

Collegio Carlo Alberto

Efficient Risk Sharing
in the Presence of a Public Good

Christine Hauser

Working Paper No. 87

December 2008

www.carloalberto.org

Efficient Risk Sharing in the Presence of a Public Good

Christine Hauser¹

September 2008²

¹Collegio Carlo Alberto and CHILD

²© 2008 by Christine Hauser. Any opinions expressed here are those of the authors and not those of the Collegio Carlo Alberto.

Abstract

This paper studies the provision of a public good between two agents under lack of commitment and applies it to the problem of children's consumption in separated couples, where children are considered to be public goods. The custodial mother controls the child's consumption, whereas the father can contribute indirectly by making monetary transfers to the mother, but has no control over how the mother spends them. Using minmax punishments, I look for the Pareto frontier of the Subgame Perfect Equilibrium payoffs, and characterize the equilibrium and long term implications of the model. As in the previous literature, agents' consumptions and continuation values covary positively with their income levels. In the case where the constraint for the public good provision binds, both agents' private consumptions increase relative to the public good provision. In the long run, if some first best allocation is sustainable, the long-term equilibrium will converge to a first best allocation. Otherwise, agents' utilities oscillate over a finite set of values. I then study the theoretical implications of one-sided enforcement when the public good provider has the authority to enforce transfers from the second agent. This is motivated by the wave of US policy reforms to enforce child support payments from fathers. The model predicts an increase in the ratio of the mother's consumption to the child's.

JEL Classification: C72, C73, D90, E21.

Keywords: insurance, lack of commitment, optimal dynamic contract, public good.

1 Introduction

I develop a model of efficient risk sharing between two agents in the presence of a public good, where one agent controls the supply of the public good. First, I characterize the Pareto frontier of Subgame Perfect Equilibrium (SPE) payoffs of the game under double sided lack of commitment. I then study the theoretical implications of *one-sided enforcement* on the set of implementable SPE and more specifically on the long run equilibrium of the game, by asking what happens when the public good provider has the authority to enforce transfers from the second agent.

The question of risk sharing in the absence of commitment has been widely studied in the dynamic contracting literature. In a seminal paper, Kocherlakota (1996) studies the problem of two risk averse agents with random endowments and a single consumption good. Agents wish to insure each other against their endowment shocks, but are unable to commit to future transfers: an agent with a high endowment one period might prefer to consume that endowment entirely instead of sharing it, and hence needs to be compensated for staying in the contract through a higher stream of consumption than what his current entitlement allows him. This paper builds on this benchmark model by introducing a public good which only one agent can supply. The question of allocation of total income between private and public consumptions is then added to the decision making process. At the beginning of each period, the state of the world realizes. The *transferring* agent makes a payment to the *providing* agent but has no control over how the transfer is spent. Given the transfer, the *providing* agent decides on how to split her income between her private consumption and the public good consumption.

Examples of such a situation are numerous. One application is to a benevolent lender and a borrowing country, where the borrowing government can spend the debt in two ways. The "corrupt way", which is to spend it on consumption and on buying individuals' votes, and the "ethical way", which is to invest it in capital and infrastructure. Assuming that the lending institution has an interest in seeing the debt spent in the "ethical way", the investment in capital and infrastructure is comparable to the public good spending above. One could also adapt the present scenario to think of an economy with firms and workers who provide the human capital necessary for production. The public good in this case is the human capital: the firm needs it for production and the worker is the only agent who can accumulate it. The firm makes state-contingent transfers that are commensurate with the worker's capital, and the worker improves her capital through effort. Moreover, the worker keeps her human capital when she repudiates, so her outside option depends on her human capital accumulation.¹

The application on which this paper focuses is taken from family economics, where the literature has traditionally considered children as public goods in divorced and separated couples. Typically, only the custodial mother can spend directly on the child while the noncustodial father can make child support payments to the mother, who then allocates

¹This is the type of environment studied by Marimon and Quadrini (2005). In their model, firms face competition from other firms, and the degree of competition affects the innovator's outside option and her incentive to invest in her human capital.

her post-transfer income between her private consumption and the child's consumption.

The low economic status of children from divorce has been a major concern for legislators and social scientists, who often cited the low child support payments from fathers as the culprit. Weiss and Willis (1985) related this fact to the mother's moral hazard problem: once the father makes a transfer to the mother, he has no control over how she spends it. During the 80's and 90's, a series of policy reforms was enacted by the state and federal governments, including setting up guidelines and enforcement agencies to collect child support from fathers in order to increase the welfare of children in single parent households. The problem with this approach is that it overlooks important dynamics in the separated couples interaction, since fixing the father's problem through child support enforcement means tampering with the mother's incentives for committing to a high level of expenditure on the child. As far as I know, this is the first paper to point out the danger of one-sided enforcement in the child support case and to formalize this statement through a dynamic contracting approach. In fact, the previous literature on child support has commonly depicted the interaction between parents as a one shot game with a Stackelberg equilibrium outcome, despite existing empirical work which suggests that parents behave strategically and adjust their transfer and expenditure choices over time according to the past actions of their partner.² Going from a noncooperative equilibrium with no child support enforcement to one with enforcement will necessarily increase children's consumption (as long as child support receipts increase as well). As will be shown, this is not necessarily true in a dynamic model.

This paper relates to two other strands of literature. The first is on the voluntary provision of public goods, where most of the recent literature has focused on the free rider problem in dynamic games of sequential contributions (Varian (1994), Marx and Matthews (2000) among others). More relevantly though, this paper builds on the contracting literature by using the dynamic programming methods of Thomas and Worrall (1990) in order to characterize the Pareto frontier of Subgame Perfect Equilibrium (SPE) payoffs. It also draws on Kocherlakota (1996) to model the efficient allocation of consumption in economies with double-sided lack of commitment. There are numerous papers that follow these seminal works. One branch focuses on the asset pricing implications of these type of contracts. Examples are Alvarez and Jermann (2000), Ábrahám and Cárceles-Poveda (2005). Another branch is on sovereign debt, such as Kehoe and Perri (2004), Kletzer and Wright (2000). However, the literature has not considered the presence of public goods in models of insurance in the absence of commitment.

The first novelty of this paper is to define and provide an incentive compatible mechanism to implement the worst possible punishments in a setting with a public good. Cooperation is usually supported by the threat of reverting to the autarkic allocation where each agent consumes his income forever. In the present context, the minmax strategies entail no transfers, and in the case where the *transferring* agent is being minmaxed, zero public

²Del Boca and Flinn (1994), Argys and Peters (1998), Knox (1996) are a few examples of studies showing that a dollar of income from child support has a larger effect on children's consumption and educational achievement than a dollar of income from any other source, a finding which cannot be accounted for in a static setting.

good provision. If agents care about the public good, these *strategies* will not be incentive compatible as the amounts that agents prefer to transfer and spend on the public good will be positive. However, given a sufficient condition on parameters, the minmax *values* will be implementable using stick-and-carrot type punishments.

As in the standard model of mutual insurance, when a *transferring* agent's incentive constraint binds, he must be compensated with higher current consumption and continuation value, while the other agent suffers lower consumption and continuation value. The new case, which underlines the public good provision problem, occurs when the *providing* agent prefers to take the transfer and spend it according to her own rule, rather than as dictated by the first best. As a result, the provision of the public good becomes very costly in terms of meeting her incentives, which leads to lower payments on the part of the *transferring* agent. So both agents' private consumptions increase at the expense of the public consumption. Moreover, one can rank ratios of individual consumption to public consumption according to continuation values and states, with that ratio being highest when an agent's incentive constraint binds. This is because agents like to be compensated by having their ratio of marginal utilities closer to their individual optimal than what the first best entails.

The long run properties of the optimal contract are similar to the standard case. If there exists a subgame perfect first best allocation, then any optimal allocation will converge to a first best allocation. Otherwise, agents' values will oscillate over a finite set that is unique and independent of their starting values. This may happen only when there is some aggregate uncertainty.

The last question this paper addresses is what happens when the public good provider has the authority to enforce transfers from the second agent. What are the consequences on the set of implementable SPE, and on the long run private and public consumptions of the two agents? *One-sided enforcement* is often thought of as a solution to the free riding problem in public goods and the purpose from this exercise is to study theoretically the consequences of one the main reforms to the child support system in the US (namely child support enforcement) on custodial mothers' consumption decisions. The enforcement equilibrium serves as the new threat point which supports the *new Pareto frontier of payoffs*. Relative to the old, no-enforcement threat point, the enforcement option delivers a lower utility for the father and a higher utility to the mother, making it harder to satisfy her incentive for providing the child with a large consumption. As a result, in the new long run equilibrium, the mother will spend on average a higher fraction of her income on herself and a lower fraction on the child.³

The paper is organized as follows: The next section lays out the environment, followed by a brief description of the first best allocation as a benchmark to the constrained problem. Section 4 presents the recursive formulation, while sections 5, 6 and 7 characterize the optimal contract and the long run equilibrium. The following section analyzes the effects of one-sided enforcement. Section 9 concludes. Appendix A describes how to implement the minmax values using stick-and-carrot type punishments. Appendix B contains proofs

³Hauser (2008) empirically tests this claim using consumer expenditure data and finds strong supportive evidence for it.

of propositions, and appendix C the proofs and illustrative cases of the equilibrium under one-sided enforcement.

