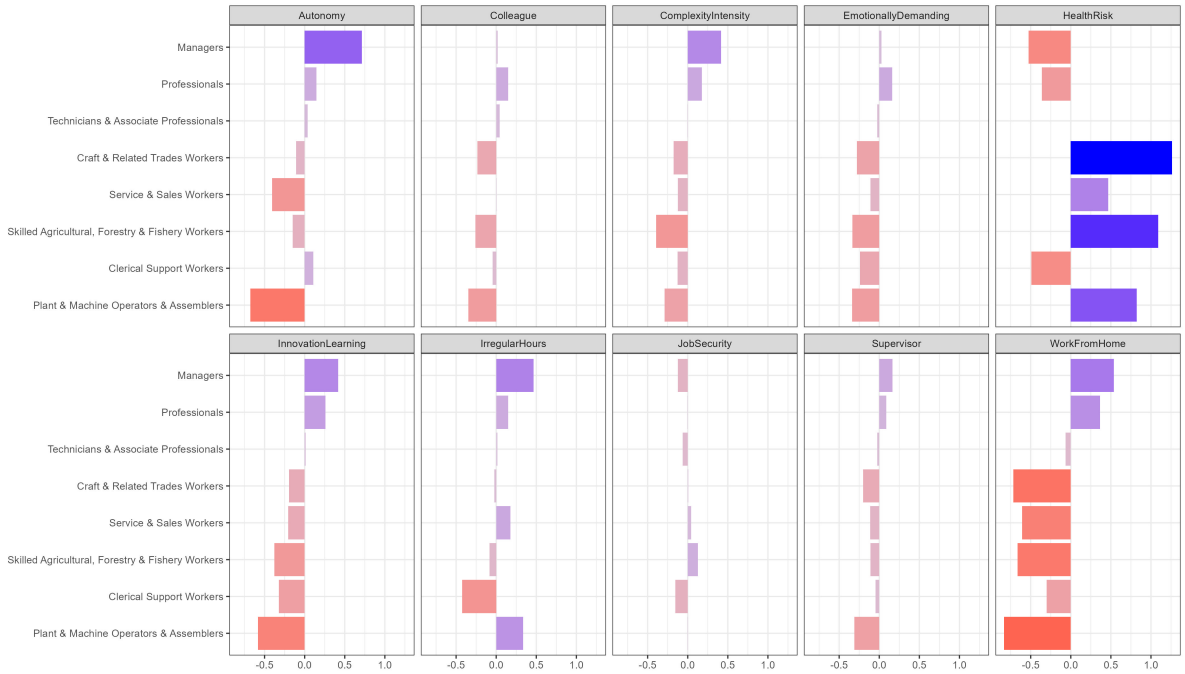

Note. Amenities by broad occupational groups. Amenity measures are standardised with mean zero.

Figure C2.3: Amenities by sector

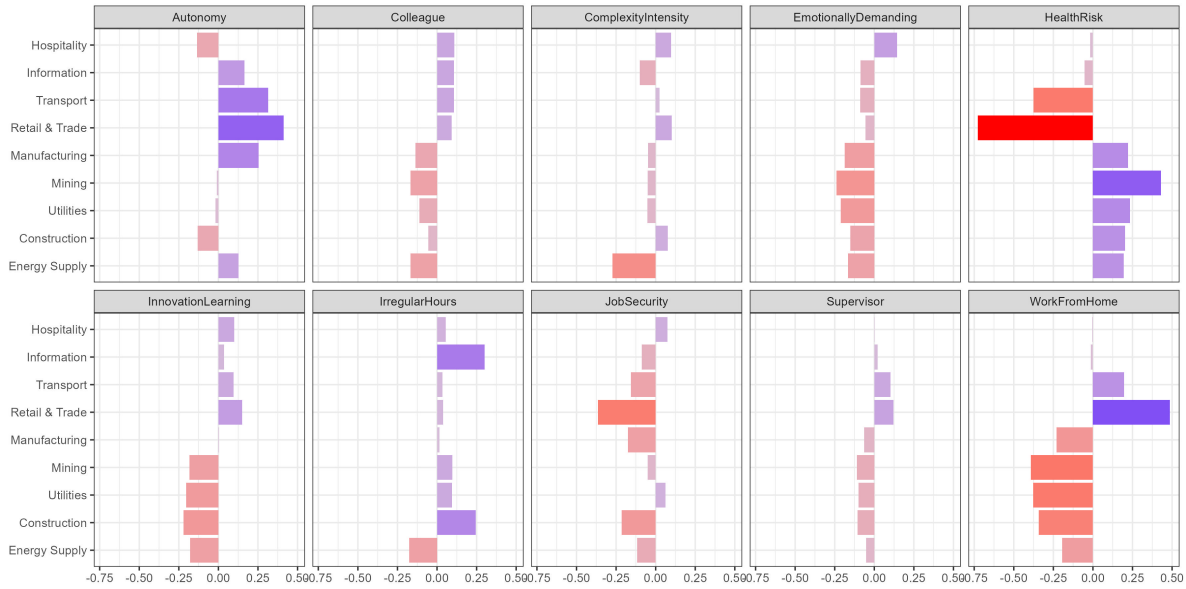

Note. Amenities by sectors. Amenity measures are standardised with mean zero.

Figure C2.4: Amenities by firm productivity

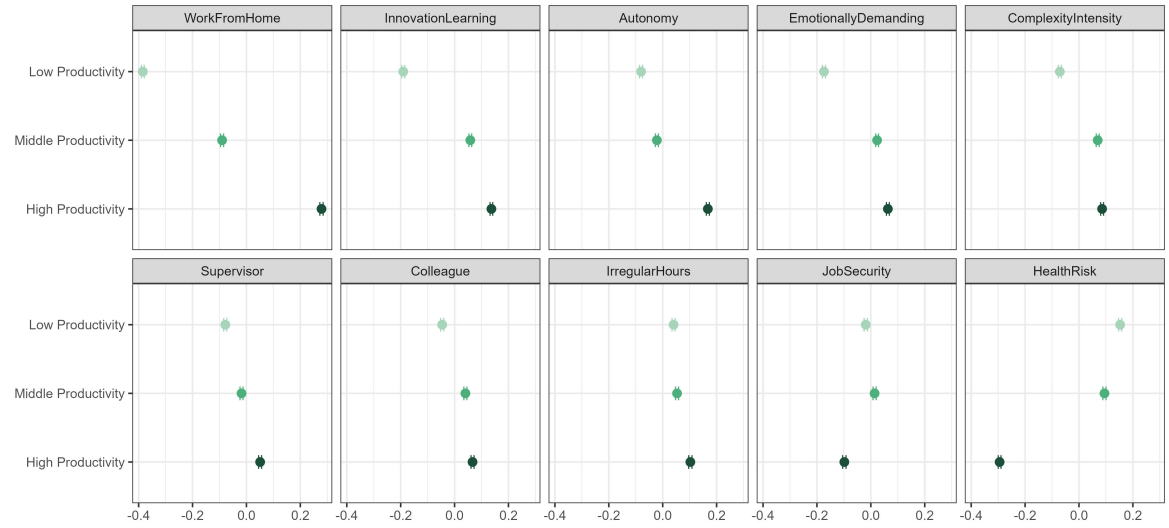

Note. Average amenity level (standardised) by firm productivity. Firm productivity is measured in firm AKM fixed effects. Amenity measures are standardised with mean zero and standard deviation one.

Figure C2.5: Time trends for selected amenities.

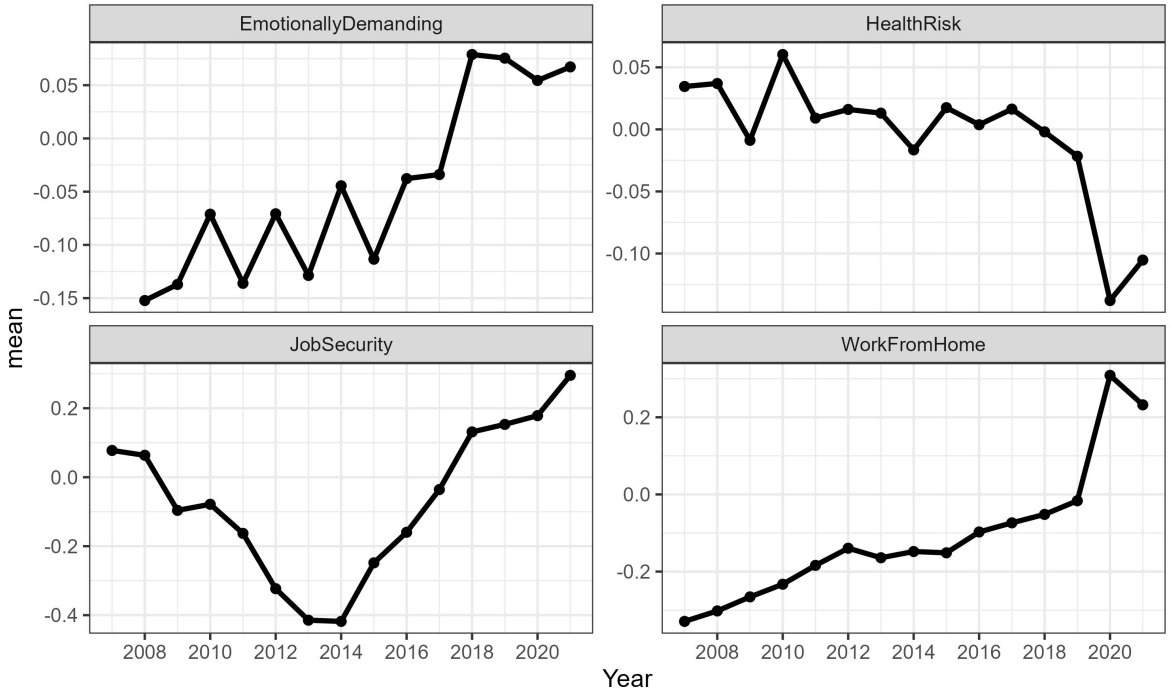

Note. Average amenity level (standardised) across years. Amenity measures are standardised with mean zero.

Figure C2.6: Sectors explain only about 1 to 7% of amenity variation

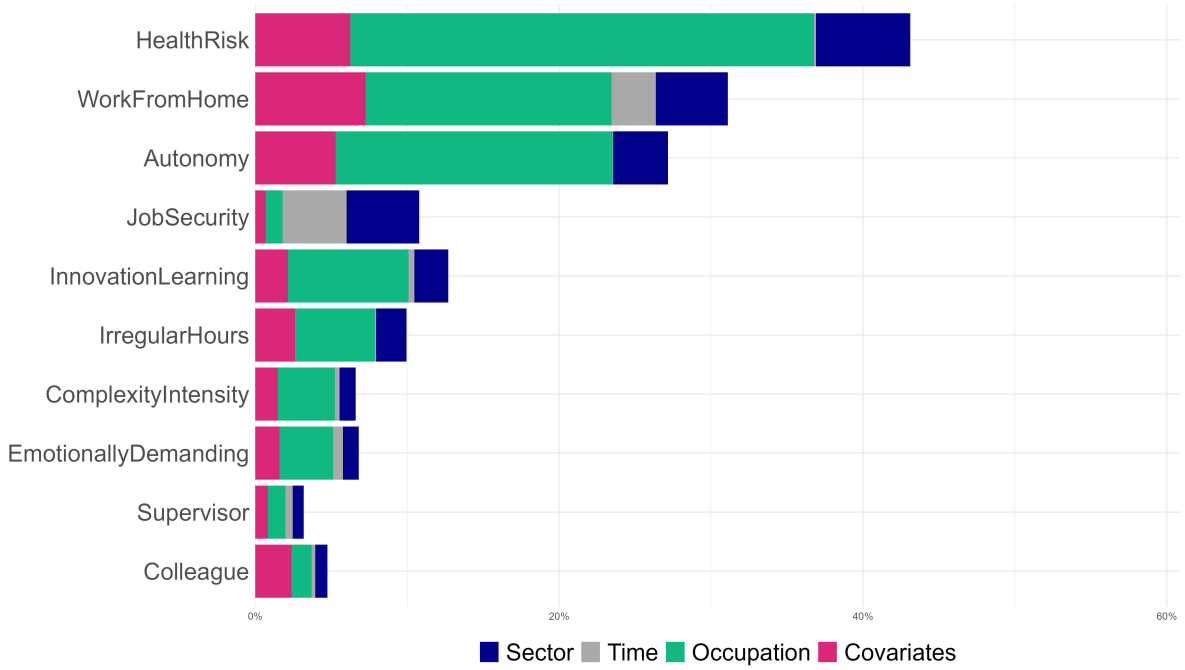

Note. Decomposition of amenity variation into sector, time (year), occupation, and individual covariates. The x-axis measures the total share of explained variation, and colors indicate the share explained by each of the variables.

