NOTES ON INTRA-HOUSEHOLD ALLOCATION

ETHAN LIGON

Contents

1. Introduction 2
2. Static Deterministic Models 2
2.1. The Altruistic Dictator 2
2.2. The Nash-Bargaining Model 3
2.3. The "Collective" Model 3
2.4. Distinguishing Between Unitary and Collective 3
3. Dynamic Stochastic Models 5
3.1. Dynamic Bargaining in Households 6
3.2. Cooperative Dynamic Bargaining 10
4. Testing Dynamic Models of Intra-household Allocation 16
References 21

Date: September 22, 2009 Printed September 26, 2009.

1. Introduction

We began this part of this course by describing a model of a "farm-household", which we supposed aggregated the preferences of multiple people, and which made decisions regarding the management of the farm, and the allocation of goods within the household.

One of the main results from our discussion of the farm-household model was the "separation" result: given an adequate set of markets, the farm-household will act like a profit-maximizing firm on the production side, and like a utility maximizing consumer on the demand side, with no interdependence between the two beyond their interaction in the budget constraint.

We've seen that when we add agricultural risk, then the separation result can fail—when output is risky, this will in general lead to a distortion in both production and consumption decisions.

However, we've also seen that if households are able to *pool* their risk, we can retrieve a version of separation. Even when an individual household bears risk, if that risk is shared across the population, then household production decisions will depend only on prices, and household consumption decisions will depend only on prices and wealth, and *won't* depend on any idiosyncratic shocks experienced by the "farm" part of the farm-household operation.

We've discussed methods of testing whether or not households consumption decision really are affected by idiosyncratic production shocks or not. In many settings where these tests have been conducted, it looks as though households typically *do* bear some idiosyncratic risk, so we should expect the separation property to fail in these settings, though what risk-sharing does exist seems likely to mitigate the distortions in production and consumption associated with the failure of separation.

Here, we narrow our gaze somewhat. If it's the case that risk isn't properly shared *across* households, we're led to a relatively complicated model of allocation across these same households. But if people aren't able to effectively share risk across households, we might ask whether people are even able to share risk *within* households. To get at this question, we're led to develop some simple models of allocation within households.

2. Static Deterministic Models

2.1. **The Altruistic Dictator.** Becker (1974)

2.2. **The Nash-Bargaining Model.** McElroy (1990) describes restrictions on demands that are implied by a sort of "Nash-bargaining" model (Nash, Jr., 1950) of the household.

Thomas (1990) offers a test which is meant to distinguish between the altruism models and the Nash-bargaining model. He finds evidence that bargaining may influence allocation within households: in particular, for households in which mothers contribute a larger share of (non-labor) income, children seem to be better nourished.

2.3. **The "Collective" Model.** Chiappori (1992) points out that the real difference between the altruism and bargaining models is in the assumptions doesn't really have to do with preferences, but rather with outside options.

The basic reason is that *both* the Nash-bargaining and altruism models either assume (Nash-bargaining) or predict (altruism) that allocations within the household will be Pareto-optimal. The only real difference between the models in terms of their *predictions* has to do with the point on the Pareto frontier they identify. And since the predictions of at least one of these models depends on something which is very difficult to observe (altruism), empirically the important element common to both models is simply that intra-household allocations should be efficient.

2.4. **Distinguishing Between Unitary and Collective.** The formulation of the collective model has evolved over the last fifteen years. We'll skip the evolution, and look at the current state of affairs, as described by Bourguignon et al. (2008).

Let's assume that individuals have utility functions which are additively separable in consumption of the different household members (note that this does *not* rule out altruism). Further, the resources available to different household members may depend on a set of variables z, which (by definition) do *not* directly influence utility.

Then the problem facing a social planner trying to maximize the household's welfare can be written

$$\max \sum_{i=1}^{n} \lambda_i U_i(c_i)$$

subject to a constraint that

$$\sum_{i=1}^{n} c_i \leq F(z_1, \ldots, z_n).$$

Bourguignon et al. (2008) would say that this models a *unitary* household so long as the weights λ_i don't depend on the variables $\{z_i\}$ (as in the altruistic model of Becker, where the weights are assumed to reflect the unalterable preferences of the household head), while if the weights *do* depend on the "distribution factors" $\{z_i\}$, then this is a model of a *collective* household. Thus, if we could observe two households that differ only in terms of their distribution factors $\{z_i\}$, then we could distinguish between the unitary and collective models by measuring whether household-level demands were influenced by having different distribution factors or not.

Thomas (1990) was among the first to try and construct a test of the unitary model. To get around the problem of not being to observe otherwise-identical households with different distribution factors, he used a cross-sectional dataset of urban Brazilian households, with data on components of household income and measures of childrens' nutrition and anthropometrics. His specific hypothesis was that mothers and fathers might value nutritional investments in children differently. Thus, he tried to relate variation in child nutrition to variaton in male and female non-labor income. Under the collective model, we'd expect an increase in females' resources to be associated with a shift in household level demands which would reflect their improved bargaining position in the household, and indeed, Thomas finds that the effect of an increase in female non-labor income on children's protein intake is significantly greater than the effect of a similar increase in male non-labor income.

Thomas' results are intriguing, but don't constitute an airtight rejection of the unitary model. The chief problem is that, with just a cross-section of data one can't know when there's been a change in distribution factors. It's perfectly possible, for example, that much of the cross-sectional variation in non-labor income in Thomas' Brazilian data existed at the time households were formed, and before children were born, and that this variation in non-labor income is correlated with demand for well-nourished children. Such a correlation might arise from individual preferences (e.g., well-nourished children might be a normal good) or from the process of forming marriages (e.g., wealthier women might seek partners who care more about children than does the average Brazilian man).