2 The Environment

I study an infinite horizon repeated game between two agents: the father F and the mother M , who each enjoy the consumption of a private good and a public good: the child's consumption. Time is discrete and agents discount future utility at rate β . The state of the world in period τ is stochastic and is determined by the realization of a discrete random variable θ , independently and identically distributed over time, with support equal to $\{1, \dots, S\}$. The probability that θ takes on the value s is denoted by π_s , where $\pi_s > 0$ for all s . Incomes at time τ are denoted by Y_τ^F and Y_τ^M for F and M respectively, and are determined by the realization of θ in every period, with aggregate income denoted by $Y_\tau^F + Y_\tau^M = \mathbf{Y}_\tau$. The mother is the only one capable of spending directly on the child, whereas the father is merely a contributor who can make monetary transfers, or child support payments, to the mother. We assume for simplicity that incomes and preferences are such that transfers are unilateral and positive in all states in the Stackelberg equilibrium of the static game.⁴ All information is public, so parents observe each other's income and expenditure.⁵

Let f_τ and m_τ be the consumptions of F and M at time τ , and let c_τ be the child's. Then F 's lifetime utility at τ can be written as

$$E_\tau \sum_{r=0}^{\infty} \beta^r u(f_{\tau+r}, c_{\tau+r})$$

and M 's utility as

$$E_\tau \sum_{r=0}^{\infty} \beta^r z(m_{\tau+r}, c_{\tau+r})$$

where $u(\cdot, \cdot)$ and $z(\cdot, \cdot)$ are general homothetic functions (not necessarily of the same form), strictly increasing in the levels of the private goods f and m respectively, and in the level of the public good c . The expectation is taken with respect to the distribution of consumption allocation $\{f_{\tau+r}, m_{\tau+r}, c_{\tau+r}\}_{r=0}^{\infty}$, conditional on the information available at time τ . Finally, the price both of a unit of private and public goods is normalized to one.

⁴This means that F 's income realizations are high enough relative to M 's, and that his altruism level for the child is greater than some minimal level, which depends on M 's altruism level and the utility functional forms. For the case with bilateral transfers, see Hauser and Uysal (2006)

⁵This assumes some degree of contact between the parents, and between the father and child. Argys and Peters (2003) report that only 13% of fathers who settled personally their child support agreement had no contact with their children. Moreover, if the mother is aware that the transfers she receives are mitigated by the fact that the father cannot observe the expenditure on the child, she could provide evidence of expenditure and income to him.

3 First Best Allocation

In order to better understand the incentive problems that parents face in this context, it is useful to characterize the first best allocations where both can commit to a sequence of state contingent consumptions. We can solve for these allocations by writing down the planner's problem where the agents' incomes are pooled in one resource constraint. This is equivalent to writing an optimization problem where F maximizes his lifetime utility subject to a reservation utility for M . Letting μ be the relative Pareto weight on M 's utility, the planner's problem at τ is

$$\begin{aligned} \max_{\{f_{\tau+r}, m_{\tau+r}, c_{\tau+r}\}_{r=0}^{\infty}} \quad & E_{\tau} \sum_{r=0}^{\infty} \beta^r (u(f_{\tau+r}, c_{\tau+r}) + \mu_{\tau+r} z(m_{\tau+r}, c_{\tau+r})) \\ \text{s.t.} \quad & f_{\tau+r} + m_{\tau+r} + c_{\tau+r} = Y_{\tau+r}^F + Y_{\tau+r}^M \quad \text{for all } r. \end{aligned}$$

The first order conditions imply the following relation holds for all dates and states:

$$\begin{aligned} u_f(f_{\tau}, c_{\tau}) &= u_c(f_{\tau}, c_{\tau}) + \underbrace{\mu_{\tau} z_c(m_{\tau}, c_{\tau})}_{\text{wedge}} \\ z_m(m_{\tau}, c_{\tau}) &= z_c(m_{\tau}, c_{\tau}) + \underbrace{\frac{1}{\mu_{\tau}} u_c(f_{\tau}, c_{\tau})}_{\text{wedge}} \end{aligned}$$

Consider a hypothetical case where each agent could decide on how much to allocate to the public good out of the available budget, without taking into account the other agent's action. M 's decision would be given by $z_m(m_{\tau}, c_{\tau}) = z_c(m_{\tau}, c_{\tau})$ for all dates and states. Similarly, F would set $u_f(f_{\tau}, c_{\tau}) = u_c(f_{\tau}, c_{\tau})$ for all dates and states. These are agents' *individual optimality conditions*, which in the first best are never satisfied, since the additional terms on the right hand side of the first order conditions will never be equal to zero simultaneously. This is a standard result in settings with public goods since the social planner internalizes the effects of public good consumption decisions on both agents' utilities. Note that as the relative Pareto weight of M increases, the first best will prescribe a consumption which is increasingly aligned with her individually optimal consumption, thus decreasing the wedge between them, and vice versa for F .

This already gives an idea why, in a setting with a lack of commitment, agents may not be able to achieve the first best allocation. Although the first best achieves the largest joint surplus for the agents, it always dictates a provision of the public good that is too high, relative to their ideal consumptions. Generally, when an agent's Pareto weight is low, he or she will be tempted to deviate from the first best and pick the consumption combination which maximizes his or her period utility. The constrained optimal contract will find a "middle ground" solution which will bring agents closest to the first best payoffs, while still satisfying their incentive constraints.

4 Subgame Perfect Equilibria

Under the assumption of no outside enforcement, the parents cannot commit to the behavior prescribed by the first best equilibrium. We then look for self-enforceable contracts, meaning

incentive compatible agreements from which agents will not want to deviate. At this point, it is important to understand the incentives and disincentives of the parents from engaging in a long term agreement, instead of playing non cooperatively. For F , the benefit from making transfers to M is to increase the provision of the public good. The price he has to pay in return is a "tax" on these transfers by M , who will privately consume a part of them. Hence, any self-sustaining agreement should ensure that F gains enough from it to still make the optimal transfers. On the other hand, the benefit of this arrangement to M is that it increases her disposable income. In return, she has to distort her expenditure choice in favor of a higher amount of child consumption. This is true regardless of the altruism levels and holds even if the mother cares more about the child than the father. So a self-sustaining agreement should guarantee that once M receives F 's transfer, she would spend it in the agreed way.

4.1 Strategies

The interaction between agents in our environment involves a two-part decision making process in each period. At the beginning of period τ , both agents observe the realization of θ . F makes a nonnegative transfer: $t_\tau \in [0, Y_\tau^F]$ and consumes his post-transfer income $f_\tau = Y_\tau^F - t_\tau$. M decides how to split her post-transfer income between the child's consumption, c_τ , and the private good, $m_\tau = Y_\tau^M + t_\tau - c_\tau$.

Define an allocation $\{t_\tau, c_\tau\}_{\tau=1}^\infty$ to be a vector of state-dependent transfers and public good consumptions. A period τ history in this game consists of a sequence of realizations for θ, t and c :

$$h^\tau = (\theta_1, t_1, c_1, \theta_2, t_2, c_2, \dots, \theta_{\tau-1}, t_{\tau-1}, c_{\tau-1}, \theta_\tau)$$

A strategy for F at τ is a mapping from possible histories at τ into a transfer. M 's strategy is a mapping from possible histories and current transfer amounts into a public good consumption. A subgame-perfect equilibrium (SPE) specifies:

1. A strategy for F such that his transfer after any history is optimal, given M 's transfer and consumption strategies;
2. A consumption strategy for M given the observed history and current period transfer.

4.2 Minmax Punishments

The aim is to characterize the Pareto frontier of the set of SPE payoffs. A critical element to the environment is the punishment each agent faces if he/she were to deviate from the equilibrium path play. The worse the threat of the punishment, the larger the set of implementable equilibria. Usually, equilibria on the Pareto frontier can be supported by reverting to the autarkic allocation as punishment, where each agent consumes his or her income every period. This is the worst SPE of the game without public goods, which also achieves the agents' minmax values.⁶ In this model, the worst possible punishments entail

⁶The autarkic equilibrium of the standard game without a public good is also the Nash equilibrium of the static game. In the present model, this would correspond to the Stackelberg equilibrium, where

no transfers between agents as in the game without a public good, but also no provision of the public good when F is punished.

M 's minmax value is achieved by F making no transfer, and M spending her income to maximize her period utility. Her lifetime utility if minmaxed is:

$$\underline{W} = E_{\tau} \sum_{r=0}^{\infty} \beta^r z(Y_{\tau+r}^M - c_{\tau+r}^*, c_{\tau+r}^*)$$

where c_{τ}^* is M 's optimal child expenditure given that her income is Y_{τ}^M . F 's minmax entails M giving the child no consumption no matter what F transfers to her. This would give F a payoff of

$$\underline{V} = E_{\tau} \sum_{r=0}^{\infty} \beta^r u(Y_{\tau+r}^F, 0)$$

Because of the presence of the public good, minmaxing one's opponent can be very costly. Take M for example. Her payoff from minmaxing F is

$$W(\underline{V}) = E_{\tau} \sum_{r=0}^{\infty} \beta^r z(Y_{\tau+r}^M, 0)$$

which is even lower than her minmax value \underline{W} , so punishing F in that way is not incentive compatible for her. So if the mother threatens not to give the child any consumption, the father will not believe her. The same holds for F . Since he cares about the child's consumption, he would still like to transfer some of his income for M to spend on the child. Hence, it is obvious that the *minmax strategies are not subgame perfect*. However, it is possible to implement the *minmax values* using stick-and-carrot punishments which deliver high values to the punisher. Appendix A illustrates in detail how to achieve that for very general preferences, but I give here an intuitive description.

