RIDGE
RESEARCH INSTITUTE FOR DEVELOPMENT, GROWTH AND ECONOMICS
SINCE 2014

Gender - specific subsidies and female empowerment in optimal taxation

Cassiano Alves, Carlos da Costa, Felipe Lobel, Katia Nishiyama

RIDGE Working Paper 015

May 2025

This paper was presented at the 2024 RIDGE December Forum

Gender-specific subsidies and female empowerment in optimal taxation*

Cassiano Alves[☆], Carlos da Costa[☆], Felipe Lobel [★], and Katia Nishiyama [☆]

[☆] *FGV EMAp School of Applied Mathematics and PUC-RJ*
[☆] *FGV EPGE Brazilian School of Economics and Finance*
[★] *Columbia University*

July 4, 2025

Abstract

Starting with an optimal income-splitting household tax schedule we assess the impact of gender-specific subsidies. Motivated by evidence that spouses' relative earnings influence their power, we let bargaining weights respond to this subsidy and household labor supply choices to vary with weights. Quantitative exploration reveals that a subsidy on women's earnings is welfare-improving, but that neglecting the empowering effect of subsidies greatly underestimates those gains. In our baseline assessment, 99.6% of all women benefit from the policy. For 78% of women, welfare gains are no smaller than 5%, and for 15%, gains exceed 10%. The optimal subsidy for women is about 16% while the benchmark of models where the power channel is neglected is close to 0% with trivial average gains. We find that it is women in the most productive households who benefit the most from this policy.
Keywords: *Intrahousehold Inequality; Joint Taxation; Collective Household* ***JEL Classification:*** *H21, H31, D13.*

*Corresponding author: Cassiano Alves. Email address: cassiano.alves@fgv.br. We thank Rafael Araújo, Andrea Flores, João Guerreiro, and the seminar participants at PUC-RJ, USP-RP, the 2023 GeFam and 2025 Ridge meetings for their comments and suggestions. Da Costa thanks CNPq project 304955/2022-1 for financial support. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. All errors are our responsibility.

1 Introduction

Intrahousehold inequality is a pervasive yet largely untamed phenomenon. Many countries have designed and implemented public policies to address this issue, often in the form of cash transfers targeted either directly or indirectly at women within households. Notable examples include Brazil's *Bolsa Família* and Mexico's *Oportunidades* programs, both of which channel financial resources primarily to women.

These policy initiatives are grounded in the evidence-based understanding that household consumption and resource allocation decisions are shaped by bargaining processes, which are closely tied to the right to command resources among household members. By altering the distribution of resources between spouses, gender-based policies aim to economically empower women, thereby addressing disparities in intra-household decision-making and well-being distribution.

Although the connection between resource command and power has informed policy formulation, it has remains largely unexplored in the optimal tax theory. It is our goal in this paper to address this gap in the literature. We incorporate to an optimal tax problem the direct impact of gender-specific policies on the intra-household distribution of power using a collective approach to household behavior.

Households in our model are 'collective' in the sense that spouses differ with respect to how they rank alternatives but follow decision protocols that lead to efficient choices – see, e.g., Chiappori (1988b); Browning and Chiappori (1998). Their behavior can, therefore, be rationalized as the solution to a Pareto problem using the spouses' individual utilities and Pareto weights, to which the relevant notion of power is attached. Because these weights are not invariant to prices, households are also collective in the sense that they are not Unitary; their behavior need not satisfy the restrictions associated with rational behavior.

We refer to the *empowerment* of a spouse as an increase in her Pareto weight. It is modeled to reflect key empirical regularities in the relationship between wage differentials and intra-household bargaining weights, as documented by, e.g., Lise and Yamada (2019); Flores (2024).[1] An essential feature of this collective view of households, namely, that

[1]In Lise and Yamada (2019), differences in expected wage profiles significantly influence intra-household bargaining weights in the cross-section; realized deviations from expected wages shift bargaining power in favor of the spouse receiving the positive wage shock; and the magnitude of this adjustment increases with the size of the shock. Moreover, there is structural heterogeneity in how bargaining weights respond across households: couples with more symmetric potential earnings display different dynamics than those with

empowerment alters not only how resources are allocated across spouses but also which allocations are preferred by the household, is taken into account in our approach.

To incorporate the current understanding of households as collective units into the optimal taxation of couples, we must face several challenges. First are the technical issues related to the multidimensional nature of household characteristics. Second, although household choices can be rationalized as the solution of a Pareto problem that one can take as representing the 'household preferences', specific policies may affect these Pareto weights, thus leading household behavior to violate the restrictions imposed by rationality – see Browning and Chiappori (1998). Finally, even for policy reforms that do not affect Pareto weights, i.e., for which household behavior satisfies revealed preferences restrictions, these preferences need not have a normative meaning under methodological individualism.

To address these issues, we proceed as follows. First, by assuming that all persons of the same gender have the same preferences, we encode the differences in household choices in a vector $\boldsymbol{\iota} = (w_f, w_m, \alpha)$ in which w_f is the wife's productivity, w_m, the husband's productivity and α the wife's Pareto weight. For income-splitting schedules, that we take as the departure point of our analysis, household preferences for income and consumption are pinned down by $\boldsymbol{\iota}$; a Mirrlees's (1971) economy with three-dimensional types. Importantly, while $\boldsymbol{\iota}$ encodes all differences in household choices, we show that these 'revealed preferences' differ from a Utilitarian normative criterion by a dissonance term that plays a crucial role in formulating the planner's program. Second, we collapse the three dimensions into a single index that represents a sufficient statistic for household choices, by assuming that preferences are iso-elastic and the same for the two genders. Third, because changes in income-splitting schedules do not affect *differently* the prices faced or the unearned income of each spouse, we assume that tax perturbations that preserve income splitting do not change α. This allows us to use $\boldsymbol{\iota}$ as the invariant household type in the derivation of the optimal income-splitting schedule. Fourth, we introduce a subsidy on the wife's earnings and show that its effect on choices and welfare is captured by a change in the household type. Crucially, because this reform increases the return to the wife's effort relative to that of the husband, we allow it to change α. It is through this latter effect that we capture the notion of empowerment. Finally, we rely on the empirical evidence relating the ratio w_f/w_m to α to quantify the importance of empowerment.

highly asymmetric wage profiles.

If the planner has no concerns with intra-household distribution, the aggregation result that allows us to collapse the multi-dimensional household type into a single household productivity transforms the planner's program in a very standard Mirrlees's (1971) program where households substitute for individuals. We might, then, summarize the main question in this paper as whether and under what conditions treating differently the two spouses' earnings help the planner promote its distributive objective. With intra-household equity concerns, differing treatment plays an extra role that we can isolate under our approach. Of course, adding instruments cannot hurt, so the relevant point is what makes the gains quantitatively meaningful. In particular, we focus on the role of empowerment and intra-household equity concerns and find that if either one is not taken into account, there are no relevant benefits from deviating from an income splitting schedule.

For our main numerical assessment, we assume that power is linked to relative productivity, following Lise and Yamada (2019). Since an income-splitting schedule equalizes marginal tax rates across spouses, relative marginal productivity remains unaffected by the tax schedule itself. However, a small direct subsidy increases wives' productivity at the margin relative to their husbands', enhancing their bargaining power through the channel evidenced by Lise and Yamada.

Under our baseline parametrization, the optimal gender-specific subsidy rate is 16%. Evaluated with a utilitarian welfare metric, this policy yields a 0.7% gain in consumption-equivalent terms. This overall gain can be decomposed into a 0.8% improvement in intra-household consumption allocation and a 0.1% efficiency loss in effort allocation. Additionally, we access the incidence by directly assessing the utility impact of the subsidies on women and find that nearly all benefit: 99.6% of women experience welfare gains, with 78% seeing improvements of at least 5% equivalent-consumption, and 15% exceeding 10%. The largest gains accrue to women in affluent households, particularly those with high productivity who are married to equally productive spouses. These findings suggest that, while the policy is effective in addressing intra-household inequality, its limited effect on inter-household inequality may pose a limitation on its scope.

Our results highlight the importance of incorporating the empowerment effect. In fact, if the impact of tax policy on intra-household power dynamics were ignored, the estimated optimal subsidy rate would drop substantially, to just 1.5%. This raises the risk that prevailing policy prescriptions may be significantly misguided. Crucially, ignoring the empowerment effect leads to starkly different conclusions: absent this channel, one would counterfactually infer that over 75% of women are worse off by the optimal subsidy policy, likely resulting in the policy being rejected on gender equity grounds. The most pronounced deviations in welfare outcomes occur among women whose productivity is lower than that of their spouses.

Finally, although we find that gender-based policies are not very powerful if there are no intra-household equity concerns, subsidies to second earners can play an important role in inter-household distribution. We find the optimal subsidy to be 12% in contrast with only 2% for the gender based subsidy. When dissonance is taken into account the optimal subsidy for secondary earners reaches 30%.

The rest of the paper is organized as follows. After a brief literature review, we present the environment in Section 2. Section 3 describes the planner's program. Our quantitative findings are described in Section 4. Robustness and extensions are presented in Section 5. In Section 6, we explain what restrictions of a generaly collective model are implied by our approach. Section 7 concludes. Longer derivations and theoretical results are collected in the Appendix.

Literature Review

Until recently, economic models of household behavior assumed a common utility function, where all resources were pooled to maximize a single objective. This "household welfare function" provided the normative foundation for optimal tax theory. Under this assumption, the key hurdle is the multi-dimensional nature of household heterogeneity. Another layer of complexity is due to the replacement of this Unitary view of households by a Collective one in which the individual preferences of spouses are explicitly considered. Indeed, following the pioneering works by Manser and Brown (1980) and McElroy and Horney (1981), conceptual shifts in family economics – (Chiappori, 1988a, 1992; Apps and Rees, 1988) – emphasized joint decision-making within households shaped by individual incentives.[2] This collective view implies not only that household choices need not satisfy revealed preferences axioms but even when they do, no normative content can be attached to these preferences. Importantly, recent contributions have permitted important

[2]For instance, Lundberg and Pollak (1996) adopted a bargaining framework to model individual agency within marriage, and many of Gary Becker's contributions laid the groundwork for understanding intra-household interactions.

advances in handling the multidimensionality of the problem, but have largely ignored the impact of policy on household dynamics.

To situate our contribution we recognize these three main issues in moving from the optimal taxation of individuals to the optimal taxation of couples, the multidimensionality of household heterogeneity, the inappropriateness of using household choices for welfare assessment, and finally the deviations from the rationality restrictions typically observed in multi-persons household choices, and discuss how the literature has dealt with each one separately.

Multidimensionality Starting from Boskin and Sheshinski (1983), the optimal taxation of couples literature has taken differences between spouses as essential to the problem, meaning that typically household heterogeneity is multi-dimensional. While this has been of little consequence for the early literature that focused on parametrized tax schedules, the move to optimal unrestricted systems has posed the difficult problem of screening with multi-dimensional types. One possibility is the use of perturbation methods. The assumption is that it does not demand restrictions on preferences or the nature of heterogeneity. Whether this is granted is beyond the scope of the present work. From a practical perspective, even if we accept that the method can be applied with all generality, the restrictions on preferences and the relationship between these and the dissonance terms that allows one to use well known sufficient statistics are essentially what we use obtain our aggregation results.

The alternative, screening with multi-dimensional types, is a very challenging problem – Rochet and Choné (1998); Spiritus et al. (2024). Recent work by Golosov and Krasikov (2025) has expanded the possibilities for the analysis of optimal taxation of couples taking this feature into account without restrictions on the space of feasible taxes. This is in contrast with Kleven et al. (2009), who examined nonlinear optimal income taxation for couples in a bargaining framework, where primary earners adjust labor supply and secondary earners only make participation decisions, which reduces the relevant set of tax schedules. Our contribution differs from Golosov and Krasikov's since we restrict taxes to be a combination of an unrestricted tax on household earnings supplemented by a tax/subsidy on women's relative contribution to household earnings, and from Kleven et al.'s since we allow both spouses to respond at the intensive margin.

Several other studies, including Cremer et al. (2012, 2016), examine the challenges of multidimensional household decision-making in taxation. Strong assumptions about preferences are made to reduce the dimensionality to a single index. We handle multi-dimensionality with the combination of income-splitting schedules and the assumption of identical iso-elastic preferences as Heathcote and Tsujiyama (2021); Alves et al. (2024). As in the case of these earlier works, the multiple dimensions collapse into a single index that rationalizes all household choices.

Dissonance By assuming that spouses have iso-elastic and similar preferences the multi-dimensionality of households' heterogeneity collapses into a single willingness-to-work index that rationalizes the cross-sectional marginal propensity to generate income, under an income-splitting schedule.[3] The 'willingness-to-pay/work' index serves as a sufficient statistic for characterizing behavior, but it does not directly inform normative analysis. As a result, standard optimal tax formulae must be adjusted accordingly (see Choné and Laroque (2010); Condorelli (2013); Alves et al. (2024); Akbarpour et al. (2024)).

It is the recognition that individual outcomes may vary in households with similar resources, e.g., Thomas (1990); Calvi (2020), combined with a commitment to methodological individualism that implies that we cannot attach normative content to household choices. To formally incorporate the planner's concern with intra-household distribution, we therefore distinguish between the household utility function as a positive rationalization of household choices and its normative role in policy evaluation—following the tradition of methodological individualism (Chiappori and Meghir (2015)). This misalignment between the household and the planner's objectives parallels the decision utility versus experience utility explored in behavioral economics. Gerritsen (2016) highlights the implications of such misalignment, in a behavioral context, introducing welfare-weighted elasticities in optimal tax formulas derived from perturbation methods. Alves et al. (2024) shows that the same 'behavioral logic' applies for households, despite the fact that all agents are rational.

With dissonance, formulas generated from perturbation methods contains both standard average elasticities and welfare-adjusted elasticities that introduce a new layer of complexity, akin to risk adjustment in finance. This relates to the "smooth bunching" found in Saez (2001), where single-crossing is not imposed. A key novelty is that dissonance generates

[3] It is the combination of our preference restrictions and the application of mechanism design to total household earnings that ensures the tractability of our model. Golosov and Krasikov (2025) adopts a more restrictive class of preferences that allows a characterization of an unrestricted schedule using a direct revelation mechanism.

two distinct aggregated elasticities: the conventional average elasticities and a welfare-weighted counterpart. In Saez's words: "It is not necessary to assume that people earning the same income have the same elasticity; the relevant parameters are simply the average elasticities at given income levels" (2001, p. 210). While government revenues depend on empirical averages, behavioral responses matter for welfare because dissonance invalidates the envelope theorem, making mechanical tax effects insufficient for welfare analysis. In this context, the relevant average elasticity is computed using a welfare-adjusted income distribution, akin to risk-adjustment methods in finance. This aggregation complexity is also central to Gerritsen's work. For our purposes, dissonance all but eliminates the usefulness of perturbation methods.

Empowerment There is now robust evidence of empowerment effects from gender-based distribution policies in cash transfer programs like *Oportunidades* in Mexico and *Bolsa Família* in Brazil, and 'from the pocket to the purse' policy reforms, e.g., Browning and Gørtz (2012). While these reforms typically comprise direct transfers, the evidence in Lise and Yamada (2019) and Flores (2024) speaks more directly to the role of relative productivity in determining spouses' relative power. The family economics literature has, in turn, been slowly replacing the Unitary approach to household behavior with a Collective one.

Golosov and Krasikov (2025), for example, derives optimal taxation for couples using a multi-dimensional mechanism design, with an equal split of utility in a transferable utility setting. By focusing on transferable utility, household market choices are invariant to power. Hence, while in Section 5.4 of their work, they allow the policy to influence intra-household distribution by assuming that the household surplus is split according to a Nash bargain where utilities as singles serve as threat points, the addition of Nash Bargaining does not change optimal tax formulae for couples. Importantly, household behavior cannot be fully accounted for by a Nash bargain approach with external threat points used in their only assessment with a policy impact on the intra-household distribution of power.[4]

Collective households are behavioral agents since their choices need not satisfy the

[4] Earlier Nash-bargain models like Manser and Brown (1980); McElroy and Horney (1981) assumed that threat points were the utilities as singles; external threat points. Lundberg and Pollak (1993) introduced the notion of non-cooperative behavior within marriage to define internal threat points. The latter approach rationalizes the evidence generated by these policies as well as the experimental evidence in Almas et al. (2018); Armand et al. (2020). It also accommodates the impact on household market choices of changes in power as evidenced through violations of income-pooling – Thomas (1990); Duflo (2003) – and symmetry and negative semi-definiteness of household compensated demand matrices – Browning and Chiappori (1998).

rationality restrictions. Taking empowerment into account means accounting for this behavioral aspect of household choices. We handle this aspect by noting that the impact of subsidies on women's earnings can be mapped into a change in household types. That is, in our framework, household heterogeneity is summarized by a vector containing spouses' productivities and relative power. We show that under suitable assumptions on the mapping from prices and incomes to power, the impact of a tax reform can be accounted for by a change in household types. While recent works that address household taxation – Alves et al. (2024); Bierbrauer et al. (2023, 2025), have largely overlooked the potential for tax policy itself to shape power dynamics within households directly, our approach fills this gap by formally introducing the empowerment effect in this mapping. Specifically, we take empowerment into account in our model by attaching a causality interpretation to the findings in Lise and Yamada (2019), following the parallel evidence from Flores (2024).

