Chapter 23
Credible Government Policies, I

23.1. Introduction

Kydland and Prescott (1977) opened the modern discussion of time consistency in macroeconomics with some examples that show how outcomes differ in otherwise identical economies when the assumptions about the timing of government policy choices are altered.[1] In particular, they compared a timing protocol in which a government chooses its (possibly history-contingent) policies once and for all at the beginning of time with one in which the government chooses sequentially. Because outcomes are worse when the government chooses sequentially, Kydland and Prescott's examples illustrate the value to a government of having access to a "commitment technology" that requires it not to choose sequentially.

Subsequent work on time consistency focused on how a reputation can substitute for a commitment technology when the government chooses sequentially.[2] The issue is whether constraints confronting the government and private sector expectations can be arranged so that a government adheres to an expected pattern of behavior because it would worsen its reputation if it did not.

A "folk theorem" from game theory states that if there is no discounting of future payoffs, then virtually any first-period payoff can be sustained as a reputational equilibrium. A main purpose of this chapter is to study how discounting might shrink the set of outcomes that are attainable with a reputational mechanism.

Modern formulations of reputational models of government policy exploit ideas from dynamic programming. Each period, a government faces choices whose consequences include a first-period return and a reputation to pass on to

[1] Consider two extensive-form versions of the "battle of the sexes" game described by Kreps (1990), one in which the man chooses first, the other in which the woman chooses first. Backward induction recovers different outcomes in these two different games. Though they share the same choice sets and payoffs, these are different games.

[2] Barro and Gordon (1983a, 1983b) are early contributors to this literature. See Kenneth Rogoff (1989) for a survey.

next period. Under rational expectations, any reputation that the government carries into next period must be one that it will want to confirm. We shall study the set of possible values that the government can attain with reputations that it could conceivably want to confirm.

This and the following chapter apply an apparatus of Abreu, Pearce, and Stacchetti (1986, 1990) (APS) to reputational equilibria in a class of macroeconomic models. APS exploit ideas from dynamic programming.[3] Their work exploits the insight that it is more convenient to work with the set of continuation values associated with equilibrium strategies than it is to work directly with the set of equilibrium strategies. We use an economic model like those of Chari, Kehoe, and Prescott (1989) and Stokey (1989, 1991) to exhibit what Chari and Kehoe (1990) call sustainable government policies and what Stokey calls credible public policies. The literature on sustainable or credible government policies in macroeconomics adapts ideas from the literature on repeated games so that they can be applied in contexts in which a single agent (a government) behaves strategically, and in which all other agents' behavior can be summarized in terms of a system of expectations about government actions together with competitive equilibrium outcomes that respond to the government's choices.[4]

[3] This chapter closely follow Stacchetti (1991), who applies Abreu, Pearce, and Stacchetti (1986, 1990) to a more general class of models than that treated here. Stacchetti also studies a class of setups in which the private sector observes only a noise-ridden signal of the government's actions.

[4] For descriptions of theories of credible government policy, see Chari and Kehoe (1990), Stokey (1989, 1991), Rogoff (1989), and Chari, Kehoe, and Prescott (1989). For applications of the framework of Abreu, Pearce, and Stacchetti, see Chang (1998), and Phelan and Stacchetti (1999).

23.1.1. Diverse sources of history dependence

The theory of credible government policy uses particular kinds of history dependence to render credible a sequence of actions chosen by a *sequence* of policy makers. Here *credible* means an action that a government decision maker *wants* to implement. By way of contrast, in chapter 19, we encountered a distinct source of history dependence in the policy of a Ramsey planner or Stackelberg. There history dependence came from the requirement that it is necessary properly to account for constraints that dynamic aspects of private sector behavior put on the time t choices of a Ramsey planner or Stackelberg leader that at time 0 makes once-and-for-all choices of intertemporal sequences. In this context, history dependence must take into account the requirement that at time t the Ramsey planner must confirm private sector expectations about those time t actions that at time 0 were partly designed to influence private sector outcomes in periods $0, \ldots, t-1$.

In settings in which private agents face genuinely dynamic decision problems having their own endogenous state variables like various forms of physical and human capital, *both* of these sources of history dependence influence a credible policy. It can be subtle to disentangle the economic forces contributing to history dependence in government policies in such settings. However, for a special classes of examples with private agents' choices sufficiently simplified that they deprive private agents' decision problem of any 'natural' state variables, we can isolate the source of history dependence coming from the requirement that a government policy be must be credible. We consider such a special class in this chapter for the avowed purpose of isolating the source of history dependence coming from credibility considerations and distinguishing it from the chapter 19 source that instead comes from the need to respect substantial dynamics coming from equilibrium private sector behavior. Having isolated one source of history dependence in chapter 19 and another in the present chapter, we proceed in chapter 24 to activate both sources of history dependence and then seek a recursive representation for a credible government policy in that more sophisticated setting.

23.2. The one-period economy

There is a continuum of households, each of which chooses an action $\xi \in X$. A government chooses an action $y \in Y$. The sets X and Y are compact. The average level of ξ across households is denoted $x \in X$. The utility of a particular household is $u(\xi, x, y)$ when it chooses ξ, when the average household's choice is x, and when the government chooses y. The payoff function $u(\xi, x, y)$ is strictly concave and continuously differentiable in ξ and y.[5]

23.2.1. Competitive equilibrium

For given levels of y and x, the representative household faces the problem $\max_{\xi \in X} u(\xi, x, y)$. Let the solution be a function $\xi = f(x, y)$. When a household believes that the government's choice is y and that the average level of other households' choices is x, it acts to set $\xi = f(x, y)$. Because all households are alike, this fact implies that the actual level of x is $f(x, y)$. For the representative household's expectations about the average to be consistent with the average outcome, we require that $\xi = x$, or $x = f(x, y)$. This makes the representative agent representative. We use the following:[6]

DEFINITION 1: A *competitive equilibrium* or a *rational expectations equilibrium* is an $x \in X$ that satisfies $x = f(x, y)$.

A competitive equilibrium satisfies $u(x, x, y) = \max_{\xi \in X} u(\xi, x, y)$.

For each $y \in Y$, let $x = h(y)$ denote the corresponding competitive equilibrium. We adopt:

DEFINITION 2: The set of competitive equilibria is $C = \{(x, y) \mid u(x, x, y) = \max_{\xi \in X} u(\xi, x, y)\}$, or equivalently $C = \{(x, y) \mid x = h(y)\}$.

[5] The discrete-choice examples given later violate some of these assumptions in non essential ways.

[6] See the definition of a rational expectations equilibrium in chapter 7.

23.2.2. The Ramsey problem

The following timing of actions underlies a *Ramsey plan.* First, the government selects a $y \in Y$. Then, knowing the government's choice of y, the aggregate of households responds with a competitive equilibrium. The government evaluates policies $y \in Y$ with the payoff function $u(x,x,y)$; that is, the government is benevolent.

We assume that in choosing y, the government has to forecast how the economy will respond. The government correctly forecasts that the economy will respond to y with a competitive equilibrium, $x = h(y)$. We use these definitions:

DEFINITION 3: The *Ramsey problem* is $\max_{y\in Y} u[h(y), h(y), y]$, or equivalently $\max_{(x,y)\in C} u(x,x,y)$.

DEFINITION 4: The policy that attains the maximum for the Ramsey problem is denoted y^R. Let $x^R = h(y^R)$. Then (y^R, x^R) is called the *Ramsey outcome* or *Ramsey plan.*

Two remarks about the Ramsey problem are in order. First, the Ramsey outcome is typically inferior to the "dictatorial outcome" that solves the unrestricted problem $\max_{x\in X,\, y\in Y} u(x,x,y)$, because the restriction $(x,y) \in C$ is in general binding. Second, the timing of actions is important. The Ramsey problem assumes that the government chooses first and must stick with its choice regardless of how private agents subsequently choose $x \in X$.

If the government were granted the opportunity to reconsider its plan *after* households had chosen $x = x^R$, the government would in general want to deviate from y^R because often there exists an $\alpha \neq y^R$ for which $u(x^R, x^R, \alpha) > u(x^R, x^R, y^R)$. The "time consistency problem" is the incentive the government would have to deviate from the Ramsey plan if it were allowed to react *after* households had set $x = x^R$. In this one-period setting, to support the Ramsey plan requires a timing protocol that forces the government to choose first.

23.2.3. Nash equilibrium

Consider an alternative timing protocol that confronts households with a forecasting problem because the government chooses after or simultaneously with the households. Assume that households forecast that, given x, the government will set y to solve $\max_{y \in Y} u(x,x,y)$. We use:

DEFINITION 5: A *Nash equilibrium* (x^N, y^N) satisfies

(1) $(x^N, y^N) \in C$;

(2) Given x^N, $u(x^N, x^N, y^N) = \max_{\eta \in Y} u(x^N, x^N, \eta)$.

Condition (1) asserts that $x^N = h(y^N)$, or that the economy responds to y^N with a competitive equilibrium. Thus, condition (1) says that given (x^N, y^N), each individual household wants to set $\xi = x^N$; that is, the representative household has no incentive to deviate from x^N. Condition (2) asserts that given x^N, the government chooses a policy y^N from which it has no incentive to deviate.[7]

We can use the solution of the problem in condition (2) to define the government's *best response* function $y = H(x)$. The definition of a Nash equilibrium can be phrased as a pair $(x, y) \in C$ such that $y = H(x)$.

There are two timings of choices for which a Nash equilibrium is a natural equilibrium concept. One is where households choose first, forecasting that the government will respond to the aggregate outcome x by setting $y = H(x)$. Another is where the government and households choose simultaneously, in which case a Nash equilibrium (x^N, y^N) depicts a situation in which everyone has rational expectations: given that each household expects the aggregate variables to be (x^N, y^N), each household responds in a way to make $x = x^N$, and given that the government expects that $x = x^N$, it responds by setting $y = y^N$.

We let the values attained by the government under the Nash and Ramsey outcomes, respectively, be denoted $v^N = u(x^N, x^N, y^N)$ and $v^R = u(x^R, x^R, y^R)$. Because of the additional constraint embedded in the Nash equilibrium, outcomes are ordered according to

$$v^N = \max_{\{(x,y) \in C:\, y = H(x)\}} u(x,x,y) \le \max_{(x,y) \in C} u(x,x,y) = v^R .$$

[7] Much of the language of this chapter is borrowed from game theory, but the object under study is not a game, because we do not specify all of the objects that formally define a game. In particular, we do not specify the payoffs to all agents for all feasible choices. We only specify the payoffs $u(\xi, x, y)$ where each private agent chooses the *same* value of ξ.

23.3. Nash and Ramsey outcomes

To illustrate these concepts, we consider two examples: taxation within a fully specified economy, and a black-box model with discrete choice sets.

23.3.1. Taxation example

Each of a continuum of households has preferences over leisure ℓ, private consumption c, and per capita government expenditures g. The utility function is

$$U(\ell, c, g) = \ell + \log(\alpha + c) + \log(\alpha + g), \qquad \alpha \in (0, 1/2).$$

Each household is endowed with one unit of time that can be devoted to leisure or labor. The production technology is linear in labor, and the economy's resource constraint is

$$\overline{c} + g = 1 - \overline{\ell},$$

where $\overline{c}$ and $\overline{\ell}$ are the average levels of private consumption and leisure, respectively.

A benevolent government wants to maximize the utility of the representative household. A benevolent government that is subject only to the constraint imposed by the technology and would choose $\ell = 0$ and $c = g = 1/2$. This "dictatorial outcome" yields welfare $W^d = 2\log(\alpha + 1/2)$.

