Labour supply decisions of somewhat caring individuals

Amelie Mennerich*
University of Oxford

November 2, 2026

Abstract

How do couples make labour supply decisions when their degree of cooperation varies? This paper applies a model of two-agent decision-making that spans the full range from non-cooperation to perfect cooperation to household labour supply. The degree of cooperation is equivalent to how much partners care about each other. Increasing caring improves individual and household utility, except when caring is very asymmetric. A rise in the lower earner's wage decreases the higher earner's utility by shifting the within-household distribution of consumption. Productivity in housework, by contrast, confers no power because its benefits are non-excludable. All else equal, the partner with the lower productivity in housework is better off. Labour supply elasticities rise with intra-household wage inequality, and more so the higher the degree of cooperation. Caring therefore matters: assuming perfect or no cooperation misses these patterns and misstates the efficiency gains from taxing labour income.

JEL codes: D13, J22, D64, H31.
Keywords: Intra-household cooperation, caring, labour supply, joint taxation.

*University of Oxford, Department of Economics. Email: amelie.mennerich@economics.ox.ac.uk. I would like to thank Abi Adams and Hamish Low for their guidance and support. This project has also benefited from helpful discussions with seminar participants at the University of Oxford.

1 Introduction

The last 50 years have seen large changes in what families look like. For example, in the UK, 52% of children are born to married parents, 32% to unmarried but cohabiting parents and 16% to parents who do not live together (Kiernan et al., 2024). The couples coordinate market work, housework and care, but to potentially different degrees. For example, couples who do not live together or are not married, may pool their income less than married spouses. How much partners cooperate therefore shapes how they divide paid and unpaid work, and how they respond to wages and taxes. Yet, most models of family labour supply assume either full cooperation or no cooperation. This paper studies couples' labour supply and housework decisions when cooperation can lie anywhere between none and full.

Household models rest on two kinds of assumptions. Static assumptions regard how much couples coordinate within a period, and dynamic assumptions concern the nature of commitment over time. Based on the assumption of static efficiency, there is an ongoing debate about how commitment should be modeled (Mazzocco, 2007; Theloudis et al., 2025). The focus of this paper is to instead assess the scope and implications of the underlying static assumption. The first class of static household models is the unitary model, in which households behave like a single agent, but this has been rejected in multiple studies (Browning and Chiappori, 1998; Lundberg et al., 1997).

In non-cooperative models, partners play a Nash equilibrium, that is self-enforcing but inefficient (Balasubramanian, 2013; Bloch and Rao, 2002; Konrad and Lommerud, 2000; Lundberg and Pollak, 1993). In the alternative collective models of household behaviour, partners are assumed to have perfect information and a long relationship outlook, such that allocations are Pareto efficient (Chiappori, 1988). As an important exception, Del Boca and Flinn (2012) implement multiple constraints that determine whether a household acts with full cooperation or non-cooperatively in a given period. Because the assumption of efficiency reduces the household optimisation problem to a single-valued objective subject to participation constraints, collective models of the household are popular even in contexts, where full cooperation is a strong assumption. For example, combined with the dynamic assumption of limited-commitment, this is the prevalent model for the study of divorce (see Voena (2015), Low et al. (2018), Fernández and Wong (2017)). This implies that divorce decisions are efficient and that labour supply and housework decisions are efficient in a static sense. The risk of future divorce can distort saving, but it cannot cause static inefficiencies like the underprovision of goods that are public within the household. From the perspective of a policy-maker, there is therefore no motive for intervention, as there is no room for Pareto

improvements.[1]

When families do not cooperate and public goods are underprovided, family policy can raise efficiency (Konrad and Lommerud, 1995). If couples coordinate partially, taxation affects not only redistribution but also efficiency. How important each of these is depends on how cooperative couples are, a parameter that is fixed by assumption in existing models.

I apply the generalisation of Cherchye et al. (2020), which nests both types of models, to a household with home production of the public good. Cherchye et al. (2020) develop a revealed preference framework that allows for a full spectrum of cooperation in a household consuming any number of private or public goods purchased in the market. Specifically, their solution is based on Lindahl prices of the public good. They apply it to an experiment with dyads of Belgian primary school children. d'Aspremont and Dos Santos Ferreira (2014) explicitly apply the Lindahl model to household behaviour and they implement the degree of cooperation via budget constraints rather than preferences. However, like Cherchye et al. (2020) they only consider public goods purchased in the market.

Home production matters for two reasons. First, household expenditure data assigned to individual members are rare and the classification of expenditure into public and private goods is, with the exception of women's and men's clothing, notoriously difficult. In contrast, data on time use is widely available. Second, many public goods require partners' time inputs by their nature. Childcare cannot be fully outsourced, and even a family meal that is held at a restaurant requires all household members to be present, i.e. invest time, else it would not be considered a family meal. In fact, the wish to spend time together may be a main reason for the decision to live together. For these reasons, collective models have long included home production (Apps and Rees, 1997; Chiappori, 1997).

Because the model has a single caring parameter, future work can estimate the degree of caring to test whether couples are closer to full or no cooperation. It could also be the case that in practice partners can best be described as caring for each other to an intermediate extent, in that case the model would be even more relevant. Alternatively, it allows the researcher to ex ante be agnostic about the level of cooperation. Other than contributing by removing the marketability assumption, to my knowledge, this study is also the first to

[1]Of course, the use of collective models could also be motivated by empirical evidence. Various authors have tested implications of the collective model and failed to reject it (e.g. Browning and Chiappori (1998), Dauphin et al. (2011), Cherchye et al. (2009)). However, Dauphin et al. (2018) have called some of this evidence into question. Few papers reject the assumption of efficiency, namely Dauphin et al. (2018) and Udry (1996). One further challenge for collective models is that, as pointed out for example by Doepke and Tertilt (2019), in order to explain why changing the recipient of Government transfers can affect household behaviour, it would be necessary and unattractive to assume that partners have different preferences.

spell out implications of this model once it is parameterised.

This paper further contributes to the literature on how couples jointly choose labour supply and housework of each partner by discussing how these decisions depend on the level of cooperation. The model yields four main results. First, increasing both partners' level of caring raises individual and household utility, but very asymmetric caring does not. Second, because earning power determines how private consumption is split, a rise in the wage of the lower earner decreases utility of the higher earner. When there is perfect cooperation, this maps into a specific distribution factor in the collective model. Third, labour market and household productivity have very different implications. Because gains in public good production are by definition non-excludable, productivity in household work confers no power. All else equal, the partner who is less productive in the household is therefore better off. And, lastly, labour supply elasticities rise with within-household wage inequality, but they rise faster the more the partners cooperate. Conversely, labour supply elasticities rise faster in the degree of caring, the larger the intra-household difference in wages. These results have implications for policy. Because partially cooperative couples spend too little time on home production and, when their wages differ, specialise too little, taxing labour income with lump-sum rebates raises welfare, and the tax system should not penalise couples for dividing market work unequally, which supports joint taxation.

In the following, Section 2 introduces the model, points out a duality of cooperation and caring and presents how the optimal behaviour varies along all parameters. Section 3 more specifically analyses how labour-supply elasticities depend on both the distribution of wages and cooperation. Section 4 considers how the different interpretations of the model relate and what the model implies for optimal household taxation and Section 5 concludes.

2 Model

2.1 Set-Up

Individuals maximise their own utility by choosing how to allocate their time between working in the market and consuming their income as a private good q_i[2], and working in the household d_i to contribute to the production of a public good Q. They face subsidies σ_j for their time investment in household production and need to pay transfers σ_i to their partners for their housework. Their utility directly depends on their partner's choices since the level

[2]Leisure can be conceptualised as being subsumed under the private consumption index q_i. Agents can spend privately consumed money on leisure by reducing their work hours.

of the public good depends on both partners' time inputs (Equation 1). Specifically, home production is additively separable in each partner's contribution, but there are decreasing returns to scale to each individual's time input (Equation 6). For ease of notation, I only present the problem for person A since the two agents face symmetric optimisation problems.

$$\max_{q_A, d_A} U_A(q_A, d_A; d_B) \tag{1}$$

s.t.

$$p' q_A + (w_A - \sigma_B)' d_A + \sigma_A' d_B \leq T * w_A \tag{2}$$

$$\sigma_A = \theta_A \ \tau_{A,k}(q_A, d_A, d_B) \tag{3}$$

$$\tau_{A,k}(q_A, d_A, d_B) = \frac{\partial U_A(q_A, Q)}{\partial Q} \ \frac{\partial Q(d_A, d_B)}{\partial d_B} \ \frac{1}{\lambda_A} \tag{4}$$

$$U_i = \log(q_i) + \log(Q) \tag{5}$$

$$Q = d_A^{\gamma_A} + d_B^{\gamma_B} \tag{6}$$

The model allows for flexible levels of cooperation through the budget constraint (Equation 2). Each partner has total time T, which tehy can spend on market work at wage w_i or on housework. Their available income is their full income Tw_i, net of the earnings foregone through time spent on home production. Crucially, the wage loss is mediated due to the fact that proportional subsidies σ_i are paid for time spent on home production. The Lindahl equilibrium, in which every beneficiary of a public good contributes their marginal value from consumption of that public good, can be implemented by making individuals face its actual price reduced by all other contributions, i.e. marginal values, to the total price of the public good. I deviate from the specification of Cherchye et al. (2020) by allowing subsidies to apply to inputs to public good production rather than the public goods themselves. There are exactly two public good inputs, d_A and d_B and each agent can procure only one of them.[3] In a two-agent model in which the price of the public good equals the forgone income, this means that every agent has to pay their own wage rate minus the other agent's

[3]This holds independently of any assumption on the functional form of production. (E.g. could make $Q = d_A^{\gamma_A} + d_B^{\gamma_B} + d_A * d_B$ but would still have exactly two public good inputs.) One could further add market purchased inputs to public good production.

marginal value for spending time on home production.[4] To allow for a spectrum of efficiency between the Lindahl equilibrium and the non-cooperative case, a parameter θ_i is introduced and the actual subsidies paid are scaled by the level of cooperation or caring (Equation 3). If $\theta = 1$, every agent pays exactly their marginal value of the public good inputs such that a Lindahl equilibrium obtains. When $\theta = 0$, no transfers are paid and the agents play a Nash equilibrium.

2.1.1 The dual problem characterised by caring

As Cherchye et al. (2020) explain, the above problem in which parameter θ was presented as a measure of the level of cooperation within and therefore efficiency of households is dual to one in which caring individuals maximise functions of each agent's egoistic utility. In this framework, agent A maximises a "caring function" $W_A(U_A, U_B; \theta_B)$ in which the same parameters θ relate to the equilibrium curvatures of the caring functions. In the Cherchye et al. (2020) caring framework agents maximise utility by buying bundles of private consumption for themselves and for their partner as well as bundles of public consumption. Whilst this contradicts standard economic intuition, with caring preferences individuals benefit themselves from buying goods for their partner. Given the endogeneity of home production, agents must also buy each other's housework time for each other. As noted above, one way to interpret this is by paying wages for housework.