It is also important to mention that we focus on gender-based policies. Most of the literature do not distinguish between genders – Golosov and Krasikov (2025); Bierbrauer et al. (2024). We discuss the differences in outcomes from subsidizing secondary earners instead of women. Important exceptions are found in Jacquet and Lehmann (2016); Spiritus et al. (2024). In both cases, the optimal taxation of couples is studied as an example of taxation with multi-dimensional types. Spiritus et al. find that the optimum with an unrestricted bi-dimensional schedule can be approximated with an income-splitting schedule augmented with a linear tax on men's earnings.

Our approach is also in the spirit of Bierbrauer et al. (2023), in the sense that we perturb a given schedule to highlight the relevance of a given feature, in our case, the difference in spouses' marginal tax rates. Unlike them, our departure point is always an optimal income-splitting schedule. Our quantitative model, however, is closest to Alves et al.' (2024), which assumes the planner controls only the labor income tax schedule. In contrast, we focus on the role of an auxiliary subsidy that enhances women's earnings and explicitly model how relative bargaining power depends on policy instruments via relative productivity, as in Lise and Yamada (2019).

2 The Environment

The economy comprises a continuum of households (or families) with measures normalized to one. Each household is formed by two spouses, denoted $i \in \{f, m\}$, with f representing the female spouse and m the male spouse. Each spouse possesses an individual labor market productivity (or wage rate) denoted w_i, falling within the range $[\underline{w}, \overline{w}] \subset \mathbb{R}_+$.

In this economy, the technology is linear: one efficient unit of labor yields one unit of consumption good. Agents receive income equal to their marginal productivity, implying that the labor income of the spouse i (with productivity w_i providing l_i hours of work) is given by $z_i = w_i l_i$.

Individuals derive utility from consuming a private good ($c_i \in \mathbb{R}_+$) and experience disutility from supplying labor ($l_i \in \mathbb{R}_+$). Their total utility is a function $\mathscr{U}_i(c_i, l_i)$, which is common to all individuals within each gender category. This utility function $\mathscr{U}_i : \mathbb{R}^2_+ \to \mathbb{R}$ is of class $\mathcal{C}^2$ and satisfies the following properties: strict quasi-concavity, strict increase in consumption, strict decrease in labor supply.

Households We take the multi-person nature of households seriously, modeling them as collective units in the sense of Chiappori (1988a, 1992); Apps and Rees (1988). In this framework, households consist of individuals with well-defined preferences who make joint efficient consumption and labor supply decisions, regardless of the underlying bargaining process. Efficiency implies that decisions lie on the utility possibility frontier, although they may still reflect inherited inequalities in decision-making dynamics, as long as no resources are left unutilized.

Consequently, without any loss of generality, it can be assumed that households, whose members are characterized by productivities (w_f, w_m), make consumption and labor supply decisions $(c_i, l_i)_{i \in \{f,m\}}$ to maximize the following Bergsonian household utility function

$$\alpha \mathscr{U}_f(c_f, l_f) + (1 - \alpha)\mathscr{U}_m(c_m, l_m), \tag{1}$$

for $\alpha \in (0, 1)$.

The parameter α is a Pareto weight, denoting the relative contribution of the female spouse to household welfare. It determines the point at the frontier of the household utility possibility set that the family chooses, thus encapsulating the relevant notion of power for our analysis; the assumption of efficiency implies there is an $\alpha \in (0,1)$ that (combined with spouses' utilities) *rationalizes household choices*. Because we assume that all women and all men have the same preferences, we identify a household with a triple $\boldsymbol{\iota} = (w_f, w_m, \alpha) \in \Lambda \equiv [\underline{w}, \overline{w}] \times [\underline{w}, \overline{w}] \times (0,1)$.

Tax policy Policies influence family choices by restricting/expanding the set of possibilities a family faces at the 'market', and, potentially, influencing α. Taxes are mappings, $\mathcal{T} : \mathbb{R}^2_+ \mapsto \mathbb{R}$, from earnings (z_f, z_m) to post-tax income, c.[5]

Let $\mathcal{B}(w_f, w_m, \mathcal{T})$ be the family's budget set, as determined by the tax schedule $\mathcal{T}$. We allow the schedule to include typical redistribution, universal basic income programs, and gender-targeted transfers.

Typically, what happens within the household, the characteristics of decision protocols, or decisions themselves are neither observable nor enforceable by policy-makers. The household economics literature often assumes that only the total household consumption ($c \equiv c_f + c_m$) is observable externally, as it is not possible to access the share of family consumption going to a specific spouse. This restricts the types of policies that can be implemented.

Income Splitting We take as the starting point of our analysis income-splitting taxes of the form $\mathcal{T}(z_f, z_m) = T(z_f + z_m)$, as in Alves et al. (2024), and later we introduce subsidies/taxes on the earnings of a specific spouse. In this case, the family's budget set may be written $\mathcal{B}(w_f, w_m, T) = \{(c, z) \quad \text{s.t.} \quad c = z - T(z), \quad z = z_f + z_m\}$.

With joint, income-splitting taxation, the government utilizes information on the total household income, $z = z_f + z_m$, to impose a nonlinear joint income tax schedule, $T(z)$. This approach is pertinent, considering its empirical relevance, as in the United States, where couples can file taxes individually or jointly. Under joint taxation, the marginal tax rates depend exclusively on the aggregate income of the household. The progressive nature of the labor income tax schedule means that filing individually is rarely optimal. This approach is also used in countries like Ireland and Germany. Yet, our main concern is how

[5]More generally, we can consider $\mathcal{T} : \mathbb{R}^2_+ \mapsto \mathbb{R}^2$, taking (z_f, z_m) to $(\hat{c}_f, \hat{c}_m)$ to allow for separate filing. This may be important in practice to determine who controls resources – e.g., da Costa and de Lima (2024). One must, however, bear in mind that the mapping from household earnings $c = \hat{c}_f + \hat{c}_m$ to individual consumption c_i, $i = f, m.$, depends only on α. Any influence that the split $(\hat{c}_f, \hat{c}_m)$ induced by separate filings, or any other exogenous assignment, may have on the split (c_f, c_m) must be due to its impact on α.

gender-specific policies that promote secondary earners can improve upon such a system.

The use of joint taxation as a departure point could be justified by its empirical importance. For instance, in the US, couples can file taxes individually or jointly. Under the joint tax option, the marginal tax rates depend exclusively on the aggregate income of the household. The overall progressivity of the labor income tax schedule means that filing individually is seldom optimal.

Household Decision Process Given $\mathscr{B}(w_f, w_m, T)$, the family's budget set implied by the policy, is $c_f + c_m \leq z_f + z_m - T(z_f + z_m)$ which we can decompose in three parts $c_f + c_m \leq c$, $z = z_f + z_m$ and $c \leq z - T(z)$.

Household choices are the solution to

$$\max_{c_f+c_m \leq z_f+z_m-T(z_f+z_m)} \left\{ \alpha \mathscr{U}_f\left(c_f, \frac{z_f}{w_f}\right) + (1-\alpha)\mathscr{U}_m\left(c_m, \frac{z_m}{w_m}\right) \right\},$$

I.e., the **optimal choices** are $c_i(\boldsymbol{\iota}), z_i(\boldsymbol{\iota})$, $i = f, m$; **individual utilities**: $\bar{U}_i(\boldsymbol{\iota}) = \mathscr{U}_i\big(c_i(\boldsymbol{\iota}), z_i(\boldsymbol{\iota})/w_i\big)$, $i = f, m$, and; the **aggregates** $c(\boldsymbol{\iota}) = \sum_{i=f,m} c_i(\boldsymbol{\iota})$, $z(\boldsymbol{\iota}) = \sum_{i=f,m} z_i(\boldsymbol{\iota})$, and

$$\bar{U}(\boldsymbol{\iota}) = \alpha \bar{U}_f(\boldsymbol{\iota}) + (1-\alpha)\bar{U}_m(\boldsymbol{\iota}),$$

for $\boldsymbol{\iota} = (w_f, w_m, \alpha)$. We have omitted $T(\cdot)$ from the mappings to focus on the role of household heterogeneity.

For future reference, let $l_i(\boldsymbol{\iota}) = z_i(\boldsymbol{\iota})/w_i$.

It will be convenient to break the household problem into two stages. First, there is the internal problem. For any market bundle (c, z), it defines the optimal way to share the benefits of consumption, c, and the burden of generating earnings, z, as the solution to

$$U(c, z, \boldsymbol{\iota}) = \max_{(c_f, z_f) \geq (0,0)} \left\{ \alpha \mathscr{U}_f\left(c_f, \frac{z_f}{w_f}\right) + (1-\alpha)\mathscr{U}_m\left(c - c_f, \frac{z - z_f}{w_m}\right) \right\}. \qquad (2)$$

Note that (2) defines, for $i = f, m$, the **policy functions**, i.e., mappings $\mathfrak{c}_i : R_+ \times R_+ \times \Lambda \mapsto R_+$ and $\mathfrak{z}_i : R_+ \times R_+ \times \Lambda \mapsto R_+$ from bundles to internal splits.

The Market Program $U(c, z, \boldsymbol{\iota})$ as defined in (2) is a proper representation of how the household ranks bundles (c, z). This function is optimized by the household facing the tax schedule $T(\cdot)$ at the market stage,

$$\bar{U}(\boldsymbol{\iota}) \max_{c \leq z - T(z)} = U(c, z, \boldsymbol{\iota}). \tag{3}$$

Solving the market problem we have the optimal choices, $c_i(\boldsymbol{\iota}) = c_i(c(\boldsymbol{\iota}), z(\boldsymbol{\iota}), \boldsymbol{\iota})$, and $z_i(\boldsymbol{\iota}) = \mathfrak{z}_i(c(\boldsymbol{\iota}), z(\boldsymbol{\iota}), \boldsymbol{\iota})$.

2.0.1 $U(c, z, \boldsymbol{\iota})$'s (lack of) normative content

Although $U(c, z, \boldsymbol{\iota})$ rationalizes household market choices, we cannot attribute normative content to $U(c, z, \boldsymbol{\iota})$ under methodological individualism; social welfare must be assessed from individual utilities. $U(\cdot, \cdot, \boldsymbol{\iota})$ accounts for intrahousehold allocation dynamics, considering how each spouse contributes to family income and how resources are distributed among family members. The resulting allocations need not reflect intrinsic normative principles. Indeed, from the methodological individualism perspective, following those orderings would lead to a deviation of horizontal equity since power distribution is heterogeneous across families.

Using a Utilitarian criterion the value assigned to a bundle (c_f, l_f, c_m, l_m) in the hands of a $\boldsymbol{\iota}$-household is $\frac{1}{2}\mathscr{U}_f(c_f, l_f) + \frac{1}{2}\mathscr{U}_m(c_m, l_m)$, which, in the context of an income-splitting schedule, leads to a normative utility,

$$V(c, z, \boldsymbol{\iota}) = \frac{1}{2}\mathscr{U}_f\left(c_f(c, z, \boldsymbol{\iota}), \frac{\mathfrak{z}_f(c, z, \boldsymbol{\iota}))}{w_f}\right) + \frac{1}{2}\mathscr{U}_m\left(c_m(c, z, \boldsymbol{\iota}), \frac{\mathfrak{z}_m(c, z, \boldsymbol{\iota})}{w_m}\right).$$

Next, recalling that $c = z - T(z)$, define the *dissonance index*, ξ, through the total derivative of V with respect to z,

$$\frac{dV}{dz} = \frac{\partial V(c, z, \boldsymbol{\iota})}{\partial c}[1 - T'(z)] + \frac{\partial V(c, z, \boldsymbol{\iota})}{\partial z} = [1 - \xi(c, z, \boldsymbol{\iota})]\,\frac{\partial V(c, z, \boldsymbol{\iota})}{\partial c}[1 - T'(z)],$$

where

$$\xi(c, z, \boldsymbol{\iota}) := \frac{\partial V(c, z, \boldsymbol{\iota})/\partial z}{\partial V(c, z, \boldsymbol{\iota})/\partial c}\left[\frac{\partial U(c, z, \boldsymbol{\iota})/\partial z}{\partial U(c, z, \boldsymbol{\iota})/\partial c}\right]^{-1}.$$

2.1 The Optimal Tax Problem

Let us disregard, for now, the use of subsidies, s. Then the planner's program is to choose $T(\cdot)$ to maximize the sum of agents' utilities. The planner's objective is

$$\int V(c(\boldsymbol{\iota}), z(\boldsymbol{\iota}), \boldsymbol{\iota}) d\Phi(\boldsymbol{\iota}),$$

where $\Phi(\boldsymbol{\iota}) = \Pr\big((w_f, w_m, \alpha) \leq \boldsymbol{\iota}\big)$ is the distribution of household types, and the budget constraint is

$$\int T\left(z(\boldsymbol{\iota})\right) d\Phi(\boldsymbol{\iota}) \geq G,$$

for some exogenously given revenue requirement, G.

At this moment, it is important to note that we used $\boldsymbol{\iota}$ as a family type in stating the planner's program. Yet, strictly speaking, only w_f and w_m are structural parameters since α is allowed to vary with the policy under the Collective Approach – Chiappori (1988b); Browning and Chiappori (1998). In this context, the assumption we are making is that, because changes in the income-splitting schedule affects equally the prices faced by both spouses it does not affect α.[6]

2.1.1 Gender-based earnings subsidies.

Starting from an optimal income-splitting schedule $T(z_f+z_m)$, we introduce a subsidy s on women's earnings. It is worth mentioning that Spiritus et al. (2024) find that a combination of a joint tax and a linear reduction of marginal taxes (i.e., a linear subsidy) on secondary earners to approximate very well the unconstrained optimum.

The implied budget set becomes

$$\mathcal{B}(w_f, w_m, T, s) = \big\{(c, z_f, z_m) \quad \text{s.t.} \quad c = z_f(1+s) + z_m - T(z_f(1+s) + z_m)\big\}.$$

Let $\hat{z}_f = l_f w_f(1+s)$, and $\hat{z} = l_f w_f(1+s) + l_m w_m$. If we ignore for the moment the potential impact of s on α then, for any couple $\boldsymbol{\iota}$ we can write (2) exactly as before with $w_f(1+s)$ substituting for w_f in $\boldsymbol{\iota}$. More generally, for every s, let $i_s : \Lambda \mapsto \Lambda$ be the mapping from original to modified types, $i_s(\boldsymbol{\iota}) = \hat{\boldsymbol{\iota}}$, for $\hat{\boldsymbol{\iota}} = \big(w_f(1+s), w_m, \hat{\alpha}\big)$, where

[6] In Section 6 we use a general Collective function for α to make our assumption explicit with reference to the literature and to discuss the empirical evidence that supports it.

we use the notation $\hat{\alpha}$ to allow the policy to affect the Pareto weights. Empowerment takes the form of an increase in the Pareto weight α that women have in the decisions taken by the family. For the moment, we are agnostic about the mechanism through which subsidies affect power, noting only that, in contrast with changes in the income-splitting schedule we are now introducing a policy that affect the relative (to the husband's) return to labor the wife's market effort; more control over the generation of household resources may affect the bargaining position of women.[7,8]

Household choices are now rationalized by $U(c, z, i_s(\boldsymbol{\iota}))$ and the planner's assessment for the same (c, z), becomes

$$V(c, z, i_s(\boldsymbol{\iota})) = \frac{1}{2}\mathscr{U}_f\left(c_f(c, z, i_s(\boldsymbol{\iota})), \frac{\mathfrak{z}_f(c, z, i_s(\boldsymbol{\iota}))}{w_f[1+s]}\right) \\ + \frac{1}{2}\mathscr{U}_m\left(c_m(c, z, i_s(\boldsymbol{\iota})), \frac{\mathfrak{z}_m(c, z, i_s(\boldsymbol{\iota}))}{w_m}\right).$$

If $\alpha < 1/2$, and if $\hat{\alpha} > \alpha$ for $s > 0$, then s can be used to bring household choices closer to the one that maximizes the planner's criterion. The cost of such policy must be taken into account in the planner's budget constrain that now has an explicit recognition of average female earnings $\mathbb{E}\big[\mathfrak{z}_f(z, c, i_s(\boldsymbol{\iota})) \mid z(i_s(\boldsymbol{\iota})) = z\big]$, for every household earnings, z. The planner's budget constraint, therefore, becomes

$$\int T\left(z(i_s(\boldsymbol{\iota})) + s\mathfrak{z}_f(c(i_s(\boldsymbol{\iota})), z(i_s(\boldsymbol{\iota})), i_s(\boldsymbol{\iota}))\right) d\Phi(\boldsymbol{\iota}) \geq \\ G + s\int \mathfrak{z}_f(c(i_s(\boldsymbol{\iota})), z(i_s(\boldsymbol{\iota})), i_s(\boldsymbol{\iota}))d\Phi(\boldsymbol{\iota}).$$

As we will explain, by relying on the $i_s(\cdot)$ mapping to transformed types, all the tractability explored by Heathcote and Tsujiyama (2021); Alves et al. (2024) is preserved with the introduction of s. This allows us to incorporate the impact of such a policy on women's empowerment, which has been mostly neglected by the literature.

[7] Supporting this view is the panel evidence from Lise and Yamada's (2019) work with Japanese data and Flores' (2024) assessment of reforms on the *Oportunidades* program in Mexico.

[8] An interesting parallel can be drawn between these subsidies and the EITC program, as both policies operate through labor supply channels. Gender-based subsidies are anticipated to enhance women's welfare by boosting their participation in the labor market. Additionally, we uncover a new channel by recognizing the potential impacts of these policies on intra-household bargaining.