Competitive equilibrium in general imposes more restrictions on the allocations attainable by a benevolent government. Here we will focus on competitive equilibria where the government finances its expenditures by levying a flat-rate tax τ on labor income. The household's budget constraint at equality is $c = (1-\tau)(1-\ell)$. Given a government policy (τ, g), an individual household's optimal decision rule for leisure is

$$\ell(\tau) = \begin{cases} \dfrac{\alpha}{1-\tau} & \text{if } \tau \in [0, 1-\alpha]; \\ 1 & \text{if } \tau > 1-\alpha. \end{cases}$$

Due to the linear technology and the fact that government expenditures enter additively in the utility function, the household's decision rule $\ell(\tau)$ is also the equilibrium value of individual leisure at a given tax rate τ. Imposing government budget balance, $g = \tau(1-\ell)$, the representative household's welfare in a competitive equilibrium can be expressed as a function of τ and is equal to

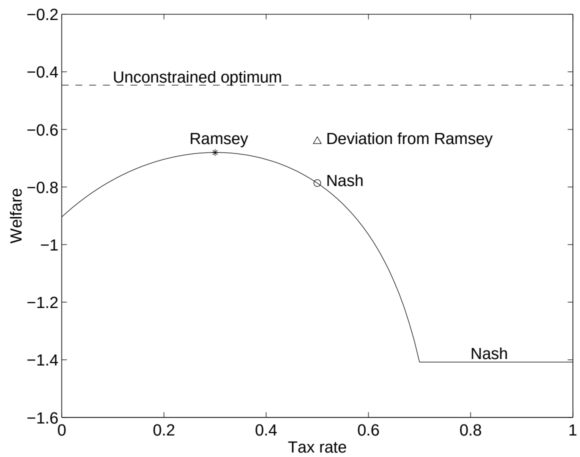

Figure 23.3.1: Welfare outcomes in the taxation example. The solid curve depicts the welfare associated with the set of competitive equilibria, $W^c(\tau)$. The set of Nash equilibria is the horizontal portion of the solid curve and the equilibrium at $\tau = 1/2$. The Ramsey outcome is marked with an asterisk. The "time inconsistency problem" is indicated with the triangle showing the outcome if the government were able to reset τ after households had chosen the Ramsey labor supply. The dashed line describes the welfare level at the unconstrained optimum, W^d. The graph sets $\alpha = 0.3$.

$$W^c(\tau) = \ell(\tau) + \log\{\alpha + (1-\tau)[1-\ell(\tau)]\} + \log\{\alpha + \tau[1-\ell(\tau)]\}.$$

The Ramsey tax rate and allocation are determined by the solution to $\max_\tau W^c(\tau)$. It can be verified that because $\alpha \in (0, .5)$, the Ramsey plan sets $\tau < .5$, which produces an allocation in which c, g, and $1-\ell$ are all positive.

By way of contrast, the government's problem in a Nash equilibrium is $\max_\tau \{\ell + \log[\alpha + (1-\tau)(1-\ell)] + \log[\alpha + \tau(1-\ell)]\}$. If $\ell < 1$, the optimizer is $\tau = .5$. There is a continuum of Nash equilibria indexed by $\tau \in [1-\alpha, 1]$ where agents choose not to work, and consequently $c = g = 0$. The only Nash equilibrium with production is $\tau = 1/2$ with welfare level $W^c(1/2)$. This conclusion follows directly from the fact that the government's best response is $\tau = 1/2$ for any $\ell < 1$. These outcomes are illustrated numerically in Figure 23.3.1. Here the time inconsistency problem surfaces in the government's incentive, if offered the choice, to reset the tax rate τ, after the household has set its labor supply.

The objects of the general setup in the preceding section can be mapped into the present taxation example as follows: $\xi = \ell$, $x = \overline{\ell}$, $X = [0,1]$, $y = \tau$, $Y = [0,1]$, $u(\xi,x,y) = \xi + \log[\alpha + (1-y)(1-\xi)] + \log[\alpha + y(1-x)]$, $f(x,y) = \ell(y)$, $h(y) = \ell(y)$, and $H(x) = 1/2$ if $x < 1$; and $H(x) \in [0,\, 1]$ if $x = 1$.

23.3.2. Black-box example with discrete choice sets

Consider a black box example with $X = \{x_L,\, x_H\}$ and $Y = \{y_L,\, y_H\}$, in which $u(x,x,y)$ assume the values given in Table 23.3.1. Assume that values of $u(\xi,x,y)$ for $\xi \neq x$ are such that the values with asterisks for $\xi = x$ are competitive equilibria. In particular, we might assume that

$$u(\xi,x_i,y_j) = 0 \quad \text{when } \xi \neq x_i \text{ and } i = j,$$
$$u(\xi,x_i,y_j) = 20 \quad \text{when } \xi \neq x_i \text{ and } i \neq j.$$

These payoffs imply that $u(x_L,x_L,y_L) > u(x_H,x_L,y_L)$ (i.e., $3 > 0$), and $u(x_H,x_H,y_H) > u(x_L,x_H,y_H)$ (i.e., $10 > 0$). Therefore, (x_L,x_L,y_L) and (x_H,x_H,y_H) are competitive equilibria. Also, $u(x_H,x_H,y_L) < u(x_L,x_H,y_L)$ (i.e., $12 < 20$), so the dictatorial outcome cannot be supported as a competitive equilibrium.

	x_L	x_H
y_L	3*	12
y_H	1	10*

Table 23.3.1: One-period payoffs $u(x_i,x_i,y_j)$; * denotes $(x,y) \in C$; the Ramsey outcome is (x_H,y_H) and the Nash equilibrium outcome is (x_L,y_L).

Figure 23.3.2 depicts a timing of choices that supports the Ramsey outcome for this example. The government chooses first, then walks away. The Ramsey outcome (x_H, y_H) is the competitive equilibrium yielding the highest value of $u(x,x,y)$.

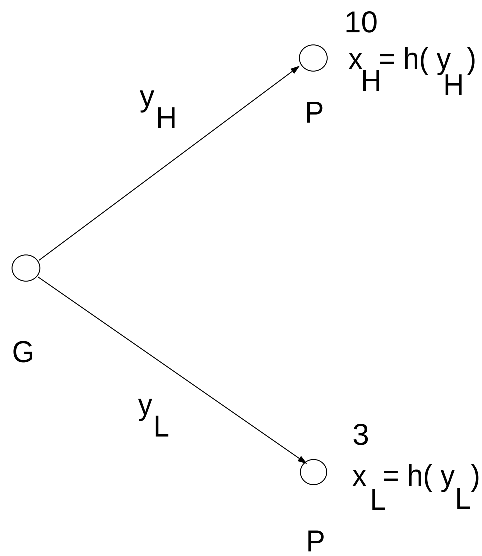

Figure 23.3.2: Timing of choices that supports Ramsey outcome. Here P and G denote nodes at which the public and the government, respectively, choose. The government has a commitment technology that binds it to "choose first." The government chooses the $y \in Y$ that maximizes $u[h(y),\ h(y), y]$, where $x = h(y)$ is the function mapping government actions into equilibrium values of x.

Figure 23.3.3 diagrams a timing of choices that supports the Nash equilibrium. Recall that by definition, every Nash equilibrium outcome has to be a competitive equilibrium outcome. We denote competitive equilibrium pairs (x, y) with asterisks. The government sector chooses after knowing that the private sector has set x, and chooses y to maximize $u(x,x,y)$. With this timing, if the private sector chooses $x = x_H$, the government has an incentive to set $y = y_L$, a setting of y that does not support x_H as a Nash equilibrium. The unique Nash equilibrium is (x_L, y_L), which gives a lower utility $u(x,x,y)$ than does the competitive equilibrium (x_H, y_H).

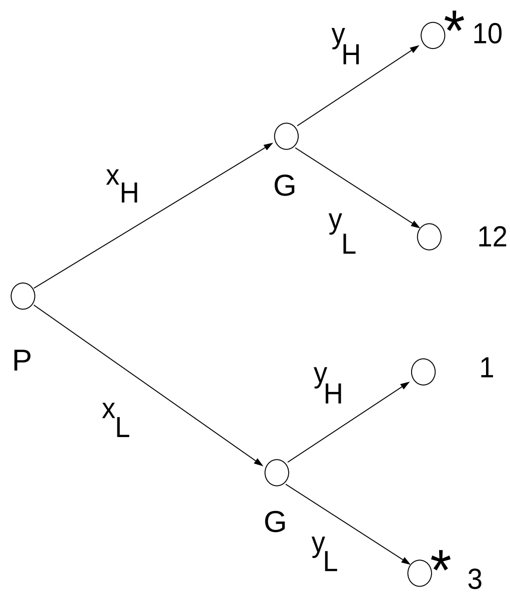

Figure 23.3.3: Timing of actions in a Nash equilibrium in which the private sector acts first. Here G denotes a node at which the government chooses and P denotes a node at which the public chooses. The private sector sets $x \in X$ before knowing the government's setting of $y \in Y$. Competitive equilibrium pairs (x,y) are denoted with an asterisk. The unique Nash equilibrium is (x_L, y_L).

23.4. Reputational mechanisms: General idea

In a finitely repeated economy, the government will certainly behave opportunistically the last period, implying that nothing better than a Nash outcome can be supported the last period. In a finite horizon economy with a unique Nash equilibrium, we won't be able to sustain anything better than a Nash equilibrium outcome in *any* earlier period.[8]

We want to study situations in which a government might sustain a Ramsey outcome. Therefore, we shall study economies repeated an infinite number of times. Here a system of history-dependent expectations interpretable as a government reputation might be arranged to sustain something better than the Nash outcome. We strive to set things up so that the government so dearly wants to confirm a reputation that it will not submit to the temptation to behave opportunistically. A reputation is said to be *sustainable* if it is always in the government's interests to confirm it.

A state variable that is capable of encoding a "reputation" is peculiar because it is both "backward looking" and "forward looking." It is backward looking because it remembers salient features of past behavior. It is forward-looking behavior because it measures something about what private agents expect the government to do in the future. We are about to study the ingenious machinery of Abreu, Pearce, and Stacchetti that astutely exploits these aspects of a reputational variable by recognizing that the ideal reputational state variable is a "promised value."

---

[8] If there are multiple Nash equilibria, it is sometimes possible to sustain a better-than-Nash equilibrium outcome for a while in a finite horizon economy. See exercise *23.1*, which uses an idea of Benoit and Krishna (1985).

23.4.1. Dynamic programming squared

A sustainable reputation for the government is one that (a) the public, having rational expectations, wants to believe, and (b) the government wants to sustain. Rather than finding all possible sustainable reputations, Abreu, Pearce, and Stacchetti (henceforth APS) (1986, 1990) used dynamic programming to characterize all *values* for the government that are attainable with sustainable reputations. This section briefly describes their main ideas, while later sections fill in many details.

First we need some language. A *strategy profile* is a pair of plans, one each for the private sector and the government. The time t components of the pair of plans maps the observed history of the economy into current-period outcomes (x, y). A *subgame perfect equilibrium* (SPE) strategy profile has a current-period outcome being a competitive equilibrium (x_t, y_t) whose y_t component the government would want to confirm at each $t \geq 1$ and for every possible history of the economy.

To characterize SPE, the method of APS is to formulate a Bellman equation that describes the value to the government of a strategy profile and that portrays the idea that the government wants to confirm the private sector's beliefs about y. For each $t \geq 1$, the government's strategy describes its first-period action $y \in Y$, which, because the public had expected it, determines an associated first-period competitive equilibrium $(x, y) \in C$. Furthermore, the strategy implies two continuation values for the government at the beginning of next period, a continuation value v_1 if it carries out the first-period choice y, and another continuation value v_2 if for any reason the government deviates from the expected first-period choice y. Associated with the government's strategy is a current value v that obeys the Bellman equation

$$v = (1-\delta)u(x, x, y) + \delta v_1, \tag{23.4.1a}$$

where $\delta \in (0,1)$ is a discount factor, $(x, y) \in C$, v_1 is a continuation value awarded for confirming the private sector's expectation that the government will choose action y in the current period, and (y, v_1) are constrained to satisfy the incentive constraint

$$v \geq (1-\delta)u(x, x, \eta) + \delta v_2, \quad \forall \eta \in Y, \tag{23.4.1b}$$

or equivalently

$$v \geq (1-\delta)u\big[x, x, H(x)\big] + \delta v_2,$$

where $H(x) = \arg\max_y u(x, x, y)$ is the government's opportunistic one period best policy in response to x. Here v_2 is the continuation value awarded to the government if it fails to confirm the private sector's expectation that $\eta = y$ this period. Because it receives the same continuation value v_2 for *any* deviation from y, if it does deviate, the government will choose the most rewarding action, which is to set $\eta = H(x)$.

Inequalities (23.4.1) define a Bellman equation that maps a *pair* of continuation values (v_1, v_2) into a value v and first-period outcomes (x, y). Figure 23.4.1 illustrates this mapping for the infinitely repeated version of the taxation example. Given a pair (v_1, v_2), the solid curve depicts v in equation (23.4.1a), and the dashed curve describes the right side of the incentive constraint (23.4.1b). The region in which the solid curve is above the dashed curve identifies tax rates and competitive equilibria that satisfy (23.4.1b) at the prescribed continuation values (v_1, v_2). As can be seen, when $\delta = .8$, tax rates below 18 percent cannot be sustained for the particular (v_1, v_2) pair we have chosen.