Figure C2.7: Regressions of amenities on job value and mobility at the firm level

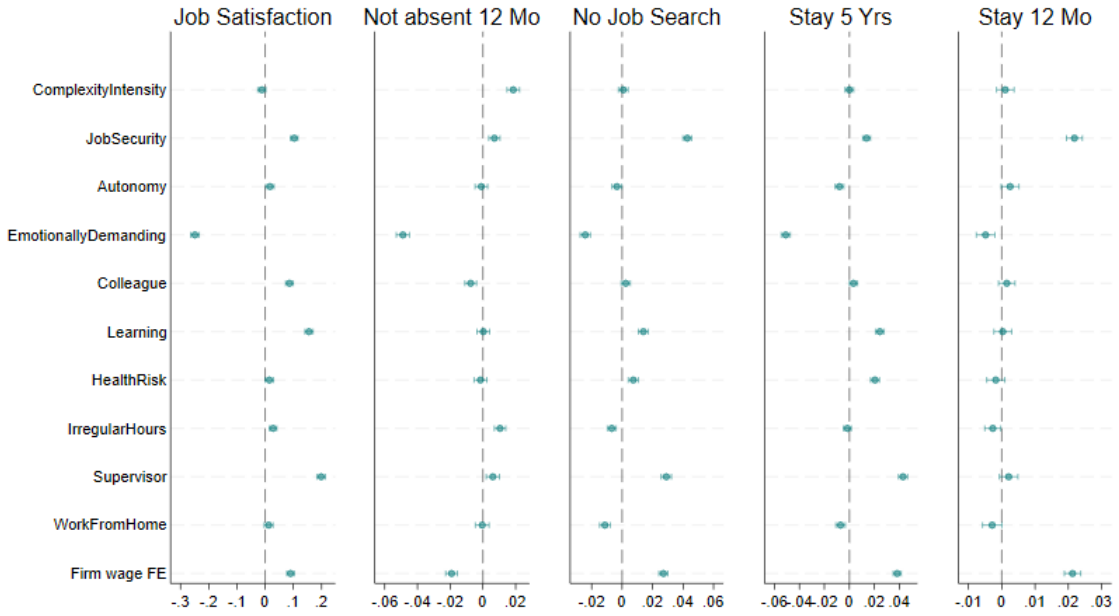

Note. This figure reports regressions of firm-level outcomes on firm-level amenities and wages. Firm-level amenities are constructed by averaging responses from firms with at least three observations, netting out demographic characteristics (age, gender, education) of the reporting workers and survey year. All amenities, including wages, and the job satisfaction measure, are standardised with mean zero and standard deviation one. Absenteeism in the past 12 months, intention to stay for five more years, no job search, and staying for the following 12 months are binary outcomes estimated using linear probability models. All regression specifications control for year fixed effects.

Chapter 3

FDI and Remittances

This chapter is co-authored with Simon Görlach (Università Bocconi).

Abstract

We aim to understand how returns to FDI and remittances contribute to cross-country risk-sharing over the business cycle. First, we document the relationship between these flows and GDP growth in the recipient economies: FDI behaves procyclically and remittances behave countercyclically with respect to receiving countries' growth. We then build a model, using these facts to inform our modelling choices. Estimating the model's underlying parameters allows us to quantitatively evaluate the effect of economic policies to further international capital markets, such as lowering the cost of sending remittances or FDI, on capital flows and the intensity of business cycle fluctuations. A key benefit of such integration is that it can help share risk across countries. Our model allows us to directly compare the welfare effects of FDI and remittances, highlighting their distinct roles in cross-country risk sharing.

Keywords: Foreign Direct Investment, Remittances, Risk sharing
JEL Codes: F21, F24

3.1 Introduction

Efficient allocation of factors, such as labour and capital, plays a critical role in driving economic growth, and it can be achieved through reallocation across sectors (Duarte and Restuccia, 2010) or within the same sector (Hsieh and Klenow, 2009). However, the optimal allocation of factors need not be limited to national borders. Migration allows individuals to work in locations where they are more productive, without the need for capital relocation. Similarly, capital can be invested abroad in the form of foreign direct investment (FDI), enhancing productivity for workers in their home countries without necessitating migration for employment.

In addition to enhancing efficiency, international migration and FDI offer a means of diversifying away from risks that can impact the domestic economy, which is the focus of this project. When residents invest money abroad or migrants send remittances – transfers to their families in their home countries –, they create separate sources of funding that are far less exposed to shocks that only hit the domestic economy. This diversification of income sources potentially leads to a form of cross-country income smoothing, bolstering the resilience of the economy and reducing the vulnerability of individuals and households to domestic economic fluctuations.

We aim to understand how returns to FDI and remittances contribute to cross-country risk-sharing over the business cycle. We document two key observations: Firstly, remittances follow a countercyclical pattern with respect to GDP growth in the migrants' home country. Conversely, FDI – like investment more generally – behaves procyclically in relation to GDP growth in the investment destination. While these facts have been documented separately in multiple contexts before (Constantinescu and Schiff, 2014; Frankel, 2011; Balli and Rana, 2015),[1] they have not yet been examined in a unified model that incorporates potential interaction or substitution effects, which matter for welfare consequences.

In light of these empirical regularities, we build an equilibrium model of the economy. Firms produce output according to a constant returns to scale production function, renting capital and employing both native and migrant workers. The two types of workers are assumed to be imperfectly substitutable and supply labour inelastically. Their wages are determined by the marginal product of labour. Native workers divide their income into consumption of domestically produced goods and imported goods. Migrant workers proceed likewise, and have the additional option of remitting part of

[1] A few country-specific studies do not find a countercyclical pattern between growth in migrants' countries of origin and remittances. Akkoyunlu and Kholodilin (2008) instead find an acyclical relationship between remittances and growth in home country and a procyclical one with growth in the host country. Lueth and Ruiz-Arranz (2007) find evidence of procyclicality of remittances for Sri Lanka.

their income to family in their home country. Sending remittances comes at a cost that is proportional to the amount remitted, in line with empirical evidence. The optimal level of remittances depends on an altruism parameter and increases in the wage rate of migrants and decreases in the wage in migrants' country of origin. Investors can invest in production at home or abroad (in the form of FDI), and allocate capital optimally in order to equate marginal returns (net of costs, as investing abroad also comes at a cost). In equilibrium, markets clear: Goods are traded and consumed internationally, and trade is balanced. Capital and labour markets also clear. As we focus on international capital flows, which can respond faster and are less constrained than international labour mobility, we hold migrant stocks fixed in our analysis.

We estimate the model using a method of moments estimator to estimate our model on twelve OECD countries.[2] Our estimates show that the cost of remittances is larger than the cost of sending FDI. Possible explanations are the generally modest amounts remitted and the market power of transfer operators. Counterfactual exercises predict that in the absence of remittances, overall average welfare would be modestly lower, but with some substantial heterogeneity by country, depending on how reliant it is on remittances.

To illustrate the mechanism behind the risk-sharing properties of FDI and remittances, consider the following two-country example: Country A experiences a recession (GDP growth is negative, say because of a decrease in productivity), which makes it less attractive as a destination for FDI, as the decline in TFP leads to a decrease in returns to capital. At the same time, lower TFP, and potentially less capital, will lower the marginal product of labour, which is what workers are paid. Therefore, lower incomes in country A will induce migrants from country A in country B to send more remittances to help family members at home smooth consumption. Thanks to migrants' remittances, demand and thus production do not drop by as much as they would have without, and FDI potentially drops by less than in the absence of remittances. Thus, whereas FDI exacerbates asymmetric shocks, remittances help smooth consumption and attenuate business cycle fluctuations. A second income smoothing channel works through capital markets: Investors from country A who have invested abroad continue to receive returns on those investments, which are not affected by the TFP shock in country A, or at least not to the same degree.

The main contribution of this project is to jointly model remittances and returns to FDI as a way of sharing idiosyncratic macroeconomic risk across countries. There have been recent interesting considerations of the risk-sharing aspects of both migration (Kohler

[2]The countries are chosen based on availability of bilateral FDI data. They comprise the Czech Republic, France, Germany, Greece, Hungary, Italy, the Netherlands, Poland, Spain, Turkey, the United Kingdom, and the United States.

et al., 2021; Franceschin and Görlach, 2023; House et al., 2018), through migration itself or remittances, and capital markets (Cimadomo et al., 2022; Dufrénot et al., 2021; Furceri et al., 2013), as returns on investment abroad tend to be less correlated with the domestic business cycle than returns on domestic capital investments. Bringing the two together allows us to evaluate their relative effectiveness in smoothing income (and thus consumption) over the business cycle, and potential interactions between the two flows. With the help of our model, we are able to take a stance on whether migration and the ensuing remittances and FDI are (partial) substitutes in cross-country risk sharing.

Our findings also relate to the literature on the direction of causality in the relationship between GDP growth and remittances and FDI. Most papers try to estimate the effects of remittances or FDI on growth (Barajas et al., 2009; Berument and Dinçer, 2004; Clemens and McKenzie, 2018; Chatterjee and Naknoi, 2010).[3] Instead, we motivate our model by showing that growth as a key macroeconomic condition is also a driver of the two capital flows.[4]

The remainder of this paper is structured as follows: In the first part of the paper, we document two empirical regularities in Section 3.2, which we then use to inform our model (Section 3.3). Following a description of our estimation strategy in Section 3.4, we present the preliminary results and discuss them in Sections 3.5 and 3.6. Section 3.7 concludes.

3.2 Empirical regularities

In this section, we describe the key empirical relationships that motivate our research question and help inform our modelling choices.

3.2.1 Data

The descriptive section is based on micro data panels. The reason for using different data for the empirical analysis is that these allow us to control for individual characteristics

[3]The conclusions regarding the effect of remittances and FDI on growth seem to be context-dependent. Barajas et al. (2009) find no impact of remittances on growth whereas Clemens and McKenzie (2018) suggest positive but hard-to-measure effects of remittances on development; Berument and Dinçer (2004) find that FDI increases growth in Turkey and Chatterjee and Naknoi (2010) also note positive but small effects of capital inflows on growth in most countries – possibly due to high investment prices.

[4]Mistura and Roulet (2019)'s focus is the effect of regulatory burdens on FDI, but their inclusion of five-year average real GDP growth in the destination country also finds that FDI is positively affected by growth.

and thus establish a more robust relationship between growth and capital flows. Below we outline the main datasets used in this paper.[5]

Remittances For the reduced-form evidence on remittances, we use micro data from the German Socio-Economic Panel (SOEP) for the years 1984 to 2019, except for 1992 and 1994, for which the relevant question was not part of the survey. Our definition of migrants is individuals who respond having a country of origin other than Germany. We further restrict the sample to individuals aged 25-64. Among this age group, the share of migrants in the SOEP for the sample period averages 22.8%. The main variable for remittances is payments to family members as this is consistently asked for throughout the survey years (with the exceptions of 1992 and 1994). Combining this information with the follow-up question on whether the family members live abroad, we create a variable for remittances, assuming that the family members abroad reside in the migrant's country of origin.

FDI The Microdatabase Direct Investment (MiDi) captures all inward and outward foreign direct investment to or from Germany between 1999 and 2020 that is above the reporting threshold specified in the German Foreign Trade and Payments Regulation (Außenwirtschaftsverordnung). It contains identifiers and country of origin for the direct investment enterprise (DIE), the direct investor, and, where different from the direct investor, the ultimate controlling institution. In particular, these identifiers allow to control for direct investor and DIE specific characteristics that are constant across years. The dataset is a full panel of investments underlying the reporting requirements between 1999 and 2020, with observations at the investor-DIE-year level. The key outcomes of interest are primary and secondary outward FDI.[6]

Other data The other key variable in our empirical analysis is GDP growth, which we take from the World Development Indicators (WDI). A detailed list of other data used in the model can be found in Table D3.1 in the Appendix.

Throughout, all monetary values are in Euro; if they were originally in USD, they were transformed into Euro at current exchange rates and deflated to 2015. Unless otherwise stated, all values used in the model are averages for the years 2015-2018.

[5] More information is included in Appendix D.
[6] Details can be found in Blank et al. (2020).