In the caring model, A's problem is given by:

$$\max_{q^{A,A},\, q^{A,B},\, d^{A,A},\, d^{A,B}} W^A(U^A(q^A, Q(d^A, d^B)), U^B(q^B, Q(d^A, d^B))) \tag{7}$$

[4] In some way, this model can be rephrased as requiring partners to pay wages on the other's housework as demanded by second wave feminists, see for example Federici (1975).

s.t.

$$p'(q^A + q^B) + w_A d^A + w_B d^B \leq T(w_A + w_B)$$
$$q^{A,A} + q^{B,A} = q^A$$
$$q^{A,B} + q^{B,B} = q^B$$
$$d^{A,A} + d^{B,A} = d^A$$
$$d^{A,B} + d^{B,B} = d^B$$
$$(d^A)^{\gamma_A} + (d^B)^{\gamma_B} = Q$$

And a "group equilibrium with caring" is defined to be an allocation $\{q^{A,A}, q^{A,B}, q^{B,A}, q^{B,B}, d^{A,A}, d^{A,B}, d^{B,A}, d^{B,B}, \}$ that simultaneously solves A's problem and the symmetric problem of B. Importantly, both A and B choose their contribution to each public good input $d^{i,j}$ as well as their contribution to each agent's private consumption $q^{i,j}$. There is an income pooling budget constraint, but each partner takes the other's expenditure as given.

Again following Cherchye et al. (2020), the transfer scaling parameter that indicates the level of cooperation θ equals the ratio of marginal values of each agent's private consumption in the caring function at equilibrium. Given the above functional forms, this implies:

$$\theta^A = \frac{\frac{\partial W^B}{\partial U^A}}{\frac{\partial W^B}{\partial U^B}} \frac{q^B}{q^A} \tag{8}$$

Whilst it may not be possible in practice to find the precise parametric specification of the caring function that is dual to the problem of non-caring, non-cooperative individuals with mandatory transfers, solutions implied by maximising a CES function of the egotistic utilities are not too far off.

2.2 Behaviour over the ranges of all parameters

In the following, I will analyse how the couple behaves over a spectrum of parameter values. A few interesting findings emerge. Firstly, an individual who cares a bit about their partner is better off than an individual who does not, whatever the level of caring of the partner.

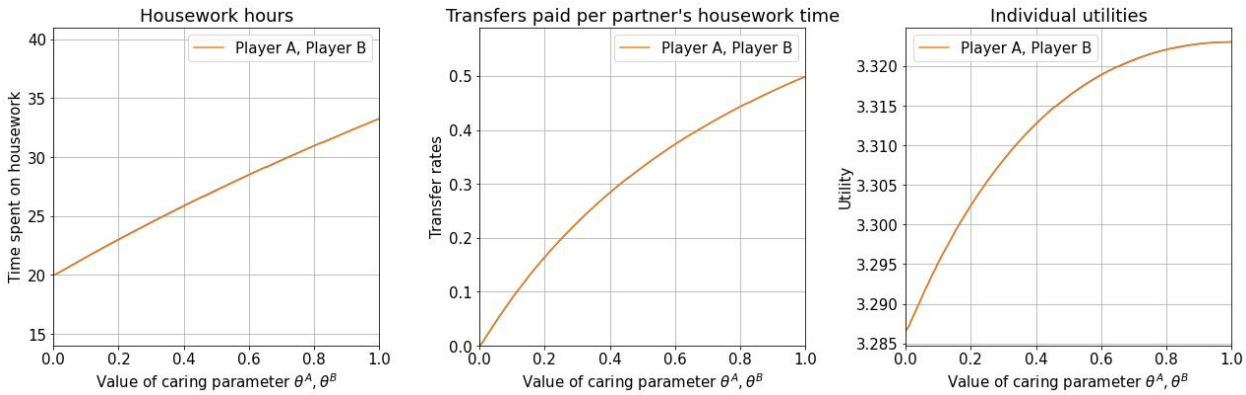


Figure 1: Caring varies for both partners

However, individuals who care a lot more about their partner than their partner cares about them, are worse off than those being only slightly more caring than their partner. At large levels of caring inequality, even the sum of utilities is detrimentally affected since the inequality in caring causes inequality in consumption and household production becomes inefficiently asymmetric. Secondly, a rise in the wage of the lower-earning partner reduces the utility of the higher-earning partner, whereas a rise in the higher earner's wage benefits all. Thirdly, household productivity and wages play very different roles in the model. Whilst both agents benefit from any rise in home productivity, when all other parameters like wages are equal for both, the partner who is less productive in the household is better off. This is because, fourthly, the distribution of the gains from cohabitation depends on wages but not household productivity. And, lastly, the distribution of wages and level of cooperation are correlated in their effects on individual utilities.

2.2.1 Caring parameter

As explained above, when the caring parameter equals zero for both partners, no transfers are paid and the agents play a Nash equilibrium characterised by relatively low levels of home production (Figure 1). As θ rises, there is more cooperation, which in this model manifests itself in higher transfers. Agents do not internalise how the transfers depend on their choice of time allocation and respond to increased transfers as they would to a fall in the relative price of time spent on housework. Therefore housework hours and public good consumption rise. This comes at the expense of a fall in consumption of the private good, but is efficiency and welfare improving.

The parameter θ_A expresses the share of the marginal value A derives from B's housework, that A transfers to B. Therefore, it correlates with the curvature of A's welfare function in a

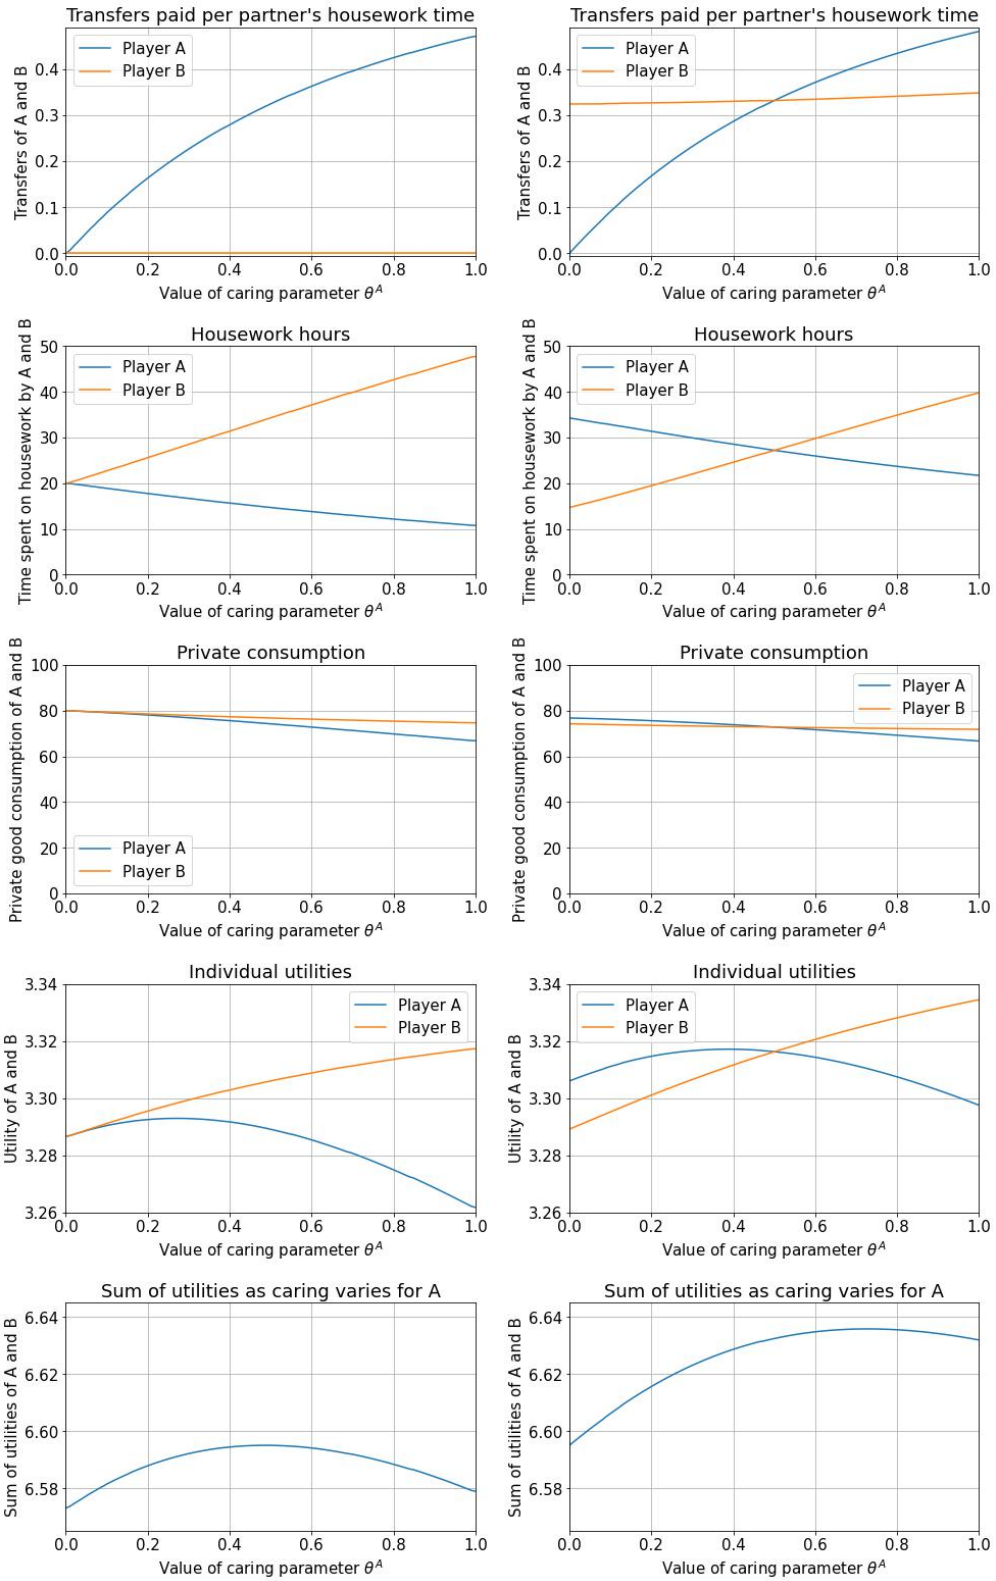


Figure 2: Caring varies for A when $\theta_B = 0$ (left column) or $\theta_B = 0.5$ (right column)

model in which both agents maximise a caring function that aggregates individual utilities. Varying θ_A whilst holding θ_B constant at for example 0 (or 0.5) as shown on the left-hand (right-hand) side of Figure 2 therefore describes what choices households make when A and B care about each other to different extents.