Imagine that F was being punished. The idea is to divide his punishment into two phases: the "stick" phase of the punishment where his period utility is lower than his minmax utility, and the "carrot" phase where his utility is higher than his minmax utility, such that on average, his expected lifetime utility in the stick-and-carrot scheme is equal to \underline{V} . For any utility level, F 's indifference curve passes through different combinations of private and public goods. Those with low private and high public consumptions correspond to higher levels of M 's utility. Pick such a point in the "stick" phase in order to minimize M 's loss from the punishment. Whenever both agents abide by the "stick" strategy, they are promised to move back to a specified point on the Pareto frontier (the "carrot") with some probability. From there on, normal play resumes. On the other hand, if F refuses to temporarily distort his consumption (so if he cheats on his punishment), the "stick" phase continues for certain. Similarly, if M deviates from F 's punishment, she starts off her own

M takes F 's transfer as given and maximizes her period utility, and F transfers accordingly. Taking the Stackelberg equilibrium to be the threat point in this game will have different implications on the optimal contract, generally reducing the set of feasible SPE payoffs, but also inducing some level of insurance off the equilibrium path.

punishment. The objective then is to find the "stick" phase allocation and the "carrot" phase continuation values which will maximize M 's value from punishing F , subject to the constraint that F 's value is \underline{V} , and to participation constraints of both agents. An analogous scheme is used for minmaxing M .

Proposition 1 *An allocation $\{t_j, c_j\}_{j=\tau}^{\infty}$ is subgame perfect if and only if it satisfies:*

$$u(Y_{\tau}^F - t_{\tau}, c_{\tau}) + E_{\tau} \sum_{r=1}^{\infty} \beta^r u(Y_{\tau+r}^F - t_{\tau+r}, c_{\tau+r}) \geq u(Y_{\tau}^F, 0) + \beta \underline{V}$$

$$z(Y_{\tau}^M + t_{\tau} - c_{\tau}, c_{\tau}) + E_{\tau} \sum_{r=1}^{\infty} \beta^r z(Y_{\tau+r}^M + t_{\tau+r} - c_{\tau+r}, c_{\tau+r}) \geq z(Y_{\tau}^M + t_{\tau} - c_{\tau}^*, c_{\tau}^*) + \beta \underline{W}$$

for all dates and states, where c_{τ}^* is the optimal deviation consumption defined as

$$c_{\tau}^* = \arg \max_C z(Y_{\tau}^M + t_{\tau} - C, C).$$

Proof. In Appendix B ■

If F deviates, he immediately sets off his punishment, whereby M would spend nothing on the child. Hence, F 's optimal deviation is to transfer zero. The right hand side of the second inequality is M 's payoff if she deviates by splitting her disposable income to maximize her period utility, and continues with her minmax value. Consider an allocation $\{t_j, c_j\}_{j=\tau}^{\infty}$ satisfying the conditions of the proposition above, and let the agents follow a strategy whereby they transfer and consume the amounts dictated by the allocation as long as both have done so in the past, otherwise, they revert to the deviating agent's punishment. These strategies define a contract.

5 Recursive Formulation

Definition *A subgame perfect allocation is efficient if and only if there is no other subgame perfect allocation that Pareto dominates it, and an optimal contract is one which implements such an allocation.*

Let \mathbf{V} be the maximal payoff F can obtain in a subgame perfect equilibrium, and \mathbf{W} be that attainable by M . Define the function $V : [\underline{W}, \mathbf{W}] \rightarrow [\underline{V}, \mathbf{V}]$ to be the following:⁷

$$\begin{aligned} V(W) &= \max_{\{t_{\tau}, c_{\tau}\}_{\tau=1}^{\infty}} E_0 \sum_{\tau=1}^{\infty} \beta^{\tau-1} u(Y_{\tau}^F - t_{\tau}, c_{\tau}) \\ \text{s.t.} & \quad \{t_{\tau}, c_{\tau}\}_{\tau=1}^{\infty} \text{ is a subgame perfect allocation.} \\ & \quad E_0 \sum_{\tau=1}^{\infty} \beta^{\tau-1} z(Y_{\tau}^M + t_{\tau}, c_{\tau}) \geq W \end{aligned}$$

⁷This definition assumes the convexity of the set of SPE. The maximization problem shows that the choice variable t_s enters the right hand side of M 's constraint positively, which makes it difficult to guarantee the convexity of the constraint set. If utility functions are homothetic of degree one, the deviation utility will be linear in the transfer, hence guaranteeing convexity. If the Pareto frontier is not concave, there may be gains from randomization. However, previous computations and results from Thomas and Worrall (1994) and Albuquerque and Hopenhayn (1994) point to the strict concavity of the frontier.

One can think of F as choosing the allocation $\{t_\tau, c_\tau\}_{\tau=0}^\infty$ to maximize his utility, while providing M with an ex-ante promised lifetime utility W , and satisfying the incentive constraints for every possible history. The function V is the Pareto frontier of subgame perfect equilibrium payoffs. As is standard in the literature, it is useful to solve the problem above recursively. Following Thomas and Worrall (1990) and Kocherlakota (1996), the frontier V can be characterized by the following recursive program:⁸

$$\begin{aligned}
V(W) &= \max_{\{t_s, c_s, W_s\}_{s=1}^S} \sum_{s=1}^S \pi_s \{u(Y_s^F - t_s, c_s) + \beta V(W_s)\} \\
\text{s.t.} & \quad \sum_{s=1}^S \pi_s \{z(Y_s^M + t_s - c_s, c_s) + \beta W_s\} = W \\
\text{and for all } s & : \\
& \quad z(Y_s^M + t_s - c_s, c_s) + \beta W_s \geq z(Y_s^M + t_s - c_s^*, c_s^*) + \beta \underline{W} \\
& \quad u(Y_s^F - t_s, c_s) + \beta V(W_s) \geq u(Y_s^F, 0) + \beta \underline{V} \\
& \quad t_s \geq 0 \\
& \quad W_s \in [\underline{W}, \mathbf{W}]
\end{aligned}$$

The first constraint is a standard promise-keeping constraint. The second and third constraints are M 's and F 's incentive constraints which ensure that the contract is self-enforceable. The non-negativity constraint on the transfers follows. Finally, the last constraint puts bounds on M 's utility from the contract, where her maximum value is determined endogenously.

The fact that M is the sole provider of the child's consumption means that F 's transfer, and subsequently the child expenditure level, are bounded above by M 's private consumption and continuation value. For example, even if F had a very high income shock in one state and wanted to split that surplus between his own and the child's consumptions, he would be restricted in doing so since any large transfer to M that is not matched by a substantial private consumption or continuation utility for her, would lead her to deviate.

Figures (1) and (2), which represent the generic Pareto frontiers of payoffs from the unconstrained and constrained problems, illustrate this fact. The origin depicts the minmax values for both agents, values on the X-axis denote M 's lifetime utility, W , while those on the Y-axis denote F 's lifetime utility, $V(W)$. The existence of the public good gives rise to the upward sloping parts of the utility possibility frontier since beyond a certain point, decreasing one agent's utility can only come at a cost to the other agent as well. The Pareto frontier is restricted to the downward sloping part, delimited by the points A and B. In the first picture, the constrained and first best frontiers partially overlap, so for some values of M , neither agent's incentive constraint binds, and the first best allocation is sustainable at these points. In the second picture, the constrained frontier lies entirely beneath the

⁸The Pareto frontier is self-generating, so continuation values will always be in the frontier. Agents will not resort to inefficient punishments on the equilibrium path, since higher continuation values always contribute positively to their incentives.

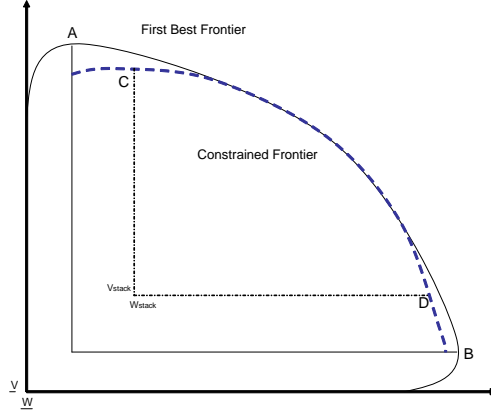


Figure 1: First Best and Constrained Pareto Frontiers 1

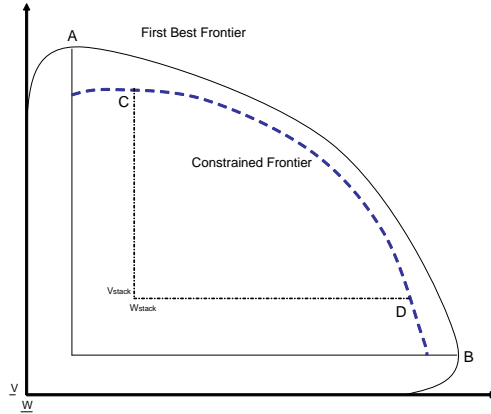


Figure 2: First Best and Constrained Pareto Frontiers 2

first best frontier, meaning that none of the values combinations which are feasible under the first best are so in the constrained problem. Finally, the two agents will only enter the contract if their ex-ante values exceed their Stackelberg values, so the initial point must lie somewhere between C and D.