3.2.2 Big picture

Our detailed micro-level data for Germany show that both FDI and remittances are correlated with GDP growth in the recipient economy (Figure 3.1). Specifically, higher GDP growth in migrants'[7] country of origin entails *lower* amounts remitted (Panel a), whereas higher GDP growth in the investment destination is *positively correlated* with FDI inflows (Panel b). Figure 3.1 shows these relationships as a pair of binned scatter plots of the residuals of remittances (FDI) after taking out individual (country) and year effects in 20 bins of equal mass of residuals of lagged GDP growth, also net of destination country and year effects. For comparability, both outcome variables have been transformed using the inverse hyperbolic sine transformation. This implies that the axes can be loosely interpreted in percentage deviations (e.g a 5 percentage point increase in GDP growth is associated with about 5% lower remittances).

Figure 3.1: Correlation of remittances and FDI with receiving country GDP growth

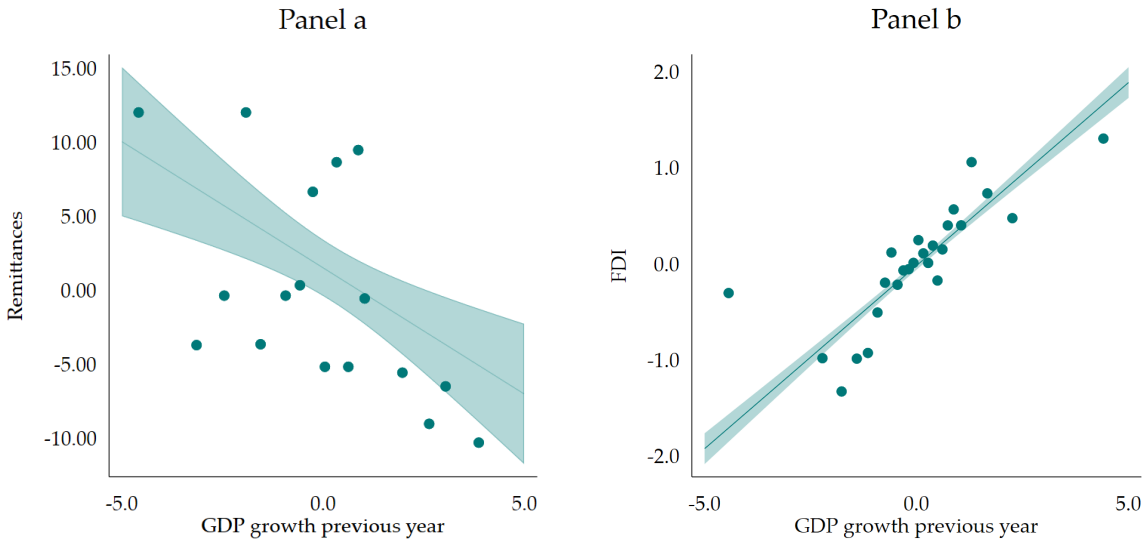

Note. Panel a: Change in remittances in percent (1984-2019, except for 1992 and 1994). Panel b: Residual primary and secondary FDI against residual GDP growth (both net of country and year fixed effects). Change in outward FDI from Germany to destination countries in percent (1999-2020). GDP growth refers to growth in the recipient economy, is denoted in percentage points, and lagged by one period with respect to remittances / FDI. Remittances are net of individual and year fixed effects, and FDI is net of year and country fixed effects. Data are transformed using the inverse hyperbolic sine transformation. *Sources*: SOEP (remittances), Bundesbank (FDI), and WDI (growth).

The data on remittances are a taken from a migrant-year panel based on the German Socio-Economic Panel (SOEP) from 1984 to 2019 while the FDI data are firm-level data

[7] Throughout, we define migrants as foreign-born individuals.

collected by the German Bundesbank and show outward FDI flows from Germany to destination countries for the years 1999-2020. Both panels in the figure are constructed based on microdata. However, the same pattern emerges when using aggregate FDI outflows from Germany to 161 countries for which data are available (Figure B3.2). Table A3.1 in the Appendix shows summary statistics for the SOEP data used for our descriptive statistics on remittances and Table A3.2 does so for the Bundesbank firm-level FDI data. These are the data used in Figure 3.1.

3.2.3 Remittances and growth

To more formally document the negative correlation between remittances and GDP growth in migrants' countries of origin, we estimate the following regression:

$$\text{remittances}_{ct} = \beta \text{growth}_{c,t-1} + \theta_c + \mu_t + \varepsilon_{ct} \tag{3.1}$$

where remittances_{ct} are remittances and $\text{growth}_{c,t-1}$ is lagged GDP growth. We control for country of origin and time fixed effects in all specifications, and in some of them additionally for other characteristics.

Table 3.1 shows that GDP growth in migrants' country of origin lowers the remittances they send from Germany, with a 1 percentage point increase in growth being associated with a 1% decrease in remittances. This relationship is robust to the inclusion of controls for education, employment status, gender, earnings, migration duration (using a full set of annual indicators), or immigration year, as well as individual fixed effects. The remittances regressions cover $10,648$ unique individuals and 129 countries of origin. Migrants who do not send remittances in a given year have a zero entry for the value of remittances.

The negative correlation between remittances and GDP growth in migrants' country of origin remains the same when adjusting remittances for the price level of migrants' home countries, which more accurately reflects their purchasing power in the country to which they are sent. If anything, coefficients are slightly larger (Table A3.4). The results also hold when adjusting standard errors for potential serial correlation between two and four years, both for the specifications in Table 3.1 and those in Table A3.4 (not shown). All regressions control for country of origin and year fixed effects.

Of course, there are factors other than growth that affect remittances. Chief among them is how much individuals earn: In Table A3.5, we show that remittances are strongly increasing in earnings.

Table 3.1: Percentage change in remittances in response to a 1 percentage point increase in growth in migrants' country of origin

	Remittances							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged GDP growth	-1.079	-1.050	-1.135	-1.004	-1.131	-1.068	-1.487	-1.096
	(0.372)	(0.315)	(0.326)	(0.370)	(0.357)	(0.373)	(0.463)	(0.337)
	[0.004]	[0.001]	[0.001]	[0.008]	[0.002]	[0.005]	[0.002]	[0.001]
Constant	126.464	128.320	128.600	164.439	62.899	231.550	137.374	126.781
	(1.226)	(1.042)	(1.079)	(72.774)	(6.497)	(19.800)	(9.917)	(1.119)
	[0.000]	[0.000]	[0.000]	[0.026]	[0.000]	[0.000]	[0.000]	[0.000]
Observations	46944	45017	45017	46010	46850	46944	28978	45298
R^2	0.064	0.067	0.072	0.072	0.097	0.079	0.096	0.529
Country of origin	✓	✓	✓	✓	✓	✓	✓	
Year	✓	✓	✓	✓	✓	✓	✓	✓
Migration duration		✓						
Immigration year			✓					
Education				✓				
Employment status					✓		✓	
Gender						✓		
Earnings							✓	
Individual								✓

Note. Remittances from migrants in Germany to their countries of origin for 1984-2019, except for 1992 and 1994. GDP growth refers to growth in the recipient economy, is denoted in percentage points, and lagged by one period with respect to remittances. Remittances are transformed using the inverse hyperbolic sine transformation, and multiplied by 100 for readability. Education is individuals' highest level of education as per the 1997 International Standard Classification of Education. Earnings refers to net labour earnings. All regressions include country of origin and year fixed effects. Standard errors are clustered at the country of origin and reported in parenthesis; p-values are reported in brackets. *Sources*: WDI (growth) and SOEP (all other variables).

Table 3.2: Propensity to remit and remittances in levels

	Propensity to remit		**Levels**	
	(1)	(2)	(3)	(4)
Lagged GDP growth	-0.149	-0.142	-2.254	-2.282
	(0.050)	(0.043)	(1.234)	(1.761)
	[0.003]	[0.001]	[0.070]	[0.197]
Constant	16.513	16.513	304.375	305.355
	(0.163)	(0.141)	(4.069)	(5.843)
	[0.000]	[0.000]	[0.000]	[0.000]
Observations	46944	45298	46944	45298
R^2	0.067	0.509	0.028	0.468
Country of origin	✓		✓	
Individual		✓		✓
Year	✓	✓	✓	✓

Note. Remittances from migrants in Germany to their countries of origin for 1984-2019, except for 1992 and 1994. GDP growth refers to growth in the recipient economy, is denoted in percentage points, and lagged by one period with respect to remittances. The propensity to remit is calculated using a linear probability model and multiplied by 100 for readability. All regressions include country and year fixed effects. Standard errors are clustered at the country of origin and reported in parenthesis; p-values are reported in brackets. *Sources*: SOEP (remittances) and WDI (growth).

While the average effects in Table 3.1 inform about the aggregate implications of growth for remittance flows, there is substantial heterogeneity in migrants' remitting behaviour. Only 20% of migrants in the SOEP sample send remittances in any given year (Table A3.1). This implies that an important reaction to growth shocks in their countries of origin may not only be whether to remit more or less, but whether to remit at all. Therefore, we also analyse the extensive margin of remittances: In the first two columns of Table 3.2, we show migrants' propensity to remit, estimated as a linear probability model of whether they remit anything or not. And indeed, higher growth in migrants' country of origin reduces the share of migrants who report sending remittances by 0.14 to 0.15 percentage points for every percentage point increase in growth.

Columns (3) and (4) show the coefficients for regressions of remittance *levels* on growth, and that these two are also negatively correlated (following a 1 percentage point increase in growth, average yearly remittances decrease by 2.25 Euro). Coefficients for both outcomes of interest remain very stable in magnitude when including individual fixed effects, though the results on levels are not particularly statistically significant.

Table 3.3: Share of earnings remitted

	Share of GDP per capita remitted			
	(1)	(2)	(3)	(4)
Lagged GDP growth receiving country	-0.205	-0.204		
	(0.063)	(0.063)		
	[0.001]	[0.001]		
Lagged GDP growth per capita receiving country			-0.154	-0.153
			(0.062)	(0.062)
			[0.013]	[0.014]
Constant	10.625	10.620	10.267	10.267
	(1.111)	(0.194)	(1.059)	(0.110)
	[0.000]	[0.000]	[0.000]	[0.000]
Mean (dep.var.)	9.993	9.993	9.993	9.993
Observations	42,386	42,386	42,386	42,386
Remittance sending country		✓		✓

Note. The dependent variable is the share of GDP per capita in the destination country that is remitted per migrant. It is multiplied by 100 for readability. We use GDP per capita as a proxy for average income. GDP growth refers to growth in the recipient economy, is denoted in percentage points, and lagged by one period with respect to remittances. The dataset includes information on 214 sending and 205 receiving countries, and one observation is a pairwise combination of the two. Standard errors are clustered by sending country and reported in parenthesis; p-values are reported in brackets. *Source*: World Bank Bilateral Remittances and Bilateral Migration matrices (remittances) and WDI (GDP per capita and growth), all for 2017.

Another way of capturing the relationship between remittances and growth is through the share of their earnings that migrants remit. This is also one of the two key moments targeted in our model. To be consistent with the data used in the model, we document this relationship using data from the World Bank's Bilateral Remittances matrix. Ideally, we would like to estimate the relationship between growth in migrants' country of origin and the share of their earnings that they remit. Absent migrants' earnings data, we approximate earnings with the GDP per capita in their host country.[8] Table 3.3 shows that migrants remit on average 10 percent of the GDP per capita, and that this share is reduced when their country of origin grows more. Each percentage point of growth in their home country is associated with a decrease in that share by 0.2 percentage points. The relationship becomes slightly weaker when using growth in GDP per capita instead of standard GDP growth (0.16 percentage points), but is relatively

[8]These numbers are about 2.5 times the size of the share of earnings that migrants in the SOEP report as remittances, which could be driven by different composition in terms of countries and time period. We show the relationship for the World Bank's Bilateral Remittances data as these are the data used in the model.

stable in size and robust to controlling for remittance sending country fixed effects. A drawback to showing the relationship between share of earnings remitted and growth on the data that is used for the model is that the aggregate nature of the bilateral remittances data does not allow us to look at effects on the intensive vs. extensive margin of remittance sending behaviour separately. Moreover, given that the dataset is a cross section, we cannot include fixed effects for migrants' countries of origin.

