The Next Step

- Computation begins where theory leaves off. Computer models enable us to study larger, more detailed and more realistic theoretical models than aren't amenable to closed-form solutions or whose solutions are hard to characterize theoretically.
- One debate that we have not emphasized is the one of Calibration vs. Estimation. Many complex models are first calibrated in order to assess if they deliver realistic characterization of the important variables in an economic model. Econometricians criticize calibration as being *ad hoc*, incoherent and an inefficient way to do inference and testing. However, econometricians have not been successful in producing computationally feasible methods for doing rigorous econometric inference for large scale general equilibrium models, leaving a gap that has been filled by then non-rigorous, ad hoc calibration approach. In some cases this may be the best we can do.
- Another debate that has been underlying during the semester is the Structural vs. Reduced form estimation controversy. It is feasible, although still very challenging, to do rigorous estimation and inference in "single agent" models that use dynamic life cycle models and dynamic programming methods to study individuals' "best responses" to different incentive systems.
- These models are partial equilibrium, but avoid the "Lucas Critique," that reduced-form econometric models are subject to. This critique basically states that the conclusions of reduced-form models are wrong if we want to extend our conclusions to situations were the current state of the incentive system is likely to change.
- On the other hand, structural models typically require a host of strong maintained optimality, rational expectations, and parametric functional form assumptions that many people find questionable. The Princeton and MIT schools of labor, public economics, and applied econometrics are skeptical of the use of structural models in empirical work and rely on much simpler regression and instrumental variables models and use policy changes as natural experiments or instruments to identify key behavioral parameters.

- Reduced-form models have a decisive limitation that they cannot generally be used to predict how individual behavior will change in out-of-sample policy forecasting experiments whereas structural models can. Policy experiments are an extremely useful way to evaluate the accuracy of structural models, since the parameters of these models can be estimated prior to the change in policy and result in verifiable forecasts of the policy change. Reduced-form models use the policy change to estimate key unknown parameters and thus policy experiments are more used for estimating rather than testing reduced form models.

How to proceed?

- Smart researchers see opportunities where others see problems or controversy. Reduced-form models and structural models are not mutually exclusive, but can be used to learn about each other. Most researchers are likely to start exploring either a theoretical question or an empirical phenomenon.
- If a theory is to be developed is likely to have some empirical implications that reduced form static models can help clarify on the way to possibly fully estimating a model. If the empirical phenomenon is the seed then a reduced form analysis can be a coherent and important first step.
- One thing that you should have realized already is how important parsimonious modeling is, we need the fewest number of states and control to explain a complicated process, without a deep analysis of an empirical phenomenon we are unlikely to be able to appropriately model complicated characterizations of reality, and it will be even more unlikely that we would convince the readers that our representation of reality is to be believed, and our results taken seriously.
- More and more questions in economics are likely to fall into a category where the appropriate characterization of the dynamics and uncertainty is to be decisive. In some fields, like IO, almost no researcher can ignore how important is to take into account the structural aspects of the problem at hand. The dynamics can be very complicated but in most cases with some simplifications it is doable. In other areas the profession seems to be more resistant, when individual decision making is involved the controversy is stronger and it will take some more years to see structural estimation as the mainstream in labor economics, public economics, health economics, and any applied micro field.

- On the other hand, in macroeconomics computational and dynamic models are mainstream and as micro-foundations become more and more important estimable general equilibrium models will go from state-of-the-art to state of the field.
- Not every question has to be answered in the framework of a behavioral model, reduced-form analysis will still produce the majority of the papers in the literature, even in 5 years time, but they will no longer be considered as representing the frontier of knowledge in most areas, that will belong to the structural estimation of discrete and continuous decision processes.