2.2 Specializing Preferences

The planner's problem presents two additional challenges when compared to Mirrlees's (1971). First, there is the multidimensionality of types. Second is the fact that the utility that appears in the planner's objective is not the one (or a monotonic transformation of the one) that represents household preferences. A convenient assumption about preferences, namely that they are separable iso-elastic, and equal for the two genders, allows us to circumvent both issues.

If preferences are separable, $\mathscr{U}_i(c,l) = u_i(c) - h_i(l)$, the family program can be split into two separate allocation problems: how to allocate resources toward the consumption of the spouses conditional on the budget for a household consumption level of c,

$$\mathcal{U}(c,\alpha) = \max_{c_f,c_m} \Big\{ \alpha u_f(c_f) + (1-\alpha)u_m(c_m) \quad \text{s.t.} \quad c_f + c_m = c \Big\}, \tag{4}$$

and how to allocate the efforts towards generating income, z, in the least costly manner,

$$\mathscr{H}(z,\boldsymbol{\iota}) = \min_{\{l_f,l_m\}} \Big\{ \alpha h_f\left(l_f\right) + (1-\alpha)h_m\left(l_m\right), \quad \text{s.t.} \quad w_f l_f + w_m l_m = z \Big\}. \tag{5}$$

With separability, the policy function $c_i(c,\alpha)$, $i = f, m.$, does not depend on z, and the policy function $\mathfrak{z}_i(z,\boldsymbol{\iota})$, $i = f, m.$, does not depend on c.

The solution to (4) satisfies $\alpha u_f'(c_f) = [1-\alpha]u_m'(c_m)$: the ratio of marginal utilities is determined by the power ratio $\alpha/(1-\alpha)$. It is an application of Borch's rule to this simple environment. Therefore, it is only through its impact on α that policy can influence the household sharing rule, $c_i(c,\alpha)$. The solution of the second program is the set of policy functions, $\ell_i(z,\boldsymbol{\iota})$, $i \in \{f,m\}$, that depends on α but also (w_f, w_m).

When preferences are not only separable but also iso-elastic and identical for the two genders, the policy functions have a closed form. Importantly, the share of consumption $c_i(c,\alpha)/c$ that spouse i gets is independent of c and the share of earnings, $\mathfrak{z}_i(z,\boldsymbol{\iota})/z$ that spouse i must generate is independent of z. That is, there are functions $\bar{c}_i(\alpha)$ and $\bar{\mathfrak{z}}_i(\boldsymbol{\iota})$ such that $c_i(c,\alpha) = \bar{c}_i(\alpha)c$ and $\mathfrak{z}_i(z,\boldsymbol{\iota}) = \bar{\mathfrak{z}}_i(\boldsymbol{\iota})z$. Of course, the latter also implies the existence of a function $\bar{\ell}_i(\boldsymbol{\iota}) = \bar{\mathfrak{z}}_i(\boldsymbol{\iota})/w_i$ such that $\ell_i(z,\boldsymbol{\iota}) = \bar{\ell}_i(\boldsymbol{\iota})z$.

Indeed, let $u_i(\cdot) = (c_i^{1-\sigma} - 1)/(1-\sigma)$, $i = f, m.$, and $h_i(l) = l^{1+\gamma}/[1+\gamma]$, $i = f, m$, the policy functions are such that

$$\bar{c}_f(\alpha) = \left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{\sigma}} \bar{c}_m(\alpha) \quad \text{and,} \quad \bar{l}_f(\boldsymbol{\iota}) = \left[\left(\frac{1-\alpha}{\alpha}\right)\left(\frac{w_f}{w_m}\right)\right]^{1/\gamma} \bar{l}_m(\boldsymbol{\iota}).$$

Relative earnings z_f/z_m are independent of z, hence, of the tax schedule in place. For all (z, c), both the allocation of effort and the allocation of consumption between spouses will always be along the same line from the origin.

The main advantage of assuming identical iso-elastic preferences is, however, that it simplifies the screening problem by collapsing the three-dimensional types into a single index. As in Alves et al. (2024); Heathcote and Tsujiyama (2021), the technical challenges imposed by multi-dimensional types that only now the literature has started to handle Golosov and Krasikov (2025); Bierbrauer et al. (2023) are overcome with the combination of income-splitting schedules and iso-elastic preferences.

With $\mathscr{U}_i(c, l) = \ln c - l^{1+\gamma}/(1+\gamma)$, $i = f, m$, (4) becomes

$$\mathcal{U}(c, \alpha) = \alpha \ln \alpha + (1-\alpha)\ln(1-\alpha) + \ln c,$$

and (5),

$$\mathscr{H}(z, \boldsymbol{\iota}) = \frac{1}{1+\gamma}\left(\frac{z}{\mathcal{w}(\boldsymbol{\iota})}\right)^{1+\gamma},$$

where

$$\mathcal{w}(\boldsymbol{\iota}) = \left[\alpha^{\frac{-1}{\gamma}} w_f^{\frac{1+\gamma}{\gamma}} + (1-\alpha)^{\frac{-1}{\gamma}} w_m^{\frac{1+\gamma}{\gamma}}\right]^{\frac{\gamma}{1+\gamma}}.$$

In other words, a $\boldsymbol{\iota}$-couple behaves as if it were an individual with preferences that are identical to those of each spouse but with a labor market productivity $\omega = \mathcal{w}(\boldsymbol{\iota})$; ω is a *sufficient statistic* for the household behavior under iso-elastic preferences and income-splitting taxes. While $\boldsymbol{\iota}$ still carries relevant information for the normative assessment of allocations, all the behavioral content is condensed in $\omega = \mathcal{w}(\boldsymbol{\iota})$, a fact explored by Alves et al. (2024) that bears direct relation to the misalignment between social value and willingness to pay that drives the results in Choné and Laroque (2010); Condorelli (2013); Akbarpour et al. (2024), among others.

Define the distribution $\Psi(\omega) = \Phi(\boldsymbol{\iota}|\mathcal{w}(\boldsymbol{\iota}) \leq \omega)$, and, for every ω, the set $\mathscr{I}(\omega) \equiv \{\boldsymbol{\iota} \in \Lambda | \mathcal{w}(\boldsymbol{\iota}) = \omega\}$ of all family types $\boldsymbol{\iota}$ whose market behavior is summarized by the same ω.

Then, for all ω,

$$\mathscr{v}(c, z, \omega) \equiv \left[V(c, z, \boldsymbol{\iota})\middle|\boldsymbol{\omega} \in \mathcal{I}(\omega)\right], \quad \text{and} \quad \mathscr{u}(c, z, \omega) \equiv \left[U(c, z, \boldsymbol{\iota})\middle|\boldsymbol{\iota} \in \mathcal{I}(\omega)\right].$$

Next, if we average $\xi(c, z, \boldsymbol{\iota})$ across all $\boldsymbol{\iota} \in \mathcal{I}(\omega)$ we define a 'ω-household average wedge',

$$\bar{\xi}(\omega) \equiv \mathbb{E}\left[\xi(c, z, \boldsymbol{\iota})\middle|\boldsymbol{\iota} \in \mathcal{I}(\omega)\right],$$

which we will refer to simply as the ω-household wedge. Similarly, if we let $k(\boldsymbol{\iota}) = \mathfrak{z}_f(z, \boldsymbol{\iota})/z$, then[9]

$$\bar{k}(\omega) \equiv \mathbb{E}\left[k(z, \boldsymbol{\iota})\middle|\boldsymbol{\iota} \in \mathcal{I}(\omega)\right].$$

For the identical iso-elastic preferences specification that we are using, we can express the planner's utility as

$$\mathscr{v}(c, z, \omega) = \mathscr{u}(c, z, \omega) + \mathcal{Q}(\omega) + \frac{1 - \bar{\xi}(\omega)}{1 + \gamma}\left(\frac{z}{\omega}\right)^{1+\gamma},$$

where $\mathcal{Q}(\omega) = \mathbb{E}\left[[0.5 - \alpha]\ln\alpha + [\alpha - 0.5]\ln(1 - \alpha)\middle|\boldsymbol{\iota} \in \mathcal{I}(\omega)\right]$. Note that while $\mathcal{Q}(\omega)$ plays no role in the planner's program when policy instruments are unable to change α – e.g. Alves et al. (2024) — this term will be important as we allow subsidies to directly affect α.

While invariant to the income-splitting schedules, relative earnings and relative consumption change when we introduce a subsidy on the wife's earnings. Studying the impact of empowerment policies is made very simple in the iso-elastic case,

$$\frac{\bar{c}_f(\hat{\alpha})}{\bar{c}_m(\hat{\alpha})} = \left(\frac{\hat{\alpha}}{1 - \hat{\alpha}}\right)^{\frac{1}{\sigma}} \quad \text{and,} \quad \frac{\bar{\ell}_f(\hat{\boldsymbol{\iota}})}{\bar{\ell}_m(\hat{\boldsymbol{\iota}})} = \left[\left(\frac{1 - \hat{\alpha}}{\hat{\alpha}}\right)\left(\frac{w_f(1 + s)}{w_m}\right)\right]^{1/\gamma}.$$

3 The Mirrlees' Program.

Following Mirrlees's (1971) lead, we can approach the problem using the mechanism design machinery. From this perspective, family characteristics $\boldsymbol{\iota}$ define its type, which is each family's private information, and the planner's objective is to achieve redistribution

[9] Recall that ω is a sufficient statistic for z.

within and between families by utilizing a direct mechanism and the subsidy rate.

The direct mechanism is as follows. The planner asks each household its type, $\boldsymbol{\iota}$, and assigns a bundle $(z(\boldsymbol{\iota}), c(\boldsymbol{\iota}))$. Telling the truth is a best strategy provided that for all $\boldsymbol{\iota}, \boldsymbol{\iota}'$, $U(c(\boldsymbol{\iota}), z(\boldsymbol{\iota}), i_s(\boldsymbol{\iota})) \geq U(c(\boldsymbol{\iota}'), z(\boldsymbol{\iota}'), i_s(\boldsymbol{\iota}))$.[10] In other, words, the planner maximizes

$$\int V(c(\boldsymbol{\iota}), z(\boldsymbol{\iota}), i_s(\boldsymbol{\iota})) d\Phi(\boldsymbol{\iota}),$$

subject to the incentive constraint specified above and the resource constraint.

While stating the problem is rather easy the multidimensionality of family types along with the non-standard fact that the planner's utility is not a monotonic transformation of the one that rationalizes the family choices poses a real challenge that is easily overcome with the assumption we have made regarding agents' preferences.

3.1 Solving the Mirrlees's Program

To explain the mechanism design problem, it is useful to refer to Alves et al. (2024). There, the optimal design of joint (income-splitting) is studied under the assumption that policies cannot be used directly to target a specific spouse. We depart from them by assuming that the planner can subsidize wives at a constant rate s.

Mechanism Design in Alves et al. (2024) At $s = 0$, considered by Alves et al., the planner's objective is

$$\max_{(x(\cdot), z(\cdot))} \int_{\underline{\omega}}^{\bar{\omega}} \left\{ \mathcal{u}(c(\omega), z(\omega), \omega) + \mathcal{Q}(\omega) + \frac{1 - \bar{\xi}(\omega)}{1+\gamma} \left(\frac{z(\omega)}{\omega} \right)^{1+\gamma} \right\} \psi(\omega) d\omega,$$

the incentive compatibility constraint is

$$\omega \in \operatorname*{argmax}_{r} \mathcal{u}(c(r), z(r), \omega).$$

[10] We use the notation $x(\boldsymbol{\iota})$ noting that, given the one to one mapping $i_s(\cdot)$, we could have also used $x(i_s(\boldsymbol{\iota}))$.

and the resource constraint,

$$\int_{\underline{\omega}}^{\bar{\omega}} [z(\omega) - c(\omega)]\, \psi(\omega) d\omega \geq G.$$

Note that if α is invariant to policy, then $\mathcal{Q}(\omega)$ does not affect the planner's program. Now, the solution of this program is the optimal schedule for the case in which the planner can only use a joint (income-splitting) schedule to address both inter- and intra-household distributive concerns.

Letting $v(\omega) := \max_r \boldsymbol{u}(c(r), z(r), \omega)$, one can write the incentive compatibility constraint as

$$\dot{v}(\omega) = \frac{z^{1+\gamma}}{\omega^{2+\gamma}},$$

plus monotonicity and solve the optimal taxation problem using Mirrlees's optimal control approach.[11]

The modified income tax program ($s \neq 0$). As previously discussed, we incorporate subsidies to Alves et al.'s (2024) program in a straightforward manner by transforming the distribution of household types through the mapping $i_s(\boldsymbol{\iota}) = (w_f(1+s), w_m, \hat{\alpha})$. That is, we write the modified program by recalling that we can use the mapping i_s to define the new distribution of types $\hat{\Phi} : \Lambda \mapsto [0,1]$ through the composition $\hat{\Phi}_s = \Phi \circ i_s$. Figure 1 displays the new cumulative distribution of household productivity induced by the subsidy s —$\hat{\Psi}_s(\omega) = \Phi\left(\boldsymbol{\iota} \middle| \boldsymbol{w}(i_s(\boldsymbol{\iota})) \leq \omega\right)$ — and without $\Psi(\omega) = \Phi\left(\boldsymbol{\iota} \middle| \boldsymbol{w}(\boldsymbol{\iota}) \leq \omega\right)$ — subsidies.

The main simplification obtained by Alves et al. (2024) is to reduce the multidimensional type $\boldsymbol{\iota}$ to a summarizing unidimensional statistic $\omega = \boldsymbol{w}(\boldsymbol{\iota})$. Because household market choices can be summarized by ω, incentive constraints are written in ω, and household preferences are shown to satisfy single-crossing. However, the fact that two different families $\boldsymbol{\iota}$ and $\boldsymbol{\iota}'$ may have the same ω, $\boldsymbol{w}(\boldsymbol{\iota}) = \boldsymbol{w}(\boldsymbol{\iota}') = \omega$, matters for welfare assessment.

For every ω, Alves et al. (2024) define $\hat{\mathcal{I}}_s(\omega) \equiv \left\{\boldsymbol{\iota} \in \Lambda \mid \boldsymbol{w}(i_s(\boldsymbol{\iota})) = \omega\right\}$, the set of all $\boldsymbol{\iota}$-family whose productivities are ω, and use it to write, $\bar{\xi}_s(\omega) = \mathbb{E}^{\hat{\psi}_s}\left[\xi(\boldsymbol{\iota}) | \boldsymbol{\iota} \in \hat{\mathcal{I}}_s(\omega)\right]$, the average dissonance for families with subsidy-adjusted productivity ω. For our purposes, it will also be important to define $\bar{\mathcal{Q}}_s(\omega) = \mathbb{E}\big[[0.5-\alpha]\ln\alpha + [\alpha - 0.5]\ln(1-\alpha)\big|\boldsymbol{\iota} \in \hat{\mathcal{I}}_s(\omega)\big]$

[11] As most of the literature we use a first-order approach, dropping the monotonicity constraint and verifying 'ex-post' whether it is satisfied.

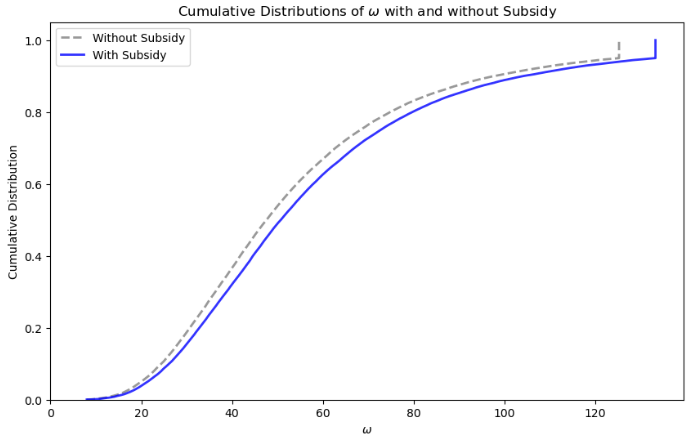

Figure 1: The figure displays the distribution of household productivity with subsidies and without subsidies, ω, obtained from the distribution of $\boldsymbol{\iota}$ (gray) and $i_s(\boldsymbol{\iota})$ (blue). The distribution with subsidies first-order stochastically dominates the distribution without subsidies.

the average consumption value misalignment for families with cum subsidy productivity ω and, $\bar{k}_s(\omega) = \mathbb{E}^{\hat{\psi}_s}\left[k(\boldsymbol{\iota})|\boldsymbol{\iota} \in \hat{\mathcal{I}}_s(\omega)\right]$, the average share of household earnings produced by wives in families with (after-subsidy) aggregate productivity ω.

With $\hat{z}$ replacing z, the program is exactly as before except for the resource constraint that becomes[12]

$$\int_{\underline{\omega}}^{\bar{\omega}} \left[\frac{\hat{z}(\omega)}{1+\bar{k}_s(\omega)s} - \mathcal{c}\left(v(\omega), \hat{z}(\omega), \omega\right)\right] \hat{\psi}_s(\omega) d\omega \geq G,$$

where $\mathcal{c}(v, z, \omega)$ is implicitly defined by $v = \mathcal{U}(\mathcal{c}(v, z, \omega), z, \omega)$, $\forall\,(v, z)$, and ω. Crucially, the additive term $\bar{\mathcal{Q}}_s(\omega)$ plays no role in the derivation of the optimal income tax schedule but is key for the optimal subsidy.