3.2.4 FDI and growth

Table 3.4: Regression of Primary and Secondary FDI (only non-zero values) on GDP growth in the previous year

	Primary and secondary FDI					
GDP growth in the investment destination (previous year)	0.252***	0.248***	0.232***	0.218***	0.247***	0.246***
	(0.047)	(0.046)	(0.034)	(0.037)	(0.052)	(0.053)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Constant	2.173***	2.171***	2.133***	2.148***	2.168***	2.160***
	(0.125)	(0.121)	(0.092)	(0.100)	(0.138)	(0.140)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Observations	421771	420559	360519	355917	412934	409877
R^2	0.044	0.095	0.229	0.293	0.192	0.198
Country & year	✓	✓			✓	✓
Investor		✓			✓	
Country & investor × year			✓			
Investor × year & investor × country				✓		
DIE					✓	
Investor × DIE						✓

Note. GDP growth refers to growth in the recipient economy and is denoted in %. *Sources*: FDI: Research Data and Service Centre (RDSC) of the Deutsche Bundesbank, MiDi, 1999-2020, own calculations. GDP growth: WDI.

Analogously to the specification for remittances, we regress FDI flows on lagged GDP growth in the destination countries:

$$\text{FDI}_{ct} = \beta \text{growth}_{c,t-1} + \theta_c + \mu_t + \varepsilon_{ct} \tag{3.2}$$

where FDI_{ct} are FDI flows and $\text{growth}_{c,t-1}$ is lagged GDP growth. We control for country of origin and time fixed effects in all specifications, and in some of them additionally for other characteristics.

In contrast to remittances, FDI is positively correlated with growth. For our analysis of this relationship we rely on administrative investor-firm level data on outward FDI from German investors to over 90,000 direct investment enterprises between 1999 and 2020. The data cover all transactions to which the reporting requirements apply, and compliance in reporting is almost perfect. The detailed nature of the data allows us to establish a robust relationship between growth and FDI similar to the one presented for growth and remittances (Table 3.4).[9] All regressions include investment destination and year fixed effects. The results in Table A3.6 exclude investor-firm-year observations with zero FDI flows. The results are robust to including these zero values (Table A3.6). When analysing the correlation between GDP growth in destination countries and different components of FDI, the results appear to be most strongly driven by equity investments (Table A3.7).

The take-away from this section is that remittances are countercyclical and FDI is procyclical with respect to the receiving country's GDP growth. This empirical pattern is robust to different specifications and the inclusion of controls. In the next section, we outline a simple model that draws on the empirical facts described so far. In particular, we model remittances as decreasing in origin country GDP. Theoretically, migrants could invest more in their home country as it grows faster – yet, this is not what we observe in the data. Even though some remittances are potentially driven by an investment motive and could behave procyclically, they are quantitatively outweighed by remittances in line with an altruistic consumption smoothing motive, so that in aggregate the observed remittances flows are countercyclical to growth in the home country.

3.3 Model setup

We build a simple equilibrium model of the economy that explicitly incorporates both FDI and remittances, taking the stocks of capital and migrants as given (i.e. there is no migration response following growth shocks). Each country produces one good, and

[9] However, even aggregate data show a positive relationship between FDI and GDP growth, but with a correlation about half as strong as the one that can be inferred from the analyses using investor-firm level data.

its production can be modelled by a representative firm that produces with the inputs (capital and labour) available in the economy but does not take any active decisions. The key agents in the model are investors, who decide how much to invest in each country, migrants, who choose how much to remit, and workers, who decide which goods to consume.

Notation Subscripts indicate the direction of flows of capital or goods: in subscript $_{ij}$, i denotes the origin and j the destination. For example, C_{ij} are goods produced in country i to be consumed in country j and K_{ij} is capital owned by investors from country i that is invested in country j as FDI.

Superscripts stand for the type of individual: K for capital owner (investor), R for remittance agent, N for native, and M for migrant.

3.3.1 Producers

Firms produce output according to a standard constant returns to scale Cobb-Douglas production function. They employ both native and migrant workers and rent capital from domestic and/or international investors. Capital inputs are perfectly substitutable for one another irrespective of the investor's origin. The production function for country i is:

$$P_i Y_i = P_i A_i K_i^{\alpha_i} L_i^{1-\alpha_i}$$

where

$K_i = K_{ii} + \sum_{j \neq i} K_{ij}$ is the capital stock available for production in country i

$L_i = \left(\delta (L_i^M)^\rho + (1-\delta)(L_i^N)^\rho\right)^{\frac{1}{\rho}}$ is effective labour, composed of immigrants L_i^M and natives L_i^N

$\frac{\delta}{1-\delta}$ is the productivity of migrants relative to natives

$\varepsilon = \frac{1}{1-\rho}$ is the elasticity of substitution between natives and migrants

A_i is total factor productivity in country i

and P_i is the price of the good produced by country i

The two types of workers are assumed to be imperfectly substitutable[10] and supply

[10] $\varepsilon \neq 1$ implies imperfect substitutability of labour supplied by migrants and natives.

labour inelastically. Workers receive their marginal product as the wage, which differs between natives and migrants. Different countries are allowed to differ in their productivities and the shares of GDP going to capital and labour.

Assumptions:

1. Perfect competition
2. Capital is perfectly mobile across borders
3. Labour supply is inelastic
4. Total factor productivity is determined exogenously
5. Firms operate using a constant returns to scale production technology

3.3.2 Investors

Investors own capital and can rent it to domestic or foreign firms in the form of FDI. Renting out capital to foreign firms entails a cost proportional to the amount invested. This cost can be interpreted as encompassing additional risk and uncertainty, or capture other factors that make cross-border investment costly, such as learning about foreign markets or additional accounting requirements. It should lead to a home bias in investing.

Ultimately, investors want to maximise utility U_i^k, which they derive from consumption:[11]

$$\begin{aligned} \max U_i^K =& (C_{ii}^K)^{\beta_{ii}} \prod_{j \neq i} (C_{ji}^K)^{\beta_{ji}} \\ \text{s.t. } & \sum_i C_{ji}^K P_j = \underbrace{K_{ii} MPK_i + \sum_{i \neq j} K_{ij} \left(MPK_j - c_{FDI}\right)}_{\text{capinc}} \\ & \sum_j \beta_{ji} = 1 \end{aligned}$$

where β_{ji} is the share of income that investors from country i devote to goods produced in country j.

In a first step, they maximise their income, which can then be spent on consumption. When deciding how to optimally allocate their capital, they equate marginal returns

[11] We have chosen a Cobb-Douglas utility function because it allows for a clear correspondence between theoretical and data moments. Qualitatively, a CES utility function should yield the same results, but will require further assumptions about the elasticity of substitution between different goods.

net of costs across investment destinations.[12]

Investment

Then, the investment decision amounts to maximising income:

$$\max_{\{K_{i1},\ldots,K_{iN}\}} \text{capinc} = K_{ii}MPK_i + \sum_{j\neq i} K_{ij}(MPK_j - c_{FDI})$$
$$\text{s.t. } \bar{K}_i = \sum_j K_{ij}$$

where c_{FDI} is the cost of investing abroad per unit of capital and $MPK_j - c_{FDI}$ is the net marginal product of capital.

Solving shows that investors allocate their capital in such a way that their net marginal returns to capital across countries are equalised:

$$MPK_i = MPK_j - c_{FDI}.$$

Consumption

Given the capital income they receive from optimally allocating their investments, investors' optimal consumption of each good j is

$$C_{ji}^K = \beta_{ji}\frac{\text{capinc}}{P_j}.$$

3.3.3 Transfer operators

Transfer operators can be imagined as remittance agents who also earn the cost of FDI. Their income is determined by how much migrant workers choose to remit, how much investors invest abroad, and the going fees; their only choice is how to optimally consume their income:

$$\max U_i^R = (C_{ii}^R)^{\beta_{ii}}\Pi_{j\neq i}(C_{ji}^R)^{\beta_{ji}}$$
$$\text{s.t. } \sum_i C_{ji}^R P_j = \underbrace{\underbrace{\sum_j c_{ij}^R R_{ij}}_{\text{remittance fee income}} + \underbrace{\sum_{i\neq j} K_{ij}c_{FDI}}_{\text{FDI fee income}}}_{\text{fee income}}$$
$$\sum_j \beta_{ji} = 1$$

[12] Allowing only the investor type to make capital investments is a simplification also used, among others, by Kleinman et al. (2023).

Optimal consumption of each good j is

$$C_{ji}^R = \beta_{ji} \frac{\text{fee income}}{P_j}.$$

3.3.4 Workers

Workers supply labour inelastically and receive their marginal product of labour as the real wage: $\frac{w_i^M}{P_i} = MPL_i^M$ for migrants and $\frac{w_i^N}{P_i} = MPL_i^N$ for natives. Native workers – like investors and transfer operators – divide their income into consumption of domestically produced goods and imported goods.

In addition to the assumption of inelastic labour supply and workers being wage-takers, we assume that individuals cannot save or borrow, so all income is spent on consumption.

Natives

Native workers consume their earnings, and choose their consumption bundle as to maximise utility U_i^N.

Optimal consumption of each good j is

$$C_{ji}^N = \beta_{ji} \frac{w_i^N}{P_j}.$$

where β_{ji} is the share of income of a worker in country i spent on goods originating in country j.

Migrants

In addition to choosing their optimal consumption basket, migrants' optimisation problem also involves deciding on the amount of earnings to be remitted to the country of origin to support consumption there. Migrants' remittance behaviour is governed by an altruism parameter $1 - \gamma$.

The optimisation problem of a migrant working in country i from country of origin j is:

$$\begin{aligned} \max U_i^M = &\left[(C_{ii}^M)^{\beta_{ii}} \Pi_{j \neq i} (C_{ji}^M)^{\beta_{ji}}\right]^{\gamma} \left[(C_{jj})^{\beta_{jj}} \Pi_{k \neq j} (C_{kj})^{\beta_{kj}}\right]^{1-\gamma} \\ \text{s.t. } & C^M P = w_i^M - R(1 + c_{ij}^R) \\ & C^F P = w_j^N + R \\ & \sum_j \beta_{ji} = 1 \end{aligned}$$

A migrant from country j in country i values their own consumption and their family's consumption in the home country – with parameter γ determining how much migrants care about themselves relative to their family. Sending remittances comes at a cost that is proportional to the amount remitted, in line with empirical evidence. If $\gamma = 1$, migrants behave identically to natives and do not send remittances. This γ is one of the parameters we estimate.

We assume that migrants across all countries share the same utility parameter γ and that remittances cannot be negative (i.e. that migrants in the domestic economy do not receive remittances from their country of origin).

Remittances In a first step, we solve for the optimal amount remitted. Based on the disposable income net of remittances sent, individuals choose their consumption varieties.

Solving for the optimal amount to remit R yields the following expression:[13]

$$R_{ij}^{*} = \max\left\{0, \frac{1-\gamma}{1+c_{ij}^{R}} w_i^M - \gamma w_j^N\right\}$$

The optimal level of remittances thus increases in the wage rate of migrants (Table A3.5) and decreases in the wage prevailing in the country of origin, as suggested by the empirical evidence presented in Tables 3.1 and 3.3. Given different incomes across countries, migrants from different countries send different optimal remittances.

Consumption Having determined the optimal amount of remittances R_{ij}^{*}, migrants are left with $w_i^M - R_{ij}^{*}(1 + c_R)$ to spend on their own consumption.

The problem of how to optimally allocate disposable income to the consumption of different goods is the same as that of natives, with the modification in the budget constraint. So optimal consumption is given by:

$$C_{ji}^{M} = \beta_{ji} \frac{\tilde{w}_i^M}{P_j}$$

where $\tilde{w}_i^M \equiv w_i^M - R_{ij}^{*}(1 + c_R)$ are migrants' earnings net of remittances and associated costs.

[13] See Appendix C for derivation.

3.3.5 Market clearing

In equilibrium, markets clear: Goods are traded and consumed internationally, and trade is balanced. Capital and labour markets also clear.

Goods market clearing Demand for goods produced in country i must equal supply:

$$Y_i = \underbrace{(C_{ii}^N + C_{ii}^M + C_{ii}^K + C_{ii}^R)}_{C_{ii}} P_i + \sum_{j \neq i} C_{ij} P_i \quad \forall i, j$$

Capital market clearing No capital can be left idle: it is used either in domestic production or invested abroad as FDI:

$$\overline{K_i} = K_{ii} + \sum_{j \neq i} K_{ij} \quad \forall i, j$$

Labour market clearing Labour markets clear by assumption of inelastic labour supply:

$$L_i^M = \overline{L_i^M}, \quad L_i^N = \overline{L_i^N} \quad \forall i$$

3.3.6 Identification

To identify γ and c_{FDI}, the parameters we estimate, we target the bilateral net FDI stocks and the share of annual earnings remitted. The intuition behind the identification is the following: The higher γ in migrants' utility function, the more they care about their own consumption relative to their family's. So a higher share of earnings remitted implies a lower γ, meaning less weight on own consumption in the utility function. Regarding the second moment, the cost of FDI captures how far returns to capital are from being equalised across countries: If there was no cost to investing abroad, returns would be expected to be equal across all countries (if they were not, there would be profitable deviations to capital allocations).