As A transfers B an increasing amount for their housework, the real price of housework falls for B, such that B engages in more housework. Because B also receives more transfers, this counteracts the fall in monetary income B would have otherwise sustained by shifting labour supply from the market to the household. Therefore, B's private consumption only falls marginally and B's welfare rises. The rate of transfers B pays to A for an extra unit of housework is flat by construction in the case of $\theta_B = 0$. When $\theta_B = 0$ as in the left column Figure 2, B generally does not make any transfers to A such that A only provides as much housework as an agent playing a Nash equilibrium and best responding to B's level of public good provision would. Therefore, this displays a case of unilateral deviation from the non-cooperative equilibrium. Changes in A's behaviour in response to the rise in θ_A are therefore like changes in response to a fall in income together with a rise in public good provision by the partner. Both of these factors cause A to reduce time spent on housework. When B cares about A to an intermediate extent (right column of 2), the rate of transfers B pays to A for their housework is nearly flat. The change in time allocation causes a marginal rise in public good consumption in either case and B's marginal value of money rises marginally due to a minor drop in B's private consumption q_B. For the case of $\theta_B > 0$ these two changes would ceteris paribus cause the transfer rate σ_B to fall but they are countered by the fact that because A engages in less home production, its marginal product rises. Overall, σ_B rises marginally as θ_A rises for a given level of $\theta_B > 0$.

This analysis has interesting implications for each agent's as well as overall welfare. As noted, B's utility unambiguously increases in the amount of caring displayed by A. Both when B does not or intermediately cares about A, the type of A with $\theta_A = 0$ has lower egoistic utility U_A than some type of A with small but positive caring.[5] When θ_A is very low, the loss of welfare due to inefficiently low public good contribution of B exceeds the benefit A derives from not paying transfers and therefore having a higher private consumption. Contrarily, at larger values of θ_A, A's welfare decreases from further rises in caring since transfers increase by more than A's gain from rising home production. This is in contrast to the case of

[5]The fact that egoistic utility is higher when $\theta_A > 0$ implies that an non-caring type of agent A would benefit from being able to commit to paying transfers to B and therefore mimicking a type with higher θ_A. If there were a general tool available by which agents could pay subsidies to home production (conditional transfers), A would choose to do so unilaterally at low levels of θ_B. Specifically, A would choose a level of caring above B's as long as $\theta_B < 0.35$ and a level below B's for larger values of θ_B.

symmetric changes in caring, any mutual rise in θ is welfare improving.

The sum of utilities is maximised when θ_A is larger than $\theta_B = 0$ or $\theta_B = 0.5$. In general, total welfare increases in the level of caring. Yet, this only holds for some degrees of caring inequality. Two factors contribute to the decline in total utility as θ_A becomes large compared to θ_B. Firstly, due to the log form of utility consumption inequality is detrimental to the sum of utilities. Since the level of public good consumption is by definition equal, the sum of logs of private consumption is maximised when income is split equally. Therefore, total utility falls as consumption inequality increases. Secondly, in the household production function (Equation 6) there are decreasing returns to scale to individual inputs. Therefore, the asymmetry in the production schedule caused by the change in marginal prices of housework that follows from unilaterally increasing transfers, is inefficient. When $\theta_A = 1$, but $\theta_B < 1$, the level of housework hours of B is higher than in the case of mutual full cooperation as presented in Figure 1.

2.2.2 Wages

As both partners' wages rise, there is no change in time allocation (see Appendix Figure 13). Due to the logarithmic shape of the utility function, income and substitution effects on time use and public good consumption balance out. Private consumption and transfers rise linearly with both wages and welfare rises at a decreasing rate.

When A's wage varies for a given level of B's wage, A again faces income and substitution effects on time allocation that go in contrary directions (left column of Figure 3). However, the income effect also causes the value of the Lagrange multiplier λ_A to fall, which causes A to pay higher subsidies to B for their time spent on public good production (by Equation 3). As B's relative price of housework falls, d_B rises. This causes A to reduce time spent on housework. A's private consumption rises linearly with their income, but for low levels of w_A, at which an interior solution is played, B's private consumption falls slightly. On the one hand, σ_A rises, so B receives a greater subsidy per unit of housework. But the rise in transfers received does not fully equal out the loss of private consumption that results from the reallocation of time. Only when the wages are so unequal that the agents completely specialise does B's private consumption increase. In that range, the transfers paid from A to B still increase with A's income, but there is no longer a loss of q_B due to reduction in work hours.

Surprisingly, the utility of B even displays a fall at very low levels of w_A. This is due to initial falls in both home production and private consumption. For intermediate levels of

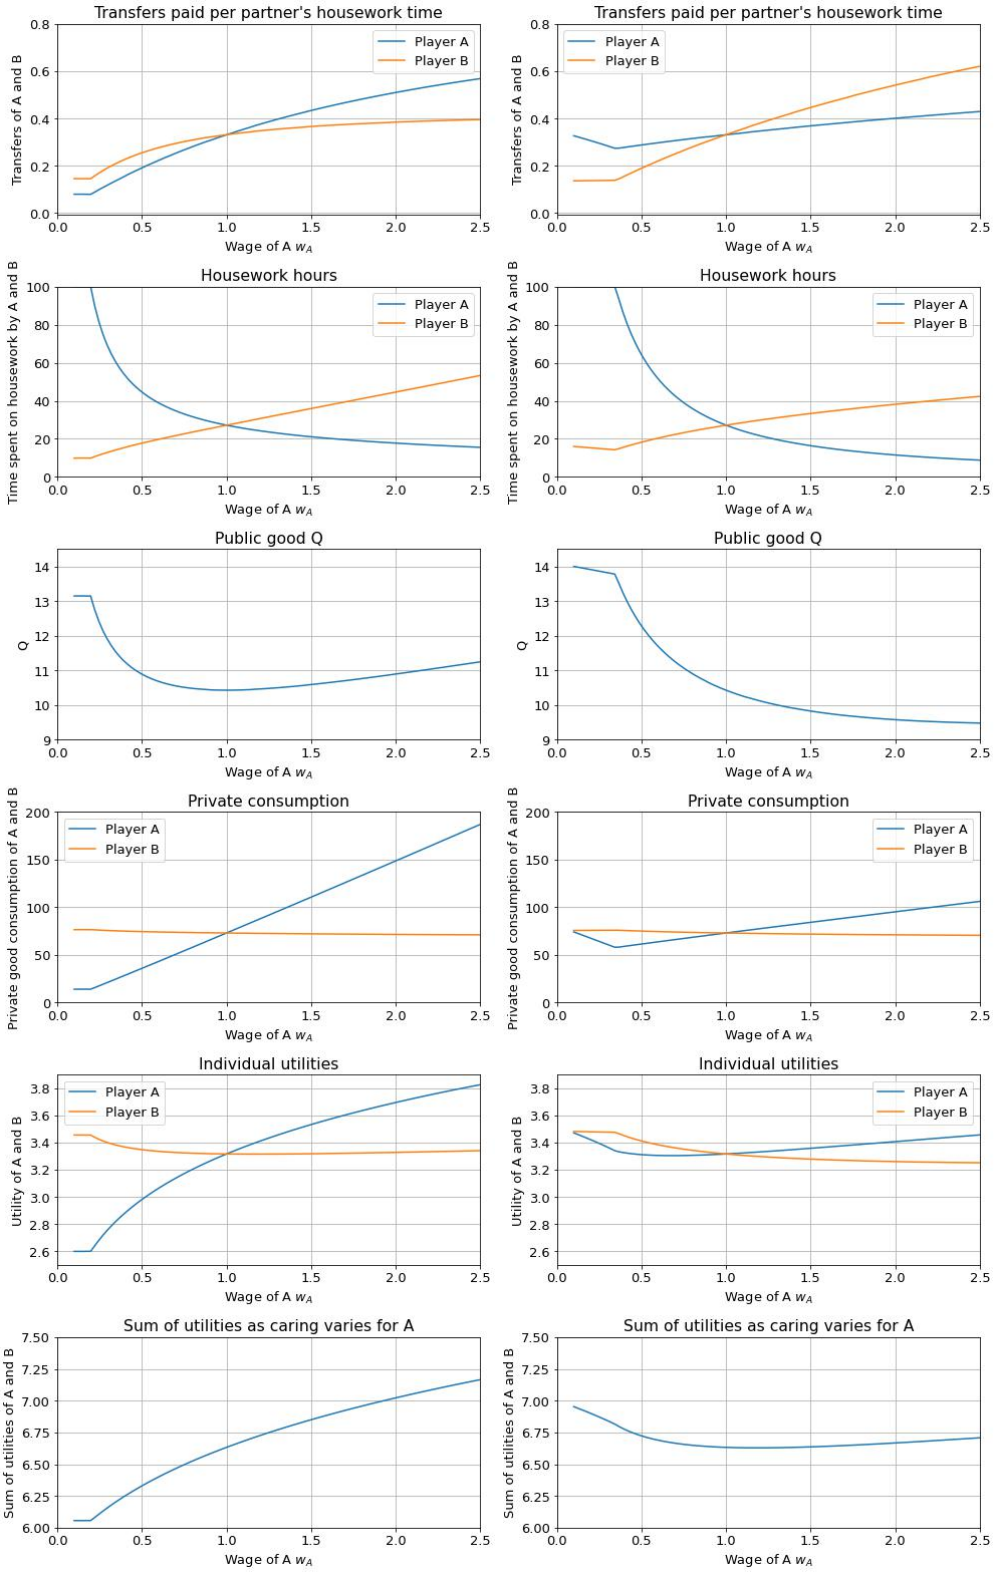


Figure 3: Wage of A varies without compensation (left column) or with Slutsky compensation (right column)

caring, not all of A's gains from specialising in market work after a rise in wages are passed on to B in the form of subsidies, such that B's extra time spent on home production initially does not fully compensate the loss of public good production resulting from A's reallocation. For intermediate levels of caring as in Figure 3, since A only transfers fraction θ_A of her marginal value to B for their homework, B's private consumption falls as they start doing more housework. Conversely, when there is full cooperation and $\theta = 1$, the fall in B's utility is only caused by a fall in public good consumption, private consumption is constant for the range in which B has not fully specialised in housework (shown in Figure 8).

The right column of Figure 3 shows what happens when A's wage rises but A's total income is held constant to the level achieved by optimal behaviour at $w_A = w_B = 1$ by a lumpsum cost (Slutsky compensation). Therefore, to the left of $w_A = 1$, the agent receives a positive lumpsum payment, and to the right of this point the transfer is negative. As a wage increase only causes a change in the relative price of homework and market work but no positive income effect for A, they engage in more market and less housework and the reverse happens for B. Since income is compensated, transfers to B increase only due to the fall in housework by A and the associated rise in marginal utility of public good consumption. B increases their time spent on housework both to pick up the slack and due to increased transfers. Hence, public good production falls but remains positive. Compared to the non-compensated case in Figure 3, the curve of A's housework levels is steeper, so specialisation happens for smaller differences in pay. This is all because there is only a substitution effect away from public good production and no income effect partially offsetting it. As described above, B's utility falls because their move from market to housework is not fully compensated by transfers from A. But now, A's private consumption and utility are flattened due to the compensation.

One further noteworthy attribute of these graphs is that A's level of public good production never falls to zero. When B has fully specialised in home production, income and substitution effects on home production balance out for A.

2.2.3 Household productivity

For the problem to be well-behaved, it should be assumed that there are decreasing returns to scale in home production since there are multiple solutions if home productivity $\gamma = 1$. For strictly decreasing returns to scale, any rise in household productivity causes the agents to spend more time on household production and therefore consume relatively more of the public and relatively less of the private good. Raising productivity is welfare improving.