The problem with Thomas' exercise is that while we can regard variation in non-labor income as capturing in part variation in distribution factors, he can't convincingly claim that this variation is uncorrelated with unobserved characteristics of households which might also influence demand for child nutrition. Unfortunately, we probably can't really hope to collect data on households which are identical save for their distribution factors and other observables (among other things, levels of altruism will presumably differ, and may be very difficult to observe).

However, we *can* hope to observe the *same* household at two different points in time, and it doesn't seem too forlorn a hope to think that important unobservable household characteristics (e.g., altruism) will be unchanged across periods. A number of papers (e.g., Lundberg et al., 1997; Bobonis, 2009) have exploited the idea that some *change* in distribution factors ought to result in a change in household demands under the collective model, but not the unitary model.

Using household panel data solves the serious problem of being able to identify *changes* in distribution factors in the data. However, this creates a second problem which can't be usefully addressed by the collective model: once one thinks of a change affecting household demands, then there's necessarily a *before* and *after*. But the standard account of the collective model is static, and can't offer a prediction about how a given household would deal with change. Further, it turns out that when one extends the collective or Nash-Bargaining models to a multi-period context, the predictions of these models can no longer be distinguished from those of the unitary model. But this is an issue we turn to in the next section.

3. Dynamic Stochastic Models

All of the models described above are static and deterministic; they all also either assume (Nash-Bargaining, Collective) or predict (Unitary) that allocations within the household will be efficient.

There have long been calls for the development of approaches to modeling the household which would permit predicted allocations to be *in*efficient (Lundberg and Pollak, 1994). But, as elsewhere in economics, saying that allocations are inefficient isn't very helpful if our aim is to understand how households work—while there's typically only a small set of efficient allocations, there's a much larger set of inefficient allocations, and we need some strategy to make reasonably sharp predictions about *what* inefficient outcomes we expect if we're to say anything interesting about the household.

One general strategy it's possible to pursue is to introduce some sort of "friction" to the operation of the household, and then to assume that allocations are efficient conditional on that friction. The most prominent examples of different kinds of frictions that one might introduce are private information within the household, or (and this is the friction we'll explore farther here) limitations on the abilities of the members of a household to commit themselves *ex ante* to a given course of action *ex post*.

This friction is only really interesting in multi-period problems. The idea is that one person can make promises regarding future behavior ("I'll gladly pay you tuesday...") in exchange for some action or payment in the current period ("...for a hamburger today"). In some circumstances those promises may be credible (e.g., if the two parties sign an enforceable contract backed up by some externally imposed threat of punishment, such as court-ordered fines, imprisonment, or other remedies for breach of the contract), but many of the agreements and promises that people make within households are outside the reach of the courts.

Even in the absence of an external authority to provide teeth to private contracts, private parties may be able to devise mechanisms which allow them to make efficient intertemporal agreements with one another (cf. the literature on "the Folk Theorem"). However, depending on the agents' patience and other aspects of their environment, these mechanisms can't generally deliver full efficiency. Instead, household members must be content with only imperfect solutions to the problem of "limited commitment" that they face.

3.1. **Dynamic Bargaining in Households.** We begin by describing a consumption-smoothing problem, which illustrates some of the issues involved in thinking about dynamic bargaining. Let $u_i(c)$ denote the utility from private consumption enjoyed by individual i at time t. We assume that every u_i is increasing and strictly concave. For the purposes of this example, we assume two individuals ($i = 1, 2$) and two time periods ($t = 1, 2$). Individuals 1 and 2 are 'married' at the beginning of the first period, and in each period each receive some endowment of the consumption good y_{it}. Should the two divorce, each receives some alienation utility $a_i(y_{it})$, which we also take to be increasing and strictly concave functions.

For the purposes of this illustration, we consider a particularly simple set of possible endowments. There are two possible states of the world, denoted by $s \in \{1, 2\}$. In state one, agent one receives an endowment of two, while agent two receives an endowment of zero. In state two, their positions are reversed; that is, agent one's endowment is zero, while agent two's endowment is two. Furthermore, these states simply alternate; in odd periods, $s = 1$, while in even periods $s = 2$.

Faced with these endowment processes, what might we expect a married couple to do? Both have concave utility functions, and *ex ante* would benefit from some sort of smoothing of consumptions over time. Thus, if both agents can commit themselves in advance (and it's natural to suppose that marriage is all about commitment), we might expect an egalitarian sharing arrangement to prevail. For example, consider a married couple, each of whom receives a bi-weekly paycheck, but on alternating weeks; here an egalitarian solution seems quite plausible. Modifying the example somewhat, if the length of a period is a decade rather than a fortnight, then the plausibility of an invariant, egalitarian split is rather lessened.

Taking this simple example as a sort of benchmark, we next compare each of several approaches to predicting the allocation of private consumption. We consider, in turn, a unitary household model, and a sort of repeated cooperative bargaining model, before turning our attention to a dynamic cooperative bargaining model.