3.3.7 Data

We use different datasets for the model estimation than for the regression-based analyses: For the cross-country model, we need data on multiple countries and use those provided by the World Bank and the OECD.

Remittances In the model, we use data on bilateral remittances from the World Bank's bilateral remittances matrix for 2017. These data describe aggregate bilateral remittance flows. Therefore, to obtain average remittances per migrant, we divide these numbers by the migrant stock, using data from the World Bank's bilateral migrant stock data for the same year.

FDI We use OECD data on bilateral FDI for 2015-2018 in the model. For 1985 to 2004, data are collected according to the 3rd benchmark definition of FDI (OECD, 1996); from 2005 to 2020, the 4th benchmark definition is available (OECD, 2008).[14] Given the often large discrepancies between reported inward and outward FDI, we use the average of the two stocks. We also show some empirical evidence using these aggregate data for 1985-2019.

3.4 Estimation

A number of parameters are taken directly from data or previous literature, or pre-estimated: We take the parameter for the cost of remittances from data on remittance prices (0.01, i.e. 1% of the amount remitted). We estimate the migrant wage penalty δ from the SOEP data on the sample of individuals aged $25-64$ who are working full time ($\delta = 0.4085$) and generalise from there as such data is not available for all countries included in our model. We borrow the elasticity of substitution between natives and migrants ρ from the literature (Brücker et al., 2014) ($\rho = 0.8510$).[15] Using data on GDP, capital stocks, and effective labour, we estimate productivity as a residual from a constant returns to scale Cobb-Douglas production function. All additional variables

[14] The respective core definitions are the following:

- BD3: "FDI is an international investment in which a resident investor acquires 10% or more of the ordinary shares or voting power of an enterprise in another economy ("lasting interest"). The definition covers subsidiaries, associates, and branches; FDI flows comprise equity, reinvested earnings and inter-company debt."
- BD4: "FDI is a cross-border investment made with the objective of establishing a lasting interest—evidenced when the investor owns at least 10% of the voting power of the enterprise abroad—and gaining a "significant degree of influence" over its management."

The differences between the two are small. The BD4 requires a more detailed breakdown of FDI into special purpose entities, mergers and acquisitions, greenfield investment, and ultimate investing country in addition to the immediate counterpart. BD4 also aligns fully with the IMF Sixth Balance of Payments Manual (BPM6) and the 2008 System of National Accounts, whereas BD3 was tied to BPM5 (1993).

[15] We use Brücker et al. (2014)'s estimate for Germany. Alternatively, we could have chosen the estimates by Ottaviano and Peri (2012) for the U.S. or Manacorda et al. (2012) for the U.K. They differ slightly: The estimate by Manacorda et al. (2012) is about a third higher than that reported by Ottaviano and Peri (2012). The estimate by Brücker et al. (2014) lies between the two and is slightly closer to the one by Ottaviano and Peri (2012).

used in the model and their sources are listed in Table D3.1. Throughout, the countries under consideration are Germany, the Czech Republic, Spain, France, the UK., Greece, Hungary, Italy, the Netherlands, Poland, Turkey, and the U.S. Our choice of countries is determined by availability of OECD data on bilateral FDI stocks.

We estimate the remaining parameter vector $\theta = [\gamma, c_{FDI}]$ using the Method of Moments. In estimating the model, we target two moments from the data: mean bilateral FDI across countries (to identify the cost of FDI), and the average share of income that migrants remit from their country of residence to their country of origin (by country of origin, to identify the utility function parameter γ). We estimate the parameters that give the model the best fit to the data, minimising the squared distance between data and model moments based on the following criterion:

$$crit(\theta) = (m_d - m_m(\theta))(m_d - m_m(\theta)),$$

where m_d are the moments observed in the data, and m_m are the moments generated by the model as a function of the parameter vector θ. We estimate two parameters based on two moments; hence, the current model is just identified. We choose the parameter combination that minimises the criterion function first using a grid search and then a solver. As our model predicts only (and at most) FDI from the country with relatively lower marginal product of capital to that with relatively higher marginal product of capital, and remittances from the country with higher wages to that in which wages are lower, we target net differences in the data, i.e. the model predicts remittance transfers only in one direction.

The model tries to find an allocation that results in an equilibrium in both the goods and capital market (labour market clearing is guaranteed by the assumption of inelastic labour supply). We iterate over prices for the goods market to clear, and inside that FDI adjusts to equate the net marginal value products of capital across countries. The returns to FDI depend on prices, so the two markets have to clear jointly.

3.5 Results

3.5.1 Estimated parameters

The cost of FDI is modelled as a share of FDI invested abroad, i.e. our estimates suggest that it amounts to 0.603% of the gross value invested abroad. This is slightly smaller than the cost of remittances at 1%, which we estimate directly from the data.

The parameter estimate for γ is 0.5276, implying that – in terms of consumption –

migrants care almost as much about their family than about themselves. This might also be explained by the fact that one migrant remitting might benefit multiple family members in their country of origin.

Table 3.5: Parameter estimates

Parameters	Point estimates
$\mathbf{c_{FDI}}$	0.00603
$\boldsymbol{\gamma}$	0.5276

Note. Point estimates of parameters.

3.5.2 Model fit

The two model moments fit the data very well: Table A3.8 summarises the model's fit as both the relative and absolute (in Euro) deviation of the targeted moments in the model from those in the data. It fits the data especially well for the share of earnings remitted, which identifies the parameter γ.

3.5.3 Counterfactuals

Welfare gains from remittances and FDI What are the welfare effects of migrant remittances and FDI? In order to compute these effects, we shut down FDI and remittances (i.e. raise the cost of FDI or remittances to 1 such that zero FDI or remittances arrives at the destination as 100% is charged in fees), so that the model generates no FDI or no remittances (or neither), while keeping all other inputs and parameters unchanged. By construction, aggregate utility decreases relative to the baseline in all three scenarios.[16]

Table 3.6: Welfare gains from remittances and FDI

	No FDI	No remittances	Neither
Losses relative to baseline utility	0.0158	0.1786	0.1922

Note. Utilities relative to model-generated utility under the baseline model in % deviation. All individuals are given equal weight in the calculations.

Overall, the welfare consequences of remittances appear an order of magnitude larger

[16] We compute utility in a utilitarian sense by weighting each individual equally.

than those of FDI. The welfare loss in a world without remittances is 0.1786% compared to one with remittances. This relatively low average welfare loss masks substantial heterogeneity by country, with some poorer remittance receiving countries such as Poland losing as much as 5% of their model-generated welfare while the U.S. sees no change in welfare (see Figures B3.3 and B3.4).

Figure 3.2: Welfare losses from no remittances or FDI

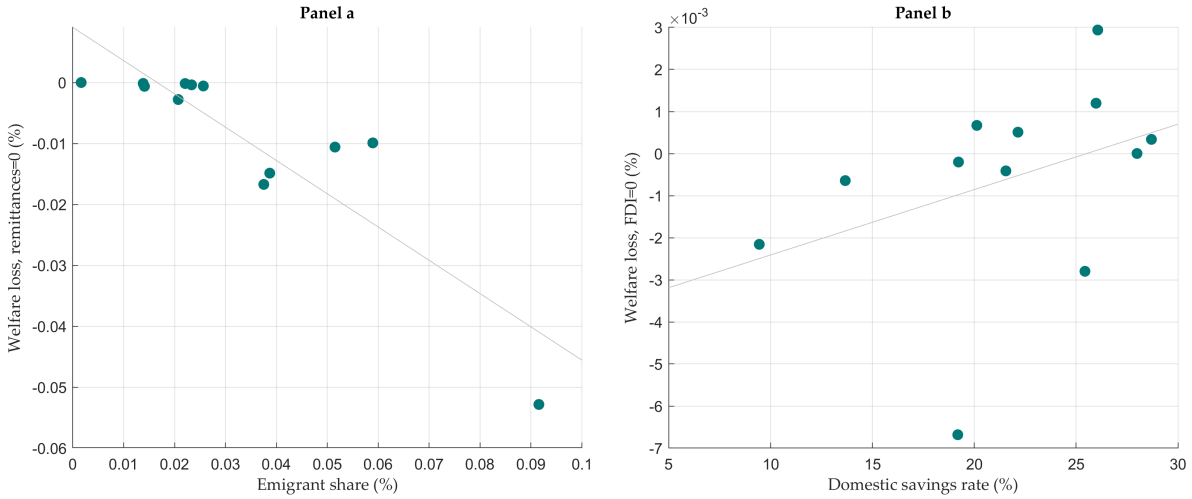

Note. Utilities under a no-remittance scenario against the emigrant share (panel a) and utilities under a no-FDI scenario against the domestic savings (panel b) relative to model-generated utility under the baseline model in % deviation. Each marker represents a country. All individuals are given equal weight in the calculations. *Sources*: UN bilateral migration matrix (emigrant share, 2017) and WDI (gross savings rate, average between 2015 and 2018).

In the scenario without remittances, the global welfare losses arise from the fact that consumption now cannot be transferred to individuals with lower incomes and higher marginal utilities of consumption. Countries that are home to remittance sending migrants see increases in domestic consumption, while remittance receiving countries consume less in the absence of remittances. Indirectly, remittance sending countries experience a decrease in demand for their exports because the formerly remittance-receiving countries can now spend less on export. This effect, however, is partly compensated by migrants' spending the wage they would have otherwise chosen to remit home, because remittances would not only be spent on goods from the country in which the migrant resides, but be spread across goods from all countries.

The welfare losses of not being open to FDI are more modest (at least in this static setting) at 0.0158% of status-quo welfare. Again, the welfare consequences are larger for some countries, such as the Czech Republic, Greece, and Poland.

The countries that experience the largest welfare losses when remittances or FDI are not available are also those that are most exposed to them: Panel a in Figure 3.2 plots the model-predicted welfare losses from not being able to send remittances for each

country against its emigrant share. As one would expect, the countries with higher emigrant shares see the largest welfare gains from remittances. Panel b instead shows each country's welfare loss from no FDI against its domestic savings rate. Countries with deeper domestic capital markets (here proxied by their savings rate) are less reliant on external financing.

The welfare effects of remittances and FDI work through different channels: Remittances benefit migrants because they channel funds to their families, for whom incomes are lower and thus marginal utility of consumption is higher. FDI is welfare enhancing because it channels funds to where they are most productive, yielding higher returns to capital for investors but also resulting in higher wages because workers' marginal product of labour increases.

Table 3.7: Consumption changes from no remittances and/ or FDI

	No FDI	No remittances	Neither
Consumption change relative to baseline utility	0.053105	-0.015868	0.040667

Note. Aggregate consumption changes relative to model-generated consumption under the baseline model in % deviation.

Table 3.7 also shows the aggregate consumption changes under scenarios of no remittances or no FDI (or neither). Without remittances, consumption would be higher. The effects on overall consumption are different from those on overall utility. The reason for this is that altruistic migrants derive greater utility from being able to send remittances, but will experience lower consumption. A no-FDI world would slightly decrease overall consumption. This is due to capital no longer being allocated where it is most productive. Figures B3.5 and B3.6 show the consumption changes by country. Countries with a higher immigrant share see increases in national consumption in a no-remittance world as former remittances are now spent domestically. Conversely, countries with high emigrant shares – those who receive a lot of remittances – experience declines in consumption. Overall, the effects of bilateral FDI and remittances are relatively small for the countries considered, but confer substantial welfare effects to those who benefit from them.

3.6 Discussion

As noted in Lucas (1990) paradox, capital often does not flow from richer countries with higher capital-to-labour ratios to poorer ones with less capital per worker. Our

model can accommodate both possible explanations: better technology in the form of greater factor productivity in richer countries (higher A in our model), and higher costs of investing in poorer countries (higher c_{FDI}), e.g. due to greater informational costs or regulatory restrictiveness. To adequately capture the second explanation, it becomes crucial to make the cost of investing specific to the destination country (estimation in progress).

This model is static, and thus does not capture migration dynamics that could occur in response to economic growth. But what are the likely effects of greater migration on remittances? If (a lot) more would-be migrants decided to emigrate, one might expect higher aggregate flows of remittances, but lower per-migrant transfers. First, more individuals will now remit, which means that each migrant can send less (think two family members sending money home instead of previously one). And second, a smaller labour force in the home country should raise workers' marginal product and thereby their wages. This in turn decreases the wage gap between the origin and destination countries and ultimately remittances.