The level of transfers per unit of home production is nearly constant (very small increase), since the rise in marginal product of time spent on home production almost fully cancels out with the fact that the marginal value of the public good falls with the level of public good consumption due to the logarithmic form of utility.

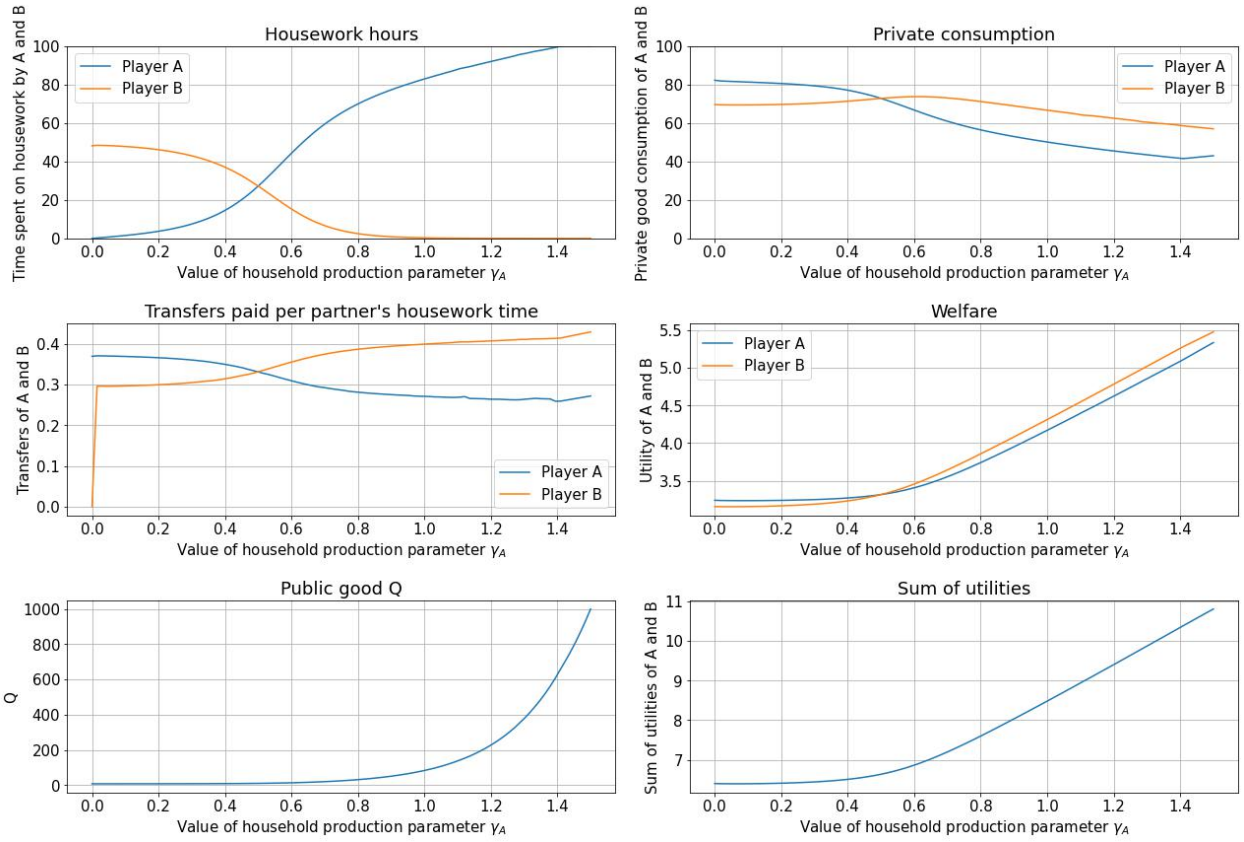


Figure 4: Productivity in household good varies for A

As only person A becomes more productive in the household, this causes the marginal product of A's housework and therefore the subsidy A receives from B to increase (Figure 4). It follows that A finds it optimal to specialise in housework and as the level of the public good rises, B finds it optimal to spend more time on market work. In the intermediate cases as for example when $\theta = 0.5$ as shown in Figure 4, gains from a rise in home productivity are not shared equally. As A's home productivity is lower, A spends more time on market work, therefore has a higher income and higher private good consumption. Therefore, surprisingly it is the partner with the lower household productivity who has higher utility. To some extent, in the range where $\gamma_A < \gamma_B$, A freerides on B's public good production because A does not transfer B the complete marginal value they derive from it.[6] Therefore, A's

[6]In this model, B cannot exclude A from benefiting from their productivity because that productivity

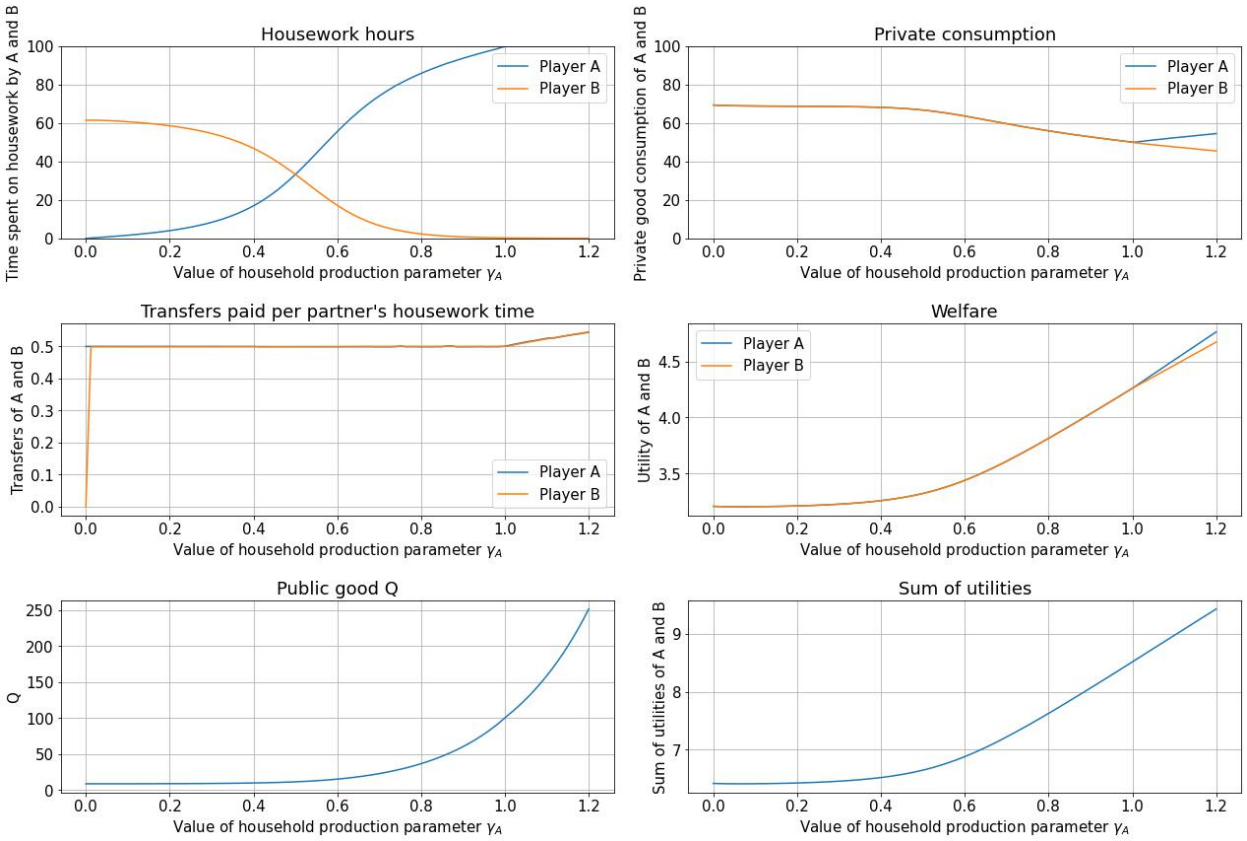


Figure 5: Productivity in household good varies for A, $\theta = 1$

welfare is higher than B's as long as A is less productive than B. Counterintuitively, for low levels of the home productivity parameter, the less productive household member even receives higher subsidies for their housework further aggravating the inequality in private consumption. This can only be explained by the fact that specialisation is so stark that overall marginal productivity for the lower household productivity person is higher since they operate at such low levels of time input.[7]

---

only manifests itself in the public, i.e. non-excludable, consumption. In a model of marriage markets or a limited commitment model, one would expect γ_B to determine B's utility when single and therefore B's outside option.

[7]The ratio of tax rates of the two agents is a function of caring ratio, marginal household productivity and marginal value of money.

$$\frac{\sigma_A}{\sigma_B} = \frac{\theta_A}{\theta_B} \frac{\frac{\partial U_B}{\partial Q}}{\frac{\partial U_B}{\partial Q}} \frac{\frac{\partial Q}{\partial d_A}}{\frac{\partial Q}{\partial d_B}} \frac{\lambda_A}{\lambda_B} = [\text{ if } \theta_A = \theta_B] \frac{\frac{\partial Q}{\partial d_A}}{\frac{\partial Q}{\partial d_B}} \frac{\lambda_A}{\lambda_B}$$

Here, cooperation is symmetric and as A consumes more of the private good for low levels of γ_A, initially $\lambda_A < \lambda_B$. Furthermore, it is the case that in this range $\frac{\partial Q}{\partial d_A} < \frac{\partial Q}{\partial d_B}$, which explains why $\sigma_A < \sigma_B$.

The higher the cooperation parameter θ, the more responsive are couples to differences in household productivity (Figure 4 vs. Figure 5). Very cooperative couples specialise greatly even for small differences in household productivity γ and share gains from household productivity equally. Completely uncooperative couples would not specialise at all.

2.3 Relation to collective model

So far, it has been claimed that the Lindahl equilibrium collapses onto the collective equilibrium as the caring parameter θ tends to one for both partners. It has been shown that when $\theta = 1$ partners pay the full personalised price for each public good input and therefore the efficient level of the public good is consumed. However, in its standard formulation (9), the collective model itself is a spectrum of Nash bargaining solutions for different power distributions. Specifically, due to the mechanics of assuming two individual budget constraints in this paper, the relative bargaining power of each partner is, for interior solutions, equal to their relative earning capacity (11) and the model can be presented as in (9).

$$\max_{q_A, q_B, d_A, d_B} \mu \; U_A(q_A, d_A, d_B) + (1 - \mu) \; U_B(q_A, d_A, d_B) \tag{9}$$

s.t.

$$p' \; (q_A + q_B) + w_A \; d_A + w_B \; d_B \leq T * (w_A + w_B) \tag{10}$$

$$\mu = \frac{w_A}{w_A + w_B} \tag{11}$$

This sheds further light on the finding that labour market and household productivity do not benefit household members in the same way. Whilst the higher earner has higher utility than a lower earner, someone who is more productive at home than their partner will have lower utility than their partner even when wages are equal. In the case of household productivity γ, the only way in which person A can gain from a rise in productivity is through public good consumption (which is by definition common to both) or through higher transfers. Therefore, if there is perfect cooperation, i.e. $\theta = 1$, A's and B's welfare are the same and this corresponds to a collective model with equal weights on both partners' preferences. However, in the case of higher market productivity, A can gain through heightened private good consumption, such that A's welfare rises above B's.