3.1.1. *Unitary Household.* In the household model of Becker (1974), links between household members are sustained by altruism. For example, in addition to the utility that agent one derives from private consumption, suppose that he receives $\theta u_2(c_{2t})$ utils from the consumption of agent two. If we allow agent one to control the allocation of consumption within the household, then in each period he will solve

$$\max_{c_{1t},c_{2t}} u_1(c_{1t}) + \theta u_2(c_{2t})$$

subject to the household resource constraint $c_{1t} + c_{2t} \leq y_{1t} + y_{2t}$ in each of periods one and two (note that, for simplicity, no savings are allowed).[1] The first order conditions from this problem imply that the ratio of marginal utilities of the two agents will be a constant, regardless of the income realization or the period, since

$$\frac{u_1'(c_{1t})}{u_2'(c_{2t})} = \theta. \tag{1}$$

This equation determines a sharing rule. Because θ is taken to be a constant determined by one's feelings toward two, this ratio of marginal utilities must be a constant, and since the functions u_i are themselves time-invariant, then the sharing rule must itself be time invariant. Each agent's share of consumption depends on the preference parameter θ; the more one 'cares' about two (the higher the value of θ), the more consumption allocated to two. Because θ is a preference parameter which is presumably difficult to observe, the important empirical consequences of the unitary model are the Pareto optimality of allocations and the invariance of the sharing rule.

[1]See Ligon et al. (2000) for a treatment of individual savings in a related model.

Notice that we've separated agent one's problem into two distinct subproblems, one for each period. For this particular problem, this separation of a dynamic allocation problem into a sequence of static problems doesn't affect the solution. This separability is due to the lack of an intertemporal technology such as savings, and due also to the invariance of the sharing rule implied by (1).

Also notice that any change in the distribution of income between agents one and two (or who receives the paycheck in a given period, per our example) will make no difference in the allocation of consumption so long as the total resources available to the household are unchanged, and so long as agent one remains in charge of the allocation decision.

3.1.2. *Cooperative Bargaining.* Following the analysis of, e.g. McElroy and Horney (1981), we consider the consequences of a change in the intra-household distribution of income in a Nash bargaining model. Nash's (1950) approach was to specify a set of axioms which any 'reasonable' solution ought to satisfy. Nash's set of axioms included a requirement of symmetry, and a requirement of Pareto optimality. Subsequent authors have investigated the consequences of modifying these axioms, in ways which we'll exploit. In particular, Harsanyi and Selten (1972) abandon Nash's symmetry axiom. With this modification, we're left with the following axioms, as applied to our simple two period example (note that feasibility is assumed).

- **INV:** Consumption allocations are *invariant* to affine transformations of agents' utilities.
- **IIA:** Consumption allocations are *independent of irrelevant alternatives*; that is, if the allocation $\{(c_{1t}, c_{2t})\}$ solves one bargaining problem when aggregate incomes are (y_1, y_2) and the same allocation is feasible in a second (otherwise identical) problem with aggregate incomes $(\hat{y}_1, \hat{y}_2) < (y_1, y_2)$, then the allocation $\{(c_{1t}, c_{2t})\}$ also solves the second problem.
- **PO:** The allocation is Pareto optimal.

Nash offers a simple algorithm to compute the allocations satisfying the axioms. McElroy-Horney applies this algorithm to the intrahousehold allocation problem by solving:

$$\max_{c_{1t}, c_{2t}} [u_1(c_{1t}) - a_1(y_{1t})]^{\alpha} [u_2(c_{2t}) - a_2(y_{2t})]^{1-\alpha}$$

with $0 < \alpha < 1$ for $i = 1, 2$, subject to the household budget constraint $c_{11} + c_{21} \leq y_{11} + y_{21}$. Then the first order conditions of this problem once again give us a set of restrictions on the ratio of marginal utilities of consumption,

$$\frac{u_1'(c_{1t})}{u_2'(c_{2t})} = \frac{1-\alpha}{\alpha} \frac{u_1(c_{1t}) - a_1(y_{1t})}{u_2(c_{2t}) - a_2(y_{2t})}. \tag{2}$$

Comparing this condition with (1), notice that although the left hand side of the two equations are identical, the right hand side of (2) is not time invariant; in particular, it will depend on the incomes realized. If, for example, there is an increase in y_{1t}, then this will result in a decrease in the marginal utility of agent one relative to agent two, and hence an increase in relative consumption. In our alternating paycheck example, each agent would take turns consuming a larger share of the total endowment.

This dependence of allocation on the *ex post* distribution of income is the feature of this model which has made it attractive to people seeking an interesting theory of intra-household allocation. However, as a consistent theory, this approach has a serious flaw; the solution to a sequence of static problems does not at all correspond to the solution to the dynamic problem if agents are forward-looking. To see this, suppose that agent i's time one utility if married is given by the time separable function $u_i(c_{i1}) + \beta u_i(c_{i2})$, while utility if divorced (in the first period) is $a_i(y_{i1}) + \beta a_i(y_{i2})$. Then the appropriate solution to the Nash bargaining problem is given by the solution to

$$\max_{c_{11},c_{12},c_{21},c_{22}} [u_1(c_{11}) + \beta u_1(c_{12}) - a_1(y_{11}) - \beta a_1(y_{12})]^{\alpha} \cdot [u_2(c_{21}) + \beta u_2(c_{22}) - a_2(y_{21}) - \beta a_2(y_{22})]^{1-\alpha} \tag{3}$$

with $0 < \alpha < 1$, and subject to the constraints $c_{1t} + c_{2t} \leq y_{1t} + y_{2t}$ for $t = 1, 2$. The first order conditions from this problem once again give us an expression for the ratio of marginal utilities between agents one and two, given this time by

$$\frac{u_1'(c_{1t})}{u_2'(c_{2t})} = \frac{1-\alpha}{\alpha} \frac{u_1(c_{11}) + \beta u_1(c_{12}) - a_1(y_{11}) - \beta a_1(y_{12})}{u_2(c_{21}) + \beta u_2(c_{22}) - a_2(y_{11}) - \beta a_2(y_{22})}, \tag{4}$$

for $t = 1, 2$. This equation once more defines a sharing rule for consumption. The remarkable thing about this equation is that the implied sharing rule is, once again, time-invariant. The distribution of consumption in each period depends on the entire sequence of income realizations of both agents.