6 Equilibrium

A standard result of two-sided lack of commitment models is that if one agent's incentive constraint binds in some particular state, she is compensated with higher consumption and continuation value. In the presence of a public good, is it better to provide higher consumption of the private good, or of the public good? Is there an optimal combination of the two that should be offered? Below are some answers using the first order conditions

of the problem. Let μ be the multiplier associated with the promise keeping constraint, and λ_s^F, λ_s^M and ϕ_s^F the multipliers associated with incentive constraints for F and M , and the nonnegativity constraint, for each state s , respectively. By homotheticity of M 's utility function, we know that her deviation utility will always involve a fixed ratio of m to c , so we can write $z(Y_s^M + t_s - c_s^*, c_s^*) = z(\alpha(Y_s^M + t_s), (1 - \alpha)(Y_s^M + t_s))$, where $0 < \alpha < 1$. The first order and envelope conditions imply the following:⁹

$$\begin{aligned} \frac{u_f(f_s, c_s)}{u_c(f_s, c_s)} &= 1 - V'(W_s) \left(\frac{z_c(m_s, c_s)}{u_c(f_s, c_s)} - \lambda_s^M \frac{\alpha z_m(m_s^*, c_s^*) + (1 - \alpha) z_c(m_s^*, c_s^*) - \phi_s^F}{u_c(f_s, c_s)} \right) \\ \frac{z_m(m_s, c_s)}{z_c(m_s, c_s)} &= 1 - \frac{1}{V'(W_s)} \frac{u_c(f_s, c_s)}{z_c(m_s, c_s)} \\ (1 + \lambda_s^F) V'(W_s) &= V'(W) - \lambda_s^M \end{aligned}$$

where m_s^* and c_s^* refer to the deviation consumption levels of M in state s . For any given promised value W to M , the states are divided into three sets: S_0^W where no incentive constraint binds, S_F^W where F 's incentive constraint binds, and S_M^W where M 's incentive constraint binds. What about the possibility of both agents' constraints binding in the same state? In the standard insurance model, the existence of a sustainable non-autarkic allocation is enough to show that there are gains from contracting, which also implies that both agents' incentive constraints do not bind simultaneously. In this case, since the minmax values are implemented using stick-and-carrot punishments which rely on values in the Pareto frontier (as shown in Appendix B), a similar statement will be meaningless. The correct assumption is: if there are initially no gains from cooperation, agents will be in the Stackelberg equilibrium of the stage game, which is also an equilibrium of the repeated game. However, if there exists an allocation which ex-ante dominates the Stackelberg allocation, agents will reach an agreement whereby any deviation by an agent will lead to minmaxing that agent.¹⁰

The following sections look at how each of the three sets of states evolves as M 's promised value W varies.

⁹Notice that this characterization relies on the differentiability of the Pareto frontier. Koepl (2003) finds sufficient conditions for the differentiability of the efficient frontier in risk sharing problems with lack of commitment. His proof is adapted to a Kocherlakota-type setting, but extends immediately to this model. His conditions are as follows. Let S_1 be the set of states where agent one's incentive constraint binds, and S_2 be the set of states where agent two's incentive constraint binds. If $S_1 \sqcup S_2 \neq S$ at W_0 , then V is differentiable at W_0 . If there exists an incentive compatible first best allocation, then V is differentiable everywhere. Another paper by Rincon-Zapatero and Santos (2007) shows that by changing the timing from an ex-ante to an ex-post optimization problem, the value function in Kocherlakota's model will always be differentiable.

¹⁰In this model of sequential interaction where M 's deviation utility depends on F 's realized transfer, one could imagine a contract which dictates in a certain state a large transfer from F to M (say when F 's income is highest relative to M 's), and a high child expenditure to match that. A natural question is whether both agents' constraints could bind sequentially in the same period. The answer is no, since the worst that the contract can do is the Stackelberg allocation in which constraints are slack. Starting from that point, increase one agent's utility until the other agent's constraint binds.

6.1 States where no incentive constraint binds

For a given promised utility, if no incentive constraint binds in a set of the states, agents will be able to achieve the first best allocation in these states in the contract. Attach the FB subscript to the resulting consumption and transfer values. The first order conditions imply that for a given value W , the ratios of consumptions should be equal in all states in S_0^W , so the consumptions f_s^{FB} , m_s^{FB} and c_s^{FB} are each a constant fraction of the total income \mathbf{Y}_s .¹¹ Moreover, as M 's relative Pareto weight μ increases, her private consumption increases both in absolute terms, and relative to the child's ($\partial (c_s^{FB}/m_s^{FB})/\partial\mu < 0$), hence shrinking the wedge between her individually optimal consumption and her actual consumption. The reverse holds for F . The envelope condition implies that the continuation values of agents are also constant for states in S_0^W .

6.2 States where F 's incentive constraint binds

When F 's incentive constraint binds, the outcome is similar to an increase in his relative Pareto weight. The envelope condition implies that $V'(W_s) > V'(W)$, so his continuation value increases. His private consumption increases both in absolute terms, and relative to the child's, while the opposite holds for M . Looking back at F 's incentive constraint and noting that for a given state s , the transfer t_s is increasing in W , one can see that the lower F 's value, the more difficult he finds it to comply with the contract allocation. And since the right hand side of the inequality depends only on F 's income in a given state and not on his current value, so will the compensation when his incentive constraint binds. The period deviation utilities of F are state-dependent, but independent of his promised value, and so will his consumption values and continuation utilities be. The following proposition formalizes the claims made in the last two subsections:

- Proposition 2** 1. If $s \in S_0^W$ and $S_0^{\tilde{W}}$ where $W < \tilde{W}$ then $m_s < \tilde{m}_s$ and $(m_s/c_s) < (\tilde{m}_s/\tilde{c}_s)$, where \tilde{m}_s and \tilde{c}_s are consumption values at \tilde{W} .
2. If $s \in S_F^W$, then $f_s > f_s^{FB}$, $(m_s/c_s) < (m_s^{FB}/c_s^{FB})$ and $W_s < W$.
3. If $s \in S_F^W$, then $s \in S_F^{\tilde{W}}$ for all $\tilde{W} \geq W$. Moreover, $f_s = \tilde{f}_s$ and $V(W_s) = V(\tilde{W}_s)$, where \tilde{f}_s is F 's consumption and $V(\tilde{W}_s)$ his continuation value in state s , for \tilde{W} .

Proof. In Appendix B ■

The sets S_0^W and S_F^W can easily be compared to those in the model without public good. The higher the relative income today, the more likely is the incentive constraint to bind. Agents are compensated through higher private consumption, allocations that are increasingly aligned with their taste for the public good, as well as through continuation utilities that covary positively with income levels. The last part of this analysis studies the set of states where M 's incentive constraint binds. This is maybe the most interesting

¹¹In the case of no aggregate uncertainty where $\mathbf{Y}_s = \mathbf{Y}$ for all s , this would correspond to complete insurance against income shocks.

case, since it captures the moral hazard problem related to the provision of the public good specifically, as opposed to that related to the insurance aspect.

6.3 States where M 's incentive constraint binds

In this set of states, the first best allocation provides M with a lower period utility than what she would get by spending the transfer as her individual optimality condition dictates. F can compensate M in three ways: by increasing her continuation utility, increasing her disposable income, or, for a given level of disposable income, by shifting her expenditure from c_s to m_s . The envelope condition implies that M 's continuation utility is indeed higher. Looking again at M 's incentive constraint, one sees that the transfer t_s enters positively on both sides of the inequality, so increasing M 's disposable income would exacerbate her incentive problem by granting her a larger income with which to abscond. As opposed to the case where F 's constraint binds, M 's problem is alleviated by allowing her a *lower* fraction of income, but still letting her increase her private consumption relative to the child's, which necessarily falls. Hence, both agents' ratios of private consumption to the child's consumption rise.¹² Notice that this is also different from the insurance case with no public good, where one agent's constraint binding always means a lower consumption for the other agent. This is an important result as it clarifies why the solution to the low consumption of children from divorce may not necessarily reside in increasing the custodial mothers' incomes, or at least not doing so unconditionally.

Proposition 3 1. If $s \in S_M^W$, then $(m_s/c_s) \geq (m_s^{FB}/c_s^{FB})$, $(f_s/c_s) > (f_s^{FB}/c_s^{FB})$, $c_s < c_s^{FB}$ and $W_s > W$.

2. If $s \in S_M^W$, then $s \in S_M^{\tilde{W}}$ for all $\tilde{W} < W$. Moreover, $\tilde{m}_s \leq m_s$ and $\tilde{W}_s \leq W_s$.

Proof. In Appendix B ■

In the standard model of risk-sharing as well as in this model, if the incentive constraint of a transferring agent binds in all states, it must be that the agent is receiving the minmax lifetime utility. Here, there may be a *set of continuation utilities* for which M 's incentive constraint binds for all states. This is a direct implication of M 's deviation utility being a function of the transfer since her disposable income in these states is $Y_s^M + t_s$.

7 Long Run Equilibrium

In the long run, agents may converge to two types of equilibria, depending on the surplus from contracting: one where some first best allocation is ultimately implementable, and one where their incentive constraints will always bind. What matters in determining the long run is the agents' continuation values, rather than their period consumptions. A quick recapitulation is useful to better visualize the long run equilibrium:

¹²Comparing m_s to m_s^{FB} analytically without further assumptions on functional forms is not possible, but numerical computations show that for CES and square root utility functions, m_s is higher than m_s^{FB} .

If F 's incentive constraint binds for a given state at some value W , it will bind in that state for all $\tilde{W} > W$. Hence, there exists a threshold value W_s^F for which F 's constraint starts binding for a given income level Y_s^F . Denote the lowest threshold, which corresponds to some income realization Y_j^F , by W_{\max} .¹³ When M 's incentive constraint binds for a given state at some value W , it will bind for that state for all her lower values. Hence, there exists a threshold value W_{\min} , above which her incentive constraint doesn't bind in any state. It is easy to show that $W_{\min} < W_{\max}$, so the two values delimit two regions and in a way shape the long run equilibrium. We can now turn to the characterization of the long run equilibrium of the model.

7.1 When Subgame Perfect First Best (SPFB) Allocations Exist

The properties of the long run equilibrium are similar to the case with no public good, and depend on the existence of an ex-ante value W of M where the first best allocation is implementable.

Proposition 4 *If there exists an interval of the Pareto frontier where the first best is implementable, the long term equilibrium pair of values belongs to this interval. If the initial promised utility to M , $W_0 < W_{\min}$, agents' values converge to $(W_{\min}, V(W_{\min}))$. If the initial promised utility to M , $W_0 > W_{\max}$, agents' values converge to the pair $(W_{\max}, V(W_{\max}))$. If M starts with any promised utility W_0 between W_{\min} and W_{\max} , agents' values stay at $(W_0, V(W_0))$.*

Proof. In Appendix B ■

Proposition 4 implies that the starting point of agents' values does matter in the Long Run. If one imagines that the starting point was reached through some bargaining between the parents, then the long run equilibrium will always reflect their initial bargaining powers. Since the first best is achievable here, the lack of commitment and the public good provision problem are both irrelevant in the long run.