APS calculate the *set* of equilibrium values by iterating on the mapping defined by the Bellman equation (23.4.1). Let W be a set of candidate continuation values. As we vary $(v_1, v_2) \in W \times W$, the Bellman equation traces out a *set* of values, say, $v \in B(W)$. Thus, the Bellman equation maps *sets* of continuation values W (from which we can draw a *pair* of continuation values $(v_1, v_2) \in W \times W$) into sets of current values $v \in$B(W). To qualify as SPE values, we require that $W \subset B(W)$, i.e., the *continuation* values drawn from W must themselves be *values* that are in turn supported by continuation values drawn from the same set W. APS reason that the largest set for which $W = B(W)$ *is* the set of all SPE values. APS show how iterations on the Bellman equation can determine the set of equilibrium values, provided that one starts with a big enough but bounded initial set of candidate continuation values. Furthermore, after that set of values has been found, APS show how to find a strategy that attains any equilibrium value in the set. The remainder of the chapter describes details of APS's formulation as applied in our setting. We shall see why APS want to get their hands on the entire set of equilibrium values.

Why do we call it 'dynamic programming squared'? There are two reasons.

1. The construction works by mapping *two* continuation values into one, in contrast to ordinary dynamic programming, which maps *one* continuation value tomorrow into one value function today.

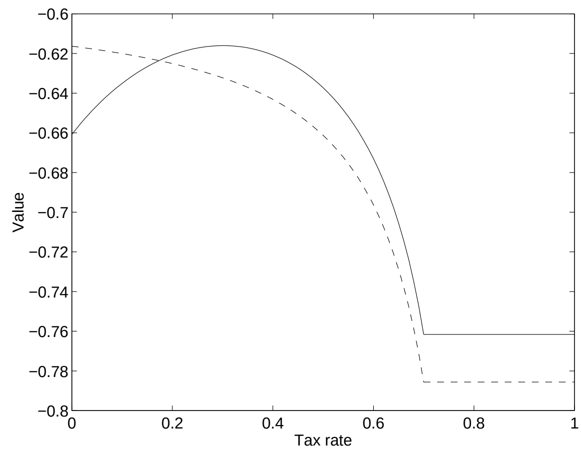

Figure 23.4.1: Mapping of continuation values (v_1, v_2) into values v in the infinitely repeated version of the taxation example. The solid curve depicts $v = (1-\delta)u[\ell(\tau), \ell(\tau), \tau] + \delta v_1$. The dashed curve is the right side of the incentive constraint, $v \geq (1-\delta)u\{\ell(\tau), \ell(\tau), H[\ell(\tau)]\} + \delta v_2$, where H is the government's best response function. The part of the solid curve that is above the dashed curve shows competitive equilibrium values that are sustainable for continuation values (v_1, v_2). The parameterization is $\alpha = 0.3$ and $\delta = 0.8$, and the continuation values are set as $(v_1, v_2) = (-0.6, -0.63)$.

2. A continuation value plays a double role, one as a promised value that summarizes expectations of the rewards associated with future outcomes, another as a state variable that summarizes the history of past outcomes. In the present setting, a subgame perfect equilibrium strategy profile can be represented recursively in terms of an initial value $v_1 \in I\!R$ and the following 3-tuple of functions:

$$x_t = z^g(v_t)$$
$$y_t = z^g(v_t)$$
$$v_{t+1} = \mathcal{V}(v_t, x_t, y_t)$$

the first two of which maps a promised value into a private sector decision and a government action, while the third maps a promised value and an action pair into a promised value to carry into tomorrow. By iterating these equations, we can deduce that the the triple of functions $z^h, z^g, \mathcal{V}$ induces a strategy profile that maps histories of outcomes into sequences of outcomes. The capacity to represent a subgame perfect equilibrium recursively affords immense simplifications in terms of the number of functions we must carry.

23.5. The infinitely repeated economy

Consider repeating our one-period economy forever. At each $t \geq 1$, each household chooses $\xi_t \in X$, with the result that the average $x_t \in X$; the government chooses $y_t \in Y$. We use the notation $(\vec{x}, \vec{y}) = \{(x_t, y_t)\}_{t=1}^{\infty}$, $\vec{\xi} = \{\xi_t\}_{t=1}^{\infty}$. To denote a *history* of (x_t, y_t) up to t, we use the notation $x^t = \{x_s\}_{s=1}^{t}$, $y^t = \{y_s\}_{s=1}^{t}$. These histories live in the spaces X^t and Y^t, respectively, where $X^t = X \times \cdots \times X$, the Cartesian product of X taken t times, and Y^t is the Cartesian product of Y taken t times.[9]

For the repeated economy, the government evaluates paths $(\vec{x}, \vec{y})$ according to

$$V_g(\vec{x}, \vec{y}) = \frac{(1-\delta)}{\delta} \sum_{t=1}^{\infty} \delta^t \, r(x_t, y_t), \tag{23.5.1}$$

where $r(x_t, y_t) \equiv u(x_t, x_t, y_t)$ and $0 < \delta < 1$.[10] A *pure strategy* is defined as a sequence of functions, the tth element of which maps the history (x^{t-1}, y^{t-1}) observed at the beginning of t into an action at t. In particular, for the aggregate of households, a strategy is a sequence $\sigma^h = \{\sigma_t^h\}_{t=1}^{\infty}$ such that

$$\begin{aligned} &\sigma_1^h \in X \\ &\sigma_t^h : X^{t-1} \times Y^{t-1} \to X \qquad \text{for each} \quad t \geq 2\ . \end{aligned}$$

[9] Marco Bassetto's work (2002, 2003) shows that this specification, which is common in the literature, excludes some interesting applications. In particular, it rules out contexts in which the set of time t actions available to the government is influenced by past actions taken by households. Such excluded examples prevail, for example, in the fiscal theory of the price level. To construct sustainable plans in those interesting environments, Bassetto (2002, 2003) refines the notion of sustainability to include a more complete theory of the government's behavior off an equilibrium path.

[10] Note that we have not defined the government's payoff when $\xi_t \neq x_t$. See footnote 7.

Similarly, for the government, a strategy $\sigma^g = \{\sigma_t^g\}_{t=1}^\infty$ is a sequence such that

$$\begin{aligned} &\sigma_1^g \in Y \\ &\sigma_t^g : X^{t-1} \times Y^{t-1} \to Y \qquad \text{for each} \quad t \geq 2. \end{aligned}$$

We call $\sigma = (\sigma^h, \sigma^g)$ a *strategy profile.* We let $\sigma_t = (\sigma_t^h, \sigma_t^g)$ be the tth component of the strategy profile.

23.5.1. A strategy profile implies a history and a value

APS begin with the insight that a strategy profile $\sigma = (\sigma^g, \sigma^h)$ evidently recursively generates a trajectory of outcomes $\{[x(\sigma)_t, y(\sigma)_t]\}_{t=1}^\infty$:

$$\begin{aligned} \left[x(\sigma)_1, y(\sigma)_1\right] &= (\sigma_1^h, \sigma_1^g) \\ \left[x(\sigma)_t,\, y(\sigma)_t\right] &= \sigma_t\left[x(\sigma)^{t-1},\; y(\sigma)^{t-1}\right]. \end{aligned}$$

Therefore, a strategy profile also generates a pair of values for the government and the representative private agent. In particular, the value for the government of a strategy profile $\sigma = (\sigma^h, \sigma^g)$ is the value of the trajectory that it generates

$$V_g(\sigma) = V_g\left[\vec{x}(\sigma),\, \vec{y}(\sigma)\right].$$

23.5.2. Recursive formulation

A key step toward APS's recursive formulation comes from defining *continuation stategies* and their associated *continuation values*. Since the value of a path $(\vec{x}, \vec{y})$ in equation (23.5.1) is additively separable in its one-period returns, we can express the value recursively in terms of a one-period economy and a continuation economy. In particular, the value to the government of an outcome sequence $(\vec{x}, \vec{y})$ can be represented

$$V_g(\vec{x}, \vec{y}) = (1-\delta)\, r(x_1, y_1) + \delta V_g\big(\{x_t\}_{t=2}^\infty\,,\ \{y_t\}_{t=2}^\infty\big) \tag{23.5.2}$$

and the value for a household can also be represented recursively. Notice how a strategy profile σ induces a strategy profile for the continuation economy, as follows. Let $\sigma|_{(x^t,y^t)}$ denote the strategy profile for a continuation economy whose first period is $t+1$ and that is initiated after history (x^t, y^t) has been observed; here $(\sigma|_{(x^t,y^t)})_s$ is the sth component of $(\sigma|_{(x^t,y^t)})$, which for $s \geq 2$ is a function that maps $X^{s-1} \times Y^{s-1}$ into $X \times Y$, and for $s = 1$ is a point in $X \times Y$. Thus, after a first-period outcome pair (x_1, y_1), strategy σ induces the continuation strategy

$$\begin{aligned}(\sigma|_{(x_1,y_1)})_{s+1}\,(\nu^s, \eta^s) &= \sigma_{s+2}\,(x_1, \nu_1, \ldots, \nu_s, y_1, \eta_1, \ldots, \eta_s)\\ &\quad \text{for all } \; (\nu^s, \eta^s) \in X^s \times Y^s \; , \quad \forall s \geq 0.\end{aligned}$$

It might be helpful to write out a few terms for $s = 0, 1, \ldots$:

$$\begin{aligned}(\sigma|_{(x_1,y_1)})_1 &= \sigma_2(x_1, y_1) = (\nu_1, \eta_1)\\ (\sigma|_{(x_1,y_1)})_2(\nu_1, \eta_1) &= \sigma_3(x_1, \nu_1, y_1, \eta_1) = (\nu_2, \eta_2)\\ (\sigma|_{(x_1,y_1)})_3(\nu_1, \nu_2, \eta_1, \eta_2) &= \sigma_4(x_1, \nu_1, \nu_2, y_1, \eta_1, \eta_2) = (\nu_3, \eta_3).\end{aligned}$$

More generally, define the continuation strategy

$$\begin{aligned}(\sigma|_{(x^t,y^t)})_1 =& \sigma_{t+1}(x^t, y^t)\\ (\sigma|_{(x^t,y^t)})_{s+1}(\nu^s, \eta^s) =& \sigma_{t+s+1}\,(x_1, \ldots, x_t, \nu_1, \ldots, \nu_s; y_1, \ldots, y_t, \eta_1, \ldots, \eta_s)\\ & \text{for all } \; s \geq 1 \quad \text{and all} \quad (\nu^s, \eta^s) \in X^s \times Y^s \; .\end{aligned}$$

Here $(\sigma|_{(x^t,y^t)})_{s+1}$ (ν^s, η^s) is the induced strategy pair to apply in the $(s+1)$th period of the continuation economy. We attain this strategy by shifting the original strategy forward t periods and evaluating it at history $(x_1, \ldots, x_t, \nu_1, \ldots, \nu_s; y_1, \ldots, y_t, \eta_1, \ldots, \eta_s)$ for the *original* economy.

In terms of the continuation strategy $\sigma|_{(x_1,y_1)}$, from equation (23.5.2) we know that $V_g(\sigma)$ can be represented as

$$V_g(\sigma) = (1-\delta) r(x_1, y_1) + \delta V_g(\sigma|_{(x_1,y_1)}). \tag{23.5.3}$$

Representation (23.5.3) decomposes the value to the government of strategy profile σ into a one-period return and the continuation value $V_g(\sigma|_{(x_1,y_1)})$ associated with the continuation strategy $\sigma|_{(x_1,y_1)}$.

Any sequence $(\vec{x}, \vec{y})$ in equation (23.5.2) or any strategy profile σ in equation (23.5.3) can be assigned a value. We want a notion of an *equilibrium* strategy profile.

23.6. Subgame perfect equilibrium (SPE)

DEFINITION 6: A strategy profile $\sigma = (\sigma^h, \sigma^g)$ is a *subgame perfect equilibrium* (SPE) of the infinitely repeated economy if for each $t \geq 1$ and each history $(x^{t-1},\, y^{t-1}) \in X^{t-1} \times Y^{t-1}$

(a) The private sector outcome $x_t = \sigma_t^h\,(x^{t-1}, y^{t-1})$ is consistent with competitive equilibrium when $y_t = \sigma_t^g\,(x^{t-1}, y^{t-1})$;

(b) For each possible government action $\eta \in Y$

$$(1-\delta) r(x_t, y_t) + \delta V_g(\sigma|_{(x^t, y^t)}) \geq (1-\delta)\, r(x_t, \eta) + \delta V_g(\sigma|_{(x^t; y^{t-1}, \eta)}).$$

Requirement a says two things. It attributes a theory of forecasting government behavior to members of the public, in particular, that they use the time t component σ_t^g of the government's strategy and information available at the end of period $t-1$ to forecast the government's behavior at t. Condition a also asserts that a competitive equilibrium appropriate to the public's forecast value for y_t is the outcome at time t. Requirement b says that at each point in time and following each history, the government has no incentive to deviate from the first-period action called for by its strategy σ^g; that is, the government always wants to behave as the public expects. Notice how in condition b, the government *contemplates* setting its time t choice η_t at something other than the value forecast by the public, but confronts consequences that deter it from choosing an η_t that fails to confirm the public's expectations of it.

