Monetary Policy Analysis when Planning Horizons are Finite

Michael Woodford

Columbia University

Macroeconomics Annual Conference
NBER
April 12, 2018

Finite-Horizon Planning

- Paper proposes to relax a particularly strong assumption of DSGE monetary models: the assumption that agents formulate **complete infinite-horizon state-contingent plans** that are optimal, under a correct understanding of how the economy will evolve [according to one's model]

Finite-Horizon Planning

- Paper proposes to relax a particularly strong assumption of DSGE monetary models: the assumption that agents formulate **complete infinite-horizon state-contingent plans** that are optimal, under a correct understanding of how the economy will evolve [according to one's model]

- This is surely not feasible in practice, no matter the extent to which one may assume agents are **motivated** and **experienced**

  — for example, even in artificial environments where set of feasible moves from any position is finite (e.g., chess or go), not even the best professional players (human or AI) can **solve the game by backward induction,** and simply execute the optimal strategy

Finite-Horizon Planning

- What the best programs (DeepMind, AlphaGo) actually do: each time one must move,

  1. **look forward** from one's current position a **finite number** of steps, calculating the possible positions that can be reached by finite sequences of moves [under a model of opponent play]

Finite-Horizon Planning

- What the best programs (DeepMind, AlphaGo) actually do: each time one must move,

  1. **look forward** from one's current position a **finite number** of steps, calculating the possible positions that can be reached by finite sequences of moves [under a model of opponent play]
  2. evaluate those possible positions, using a **value function** that assigns an estimated probability of winning from that position

Finite-Horizon Planning

- What the best programs (DeepMind, AlphaGo) actually do: each time one must move,

1. **look forward** from one's current position a **finite number** of steps, calculating the possible positions that can be reached by finite sequences of moves [under a model of opponent play]
2. evaluate those possible positions, using a **value function** that assigns an estimated probability of winning from that position
3. by backward induction from the nodes at which the tree search has been terminated [and value function applied], **assign a value** to each of the possible **initial moves** from the current position

Finite-Horizon Planning

- What the best programs (DeepMind, AlphaGo) actually do: each time one must move,

1. **look forward** from one's current position a **finite number** of steps, calculating the possible positions that can be reached by finite sequences of moves [under a model of opponent play]
2. evaluate those possible positions, using a **value function** that assigns an estimated probability of winning from that position
3. by backward induction from the nodes at which the tree search has been terminated [and value function applied], **assign a value** to each of the possible **initial moves** from the current position
4. **select** the move with highest estimated value

Finite-Horizon Planning

- A crucial observation about such a strategy: in practice, the value function that is used cannot be exactly correct

  — can only evaluate possible positions on the basis of **a few features,** the average values of which can be estimated from some finite database of prior (or simulated) play — not a **complete description** of the state

Finite-Horizon Planning

- A crucial observation about such a strategy: in practice, the value function that is used cannot be exactly correct

  — can only evaluate possible positions on the basis of **a few features,** the average values of which can be estimated from some finite database of prior (or simulated) play — not a **complete description** of the state

- This is in fact why forward planning improves the algorithm, even when forward planning is only possible for a modest number of steps ahead

Goals of the Paper

- Show that decision makers in macro models can be modeled as thinking like these AI programs

  — and to see to what extent conclusions are similar or different from rational expectations equilibrium analysis

Goals of the Paper

- Show that decision makers in macro models can be modeled as thinking like these AI programs

  — and to see to what extent conclusions are similar or different from rational expectations equilibrium analysis

- Will illustrate the method in the context of a basic New Keynesian DSGE model

Goals of the Paper

- Show that decision makers in macro models can be modeled as thinking like these AI programs

  — and to see to what extent conclusions are similar or different from rational expectations equilibrium analysis

- Will illustrate the method in the context of a basic New Keynesian DSGE model

- Application: consider predicted effects of "forward guidance" about monetary policy years into the future

Household with k-Period Planning Horizon

- Household i problem in period t: choose spending plan $\{c^i_\tau(s_\tau)\}$ for periods $t \leq \tau \leq t+k$ to maximize

$$\hat{\mathrm{E}}^i_t \sum_{\tau=t}^{t+k} \beta^{\tau-t} u(c^i_\tau; \xi_\tau) + \beta^{k+1} v(b^i_{t+k+1}; s_{t+k})$$

subject to constraints

$$b^i_{\tau+1} = (1+i_\tau)\left[b^i_\tau (P_{\tau-1}/P_\tau) + y_\tau - c^i_\tau\right]$$

for all $t \leq \tau \leq t+k$,

- Here $v(b^i_{\tau+1}; s_\tau)$ is the **value function** used to evaluate possible situations in a terminal state s_τ