Note that the bargaining weights adjust according to 11 only for interior solutions as can be seen on the left-hand side of Figure 6. As either partner hits their time constraint and

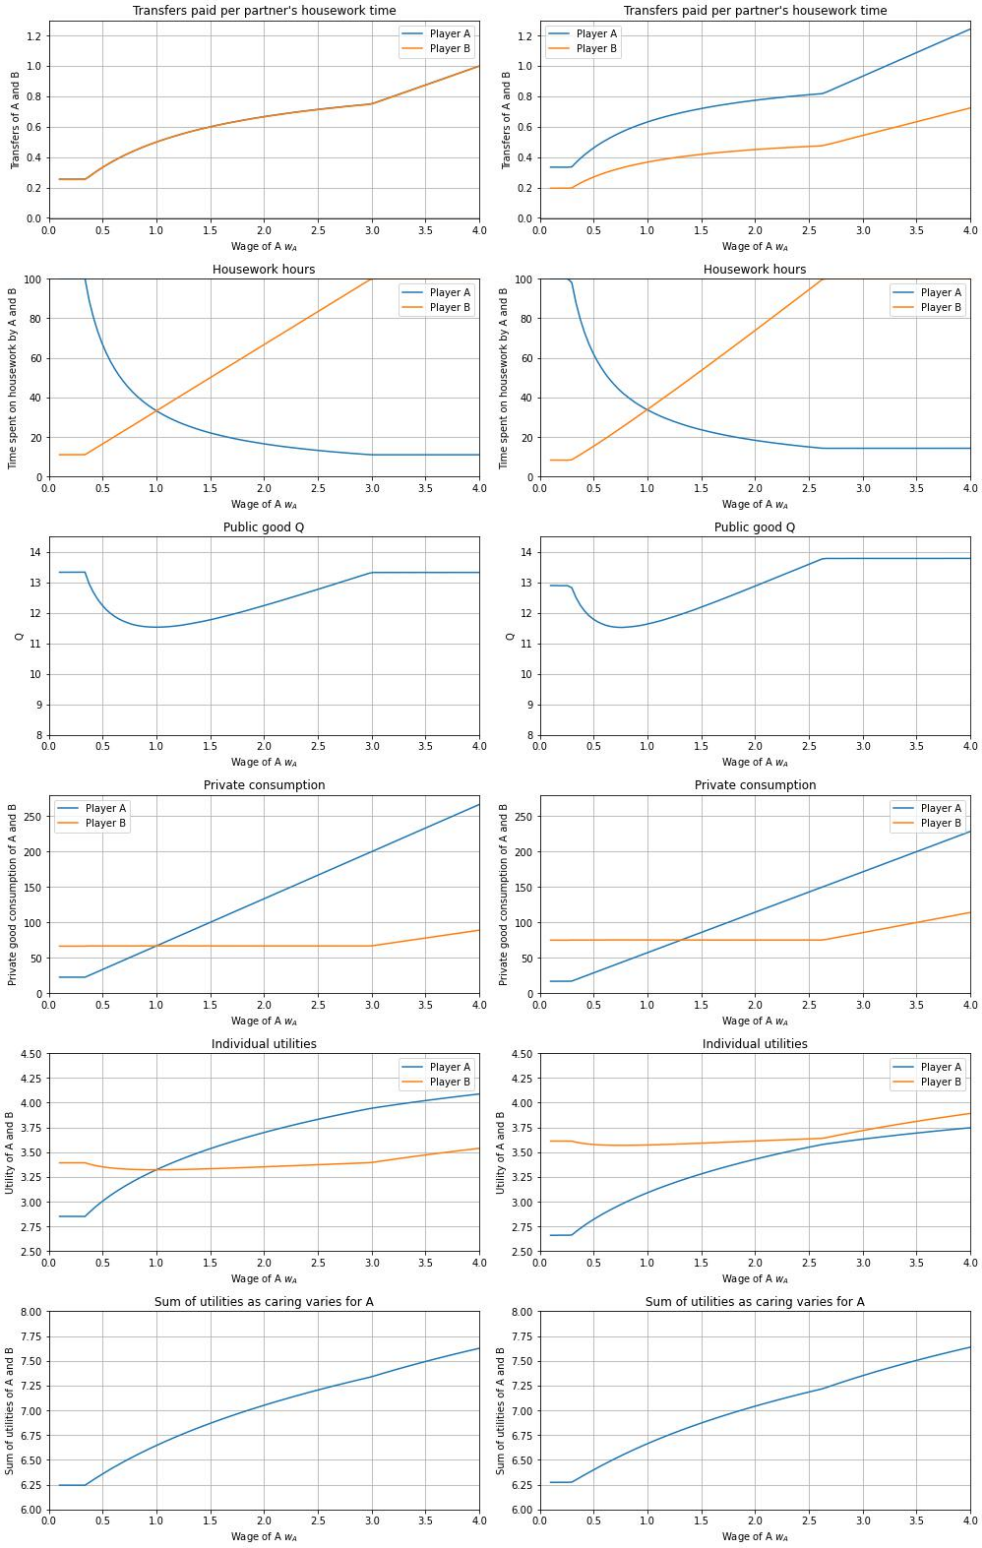


Figure 6: Wage of A varies and $\theta = 1$ (left-hand side) and preferences differ $\alpha_A = 0.4, \alpha_B = 0.6$ (right-hand side)

fully specialises in housework, further changes in wage rates have no additional impact on power. For example, if B spends all their time doing housework at high levels of w_A, any increase in w_A does not cause further shifts in the distribution of private consumption as noted before, but instead consumption shares remain constant, that is, private consumption is parallel for both partners in this range.

When preferences differ, private consumption and welfare differences are not the same and the bargaining weights do not only affect the distribution of the private good but also the production schedule. The right-hand side of Figure 6 also shows optimal time and consumption allocations for perfectly cooperating couples, but for the case in which B cares relatively more about the public good. When preferences differ, it is generally the case that the person caring relatively less about the public good is better off. This is because they benefit from their partner's higher valuation and therefore higher provision of and marginal willingness to pay for the public good. But as A, who in this example cares relatively less about private consumption and more about public consumption, has a higher relative-earning potential, the total level of the public good consumed by the household is higher. Unlike in the left-hand side graph, the right-hand side graph does not show the same maximum levels of public good consumption at both ends of the range of A's wage.

2.3.1 A generalisation

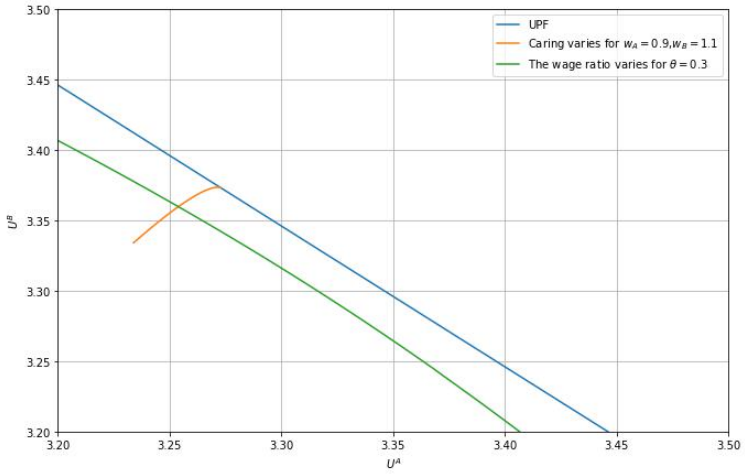


Figure 7: Utility possibility frontier and location within it

The fact that the bargaining weight shifts with individual wages explains why in Figure 3 the utility of B falls as A's wage rises from low levels. Even though total private consumption of the household increases, B's level of private consumption falls as their bargaining power

$1 - \mu = \frac{w_B}{w_A + w_B}$ declines. To illustrate how utilities depend both on bargaining power and level of cooperation, Figure 7 displays how couples move within the utility possibility frontier. Higher levels of caring θ move couples towards the frontier along rays implied by the wage ratio. And for a given level of caring, changes in the wage ratio move couples along an arc.

3 Implications for labour supply elasticities

3.1 Effects of wage change by level of cooperation

Figure 8 shows that the level of cooperation matters for the effects of one partner's wages on the other's labour supply. For the following plots, it is assumed that $w_B = 1$ such that wages are equal at $w_A = 1$. In the extreme case of no cooperation, labour supply decisions are made independently, such that there can for example be no added-worker effect. For this reason, Figure 8 displays the case of a low but positive amount of cooperation and that of full cooperation. The first subplot shows that the higher the level of cooperation, the more responsive is A's labour supply to their own wage. The second subplot shows that B is also more responsive to A's wage in the case of greater cooperation.

Private consumption of A rises with own wage, but for high levels of cooperation it rises less since transfers paid to B are higher. They are higher per unit of time spent on household production (bottom left subplot), but because B responds way more to A's wage rise by specialising in housework, the difference is even bigger in absolute transfers paid.

Again related to what was explained above, private consumption of B is lower in the more cooperative model up to some level of wage inequality. It is lower than in the less cooperative case because the more cooperative couple consumes relatively more of the public compared to the private good (this is the same for q_A, q_B). However, there comes a point at which in the case of perfect cooperation, B fully specialises in housework such that any further changes in A's wage only cause transfers to rise and no longer bring about a reduction in B's market income. At this point, there is a kink in B's private consumption function and eventually B's private consumption becomes higher than in the absence of a high level of cooperation.