How should we interpret this invariance? It is as if the two agents, each knowing their income sequence, engaged in some borrowing-lending arrangement between them. In our alternating paycheck example, though agent two has no income in the first period, she will receive the entire endowment in the second. Knowing this, the first agent is perfectly happy to enter into an egalitarian sharing relationship in each period.

This example also makes it clear that the sharing rule implied by (2) is not in fact the solution to the Nash bargaining problem if agents are forward looking. The reason for this is that the sharing rule of (2) is *not* Pareto optimal; a move to the sharing rule implied by (4) makes each forward looking agent better off. Thus this sort of solution violates one of the key axioms underlying the Nash bargaining approach—the axiom of Pareto optimality.

Although the violation of Pareto optimality loses us the interpretation of solutions to (2) as Nash bargaining solutions, one might well argue that the abandonment of Pareto optimality represents a gain terms of realism (cf. Lundberg and Pollak (1996)). In the next section we offer an alternative axiomatic development of a dynamic cooperative bargaining model in which cooperative outcomes are not generally Pareto optimal.

3.2. **Cooperative Dynamic Bargaining.** This section describes an adaptation of the static Nash bargaining problem to a dynamic setting, and is drawn from Ligon (2002), which uses this framework to construct examples of how we might expect the introduction of Grameen-Bank style micro-lending schemes to affect intra-household allocation. Mazzocco (2007) has subsequently used a similar model framework to examine intra-household allocations in the U.S.; we'll discuss his approach to estimation and testing below.

The appeal of a sequence of static Nash bargaining problems as a description of intra-household allocation seems to be due to the idea that binding legal contracts seldom constrain the members of a household with respect to their actions toward one another—within a household, the ability to *commit* to a sharing rule of the form suggested by (4) is limited. For example, we might think of a marriage in which the prospects of the wife unexpectedly turn out to be much better than those of her husband. After this state of affairs is realized, it seems likely that the wife might thenceforth claim a larger share of household resources.

This section develops a dynamic bargaining problem in which household members strike a bargain over the intra-household allocation of resources. A solution to the bargaining problem takes the form of an agreement regarding a sharing rule to use after any history. Realized allocations are *ex post* Pareto optimal, in the sense that once the state is known and the consumption good allocated, there is no alternative

Pareto-improving allocation of resources in that state (this is similar to the "collective efficiency" of Bourguignon and Chiappori (1992)). However, these allocations are *not* generally Pareto optimal *ex ante*, because of the lack of commitment available to household members. In a state of the world in which agent one realizes a very high endowment and agent two a very low endowment, agent one can, in effect, demand a renegotiation of the sharing rule in force from previous periods.

The bargaining environment has two infinitely-lived agents, each with a time-separable von Neumann-Morgenstern utility function. Each agent discounts future utility at some positive rate $1/\beta - 1$. In any given period, one of a finite number of possible states $S = \{1, 2, \ldots, m\}$ occurs. If the current state is $s \in S$, then the probability that the state in the subsequent period is $r \in S$ is some number π_{sr}. For the sake of simplicity, we assume that $\pi_{sr} > 0$ for all s and r in S.

In this environment, bargaining involves deciding how to assign surpluses both now and in the future. Denote the set of possible (momentary) surpluses in state s by $W_s \in \mathbb{R}^2$, and let $\mathbf{W} = \{W_s\}_{s \in S}$. We assume that each W_s is convex, compact, and contains the origin. Let $\rho_s(w_1) = \max\{w_2 | (w_1, w_2) \in W_s\}$ trace out the Pareto frontier of the set W_s; we assume that each function ρ_s is monotone decreasing and strictly concave. Let $\mathcal{W}$ denote the set of $\mathbf{W}$ having surplus possibilities which satisfy these conditions.

A bargaining *problem* is completely described by the current state s and the set of surplus possibilities $\mathbf{W}$, so that a problem is simply some $(s, \mathbf{W})$. Thus, $\mathcal{B} = \{(s, \mathbf{W}) \in S \times \mathcal{W}\}$ denotes the set of bargaining problems we explore in this paper.

Any solution to the bargaining problem must assign surpluses in every possible sequence of states. Borrowing an idea from the literature on dynamic contracting (Green, 1987; Spear and Srivastava, 1987), we use surpluses in the previous period to summarize histories. Because of the Markovian nature of the environment, a solution to the bargaining problem $(s, \mathbf{W})$ is some initial assignment of surplus $w \in W_s$ and a collection of mappings $\Phi(\mathbf{W}) = \{\phi_{sr}\}_{(s,r) \in S \times S}$, with $\phi_{sr} : W_s \to W_r$. Thus, the collection $\Phi(\mathbf{W})$ defines a law of motion for surpluses, and thus a rule for assigning surpluses in every possible sequence of states. We write a solution to the problem $(s, \mathbf{W})$ as a pair (w, Φ); the pair

w is the assignment of surpluses in the initial state s,[2] while the collection Φ gives rules for assigning all future surpluses, given the initial assignment.[3]

To make this notation clear, we describe our earlier alternating paycheck example using this new notation. First, the possible states are simply $S = \{1, 2\}$, with probabilities $(\pi_{11}, \pi_{12}, \pi_{21}, \pi_{22})$ equal to $(0, 1, 1, 0)$, to capture the strict alternation of paychecks.[4] Momentary surpluses are given by $W_1 = \{(u_1(c) - a_1(2), u_2(2-c) - a_2(0)) | c \in [0, 2]\}$ and $W_2 = \{(u_1(c) - a_1(0), u_2(2-c) - a_2(2)) | c \in [0, 2]\}$. Notice that the primitive objects of the bargaining problem (the sets of surpluses) are here described in terms of momentary utilities associated with marriage and divorce.