3.7 Conclusion

In this paper, we first document two empirical facts: that foreign direct investment – like investment more broadly – moves procyclically, and that remittances move countercyclically. We then build and estimate a joint model of FDI and remittances for 12 OECD countries and estimate the welfare and consumption effects of both capital flows.

Our findings speak to a broader macro-development debate: in an integrated world economy, private transfers by households and firms constitute risk-sharing mechanisms. We show that shutting down remittances would extinguish roughly an eighth of one per cent of aggregate welfare on average, an order of magnitude larger than the loss from eliminating FDI flows, and in heavily emigrant economies such as Poland the welfare loss would approach five per cent. These magnitudes put the stakes of seemingly technical policies that raise the cost of remittances or FDI – like setting ceilings on money-transfer fees or tightening outward-investment controls – on par with much-debated fiscal stabilisers.

Moreover, against this backdrop, proposals currently debated in the United States Congress to levy a federal remittance tax of between 1% and 3.5% – a measure intended primarily to curb informal migration – would directly erode the counter-cyclical buffer that remittances provide. Given that the United States accounts for roughly one-quarter of all global remittance outflows, the welfare costs would be substantial.

Bibliography

Akkoyunlu, Ş. and K. A. Kholodilin (2008): "A link between workers' remittances and business cycles in Germany and Turkey," *Emerging Markets Finance and Trade*, 44, 23–40.

Balli, F. and F. Rana (2015): "Determinants of risk sharing through remittances," *Journal of Banking & Finance*, 55, 107–116.

Barajas, A., M. T. Gapen, R. Chami, P. J. Montiel, and C. Fullenkamp (2009): "Do Workers' Remittances Promote Economic Growth?" *IMF Working Papers*, 2009.

Berument, H. and N. N. Dinçer (2004): "Do capital flows improve macroeconomic performance in emerging markets?: the Turkish experience," *Emerging Markets Finance and Trade*, 40, 20–32.

Blank, S., A. Lipponer, C.-J. Schild, and D. Scholz (2020): "Microdatabase Direct Investment (MiDi)–A full survey of German inward and outward investment," *German Economic Review*, 21, 273–311.

Brücker, H., A. Hauptmann, E. J. Jahn, and R. Upward (2014): "Migration and imperfect labor markets: Theory and cross-country evidence from Denmark, Germany and the UK," *European Economic Review*, 66, 205–225.

Chatterjee, S. and K. Naknoi (2010): "The Marginal Product of Capital, Capital Flows, and Convergence," *American Economic Review*, 100, 73–77.

Cimadomo, J., E. Gordo Mora, and A. A. Palazzo (2022): "Enhancing private and public risk sharing," *ECB Occasional Paper*.

Clemens, M. A. and D. McKenzie (2018): "Why don't remittances appear to affect growth?" *The Economic Journal*, 128, F179–F209.

Constantinescu, I. C. and M. Schiff (2014): "Remittances, FDI and ODA: Stability, cyclicality and stabilising impact in developing countries," *International Journal of Migration and Residential Mobility*, 1, 84–106.

Duarte, M. and D. Restuccia (2010): "The Role of the Structural Transformation in Aggregate Productivity*," *The Quarterly Journal of Economics*, 125, 129–173.

Dufrénot, G., J.-B. Gossé, and C. Clerc (2021): "Risk sharing in Europe: new empirical evidence on the capital markets channel," *Applied Economics*, 53, 262–276.

Franceschin, R. and S. Görlach (2023): "Asymmetric Shocks and Heterogeneous Worker Mobility in the Euro Zone," .

Frankel, J. (2011): "Are bilateral remittances countercyclical?" *Open Economies Review*, 22, 1–16.

Furceri, D., A. Zdzienicka, and O. J. Blanchard (2013): "The Euro Area Crisis: Need for a Supranational Fiscal Risk Sharing Mechanism? 1," *IMF Working Papers*, 2013.

House, C. L., C. Proebsting, and L. L. Tesar (2018): "Quantifying the Benefits of Labor Mobility in a Currency Union," Tech. rep., National Bureau of Economic Research.

Hsieh, C.-T. and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India*," *The Quarterly Journal of Economics*, 124, 1403–1448.

Kleinman, B., E. Liu, and S. J. Redding (2023): "Dynamic spatial general equilibrium," *Econometrica*, 91, 385–424.

Kohler, W. K., G. Müller, and S. Wellmann (2021): "Risk Sharing in Currency Unions: The Migration Channel," *CEPR Press Discussion Paper No. 16178*.

Lucas, R. E. (1990): "Why Doesn't Capital Flow from Rich to Poor Countries?" *The American Economic Review*, 80, 92–96.

Lueth, E. and M. Ruiz-Arranz (2007): "Are Workers' Remittances a Hedge Against Macroeconomic Shocks? the Case of Sri Lanka," *IMF Working Papers*, 2007.

Manacorda, M., A. Manning, and J. Wadsworth (2012): "The impact of immigration on the structure of wages: theory and evidence from Britain," *Journal of the European Economic Association*, 10, 120–151.

Mistura, F. and C. Roulet (2019): "The determinants of Foreign Direct Investment: Do statutory restrictions matter?" Tech. rep., OECD Publishing.

OECD (1996): "OECD Benchmark Definition of Foreign Direct Investment: Third Edition," .

——— (2008): "OECD Benchmark Definition of Foreign Direct Investment: Fourth Edition," .

Ottaviano, G. I. and G. Peri (2012): "Rethinking the Effect of Immigration on Wages," *Journal of the European Economic Association*, 10, 152–197.

Appendix

Appendix A: Additional tables

Table A3.1: Summary statistics of remittances

	Mean (1)	**Std. dev.** (2)	**1st pctl** (3)	**Median** (4)	**99th pctl** (5)	**N** (6)
Remittances sent to family members (yes/no)	0.20	0.40	0.00	0.00	1.00	63,068
Remittances to family members (amount)	328.11	1,291.06	0.00	0.00	5,113.00	61,600
Share of earnings remitted	0.04	0.87	0.00	0.00	0.50	38,378
Annual GDP growth (in percent)	3.36	5.33	-7.28	3.19	12.89	57,686
Earnings	16,056.32	12,090.20	1,800.00	13,800.00	60,000.00	39,205

Note. Remittances from migrants in Germany to their countries of origin in Euro (1984-2019, except for 1992 and 1994). GDP growth refers to growth in the recipient economy, is denoted in %, and lagged by one period with respect to remittances. *Sources*: SOEP (remittances), and WDI (growth).

Table A3.2: Summary statistics of German outward FDI

	Mean (1)	**Std. dev.** (2)	**1st pctl** (3)	**Median** (4)	**99th pctl** (5)	**N** (6)
Primary FDI	35458170	456600000	-2097000	3892000	542600000	562928
Primary and Secondary FDI	38763247	444400000	-15310000	5225000	643800000	513734
Primary FDI flow	1489625	152900000	-60390000	97000	85720960	457007
Primary and Secondary FDI flow	1737886	164100000	-69950000	196000	103500000	425549
GDP growth previous year	2.771	2.883	-5.281	2.483	10.636	667961

Note. GDP growth refers to growth in the recipient economy and is denoted in %. There are 18917 unique investors, 90716 direct investment enterprises, and 112215 investor-DIE pairs in the period from 1999 to 2020. *Sources*: FDI: Research Data and Service Centre (RDSC) of the Deutsche Bundesbank, MiDi, 1999-2020, own calculations. GDP growth: WDI.

Table A3.3: Summary statistics of aggregate FDI (OECD)

	Mean (1)	**Std. dev.** (2)	**1st pctl** (3)	**Median** (4)	**99th pctl** (5)	**N** (6)
FDI flow	645.64	3,295.73	-3,144.60	18.95	13,167.98	2,668
FDI net position	10,363.63	26,278.85	0.00	1,139.02	150959.37	1,978
Annual GDP growth (in percent)	3.72	5.31	-10.00	3.66	14.72	2,668

Note. Bilateral FDI between Germany and 161 countries in millions of Euro (1984-2021). GDP growth refers to growth in the destination economy, is denoted in %, and lagged by one period with respect to FDI. *Sources*: OECD (FDI), and WDI (growth).

Table A3.4: Growth in migrants' country of origin lowers their remittances (valued at PPP)

		Remittances							
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged growth	GDP	-1.608	-1.571	-1.674	-1.495	-1.646	-1.593	-1.975	-1.773
		(0.375)	(0.342)	(0.360)	(0.400)	(0.354)	(0.375)	(0.511)	(0.628)
		[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.006]
Constant		149.201	151.244	151.586	198.401	76.367	269.752	171.375	150.242
		(1.240)	(1.134)	(1.195)	(85.682)	(5.512)	(17.014)	(12.164)	(2.084)
		[0.000]	[0.000]	[0.000]	[0.022]	[0.000]	[0.000]	[0.000]	[0.000]
Observations		46453	44537	44537	45559	46359	46453	28758	44847
R^2		0.063	0.065	0.070	0.068	0.092	0.075	0.087	0.519
Country of origin		✓	✓	✓	✓	✓	✓	✓	
Year		✓	✓	✓	✓	✓	✓	✓	✓
Migration duration			✓						
Immigration year				✓					
Education					✓				
Employment status						✓		✓	
Gender							✓		
Earnings								✓	
Individual									✓

Note. Remittances from migrants in Germany to their countries of origin for 1984-2019, except for 1992 and 1994. GDP growth refers to growth in the recipient economy, is denoted in %, and lagged by one period with respect to remittances. Remittances are transformed using the inverse hyperbolic sine transformation, and multiplied by 100 for readability. Education is individuals' highest level of education as per the 1997 International Standard Classification of Education. Earnings refer to net earnings. All regressions include country of origin and year fixed effects. Standard errors are clustered at the country of origin and reported in parenthesis; p-values are reported in brackets. *Sources*: WDI (growth) and SOEP (all other variables).

Table A3.5: Migrants' remittances are increasing in their earnings and decreasing in the growth in their country of origin

	Remittances	
	(1)	(2)
Lagged GDP growth	-1.478	-1.154
	(0.463)	(0.448)
	[0.002]	[0.011]
Log earnings previous year	53.525	22.613
	(12.092)	(10.593)
	[0.000]	[0.035]
Constant	-341.597	-45.564
	(117.275)	(101.812)
	[0.004]	[0.655]
Observations	28978	27496
R^2	0.096	0.572
Country of origin	✓	
Year	✓	✓
Employment status	✓	✓
Earnings	✓	✓
Individual		✓

Note. Remittances from migrants in Germany to their countries of origin for 1984-2019, except for 1992 and 1994. GDP growth refers to growth in the recipient economy, is denoted in %, and lagged by one period with respect to remittances. Remittances are transformed using the inverse hyperbolic sine transformation, and multiplied by 100 for readability. Earnings refer to net earnings. All regressions include year fixed effects. Standard errors are clustered at the country of origin and reported in parenthesis; p-values are reported in brackets. *Sources*: WDI (growth) and SOEP (all other variables).

Table A3.6: Regression of Primary and Secondary FDI (including zero values) on GDP growth in the previous year

	Primary and secondary FDI					
	(1)	(2)	(3)	(4)	(5)	(6)
GDP growth in the investment destination (previous year)	0.229***	0.225***	0.212***	0.204***	0.224***	0.225***
	(0.044)	(0.043)	(0.032)	(0.034)	(0.049)	(0.050)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Constant	1.674***	1.689***	1.640***	1.680***	1.726***	1.734***
	(0.121)	(0.117)	(0.089)	(0.095)	(0.134)	(0.136)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Observations	521088	519524	449695	443739	508314	504355
R^2	0.043	0.087	0.222	0.281	0.180	0.187
Country & year	✓	✓			✓	✓
Investor		✓			✓	
Country & investor × year			✓			
Investor × year & investor × country				✓		
DIE					✓	
Investor × DIE						✓

Note. *Sources*: FDI: Research Data and Service Centre (RDSC) of the Deutsche Bundesbank, MiDi, 1999-2020, own calculations. GDP growth: WDI.

Table A3.7: Regression of Primary and Secondary equity FDI on GDP growth in the previous year

	Primary and secondary equity FDI					
	(1)	(2)	(3)	(4)	(5)	(6)
GDP growth in the investment destination (previous year)	0.271***	0.261***	0.246***	0.220***	0.240***	0.235***
	(0.049)	(0.047)	(0.039)	(0.045)	(0.054)	(0.055)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Constant	3.046***	3.066***	3.074***	3.138***	3.147***	3.159***
	(0.130)	(0.126)	(0.105)	(0.119)	(0.144)	(0.145)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Observations	456966	455745	396447	391738	446386	442781
R^2	0.028	0.088	0.209	0.286	0.218	0.227
Country & year	✓	✓			✓	✓
Investor		✓			✓	
Country & investor × year			✓			
Investor × year & investor × country				✓		
DIE					✓	
Investor × DIE						✓

Note. *Sources*: FDI: Research Data and Service Centre (RDSC) of the Deutsche Bundesbank, MiDi, 1999-2020, own calculations. GDP growth: WDI.