Decisions with a Finite Planning Horizon

- Expectations about periods $t \leq \tau \leq t + k$ used in planning exercise:
  - **deduced from structural equations of model** (including monetary policy rule) for periods t through $t + k$
  - hence **take account** of any announced changes in policy, over the planning horizon

Decisions with a Finite Planning Horizon

- Expectations about periods $t \leq \tau \leq t+k$ used in planning exercise:
  - **deduced from structural equations of model** (including monetary policy rule) for periods t through $t+k$
  - hence **take account** of any announced changes in policy, over the planning horizon
  - but no consideration of branches **beyond horizon** $t+k$ means aggregate conditions in period $t+j$ assumed to be determined by decisions of people who plan **only** $k-j$ **periods ahead**

Decisions with a Finite Planning Horizon

- Expectations about periods $t \leq \tau \leq t + k$ used in planning exercise:
  - **deduced from structural equations of model** (including monetary policy rule) for periods t through $t + k$
  - hence **take account** of any announced changes in policy, over the planning horizon
  - but no consideration of branches **beyond horizon** $t + k$ means aggregate conditions in period $t + j$ assumed to be determined by decisions of people who plan **only** $k - j$ **periods ahead**
- Just as household models **own** behavior in future period $t + j$ as if will only have horizon of length $k - j$ then, models all **other households and firms** as optimizing, but only having horizons of length $k - j$ in period $t + j$

Decisions with a Finite Planning Horizon

- Assumptions about **value function** used by households:
  - depends only on a **coarse** description of state s_{t+k}
  - in particular: assumed **not** to take into account any announcements about **new policies** that may apply after date $t + k$

Decisions with a Finite Planning Horizon

- Assumptions about **value function** used by households:
  - depends only on a **coarse** description of state s_{t+k}
  - in particular: assumed **not** to take into account any announcements about **new policies** that may apply after date $t + k$
  - in specific results shown here: assume $v(b^i_{t+k+1})$ doesn't depend on **any** other state variables

Decisions with a Finite Planning Horizon

- Assumptions about **value function** used by households:
  - depends only on a **coarse** description of state s_{t+k}
  - in particular: assumed **not** to take into account any announcements about **new policies** that may apply after date $t+k$
  - in specific results shown here: assume $v(b^i_{t+k+1})$ doesn't depend on **any** other state variables
- Value function is **learned from past experience**

Learning the Household Value Function

- As part of its finite-horizon planning exercise in period t, each household i computes an **estimate** of the value of its objective [expected discounted utility from t onward], for any (counterfactual) level of wealth $b_t^i = b$ that it **might** have brought into the period

  — call this function $v_t^{est}(b)$

Learning the Household Value Function

- As part of its finite-horizon planning exercise in period t, each household i computes an **estimate** of the value of its objective [expected discounted utility from t onward], for any (counterfactual) level of wealth $b_t^i = b$ that it **might** have brought into the period

  — call this function $v_t^{est}(b)$

- Let the value function used in its period t planning exercise be $v_t(b)$; this is then **updated** for the next period's exercise using the **error-correction** rule ["constant-gain learning"]

$$v_{t+1}(b) \;=\; v_t(b) + \gamma\,(v_t^{est}(b) - v_t(b)) \qquad \text{for all } b$$

  for some $0 < \gamma \leq 1$.

Learning the Household Value Function

- As part of its finite-horizon planning exercise in period t, each household i computes an **estimate** of the value of its objective

- Let the value function used in its period t planning exercise be $v_t(b)$; this is then **updated** for the next period's exercise using the **error-correction** rule ["constant-gain learning"]

$$v_{t+1}(b) \;=\; v_t(b) + \gamma\left(v_t^{est}(b) - v_t(b)\right) \qquad \text{for all } b$$

  for some $0 < \gamma \leq 1$.

  - in a constant environment, this converges to the $v(b)$ that solves the Bellman equation for infinite-horizon optimization [essentially, solution through **value-function iteration**]

Log-Linearized Dynamics: Aggregate Demand

- Log-linearize household decision rule around perfect foresight steady state with constant inflation $\bar{\Pi}$, and assume value function $v(b)$ that is optimal for that steady state

Log-Linearized Dynamics: Aggregate Demand

- Log-linearize household decision rule around perfect foresight steady state with constant inflation $\bar{\Pi}$, and assume value function $v(b)$ that is optimal for that steady state

- Let y_t^j be aggregate real spending (and income) in period t if all households have planning horizon j [any $j \geq 0$]; similarly, π_t^j overall inflation rate if all firms have planning horizon j, $\hat{\imath}_t^j$ the nominal interest rate if CB reacts to π_t^j and y_t^j