Any solution (w, Φ) which assigns a sequence of surpluses also induces *values* $\{U_s(w)\}_{s \in S}$, which solve the equations

$$U_s(w) = w + \beta \sum_{r \in S} \pi_{sr} U_r(\phi_{sr}(w)),$$

for all $s \in S$. We interpret $U_s(w)$ as the values each agent derives from continuing the bargaining relationship, given that the current state is s and the current surplus assignment is w.[5]

We next adapt Nash's axioms to the stationary dynamic environment described here. First is the axiom of invariance (INV). The only real difference from our earlier formulation is that we now work in a space of surpluses, rather than consumption allocations.

---

[2]Since the initial w may be chosen arbitrarily, the solutions given here are not generally unique. As in Nash, Jr. (1950), it would be a simple matter to also require that a solution to the bargaining problem satisfy some sort of *symmetry* which would guarantee uniqueness, though in the present case this requirement of symmetry would apply only to the *ex ante* discounted expected surpluses.

[3]We will often be working with pairs of real numbers. If a function $f(x)$ maps x into $\mathbb{R}^2$, then let $(f^1(x), f^2(x)) = f(x)$ denote the coordinates of the mapping. For any two real pairs (a, b) and (c, d), the strict inequality $(a, b) > (c, d)$ should be read as $a > c$ and $b > d$, while the weak inequality $(a, b) \geq (c, d)$ should be read as $a \geq c$ and $b \geq d$.

[4]This violates our assumption that all probabilities ought to be strictly positive, but the violation does no harm here.

[5]In this stochastic, dynamic environment one might think it to be more natural to define the bargaining problem in terms of these total values, rather than in terms of momentary surpluses. The problem with this is that, as our construction of these values should make clear, the set of values (discounted expected surpluses) is co-determined with the solution to the bargaining problem; the set of achievable values isn't known in advance.

Let $\mathbf{W} = \{W_s\}_{s \in S} \in \mathcal{W}$, and define the operation of scalar multiplication on $\mathcal{W}$ by $\theta\mathbf{W} = \{\theta W_s\}_{s \in S}$ for any scalar θ.

INV: For all $\mathbf{W} \in \mathcal{W}$ and all $\theta > 0$, $\Phi(\theta\mathbf{W}) = \theta\Phi(\mathbf{W})$.

Since utilities are von Neumann-Morgenstern, surpluses are invariant only up to a positive affine transformation. Accordingly, any solution should be similarly invariant.

Next is independence of irrelevant alternatives. For the purposes of stating the axiom, let $\phi_{sr}(\cdot)$ denote part of the solution (w, Φ) to the problem $(s, \mathbf{W})$.

IIA: For any $\mathbf{W}$ and $\hat{\mathbf{W}}$ such that $\phi_{sr}(w) \in \hat{W}_r \subseteq W_r$ for all $w \in \hat{W}_s$ and for all $(s, r) \in S \times S$, $\Phi(\hat{\mathbf{W}}) = \Phi(\mathbf{W})$.

The idea here is simply that if (w, Φ) is the solution to one problem, then if feasible it should also be the solution to a more constrained problem.

Both the INV and IIA conditions are essentially the same as conditions advanced by Nash. In particular, we can think of any of the values $U_s(w)$ as static bargaining surpluses. If we then work with the set of possible values, we can apply Nash's invariance and independence axioms directly to any of these sets. His axioms are then easily seen to be equivalent to the versions of INV and IIA given here.[6]

As discussed above, the axiom of Nash's which is problematic from our point of view is Pareto optimality. This implies that agents can always commit *ex ante* to never renegotiate *ex post*, which seems too strong. We replace Nash's axiom with two alternative requirements. The first of these is a notion of *individual rationality*:

IR: For any problem $(s, \mathbf{W})$, $U_r(\phi_{sr}(w)) \geq 0$ for all $r \in S$; also, if there exists a state $q \in S$ and $\hat{w} > 0$ for some $\hat{w} \in W_q$, then for any $r \in S$ $U_r^i(\phi_{sr}(w)) > 0$ for some $i \in \{1, 2\}$.

This requirement has two parts. First is the idea that neither agent can be made strictly worse off by continuing the bargaining relationship than by discontinuing it. Second, if positive surpluses for either agent are possible in some state, then any solution to the bargaining problem must have the property that at least some of this surplus is enjoyed by at least one of the agents.

[6] Of course, the IIA axiom has proven to be controversial in the static bargaining literature; however, as we've already assumed that agents are expected utility maximizers (the usual axiomatic development of this framework relies on an IIA type of axiom), our use of IIA here should be unobjectionable.

Though IR is implied by Pareto optimality, Roth (1977) shows that in the static case ($\beta = 0$) INV, IIA, and IR imply Pareto optimality.[7] This does not carry over to the dynamic problem; in particular, an *ex ante* Pareto optimal solution may not satisfy IR. We recover Nash's sentiment that agents ought to always be able to negotiate until they reach an efficient outcome by using a notion of *constrained Pareto optimality*:

CPO: Let $\{U_s\}$ be induced by the solution (w, Φ), and let $\{\hat{U}_s\}$ be induced by $(w, \hat{\Phi})$. If $\hat{U}_s(w) \geq U_s(w)$ with $\hat{U}_s^i(w) > U_s^i(w)$ for some $s \in S$ and some $i \in \{1, 2\}$, then $(w, \hat{\Phi})$ fails to satisfy at least one of INV, IIA, and IR.

The idea here is simple; we simply require that the solution be optimal within the class of surplus assignments satisfying the other three axioms.