Both spouses' welfare is higher in the cooperative case, but the extent of the welfare effect of reduced cooperation depends on the wage inequality. As Figure 9 shows, when wages are equal at $w_A = 1$ the unweighted sum of utilities is only slightly smaller when cooperation is low compared to high. In this case, couples spend inefficiently little time

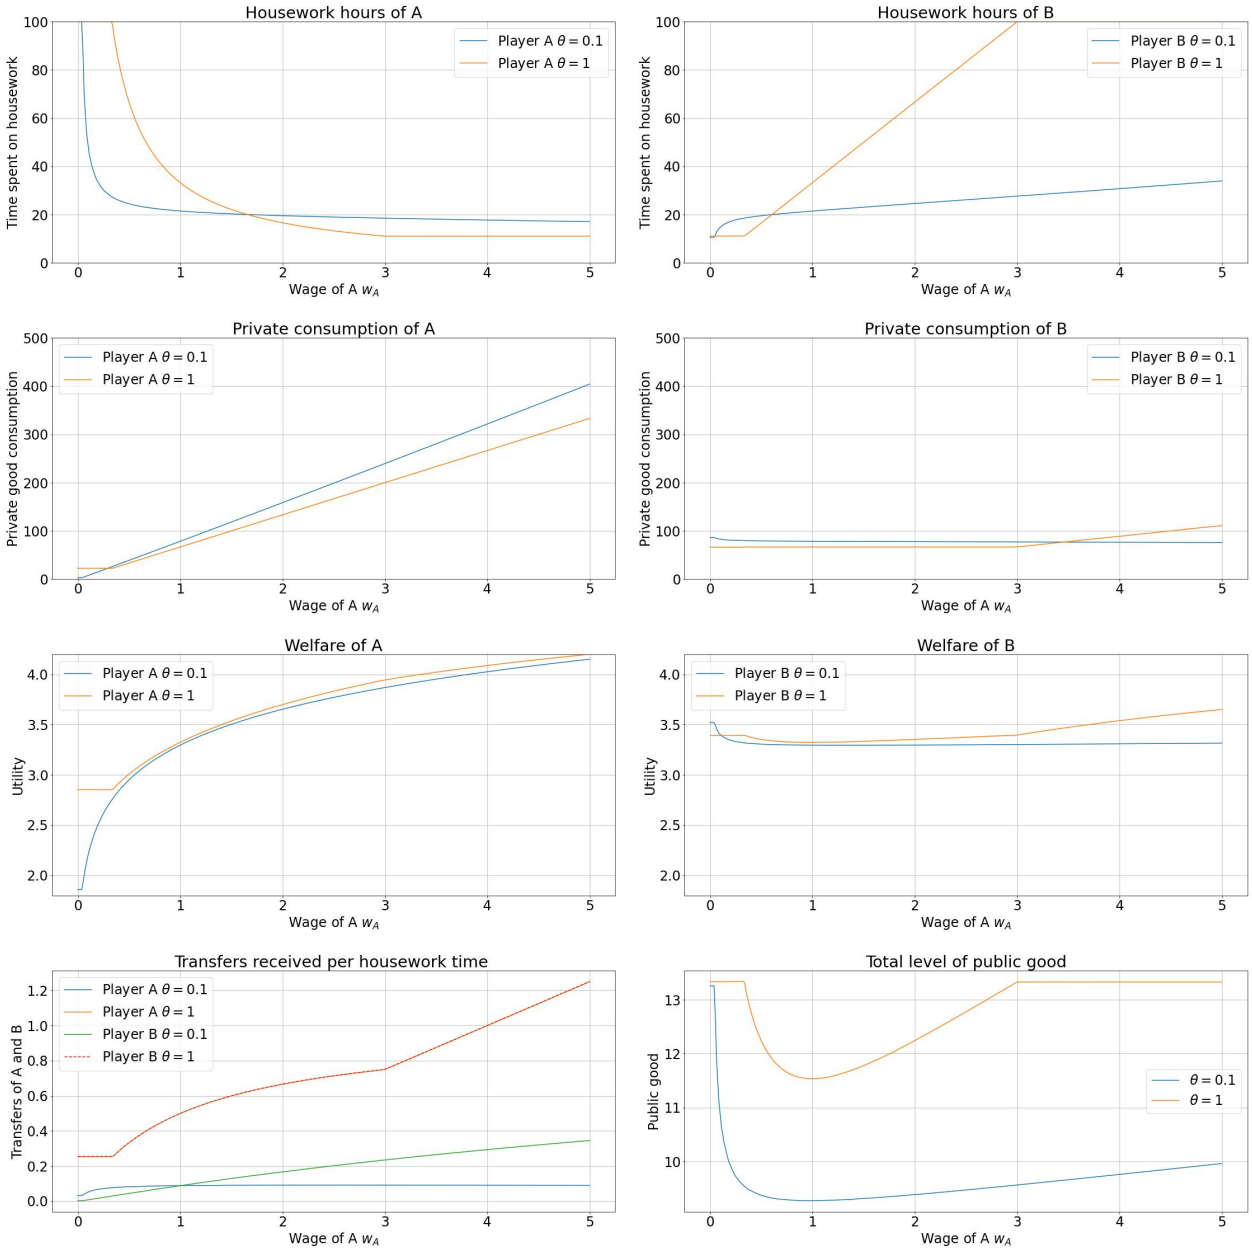
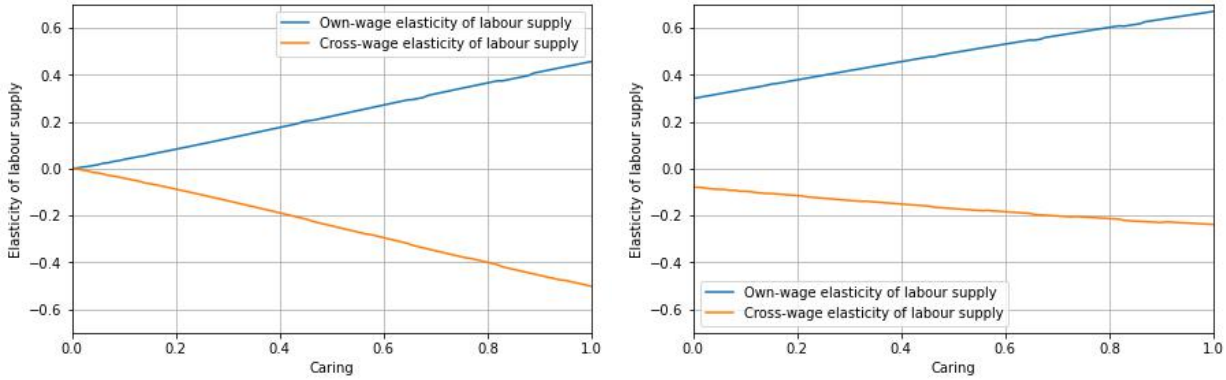


Figure 8: Effects of one partner's wage on all behaviour by level of cooperation



Figure 9: Effects of one partner's wage on the sum of utilities by level of cooperation

on public good production because they fail to internalise their partner's benefit from it. However, the welfare difference becomes much larger when w_A is either small or large. In this case, there is a further inefficiency of too little specialisation. There are furthermore two kinks in the difference corresponding to the points in which fully cooperative couples completely specialise (or rather: one partner specialises since there is a lower bound on home production, but not on market labour).



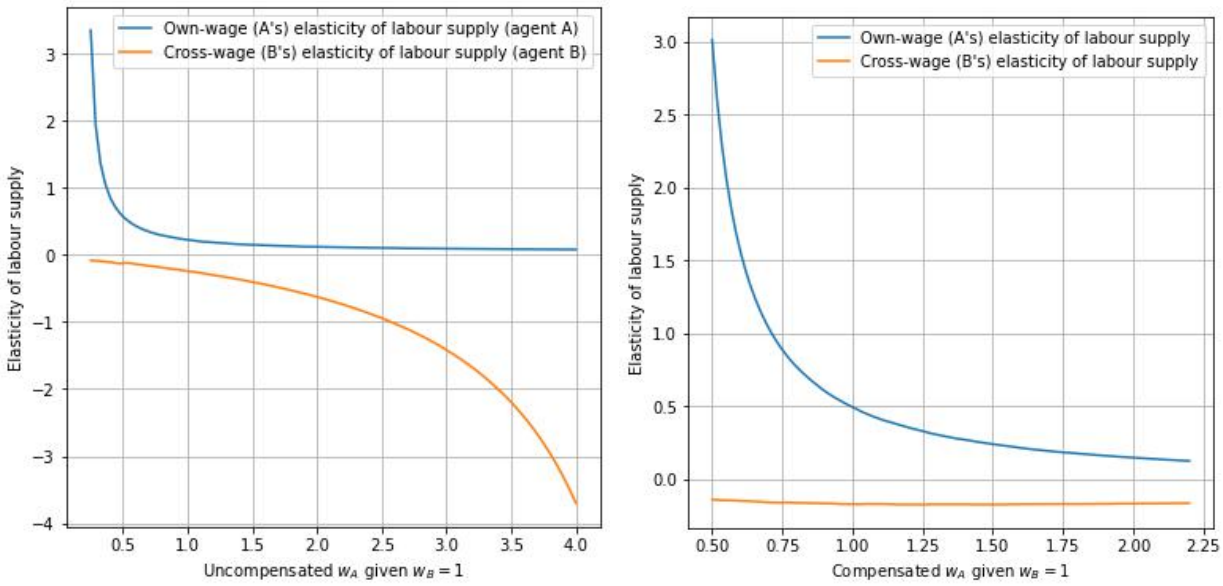
(a) Uncompensated (b) Compensated

Figure 10: Wage elasticities of labour supply

Figure 10a shows the uncompensated own- and cross-wage elasticities as a function of the level of cooperation when both partners' wages are held constant and are equal. Labour supply responds positively to a rise in own wage, implying that the substitution effect outweighs the income effect for the privately consumed good. Players in this model reduce their labour supply as their partner's wage rises because they shift their labour supply towards

domestic work. They do that because in public good contribution with few agents, it is optimal to pick up some slack that the other partner leaves by shifting towards market work. But on top of that, their partner starts transferring them more because in the absence of the behavioural change the level of public good consumption would fall and thus its marginal value rise, and also because transfers are scaled by the marginal value of income. These three effects compound and cause these elasticities to be symmetric. Figure 10b also plots own- and cross-wage elasticities by level of cooperation but when total potential income of agents is held constant by lumpsum taxes. Compensated own-wage elasticities are much larger but cross-wage elasticities are no longer symmetric. They are smaller because without the positive income effect to one partner, say A, transfers they make to the other person do not scale up as much. If B receives no higher transfers from A, they will still shift towards more housework as A's contribution to the public good falls, but not to the same extent as in the uncompensated case. Conversely, the fact that the own-wage elasticity is much larger in Figure 10b compared to Figure 10a can be explained by the presence of income effects.

3.2 Elasticities by wage ratio



(a) Uncompensated (b) Compensated

Figure 11: Elasticities of labour supply by relative wage

Figure 11 shows that A responds more to changes in their own wage when their own wage is very low compared to B's wage both when total income is held constant (Figure 11b) and when it is not (Figure 11a). As noted above, cross-wage elasticities are small when potential

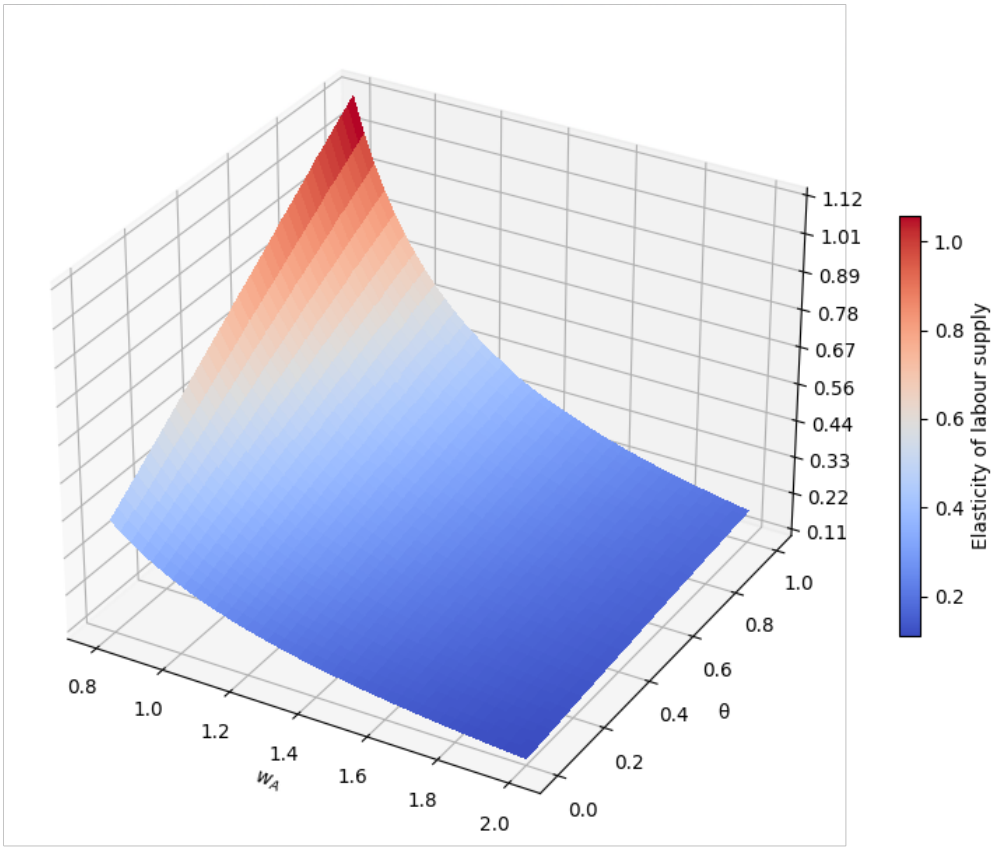


Figure 12: Own-wage elasticity of labour supply

income is held constant. The behavioural response to the partner shifting their time use towards market work is still a partial picking up of the slack in public good production, but the extent of this does not depend much on relative wages. Regarding uncompensated labour supply elasticities, B responds more to changes in A's wage when A's wage is very high compared to B's wage.