Table A3.8: Model Fit

	Data moment	Model moment	Absolute deviation	Relative deviation
Average share of earnings remitted	0.108786751	0.108786328	−0.00000042	−0.0000039
Total bilateral FDI	10049984615	10047392195	−2592421	−0.00026

Note. Comparing data moments and model moments, as well as their absolute (in Euro) and relative deviations.

Appendix B: Additional figures

Figure B3.1: Outward FDI from Germany, 1999-2020

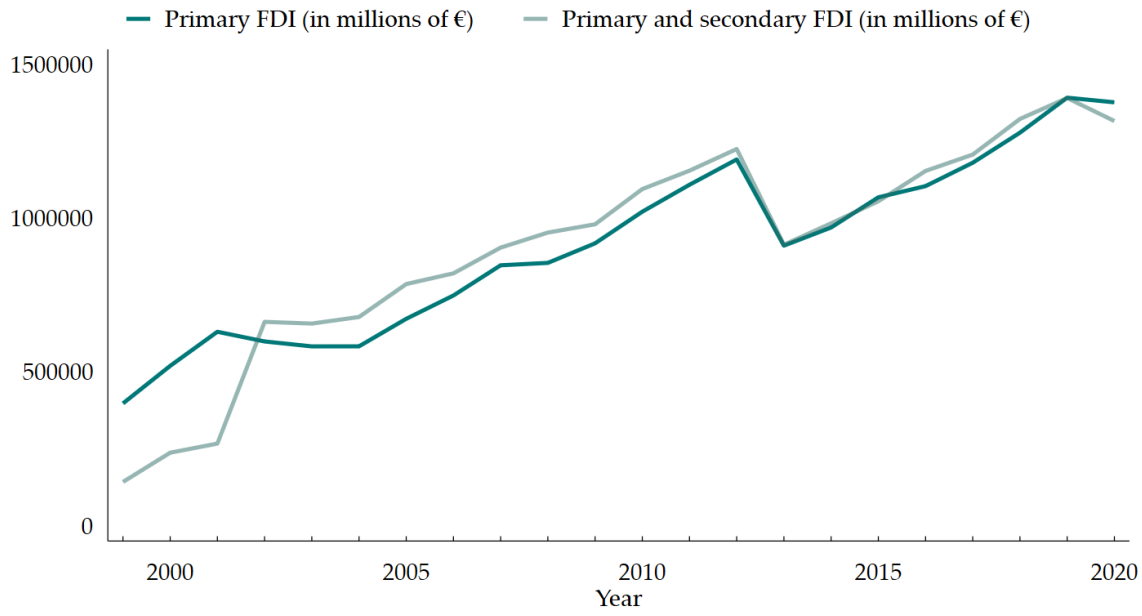

Note. Primary and secondary outward FDI from Germany between 1999 and 2020. *Source*: Research Data and Service Centre (RDSC) of the Deutsche Bundesbank, MiDi, 1999-2020, own calculations.

Figure B3.2: Correlation of FDI with destination country GDP growth

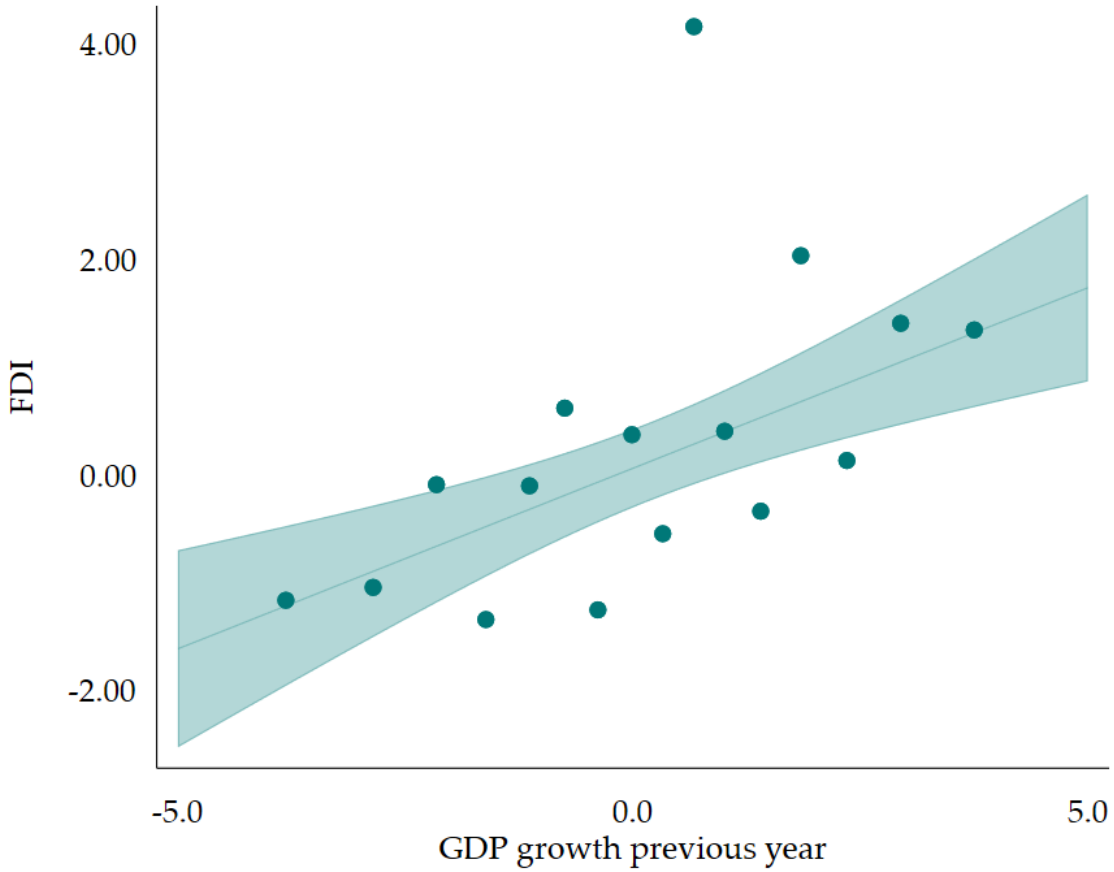

Note. Change in bilateral FDI between Germany and 161 countries in percent for the years 1985-2019. GDP growth refers to growth in the recipient economy, is denoted in percentage points, and lagged by one period with respect to remittances / FDI. FDI is net of year and country fixed effects. Data are transformed using the inverse hyperbolic sine transformation. *Sources*: OECD (FDI), and WDI (growth).

Figure B3.3: Welfare losses by country (I)

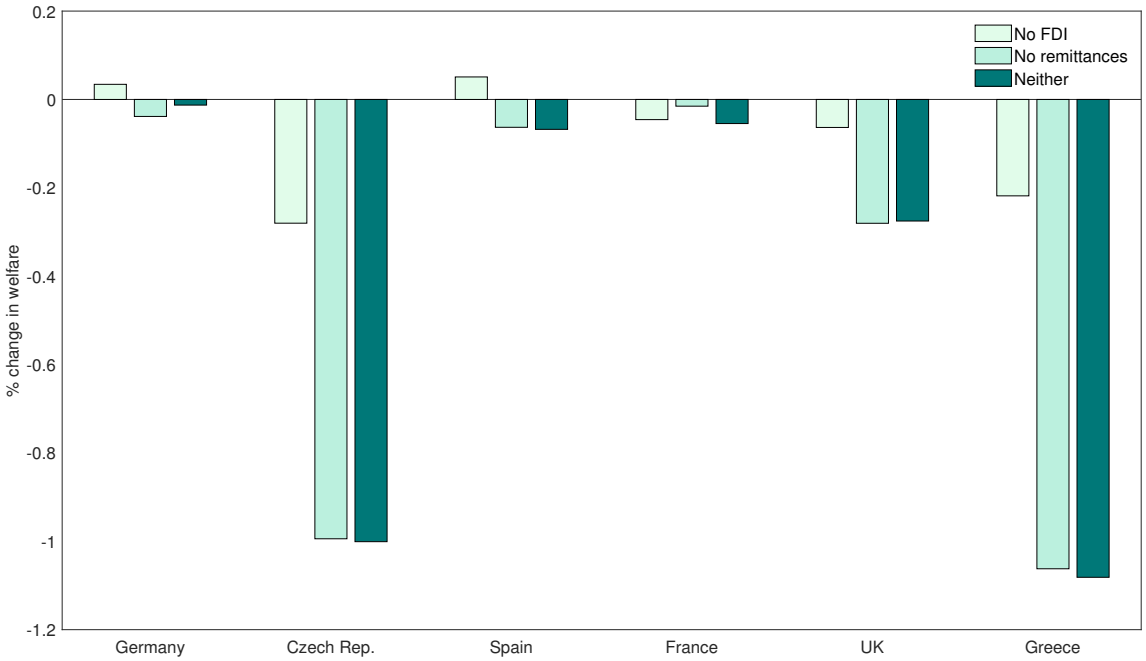

Note. Utilities by country, relative to model-generated utility under the baseline model in % deviation. All individuals are given equal weight in the calculations.

Figure B3.4: Welfare losses by country (II)

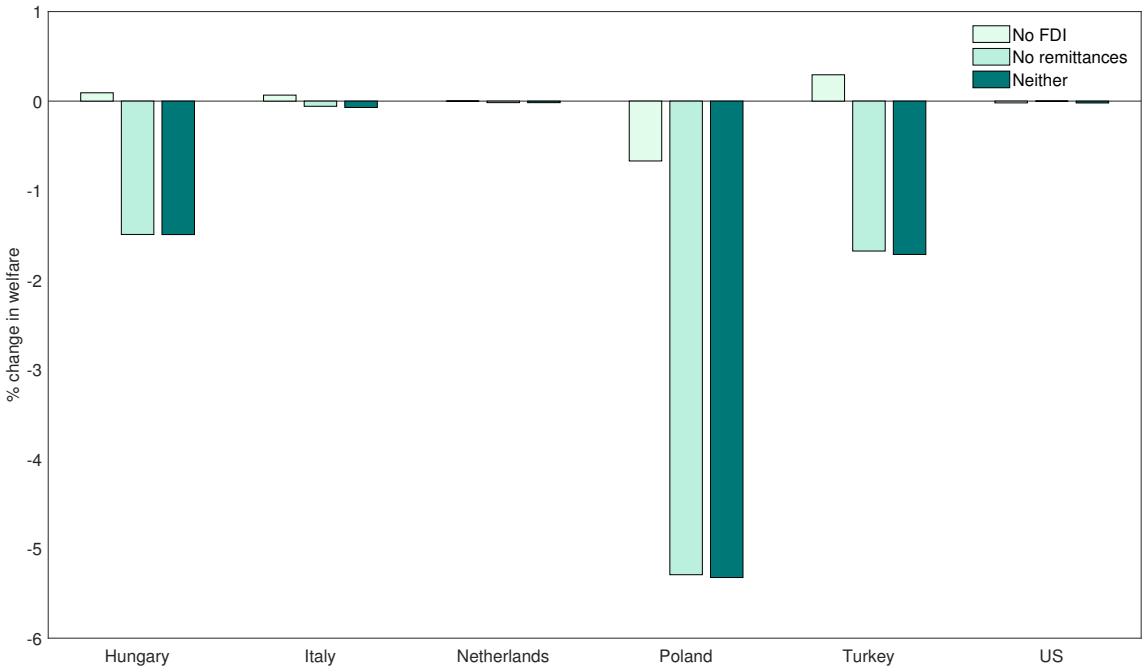

Note. Utilities by country, relative to model-generated utility under the baseline model in % deviation. All individuals are given equal weight in the calculations.