In section 23.14 we'll discuss the following question: who *chooses* σ^g, the government or the public? This question arises naturally because σ^g is *both* the government's sequence of policy functions *and* the private sector's rule for forecasting government behavior. Condition b of definition 6 says that the government chooses to confirm the public's forecasts.

Definition 6 implies that for each $t \geq 2$ and each $(x^{t-1}, y^{t-1}) \in X^{t-1} \times Y^{t-1}$, the continuation strategy $\sigma|_{(x^{t-1}, y^{t-1})}$ is itself an SPE. We state this formally for $t = 2$.

PROPOSITION 1: Assume that σ is an SPE. Then for all $(\nu, \eta) \in X \times Y$, $\sigma|_{(\nu,\eta)}$ is an SPE.

PROOF: Write out requirements a and b that Definition 6 asserts that the continuation strategy $\sigma|_{(\nu,\eta)}$ must satisfy to qualify as an SPE. In particular, for all $s \geq 1$ and for all $(x^{s-1}, y^{s-1}) \in X^{s-1} \times Y^{s-1}$, we require

$$(x_s, y_s) \in C, \tag{23.6.1}$$

where $x_s = \sigma^h|_{(\nu,\eta)}(x^{s-1}, y^{s-1}), y_s = \sigma^g|_{(\nu,\eta)}(x^{s-1}, y^{s-1})$. We also require that for all $\tilde{\eta} \in Y$,

$$(1-\delta) r(x_s, y_s) + \delta V_g(\sigma|_{(\eta,x^s;\nu,y^s)}) \geq (1-\delta) r(x_s, \tilde{\eta}) + \delta V_g(\sigma|_{(\nu,x^s;\eta,y^{s-1},\tilde{\eta})}) \tag{23.6.2}$$

Notice that requirements a and b of Definition 6 for $t = 2, 3, \ldots$ imply expressions (23.6.1) and (23.6.2) for $s = 1, 2, \ldots$. ∎

The statement that $\sigma|_{(\nu,\eta)}$ is an SPE for all $(\nu, \eta) \in X \times Y$ ensures that σ is *almost* an SPE. If we know that $\sigma|_{(\nu,\eta)}$ is an SPE for all $(\nu, \eta) \in (X \times Y)$, we must only add two requirements to ensure that σ is an SPE: first, that the $t = 1$ outcome pair (x_1, y_1) is a competitive equilibrium, and second, that the government's choice of y_1 satisfies the time 1 version of the incentive constraint b in Definition 6.

This reasoning leads to the following lemma that is at the heart of the APS analysis:

LEMMA: Consider a strategy profile σ, and let the associated first-period outcome be given by $x = \sigma_1^h, y = \sigma_1^g$. The profile σ is an SPE if and only if

(1) for each $(\nu, \eta) \in X \times Y, \quad \sigma|_{(\nu,\eta)}$ is an SPE;

(2) (x, y) is a competitive equilibrium;

(3) $\forall\, \eta \in Y,\ (1-\delta)\, r(x, y) + \delta\, V_g(\sigma|_{(x,y)}) \geq (1-\delta)\, r(x, \eta) + \delta V_g(\sigma|_{(x,\eta)})$.

PROOF: First, prove the "if" part. Property a of the lemma and properties (23.6.1) and (23.6.2) of Proposition 1 show that requirements a and b of Definition 6 are satisfied for $t \geq 2$. Properties (2) and (3) of the lemma imply that requirements a and b of Definition 6 hold for $t = 1$.

Second, prove the "only if" part. Part (1) of the lemma follows from Proposition 1. Parts (2) and (3) of the lemma follow from requirements a and b of Definition 6 for $t = 1$. ∎

The lemma is very important because it characterizes SPEs in terms of a first-period competitive equilibrium outcome pair (x, y), and a *pair* of continuation values: a value $V_g(\sigma|_{(x,y)})$ to be awarded to the government next period if it adheres to the y component of the first-period pair (x, y), and a value $V_g(\sigma|_{(x,\eta)})$, $\eta \neq y$, to be awarded to the government if it deviates from the expected y component. Each of these values has to be selected from a set of values $V_g(\sigma)$ that are associated with some SPE σ.

23.7. Examples of SPE

23.7.1. Infinite repetition of one-period Nash equilibrium

It is easy to verify that the following strategy profile $\sigma^N = (\sigma^h, \sigma^g)$ forms an SPE: $\sigma_1^h = x^N, \sigma_1^g = y^N$ and for $t \geq 2$

$$\begin{aligned} \sigma_t^h &= x^N \qquad \forall\ t\ , \quad \forall\ (x^{t-1}, y^{t-1}); \\ \sigma_t^g &= y^N \qquad \forall\ t\ , \quad \forall\ (x^{t-1}, y^{t-1}). \end{aligned}$$

These strategies instruct the households and the government to choose the static Nash equilibrium outcomes for all periods for all histories. Evidently, for these strategies, $V_g(\sigma^N) = v^N = r(x^N, y^N)$. Furthermore, for these strategies the continuation value $V_g(\sigma|_{(x^t; y^{t-1}, \eta)}) = v^N$ for all outcomes $\eta \in Y$. These strategies satisfy requirement a of Definition 6 because (x^N, y^N) is a competitive equilibrium. The strategies satisfy requirement b because $r(x^N, y^N) = \max_{y \in Y} r(x^N, y)$ and because the continuation value $V_g(\sigma) = v^N$ is independent of the action chosen by the government in the first period. In this SPE, $\sigma_t^N = \{\sigma_t^h, \sigma_t^g\} = (x^N, y^N)$ for all t and for all (x^{t-1}, y^{t-1}), and the value $V_g(\sigma^N)$ and the continuation values for each history (x^t, y^t), $V_g(\sigma^N|_{(x^t, y^t)})$, all equal v^N.

It is useful to think about this SPE in terms of the lemma. To verify that σ^N is a SPE, we work with the first-period outcome pair (x^N, y^N) and the pair of values $V_g(\sigma|_{(x^N, y^N)}) = v^N, V_g(\sigma|_{(x,\eta)}) = v^N$, where $v^N = r(x^N, y^N)$. With these settings, we can verify that (x^N, y^N) and v^N satisfy requirements (1), (2), and (3) of the lemma.

23.7.2. Supporting better outcomes with trigger strategies

The public can have a system of expectations about the government's behavior that induces the government to choose a better-than-Nash outcome $(\tilde{x}, \tilde{y}) \in C$. Thus, suppose that the public expects that as long as the government chooses $\tilde{y}$, it will continue to do so in the future, but that once the government deviates from this choice, the public expects that the government will choose y^N thereafter, prompting the public (or what we can call "the market") to react with $x^N = h(y^N)$. This system of expectations confronts the government with the prospect of being "punished by the market's expectations" if it chooses to deviate from $\tilde{y}$.

To formalize this idea, we shall use the SPE σ^N as a continuation strategy and the value v^N as a continuation value on the right side of part (b) of Definition 6 of an SPE (for $\eta \neq y_t$); then by working backward one step, we shall try to construct *another* SPE $\tilde{\sigma}$ with first-period outcome $(\tilde{x}, \tilde{y}) \neq (x^N, y^N)$. In particular, for our new SPE $\tilde{\sigma}$ we propose to set

$$\begin{aligned} \tilde{\sigma}_1 &= (\tilde{x}, \tilde{y}) \\ \tilde{\sigma}|_{(x,y)} &= \begin{cases} \tilde{\sigma} & \text{if } (x,y) = (\tilde{x}, \tilde{y}) \\ \sigma^N & \text{if } (x,y) \neq (\tilde{x}, \tilde{y}) \end{cases} \end{aligned} \tag{23.7.1}$$

where $(\tilde{x}, \tilde{y})$ is a competitive equilibrium that satisfies the following particular case of part b of Definition 6:

$$\tilde{v} = (1-\delta)\, r(\tilde{x}, \tilde{y}) + \delta \tilde{v} \geq (1-\delta)\, r(\tilde{x}, \eta) + \delta v^N , \tag{23.7.2}$$

for all $\eta \in Y$. Inequality (23.7.2) is equivalent with

$$\max_{\eta \in Y} r(\tilde{x}, \eta) - r(\tilde{x}, \tilde{y}) \leq \frac{\delta}{1-\delta} \, (\tilde{v} - v^N). \tag{23.7.3}$$

For any $(\tilde{x}, \tilde{y}) \in C$ that satisfies expression (23.7.3) with $\tilde{v} = r(\tilde{x}, \tilde{y})$, strategy (23.7.1) is an SPE with value $\tilde{v}$.

If $(\tilde{x}, \tilde{y}) = (x^R, y^R)$ satisfies inequality (23.7.3) with $\tilde{v} = r(x^R, y^R)$, then repetition of the Ramsey outcome (x^R, y^R) is supportable by a subgame perfect equilibrium of the form (23.7.1).

This construction uses the following objects:

1. A proposed first-period competitive equilibrium $(\tilde{x}, \tilde{y}) \in C$;

2. An SPE σ^2 with value $V_g(\sigma^2)$ that is used as the continuation strategy in the event that the first-period outcome does not equal $(\tilde{x}, \tilde{y})$, so that $\tilde{\sigma}|_{(x,y)} = \sigma^2$, if $(x,y) \neq (\tilde{x}, \tilde{y})$. In the example, $\sigma^2 = \sigma^N$ and $V_g(\sigma^2) = v^N$.
3. An SPE σ^1, with value $V_g(\sigma^1)$, used to define the continuation value to be assigned after first-period outcome $(\tilde{x}, \tilde{y})$; and an associated continuation strategy $\tilde{\sigma}|_{(\tilde{x},\tilde{y})} = \sigma^1$. In the example, $\sigma^1 = \tilde{\sigma}$, which is defined recursively (and self-referentially) via equation (23.7.1).
4. A candidate for a new SPE $\tilde{\sigma}$ and a corresponding value $V_g(\tilde{\sigma})$. In the example, $V_g(\tilde{\sigma}) = r(\tilde{x}, \tilde{y})$.

In the example, objects 3 and 4 are equated.

Note how we have used the lemma in verifying that $\tilde{\sigma}$ is an SPE. We start with the SPE σ^N with associated value v^N. We guess a first-period outcome pair $(\tilde{x}, \tilde{y})$ and a value $\tilde{v}$ for a new SPE, where $\tilde{v} = r(\tilde{x}, \tilde{y})$. Then we verify requirements (2) and (3) of the lemma with $(\tilde{v}, v^N)$ as continuation values and $(\tilde{x}, \tilde{y})$ as first-period outcomes.

23.7.3. When reversion to Nash is not bad enough

For discount factors δ sufficiently close to one, it is typically possible to support repetition of the Ramsey outcome (x^R, y^R) with a section 23.7.2 trigger strategy of form (23.7.1). This finding conforms with a version of the folk theorem about repeated games. However, there exist discount factors δ so small that infinite reversion to repetition of the one-period Nash outcome is not a bad enough consequence to support repetition of Ramsey. In that case, anticipating that it will revert to repetition of Nash after a deviation can at best support a value for the government that is less than that associated with repetition of Ramsey outcome, although perhaps better than repetition of the Nash outcome.

It is natural at this point to ask whether in this circumstance there is a better SPE? To support something *better* evidently requires finding an SPE that has a value *worse* than that associated with repetition of the one-period Nash outcome. This kind of reasoning directed APS to find the set of values associated with *all* SPEs. Following APS, we shall soon see that the best and worst equilibrium values are linked.