3.3 Elasticities along both dimensions

Figure 12 displays how own-wage elasticity depends on both the level of cooperation and the wage distribution. It shows that the positive effect that a higher level of cooperation has on the elasticity of labour supply is larger for a person who earns relatively less than their spouse. Conversely, the negative effect that having a relatively higher wage has on

labour-supply elasticity is larger at higher levels of cooperation. As can be inferred from Figure 11b, the compensated cross wage elasticity varies with the level of caring to an extent that is independent of the wage distribution within the household.

These findings are similar in magnitude to those of Attanasio et al. (2018). They find that women with lower wages and lower weekly hours of work in the labour market have higher income elasticities. This model shows that this can in parts be explained by the fact that these women are more likely to be second earners and are more likely to have spouses with much higher wages than them. However, the variance introduced in labour-supply elasticity through caring is larger for the spouses earning less than their partners. Empirically, Attanasio et al. (2008) find larger variances in elasticity estimates for women in higher wage-percentiles.

4 Interpretation and taxation

Caring and cooperation are different concepts. Intuitively, a level of cooperation seems like a concept of mutual behaviour and a mutual level of cooperation is something the government can act on. For example, legal frameworks should be such that contracts between private firms are enforceable. A tax system and government expenditure on public good also force households to cooperate on public good provision. The option of getting married which involves splitting assets upon separation allows spouses to cooperate and divide labour within the household. Cooperation can also be an asymmetric concept, for example, A could behave more cooperatively than B. A may be more cooperative than B in general, as a character trait, or there could be an asymmetry specific to their relationship.

Caring describes individual preferences and there is no reason to assume that partners care about each other to the same extent. Policy-makers should not interfere with preferences. But suppose couples have low levels of cooperation, then, any intervention that increases the level of cooperation increases both egoistic utilities. Even with unchanged caring preferences this increases the value of the caring function, since the caring function is non-decreasing in both egoistic utilities. So, if both partners behaved as though they were more caring than they are, they would be better off. As described above, there is an equilibrium of caring at just above one third which both partners would choose if they could choose their caring parameter. A higher level of cooperation and more caring preferences may be formally equivalent, but conceptually they are very different and they have very different implications for Government intervention.

Commitment, in the household literature, is an intertemporal concept. To some extent a higher level of cooperation relates to intertemporal commitment. In couples with unequal wages, more cooperation brings about more specialisation. Specialisation carries intertemporal risks when there is divorce and returns to experience in the labour market. Governments can partially mitigate this risk through divorce laws, social insurance and benefits. However, in this model higher levels of cooperation do not cause specialisation in couples in which both partners have the same wages. For these couples, a higher level of cooperation means reduced labour supply for both partners. If there are returns to experience in the labour market but not in the household, the intertemporal implications are again reduced future payoffs in the case of divorce. In this case, divorce law does not mitigate risks since there is no wealth to be redistributed. But social insurance and benefits still do by providing a minimum utility floor. In either case what translates cooperation into intertemporal commitment is the assumption of future returns to present labour supply, which would be an extension to the model.

Therefore, compared to a model that does not take into account the effects on cooperation, this model predicts that if there is less than full cooperation in households, there is a welfare gain from taxing labour income with lumpsum rebates such that more time is allocated to public good production and less to market work. Furthermore, in this model, couples should not be "penalised" by the tax system for distributing levels of market work unequally since specialisation expresses commitment. Thus, couples with the same household income should pay the same taxes no matter how it is split, so there is a case for joint taxation.

5 Conclusion

Generalising commonly used models of the household to allow for varying degrees of cooperation is an achievable and worthwhile goal. It can be implemented by making agents behave like players of a Nash game but subject to transfers for public good inputs that depend on a univariate index of the level of cooperation. One resulting feature of this model is that whilst general increases in caring are always beneficial to both partners' utilities, unilateral increases in caring become detrimental to overall utility when they are very large. Unilateral wage rises of the lower-earning partner can cause the other spouse's utility to fall, since in this framework the wage ratio determines power and hence intra-household distribution of private consumption. When one partner is more productive in the household, their utility is lower whenever cooperation is less than full, since the gains of their household productivity are non-excludable but remuneration for their increased level of and more productive

housework is not complete. These results have particularly relevant implications for labour-supply elasticities. Both own-wage and cross-wage elasticities depend on both the level of cooperation within the household, since greater cooperation allows for greater specialisation and hence greater labour supply flexibility, and the wage ratio. In an intertemporal extension with returns to work experience, divorce law could support cooperation by insuring the partner who specialises. This model can also be used to support a case for joint taxation of couples.

References

Apps, P. F., & Rees, R. (1997). Collective Labor Supply and Household Production. *Journal of Political Economy*, *105*(1), 178–190.

Attanasio, O., Levell, P., Low, H., & Sánchez-Marcos, V. (2018). Aggregating elasticities: Intensive and extensive margins of women's labor supply. *Econometrica*, *86*(6), 2049–2082.

Attanasio, O., Low, H., & Sánchez-Marcos, V. (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review*, *98*(4), 1517–52.

Balasubramanian, S. (2013). Why micro-credit may leave women worse off: Non-cooperative bargaining and the marriage game in south asia. *The Journal of Development Studies*, *49*(5), 609–623.

Bloch, F., & Rao, V. (2002). Terror as a bargaining instrument: A case study of dowry violence in rural india. *American Economic Review*, *92*(4), 1029–1043.

Browning, M., & Chiappori, P.-A. (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica*, *66*(6), 1241–1278.

Cherchye, L., Cosaert, S., Demuynck, T., & De Rock, B. (2020). Group consumption with caring individuals. *The Economic Journal*, *130*(627), 587–622.

Cherchye, L., De Rock, B., & Vermeulen, F. (2009). Opening the black box of intrahousehold decision making: Theory and nonparametric empirical tests of general collective consumption models. *Journal of Political Economy*, *117*(6), 1074–1104.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica*, 63–90.

Chiappori, P.-A. (1997). Introducing Household Production in Collective Models of Labor Supply. *Journal of Political Economy*, *105*(1), 191–209.

d'Aspremont, C., & Dos Santos Ferreira, R. (2014). Household behavior and individual autonomy: An extended lindahl mechanism. *Economic Theory*, *55*(3), 643–664.

Dauphin, A., El Lahga, A.-R., Fortin, B., & Lacroix, G. (2011). Are children decision-makers within the household? *The Economic Journal*, *121*(553), 871–903.

Dauphin, A., Fortin, B., & Lacroix, G. (2018). Is consumption efficiency within households falsifiable? *Review of Economics of the Household*, *16*, 737–766.

Del Boca, D., & Flinn, C. (2012). Endogenous household interaction. *Journal of Econometrics*, *166*(1), 49–65.

Doepke, M., & Tertilt, M. (2019). Does female empowerment promote economic development? *Journal of Economic Growth*, *24*, 309–343.

Federici, S. (1975). *Wages against housework*. Falling Wall Press Bristol.

Fernández, R., & Wong, J. C. (2017). Free to leave? a welfare analysis of divorce regimes. *American Economic Journal: Macroeconomics*, *9*(3), 72–115.

Kiernan, K., Crossman, S., & Phimister, A. (2024). Families and inequalities. *Oxford Open Economics*, *3*, i645–i677.

Konrad, K. A., & Lommerud, K. E. (1995). Family Policy with Non-Cooperative Families. *The Scandinavian Journal of Economics*, *97*(4), 581–601.

Konrad, K. A., & Lommerud, K. E. (2000). The bargaining family revisited. *Canadian Journal of Economics/Revue canadienne d'économique*, *33*(2), 471–487.

Low, H., Meghir, C., Pistaferri, L., & Voena, A. (2018). *Marriage, labor supply and the dynamics of the social safety net* (WP No. 24356). National Bureau of Economic Research.

Lundberg, S., & Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of political Economy*, *101*(6), 988–1010.

Lundberg, S. J., Pollak, R. A., & Wales, T. J. (1997). Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *The Journal of Human Resources*, *32*(3), 463–480.

Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies*, *74*(3), 857–895.

Theloudis, A., Velilla, J., Chiappori, P.-A., Giménez-Nadal, J. I., & Molina, J. A. (2025). Commitment and the dynamics of household labour supply. *The Economic Journal*, *135*(665), 354–386.

Udry, C. (1996). Gender, agricultural production, and the theory of the household. *Journal of political Economy*, *104*(5), 1010–1046.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, *105*(8), 2295–2332.