A natural question to ask is when does a SPFB allocation exist? Kocherlakota (1996) shows that a necessary and sufficient condition is for the symmetric first best allocation to be Subgame Perfect in the case of agents with identical preferences and income processes and no public good. In this model with income and preference heterogeneity, a condition as simple and concise does not apply. Evidently though, there exists a $\bar{\beta}$ such that a first best allocation is sustainable for all $\beta \geq \bar{\beta}$. So as agents become sufficiently patient, any constrained allocation will converge to a first best allocation.

7.2 When no SPFB Allocation Exists

What happens when there is no value W where the first best allocation is implementable? Whatever the agents' values, there is at least one state for which some incentive constraint binds, and a positive probability that continuation values will be different from current

¹³ $V(W_{\max})$ is also F 's continuation value when his incentive constraint binds for income Y_j^F at any value W of M .

values. The next proposition says that agents' values will oscillate over a set of values, and that the set is independent of the agents' starting point.

Proposition 5 *If there exists no SPFB allocation, the conditional distribution of M 's continuation value will converge to a nondegenerate distribution. Moreover, this distribution is unique and does not depend on the initial value W_0 .*

Proof. In Appendix B ■

Under this type of equilibrium, the set of agents' possible values will be the same, regardless of their starting points, and the continuing fluctuation in the consumption choices will always reflect the lack of commitment problem.

One would like to know what novelty the presence of the public good brings to the analysis, and how it affects the long run behavior. The answer to that question depends on the assumption of whether there is aggregate uncertainty (so total income is constant in all states) or not. Given that in the absence of aggregate uncertainty the first best allocations display perfect insurance, the right-hand side of M 's incentive constraint is constant across states. Hence, if her constraint binds in one state, it will bind in all states, and the consumption allocations and the continuation values will be equal across states (a kind of suboptimal but perfect risk sharing). This also means that there exists a value (W_{\min}) below which M 's constraint binds for all values and all states and above which it never binds. In the absence of aggregate uncertainty, M 's incentive constraint will not bind in the long run: once her value goes above W_{\min} , it will never go again below it. On the other hand, F 's incentive constraint starts binding for high income levels first and high values of W , then it binds for all states when he's at his minmax value. Since both agents' incentive constraints cannot bind simultaneously in a state, there should exist a value where a SPFB allocation is implementable. If on the other hand there is aggregate uncertainty, for a given value W , M 's incentive constraint may bind in some states but not in others and the ranges of values over which each agent's constraint may bind can overlap, so both types of equilibria are possible.

8 Efficient Equilibria Under Strict Enforcement:

Equipped with the characterization of the set of SPE from the previous sections, we are set to ask policy questions and conduct relevant comparative statics. What if information was worse, or incomes were more unequal, or preferences more similar? Would better enforcement solve the lack of commitment problem? The next section addresses precisely that last question by simulating a policy change which enforces transfers from F to M that are equal to a fixed percentage of F 's income.

8.1 US Institutional Background

Federal endeavors to institute a national child support program started in 1975 with the passage of title IV-D of the Social Security Act, which aimed at supporting states in establishing paternity and locating absent parents, as well as determining and enforcing child

support obligations. Since then, total IV-D expenditure has grown steadily,¹⁴ and strict laws governing child support awards and enforcement of payments have emerged.

Child Support Guidelines: Until the mid 80's, child support awards were left to the judge's discretion and decided on a case by case basis. The Child Support Amendments of 1984 required all states to adopt guidelines by 1987, based on a numerical formulation. , in the aim of standardizing and increasing child support awards. Guidelines are state specific, and are usually calculated as a function of the noncustodial parent's income, the number of children, and rarely, the custodial parent's income. Many awards are simply set to be 15 to 20% of the noncustodial parent's gross income at the time of separation, and adjust only every few years.¹⁵

Child Support Enforcement: The second obstacle to tackle was the amount actually paid by noncustodial fathers. Previously, mothers had to file separately for enforcement through court, a too costly and time consuming procedure. The Amendments of 1984 established wage withholding when payments were 1 month overdue, whereby the child support amount is deducted from the father's paycheck and transferred directly to the mother. By 1996, all states had adopted an automatic wage withholding policy from the time of the order.¹⁶

The benefits from child support enforcement are clear in the sense of guaranteeing that noncustodial parents contribute their share in the upbringing expenses of their children. Many studies have looked at the effectiveness of these policies and have found that, despite the majority of mothers relying on income withholding for their child support receipts, the total amounts paid did not increase dramatically over the last twenty years.¹⁷ Moreover, there are no studies, as far as we know, which have examined specifically child welfare as a result of those laws. A static analysis of the problem would necessarily imply that increased child support payments will result in higher children consumption. The analysis below shows that in a dynamic setting, the answer is ambiguous, since enforcement will generally weaken the mother's incentive to spend a high fraction of their income on the child.

8.2 New Outside Values

When we introduce one-sided enforcement into the picture, the outside options of the agents are no longer equal to their old minmax values. In fact, letting t_{enf} be the mandatory

¹⁴In 2006, \$3.9 billion were available for child support enforcement alone.

¹⁵The Wisconsin standard dictates 17% of the father's gross income for one child.

¹⁶Among other enforcement measures are income tax refund and unemployment benefits interception, suspension of professional licenses, freezing and seizing of accounts, passport denial.

¹⁷See Cancian and Meyer (2005) or Hauser (2008b) for more on that issue. The main reasons put forward are the changing composition of single mothers (from less than 20% who were never married in 1976 to 45% in 1997) and the difficulty of tracking down fathers who move to different states and change jobs. Another reason is that voluntary agreements in reality resulted in higher awards, payments and compliance rates than court orders.

transfer dictated by the law, the new minmax values V_{enf} and W_{enf} are equal to:

$$V_{enf} = \frac{1}{1-\beta} \sum_{s=1}^S \pi_s u(Y_s^F - t_{enf}, 0)$$

$$W_{enf} = \frac{1}{1-\beta} \sum_{s=1}^S \pi_s z(Y_s^M + t_{enf} - c_s^*, c_s^*)$$

It is evident that $\underline{W} < W_{enf}$ since M is now guaranteed some minimal payments, as opposed to none in her old punishment. On the other hand, $\underline{V} > V_{enf}$ since F 's minmax payoff is equivalent to what he *would receive* if he transferred some of his income to M , and she spent nothing on the public good. Having the possibility of privately contracting, the agents will resort to the enforcement option only if they fail to come to an agreement. Once in a self-sustainable contract, agents will never turn to that option in equilibrium. Still, the possibility of enforcement, by altering the outside values, will alter the whole equilibrium play.

First, M 's individual rationality constraint means that a set of her lower lifetime utilities (from \underline{W} to W_{enf}) cannot be sustained anymore. Moreover, by raising her outside option, her incentive constraint will start to bind in some states for some values where the first best was achievable before. Finally, for states and values where her incentive constraint was binding, the old allocation will not be incentive compatible any longer. The opposite is true for F . A set of new lower lifetime utilities (from V_{enf} to \underline{V}) which was not sustainable before will become so. Lowering F 's outside option means that his incentive constraint will stop binding for some states and values where it was binding. The overall expected effect is a *rise in the ratio of M 's private to public consumption*.

An important point here is that the new set of SPE does not correspond exactly with the old one. Indeed, there will be a shift in the Pareto frontier, making the comparison between the old equilibrium and the new one rather tricky. As the next two subsections show, the result will depend on the type of long run equilibrium which was in place before the enforcement laws came in.

8.3 Case where SPFB Allocations Exist Before and After Enforcement

Figure (3) illustrates digrammatically the Pareto frontier before and after the introduction of the enforcement policies. The dotted blue line depicts the old constrained frontier, while the dashed red line depicts the new frontier. On the old frontier, for values of M smaller than W_{\min} , her incentive constraint binds for some states, and for values greater than W_{\max} , F 's binds for some states. On the new frontier, the values W'_{\min} and W'_{\max} delimit the new interval of values where the first best is implementable. Letting W_{LR} and W'_{LR} be the old and new long run equilibrium values, the following proposition holds:

Proposition 6 *In the case where SPFB allocations exist before and after enforcement, $W'_{\min} > W_{\min}$, $W'_{\max} > W_{\max}$ and $W'_{LR} \geq W_{LR}$. In terms of consumption, $(m_s/c_s)' \geq (m_s/c_s)$ for all states, where $(m_s/c_s)'$ is the ratio of consumptions at W'_{LR} , and (m_s/c_s) that at W_{LR} .*




Figure 3: Pareto Frontier Before and After Enforcement

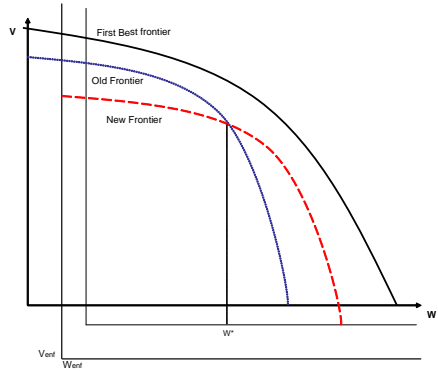


Figure 4: Pareto Frontier Before and After Enforcement

Proof. In Appendix C ■

8.4 Case where no SPFB Allocation Exists Before and After Enforcement

As in the previous case, the two frontiers will intersect at some value W^* , where for each $W < W^*$, the old frontier is closer to the first best Pareto frontier and for any $W > W^*$ the new frontier is closer. This reflects on one hand the loss from a weaker threat to M and on the other hand the benefit from a harsher punishment on F . Figure (4) illustrates this. One key property to remember is the following ranking of the ratio m/c :

$$\begin{aligned} \text{Individual optimality} &> \text{When } M\text{'s constraint binds} > \\ \text{First Best} &> \text{When } F\text{'s constraint binds} \end{aligned}$$

Out of the contract, the ratio m/c is the highest. Once in the contract, M will have an incentive to spend a higher fraction of her disposable income on the public good. Moreover, for a given value W , m/c will be greater in the states where M 's incentive constraint binds

than in states where no constraint binds, and finally it is lowest when F 's constraint binds. This leads to the following result:

Proposition 7 *The non-weighted average ratio $(m/c)'$ along the new constrained frontier is strictly greater than the ratio (m/c) along the old frontier.*

Proof. In Appendix C ■

Remember that under this type of equilibrium, agents' values will oscillate over a set of values which is independent of the starting point. In order to compare the average (m/c) between the old and the new equilibrium, one needs to know the probability distribution over the different states and values in the two equilibria. This is a difficult task and I don't have a formal proof to the claim that (m/c) will indeed be strictly higher after the enforcement laws, but in appendix C, I present illustrative examples with two and three states, showing why the claim is true in those cases.