  — then finite-horizon planning implies that

$$y_t^j - g_t \;=\; \mathrm{E}_t[y_{t+1}^{j-1} - g_{t+1}] \;-\; \sigma(\hat{\imath}_t^j - \mathrm{E}_t \pi_{t+1}^{j-1})$$

for each $j \geq 1$, and

$$y_t^0 - g_t \;=\; -\sigma\, \hat{\imath}_t^0$$

Log-Linearized Dynamics: Aggregate Demand

- Let y_t^j be aggregate real spending (and income) in period t if all households have planning horizon j [any $j \geq 0$]; similarly, π_t^j overall inflation rate if all firms have planning horizon j, $\hat{\imath}_t^j$ the nominal interest rate if CB reacts to π_t^j and y_t^j

  — then finite-horizon planning implies that

$$y_t^j - g_t \;=\; \mathrm{E}_t[y_{t+1}^{j-1} - g_{t+1}] \;-\; \sigma(\hat{\imath}_t^j - \mathrm{E}_t \pi_{t+1}^{j-1})$$

  for each $j \geq 1$, and

$$y_t^0 - g_t \;=\; -\sigma\,\hat{\imath}_t^0$$

- Compare to prediction with infinite-horizon optimization:

$$y_t - g_t \;=\; \mathrm{E}_t[y_{t+1} - g_{t+1}] \;-\; \sigma(\hat{\imath}_t - \mathrm{E}_t \pi_{t+1})$$

Price Setting

Can apply a similar analysis to the decisions of price-setting firms:

- Assume a Calvo-Yun model of staggered price adjustment by monopolistic competitors

- But suppose that a firm f that reoptimizes its price has only a **k-period planning horizon**

  — assigns value $\tilde{v}(P_t^f / P_{t+k})$ to continuation profits, if newly chosen price P_t^f still applies in period $t + k + 1$

  — and again this value function is **learned** using an error-correction rule

Log-Linearized Dynamics: Aggregate Supply

- Linearizing optimal finite-horizon pricing rule, and aggregating prices of all firms, yields

$$\pi_t^j = \kappa\,(y_t^j - y_t^*) + \beta\,\mathrm{E}_t \pi_{t+1}^{j-1}$$

for all $j \geq 1$, and

$$\pi_t^0 = \kappa\,(y_t^0 - y_t^*)$$

Log-Linearized Dynamics: Aggregate Supply

- Linearizing optimal finite-horizon pricing rule, and aggregating prices of all firms, yields

$$\pi_t^j = \kappa\,(y_t^j - y_t^*) + \beta\,\mathrm{E}_t \pi_{t+1}^{j-1}$$

for all $j \geq 1$, and

$$\pi_t^0 = \kappa\,(y_t^0 - y_t^*)$$

- Compare to prediction with infinite-horizon optimization:

$$\pi_t = \kappa\,(y_t - y_t^*) + \beta\,\mathrm{E}_t \pi_{t+1}$$

Monetary Policy

- Closing the model: assume monetary policy specified by a **reaction function**

$$\hat{\imath}_t = i_t^* + \phi_{\pi,t}\pi_t + \phi_{y,t}y_t$$

with possibly time-varying coefficients $\phi_{\pi,t}, \phi_{y,t} \geq 0$.

- Hence for any planning horizon $j \geq 0$,

$$\hat{\imath}_t^j = i_t^* + \phi_{\pi,t}\pi_t^j + \phi_{y,t}y_t^j$$

A Forward Guidance Experiment

- Suppose that for some time prior to the policy experiment, economy has been in steady state with inflation rate $\bar{\Pi}$, and value functions optimal for that environment have been learned

A Forward Guidance Experiment

- Suppose that for some time prior to the policy experiment, economy has been in steady state with inflation rate $\bar{\Pi}$, and value functions optimal for that environment have been learned

- But suppose that at some date t_0, it is **announced** that monetary policy will be determined by a **different** reaction function (consistent with a different inflation target, and possibly different response coefficients as well) until some date T

  — from date T onward, policy will instead revert to "normal" reaction function, a Taylor rule consistent with inflation rate $\bar{\Pi}$ once again

Equilibria with Forward Guidance

- For any assumed planning horizon k, we can **uniquely** solve for model's predictions for dynamics

— horizon $j = 0$ variables involve no forward planning, so uniquely determined [by backward-looking value function]

— horizon $j \geq 1$ variables uniquely determined if horizon $j - 1$ variables are uniquely determined

Equilibria with Forward Guidance

- For any assumed planning horizon k, we can **uniquely** solve for model's predictions for dynamics

  — horizon $j = 0$ variables involve no forward planning, so uniquely determined [by backward-looking value function]