Appendix

A Symmetric productivity changes

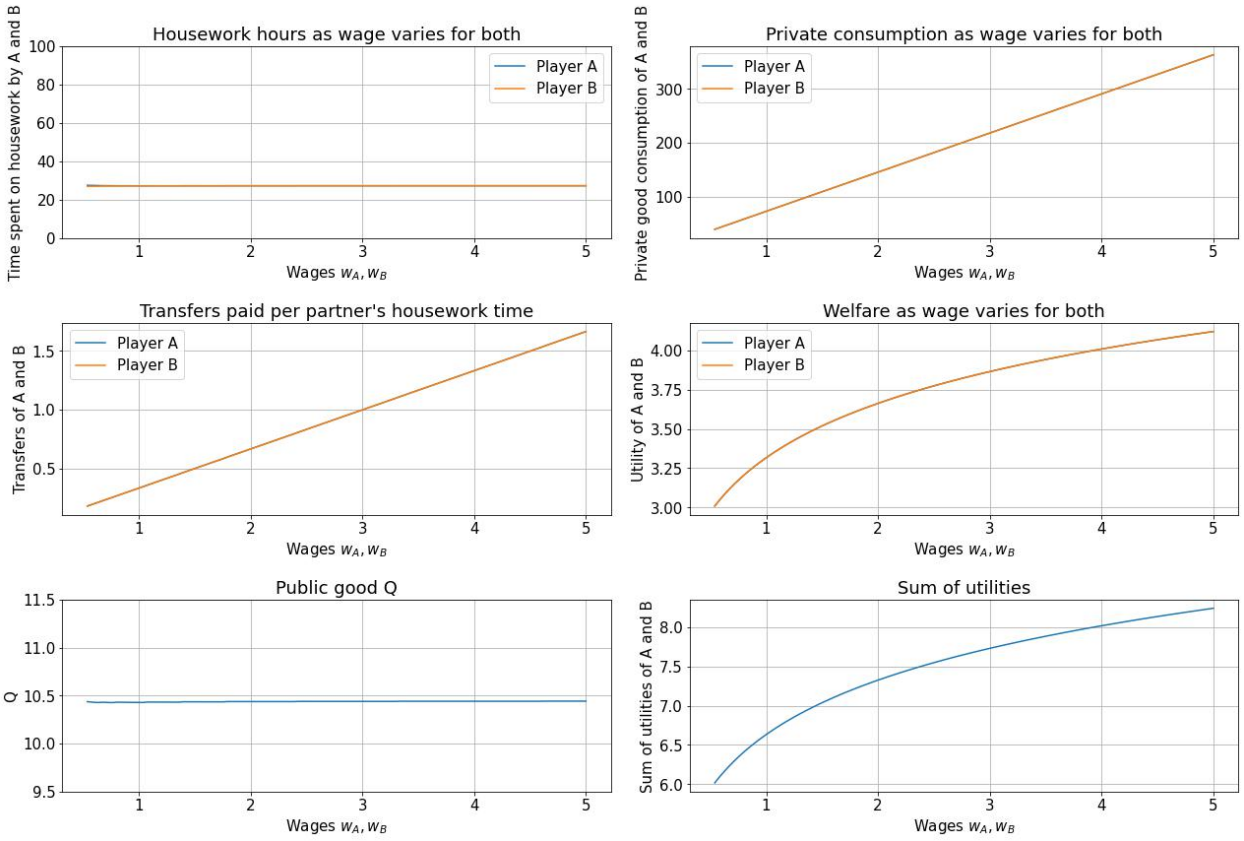


Figure 13: Wages of A and B vary

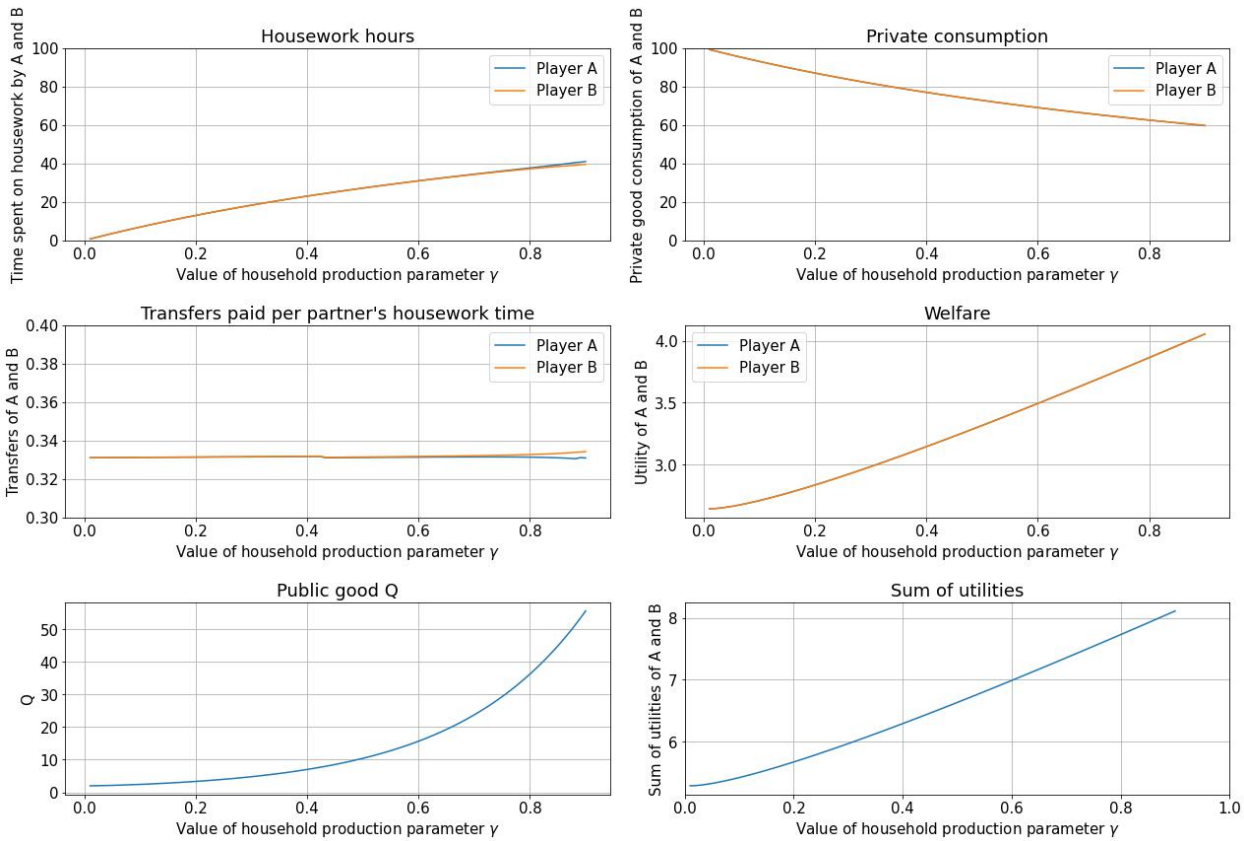


Figure 14: Productivity in household good varies for both partners

B Utility possibility frontier of collective model with household budget constraint

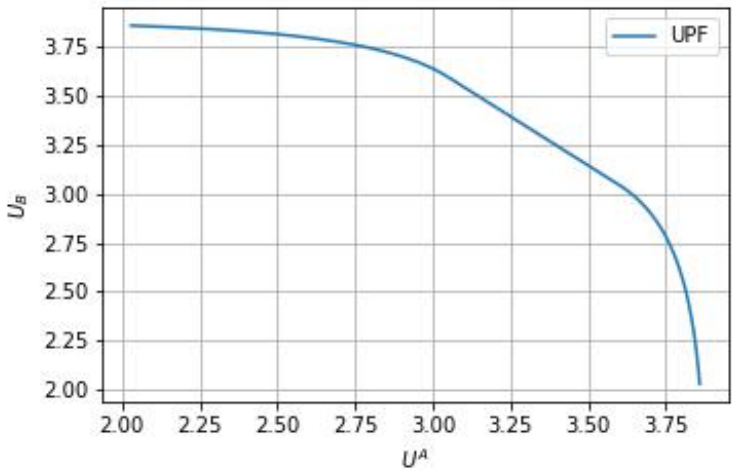


Figure 15: UPF for collective model

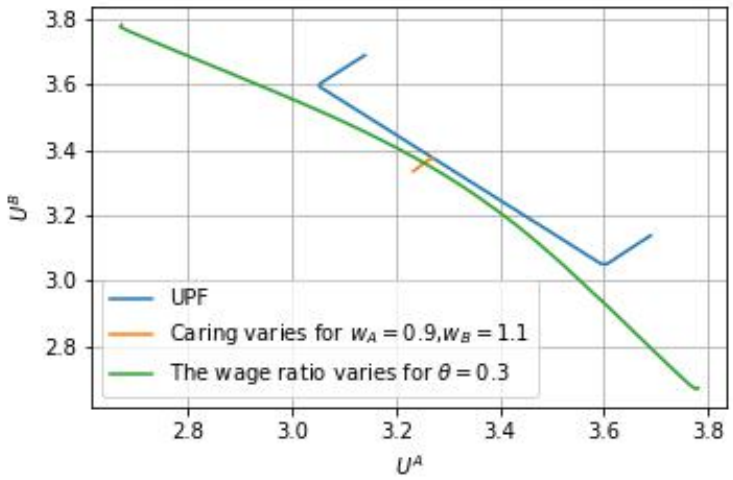
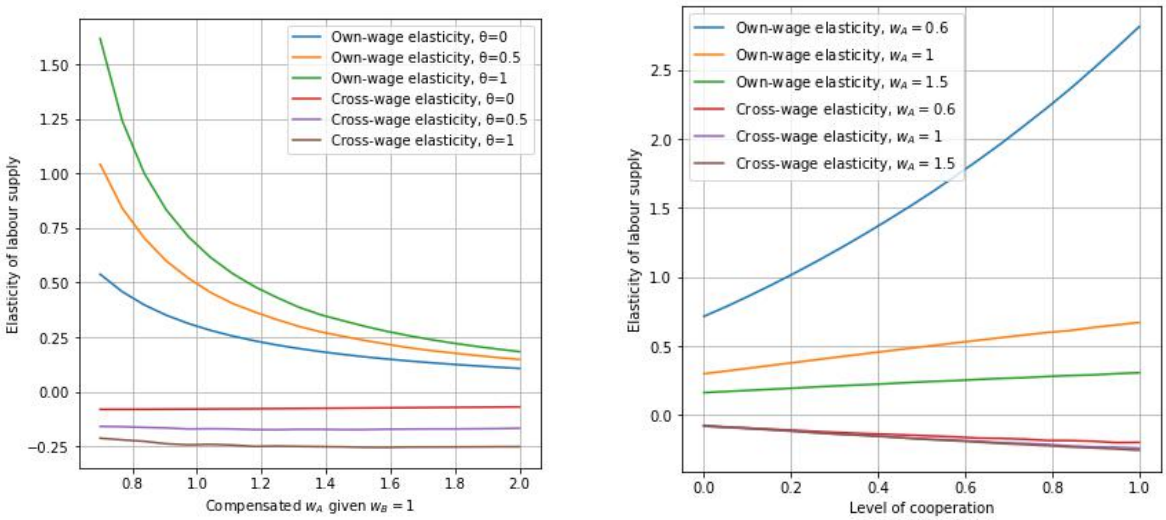


Figure 16: UPF for Lindahl model with wider range of wage ratios

Figure 15 shows the utility possibility frontier for a collective model in which the bargaining weight is equal to the wage ratio. This difference between the Lindahl model and the here presented collective model is that in the Lindahl framework the bargaining weight only adjusts for interior solutions. Therefore, as shown in Figure 16 the utility possibility frontier is kinked at the point at which corner solutions are played. Corner solutions are played in a much wider range of wage ratios when the level of cooperation is high, compared to when it is low.

C Relation between effects of level of cooperation and wage distribution on labour supply elasticities



(a) Level of cooperation

(b) Wage

Figure 17: Own- and cross-wage elasticities by wage distribution and cooperation level

Figure 17a shows the same as Figure 12 but two-dimensionally. The slope and levels of the curves of labour-supply elasticities in the relative wage space depend on the level of cooperation.