23.8. Values of all SPEs

The role played by the lemma in analyzing our two examples hints at the central role that it plays in the methods that APS developed for describing and computing values for *all* the subgame perfect equilibria. APS build on the way that the lemma characterizes SPE values in terms of a first-period competitive equilibrium outcome, along with a pair of continuation values, each element of which is itself a value associated with some SPE. The lemma directs APS's attention away from a *set of strategy profiles* σ and toward a *set of values* $V_g(\sigma)$ associated with those profiles. They define the set V of values associated with subgame perfect equilibria:

$$V = \{V_g\,(\sigma) \mid \sigma \ \text{ is an SPE}\}.$$

Evidently, $V \subset I\!\!R$. From the lemma, for a given competitive equilibrium $(x, y) \in C$, there exists an SPE σ for which $x = \sigma_1^h, y = \sigma_1^g$ if and only if there exist two values $(v_1, v_2) \in V \times V$ such that

$$(1-\delta)\,r(x,y) + \delta v_1 \geq (1-\delta)\,r(x,\eta) + \delta v_2 \quad \forall\, \eta \in Y. \tag{23.8.1}$$

Let σ^1 and σ^2 be subgame perfect equilibria for which $v_1 = V_g(\sigma^1)$, $v_2 = V_g(\sigma^2)$. The SPE σ that supports $(x, y) = (\sigma_1^h, \sigma_1^g)$ is completed by specifying the continuation strategies $\sigma|_{(x,y)} = \sigma^1$ and $\sigma|_{(\nu,\eta)} = \sigma^2$ if $(\nu, \eta) \neq (x, y)$.

This construction uses two continuation values $(v_1, v_2) \in V \times V$ to create an SPE σ with value $v \in V$ given by

$$v = (1-\delta)\,r(x,y) + \delta v_1 \ .$$

Thus, the construction maps *pairs* of continuation values (v_1, v_2) into a strategy profile σ with first-period competitive equilibrium outcome (x, y) and a value $v = V_g(\sigma)$.

APS characterize subgame perfect equilibria by studying a mapping from pairs of continuation values $(v_1, v_2) \in V \times V$ into values $v \in V$. They use the following definitions:

Definition 7: Let $W \subset I\!\!R$. A 4-tuple (x, y, w_1, w_2) is said to be *admissible with respect to* W if $(x, y) \in C$, $(w_1, w_2) \in W \times W$, and

$$(1-\delta)\,r(x,y) + \delta w_1 \geq (1-\delta)\,r(x,\eta) + \delta w_2\,, \quad \forall\, \eta \in Y. \tag{23.8.2}$$

Notice that when $W \subset V$, the admissible 4-tuple (x, y, w_1, w_2) determines an SPE with strategy profile

$$\sigma_1 = (x, y),\ \sigma|_{(x,y)} = \sigma^1,\ \sigma|_{(\nu,\eta)} = \sigma^2 \ \text{ for } (\nu, \eta) \neq (x, y)$$

where σ_1 is a continuation strategy that yields value $w_1 = V_g(\sigma^1)$ and σ_2 is a strategy that yields continuation value $w_2 = V_g(\sigma^2)$. The value of the SPE is $V_g(\sigma) = w = (1-\delta)\, r(x, y) + \delta w_1$.

We want to find the set V.

23.8.1. The basic idea of dynamic programming squared

In Definition 7, W serves as a set of candidate continuation values. The idea is to pick an $(x, y) \in C$, then to check whether you can find $(w_1, w_2) \in W \times W$ that would make the government want to adhere to the y component if w_1 and w_2 could be used as continuation values for adhering to and deviating from y, respectively. If the answer is yes, we say that the 4-tuple (x, y, w_1, w_2) is "admissible with respect to W". Because we have verified that the incentive constraints are satisfied, a yes answer allows us to calculate the *value* (i.e., the left side of (23.8.2)) that can be supported with w_1, w_2 as continuation values. Thus, the idea is to use (23.8.2) to define a mapping from values tomorrow to values today, like that used in dynamic programming. In the next section, we'll define $B(W)$ as the set of possible values attained with admissible pairs of continuation values drawn from $W \times W$. Then we'll view B as an operator mapping sets of continuation values W into sets of values $B(W)$. This operator is the counterpart to the T operator associated with ordinary dynamic programming.

To pursue this analogy, recall the Bellman equation associated with the McCall model of chapter 6:

$$Q = \int \max\left\{\frac{w}{1-\beta}, c + \beta Q\right\} d\,F(w).$$

Here $Q \in I\!R$ is the expected discounted value of an unemployed worker's income *before* he has drawn a wage offer. The right side defines an operator $T(Q)$, so that the Bellman equation is

$$Q = T(Q). \tag{23.8.3}$$

This equation can be solved by iterating to convergence starting from any initial Q.

Just as the right side of (23.8.3) takes a candidate continuation value Q for tomorrow and maps it into a value $T(Q)$ for today, APS define a mapping $B(W)$ that, by considering only admissible 4-tuples, maps a *set* of values W tomorrow into a new *set* $B(W)$ of values today. Thus, APS use admissible 4-tuples to map candidate continuation values tomorrow into new candidate values today. In the next section, we'll iterate to convergence on $B(W)$, but as we'll see, it won't work to start from just any initial set W. We have to start from a big enough set.

DEFINITION 8: For each set $W \subset I\!R$, let $B(W)$ be the set of possible values $w = (1-\delta)\, r(x,y) + \delta w_1$ associated with admissible tuples (x, y, w_1, w_2).

Think of W as a set of potential continuation values and $B(W)$ as the set of values that they support. From the definition of admissibility it immediately follows that the operator B is *monotone*.

PROPERTY (monotonicity of B): If $W \subseteq W' \subseteq R$, then $B(W) \subseteq B(W')$.

PROOF: It can be verified directly from the definition of admissible 4-tuples that if $w \in B(W)$, then $w \in B(W')$: simply use the (w_1, w_2) pair that supports $w \in B(W)$ to support $w \in B(W')$. ∎

It can also be verified that $B(\cdot)$ maps compact sets W into compact sets $B(W)$.

The self-supporting character of subgame perfect equilibria is referred to in the following definition:

DEFINITION 9: The set W is said to be *self-generating* if $W \subseteq B(W)$.

Thus, a set of continuation values W is said to be self-generating if it is contained in the set of *values* $B(W)$ that are generated by pairs of continuation values selected from W. This description makes us suspect that if a set of values is self-generating, it must be a set of SPE values. Indeed, notice that by virtue of the lemma, the set V of SPE values $V_g(\sigma)$ is self-generating. Thus, we can write $V \subseteq B(V)$. APS show that V is the *largest* self-generating set. The key to showing this point is the following theorem:[11]

[11] The *unbounded* set $I\!R$ (the extended real line) is self-generating but not meaningful. It is self-generating because any value $v \in I\!R$ can be supported if there are no limits on the

Theorem 1 (Self-Generation): If $W \subset I\!\!R$ is bounded and self-generating, then $B(W) \subseteq V$.

The proof is based on "forward induction" and proceeds by taking a point $w \in B(W)$ and constructing an SPE with value w.

Proof: Assume $W \subseteq B(W)$. Choose an element $w \in B(W)$ and transform it as follows into a subgame perfect equilibrium:

Step 1. Because $w \in B(W)$, we know that there exist outcomes (x, y) and values w_1 and w_2 that satisfy

$$\begin{aligned} w = (1-\delta)\, r(x,y) + \delta w_1 &\geq (1-\delta)\, r(x,\eta) + \delta w_2 \quad \forall \eta \in Y \\ (x,y) &\in C \\ w_1, w_2 &\in W \times W. \end{aligned}$$

Set $\sigma_1 = (x, y)$.

Step 2. Since $w_1 \in W \subseteq B(W)$, there exist outcomes $(\tilde{x}, \tilde{y})$ and values $(\tilde{w}_1, \tilde{w}_2) \in W$ that satisfy

$$\begin{aligned} w_1 = (1-\delta)\, r(\tilde{x},\tilde{y}) + \delta \tilde{w}_1 &\geq (1-\delta)\, r(\tilde{x},\eta) + \delta \tilde{w}_2, \quad \forall\ \eta \in Y \\ (\tilde{x},\tilde{y}) &\in C. \end{aligned}$$

Set the first-period outcome in period 2 (the outcome to occur *given* that y was chosen in period 1) equal to $(\tilde{x}, \tilde{y})$; that is, set $(\sigma|_{(x,y)})_1 = (\tilde{x}, \tilde{y})$.

Continuing in this way, for each $w \in B(W)$, we can create a sequence of continuation values $w_1, \tilde{w}_1, \tilde{\tilde{w}}_1, \ldots$ and a corresponding sequence of first-period outcomes (x, y), $(\tilde{x}, \tilde{y})$, $(\tilde{\tilde{x}}, \tilde{\tilde{y}})$.

At each stage in this construction, policies are *unimprovable*, which means that given the continuation values, one-period deviations from the prescribed policies are not optimal. It follows that the strategy profile is optimal. By construction $V_g(\sigma) = w$. ∎

Collecting results, we know that

1. $V \subseteq B(V)$ (by the lemma).
2. If $W \subseteq B(W)$, then $B(W) \subseteq V$ (by self-generation).
3. B is monotone and maps compact sets into compact sets.

continuation values. It is not meaningful because most points in $I\!\!R$ are values that cannot be attained with *any* strategy profile.

Facts 1 and 2 imply that $V = B(V)$, so that the set of equilibrium values is a "fixed point" of B, in particular, the *largest* bounded fixed point.

Monotonicity of B and the fact that it maps compact sets into compact sets provides an algorithm for computing the set V, namely, to start with a set W_0 for which $V \subseteq B(W_0) \subseteq W_0$, and to iterate to convergence on B. In more detail, we use the following steps:

1. Start with a set $W_0 = [\underline{w}_0, \overline{w}_0]$ that we know is bigger than V, and for which $B(W_0) \subseteq W_0$. It will always work to set $\overline{w}_0 = \max_{(x,y)\in C} r(x,y)$, $\underline{w}_0 = \min_{(x,y)\in C} r(x,y)$.

2. Compute the boundaries of the set $B(W_0) = [\underline{w}_1, \overline{w}_1]$. The value $\overline{w}_1$ solves the problem

$$\overline{w}_1 = \max_{(x,y)\in C} (1-\delta)\, r(x,y) + \delta \overline{w}_0$$

subject to

$$(1-\delta)\, r(x,y) + \delta \overline{w}_0 \geq (1-\delta)\, r(x,\eta) + \delta \underline{w}_0 \quad \text{for all} \quad \eta \in Y.$$

The value $\underline{w}_1$ solves the problem

$$\underline{w}_1 = \min_{(x,y)\in C;\,(w_1,w_2)\in[\underline{w}_0,\overline{w}_0]^2} (1-\delta)\, r(x,y) + \delta w_1$$

subject to

$$(1-\delta)\, r(x,y) + \delta w_1 \geq (1-\delta)\, r(x,\eta) + \delta w_2 \quad \forall\ \eta \in Y.$$

With $(\underline{w}_0, \overline{w}_0)$ chosen as before, it will be true that $B(W_0) \subseteq W_0$.

3. Having constructed $W_1 = B(W_0) \subseteq W_0$, continue to iterate, producing a decreasing sequence of compact sets $W_{j+1} = B(W_j) \subseteq W_j$. Iterate until the sets converge.

Later, we'll present a direct way to compute the best and worst SPE values, one that evades having to iterate to convergence on the B operator.

23.9. Self-enforcing SPE

A subgame perfect equilibrium with a *worst* value $v \in V$ has the remarkable property that it is "self-enforcing." We use the following definition:

DEFINITION 10: A subgame perfect equilibrium σ with first-period outcome $(\tilde{x}, \tilde{y})$ is said to be *self-enforcing* if

$$\sigma|_{(x,y)} = \sigma \qquad \text{if } (x,y) \neq (\tilde{x}, \tilde{y}). \tag{23.9.1}$$

A strategy profile satisfying equation (23.9.1) is called self-enforcing because after a one-shot deviation the consequence is simply to restart the equilibrium.