  — horizon $j \geq 1$ variables uniquely determined if horizon $j - 1$ variables are uniquely determined

- If the planning horizons of all households and firms satisfy $k \geq T - t_0 - 1$, then model predictions **coincide** with a rational expectations equilibrium

  — the **specific** RE solution in which economy returns to steady state from date T onward

Equilibria with Forward Guidance

- If the newly announced reaction function conforms to the **"Taylor Principle"** $[\phi_\pi + (1 - \beta/\kappa)\phi_y > 1]$, then this RE solution **converges** for $T \to \infty$

  $\Rightarrow$ even in the case of a **permanent** change in policy, solution with finite-horizon planning will **approximate** the RE solution, if sufficiently many have sufficiently long horizons

Equilibria with Forward Guidance

- If the newly announced reaction function conforms to the **"Taylor Principle"** $[\phi_\pi + (1 - \beta/\kappa)\phi_y > 1]$, then this RE solution **converges** for $T \to \infty$

  $\Rightarrow$ even in the case of a **permanent** change in policy, solution with finite-horizon planning will **approximate** the RE solution, if sufficiently many have sufficiently long horizons

  - thus finite-horizon analysis can justify use of RE analysis (as a simplifying approximation) in this case
  - and solves the **equilibrium selection** problem that bedevils RE analysis

Equilibria with Forward Guidance

- If the newly announced reaction function conforms to the **"Taylor Principle"** $[\phi_\pi + (1 - \beta/\kappa)\phi_y > 1]$, then this RE solution **converges** for $T \to \infty$

  $\Rightarrow$ even in the case of a **permanent** change in policy, solution with finite-horizon planning will **approximate** the RE solution, if sufficiently many have sufficiently long horizons

  - thus finite-horizon analysis can justify use of RE analysis (as a simplifying approximation) in this case
  - and solves the **equilibrium selection** problem that bedevils RE analysis

- But the situation is very different if Taylor Principle **not** satisfied, as in case of an interest-rate peg

Equilibria with an Interest-Rate Peg

- RE results, for temporary peg [e.g., commit to keep interest rate at lower bound], if select equilibrium which returns to steady state from $t = T$ onward:
  - effects on output and inflation predicted to **grow explosively** as $T \to \infty \Rightarrow$ forward guidance should be **extremely** effective (if credible)
  - but this is often regarded as an implausible prediction ["the forward guidance puzzle" (Del Negro *et al.*, 2015)]

Equilibria with an Interest-Rate Peg

- RE results, for temporary peg [e.g., commit to keep interest rate at lower bound], if select equilibrium which returns to steady state from $t = T$ onward:
  - effects on output and inflation predicted to **grow explosively** as $T \to \infty \Rightarrow$ forward guidance should be **extremely** effective (if credible)
  - but this is often regarded as an implausible prediction ["the forward guidance puzzle" (Del Negro *et al.*, 2015)]
- Instead, with any finite horizon k, commitment to a low-rate peg is predicted to be **stimulative,** but effects remain **bounded** no matter how long the commitment

— and may be quite modest, if k is not too large

Equilibria with an Interest-Rate Peg

- RE results, for a **permanent** peg: all non-explosive RE solutions converge in long run to inflation rate consistent with **Fisher Equation**

  — implying that commitment to maintaining a lower nominal interest rate should **lower** inflation, at least eventually

Equilibria with an Interest-Rate Peg

- RE results, for a **permanent** peg: all non-explosive RE solutions converge in long run to inflation rate consistent with **Fisher Equation**

  — implying that commitment to maintaining a lower nominal interest rate should **lower** inflation, at least eventually

- Instead, with any finite horizon k, commitment to lower rate (even if permanent) must **increase** inflation

  — **none** of the RE solutions are similar to the finite-horizon solution, for **any** distribution of planning horizons

  — suggesting that RE analysis may be quite misleading in this case

Conclusions

- Care must be used in drawing conclusions about contemplated monetary policies using RE analysis

Conclusions

- Care must be used in drawing conclusions about contemplated monetary policies using RE analysis

- In some cases, RE outcome (with suitable equilibrium selection) **is** a decent approximation to what a model of finite-horizon forward planning would imply

Conclusions

- Care must be used in drawing conclusions about contemplated monetary policies using RE analysis

- In some cases, RE outcome (with suitable equilibrium selection) **is** a decent approximation to what a model of finite-horizon forward planning would imply

  — but in other cases (e.g., commitment to maintain a fixed nominal interest rate for a long time), conclusions are **very different,** even if one assumes highly sophisticated forward planning for a long distance into future

- Hence checking the **robustness** of conclusions to modest departures from perfect rationality is important for monetary policy analysis