A Theoretical Note

- One of the most important issues in the structural estimation is identification. We want to know under what conditions agents' primitives are *identified.*
- We can pose the latter question as a dynamic form of "revealed preference": given infinite observations $\{s_t, d_t\}$ on the states and decisions of an agent, under what conditions can we use this data to uniquely recover (u, p, β)?
- Unfortunately, without strong *a priori* restrictions on (u, p, β), the answer to this question is negative: the MDP model is "non-parametrically unidentifed" in the sense that there is an equivalence class containing infinitely many distinct primitives (u, p, β) that are consistent with the observations $\{s_t, d_t\}$, and consequently Bellman's principle *per se* has no empirical content in the sense that we can always find primitives that "rationalize" any set of observations.
- Note that these negative findings stand in contrast to the case of static choice models where we know that the hypothesis of optimization *per se* does imply testable restrictions.
- The absence of restrictions in the dynamic case may seem surprising given that the structure of the MDP problem already imposes a number of strong restrictions such as time additive preferences, constant intertemporal discount factors, as well as the expected utility hypothesis itself. However many economists will probably not be surprised by this result since similar results have appeared in other areas such as the literature on choice under uncertainty, game theory. and general equilibrium theory.

- To simplify notation, we establish these results in the context of stationary infinite-horizon MDP's, although the argument carries over almost without modification to the nonstationary, finite horizon case. First we will need some definitions.

  1. The *reduced-form* of an MDP model is the agent's optimal decision rule, δ.
  2. The *structure* of an MDP model is the mapping: $\Lambda(u,p,\beta)=\delta$ defined by:
  $$\delta(s)=arg\max_{d\in D(s)}[v(s,d)], \tag{1}$$
  where v is the unique fixed point to:
  $$v(s,d)=u(s,d)+\beta\int\max_{d'\in D(s')}[v(s',d')]p(ds'|s,d). \tag{2}$$
  The rationale for identifying δ as the reduced-form of the MDP is that it embodies all observable implications of the theory and can be consistently estimated by non-parametric regression given sufficient number of observations $\{s_t,d_t\}$. We can use the reduced-form δ to define an equivalence relation over the space of primitives:
  3. Primitives (u,p,β) and $(\overline{u},\overline{p},\overline{\beta})$ are *observationally equivalent* if:
  $$\Lambda(u,p,\beta)=\Lambda(\overline{u},\overline{p},\overline{\beta}). \tag{3}$$
  Thus $\Lambda^{-1}(\delta)$ is the equivalence class of primitives consistent with decision rule δ. Expected utility theory implies that $\Lambda(u,p,\beta)=\Lambda(au+b,p,\beta)$ for any constants a and b satisfying $a>0$, so at best we will only be able to identify an agent's preferences u up to positive linear transformations.
  4. The stationary MDP problem in (1) and (2) is *non-parametrically identified* if given any reduced-form δ in the range of Λ, and any primitives (u,p,β) and $(\overline{u},\overline{p},\overline{\beta})$ in $\Lambda^{-1}(\delta)$ we have:
  $$\begin{aligned}\beta&=\overline{\beta}\\ p&=\overline{p}\\ u&=a\overline{u}+b,\end{aligned} \tag{4}$$
  for some constants a and b satisfying $a>0$.

Lemma 1.: The MDP problem in (1) and (2) is non-parametrically unidentified.

- The proof of this result is quite simple, but beyond the scope of this lecture. See Rust (1994).
- We now ask whether Bellman's principle places any restrictions on the decision rule δ. In the case of infinite-horizon MDP's, Blackwell's Theorem does provide two general restrictions: δ^* is Markovian and deterministic. In practice, it is extremely difficult to test these restrictions empirically.
- Presumably we could test the first restriction by seeing whether agents' decisions depend on lagged states s_{t-k} for $k = 1, 2, \ldots$. However given that we have not placed any *a priori* bounds on the dimensionality of S, the well-known trick of "expanding the state space" can be used to transform an N^{th} order Markov process into a 1^{st} order Markov process. For example, by defining a new state x_t by $x_t = (H_{t-1}, s_t)$, we see that the general SDP is actually a special case of an MDP where the objective is to maximize the utility of the terminal state and decision, $U(x_T, d_T) = U(H_{T-1}, s_T, d_T) = U(s, d)$.
- The second restriction might be tested by looking for agents who make different choices in some state s in two different time periods: $\delta(s_t) \neq \delta(s_{t+k})$, for some state $s_t = s_{t+k} = s$. However this behavior can be rationalized by a model where the agent is indifferent between several alternatives available in state s and simply chooses one at random.
- The following lemma shows that Bellman's principle implies no other restrictions beyond the two essentially untestable restrictions of Blackwell's Theorem:

  Lemma 2: Given an arbitrary measurable mapping $\delta : S \to D$ there exists primitives (u, p, β) such that $\delta = \Lambda(u, p, \beta)$.