Solving the planner's program, we obtain the **optimal labor income tax schedule**,

$$\frac{\tilde{T}'\left(z(\omega)\right)}{1-\tilde{T}'\left(z(\omega)\right)} = \frac{1+\gamma}{\hat{\psi}_s(\omega)\omega} \int_{\omega}^{\bar{\omega}} \left\{\frac{\tilde{c}(t)}{\tilde{c}(\omega)} - \mathbb{E}^{\hat{\psi}_s}\left[\frac{\tilde{c}(t)}{\tilde{c}(\omega)}\right]\right\} \hat{\psi}_s(t) dt - \mathbb{E}^{\hat{\psi}_s}\left[\frac{\tilde{c}(t)}{\tilde{c}(\omega)}\right]\left[1-\bar{\xi}(\omega)\right],$$

where

$$\tilde{T}'(z(\omega)) = 1 - \left[1+\bar{k}_s(\omega)s\right]\left[1-T'(z(\omega))\right],$$

and

$$\tilde{c}(\omega) = c(\omega)\left[1+\bar{k}_s(\omega)s\right].$$

At $s = 0$, the formula is exactly that in Alves et al. (2024). As for the optimal subsidy,

[12]Bearing this in mind, we slightly abuse notation by keeping the letter z to represent earnings. We let $\hat{\psi}_s$ represent the density associated with the distribution $\hat{\Psi}_s$.

again, it is worth emphasizing the role of $\mathcal{Q}(\omega)$: it does not change the design of an optimal income tax schedule, but it is crucial in the choice of an optimal subsidy, s.

3.1.1 Choosing s

Start from the solution to the planner's program as posed in Alves et al. (2024). What is the impact of introducing a subsidy? We can characterize this impact by differentiating the Lagrangian at $s = 0$,

$$\left.\frac{d\mathcal{L}}{ds}\right|_{s=0} = \int_{\underline{\omega}}^{\bar{\omega}} \left\{ v(\omega) + \mathcal{Q}(\omega) + \left[1 - \bar{\xi}(\omega)\right] \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \frac{d\psi(\omega)}{ds} d\omega$$
$$+ \lambda \int_{\underline{\omega}}^{\bar{\omega}} \left[z(\omega) - \mathcal{c}\left(v(\omega), z(\omega), \omega\right)\right] \frac{d\psi(\omega)}{ds} d\omega$$
$$+ \int_{\underline{\omega}}^{\bar{\omega}} \left\{ \frac{d\mathcal{Q}(\omega)}{ds} - \frac{d\bar{\xi}(\omega)}{ds} \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \psi(\omega) d\omega - \lambda \int_{\underline{\omega}}^{\bar{\omega}} z(\omega)\bar{k}_0(\omega)\psi(\omega) d\omega. \quad (6)$$

For any income-splitting schedule, $\boldsymbol{\iota}$ contains all information for the choices a $\boldsymbol{\iota}$-household makes. Assumption A guarantees that one may freely vary $T(\cdot)$ without changing $\boldsymbol{\iota}$. The impact of a small subsidy on the wife's earnings, in contrast, changes household choices as if one had changed the household type through the mapping i_s. Hence, through its impact on the distribution of $\boldsymbol{\iota}$ from Φ to $\hat{\Phi} = \Phi \circ i_s$, we can trace all the changes in the economy caused by the use of s.

First, since $\Psi_s(\omega) = \Phi\big(\boldsymbol{\iota} \mid \mathcal{w}(i_s(\boldsymbol{\iota})) \leq \omega\big)$, by introducing s the planner causes the distribution of ω to change – see Figure 1 – and this has a direct impact on the value of the Mirrlees' program that is displayed in the first two lines of (6). Next, because we index spouses by the aggregator $\omega = \mathcal{w}(\boldsymbol{\iota})$, we have to take into account the heterogeneous responses of families that share the same ω despite having different $\boldsymbol{\iota}$. The terms $d\mathcal{Q}(\omega)/ds$ and $d\bar{\xi}(\omega)/ds$ in the last line of (6) play this role by capturing the average change in the expected value of each of these variable across $\boldsymbol{\iota}$-families sharing the same ω.[13]

The last term of the third line captures the mechanical cost of increasing s converted in social utility by λ.

A common presumption in gender-based policies is that greater control over earnings

[13]Or, more simply, while the choices conditional on $\boldsymbol{\iota}$ do not change, the distribution of $\boldsymbol{\iota}$-families sharing the same ω does.

goes to whoever the earnings accrue. The evidence supporting this view is overwhelming – Thomas (1990); Duflo (2003); Browning and Gørtz (2012); Almas et al. (2018); Flores (2024). Policies aimed at exploring this aspect of household behavior are equally pervasive – Oportunidades (Mexico), Bolsa Família (Brazil). We take this empirical regularity into account by letting the Pareto weight $\hat{\alpha}$ vary with s. If, however, one disregards the empowerment effect, the expression is very similar to (6): only the $d\mathcal{Q}(w)/ds$ disappears. Besides, by ignoring changes in α one would use wrong values for $d\psi(w)/ds$ and $d\bar{\xi}(\omega)/ds$.In Section 4, after we introduce a formal model for the mapping from subsidies to α, we compare the semi-elasticity of ω to s at $s = 0$ for the two cases to provide a sense of this effect.

Inter-household Redistribution Only If we disregard for the moment dissonance, i.e., let $\bar{\xi}(\omega) = 1$ and $\mathcal{Q}(\omega) = 0$ for all ω, then, the first two terms define the total change in the value of the planner's program that results from a cost-free exogenous change in the distribution, $\Psi_s(\omega)$ parametrized by s. The first line is the direct impact on utilities, whereas the second is the impact on resources converted into utility by λ. In practice, the reform is not cost-free, and the first term in the last line captures the direct cost of introducing a small subsidy. Finally, the last term captures the impact of s on dissonance.

Now, it may still be possible for a subsidy to facilitate redistribution across households. This occurs when the condition 7, below, does not hold.

$$\frac{1}{\lambda}\int_{\underline{\omega}}^{\bar{\omega}} v(\omega)\frac{d\psi(\omega)}{ds}d\omega + \int_{\underline{\omega}}^{\bar{\omega}} \left[z(\omega) - \boldsymbol{x}\left(v(\omega), z(\omega), \omega\right)\right]\frac{d\psi(\omega)}{ds}d\omega$$
$$= \int_{\underline{\omega}}^{\bar{\omega}} z(\omega)\bar{k}_0(\omega)\psi(\omega)d\omega. \quad (7)$$

The Optimal Subsidy Differentiating in s the planner's Lagrangian and setting it to 0 yields the following necessary condition for an optimum,

$$\begin{aligned}\frac{d\mathcal{L}}{ds} = \int_{\underline{\omega}}^{\bar{\omega}} \left\{ v(\omega) + \mathcal{Q}(\omega) + \left[1 - \bar{\xi}(\omega)\right] \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \frac{d\psi(\omega)}{ds} d\omega \\ + \lambda \int_{\underline{\omega}}^{\bar{\omega}} \left[\frac{z(\omega)}{1+\bar{k}_s(\omega)s} - \mathcal{c}\left(v(\omega), z(\omega), \omega\right)\right] \frac{d\psi(\omega)}{ds} d\omega \\ - \lambda \int_{\underline{\omega}}^{\bar{\omega}} \left[\frac{z(\omega)}{\left[1+\bar{k}_s(\omega)s\right]^2} \left[\bar{k}_s(\omega) + \frac{d\bar{k}_s(\omega)}{ds} s\right]\right] \psi(\omega) d\omega \\ + \int_{\underline{\omega}}^{\bar{\omega}} \left\{ \frac{d\mathcal{Q}(\omega)}{ds} - \frac{d\bar{\xi}(\omega)}{ds} \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \psi(\omega) d\omega = 0. \quad (8)\end{aligned}$$

The term $d\bar{k}_s(\omega)/ds$ that appears in the third line captures the average change in women's participation in household earnings across families that share the same ω. At $s \neq 0$ we must take into account the fiscal effect of changes in the female share of earnings. The last line measures the impact on intra-household inequality that arises due to changes in how consumption and effort are shared between spouses. $\mathcal{Q}(\omega)$ and $\bar{\xi}(\omega)$ both change with s since, for all ω the distribution of $\boldsymbol{\iota}$ for which $\mathcal{w}(\boldsymbol{\iota}) = \omega$ changes when we apply $i_s(\cdot)$.

Inter-household Redistribution Only For a planner who is not concerned with intra-household inequality, the necessary condition for an optimum is

$$\begin{aligned}\frac{d\mathcal{L}}{ds} = \int_{\underline{\omega}}^{\bar{\omega}} v(\omega) \frac{d\psi(\omega)}{ds} d\omega + \lambda \int_{\underline{\omega}}^{\bar{\omega}} \left[\frac{z(\omega)}{1+\bar{k}_s(\omega)s} - \mathcal{c}\left(v(\omega), z(\omega), \omega\right)\right] \frac{d\psi(\omega)}{ds} d\omega \\ - \lambda \int_{\underline{\omega}}^{\bar{\omega}} \left[\frac{z(\omega)}{\left[1+\bar{k}_s(\omega)s\right]^2} \left[\bar{k}_s(\omega) + \frac{d\bar{k}_s(\omega)}{ds} s\right]\right] \psi(\omega) d\omega = 0. \quad (9)\end{aligned}$$

This expression is useful to disentangle the role of equity concerns.

4 Quantitative Exercises

To implement the approach we have presented we need to specify the mapping $i_s(\cdot)$; more precisely, how for every $\boldsymbol{\iota} = (w_f, w_m, \alpha)$, $\hat{\alpha}$ in $i_s(\boldsymbol{\iota}) = (w_f[1+s], w_m, \hat{\alpha})$ is determined.

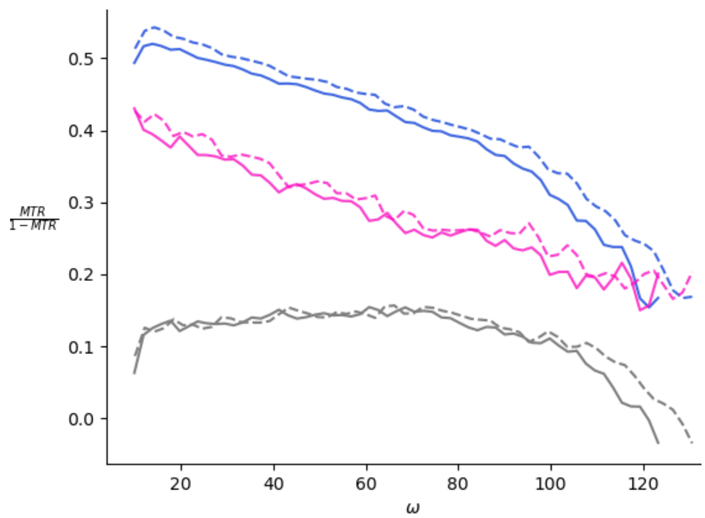

*Figure 2: **Optimal Taxes:** The figure displays in blue the marginal optimal tax rate for the cases with (continuous line) and without (dashed line) subsidies. The pink line captures the Pigouvian correction, and the gray line, the Mirrleesian term.*

A 'Theory' of Power Lise and Yamada (2019), assume that relative power is determined by the ratio of wages, (potentially) a vector of controls $\boldsymbol{X}$, and a shock $\epsilon \sim \mathcal{N}(0, \sigma^2)$, orthogonal to $\boldsymbol{X}$ and $\ln w_f - \ln w_m$,[14]

$$\ln \frac{\alpha}{1-\alpha} = \beta_0 + \beta_1 \ln \frac{w_f}{w_m} + \beta_2' \boldsymbol{X} + \epsilon$$

Equivalently, household weights, α, are drawn from a distribution,

$$\alpha = \frac{\exp\{\beta_0 + \beta_1 \ln(w_f/w_m) + \beta_2' \boldsymbol{X} + \varepsilon\}}{1 + \exp\{\beta_0 + \beta_1 \ln(w_f/w_m) + \beta_2' \boldsymbol{X} + \varepsilon\}}$$

whose parameters β_0, β_1, and β_2 one estimates from the data.

We adopt Lise and Yamada's (2019) specification and estimate this equation using the March 2016 CPS to estimate the model in the status quo, $s = 0$, and for each family in our sample, we replicated the other 29 with different love shocks ($\sim$ 450,000 families). In our estimation, we use no controls $\boldsymbol{X}$. At mean α we find that a 10% increase in $\ln w_f - \ln w_m$ leads to a 4.9% increase in α. In Section 5.1 we compare our findings with the rest of the literature and run some robustness exercises.

Our statistical approach clarifies that the productivity ratio predicts the relative power with a positive coefficient. A subsidy changes the ratio from w_f/w_m to $w_f(1+s)/w_m$. We adopt a causal interpretation for the relationship so that, by changing the productivity ratio,

[14] Under the specification for preferences that we are using, this is equivalent to assuming that $\alpha/(1-\alpha)$ is a function of z_f/z_m, if we suitably re-interpret the coefficients in the regression.

the subsidy, s, changes the ratio $\alpha/(1-\alpha)$ through $w_f(1+s)/w_m$. More precisely, we assume that

$$\ln\frac{\hat{\alpha}}{1-\hat{\alpha}} - \ln\frac{\alpha}{1-\alpha} = \beta_1 \ln[1+s]. \tag{10}$$

Thus, for every couple, $\boldsymbol{\iota} = (w_f, w_m, \alpha)$ we have $\imath_s(\boldsymbol{\iota}) = (w_f[1+s], w_m, \hat{\alpha})$, where $\hat{\alpha}$ is related to α through (10).[15]

The relationship between $\hat{\alpha}$ and α described in (10) allows us to express the semi-elasticity of ω with respect to s at $s=0$ as

$$\left.\frac{d\hat{\omega}}{ds}\right|_{s=0}\frac{1}{\omega} = (1+\gamma)\alpha\beta_1 + k(\boldsymbol{\iota})\frac{1+\gamma-\beta_1}{1+\gamma}, \tag{11}$$

where $k(\boldsymbol{\iota})$ is the share of household earnings generated by the wife. If $\beta_1 = 0$, the right-hand side of (11) is simply $k(\boldsymbol{\iota})$, which reminds us that, with no empowerment effect, such a policy has a higher impact on families in which the wife's earnings are relatively more important. In fact, provided that $\beta_1 < 1+\gamma$, the empirically relevant case, the semi-elasticity is still increasing in k, but the slope is less than one. For two couples with the same k, the impact is larger when the wife has more power, for this implies that the wage ratio w_f/w_m is higher.

It is fair one may question our causal interpretation. As an extreme example, assume that productivity depends on physical strength and that the latter defines power. By subsidizing women's work we cannot make them stronger, so there is no empowerment. The evidence uncovered by Lise and Yamada (2019) using panel data that *changes* in productivity change power for married agents suggests this cannot be the whole story. Compelling evidence of a causal effect is also provided by Flores (2024) using data from Mexico's *Oportunidades*. We adopt the causal interpretation while bearing in mind these caveats.

A Mechanism for the Household The main motivation for our work is the recognition that even when we can represent household choices as if made by a single rational individual, couples are still comprised of two different persons who order the different bundles and the possibilities for distributing consumption goods and effort among themselves differently. This being the case, what do we mean by 'asking the household its type'? The

[15] In our numeric experiment, for each household, $\boldsymbol{\iota}$, we hold ε fixed and substitute $w_f(1+s)$ for w_f. Note that adding controls does not change (10), if $\ln(w_f/w_m)$ is orthogonal to $\mathbf{X}$.

household is not a well-defined person, but an organization not too different from a firm. This begs the question of who we are asking this question. The wife? The husband? Does it matter?

While we believe this is an important issue, we will not attempt to answer it. We follow the rest of the literature – e.g. Alves et al. (2024); Golosov and Krasikov (2025) – by assuming that whoever is contacted by the planner acts on behalf of this collective household whose preferences are represented by $U(\cdot, \cdot, \boldsymbol{\iota})$.

4.1 Optimal Subsidy and Aggregate Welfare Gains

The design of gender-focused policies requires a careful assessment of their impact on both aggregate welfare and intra-household inequality. In this section, we present the optimal subsidy levels derived from our model and evaluate their effects on overall social welfare. Our analysis highlights the importance of considering intra-household bargaining dynamics when designing policies aimed at reducing gender disparities.

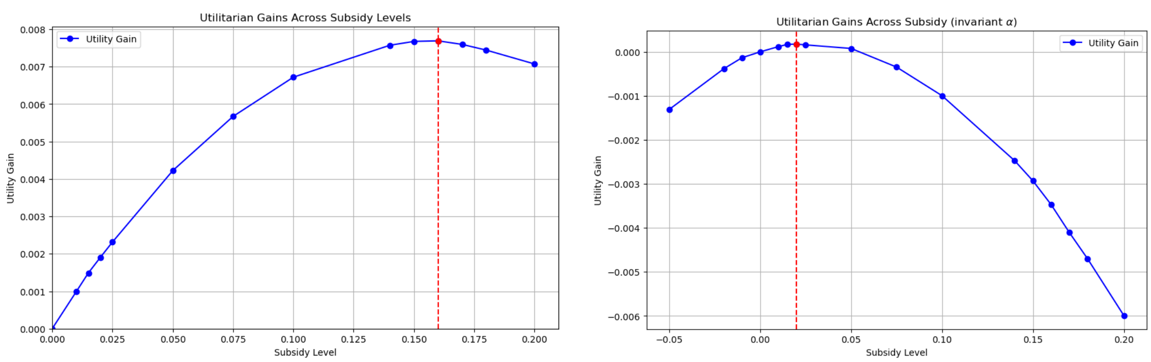

Figure 3: The figure displays the Utilitarian Gains for different subsidy levels (s). The right panel displays the utilitarian gains in the benchmark economy, where subsidies affect both household resources and the intra-household bargaining power (α). The left panel shows the same analysis under the counterfactual assumption that subsidies do not influence α, isolating the resource redistribution effect. Comparing both panels highlights the critical role of the empowerment channel in reducing intra-household inequality and improving social welfare.