We next propose an algorithm for computing a solution satisfying INV, IIA, IR, and CPO. Given a problem $(s, \mathbf{W})$, we first define a set of functions for assigning surpluses in the current period. Let $C_s : [0, 1] \to W_s$ for $s \in S$, with

$$C_s(\lambda) = \operatorname*{argmax}_{w \in W_s} \lambda w_1 + (1 - \lambda) w_2. \tag{5}$$

Thus, in any state s, $C_s(\lambda)$ tells us how to maximize a sum of current period surpluses weighted by λ. Since W_s is compact, convex, and has a strictly concave frontier, $C_s(\lambda)$ is guaranteed to be single-valued for any $\lambda \in [0, 1]$. $C_s(\lambda)$ can be regarded simply as a solution to a static bargaining problem. Note that when $W_s \subset \mathbb{R}_+^2$, this function also gives the solution to the asymmetric static bargaining problem (seen in the previous illustration):

$$N_s(\alpha) = \operatorname*{argmax}_{w \in W_s} w_1^{\alpha} w_2^{1-\alpha}. \tag{6}$$

Because we're interested in bargaining problems in which one agent may make a sacrifice in the current period (so that, e.g., w_1 may be negative) in exchange for higher future surpluses, it's necessary to work with C_s, which is well defined for any W_s.

[7]Our formulation of the IR requirement is slightly weaker than Roth's, who requires (in our notation) $U_r(\phi_{sr}(w)) > 0$. Thus, Roth's version of IR is not implied by Pareto optimality. Nonetheless, our weaker version of IR along with INV and IIA still implies Pareto optimality in the static case.

Next, we define two additional families of functions, designed to let us use surplus *weights* rather than surpluses as state variables.[8] The first, $\psi_s : [0,1] \to [0,1]$, maps the surplus weight λ in one state into a new weight given that the subsequent state is s. The second family, $V_s : [0,1] \to \mathbb{R}^2$, maps the current surplus weight into values. In particular, the $\{(\psi_s, V_s)\}_{s\in S}$ solve the set of functional equations

$$V_s(\lambda) = C_s(\lambda) + \beta \sum_{r\in S} \pi_{sr} V_r(\psi_r(\lambda)) \tag{7}$$

for $s \in S$, and

$$\psi_r(\lambda) = \begin{cases} \underline{\lambda}_r & \text{if } V_r^1(\lambda) < 0 \\ \overline{\lambda}_r & \text{if } V_r^2(\lambda) < 0 \\ \lambda & \text{otherwise} \end{cases} \tag{8}$$

where $\underline{\lambda}_r$ satisfies $V_r^1(\underline{\lambda}_r) = 0$, and $\overline{\lambda}_r$ satisfies $V_r^2(\overline{\lambda}_r) = 0$. The assumption that the Pareto frontier of each W_s is monotone decreasing along with the assumption that W_s includes the origin ensures that the $\{(\underline{\lambda}_r, \overline{\lambda}_r)\}$ exist, and so a solution to (7) and (8) exists.[9]

Very simply, the proposed class of solutions has the following character. Agents will efficiently divide any momentary surplus according to an invariant sharing rule until they reach a point such that continuing to use this rule would make one of the agents worse off than she would be on her own. At this point, the division of surplus will change so as to make this agent just indifferent between remaining in the relationship and leaving. Then this new sharing rule will remain invariant until it, in its turn, once more delivers too small a value to induce one of the agents to remain.

Although at this level of generality our specification of the bargaining problem may suggest that its solution is complicated, in fact there's an easy intuition to describe the solution. The agents begin by agreeing to some initial division of momentary surplus, characterized by the sharing parameter λ. As history unfolds, the two agents stick to the sharing rule implied by λ until a state (say r) is realized such that one of the agents (say agent one) would be better off leaving the relationship rather than continuing with the same share of surplus. At that point, both agents

[8]A number of earlier authors have used this same trick of using utility or surplus weights in models of efficient contracts, including Hayashi (1996) and Marcet and Marimon (1999).

[9]More precisely, given the stated restrictions on the $\{W_s\}$, (8) and (7) together define an isotone mapping from a subset of the set of bounded functions on $\mathbb{R}^2$; existence of a fixed point to this mapping is then guaranteed by a theorem of Knaster-Tarkski (Moore, 1985).

find it in their best interest to renegotiate the share parameter in agent one's favor, to the point where agent one is just indifferent between leaving and staying: this will be the number $\underline{\lambda}_r$. The two agents then continue as before, sharing according to the last value of λ negotiated, until they again reach a state in which one or the other would be better off by terminating the relationship.

4. Testing Dynamic Models of Intra-household Allocation

Suppose, as is standard, that the period-by-period utility possibility set for the members of a household has a concave frontier, that household members value both present and future utility using a time-separable von Neumann-Morgenstern utility function, and agree on the probability of different future events. Then when intra-household consumption allocations are efficient, as assumed by the Nash bargaining model and the static "collective" model, or as predicted by a Becker-style "unitary" model, the rule governing allocations will generally eliminate any idiosyncratic variation—household members will share any future variation in resources, whether "for richer or for poorer, in sickness as in health." For this reason, the predictions of dynamic models with efficient allocations aren't very interesting when it comes to intra-household allocation.

In the previous section, we've sketched a particular model incorporating a sort of friction (limited commitment) which will prevent households from implementing a fully *ex ante* efficient intra-household allocation rule (provided that they're not too patient). This model yields interesting predictions regarding dynamics—in particular, when one person's bargaining position improves, that can actually influence subsequent allocations.