9 Conclusion

This paper extends the problem of efficient risk sharing to include a public good. I consider the special case where one agent controls the public good provision and the play in the stage game is sequential. The public good provider may find it too costly to supply a high level of public consumption even after she's received a transfer, which would lead to the underprovision of the good. I characterize the Pareto frontier of SPE payoffs using the minmax payoffs as punishment threats and compare the equilibria under double-sided lack of commitment and after introducing a transfer enforcement policy. A consequence of the equilibrium with mandatory transfers is that the providing agent will spend a higher fraction of her income on her private consumption and a lower fraction on the public good than in the absence of enforcement. This is not surprising, given that the optimal contract specified a set of contingent expenditure which usually differed from the allocation she would have picked, were it not for the threat of punishment.

This exercise is motivated partly by the response of US lawmakers to the problem of noncustodial fathers' lack of compliance with child support orders. As a way of guaranteeing the welfare of children (which are considered to be public goods in the couple) from divorced and separated parents, strict laws governing child support awards and enforcement of payments have emerged which generally withhold a percentage of the father's income and give it automatically to the custodial mother. If indeed the fathers' low compliance is due to their lack of control over how mothers spend their income, the one-sided enforcement of child support contracts may worsen the mothers' moral hazard problem and lead to even more suboptimal results. Hauser (2008) takes precisely the theoretical prediction from this model and finds strong empirical support for it. Using Consumer Expenditure Survey data, it finds that the ratio of mother to child consumption has increased for mothers who went from not having access to the enforcement mechanism to having it, while that ratio did not change for single mothers who did not receive child support, and hence, did not benefit from the enforcement laws. Without further assumptions on the model primitives (such as parents utility functions and altruism levels, their income processes, etc) it is impossible

to conduct a welfare analysis of the policy in question. Still, this highlights the balance obtained through the double threat in a situation with lack of commitment, and questions the suitability of a one-sided fix with no accountability requirement on the second party.

References

- [1] Abraham, A. and E. Carceles-Poveda (2006), "Endogenous Trading Constraints with Incomplete Asset Markets," *University of Rochester, mimeo*.
- [2] Abreu, D. (1998), "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica* 56, 383-39.
- [3] Admati, A.R. and M. Perry (1991), "Joint Projects Without Commitment," *Review of Economic Studies* 58:2, 259-276.
- [4] Albuquerque, R. and H. Hopenhayn (1994), "Optimal Lending Contracts and Firm Dynamics," *Review of Economic Studies* 71:2, 285-315.
- [5] Alvarez, F. and U. Jermann (2000), "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica* 68:4, 775-797.
- [6] Argys, L., J. Brooks-Gunn, H. Peters and J. Smith (1998), "The impact of Child Support on Cognitive Outcomes of Young Children," *Demography* 35:2, 159-173.
- [7] Atkeson, A. (1991), "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica* 59, 1069-1090.
- [8] Bergstrom, T., L. Blume and H. Varian. (1986), "On the Private Provision of Public Goods," *Journal of Public Economics*, 29:1, 25-49.
- [9] Cancian, M. and D. Meyer (2005), "Child Support and the Economy," *National Poverty Center Working Paper Series #05-17*.
- [10] Clementi, G.L. and H. Hopenhayn (2006), "A Theory of Financing Constraints and Firm Dynamics," *Quarterly Journal of Economics*, 121:1, 229-265.
- [11] Del Boca, D. and C. Flinn (1994), "Expenditure Decisions of Divorced Mothers and Income Composition," *Journal of Human Resources* 29: 3, 742-761.
- [12] Flinn, C. (2000), "Modes of Interaction Between Divorced Parents," *International Economic Review* 41:3, 545-578.
- [13] Hauser, C. (2008), "Child Support Enforcement and Children's Consumption," *Working Paper, Collegio Carlo Alberto*.
- [14] Hauser, C., and G. Uysal (2006), "On the Dynamic Provision of Public Goods with Lack of Commitment," *University of Rochester, mimeo*.

- [15] Kehoe, P. and F. Perri (2004), "Competitive Equilibria with Limited Enforcement," *Journal of Economic Theory* 119:1, 184-206.
- [16] Kletzer, K.M. and B.D. Wright (2000), "Sovereign Debt as Intertemporal Barter," *American Economic Review* 70:3, 90:3, 621-639.
- [17] Knox, V. W. (1996), "The Effects of Child Support Payments on Developmental Outcomes for Elementary School-Age Children," *Journal of Human Resources* 31:4, 816-840.
- [18] Kocherlakota, N. (1994), "Efficient Bilateral Risk Sharing without Commitment." *University of Iowa working paper*.
- [19] Kocherlakota, N. (1996), "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies* 63:4, 595-609.
- [20] Koepl, T. (2004) "Differentiability of the Efficient Frontier when Commitment to Risk Sharing is Limited," *mimeo*.
- [21] Marimon, R. and V. Quadrini (2005), "Competition, Innovation and growth with Limited Commitment," *mimeo*.
- [22] Marx, L.M. and S.M. Matthews (2000), "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies* 67:2, 327-358.
- [23] Rincon-Zapatero, J.P. and M. Santos (2007), "Differentiability of the value function without interiority assumptions," *mimeo*.
- [24] Sorensen, E. and A. Hill (2004), "Single Mothers and their Child-Support Receipt. How Well Is Child Support Doing?," *Journal of Human Resources* 39:1, 135-154.
- [25] Thomas, J. and T. Worrall (1988), "Self-Enforcing Wage Contracts," *Review of Economic Studies* 55, 541-554.
- [26] Thomas, J. and T. Worrall (1990), "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem," *Journal of Economic Theory* 51, 367-390.
- [27] Thomas, J. and T. Worrall (1994), "Foreign Direct Investment and the Risk of Expropriation," *Review of Economic Studies* 61:1, 81-108.
- [28] Varian, H.R. (1994), "Sequential Provision of Public Goods," *Journal of Public Economics*, 53, 165-186.
- [29] Warr, P. (1983), "The Private Provision of a Public Good is Independent of the Distribution of Income," *Economics Letters*, 13(2-2), 207-211.
- [30] Weiss, Y. and R. Willis, (1985) "Children as Collective Goods and Divorce Settlements," *Journal of Labor Economics*, 3:3, 268-292.

- [31] Weiss, Y. and R. Willis, (1993) "Transfers among Divorced Couples: Evidence and Interpretation," *Journal of Labor Economics*, 11:4, 629-679.

10 Appendix A: Implementing the Minmax Payoffs

10.1 F 's Punishment:

M 's minmax requires F to make no transfer to M , and M to spend her income to maximize her period utility. Her lifetime utility if minmaxed is:

$$\underline{W} = \frac{1}{1-\beta} \sum_{s=1}^S \pi_s z(Y_s^M - c_s^*, c_s^*)$$

where c_s^* is M 's utility maximizing level of the public good, given her income Y_s^M . F 's minmax value is achieved by M not spending anything on the public good no matter what F transfers to her. This would give F and M respectively a payoff of

$$\begin{aligned} \underline{V} &= \frac{1}{1-\beta} \sum_{s=1}^S \pi_s u(Y_s^F, 0) \\ W(\underline{V}) &= \frac{1}{1-\beta} \sum_{s=1}^S \pi_s z(Y_s^M, 0) \end{aligned}$$

As said before, this value is lower than M 's own minmax value, so minmaxing F in that way is not incentive compatible for her. However one could implement F 's value in a different way, by distorting F 's consumption such that he still receives his minmax, while ensuring M gets as high a utility as possible. For that, stick-and carrot punishments are used. The objective is to deliver the expected value \underline{V} to F in two phases: the stick phase, where F 's per-period utility $v_s < (1-\beta)\underline{V}$, and the carrot phase, where F 's per-period utility $v_c > (1-\beta)\underline{V}$, such that F 's expected value is

$$\underline{V} = E \left(\sum_{\tau=0}^r \beta^\tau v_s + \sum_{\tau=r+1}^{\infty} \beta^\tau v_c \right)$$

and M gets a high enough payoff so that she would want to participate in F 's punishment. For that, we can write the following maximization problem for M

$$\begin{aligned} \overline{W} &= \max_{t_s, c_s, W} \sum_{s=1}^S \pi_s z(Y_s^M + t_s - c_s, c_s) + \beta(\rho W + (1-\rho)\overline{W}) \\ \text{st} \quad &\sum_{s=1}^S \pi_s u(Y_s^F - t_s, c_s) + \beta(\rho V(W) + (1-\rho)\underline{V}) = \underline{V} \end{aligned}$$

Let \overline{W} be M 's value from F 's punishment. While in the stick phase, agents' values lie below the Pareto frontier of payoffs (and outside the set of SPE payoffs). Every period, there is a probability ρ of ending the stick phase and turning to the carrot phase of the punishment. In that phase, agents return to the Pareto frontier and resume the contract: M 's value jumps to W , while F 's is $V(W)$. The punishment specifies state contingent transfers and

consumption values, as well as a continuation value W , subject to the constraint that F 's value from the punishment path is equal to his minmax value.¹⁸ This means that during the stick phase, F 's utility is lower than his minmax utility: while the minmax allows F to privately consume all his income, the stick phase of the punishment requires F to transfer part of his income, with possibly a very small part of that transfer going toward the provision of the public good, and the largest part being consumed by M .