Figure 3 illustrates the optimal subsidy levels in both the benchmark case (right panel) and the counterfactual economy where the power channel is deactivated (left panel). Note that after changing the distribution of $\boldsymbol{\iota}$, we adjust $T(\cdot)$ optimally.

Our findings highlight the empowerment channel as the most effective mechanism for reducing intra-household inequality.

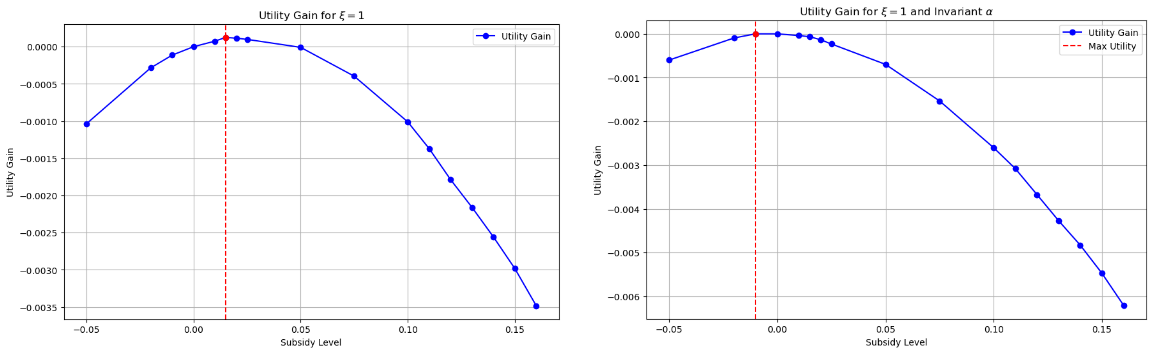

Figure 4: The figure on the left displays the utility gains for a planner who assigns welfare weights α' for women ($1-\alpha'$ for men) in every household with $\boldsymbol{\iota} = (w_f, w_m, \alpha')$. The figure on the right does the same, holding α fixed.

The key takeaway from our analysis is the crucial role of power dynamics in shaping policy outcomes. Accounting for these dynamics is not only essential for assessing the optimality and desirability of gender-focused policies but also for understanding the behavioral responses to traditional tax policies.

Pure efficiency concerns A key feature of how we state the planner's problem is the misalignment between the planner and the household's objective, dissonance. If we do not consider dissonance, that is, if we assume that the weight placed by the planner for a woman in a household with Pareto weights α and $1-\alpha$ is exactly α, then $\xi(\omega) = 1$ for all ω and $\mathcal{Q}(\omega) = 0$.

We examine this possibility – see Figure 4 – and find that the optimal subsidy is positive but small. If one counterfactually ignores the impact of subsidies on α, then one would suggest a small tax, $s < 0$, instead.

The welfare gains from optimal policy design are evaluated from the planner's perspective, but it is also important to consider the incidence effects on the targeted group of women. Figure 5 presents the cumulative distribution of women's gains, measured as the equivalent variation in consumption, under both the baseline and counterfactual scenarios. Our results indicate that virtually all women in our sample benefit from the empowerment policy. In our baseline assessment, at the optimal subsidy level, 99.6% of women experience welfare gains. Moreover, 78% gain at least 5%, while 15% see gains exceeding 10%. The top 1% of beneficiaries achieve gains as high as 16%. The optimal subsidy rate for women is approximately 16%, whereas, in models that ignore the bargaining power channel, the estimated optimal subsidy is close to 0%, yielding only trivial average gains. The right panel depicts the counterfactual scenario where the policy does not influence α. Here, the distribution shifts downward, with 25% of women experiencing a utility loss of 3% and only 1% achieving a modest gain of 5%.

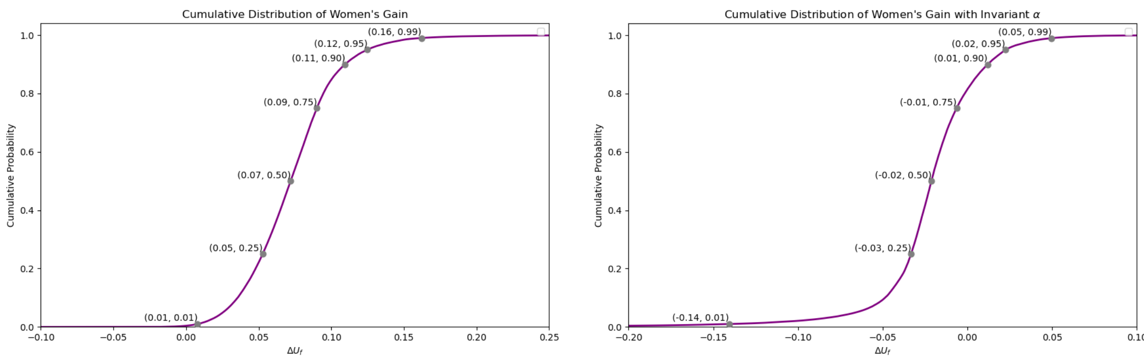

Figure 5: The figures display the Cumulative Distribution of Women's Utility Gains in both scenarios. The left panel shows the distribution when the policy influences intra-household bargaining power (α) with key percentiles highlighted: 25% of women gain at least 5%, 50% gain at least 7%, and the top 1% gain up to 16%. The right panel presents the counterfactual scenario in which policy does not affect α, leading to lower or even negative utility gains for some women – for instance, 25% experience a loss of 3%, and the top 1% gain only up to 5%. The comparison highlights the role of the empowerment channel in generating significant welfare improvements for women.

Figure 6 provides a stark account of what drives the difference. Empowerment allows women to command a higher share of household consumption. This has a direct consequence on how utility varies with s. This is not the whole story, of course. Empowerment is also consequential for how effort is shared, but it is however confounded by the increase in productivity.

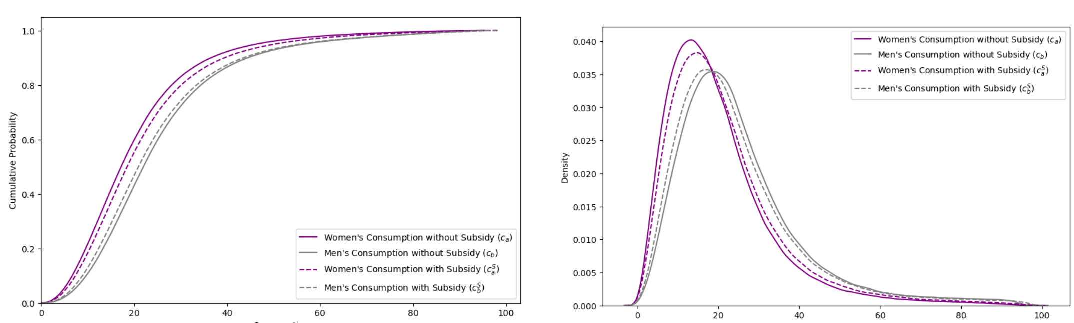

Figure 6: The figure displays cumulative distribution (left panel) and density (right panel) of consumption for men and women with (dotted line) and without 16% subsidies (straight line).

4.2 Household heterogeneity and the effect of Subsidies

While the cumulative distribution of women's utility gains provides valuable insights into the policy's impact, a deeper understanding of intra-household dynamics is essential to uncover the mechanisms driving these results. In this section, we examine how the subsidy affects spouses in different households, including those that may share a common ω. This micro-level perspective allows us to better understand the nuanced ways in which the empowerment channel operates.

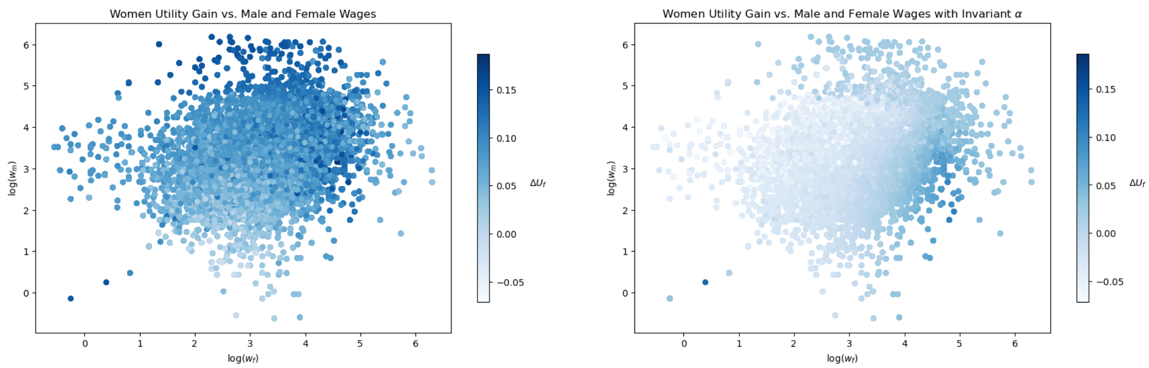

*Figure 7: **Distribution of Women's Utility Gains.** The left panel displays women's utility gains when the policy affects intra-household bargaining power (α). Gains are concentrated among women married to higher-wage husbands, as they can negotiate a larger share of resources in households with greater economic surplus. The right panel shows the counterfactual scenario where the policy does not affect α.*

The left panel of Figure 7 reveals that the utility gains from the policy are not evenly distributed across all women. Instead, they are predominantly concentrated among women married to high-productivity husbands. This pattern arises because households with higher-productivity men typically have a larger economic surplus, enabling women to secure a more significant share of resources when their bargaining power is enhanced.

In contrast, the right panel shows the counterfactual scenario where the policy does not affect intra-household bargaining power (α). In this case, the distribution of women's utility gains shifts significantly. Women married to lower-productivity husbands who benefited little from the policy now experience even smaller gains or losses. The results highlight the critical role of the empowerment channel: when α is fixed, the policy fails to benefit women in households with a lower economic surplus, emphasizing the need to consider both resource allocation and bargaining power in empowerment policies.

Figure 13 presents the net difference in utility gains for women when intra-household

Figure 8: ***Changes in Bargaining Power (α) Across Household Wage Structures.*** *The figure illustrates how changes in intra-household bargaining power (α) vary across households with different wage structures. When male and female wages are more balanced, changes in bargaining power are more pronounced. In contrast, households with large wage disparities exhibit smaller changes in bargaining power.*

bargaining power (α) is allowed to vary compared to when it remains fixed. The largest (positive) differences occur for women who are relatively less productive than their husbands. If we ignore adjustments to bargaining power, we underestimate relatively more the impact of subsidies for women who earn significantly less than their spouses. Interestingly, the most important changes in α — Figure 8 — take place along the $45°$ line, i.e., for women whose productivities are very similar to their husbands': as one empowers the wife of a rich man, she benefits relatively more than a woman of identical productivity married to a poorer man.

Figure 15 displays the impact on women's labor supply while Figure 16 illustrates the impact of the subsidy on women's consumption (Δc_f) under two scenarios: the baseline, where the subsidy affects intra-household bargaining power (α), and the counterfactual, where α remains constant. It is mostly highly productive women married to not-so-productive men who increase their labor supply. If one ignores empowerment, then one predicts that most women will work harder, even those married to more productive men, while most men would reduce their labor supply — Figure 18.

As for consumption, the largest absolute increases are attained by wealthy women. The concentration among these women is expected and mirrors the decrease in their husbands' consumption — Figure 17. Interestingly, if one were to ignore empowerment, then only women in households that are more highly subsidized, i.e., those with highly productive women, would experience relevant increases.

Finally, it is worth mentioning that a common criticism of joint taxation is that it discourages the labor supply of secondary earners, typically women. The presumption is that higher earnings are related to a higher utility. The left panel in Figure 19 casts some doubt on this presumption. Most women lower their labor supply and experience utility gains. Moreover, it is exactly those women who increased labor supply the least that gained more utility. Increases in α dampen the impact that a woman's productivity gain has on her labor supply (most women end up working less) and increase their consumption. An increase in effort is a strong signal of a small change in α. Interestingly, in the counterfactual world, women who increase labor supply are more likely to have experienced positive utility gains.

5 Extensions and Robustness

5.1 Robustnes: How s changes α.

A critical parameter in our approach is the elasticity of α with respect to (the log of) spouses' relative net productivity. In our baseline specification, we used a value of 0.88 relating to the cross-sectional coefficient for the regression of log Pareto weights ratio on log relative wages.

Table 1 displays the optimal subsidy for β's for various values of β. The optimal subsidy increases with β but remain high even for the lowest values we consider.

Table 1: Optimal subsidy and Utilitarian gain for different values of β

β	Optimal Subsidy	Utilitarian Gain
0.00	0.03	0.0002
0.30	0.08	0.0018
0.40	0.11	0.0028
0.50	0.13	0.0037
0.60	0.14	0.0047
0.70	0.15	0.0057
0.80	0.15	0.0068
0.88	0.16	0.0077
0.90	0.16	0.0079
1.00	0.16	0.0089
1.10	0.16	0.0098

The choice of range for β is largely influenced by the findings in Lise and Yamada

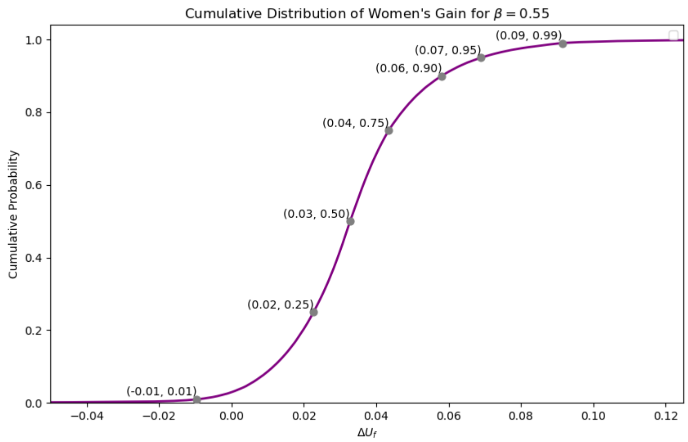

Figure 9: The figure depicts the distribution of welfare gains for wives under the optimal subsidy, $s = 0.13$, for $\beta = 0.55$. Compared to the case with no empowerment ($\beta = 0$), the distribution shifts substantially to the right. Nearly all women benefit: 99% experience positive welfare gains, and 50% gain at least 3%. While the top 1% of women gain up to 9%, even women in the lower tail face only modest losses of about 1%.

(2019). Lise and Yamada decomposes the total effect of wage ratio into ratio at the time of marriage, expected ratio growth, and unexpected changes in ratio. The estimated coefficients, 0.48, 0.11, and 0.38, respectively. The combined coefficient, 0.97, is slightly higher than what we have used.

One may, however, consider that only unexpected changes should be included. This short-run perspective is consistent with holding the distribution of couples fixed, whereas the other two are expected to influence who marries whom and whether people marry in the first place. This would lead one to use $\beta = 0.38$. The coefficients in Table 1 which include the lowest (short run) and highest end (combined value) coefficients found in Lise and Yamada (2019) makes it clear that even for the short run impact the optimal subsidy, $\approx 11\%$, is substantially higher than what one would obtain if one assumed no empowerment.

Flores (2024) provides what is perhaps the closest experiment for the short-run impact that our model captures. The elasticity at the average value of α corresponds to $\beta = 0.55$, for which the estimated optimal subsidy is $s = 0.13$. Figure 9 exhibits welfare gains for women associated with this value of β. In contrast, the no-empowerment case - illustrated in Figure 5 (right panel), where α remains constant and correspond to $\beta = 0$ - yields a substantially lower optimal subsidy, and the associated welfare gains are minimal. Over 25% of women experience utility losses in that scenario, and fewer than 1% gain more than 5%. By contrast, with $\beta = 0.55$, the distribution shifts decisively: the entire mass of women moves into positive gains, and substantial improvements become more widespread. It is apparent that ignoring the empowerment channel leads to sub-optimal policies.

5.2 Secondary Earners

We have focused in this paper on gender-based policies. In contrast, many of our optimal tax references are to works that address gender-blind policies that treat agents differently depending on their relative earnings but not their gender. To assess how sensitive our findings are to this alternative, in this section, we consider subsidies to secondary earners, independently of gender. That is, we split the subsidy between spouses in such a way that the subsidy always accrues to the spouse with lower earnings – see Appendix B.

We start our discussion by comparing the outcomes of the two policies at the same rate $s = 0.16$, the optimum for the gender-based subsidy. In contrast with the gender-based subsidy, some women experience a large welfare loss: primary earner wives. At the top, i.e., for those who experience gains one does not observe difference. This is somewhat expected, as women who are secondary earners are equally empowered in both policies. However, the lower tail differs sharply. As shown in Figure 5, under the gender-based subsidy, nearly all women experience welfare gains. In contrast, under the secondary earner policy, more than 25% of women experience a welfare loss, driven by the reduction in their Pareto weight when their spouses receive the transfer. It is important to note that the behavior we describe is despite the income tax schedules differing for the two policies.