Recall our earlier characterization of a competitive equilibrium as a pair $(h(y), y)$, where $x = h(y)$ is the mapping from the government's action y to the private sector's equilibrium response. The value $\underline{v}$ associated with the worst subgame perfect equilibrium σ satisfies

$$\underline{v} = \min_{y,v} \left\{ (1-\delta)\, r(h(y), y) + \delta v \right\} \tag{23.9.2}$$

where the minimization is subject to $y \in Y$, $v \in V$, and the incentive constraint

$$(1-\delta)\, r(h(y), y) + \delta v \geq (1-\delta)\, r(h(y), \eta) + \delta \underline{v} \quad \text{for all } \eta \in Y. \tag{23.9.3}$$

Let $\tilde{v}$ be a continuation value that attains the right side of equation (23.9.2), and let $\sigma_{\tilde{v}}$ be a subgame perfect equilibrium that supports continuation value $\tilde{v}$. Let $(\tilde{x}, \tilde{y})$ be the first-period outcome that attains the right side of equation (23.9.2). Thus, $\underline{v} = (1-\delta) r(\tilde{x}, \tilde{y}) + \delta \tilde{v}$. Since $\underline{v}$ is both the continuation value when first-period outcome $(x,y) \neq (\tilde{x}, \tilde{y})$ *and* the value associated with subgame perfect equilibrium σ, it follows that

$$\begin{aligned} \sigma_1 &= (\tilde{x}, \tilde{y}) \\ \sigma|_{(x,y)} &= \begin{cases} \sigma & \text{if } (x,y) \neq (\tilde{x}, \tilde{y}) \\ \sigma_{\tilde{v}} & \text{if } (x,y) = (\tilde{x}, \tilde{y}). \end{cases} \end{aligned} \tag{23.9.4}$$

Because of the double role played by $\underline{v}$, i.e., $\underline{v}$ is both the value of equilibrium σ and the "punishment" continuation value of the right side of the incentive constraint (23.9.3), the equilibrium strategy σ that supports $\underline{v}$ is self-enforcing.[12]

[12] As we show below, the structure of the programming problem, with the double role played by $\underline{v}$, makes it possible to compute the worst value directly.

The preceding argument thus establishes

PROPOSITION 2: A subgame perfect equilibrium σ associated with $\underline{v} = \min\{v : v \in V\}$ is self-enforcing.

23.9.1. The quest for something worse than repetition of Nash outcome

Notice that the first subgame perfect equilibrium that we computed, whose outcome was infinite repetition of the one-period Nash equilibrium, is a self-enforcing equilibrium. However, in general, the infinite repetition of the one-period Nash equilibrium is not the *worst* subgame perfect equilibrium. This fact opens up the possibility that even when reversion to Nash after a deviation is *not* able to support repetition of Ramsey as an SPE, we might still support repetition of the Ramsey outcome by reverting to a SPE with a value worse than that associated with repetition of the Nash outcome whenever the government deviates from an expected one-period choice.

23.10. Recursive strategies

This section emphasizes similarities between credible government policies and the recursive contracts appearing in chapter 20. We will study situations where the household's and the government's strategies have recursive representations. This approach substantially restricts the space of strategies because most history-dependent strategies cannot be represented recursively. Nevertheless, this class of strategies excludes no equilibrium payoffs $v \in V$. We use the following definitions:

DEFINITION 11: Households and the government follow *recursive strategies* if there is a 3-tuple of functions $\phi = (z^h, z^g, \mathcal{V})$ and an initial condition v_1 with the following structure:

$$\begin{aligned} v_1 &\in I\!R \text{ is given} \\ x_t &= z^h(v_t) \\ y_t &= z^g(v_t) \\ v_{t+1} &= \mathcal{V}(v_t, x_t, y_t), \end{aligned} \tag{23.10.1}$$

where v_t is a state variable designed to summarize the history of outcomes before t.

This recursive form of strategies operates much like an autoregression to let time t actions (x_t, y_t) depend on the history $\{y_s, x_s\}_{s=1}^{t-1}$, as mediated through the state variable v_t. Representation (23.10.1) induces history-dependent government policies, and thereby allows for reputation. We shall soon see that beyond its role in keeping track of histories, v_t also summarizes the future.[13]

A strategy (ϕ, v) recursively generates an outcome path expressed as $(\vec{x}, \vec{y}) = (\vec{x}, \vec{y})(\phi, v)$. By substituting the outcome path into equation (23.5.3), we find that (ϕ, v) induces a value for the government, which we write as

$$V^g\big[(\vec{x}, \vec{y})(\phi, v)\big] = (1-\delta)\, r\big[z^h(v), z^g(v)\big] + \delta\, V^g\Big((\vec{x}, \vec{y})\big\{\phi, \mathcal{V}[v, z^h(v), z^g(v)]\big\}\Big). \tag{23.10.2}$$

So far, we have not interpreted the state variable v, except as a particular measure of the history of outcomes. The theory of credible policy ties past and future together by making the state variable v a promised value, an outcome to be expressed

$$v = V^g\big[(\vec{x}, \vec{y})(\phi, v)\big]. \tag{23.10.3}$$

Equations (23.10.1), (23.10.2), and (23.10.3) assert a dual role for v. In equation (23.10.1), v accounts for past outcomes. In equations (23.10.2) and (23.10.3), v looks forward. The state v_t is a discounted future value with which the government enters time t based on past outcomes. Depending on the outcome (x, y) and the entering promised value v, $\mathcal{V}$ updates the promised value with which the government leaves the period. Later we shall struggle with which of two valid interpretations of the government's strategy should be emphasized: something chosen by the government, or a description of a system of public expectations to which the government conforms.

Evidently, we have the following:

---

[13] By iterating equations (23.10.1), we can construct a pair of sequences of functions indexed by $t \geq 1$ $\{Z_t^h(I_t), Z_t^g(I_t)\}$, mapping histories that are augmented by initial conditions $I_t = (\{x_s, y_s\}_{s=1}^{t-1}, v_1)$ into time t actions $(x_t, y_t) \in X \times Y$. Strategies for the repeated economy are a pair of sequences of such functions without the restriction that they have a recursive representation.

DEFINITION 12: Let V be the set of SPE values. A recursive strategy (ϕ, v) in equation (23.10.1) is a *subgame perfect equilibrium* (SPE) if and only if $v \in V$ and

(1) The outcome $x = z^h(v)$ is a competitive equilibrium, given that $y = z^g(v)$.

(2) For each $\eta \in Y$, $\mathcal{V}(v, z^h(v), \eta) \in V$.

(3) For each $\eta \in Y$,

$$\begin{aligned} v &= (1-\delta) r\big[z^h(v), z^g(v)\big] + \delta \mathcal{V}\big[v, z^h(v), z^g(v)\big] \\ &\geq (1-\delta) r\big[z^h(v), \eta\big] + \delta \mathcal{V}\big[v, z^h(v), \eta\big]. \end{aligned} \tag{23.10.4}$$

Condition (1) asserts that the first-period outcome pair (x, y) is a competitive equilibrium. Each member of the private sector forms an expectation about the government's action according to $y_t = z^g(v_t)$, and the "market" responds with a competitive equilibrium x_t,

$$x_t = h(y_t) = h\big[z^g(v_t)\big] \equiv z^h(v_t). \tag{23.10.5}$$

This construction builds in rational expectations, because the private sector knows both the state variable v_t and the government's decision rule z^g.

Besides the first-period outcome (x, y), conditions (2) and (3) associate with a subgame perfect equilibrium three additional objects: a promised value v, a continuation value $v' = \mathcal{V}[v, z^h(v), z^g(v)]$ if the required first-period outcome is observed, and another continuation value $\tilde{v}(\eta) = \mathcal{V}[v, z^h(v), \eta]$ if the required first-period outcome is not observed but rather some pair (x, η). All of the continuation values must themselves be attained as subgame perfect equilibria. In terms of these objects, condition (3) is an incentive constraint inspiring the government to adhere to the equilibrium

$$\begin{aligned} v &= (1-\delta) r(x, y) + \delta v' \\ &\geq (1-\delta) r(x, \eta) + \delta \tilde{v}(\eta), \quad \forall \eta \in Y. \end{aligned}$$

This formula states that the government receives more if it adheres to an action called for by its strategy than if it departs. To ensure that these values constitute "credible expectations," part (2) of Definition 12 requires that the continuation values be values for subgame perfect equilibria. The definition is circular, because members of the same class of objects, namely, equilibrium values v, occur on each side of expression (23.10.4). Circularity comes with recursivity.

One implication of the work of APS (1986, 1990) is that recursive equilibria of form (23.10.1) can attain *all* subgame perfect equilibrium values. As we have seen, APS's innovation was to shift the focus away from the set of equilibrium strategies and toward the set of values V attainable with subgame perfect equilibrium strategies.

23.11. Examples of SPE with recursive strategies

Our two earlier examples of subgame perfect equilibria were already of a recursive nature. But to highlight this property, we recast those SPE in the present notation for recursive strategies. Equilibria are constructed by using a guess-and-verify technique. First, guess $(v_1, z^h, z^g, \mathcal{V})$ in equations (23.10.1), then verify parts (1), (2), and (3) of Definition 12.

The examples parallel the historical development of the theory. (1) The first example is infinite repetition of a one-period Nash outcome, which was Kydland and Prescott's (1977) time-consistent equilibrium. (2) Barro and Gordon (1983a, 1983b) and Stokey (1989) used the value from infinite repetition of the Nash outcome as a continuation value to deter deviation from the Ramsey outcome. For sufficiently high discount factors, the continuation value associated with repetition of the Nash outcome can deter the government from deviating from infinite repetition of the Ramsey outcome. This is not possible for low discount factors. (3) Abreu (1988) and Stokey (1991) showed that Abreu's "stick-and-carrot" strategy induces more severe consequences than repetition of the Nash outcome.

23.11.1. Infinite repetition of Nash outcome

It is easy to construct an equilibrium whose outcome path forever repeats the one-period Nash outcome. Let $v^N = r(x^N, y^N)$. The proposed equilibrium is

$$\begin{aligned} v_1 &= v^N, \\ z^h(v) &= x^N \ \forall\ v, \\ z^g(v) &= y^N \ \forall\ v, \text{ and} \\ \mathcal{V}(v,x,y) &= v^N, \ \forall\ (v,x,y). \end{aligned}$$

Here v^N plays the roles of all *three* values in condition (3) of Definition 12. Conditions (1) and (2) are satisfied by construction, and condition (3) collapses to

$$r(x^N, y^N) \geq r\big[x^N, H(x^N)\big],$$

which is satisfied at equality by the definition of a best response function.

23.11.2. Infinite repetition of a better-than-Nash outcome

Let v^b be a value associated with outcome (x^b, y^b) such that $v^b = r(x^b, y^b) > v^N$, and assume that (x^b, y^b) constitutes a competitive equilibrium. Suppose further that

$$r\big[x^b, H(x^b)\big] - r(x^b, y^b) \leq \frac{\delta}{1-\delta}(v^b - v^N). \tag{23.11.1}$$

The left side is the one-period return to the government from deviating from y^b; it is the gain from deviating. The right side is the difference in present values associated with conforming to the plan versus reverting forever to the Nash equilibrium; it is the cost of deviating. When the inequality is satisfied, the equilibrium presents the government with an incentive not to deviate from y^b. Then an SPE is

$$\begin{aligned} v_1 &= v^b \\ z^h(v) &= \begin{cases} x^b & \text{if } v = v^b; \\ x^N & \text{otherwise;} \end{cases} \\ z^g(v) &= \begin{cases} y^b & \text{if } v = v^b; \\ y^N & \text{otherwise;} \end{cases} \\ \mathcal{V}(v,x,y) &= \begin{cases} v^b & \text{if } (v,x,y) = (v^b, x^b, y^b); \\ v^N & \text{otherwise.} \end{cases} \end{aligned}$$

This strategy specifies outcome (x^b, y^b) and continuation value v^b as long as v^b is the value promised at the beginning of the period. Any deviation from y^b generates continuation value v^N. Inequality (23.11.1) validates condition (3) of Definition 12.

Barro and Gordon (1983a) considered a version of this equilibrium in which inequality (23.11.1) is satisfied with $(v^b, x^b, y^b) = (v^R, x^R, y^R)$. In this case, anticipated reversion to Nash supports the Ramsey outcome forever. When inequality (23.11.1) is *not* satisfied for $(v^b, x^b, y^b) = (v^R, x^R, y^R)$, we can solve for the best SPE value v^b, with associated actions (x^b, y^b), supportable by infinite reversion to Nash from

$$v^b = r(x^b, y^b) = (1-\delta) r\big[x^b, H(x^b)\big] + \delta v^N > v^N. \tag{23.11.2}$$

The payoff from following the strategy equals that from deviating and reverting to Nash. Any value lower than this can be supported, but none higher.

When $v^b < v^R$, Abreu (1988) searched for a way to support something better than v^b. First, one must construct an equilibrium that yields a value *worse* than permanent repetition of the Nash outcome. The expectation of reverting to this equilibrium supports something better than v^b in equation (23.11.2).