  The proof of this result is straightforward. Given an arbitrary discount factor $\beta \in (0, 1)$ and transition probability p, define u by:

$$u(s, d) = I\{d = \delta(s)\} - \beta. \tag{5}$$

- Then it is easy to see that $v(s,d) = I\{d = \delta(s)\}$ is the unique solution to Bellman's equation (2), so that δ is the optimal decision rule implied by (u, p, β).

- If we are unwilling to place any restrictions on (u, p, β), then Lemma 1 shows that the resulting MDP model is non-parametrically unidentified and Lemma 2 shows that it has no testable implications, in the sense that we can devise an MDP model to "rationalize" any decision rule δ.

- Researchers suggest two ways around the problem. First, Bellman's principal does lead to empirically testable implications in the case of *laboratory experiments* where we have at least partial control over an agent's preferences or beliefs. A famous example of such an experiment is the *Allais paradox* which usually succeeds in rejecting the hypothesis that subjects are expected utility maximizers.

- Second, it is clear that from an economic standpoint, many of the utility functions $u + f - \beta E f$ will be completely unreasonable, as is the utility function $u(s,d) = I\{d = \delta(s)\}$. By imposing additional *identifying restrictions* on (u, p, β) we can usually succeed in identifying a unique set of primitives that are consistent with the data and are plausible on *a priori* grounds.

- At the present time, almost all identifying restrictions are imposed by assuming that (u, p, β) belongs to a parametric family (i.e. u is quadratic, Cobb-Douglas, CES, etc.). The use of laboratory experiments as an additional source of identifying restrictions for structural models is still in its infancy. Heckman (1991) has begun pioneering work on integration of experimental results and *a priori* identifying restrictions in structural estimation problems. It seems unlikely, however, that even extensive experimentation will be sufficient by itself to identify the structure of agents' decision making processes.

- The point is to show that there are strict limits on what we can learn from the data alone: the empirical content of any MDP theory is entirely a result of *a priori* identifying restrictions on (u, p, β).

- Although the results point out the futility of non-parametric estimation of (u, p, β), we definitely do not view them as suggesting that structural estimation and testing of parametric models is a futile exercise, and given the huge and rapidly growing empirical literature on structural estimation, it is clear that most economists do not view it as a futile exercise either.

- The proofs of the lemmas do indicate that in order to obtain testable implications, we can get by with weaker identifying restrictions on p than on u and β. To see this, suppose we invoke the hypothesis of *rational expectations*: i.e. agents' subjective beliefs about the evolution of the state variables coincide with the objectively measurable population probability measure.

- This identifying restriction allows us to do structural estimation by *semi-parametric* methods: i.e. we can use nonparametric methods to consistently estimate $p(ds_{t+1}|s_t,d_t)$ using observed realizations of the controlled process $\{s_t,d_t\}$, and parametric methods to estimate (u,β). However looking at the proofs of Lemmas 1 and 2 we can see that it's not possible to go non-parametric on (u,β) since the identification problem persists even if we assume that p is known *a priori.* Thus we must impose strong identifying restrictions on u and β: the usual way this is done is to assume that u and β are smooth functions of a vector of unknown parameters θ. This is sufficient to produce strong, empirically testable restrictions on the controlled process $\{s_t,d_t\}$.

- Identification of β is particularly problematic and in many cases its value is specified *a priori.* For example in life-cycle models discussed in class, it is very difficult to discriminate between a model with a high β without bequests from a model with a low β that allows for bequests: both models imply a slow rate of asset decumulation that we observe in the data. If we assume a particular parametric functional form for the bequest function B we may be able to identify β "by virtue of functional form", but the estimates are likely to be very sensitive to the specification of B.