To interpret the results pertaining to the optimal policy it is important to note that the same subsidy rate, s, on the secondary earner translates to a very different change in the distribution of household productivity than a subsidy on wives. That is, everything constant, for a household that has the wife as secondary earner the two policies are equivalent, but for families in which the primary earner is the wife the gender-based subsidy is strictly better.

Baring this in mind, we find a substantially higher optimal subsidy rate under the secondary earner rule, at $s = 0.33$. The distribution of welfare gains is considerably more dispersed relative to the gender-based case. Among women, the bottom 25% experience losses exceeding 7%, with the most negatively affected losing up to 27% of their consumption-equivalent utility. These are primarily high-productivity women who are also primary earners and see their Pareto weights decline when the transfer is directed to their lower-earning spouses. On the other extreme, such high subsidy on secondary earner wives leads 1% of women experience gains over 30%. The policy thus redistributes power and welfare asymmetrically within couples, depending on who the secondary earner is. Compared to the gender-based design, this policy improves gender symmetry in access to the subsidy, but

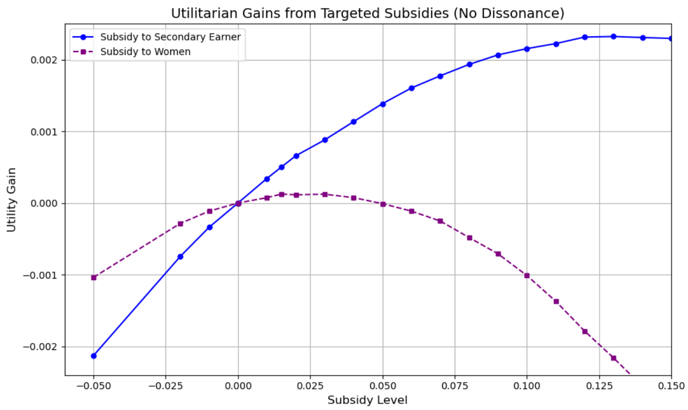

Figure 10: The figure displays the welfare gains for different level of subsidies for women and for secondary earners. Here Utilitarian means equal weight for households respecting the household internal weights.

at the cost of greater within-gender inequality and a more polarized distribution of welfare changes.

We next ask what the optimal subsidy would be in the absence of intra-household equity concerns. That is, we consider a planner who fully disregards the distribution within couples by setting $\xi = 1$ and $\mathcal{Q} = 0$, while still accounting for the empowerment channel. In this case, the optimal subsidy to secondary earners is $s = 13\%$, substantially higher than the $s = 1.5\%$ found for the gender-based subsidy under the same assumptions. This illustrates that even in a purely inter-household redistribution problem, policies that target the secondary earner remain effective when bargaining responses are present. Moreover, the welfare gains are substantially higher when subsidies are given to secondary earners independently of gender.

To further isolate the role of empowerment, we consider a counterfactual in which the subsidy does not affect intra-household bargaining, that is, $\beta = 0$, under a purely utilitarian planner with $\xi = 1$ and $\mathcal{Q} = 0$. In this case, the optimal subsidy is substantially lower, at $s = 8\%$. The associated welfare gains are modest: the median woman gains approximately 1%, and only 1% of women experience losses exceeding 6%.

This exercise highlights a central message of our analysis: models that abstract from shifts in intra-household bargaining, as in the case of $\beta = 0$, substantially underestimate both the optimal level of intervention and its distributional impact. In our framework, the empowerment channel is not an add-on but a key mechanism. The contrast between this counterfactual and our baseline results illustrates the importance of explicitly modeling shifts in Pareto weights when evaluating family-level policies.

6 Tax Reforms and Empowerment

In this Section, we rephrase our approach in a more standard Collective language to better situate our contributions within the broader household economics literature.

$\boldsymbol{\iota}$'s invariance to $T(\cdot)$. The questions we address here are related to how taxes change prices and incomes and how these impact power. That is, can we hold $\boldsymbol{\iota}$ fixed as we change $T(\cdot)$ in the optimal tax exercise? How about s? For the first question, recall that spouses' productivities are invariant parameters, so the substantive assumption we are making in writing the planner's program is that α does not vary with changes in the income-splitting schedule $T(\cdot)$. The question we ask is whether this is granted.

The tradition for the Collective approach is to be agnostic as to what determines the household Pareto weights, allowing α to be a function of prices, income, and distribution factors. Browning and Chiappori (1998), for example, provides evidence that prices can affect α, causing household demand systems to violate the rationality restrictions of individuals and household unitary models.[16] For our purposes, this is crucial for it suggests that α will be, at least, in part dependent on the policy we study. Yet, the collective approach, in all its generality, lacks the structure for an assessment of arbitrarily non-linear tax policies, which may explain why most inquiries into the optimal taxation of couples ignore or overly simplify the empowerment effect.

To understand what is at stake for the problem we study, we start by noting that a household consumes three goods, or more precisely one good, c, and two bads, l_f and l_m, whose prices are 1, $w_f(1-\mathcal{T}_f')$ and $w_m(1-\mathcal{T}_m')$, respectively. Note that when the tax system is non-linear, the prices of effort are endogenous. Moreover, for this discussion, we have allowed the marginal tax rates to differ between spouses.

We can, therefore, write α as a function of the price vector $(1, w_f(1-\mathcal{T}_f'), w_m(1-\mathcal{T}_m'))$, the households' full income, b, and a vector of distribution factors, $\mathbf{d}$, e.g., sex ratio, spouses unearned income ratio, etc. We let a denote the mapping from full income, prices, and

[16] Browning and Chiappori (1998) shows that household choices are *not* characterized by symmetry and negative semi-definiteness of the compensated demand matrix, but instead by SR1, the sum of a symmetric semi-definite and a rank one matrix.

distribution factors to the wife's Pareto weight,[17]

$$\alpha = \mathcal{a}\left(b, 1, w_f(1-\mathcal{T}_f'), w_m(1-\mathcal{T}_m'), \mathbf{d}\right).$$

While the Collective approach does not impose many restrictions on which variables or how the variables affect α, a natural (innocuous) assumption generally adopted is the absence of money illusions, i.e., only real variables matter -see Almås et al. (2023).

Beyond this, we note that only the prices $w_i(1-\mathcal{T}_i')_{i=f,m.}$, are spouse-specific and that spouse-specific incomes are all incorporated to $\mathbf{d}$ and the substantive assumption that it is only those that affect α.

Assumption A Let $\alpha = \mathcal{a}\left(b, 1, w_f(1-\mathcal{T}_f'), w_m(1-\mathcal{T}_m'), \mathbf{d}\right)$ then, for all $(w_f(1-\mathcal{T}_f'), w_m(1-\mathcal{T}_m'), \mathbf{d})$,

$$\mathcal{a}\left(\hat{b}, p, w_f(1-\mathcal{T}_f'), w_m(1-\mathcal{T}_m'), \mathbf{d}\right) = \alpha,$$

for all $\hat{b}$, $p > 0$.

In the absence of money illusion, this means that neither real unearned income (or full income) nor the relative (to aggregate effort) price of consumption affects Pareto weights if the productivity ratio $w_f(1-\mathcal{T}_f')/w_m(1-\mathcal{T}_m')$ and the distribution factors $\mathbf{d}$ are held fixed. We are assuming that anything that affects prices or incomes for the two spouses *equally* does not change their bargaining power.[18] This allows us to write, with some abuse,

$$\alpha = \mathcal{a}\left(\frac{w_f(1-\mathcal{T}_f')}{w_m(1-\mathcal{T}_m')}, \mathbf{d}\right). \tag{12}$$

When marginal tax rates are allowed to differ between the two spouses, the assessment of any tax reform must take the impact on α into account. If the joint schedule is, however,

[17]This formulation accommodates different full incomes for the two spouses in $\mathbf{d}$. Or, to take the example of individual filing, one could imagine that the ratio of spouses' virtual incomes $(T'(z_f)z_f - T(z_f))/(T'(z_m)z_m - T(z_m))$ would be important. By including this ratio in $\mathbf{d}$, we do not rule out its role in affecting α.

[18]Using data from the Japanese Panel Survey of Consumers (JPSC), Lise and Yamada (2019) find that the impact of full income is small – a 10% increase in full income leads to a 0.16% change in the wife's Pareto weight, from 0.0438 to 0.0439. As for the relative price of consumption (marginal income taxes), they do not investigate this possibility.

of the income-splitting type, $\mathcal{T}(z_f, z_m) = T(z_f + z_m)$, then we can rewrite (12) as $\alpha = \mathscr{a}(w_f/w_m)$. Reforms in the income schedule that preserve income-splitting, e.g. Alves et al. (2024), may be analyzed holding household preferences fixed.

The role of Empowerment Assumption A allowed us to hold α fixed as we varied the income-splitting schedule, $T(\cdot)$, but it does not grant our holding α fixed when we introduce s. That is, we allow

$$\hat{\alpha} = a\left(b, p, \frac{w_f(1+s)}{w_m}, \mathbf{d}\right) \neq a\left(b, p, \frac{w_f}{w_m}, \mathbf{d}\right) = \alpha$$

to accommodate the evidence that power is related to the ratio of earnings ability.

To summarize, we recognize that α may be a function of taxes but assume that the impact is zero if the relative price of spouses' leisure is not affected. Because spouses face the same marginal tax rates under an income-splitting tax, the above assumption guarantees that we may hold α fixed when we change the income-splitting schedule, but not necessarily when we introduce the gender-specific subsidy. Inspired by Lise and Yamada (2019), we allow relative (net) productivities to determine relative power.

7 Conclusion

We provide a theoretical and quantitative analysis of a policy that subsidizes married women's earnings with a focus on its empowerment role. Neglecting the direct channel of empowerment can lead to a significant underestimation of both the impact and desirability of such policies. Our work contributes to both the practical and theoretical literature, which has yet to fully explore these issues.

Our findings are particularly relevant since they complement most studies on optimal household taxation — including state-of-the-art works such as Golosov and Krasikov (2025); Bierbrauer et al. (2023) — that abstract from the role of empowerment on household behavior.

References

Mohammad Akbarpour, Piotr Dworczak, and Scott Duke Kominers. Redistributive allocation mechanisms. *Journal of Political Economy*, 132(6):1831 – 1875, 2024. 6, 16

Ingvild Almas, Alex Armand, Orazio Attanasio, and Pedro Carneiro. Measuring and changing control: Women's empowerment and targeted transfers. *Economic Journal*, 128(612):F609–F639, 2018. 7, 22

Ingvild Almås, Orazio Attanasio, and Pedro Carneiro. Chapter 3 - household decisions and intra-household distributions. In Shelly Lundberg and Alessandra Voena, editors, *Handbook of the Economics of the Family, Volume 1*, volume 1 of *Handbook of the Economics of the Family*, pages 111–149. North-Holland, 2023. doi: https://doi.org/10.1016/bs.hefam.2023.01.008. URL https://www.sciencedirect.com/science/article/pii/S2949835X23000083. 36

Cassiano Alves, Carlos E. da Costa, Felipe Lobel, and Humberto Moreira. Intrahousehold inequality and the joint taxation of household earnings. *Journal of Public Economics*, XX(X):xxx – xxx, 2024. 6, 8, 10, 14, 16, 17, 18, 19, 20, 21, 26, 37

Patricia Apps and Ray Rees. Taxation and the household. *Journal of Public Economics*, 35:355–369, 1988. 4, 9

Alex Armand, Orazio Attanasio, Pedro Carneiro, and Valérie Lechene. The effect of gender-targeted conditional cash transfers on household expenditures: Evidence from a randomized experiment. *The Economic Journal*, 130(631):1875–1897, 2020. 7

Felix Bierbrauer, Pierre Boyer, Andreas Peichl, and Daniel Weishaar. Pareto-improvements, welfare trade-offs and the taxation of couples. Rationality and Competition Discussion Paper Series 518, CRC TRR 190 Rationality and Competition, 2024. URL https://EconPapers.repec.org/RePEc:rco:dpaper:518. 8

Felix J. Bierbrauer, Pierre Boyer, Andreas Peichl, Daniel Weishaar, and Felix Bierbrauer. The taxation of couples. CESifo Working Paper Series 10414, CESifo, 2023. URL https://EconPapers.repec.org/RePEc:ces:ceswps:_10414. 8, 16, 37

Felix J. Bierbrauer, Pierre Boyer, Andreas Peichl, and Daniel Weishaar. The political economy of joint taxation. CESifo Working Paper Series 11728, CESifo, 2025. URL https://EconPapers.repec.org/RePEc:ces:ceswps:_11728. 8

Michael J. Boskin and Eytan Sheshinski. Optimal tax treatment of the family: Married couples. *Journal of Public Economics*, 20:281–297, 1983. 5

Martin Browning and Pierre-André Chiappori. Efficient intra-household allocation: a characterization and tests. *Econometrica*, 66(6):1241–1278, 1998. 1, 2, 7, 13, 35

Martin Browning and Mette Gørtz. Spending time and money within the household. *Scandinavian Journal of Economics*, 114(3):681–704, 2012. 7, 22

Rossella Calvi. Why are older women missing in india? the age profile of bargaining power and poverty. *Journal of Political Economy*, 128(7):2453–2501, 2020. doi: 10.1086/706983. URL https://doi.org/10.1086/706983. 6

Pierre-André Chiappori. Nash-bargained household decisions. *International Economic Review*, 32:791–796, 1988a. 4, 9

Pierre-André Chiappori. Rational household labor supply. *Econometrica*, 56:63–89, 1988b. 1, 13

Pierre-André Chiappori. Collective labor supply and welfare. *Journal of Political Economy*, 100:437–467, 1992. 4, 9

Pierre-André Chiappori and Costas Meghir. Chapter 16 - intrahousehold inequality. In Anthony B. Atkinson and François Bourguignon, editors, *Handbook of Income Distribution*, volume 2 of *Handbook of Income Distribution*, pages 1369–1418. Elsevier, 2015. doi: https://doi.org/10.1016/B978-0-444-59429-7.00017-0. 6

Philippe Choné and Guy Laroque. Negative marginal tax rates and heterogeneity. *American Economic Review*, 100(5):2532–47, 2010. 6, 16

Daniele Condorelli. Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior*, 82(C):582–591, 2013. 6, 16

Helmuth Cremer, Jean-Marie Lozachmeur, and Pierre Pestieau. Income taxation of couples and the tax unit choice. *Journal of Population Economics*, 25(2):763–778, 2012. ISSN 09331433, 14321475. 5

Helmuth Cremer, Jean-Marie Lozachmeur, Dario Maldonado, and Kerstin Roeder. Household bargaining and the design of couples' income taxation. *European Economic Review*, 89:454–470, 2016. 5

Carlos E. da Costa and Lucas A. de Lima. A mechanism for the household. Mimeo FGV EPGE, 2024. 10

Esther Duflo. Grandmothers and granddaughters: Old-age pensions and intra-household allocations in south africa. *World Bank Economic Review*, 17:1 – 25, 2003. 7, 22

Andrea Flores. Are cash transfers effective at empowering mothers? a structural evaluation of mexico's oportunidades. Mimeo FGV EPGE, 2024. URL https://afloresm3.github.io/JM-Files-AMFlores/2023_StructuralOportunidades_AMF.pdf. 1, 7, 8, 14, 22, 25, 32

Aart Gerritsen. Optimal taxation when people do not maximize well-being. *Journal of Public Economics*, 144:122–139, 2016. 6, 7

Mikhail Golosov and Ilia Krasikov. The optimal taxation of couples. *Quarterly Journal of Economics*, forthcoming(–):xxx–xxx, 2025. 5, 6, 7, 8, 16, 26, 37

Jonathan Heathcote and Hitoshi Tsujiyama. Optimal income taxation: Mirrlees meets ramsey. *Journal of Political Economy*, 129(11):3141–3184, 2021. 6, 14, 16

Laurence Jacquet and Etienne Lehmann. Optimal taxation with heterogeneous skills and elasticities: Structural and sufficient statistics approaches. In *Working Paper*. 2016. 8

Henrik Jacobsen Kleven, Claus Thustrup Kreiner, and Emmanuel Saez. The optimal income taxation of couples. *Econometrica*, 77(2):537–560, 2009. ISSN 1468-0262. doi: 10.3982/ECTA7343. 5

Jeremy Lise and Ken Yamada. Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *Review of Economic Studies*, 86(5):2184–2219, 2019. 1, 3, 7, 8, 14, 24, 25, 31, 32, 36, 37

Shelly Lundberg and Robert A. Pollak. Separate spheres bargaining and the marriage market. *Journal of Political Economy*, 101(6):pp. 988–1010, 1993. ISSN 00223808. 7

Shelly Lundberg and Robert A Pollak. Bargaining and distribution in marriage. *The Journal of Economic Perspectives*, 10(4):139–158, 1996. 4

Marilyn Manser and Murray Brown. Marriage and household decision making: a bargaining analysis. *International Economic Review*, 21(1):31–44, 1980. 4, 7

Marjorie B McElroy and Mary Jean Horney. Nash-bargained household decisions: Toward a generalization of the theory of demand. *International Economic Review*, 22(2):333–349, 1981. 4, 7

J. A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971. ISSN 00346527. 2, 3, 15, 17, 18, 19, 21

Jean-Charles Rochet and Philippe Choné. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, 1998. 5

Emmanuel Saez. Using elasticities to derive optimal income tax rates. *Review of Economic*

Studies, 68:205–229, 2001. 6, 7

Kevin Spiritus, Etienne Lehmann, Sander Renes, and Floris T. Zoutman. Optimal taxation with multiple incomes and types. *Theoretical Economics*, 2024. URL https://EconPapers.repec.org/RePEc:the:publsh:5479. 5, 8, 13

Duncan Thomas. Intra-household resource allocation: An inferential approach. *The Journal of Human Resources*, 25(4):635 – 664, 1990. 6, 7, 22

Appendix

A The Planner's Program

Consider the following program

$$\min \hat{\alpha} h\left(\frac{z_f}{w_f}\right) + [1-\hat{\alpha}] h\left(\frac{z_m}{w_m}\right)$$

s.t.,

$$z_f(1+s) + z_m = \hat{z}$$

which is equivalent to

$$\min_{\hat{z}_f} \left\{ \hat{\alpha} h\left(\frac{\hat{z}_f}{w_f(1+s)}\right) + [1-\hat{\alpha}] h\left(\frac{\hat{z}-\hat{z}_f}{w_m}\right) \right\}.$$

So, in all that follows we work with the transformed program.