Mazzocco (2007) describes a similar model, and devises a way to test it using only household level data from a panel. A critical maintained hypothesis for Mazzocco is that individual utility functions permit (Gorman) aggregation—that is, that they belong to what Mazzocco terms the "ISHARA" (Identically Shaped Hyperbolic Absolute Risk Aversion) class. Given this maintained hypothesis, his test is basically designed to distinguish between dynamic models in which changes in bargaining position *do not* result in changes in allocations (as in the efficient, "unitary" models), and dynamic models in which such changes in bargaining position *do* produce changes in allocations. The dynamic limited commitment model described above and in Mazzocco (2007) is an example of the latter sort of model, but there may be many other models with frictions in which changes in bargaining result in changed allocations, and Mazzocco's test will not distinguish among models within this class.

Let us illustrate Mazzocco's idea with a simple example. Suppose that there's a household with two people, indexed by $i = 1, 2$. There are exactly two periods, indexed by $t = 1, 2$. Each person cares only about utility derived from consumption in each of the two periods, and discounts future utility by a factor $\beta \in (0, 1)$. Each person i has a "weight" related to their bargaining position of λ_{it} in period t. In either period t the household has a "pooled" income of y_t, and (collectively) has access to credit markets, and can borrow or lend at an interest rate of r.

The problem facing the household is to find a way to most efficiently allocate its resources across both time and people. To do this, it solves

$$\max_{\{c_{it}\},b} \lambda_{11}u_1(c_{11}) + \beta\lambda_{12}u_1(c_{12}) + \lambda_{21}u_2(c_{21}) + \beta\lambda_{22}u_2(c_{22}) \tag{9}$$

subject to period-by-period resource constraints

$$c_{11} + c_{21} \leq y_1 + b \tag{10}$$

and

$$c_{12} + c_{22} \leq y_2 - b(1 + r), \tag{11}$$

where b is the amount of period 1 borrowing (or saving, if negative).

Associating the Lagrange multiplier μ_1 with the first-period resource constraint (10), and μ_2 with the second-period resource constraint (11), we obtain the following first order conditions from this problem:

$$\lambda_{11}u_1'(c_{11}) = \mu_1, \tag{12}$$

$$\lambda_{21}u_2'(c_{21}) = \mu_1, \tag{13}$$

$$\lambda_{12}u_1'(c_{12}) = \mu_2/\beta, \tag{14}$$

$$\lambda_{22}u_2'(c_{22}) = \mu_2/\beta; \text{and} \tag{15}$$

$$\mu_1 = (1 + r)\mu_2. \tag{16}$$

Let us take the last first order condition first: in this setting, the multipliers μ_t can be interpreted as the *household*'s marginal utility of income in either period, and thus (16) can simply be regarded as the household's collective Euler equation, showing how the household will choose to allocate resources over time to equate its collective intertemporal marginal rate of substitution (μ_1/μ_2) to the intertemporal marginal rate of transformation $(1 + r)$.

Now, dividing either (12) by (14) or (13) by (15) yields a similar sort of Euler equation, but instead of holding for the household collectively, this is for each individual separately:

$$\frac{\lambda_{i1}}{\lambda_{i2}}\frac{u_i(c_{i1})}{u_i(c_{i2})} = \beta(1+r). \tag{17}$$

Consider the following. If

1. The weights are constant (i.e., $\lambda_{i1} = \lambda_{i2}$), then the individual Euler equation (17) takes the "usual" form
$$u_i(c_{i1}) = \beta(1+r)u_i(c_{i2}).$$
2. If, in addition, $u_i(c) = \log c$ (satisfying Mazzocco's ISHARA assumption), then

$$c_{i2} = \beta(1+r)c_{i1}. \tag{18}$$

Our purpose in assuming logarithmic utility here is simply to make consumptions aggregable across household members. We now illustrate this aggregability by summing equation (18) over individuals:

$$c_{12} + c_{22} = \boxed{c_2 = \beta(1+r)c_1} = \beta(1+r)(c_{11} + c_{21}).$$

Here, the boxed expression corresponds to the collective Euler equation (16). Thus, when weights are unchanging and utilities are Gorman aggregable, the household's intertemporal allocation of resources is the same as it would be if the problem was solved by a single representative agent.

Now let us imagine that intra-household allocations aren't efficient, and don't satisfy the unitary household assumption. In particular, let us simply suppose that between periods one and two, person two experiences an increase in her bargaining weight, so that $\lambda_{21} < \lambda_{22}$. This is simply a sort of 'reduced-form' way of getting at the intra-household problem, for we haven't given any reason that we might expect such a change. The model outlined above in Section 3.2 could be used to understand the deeper connections between economic resources, outside options, and the increase in λ_{22}, but we make no attempt here to achieve such an understanding.

Instead, in something of the spirit of the static collective model, we simply suppose that something (e.g., one of the "distribution factors" of Bourguignon et al. (2008) has changed in the environment, and this has had the consequence of increasing λ_{22}. Now, for person two this has the consequence of changing how she'd like to time her consumption. In particular, from (18) we know that

$$c_{22} > \frac{\lambda_{21}}{\lambda_{22}}c_{22} = \beta(1+r)c_{21},$$

so that two's consumption in period two (relative to period one) will is larger than it would be in the unitary case—this is simply a consequence of her claim to a larger share of household resources in the second period.

In contrast, one's Euler equation is unchanged. We can then do the same sort of summing over individuals that we did above:

$$c_{12} + c_{22} = \boxed{c_2 > \beta(1+r)c_1} = \beta(1+r)(c_{11} + c_{21}). \tag{19}$$

So we see that when we aggregate up to household-level demands, changes in bargaining weights will lead to what appear to be violations of the collective Euler equation. Appearances are deceiving, of course—the *real* collective Euler equation (16) isn't violated. It's just that the relationship between the multipliers μ_t and individual consumption has changed.