In order for the punishment to be sustainable, the incentive constraint for M should specify that in the case of F deviating by transferring any amount other than the punishment amount, M will prefer to stick to the punishment path, than to cheat on it by spending according to her optimal rule, and continuing with her minmax value. Thus, the retribution for cheating on the other agent's punishment is by switching to the cheater's own punishment. This idea can be expressed in two ways. A simple (albeit restrictive) way, is to say that M 's incentive constraint should hold for any transfer that F may make, that is

$$\begin{aligned} \text{for all } s, \quad z(Y_s^M + T_s - C_s, C_s) + \beta \bar{W} &\geq z(Y_s^M + T_s - C_s^*, C_s^*) + \beta \underline{W} \\ \text{where } C_s^* &= \arg \max_C z(Y_s^M + T_s - C, C) \\ \text{and } u(Y_s^F - T_s - C_s) &\leq u(Y_s^F - t_s - c_s) \end{aligned}$$

In reality, this last constraint is more stringent than needed since all we need to make sure of is that M 's IC holds for any transfer that F would credibly make, even if that should lead her to deviate and revert to her punishment path.

10.2 M 's Punishment:

M 's minmax entails F not transferring any amount to her, and her spending her income according to her own rule. Again here, minmaxing M may not be incentive compatible for F . So I write a maximization problem for F in order to ensure he gets the maximum utility while giving M her minmax value, so he would participate in the punishment. One restriction is that the "stick" phase of the punishment cannot contain any transfers, since M would then deviate and score a utility higher than her minmax value. So the only way in which this stick-and-carrot punishment path differs from the minmax strategy is by having M spend a higher amount on the public good than she would on her own (so giving her momentarily a utility lower than her minmax utility), in return for a promise of a higher continuation value W' in the carrot phase. The maximization problem for F is the following:

$$\begin{aligned} \bar{V} &= \max_{c_s, W'} \sum_{s=1}^S \pi_s u(Y_s^F, c_s) + \beta(\rho V(W'_s) + (1 - \rho)\bar{V}) \\ \text{st} \quad &\sum_{s=1}^S \pi_s z(Y_s^M - c_s, c_s) + \beta(\rho W'_s + (1 - \rho)\underline{W}) = \underline{W} \end{aligned}$$

¹⁸Alternatively, one can specify W and pick ρ optimally.

The incentive constraint for M looks similar to her constraint in F 's punishment:

$$\text{for all } s, \quad z(Y_s^M + T_s - C_s, C_s) + \beta \overline{W} \geq z(Y_s^M + T_s - C_s^*, C_s^*) + \beta \underline{W}$$

where \overline{W} is M 's value from F 's punishment, as calculated above.

To compute the values \underline{W} and \overline{W} , one needs to solve first for the Pareto frontier, assuming that \underline{W} and \underline{V} are implementable. Then, the consumption values in the stick and carrot phases of F 's punishment can be calculated, followed by the consumption values in M 's punishment. The Folk Theorem for normal games says that it is possible to achieve the minmax values with appropriate punishments if agents are patient enough. The equivalent statement in this case is: there exists a minimum $(\widehat{\beta\rho})$, such that for all $(\beta\rho) \geq (\widehat{\beta\rho})$, the minmax payoffs are implementable.

Two issues may be of concern here. One is: are these punishments hopelessly complex? And second, are they renegotiation proof? For example, in the case of F 's punishment, if the value pair $(\underline{V}, W(\underline{V}))$ is Pareto dominated, both agents may benefit from renegotiating the punishment and resetting their values to a pair on the Pareto frontier. In general, renegotiation proofness will not be an issue if M 's value $W(\underline{V})$ is high enough, which is aided by M 's utility in the stick phase being high, or the stick phase being relatively short (which is equivalent to F 's utility being very low). I do not offer a technical proof of renegotiation proofness, but suggest examples of mechanisms which may answer both concerns:

Take again F 's punishment. We know the following hold:

$$\begin{aligned} u(0, c_s^*) &\leq v_s < u(Y_s^F, 0) < v_c \\ z(Y_s^M, 0) < z(Y_s^M + Y_s^F, c_s^*) &\quad \text{and} \quad z(Y_s^M, 0) < z(v_s) \end{aligned}$$

where $c_s^* = \arg \max z(Y_s^M + Y_s^F, c)$ and $z(v_s)$ is M 's utility from the stick phase in state s . One simple possibility is to pick $v_s = u(0, c_s^*)$, and in the carrot phase move to the corner of the Pareto frontier which gives M her highest in the contract, and finally choose ρ such that

$$\underline{V} = \sum_{s=1}^S \pi_s u(0, c_s^*) + \beta (\rho V_c + (1 - \rho) \underline{V})$$

V_c being the lowest value F can get in the contract. This ensures M a payoff higher than any payoff on the frontier so the punishment is renegotiation proof. In general, this scheme will work best when the complementarity between private consumption and public consumption is high for both agents. For the case where the substitutability between private and public consumption is high for both agents, the stick phase of the punishment will simply prescribe a low level of both F 's private consumption and of the public good, along with a high private consumption for M . The same applies when the public and private goods are highly complementary for F and highly substitutable for M . The case which is potentially the most problematic is when the public and private goods are highly substitutable for F and highly complementary for M . In this instance, it will be very difficult to grant M a high utility while achieving a low value for F .

11 Appendix B: Proofs of Propositions:

Proposition 1

Proof. The proof is similar to the proof of proposition 2.1 of Kocherlakota (1996). The main points to observe are that \underline{V} and \underline{W} are indeed the worst punishments and that the deviation utilities on the right hand sides of the inequalities are the results of best responses. ■

Proposition 2

Proof. **1) and 2)** In the case of homothetic preferences, the ratio of marginal utilities depends only on the ratio of consumptions, and not on the individual levels of consumption. This means that we can write

$$\frac{u_f(f_s, c_s)}{u_c(f_s, c_s)} = h\left(\frac{c_s}{f_s}\right) \quad \text{and} \quad \frac{z_m(m_s, c_s)}{z_c(m_s, c_s)} = g\left(\frac{c_s}{m_s}\right)$$

where $h()$ and $g()$ are increasing in their argument. The first order conditions in the states where no incentive constraint or F 's incentive constraint binds can be written as:

$$\begin{aligned} h\left(\frac{c_s}{f_s}\right) &= 1 - V'(W_s) \frac{z_c(m_s, c_s)}{u_c(f_s, c_s)} \\ g\left(\frac{c_s}{m_s}\right) &= 1 - \frac{1}{V'(W_s)} \frac{u_c(f_s, c_s)}{z_c(m_s, c_s)} \\ (1 + \lambda_s^F) V'(W_s) &= V'(W) \end{aligned}$$

which implies that

$$g\left(\frac{c_s}{m_s}\right) = 1 + 1 / \left(h\left(\frac{c_s}{f_s}\right) - 1 \right)$$

meaning that c_s/m_s and c_s/f_s , and hence m_s and f_s move in opposite directions. Moreover, dividing $h\left(\frac{c_s}{f_s}\right)$ by $g\left(\frac{c_s}{m_s}\right)$ gives

$$\frac{u_f(f_s, c_s)}{z_m(m_s, c_s)} = -V'(W_s)$$

These two conditions mean that when $-V'(W_s)$ increases (so W decreases), f_s increases and m_s decreases, which is the case when F 's incentive constraint binds, and part (2) of the proposition follows.

In the case where no incentive constraint binds, $V'(W_s) = V'(W)$, so applying the same reasoning, part (1) of the proposition follows.

3) For the first part, note that for F 's incentive constraint to bind in state s , the following has to be true:

$$u(Y_s^F - t_s^{FB}, c_s^{FB}) + \beta V(W) < u(Y_s^F, 0) + \beta \underline{V}$$

The right hand side is independent of W , while the left hand side is decreasing in W by concavity of the frontier. For the second part, since the deviation utility depends only on current income and not on the current value, so should the efficient compensation be. ■

Proposition 3

Proof. 1) First notice that by the homotheticity of $z(\cdot, \cdot)$, the function $Z_{dev} = (z(Y_s^M + t_s - c_s^*, c_s^*) - z(Y_s^M + t_s - c_s^{FB}, c_s^{FB}))$ displays the supermodularity property in the total budget $(Y_s^M + t_s)$. Then either increasing the transfer ($t_s \geq t_s^{FB}$) or decreasing the ratio of private to public consumption ($m_s/c_s \leq m_s^{FB}/c_s^{FB}$) will only make the difference between the right and left hand sides of M 's incentive constraint larger. The remaining possibilities are: ($t_s < t_s^{FB}$ and $m_s/c_s = m_s^{FB}/c_s^{FB}$), or ($t_s = t_s^{FB}$ and $m_s/c_s > m_s^{FB}/c_s^{FB}$) or ($t_s < t_s^{FB}$ and $m_s/c_s > m_s^{FB}/c_s^{FB}$). In the first case, we will have $f_s > f_s^{FB}$, $m_s < m_s^{FB}$ and $c_s < c_s^{FB}$. In the second case, $f_s = f_s^{FB}$, $m_s > m_s^{FB}$ and $c_s < c_s^{FB}$. The third case is basically a mix of the two, with $f_s > f_s^{FB}$, $m_s > m_s^{FB}$ and $c_s < c_s^{FB}$. In the first case, M 's present period utility is lower than her utility in the first best and lower than her utility in the second case. Having both M 's IC and promise keeping constraint binding, F will choose the most economical way to satisfy them. This will depend on the promised utility and on further specifications of the utility functions. For example, for the case of additively separable functions, (z_c/u_c) is constant and equal to $(z_c/u_c)^{FB}$, which implies that $m_s/c_s > m_s^{FB}/c_s^{FB}$ and $f_s/c_s \geq f_s^{FB}/c_s^{FB}$. Finally, $W_s > W$ follows directly from the envelope condition.