Let then[19]

$$\mathcal{H}(\hat{z}, \boldsymbol{\iota}) = \min_{\hat{z}_f} \left\{ \hat{\alpha} h\left(\frac{\hat{z}_f}{\hat{w}_f}\right) + [1-\hat{\alpha}] h\left(\frac{\hat{z}-\hat{z}_f}{w_m}\right) \right\},$$

where $\hat{w}_f = w_f(1+s)$, and

$$\mathcal{U}(c, \hat{\alpha}) = \max_{c_f} \left\{ \hat{\alpha} u(c_f) + [1-\hat{\alpha}] u(c-c_f) \right\}.$$

[19]The normalization $(1+\hat{\beta})^{-1}$ will be convenient when we define the planner's objective.

Then, the household earnings choice problem may be written

$$\max_{\hat{z}} \left\{ \mathcal{U}\left(\hat{z} - T\left(\hat{z}\right), \hat{\alpha}\right) - \mathcal{H}(\hat{z}, \hat{\boldsymbol{\iota}}) \right\}$$

For the iso-elastic case,

$$\mathcal{H}(\hat{z}, \hat{\boldsymbol{\iota}}) = \frac{1}{1+\gamma} \min_{\hat{z}_f} \left\{ \hat{\alpha} \left(\frac{\hat{z}_f}{\hat{w}_f} \right)^{1+\gamma} + [1-\hat{\alpha}] \left(\frac{\hat{z} - \hat{z}_f}{w_m} \right)^{1+\gamma} \right\},$$

thus,

$$\hat{z}_f = \frac{[1-\hat{\alpha}]^{\frac{1}{\gamma}} \hat{w}_f^{\frac{1+\gamma}{\gamma}}}{[1-\hat{\alpha}]^{\frac{1}{\gamma}} \hat{w}_f^{\frac{1+\gamma}{\gamma}} + \hat{\alpha}^{\frac{1}{\gamma}} w_m^{\frac{1+\gamma}{\gamma}}} \hat{z}.$$

$$\mathcal{H}(\hat{z}, \hat{\boldsymbol{\iota}}) = \frac{1}{1+\gamma} \left(\frac{\hat{z}}{\omega(\hat{\boldsymbol{\iota}})} \right)^{1+\gamma} = \mathcal{A}_H(\hat{\boldsymbol{\iota}}) \hat{z}^{1+\gamma}.$$

where

$$\omega(\hat{\boldsymbol{\iota}}) = \left[\hat{\alpha}^{\frac{-1}{\gamma}} \hat{w}_f^{\frac{1+\gamma}{\gamma}} + [1-\hat{\alpha}]^{\frac{-1}{\gamma}} w_m^{\frac{1+\gamma}{\gamma}} \right]^{\frac{\gamma}{1+\gamma}},$$

and

$$\mathcal{U}(c, \hat{\boldsymbol{\iota}}) = \max \left\{ \hat{\alpha} \ln c_f + [1-\hat{\alpha}] \ln(c - c_f) \right\},$$

which yields

$$c_f = \hat{\alpha} c,$$

and

$$\mathscr{U}(c, \hat{\boldsymbol{\iota}}) = \hat{\alpha} \ln \hat{\alpha} + [1-\hat{\alpha}] \ln[1-\hat{\alpha}] + \ln c = \mathcal{A}_U(\hat{\boldsymbol{\iota}}) + \ln c.$$

With some abuse, we can therefore write the worker's program as

$$\max_{\hat{z}} \left\{ \mathcal{U}\left(\hat{z} - T\left(\hat{z}\right), \hat{\alpha}\right) - \mathcal{H}(\hat{z}, \hat{\boldsymbol{\iota}}) \right\}.$$

As for the planner,

$$\mathcal{V}(c, \hat{\alpha}) = \frac{1}{2} \ln \hat{\alpha} + \frac{1}{2} \ln \hat{\alpha} + \ln c = \mathcal{A}_V(\hat{\boldsymbol{\iota}}) + \ln c,$$

and

$$\mathscr{K}(\hat{z},\hat{\boldsymbol{\iota}}) = \frac{1}{1+\gamma}\frac{[1-\hat{\alpha}]^{\frac{1+\gamma}{\gamma}}\hat{w}_f^{\frac{1+\gamma}{\gamma}} + \hat{\alpha}^{\frac{1+\gamma}{\gamma}} w_m^{\frac{1+\gamma}{\gamma}}}{2\left[[1-\hat{\alpha}]^{\frac{1}{\gamma}}\hat{w}_f^{\frac{1+\gamma}{\gamma}} + \hat{\alpha}^{\frac{1}{\gamma}} w_m^{\frac{1+\gamma}{\gamma}}\right]^{1+\gamma}}\hat{z}^{1+\gamma} = \mathscr{A}_K(\hat{\boldsymbol{\iota}})\hat{z}^{1+\gamma}.$$

Writing the Planner's Objective We show now that the planner's program can be viewed as Utilitarian (with respect to households) with a correction for dissonance. To formulate the planner's problem, recall that when agents' utilities are $\ln$-isoelastic, then $\partial\mathscr{U}(c,\boldsymbol{\iota})/\partial c = c^{-1} = \partial\mathscr{V}(c,\boldsymbol{\iota})/\partial c$. Moreover, we can write $\mathscr{K}(z,\boldsymbol{\iota}) = \mathscr{A}_K(\boldsymbol{\iota})z^{1+\gamma}$ and $\mathscr{H}(z,\boldsymbol{\iota}) = \mathscr{A}_H(\boldsymbol{\iota})z^{1+\gamma}$.[20] In this case, the dissonance term is independent of (c,z),

$$\xi(c,z,\boldsymbol{\iota}) = \frac{\partial\mathscr{K}(z,\boldsymbol{\iota})/\partial z}{\partial\mathscr{H}(z,\boldsymbol{\iota})/\partial z} = \frac{\mathscr{A}_K(\boldsymbol{\iota})}{\mathscr{A}_H(\boldsymbol{\iota})} \equiv \hat{\xi}(\boldsymbol{\iota}),$$

and, conveniently,

$$\begin{aligned}\mathscr{H}(z,\boldsymbol{\iota}) - \mathscr{K}(z,\boldsymbol{\iota}) &= [\mathscr{A}_H(\boldsymbol{\iota}) - \mathscr{A}_K(\boldsymbol{\iota})]\, z^{1+\gamma} = \left[1 - \frac{\mathscr{A}_K(\boldsymbol{\iota})}{\mathscr{A}_H(\boldsymbol{\iota})}\right]\mathscr{A}_H(\boldsymbol{\iota})z^{1+\gamma} \\ &= \left[1-\hat{\xi}(\boldsymbol{\iota})\right]\mathscr{H}(z,\boldsymbol{\iota}).\end{aligned}$$

Since

$$\mathscr{H}(z,\boldsymbol{\iota}) = \frac{1}{1+\gamma}\left(\frac{z}{\omega}\right)^{1+\gamma}, \quad \forall\boldsymbol{\iota}\in\mathscr{I}(\omega),$$

then

$$\mathbb{E}\left[\mathscr{H}(z,\boldsymbol{\iota})[1-\xi(c,z,\boldsymbol{\iota})]\big|\boldsymbol{\iota}\in\mathscr{I}(\omega)\right] = \frac{1}{1+\gamma}\left(\frac{z}{\omega}\right)^{1+\gamma}\mathbb{E}\left[1-\xi(c,z,\boldsymbol{\iota})\big|\boldsymbol{\iota}\in\mathscr{I}(\omega)\right].$$

Similarly, for every $\hat{\boldsymbol{\iota}}$,

$$\mathscr{V}(c,\hat{\boldsymbol{\iota}}) = \mathscr{U}(c,\hat{\boldsymbol{\iota}}) + \mathscr{A}_V(\hat{\boldsymbol{\iota}}) - \mathscr{A}_U(\hat{\boldsymbol{\iota}}).$$

[20] $\mathscr{A}_H(\boldsymbol{\iota}) = \omega(\boldsymbol{\iota})^{-1-\gamma}/(1+\gamma)$. As for $\mathscr{A}_K(\boldsymbol{\iota})$, its expression can be derived by plugging k_a and k_b in (**??**) in the general formula,

$$\mathscr{K}(z,\boldsymbol{\iota}) := \frac{z^{1+\gamma}}{2(1+\gamma)}\left[\left(\frac{k_a}{w_a}\right)^{1+\gamma} + \left(\frac{k_b}{w_b}\right)^{1+\gamma}\right].$$

For every, ω we may finally define

$$\mathcal{Q}(\omega) = \mathbb{E}\left[\mathscr{A}_V(\hat{\boldsymbol{\iota}}) - \mathscr{A}_U(\hat{\boldsymbol{\iota}}) \middle| \hat{\boldsymbol{\iota}} \in \mathcal{I}(\omega)\right].$$

Define the new distribution

$$\hat{\Psi}_s(\omega) \equiv \Phi\left(\boldsymbol{\iota} | \boldsymbol{w}\left(i_s(\boldsymbol{\iota})\right) \leq \omega\right),$$

and the associated density, $\hat{\psi}_s(\cdot)$.

The planner's program is, then, to maximize

$$\max_{(c(\cdot),z(\cdot))} \int_{\underline{\omega}}^{\bar{\omega}} \left\{ \mathcal{U}(c(\omega), z(\omega), \omega) + \mathcal{Q}(\omega) + \left[1 - \bar{\xi}(\omega)\right] \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \hat{\psi}_s(\omega) d\omega,$$

The Resource Constraint Hence, were it not for the budgetary impact of this subsidy to solve the optimal income taxation, the problem would be the same as before with the change in distribution as the only adaptation.

To consider the budgetary impact, recall the definition,

$$\bar{k}_s(\omega) = \mathbb{E}\left[k(\boldsymbol{\iota}) \middle| \boldsymbol{\iota} \in \hat{\mathcal{I}}_s(\omega)\right],$$

where

$$\hat{\mathcal{I}}_s(\omega) = \left\{\boldsymbol{\iota} \in \Lambda | \, \boldsymbol{w}\left(i_s(\boldsymbol{\iota})\right) = \omega\right\}.$$

Now the resource constraint of the economy is still

$$\int_{\underline{\hat{\omega}}}^{\bar{\hat{\omega}}} \left[z(\omega) - c\left(\omega\right)\right] \psi_s(\omega) d\omega \geq G.$$

Because we are using the transformed variable $\hat{z}(\omega) = z(\omega)[1 + \bar{k}(\omega)s]$, the resource constraint will be written,

$$\int_{\underline{\hat{\omega}}}^{\bar{\hat{\omega}}} \left[\frac{\hat{z}(\omega)}{1 + \bar{k}(\omega)s} - c\left(\omega\right)\right] \psi_s(\omega) d\omega \geq G.$$

A.1 Solving the Planner's Program

Recall that the planner's program is

$$\max_{(c(\cdot),z(\cdot))} \int_{\underline{\omega}}^{\bar{\omega}} \left\{ \mathcal{U}(c(\omega), z(\omega), \omega) + \mathcal{K}(\omega) + \left[1 - \bar{\xi}(\omega)\right] \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \psi(\omega) d\omega,$$

subject to

$$\omega \in \operatorname*{argmax}_{r} \mathcal{U}(c(r), z(r), \omega),$$

and

$$\int_{\hat{\underline{\omega}}}^{\bar{\hat{\omega}}} \left[\frac{z(\omega)}{1 + \bar{k}_s(\omega) s} - \mathcal{c}\left(v(\omega), z(\omega), \omega\right) \right] \psi(\omega) d\omega \geq G.$$

Letting,

$$v(\omega) := \mathcal{U}(c(\omega), z(\omega), \omega)$$

we have the envelope condition

$$\dot{v}(\omega) = \frac{z(\omega)^{1+\gamma}}{\omega^{2+\gamma}},$$

that, with the monotonicity constraint

$$z(\omega) \quad \text{increasing in } \omega,$$

is equivalent to the I.C. constraint.

Finally defining, $\mathcal{c}\left(v, z, \omega\right)$ through

$$v := \mathcal{U}(\mathcal{c}\left(v, z, \omega\right), z, \omega), \quad \forall \left(v, z, \omega\right),$$

allows us to write the Lagrangian,

$$\begin{aligned}
\mathcal{L} = \int_{\underline{\omega}}^{\bar{\omega}} & \left\{ v(\omega) + \mathcal{K}(\omega) + \left[1 - \bar{\xi}(\omega)\right] \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma} \frac{1}{1+\gamma} \right\} \psi(\omega) d\omega \\
& - \int_{\underline{\omega}}^{\bar{\omega}} \left[\dot{\mu}(\omega) v(\omega) + \mu(\omega) \frac{z^{1+\gamma}}{\omega^{2+\gamma}} \right] d\omega \\
& + \lambda \int_{\hat{\underline{\omega}}}^{\bar{\hat{\omega}}} \left[\frac{z(\omega)}{1 + \bar{k}_s(\omega) s} - \mathcal{c}\left(v(\omega), z(\omega), \omega\right) \right] \psi(\omega) d\omega - \lambda G
\end{aligned}$$

The first-order conditions are

$$\psi(\omega) - \dot{\mu}(\omega) - \lambda \boldsymbol{x}_v\left(v(\omega), z(\omega), \omega\right)\psi(\omega) = 0,$$

and

$$\left[1-\bar{\xi}(\omega)\right]\frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}\psi(\omega) - \mu(\omega)\left[1+\gamma\right]\frac{z(\omega)^{1+\gamma}}{\omega^{2+\gamma}} \\ + \lambda\left[\frac{1}{1+\bar{k}_s(\omega)s} - \boldsymbol{c}_z\left(v(\omega), z(\omega), \omega\right)\right]\psi(\omega) = 0.$$

Next, we note that for our specification of preferences,

$$\ln \boldsymbol{c}\left(v(\omega), z(\omega), \omega\right) - \left(\frac{z(\omega)}{\omega}\right)^{1+\gamma}\frac{1}{1+\gamma} = v(w)$$

which implies

$$\frac{\boldsymbol{c}_v\left(v(\omega), z(\omega), \omega\right)}{\boldsymbol{c}\left(v(\omega), z(\omega), \omega\right)} = 1,$$

and

$$\boldsymbol{c}_z\left(v(\omega), z(\omega), \omega\right) = \frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}c\left(\omega\right).$$

Hence,

$$\psi(\omega) - \dot{\mu}(\omega) - \lambda x\left(\omega\right)\psi(\omega) = 0$$

and

$$\left[1-\bar{\xi}(\omega)\right]\frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}\zeta(\omega)\psi(\omega) - \mu(\omega)\left[1+\gamma\right]\frac{z(\omega)^{\gamma}}{\omega^{2+\gamma}}\zeta(\omega) \\ + \lambda\left[1 - \frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}c\left(\omega\right)\zeta(\omega)\right]\psi(\omega) = 0,$$

where $\zeta(\omega) = \left[1+\bar{k}_s(\omega)s\right]^{-1}$.