Somehow the government must be induced temporarily to take an action $y^\#$ that yields a worse period-by-period return than the Nash outcome, meaning that the government in general would be tempted to deviate. An equilibrium system of expectations has to be constructed that makes the government expect to do better in the future only by conforming to expectations that it temporarily adheres to the bad policy $y^\#$.

23.11.3. Something worse: a stick-and-carrot strategy

To get something worse than repetition of the one-period Nash outcome, Abreu (1988) proposed a "stick-and-carrot punishment." The "stick" part is an outcome $(x^\#, y^\#) \in C$, which relative to (x^N, y^N) is a bad competitive equilibrium from the government's viewpoint. The "carrot" part is the Ramsey outcome (x^R, y^R), which the government attains forever after it has accepted the stick in the first period of its punishment.

We want a continuation value v^* for deviating to support the first-period outcome $(x^\#, y^\#)$ and attain the value[14]

$$\tilde{v} = (1-\delta) r(x^\#, y^\#) + \delta\ v^R \geq (1-\delta) r\big[x^\#, H(x^\#)\big] + \delta\ v^*. \tag{23.11.3}$$

Abreu proposed to set $v^* = \tilde{v}$ so that the continuation value from deviating from the first-period action equals the original value. If the stick part is severe enough, the associated strategy attains a value worse than repetition of Nash. The strategy induces the government to accept the temporarily bad outcome by promising a high continuation value.

An SPE featuring stick-and-carrot punishments that attains $\tilde{v}$ is

$$\begin{aligned} v_1 &= \tilde{v} \\ z^h(v) &= \begin{cases} x^R & \text{if } v = v^R; \\ x^\# & \text{otherwise;} \end{cases} \\ z^g(v) &= \begin{cases} y^R & \text{if } v = v^R; \\ y^\# & \text{otherwise;} \end{cases} \\ \mathcal{V}(v, x, y) &= \begin{cases} v^R & \text{if } (x, y) = [z^h(v), z^g(v)]\ ; \\ \tilde{v} & \text{otherwise.} \end{cases} \end{aligned} \tag{23.11.4}$$

When the government deviates from the bad prescribed first-period action $y^\#$, the consequence is to restart the equilibrium. In other words, the equilibrium is self-enforcing.

14 This is a "one-period stick." The worst SPE can require more than one period of a worse-than-one-period Nash outcome.

23.12. The best and the worst SPE values

The value associated with Abreu's stick-and-carrot strategy might still not be bad enough to deter the government from deviating from repetition of the Ramsey outcome. We are therefore interested in finding the worst SPE value. We now display a pair of simple programming problems to find the best and worst SPE values. APS (1990) showed how to find the entire set of equilibrium values V. In the current setting, their ideas imply the following:

1. The set of equilibrium values V attainable by the government is a compact subset $[\underline{v}, \overline{v}]$ of $[\min_{(x,y)\in C} r(x,y),\ r(x^R, y^R)]$.

2. The worst equilibrium value $\underline{v}$ can be computed from a simple programming problem.

3. Given the worst equilibrium value $\underline{v}$, the best equilibrium value $\overline{v}$ can be computed from a programming problem.

4. Given a $v \in [\underline{v}, \overline{v}]$, it is easy to construct an equilibrium that attains it.

Recall from Proposition 2 that the worst equilibrium is self-enforcing, and here we repeat versions of equations (23.9.2) and (23.9.3),

$$\underline{v} = \min_{y\in Y,\, v_1\in V} \big\{(1-\delta)\, r\big[h(y), y\big] + \delta v_1\big\} \tag{23.12.1}$$

where the minimization is subject to the incentive constraint

$$(1-\delta)\, r[h(y), y] + \delta v_1 \geq (1-\delta)\, r\big\{h(y), H[h(y)]\big\} + \delta \underline{v}. \tag{23.12.2}$$

In expression (23.12.2), we use the worst SPE as the continuation value in the event of a deviation. The minimum will be attained when the constraint is binding, which implies that[15] $\underline{v} = r\{h(y), H[h(y)]\}$, for some government action y. Thus, the problem of finding the worst SPE reduces to solving

$$\underline{v} = \min_{y\in Y} r\big\{h(y), H[h(y)]\big\},$$

then computing v_1 from $(1-\delta) r[h(\underline{y}), \underline{y}] + \delta v_1 = \underline{v}$, where $\underline{y} = \arg\min r\{h(y), H[h(y)]\}$, and finally checking that v_1 is itself a value associated with an SPE. To check this condition, we need to know $\overline{v}$.

[15] An equivalent way to express $\underline{v}$ is $\underline{v} = \min_{y\in Y} \max_{\eta\in Y} r(h(y), \eta)$.

The computation of $\overline{v}$ utilizes the fact that the best SPE is self-rewarding; that is, the best SPE has continuation value $\overline{v}$ when the government follows the prescribed equilibrium strategy. Thus, after we have computed a candidate for the worst SPE value $\underline{v}$, we can compute a candidate for the *best* value $\overline{v}$ by solving the programming problem

$$\overline{v} = \max_{y \in Y} \ r\big[h(y), y\big]$$
$$\text{subject to} \quad r\big[h(y), y\big] \geq (1-\delta) r\big\{h(y), H[h(y)]\big\} + \delta \underline{v}.$$

Here we are using the fact that $\overline{v}$ is the maximum continuation value available to reward adherence to the policy, so that $\overline{v} = (1-\delta) r[h(y), y] + \delta \overline{v}$. Let y^b be the maximizing value of y. Once we have computed $\overline{v}$, we can check that the continuation value v_1 for supporting the worst value is within our candidate set $[\underline{v}, \overline{v}]$. If it is, we have succeeded in constructing V.

23.12.1. When v_1 is outside the candidate set

If our candidate v_1 is not within our candidate set $[\underline{v}, \overline{v}]$, we have to seek a smaller set. We could find this set by pursuing the following line of reasoning. We know that

$$\underline{v} = r\left\{h(\underline{y}), H[h(\underline{y})]\right\} \tag{23.12.3}$$

for *some* $\underline{y}$, and that for $\underline{y}$ the continuation value v_1 satisfies

$$(1-\delta) r[h(\underline{y}), \underline{y}] + \delta v_1 = (1-\delta) r\left\{h(\underline{y}), H[h(\underline{y})]\right\} + \delta \underline{v}.$$

Solving this equation for v_1 gives

$$v_1 = \frac{1-\delta}{\delta}\Big(r\left\{h(\underline{y}), H[h(\underline{y})]\right\} - r[h(\underline{y}), \underline{y}]\Big) + r\left\{h(\underline{y}), H[h(\underline{y})]\right\} \tag{23.12.4}$$

The term in large parentheses on the right measures the one-period temptation to deviate from $\underline{y}$. It is multiplied by $\frac{1-\delta}{\delta}$, which approaches $+\infty$ as $\delta \searrow 0$. Therefore, as $\delta \searrow 0$, it is necessary that the term in braces approach 0, which means that the required $\underline{y}$ must approach y^N.

For discount factors that are so small that v_1 is outside the region of values proposed in the previous subsection because the implied v_1 exceeds the candidate $\overline{v}$, we can proceed in the spirit of Abreu's stick-and-carrot policy, but instead of using v^R as the continuation value to reward adherence (because that is too much to hope for here), we can simply reward adherence to the worst with $\overline{v}$, which we must solve for. Using $\overline{v} = v_1$ as the continuation value for adherence to the worst leads to the following four equations to be solved for $\overline{v}, \underline{v}, \overline{y}, \underline{y}$:

$$\underline{v} = r\left\{h(\underline{y}), H[h(\underline{y})]\right\} \tag{23.12.5}$$

$$\overline{v} = \frac{1-\delta}{\delta}\Big(r\left\{h(\underline{y}), H[h(\underline{y})]\right\} - r[h(\underline{y}), \underline{y}]\Big) + r\left\{h(\underline{y}), H[h(\underline{y})]\right\} \tag{23.12.6}$$

$$\overline{v} = r[h(\overline{y}), \overline{y}] \tag{23.12.7}$$

$$\overline{v} = (1-\delta) r\left\{h(\overline{y}), H[h(\overline{y})]\right\} + \delta \underline{v}. \tag{23.12.8}$$

In exercise *23.3*, we ask the reader to solve these equations for a particular example.

23.13. Examples: alternative ways to achieve the worst

We return to the situation envisioned before the last subsection, so that the candidate v_1 belongs to the required candidate set $[\underline{v}, \overline{v}]$. We describe examples of some equilibria that attain value $\underline{v}$.

23.13.1. Attaining the worst, method 1

We have seen that to evaluate the best sustainable value $\overline{v}$, we want to find the worst value $\underline{v}$. Many SPEs attain the worst value $\underline{v}$. To compute one such SPE strategy, we can use the following recursive procedure:

1. Set the first-period promised value $v_0 = \underline{v} = r\{h(y^\#), H[h(y^\#)]\}$, where $y^\# = \arg\min r\{h(y), H[h(y)]\}$. The competitive equilibrium with the worst one-period value gives value $r[h(y^\#), y^\#]$. Given expectations $x^\# = h(y^\#)$, the government is tempted toward $H(x^\#)$, which yields one-period utility to the government of $r\{h(y^\#), H[h(y^\#)]\}$. Then use $\underline{v}$ as continuation value in the event of a deviation, and construct an increasing sequence of continuation values to reward adherence, as follows:

2. Solve $\underline{v} = (1-\delta) r[h(y^{\#}), y^{\#}] + \delta v_2$ for continuation value v_1.

3. For $j = 1, 2, \cdots$, continue solving $v_j = (1-\delta) r[h(y^{\#}), y^{\#}] + \delta v_{j+1}$ for the continuation values v_{j+1} as long as $v_{j+1} \leq \overline{v}$. If v_{j+1} threatens to violate this constraint at step $j = \overline{j}$, then go to step 4.

4. Use $\overline{v}$ as the continuation value, and solve $v_j = (1-\delta) r[h(\tilde{y}), \tilde{y}] + \delta \overline{v}$ for the prescription $\tilde{y}$ to be followed if promised value v_j is encountered.

5. Set $v_{j+s} = \overline{v}$ for $s \geq 1$.

23.13.2. Attaining the worst, method 2

To construct another equilibrium supporting the worst SPE value, follow steps 1 and 2, and follow step 3 also, except that we continue solving $v_j = (1 - \delta) r[h(y^{\#}), y^{\#}] + \delta v_{j+1}$ for the continuation values v_{j+1} only so long as $v_{j+1} < v^N$. As soon as $v_{j+1} = v^{**} > v^N$, we use v^{**} as both the promised value and the continuation value thereafter. In terms of our recursive strategy notation, whenever $v^{**} = r[h(y^{**}), y^{**}]$ is the promised value, $z^h(v^{**}) = h(y^{**})$, $z^g(v^{**}) = y^{**}$, and $v'[v^{**}, z^h(v^{**}), z^g(v^{**})] = v^{**}$.

23.13.3. Attaining the worst, method 3

Here is another subgame perfect equilibrium that supports $\underline{v}$. Proceed as in step 1 to find the initial continuation value v_1. Now set all subsequent values and continuation values to v_1, with associated first-period outcome $\tilde{y}$ that solves $v_1 = r[h(\tilde{y}), \tilde{y}]$. It can be checked that the incentive constraint is satisfied with $\underline{v}$ the continuation value in the event of a deviation.

23.13.4. Numerical example

We now illustrate the concepts and arguments using the infinitely repeated version of the taxation example. To make the problem of finding $\underline{v}$ nontrivial, we impose an upper bound on admissible tax rates given by $\overline{\tau} = 1 - \alpha - \epsilon$, where $\epsilon \in (0, 0.5 - \alpha)$. Given $\tau \in Y \equiv [0, \overline{\tau}]$, the model exhibits a unique Nash equilibrium with $\tau = 0.5$. For a sufficiently small ϵ, the worst one-period competitive equilibrium is $[\ell(\overline{\tau}), \overline{\tau}]$.

Set $[\,\alpha \quad \delta \quad \overline{\tau}\,] = [\,0.3 \quad 0.8 \quad 0.6\,]$. Compute

$$[\,\tau^R \quad \tau^N\,] = [\,0.3013 \quad 0.5000\,],$$
$$[\,v^R \quad v^N \quad \underline{v} \quad v_{\rm abreu}\,] = [\,-0.6801 \quad -0.7863 \quad -0.9613 \quad -0.7370\,].$$

In this numerical example, Abreu's "stick-and-carrot" strategy fails to attain a value lower than the repeated Nash outcome. The reason is that the upper bound on tax rates makes the least favorable one-period return (the "stick") not so bad.