2) The proof consists in showing that the difference for M between her first best utility plus continuation value on one hand, and her deviation utility plus minmax value on the other, in any state, is increasing in her value. More precisely, that

$$\begin{aligned} & \frac{\partial z}{\partial m_s^{FB}} \frac{\partial m_s^{FB}}{\partial \mu} + \frac{\partial z}{\partial c_s^{FB}} \frac{\partial c_s^{FB}}{\partial \mu} - \frac{\partial z^*}{\partial m_s^*} \frac{\partial m_s^*}{\partial \mu} - \frac{\partial z^*}{\partial c_s^*} \frac{\partial c_s^*}{\partial \mu} \\ & + \frac{\partial z_m}{\partial \mu} m_s^{FB} + \frac{\partial z_c}{\partial \mu} c_s^{FB} - \frac{\partial z_m^*}{\partial \mu} m_s^* - \frac{\partial z_c^*}{\partial \mu} c_s^* + \beta \frac{\partial W}{\partial \mu} > 0 \end{aligned}$$

First observation, $m_s^{FB} + c_s^{FB} = m_s^* + c_s^*$, so $\partial m_s^{FB} + \partial c_s^{FB} = \partial m_s^* + \partial c_s^*$. As established in proposition 2, the ratio m_s^{FB}/c_s^{FB} is increasing in W , whereas m_s^*/c_s^* is always constant. This means that $\frac{\partial m_s^{FB}}{\partial \mu} > \frac{\partial m_s^*}{\partial \mu}$ and $\frac{\partial c_s^{FB}}{\partial \mu} < \frac{\partial c_s^*}{\partial \mu}$. From the first order conditions for the first best allocation, $\frac{\partial z}{\partial m_s^{FB}} > \frac{\partial z^*}{\partial m_s^*} = \frac{\partial z^*}{\partial c_s^*} > \frac{\partial z}{\partial c_s^{FB}}$ ■

Long Run Equilibrium

The proofs here are adapted from Kocherlakota (1996) and can be found in the online appendix.

12 Appendix C: Efficient Equilibria Under Enforcement

Proposition 6

Proof. From F 's incentive constraint, it is clear that $W_s^{F'} \geq W_s^F$ for all s with equality only if $V_{enf} = \underline{V}$, where $W_s^F, W_s^{F'}$ are the threshold values for which F 's incentive constraint stops binding for state s . Hence $W'_{\max} \geq W_{\max}$. A symmetric argument holds for M , so $W'_{\min} > W_{\min}$. There are two cases: $W'_{\min} > W_{\max}$ and $W'_{\min} < W_{\max}$. In the first case, applying a similar argument as for proposition 4, for any long run equilibrium value of M in the interval $[W_{\min}, W_{\max}]$, agents will converge to the new long run equilibrium $(W'_{\min}, V(W'_{\min}))$. That $(m_s/c_s)' > (m_s/c_s)$ follows directly from part 1) of proposition 2.

In the second case, if the long run equilibrium value of M was in the interval $[W'_{\min}, W_{\max}]$, the new enforcement minmax values don't affect the equilibrium, and $(m_s/c_s)' = (m_s/c_s)$. Otherwise, agents again converge to $(W'_{\min}, V(W'_{\min}))$ and it follows in this case too that $(m_s/c_s)' > (m_s/c_s)$. ■

Proposition 7

Proof. The proof is straightforward. Divide the values W into four intervals: $[\underline{W}, W_{enf})$, $[W_{enf}, \mathbf{W}]$ and $(\mathbf{W}, \mathbf{W}']$, where \mathbf{W}' is the new highest implementable value for M . The old and new frontiers share the second interval only. For each $W \in [W_{enf}, \mathbf{W}]$ and s , the possible events are the following: No incentive constraint binded before and M 's constraint binds now, F 's constraint binded and no constraint or M 's constraint binds now, or no change occurred regarding binding constraints. In each of these events, $(m_s/c_s)' \geq (m_s/c_s)$. For each $W \in [\underline{W}, W_{enf})$ and state s , (m_s/c_s) on the old frontier is increasing with W and is less than (m_s/c_s) at W_{enf} . Finally, for each $W \in (\mathbf{W}, \mathbf{W}']$ and s , $(m_s/c_s)'$ on the new frontier is increasing with W and is greater than $(m_s/c_s)'$ at \mathbf{W} . Putting equal probability on being at any value on the frontier, the claim in the proposition follows directly. ■

Case where no SPFB Allocation Exists Before and After Enforcement: Illustrative Examples:

Example with two states:

In figure 5 M 's values lie on the X-axis, while the Y-axis depicts the two states. The solid blue line at the level of state 1 going from \underline{W} up to W_M means that M 's incentive constraint binds in state 1 for all values in that interval. Similarly, the solid red line at the level of state 2 going from W_F up to \mathbf{W} means that F 's incentive constraint binds in state 2 over those values. The fact that for every value there is at least some incentive constraint that binds means that there exists no SPFB allocation. The set of values over which the agents play in the long run is contained in the interval $[W_F, W_M]$. Any time state 2 occurs, F 's constraint binds and the continuation value is W_F . If state 1 occurs, M 's constraint binds and her continuation value is somewhere in that interval. After the introduction of perfect enforcement, the dashed lines represent the new values where agents' constraints bind. The new set of values over which the agents play in the long run is contained in the interval $[W'_F, W'_M]$. Here, it is clear that no matter what the probability distribution over the two states is, the average $(m/c)' > (m/c)$.

Example with three states:

Figures 6 and 7 depict two possible situations.

Figure 6 is different from the two state example in the sense that the long run intervals of values over which agents play $[W_{F3}, W_{M1}]$ and $[W'_{F3}, W'_{M1}]$ actually intersect (for simplicity, I take a case where F 's constraint never binds in states 1 and 2). However, it illustrates the fact that what matters is whether the switch of values happens over the same states. For instance, suppose the number of steps it takes to go from W_{F3} to W_{M1} is two. So that when M 's value is at W_{F3} and state 1 occurs, her continuation value is some W^1_{M1} and if state 1 occurs again, M 's continuation value jumps to W_M . Then the limiting probability distribution over these values is the following:

$$\Pr(W_{F3}) = \frac{\pi_3}{1 - \pi_2}, \quad \Pr(W^1_{M1}) = \frac{\pi_1\pi_3}{(1 - \pi_2)^2}, \quad \Pr(W_{M1}) = \frac{\pi_1^2\pi_3}{(1 - \pi_2)^3}$$




Figure 5: Two States

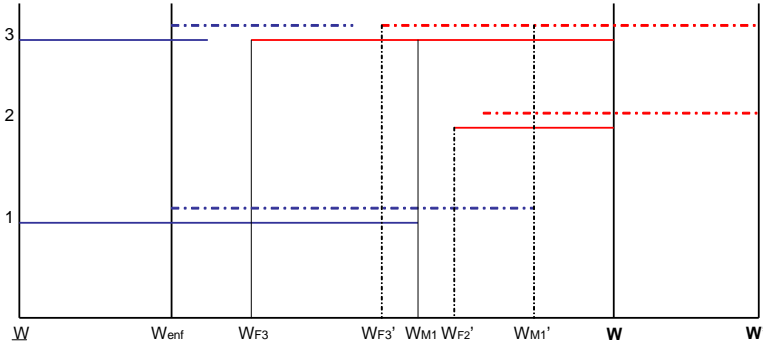


Figure 6: Three States (1)

Now suppose that it takes n steps to go from W'_{F3} to W'_{M1} . Then the limiting distribution over the values is

$$\Pr(W'_{F3}) = \frac{\pi_3}{1 - \pi_2}, \quad \Pr(W'_{M1}) = \frac{\pi_1 \pi_3}{(1 - \pi_2)^2}, \dots, \quad \Pr(W'_{M1}) = \frac{\pi_1^n \pi_3}{(1 - \pi_2)^{n+1}}$$

It remains true that the average $(m/c)' > (m/c)$.

Figure 7 is more complicated since not only intervals of values intersect, but also the states over which changes of values occur are different (for simplicity, M 's constraint in this case doesn't bind in states 1 and 2). So on the new frontier, when state 2 occurs, instead of staying the same, M 's value jumps to W'_{F2} . The problem comes from finding out (m_2/c_2) at W'_{F2} . As long as this ratio is not too high, it should not offset the difference between (m_1/c_1) on the interval $[W_{F3}, W_{M1}]$ and (m_3/c_3) at W_{F3} on one hand, and (m_1/c_1) on the interval $[W'_{F3}, W'_{M1}]$ and (m_3/c_3) at W'_{F3} on the other. In fact, since (m_2/c_2) at W'_{F2} is less than (m_3/c_3) at W'_{F3} , unless π_2 is very high, we can safely assume that it does not.




Figure 7: Three States (2)