The first order condition with respect to $z(\omega)$ can, therefore, be written

$$\left[\frac{1-\bar{\xi}(\omega)}{c(\omega)}\right]c(\omega)\frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}\zeta(\omega)\psi(\omega) - \mu(\omega)\left[\frac{1+\gamma}{\omega c(\omega)}\right]\frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}\zeta(\omega)c(\omega) \\ + \lambda\left[1 - \frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}c\left(\omega\right)\zeta(\omega)\right]\psi(\omega) = 0$$

which we can equivalently write as

$$\left[\frac{1-\bar{\xi}(\omega)}{c(\omega)}\right]\left[1-\hat{\tau}(\omega)\right]\psi(\omega)-\mu(\omega)\left[\frac{1+\gamma}{\omega c(\omega)}\right]\left[1-\hat{\tau}(\omega)\right]c(\omega)+\lambda\hat{\tau}(\omega)\psi(\omega)=0$$

for

$$1-\hat{\tau}(\omega)=c(\omega)\frac{z(\omega)^{\gamma}}{\omega^{1+\gamma}}\zeta(\omega).$$

Rearranging,

$$\left[\frac{1-\bar{\xi}(\omega)}{c(\omega)}\right]\psi(\omega)-\mu(\omega)\left[\frac{1+\gamma}{\omega c(\omega)}\right]=-\lambda\psi(\omega)\frac{\hat{\tau}(\omega)}{1-\hat{\tau}(\omega)}$$

Finally, note that

$$\int_{\underline{\omega}}^{\omega}\psi(w)dw-\lambda\int_{\underline{\omega}}^{\omega}c\left(w\right)\psi(w)dw=\mu(\omega),$$

and

$$1=\lambda\int_{\underline{\omega}}^{\bar{\omega}}c\left(\omega\right)\psi(\omega)d\omega\Longrightarrow\lambda=\mathbb{E}\left[x\right]^{-1}$$

which gives an expression for $\mu(\omega)$,

$$\Psi(\omega)-\mathbb{E}\left[c\right]^{-1}\int_{\underline{\omega}}^{\omega}c\left(w\right)\psi(w)dw=\mu(\omega).$$

This finally allows us to write

$$\left[\frac{1-\bar{\xi}(\omega)}{c(\omega)}\right]\psi(\omega)-\Psi(\omega)\left\{1-\frac{\mathbb{E}\left[c|c<c(\omega)\right]}{\mathbb{E}\left[c\right]}\right\}\left[\frac{1+\gamma}{\omega c(\omega)}\right]=-\lambda\psi(\omega)\frac{\hat{\tau}(\omega)}{1-\hat{\tau}(\omega)},$$

or, more intuitively,

$$\frac{\hat{\tau}(\omega)}{1-\hat{\tau}(\omega)}\;=\;\frac{\Psi(\omega)[1+\gamma]}{\omega\psi(\omega)}\left\{\mathbb{E}\left[\frac{c(t)}{c(\omega)}\right]-\mathbb{E}\left[\frac{c(t)}{c(\omega)}\middle|t<\omega\right]\right\}-\left[1-\bar{\xi}(\omega)\right]\mathbb{E}\left[\frac{c(t)}{c(\omega)}\right]$$

B Taxing Secondary Earners

Consider subsidizing the secondary earner. The main concern with such a policy is that who is primary and who is secondary is a household choice, therefore, something endogenous do policy. We will consider a policy that does not invert the identities of secondary and primary earners.

To do that, define s_f, the subsidy that goes to the woman. Then we have the subsidy for the man defined as $s - s_f$. With this policy we have,

$$\frac{\hat{z}_f}{\hat{z}_m} = \left(\frac{1-\hat{\alpha}}{b_f\hat{\alpha}}\right)^{\frac{1}{\gamma}} \left(\frac{w_f(1+s_f)}{w_m(1+s-s_f)}\right)^{\frac{1+\gamma}{\gamma}},$$

where

$$\frac{\hat{\alpha}}{1-\hat{\alpha}} = \left(\frac{1+s_f}{1+s-s_f}\right)^{\beta_1} \frac{\alpha}{1-\alpha}.$$

Hence,

$$\frac{\hat{z}_f}{\hat{z}_m} = \underbrace{\left(\frac{1-\alpha}{b_f\alpha}\right)^{\frac{1}{\gamma}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{\gamma}}}_{z_f/z_m} \left(\frac{1+s_f}{1+s-s_f}\right)^{\frac{1+\gamma-\beta_1}{\gamma}}$$

Define the following conditions,

$$\left(\frac{1-\alpha}{b_f\alpha}\right)^{\frac{1}{\gamma}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{\gamma}} < 1 \tag{13}$$

$$\left(\frac{1-\alpha}{b_f\alpha}\right)^{\frac{1}{\gamma}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{\gamma}} \geq 1 \tag{14}$$

$$\left(\frac{1-\alpha}{b_f\alpha}\right)^{\frac{1}{\gamma}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{\gamma}} (1+s)^{\frac{1+\gamma-\beta_1}{\gamma}} < 1 \tag{15}$$

$$\left(\frac{1-\alpha}{b_f\alpha}\right)^{\frac{1}{\gamma}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{\gamma}} (1+s)^{\frac{\beta_1-\gamma-1}{\gamma}} \geq 1 \tag{16}$$

If (13) and (15) are true, then $s_f = s$.

If (14) and (16) are true, then $s_f = 0$.

Note that either (13) or (14) are true. So, both possibilities to be false, it is either because, (13) is true but (15) is not true or because (14) is true but (16) is not. In either case, choose

$$\left(\frac{z_m}{z_f}\right)^{\frac{\gamma}{1+\gamma-\beta}} = \frac{1+s_f}{1+s-s_f},$$

or

$$s_f = \frac{1 + s - \left(\frac{1-\alpha}{b_f \alpha}\right)^{\frac{1}{1+\gamma-\beta}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{1+\gamma-\beta}}}{1 + \left(\frac{1-\alpha}{b_f \alpha}\right)^{\frac{1}{1+\gamma-\beta}} \left(\frac{w_f}{w_m}\right)^{\frac{1+\gamma}{1+\gamma-\beta}}}.$$

This guarantees that secondary earners will at most attain the same earnings of primary earners.

C Additional Figures

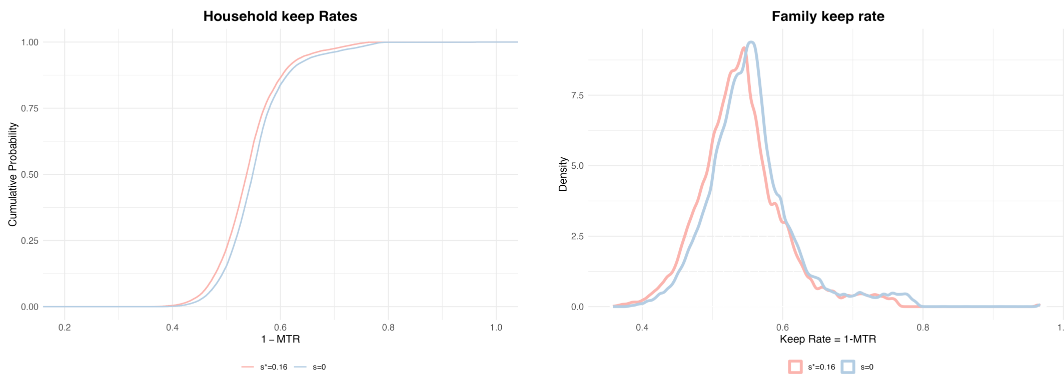

Figure 11: ***Distribution of marginal keep rates*** $1 - T'$***.*** *The figure displays the cumulative (left panel) and density (right panel) functions for* $1 - T'$*.*

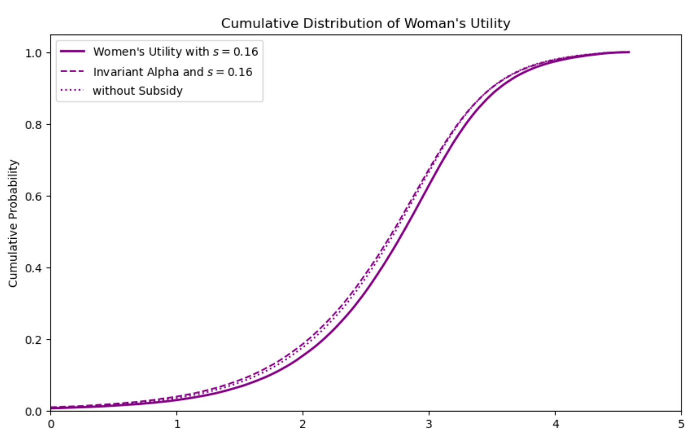

Figure 12: This figure compares the cumulative utility distribution for women in the baseline and under the optimal policy. It also displays the distribution for $s = 0.16$ *under the assumption that* s *does not influence* α*.*

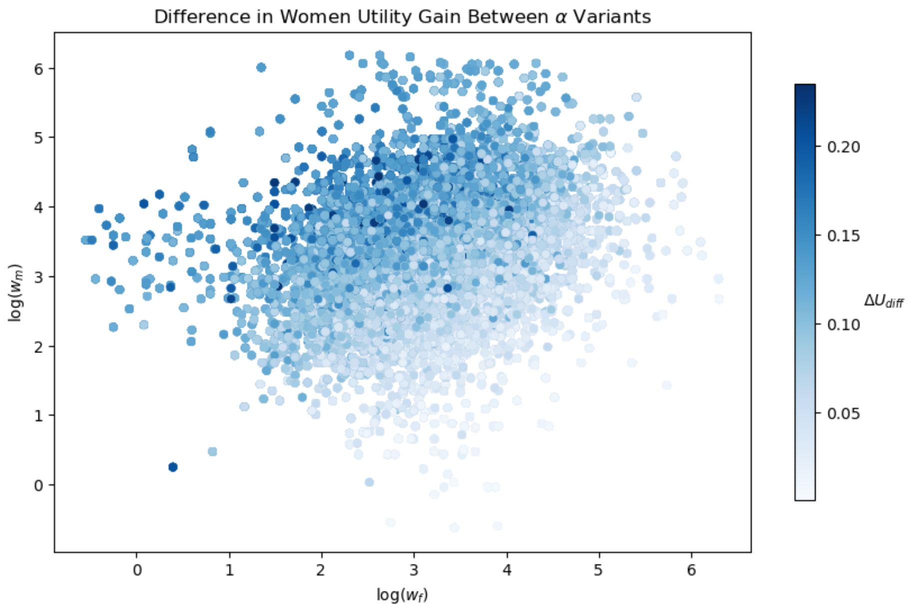

Figure 13: This figure illustrates the difference in women's utility gains between the baseline scenario, where intra-household bargaining power (α) adjusts, and the counterfactual scenario, where α remains fixed. Positive values indicate that women benefit more when the policy enhances bargaining power, particularly in households with higher-wage husbands.

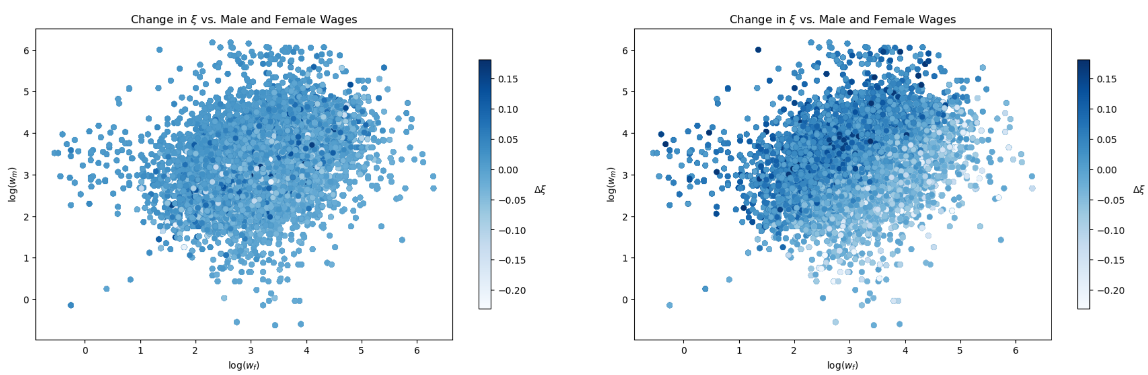

*Figure 14: **Variation in the Dissonance Term (ξ) Relative to Male and Female Wages.** The left panel shows the change in the dissonance term (ξ) when bargaining power (α) is allowed to adjust. In this scenario, the variation in ξ is more evenly distributed across households, with no clear pattern tied to wage disparities. The right panel presents the counterfactual scenario where α remains constant. Here, the change in ξ is more pronounced in households where men earn significantly more than women, suggesting that dissonance is more sensitive to wage disparities when bargaining power is fixed.*

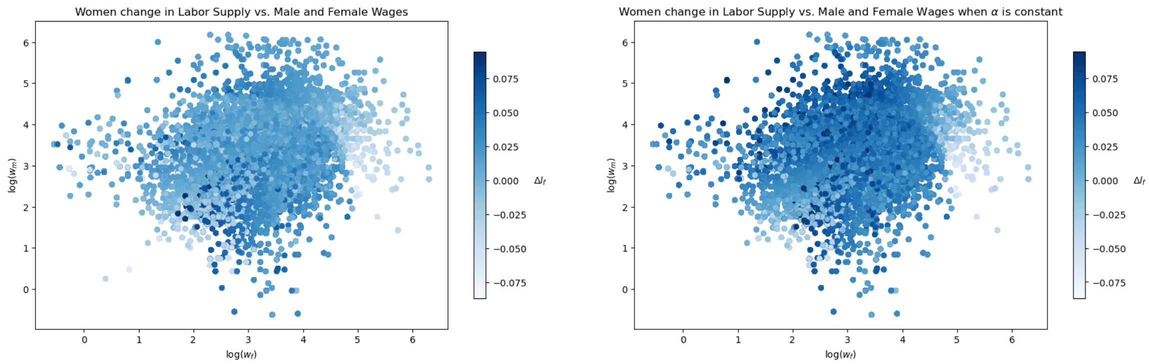

Figure 15: The left panel shows the change in women's labor supply (Δl_f) in our baseline scenario. In this case, women married to lower-productivity husbands generally reduce their labor supply, as they can capture a larger share of household resources through improved bargaining power, reducing the need to work more. The right panel presents the counterfactual scenario where α remains constant. Here, women increase their labor supply across most households, as the subsidy alone is insufficient to improve their bargaining position, forcing them to work more to capture additional resources. The comparison highlights how the empowerment channel allows women to reallocate time away from work when their bargaining power improves.

Figure 16: The left panel shows the change in women's consumption (Δc_f) in our baseline scenario. Women with higher wages experience larger consumption gains, as the subsidy strengthens even more their bargaining position, allowing them to secure a greater share of household resources. The right panel presents the counterfactual scenario where α remains constant, isolating the effect of the empowerment channel. In this case, consumption gains are more evenly distributed, and some women even experience losses.

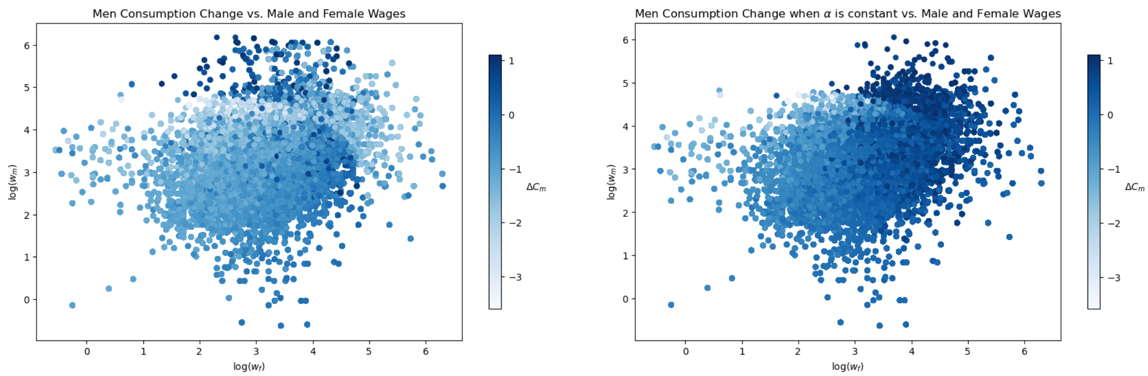

Figure 17: In our baseline case, variation in male consumption is more pronounced among households where men have higher wages. The right panel presents the counterfactual scenario where α remains fixed. Here, male consumption increases more in households where women have higher wages, suggesting that, when bargaining power does not adjust, men capture a larger share of the subsidies received by women, regardless of their productivity.

Figure 18: The left panel shows the variation in male labor supply in our baseline. In this case, most men maintain their labor supply, with reductions concentrated among a small subset – primarily those married to highly productive women. The right panel presents the counterfactual scenario where bargain remains constant, and most men reduce their labor supply.

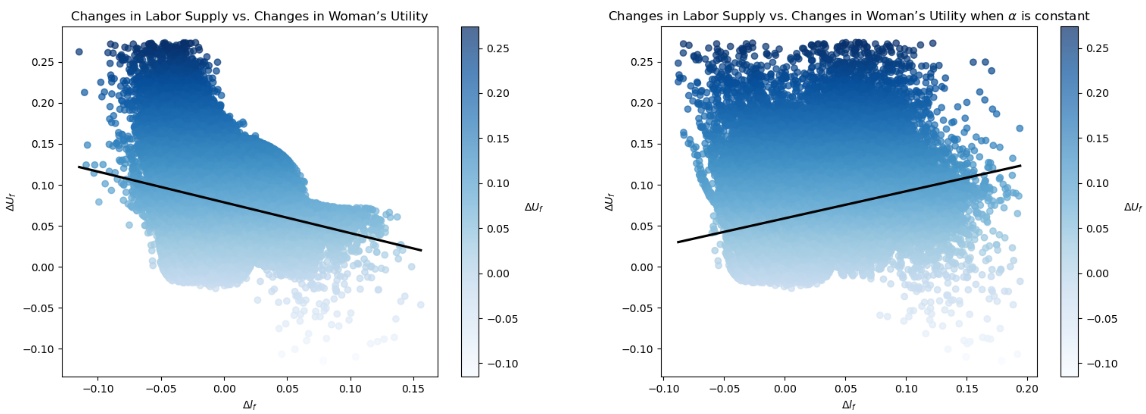

Figure 19: The figure displays changes in female welfare against changes in her labor supply. The right panel considers the counterfactual world in which α is held fixed as we change s.