This may seem surprising—since two gets more consumption in period two, shouldn't one get less? And in fact one *does* get less in period two. But because he *anticipates* getting less in period two, he also consumes less in the first period, saving more so as to smooth his consumption over time.

In the context of our little reduced-form exercise, one might well suppose that apparent violations of the collective Euler equation could stem from *decreases* in bargaining weights, just as they could from increases. Such a decrease would reverse the inequality in (19), and suggest that non-unitary households might have aggregate consumption profiles which fell over time. However, this supposition is less evidence that aggregate household consumption might tend to fall than it is evidence of the shortcomings of the reduced-form approach we've taken here. To see this, let us go just a little bit deeper into thinking about how the bargaining process might actually work.

Suppose that everything is as in the previous example, save that the weights are constant over time, as in the unitary model. Changes in bargaining position may still influence allocations, but we'll be more explicit as to the mechanism, as in Section 3.2. In this baby example, suppose that person two's outside option is a_{22}. Then the household solves:

$$\max_{\{c_{it}\},b} \lambda_1[u_1(c_{11}) + \beta u_1(c_{12})] + \lambda_2[u_2(c_{21}) + \beta\lambda_{22}u_2(c_{22})] \tag{20}$$

subject to the same period-by-period resource constraints (10) and (11) as previously. Solving this would just give an efficient unitary solution, with a fixed sharing rule. But here the household also has to solve another constraint: to keep person two from leaving in period two, the allocation must satisfy

$$u_2(c_{22}) \geq a_{22}. \tag{21}$$

Now, the first order conditions from this problem are

$$\lambda_1 u_1'(c_{11}) = \mu_1, \tag{22}$$
$$\lambda_2 u_2'(c_{21}) = \mu_1, \tag{23}$$
$$\lambda_1 u_1'(c_{12}) = \mu_2/\beta, \tag{24}$$
$$(\lambda_2 + \theta) u_2'(c_{22}) = \mu_2/\beta; \text{and} \tag{25}$$
$$\mu_1 = (1+r)\mu_2, \tag{26}$$

where θ is the Lagrange multiplier associated with the "don't leave me" constraint (21). Comparing these with the first order conditions (12)–(16), there are just two differences. The first is simply that the λs are now fixed, and don't vary over time. The second is that there's now the additional multiplier θ.

This more "primitive" specification of the household's problem reveals something important about the way in which distribution factors can effect weights. Comparing the reduced-form weight λ_{22} to the current formulation of the problem, we can see that $\lambda_{22} = \lambda_2 + \theta_2$. And since $\theta_2 \geq 0$ by the Karush-Kuhn-Tucker theorem, it follows that if the reduced-form weights are determined by a process along the lines of the present example, then it must be that weights can only ever *increase*, never *decrease*.

References

Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy 82*, 1063–1094.

Bobonis, G. (2009). Is the allocation of resources within the household efficient? new evidence from a randomized experiment. *Journal of Political Economy 117*(3), 453–503.

Bourguignon, F., M. Browning, and P. A. Chiappori (2008). Efficient intra-household allocations and distribution factors: Implications and identification. *Review of Economic Studies Forthcoming*.

Bourguignon, F. and P.-A. Chiappori (1992). Collective models of household behavior. *European Economic Review 36*, 355–364.

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy 100*(3), 437–67.

Green, E. J. (1987). Lending and the smoothing of uninsurable income. In E. C. Prescott and N. Wallace (Eds.), *Contractual Arrangements for Intertemporal Trade*. Minneapolis: University of Minnesota Press.

Harsanyi, J. C. and R. Selten (1972). A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science 18*, 80–106.

Hayashi, F. (1996). Analysis of household saving: Past, present, and future. *The Japanese Economic Review 47*(1), 21–33.

Ligon, E. (2002). Dynamic bargaining in households (with an application to Bangladesh). Giannini Foundation Working Paper.

Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics 3*(2), 216–246.

Lundberg, S. and R. A. Pollak (1994). Noncooperative bargaining models of marriage. *American Economic Review 84*(2), 132–137.

Lundberg, S. and R. A. Pollak (1996). Bargaining and distribution in marriage. *Journal of Economic Perspectives 10*(4), 139–158.

Lundberg, S. J., R. A. Pollak, and T. J. Wales (1997). Do husbands and wives pool their resources? *The Journal of Human Resources XXXII*(3), 462–480.

Marcet, A. and R. Marimon (1999). Recursive contracts. Manuscript.

Mazzocco, M. (2007). Household inter-temporal behavior: A collective characterization and a test of commitment. *Review of Economic Studies 74*(3), 857–895.

McElroy, M. (1990). The empirical content of Nash-bargained household behavior. *Journal of Human Resources 25*, 559–583.

McElroy, M. B. and M. J. Horney (1981). Nash-bargained household decisions: Toward a generalization of the theory of demand. *International Economic Review 22*(2), 333–350.

Moore, R. E. (1985). *Computational Functional Analysis.* Ellis Horwood Series in Mathematics and Its Applications. New York: Halsted Press.

Nash, Jr., J. F. (1950). The bargaining problem. *Econometrica 18*, 155–162.

Roth, A. E. (1977, February). Individual rationality and Nash's solution to the bargaining problem. *Mathematics of Operations Research 2*(1), 64–65.

Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *Review of Economic Studies 54*, 599–618.

Thomas, D. (1990). Intra-household resource allocation: An inferential approach. *Journal of Human Resources 25*, 635–64.

Agricultural & Resource Economics, University of California, Berkeley