Figure B3.5: Consumption changes by country (I)

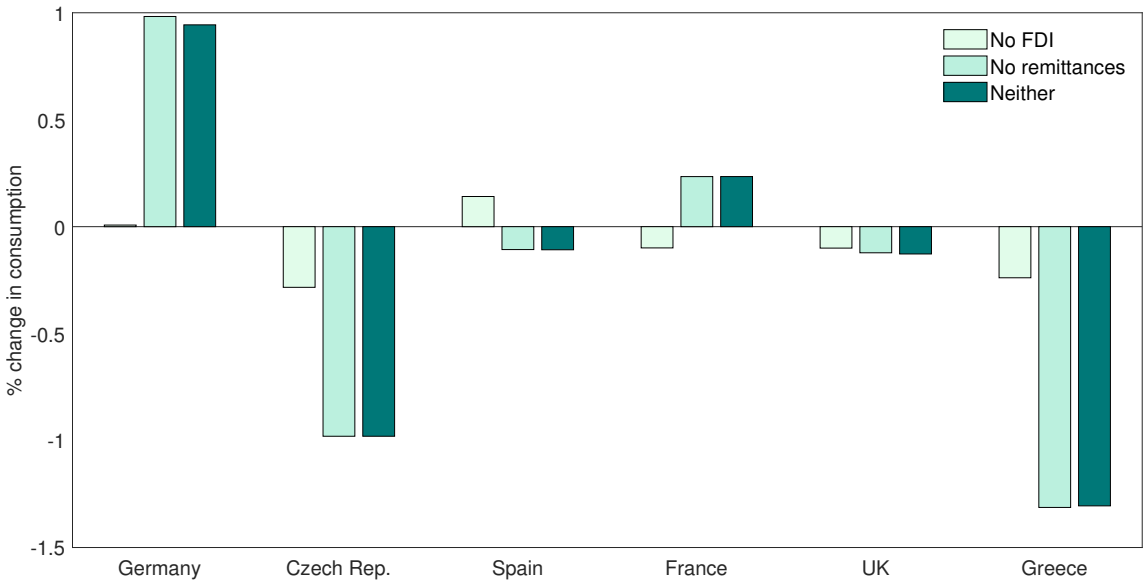

Note. Consumption changes by country, relative to model-generated consumption under the baseline model in % deviation.

Figure B3.6: Consumption changes by country (II)

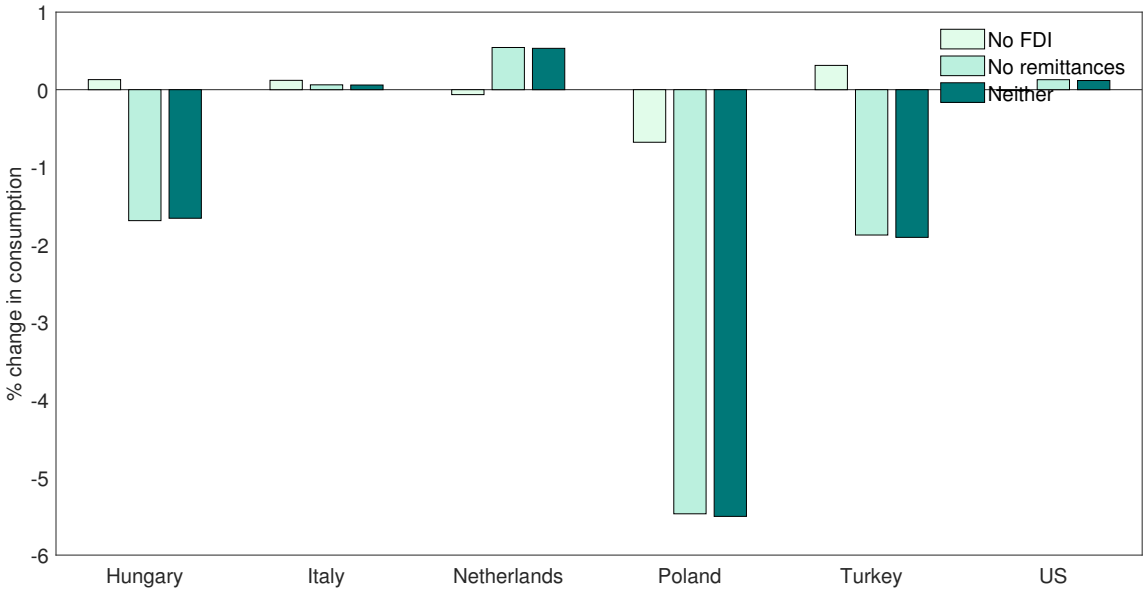

Note. Consumption changes by country, relative to model-generated consumption under the baseline model in % deviation.

Appendix C: Detailed model derivations

Investment

Investors maximise capital income net of costs:

$$\begin{aligned}
\max capinc =& \sum_j K_{ij} NMPK_j \quad \text{where } NMPK \text{ is net marginal product of capital} \\
\text{s.t. } & \bar{K}_i = \sum_j K_{ij} \\
\max capinc =& K_{ii} MPK_i + \sum_{j \neq i} K_{ij}(MPK_j - c_{FDI}) \\
\text{s.t. } & \bar{K}_i = \sum_j K_{ij} \\
\max capinc =& K_{ii} MPK_i + \sum_{j \neq i} K_{ij}(MPK_j - c_{FDI}) - \lambda \left[\sum_j K_{ij} - \bar{K}_i \right]
\end{aligned}$$

where c_{FDI} is the cost of investing abroad per unit of capital.

Investors allocate their capital in such a way that their net marginal returns to capital across countries are equalised.

$$MPK_i = MPK_j - c_{FDI}$$

Remittances

The optimisation problem of a migrant working in country i from country of origin j is:

$$\begin{aligned}
\max U_i^M =& \left[(C_{ii}^M)^{\beta_{ii}} \Pi_{j \neq i} (C_{ji}^M)^{\beta_{ji}} \right]^{\gamma} \left[(C_{jj})^{\beta_{jj}} \Pi_{k \neq j} (C_{kj})^{\beta_{kj}} \right]^{1-\gamma} \\
\text{s.t. } & C^M P = w_i^M - R(1 + c_{ij}^R) \\
& C^F P = w_j^N + R \\
& \sum_j \beta_{ji} = 1
\end{aligned}$$

In a first step, we solve for the optimal amount remitted. Based on the disposable income net of remittances sent, individuals choose their consumption varieties.

$$\begin{aligned}
\max U_i^M =& \left[(C_{ii}^M)^{\beta_{ii}} \Pi_{j \neq i} (C_{ji}^M)^{\beta_{ji}} \right]^{\gamma} \left[(C_{jj})^{\beta_{jj}} \Pi_{k \neq j} (C_{kj})^{\beta_{kj}} \right]^{1-\gamma} \\
& - \lambda \left[C^M P - w_i^M + R(1 + c_{ij}^R) \right] - \mu \left[C^F P - w_j^N - R \right] \\
\max U_i^M =& (C^M)^{\gamma} (C^F)^{1-\gamma} - \lambda \left[C^M P - w_i^M + R(1 + c_{ij}^R) \right] - \mu \left[C^F P - w_j^N - R \right]
\end{aligned}$$

An optimal solution must fulfill the following first-order conditions:

$$\frac{\partial}{\partial C^M} : \gamma (C^M)^{\gamma-1}(C^F)^{1-\gamma} - \lambda P = 0 \Longrightarrow \lambda = \frac{\gamma}{P}\left(\frac{C^F}{C^M}\right)^{1-\gamma}$$
$$\frac{\partial}{\partial C^F} : (1-\gamma)(C^M)^{\gamma}(C^F)^{-\gamma} - \mu P = 0 \Longrightarrow \mu = \frac{1-\gamma}{P}\left(\frac{C^M}{C^F}\right)^{\gamma}$$
$$\frac{\partial}{\partial R} : -\lambda(1+c_R) + \mu = 0 \Longrightarrow \frac{1-\gamma}{P}\left(\frac{C^M}{C^F}\right)^{\gamma} = \frac{\gamma}{P}\left(\frac{C^F}{C^M}\right)^{1-\gamma}(1+c_{ij}^R)$$

Solving for R yields the following expression for the optimal amount to remit:

$$(1-\gamma) = \gamma \frac{C^F}{C^M}(1+c_{ij}^R)$$
$$\frac{1-\gamma}{\gamma} = \frac{(w_j^N + R)(1+c_{ij}^R)}{w_i^M - R(1+c_{ij}^R)}$$
$$\frac{1-\gamma}{\gamma} w_i^M - (1+c_{ij}^R) w_j^N = R(1+c_{ij}^R)\frac{1}{\gamma}$$
$$\Longleftarrow R_{ij}^* = \max\left\{0, \frac{1-\gamma}{1+c_{ij}^R} w_i^M - \gamma w_j^N\right\}$$

Appendix D: Details on data

Remittances

Definition "Personal remittances comprise personal transfers and compensation of employees. Personal transfers consist of all current transfers in cash or in kind made or received by resident households to or from nonresident households. Personal transfers thus include all current transfers between resident and nonresident individuals. Compensation of employees refers to the income of border, seasonal, and other short-term workers who are employed in an economy where they are not resident and of residents employed by nonresident entities. Data are the sum of two items defined in the sixth edition of the IMF's Balance of Payments Manual: personal transfers and compensation of employees. Data are in current U.S. dollars" (World Bank).

Data used in the empirical analysis: German Socio-Economic Panel (SOEP) for 1984-2020. The sample of migrants consists of all foreign-born individuals.

Data used in the model estimation: Bilateral Remittances Matrix and Bilateral Migration Matrix, both for 2017. Average share of earnings remitted by migrants in and from the 12 countries for which there is bilateral FDI data.

FDI

Definition "Foreign direct investment reflects the objective of obtaining a lasting interest by a resident entity in one economy ("direct investor") in an entity resident in an economy other than that of the investor ("direct investment enterprise"). The lasting interest implies the existence of a long-term relationship between the direct investor and the enterprise and a significant degree of influence on the management of the enterprise. Direct investment involves both the initial transaction between the two entities and all subsequent capital and income transactions between them and among affiliated enterprises, both incorporated and unincorporated" (OECD).[17]

The OECD recommends reporting countries use a market value approach to valuing FDI stocks (positions). However, in practice, much reporting of FDI stocks is based on balance sheet values.

[17]The Bundesbank defines FDI as "cross-border investment in enterprises with the objective of establishing a lasting influence over business activities."

Data used in the empirical analysis: The MiDi data by the Bundesbank contains information on inward and outward foreign direct investments to and from Germany. As per the German Foreign Trade and Payments Regulation (Außenwirtschaftsverordnung) K3-Report, which includes information that characterises outwards FDI, firms with investments over 3 million € have been required to report on them yearly.[18] The data resulting from these reports are "balance sheet data of enterprises in which investors hold cross-border shares (in directly and indirectly held direct investment enterprises (DIE))" (Blank et al., 2020). We only use data for outward FDI from German investors to direct investment enterprises abroad.

The thresholds for reporting are based on the size of the direct investment enterprise and the share of the investment. Investors meeting the following criteria are required to submit a K3-Report are[19]

- German investors owning at least 10% of shares or voting rights in a foreign company with a balance sheet of at least 3 million €.
- German investors owning a branch / subsidiary abroad with at least 3 million € in assets.
- German investors owning a controlling shares indirectly via a DIE abroad with a balance sheet of at least 3 million € of which they own at least 50%.

Primary and secondary FDI are not stated as such by investors under the reporting requirements, but are derived from balance sheet information. Specifically, they are constructed according to the Extended Directional Principle, which calculates FDI as adds up the investor's equity capital and loans and subtracts claims of the DIE towards the investor (reverse investment). Special care is taken to include credit relationships with sister companies, whereby the location of the ultimate controlling institution determines whether these relationships count as inwards or outward FDI.

- Primary FDI is aggregated according to international standards. It comprises directly-owned equity and credit.
- In addition, primary and secondary FDI encompasses direct investment in holding companies. This is the main aggregate of FDI that we use. The results also hold for primary FDI only.

[18] There was a small change in reporting regulations between 2001 and 2002 and a more substantial one between 2006 and 2007.

[19] See Blank et al. (2020) for more details

Note that the FDI measures also include reverse investments such as loans from the DIE to the investor, which lower the FDI value. This is why combined primary and secondary FDI can be lower than only primary FDI (e.g. Figure B3.1).

Data used in the model estimation: Country level data on FDI stocks by the OECD averaged over the period of 2015-2018.

Other data in the model

Table D3.1: Data sources of variables used in the model

Variable	**Description**
GDP	World Development Indicators (WDI)
Capital stock	Private, government, and public-private partnership capital stocks, IMF
Labour force	WDI
Labour share of GDP	ILO
Wages	Constructed based on GDP, labour share of GDP, and labour force; average wages are estimated as: labour share of GDP × GDP level / labour force
Migrant stock	Migrant population by country of origin for 2015 (UN)
Migrant share	Share of total population
Imports	OECD bilateral imports
Consumption shares	Constructed based on imports and domestic GDP net of imports