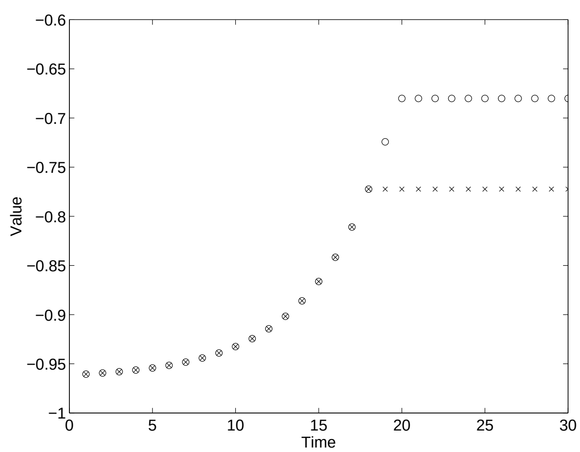

Figure 23.13.1: Continuation values (on coordinate axis) of two SPE that attain $\underline{v}$.

Figure 23.13.1 describes two SPEs that attain the worst SPE value $\underline{v}$ with the depicted sequences of time t (promised value, tax rate) pairs. The circles represent the worst SPE attained with method 1, and the x-marks correspond to method 2. By construction, the continuation values of method 2 are less than or equal to the continuation values of method 1. Since both SPEs attain the same promised value $\underline{v}$, it follows that method 2 must be associated with higher one-period returns in some periods. Figure 23.13.2 indicates that method 2 delivers those higher one-period returns around period 20 when the prescribed tax rates are closer to the Ramsey outcome $\tau^R = 0.3013$.

When varying the discount factor, we find that the cutoff value of δ below which reversion to Nash fails to support Ramsey forever is 0.2194.

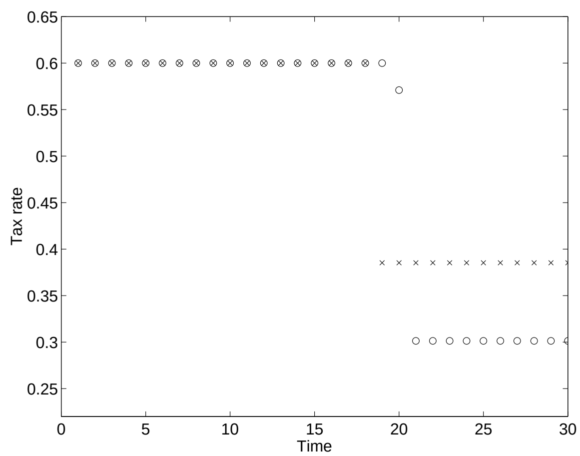

Figure 23.13.2: Tax rates associated with the continuation values of Figure 23.13.1.

23.14. Interpretations

The notion of credibility or sustainability emerges from a ruthless and complete application of two principles: rational expectations and self-interest. At each moment and for each possible history, individuals and the government act in their own best interests while expecting everyone else always to act in their best interests. A credible government policy is one that it is in the interest of the government to implement on every occasion.

The structures that we have studied have multiple equilibria that are indexed by different systems of rational expectations. Multiple equilibria are essential because what sustains a good equilibrium is a system of expectations that raises the prospect of reverting to a bad equilibrium if the government chooses to deviate from the good equilibrium. For reversion to the bad equilibrium to be credible – meaning that it is something that the private agents can expect because the government will want to act accordingly – the bad equilibrium must itself be an equilibrium. It must always be in the self-interest of all agents to behave as they are expected to. Supporting a Ramsey outcome hinges on finding an equilibrium with outcomes bad enough to deter the government from surrendering to a temporary temptation to deviate.[16]

Is the multiplicity of equilibria a strength or a weakness of such theories? Here descriptions of preferences and technologies, supplemented by the restriction of rational expectations, don't pin down outcomes. There is an independent role for expectations not based solely on fundamentals. The theory is silent about which equilibrium will prevail; the theory contains no sense in which the government *chooses* among equilibria.

Depending on the purpose, the multiplicity of equilibria can be regarded either as a strength or as a weakness of these theories. In inferior equilibria, the government is caught in an "expectations trap,"[17] an aspect of the theory that highlights how the government can be regarded as simply resigning itself to affirm the public's expectations about it. Within the theory, the government's

---

[16] This statement means that an equilibrium is supported by beliefs about behavior at prospective histories of the economy that might never be attained or observed. Part of the literature on learning in games and dynamic economies studies situations in which it is not reasonable to expect "adaptive" agents to learn so much. See Fudenberg and Kreps (1993), Kreps (1990), and Fudenberg and Levine (1998). See Sargent (1999, 2008) for macroeconomic counterparts.

[17] See Chari, Christiano, and Eichenbaum (1998).

strategy plays a dual role, as it does in any rational expectations model: one summarizing the government's choices, the other describing the public's rule for forecasting the government's behavior. In inferior equilibria, the government wishes that it could use a different strategy but nevertheless affirms the public's expectation that it will adhere to an inferior rule.

23.15. Extensions

In chapter 24, we shall describe how Chang (1998) and Phelan and Stacchetti (2001) have extended the machinery of this chapter to settings in which private agents' problems have natural state variables like stocks of real balances or physical capital so that their best responses to government policies require that Euler equations (or costate equations) be satisfied. This will turn on an additional source of history dependence. The approach merge features of the method described in chapter 19 with that of this chapter.

Exercises

Exercise 23.1 Consider the following one-period economy. Let (ξ, x, y) be the choice variables available to a representative agent, the market as a whole, and a benevolent government, respectively. In a rational expectations equilibrium or competitive equilibrium, $\xi = x = h(y)$, where $h(\cdot)$ is the "equilibrium response" correspondence that gives competitive equilibrium values of x as a function of y; that is, $[h(y), y]$ is a competitive equilibrium. Let C be the set of competitive equilibria.

Let $X = \{x_M, x_H\}, Y = \{y_M, y_H\}$. For the one-period economy, when $\xi_i = x_i$, the payoffs to the government and household are given by the values of $u(x_i, x_i, y_j)$ entered in the following table:

One-period payoffs $u(x_i, x_i, y_j)$

	x_M	x_H
y_M	10*	20
y_H	4	15*

* Denotes $(x,y) \in C$.

The values of $u(\xi_k, x_i, y_j)$ not reported in the table are such that the competitive equilibria are the outcome pairs denoted by an asterisk (*).

a. Find the *Nash equilibrium* (in pure strategies) and *Ramsey outcome* for the one-period economy.

b. Suppose that this economy is repeated twice. Is it possible to support the Ramsey outcome in the first period by reverting to the Nash outcome in the second period in case of a deviation?

c. Suppose that this economy is repeated three times. Is it possible to support the Ramsey outcome in the first period? In the second period?

Consider the following expanded version of the preceding economy. $Y = \{y_L, y_M, y_H\}$, $X = \{x_L, x_M, x_H\}$. When $\xi_i = x_i$, the payoffs are given by $u(x_i, x_i, y_j)$ entered here:

One-period payoffs $u(x_i, x_i, y_j)$

	x_L	x_M	x_H
y_L	3*	7	9
y_M	1	10*	20
y_H	0	4	15*

* Denotes $(x,y) \in C$.

d. What are Nash equilibria in this one-period economy?

e. Suppose that this economy is repeated twice. Find a subgame perfect equilibrium that supports the Ramsey outcome in the first period. For what values of δ will this equilibrium work?

f. Suppose that this economy is repeated three times. Find an SPE that supports the Ramsey outcome in the first two periods (assume $\delta = 0.8$). Is it unique?

Exercise 23.2 Consider a version of the setting studied by Stokey (1989). Let (ξ, x, y) be the choice variables available to a representative agent, the market as a whole, and a benevolent government, respectively. In a rational expectations or competitive equilibrium, $\xi = x = h(y)$, where $h(\cdot)$ is the "equilibrium response" correspondence that gives competitive equilibrium values of x as a function of y; that is, $[h(y), y]$ is a competitive equilibrium. Let C be the set of competitive equilibria.

Consider the following special case. Let $X = \{x_L, x_H\}$ and $Y = \{y_L, y_H\}$. For the one-period economy, when $\xi_i = x_i$, the payoffs to the government are given by the values of $u(x_i, x_i, y_j)$ entered in the following table:

One-period payoffs $u(x_i, x_i, y_j)$

	x_L	x_H
y_L	0*	20
y_H	1	10*

* Denotes $(x, y) \in C$.

The values of $u(\xi_k, x_i, y_j)$ not reported in the table are such that the competitive equilibria are the outcome pairs denoted by an asterisk (*).

a. Define a *Ramsey plan* and a *Ramsey outcome* for the one-period economy. Find the Ramsey outcome.

b. Define a *Nash equilibrium* (in pure strategies) for the one-period economy.

c. Show that there exists no Nash equilibrium (in pure strategies) for the one-period economy.

d. Consider the infinitely repeated version of this economy, starting with $t = 1$ and continuing forever. Define a *subgame perfect equilibrium.*

e. Find the value to the government associated with the *worst* subgame perfect equilibrium.

f. Assume that the discount factor is $\delta = .8913 = (1/10)^{1/20} = .1^{.05}$. Determine whether infinite repetition of the Ramsey outcome is sustainable as an SPE. If it is, display the associated subgame perfect equilibrium.

g. Find the value to the government associated with the *best* subgame perfect equilibrium.

h. Find the outcome path associated with the *worst* subgame perfect equilibrium.

i. Find the one-period continuation value v_1 and the outcome path associated with the one-period continuation strategy σ^1 that induces adherence to the worst subgame perfect equilibrium.

j. Find the one-period continuation value v_2 and the outcome path associated with the one-period continuation strategy σ^2 that induces adherence to the first-period outcome of the σ^1 that you found in part i.

k. Proceeding recursively, define v_j and σ^j, respectively, as the one-period continuation value and the continuation strategy that induces adherence to the first-period outcome of σ^{j-1}, where (v_1, σ^1) were defined in part i. Find v_j for $j = 1, 2, \ldots,$ and find the associated outcome paths.

l. Find the lowest value for the discount factor for which repetition of the Ramsey outcome is an SPE.

Exercise 23.3 **Finding the worst and best SPEs**

Consider the following model of Kydland and Prescott (1977). A government chooses the inflation rate y from a closed interval $[0, 10]$. There is a family of Phillips curves indexed by the public's expectation of inflation x:

$$U = U^* - \theta(y - x) \tag{1}$$

where U is the unemployment rate, y is the inflation rate set by the government, and $U^* > 0$ is the natural rate of unemployment and $\theta > 0$ is the slope of the Phillips curve, and where x is the average of private agents' setting of a forecast of y, called ξ. Private agents' only decision in this model is to forecast inflation. They choose their forecast ξ to maximize

$$-.5(y - \xi)^2. \tag{2}$$

Thus, if they know y, private agents set $\xi = y$. All agents choose the same ξ, so that $x = \xi$ in a rational expectations equilibrium. The government has one-period return function

$$r(x, y) = -.5(U^2 + y^2) = -.5[(U^* - \theta(y - x))^2 + y^2]. \tag{3}$$

Define a *competitive equilibrium* as a 3-tuple U, x, y such that given y, private agents solve their forecasting problem and (1) is satisfied.

a. Verify that in a competitive equilibrium, $x = y$ and $U = U^*$.

b. Define the government best response function in the one-period economy. Compute it.

c. Define a Nash equilibrium (in the spirit of Stokey (1989) or the text of this chapter). Compute one.

d. Define the Ramsey problem for the one-period economy. Define the Ramsey outcome. Compute it.

e. Verify that the Ramsey outcome is better than the Nash outcome.

Now consider the repeated economy where the government cares about

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}r(x_t,y_t), \tag{4}$$

where $\delta \in (0,1)$.

f. Define a *subgame perfect equilibrium*.

g. Define a *recursive* subgame perfect equilibrium.

h. Find a recursive subgame perfect equilibrium that sustains infinite repetition of the one-period Nash equilibrium outcome.

i. For $\delta = .95$, $U^* = 5, \theta = 1$, find the value of (4) associated with the worst subgame perfect equilibrium. Carefully and completely show your method for computing the worst subgame perfect equilibrium value. Also, compute the values associated with the repeated Ramsey outcome, the Nash equilibrium, and Abreu's simple stick-and-carrot strategy.

j. Compute a recursive subgame perfect equilibrium that attains the worst subgame perfect equilibrium value (4) for the parameter values in part i.

k. For $U^* = 5, \theta = 1$, find the cutoff value δ_c of the discount factor δ below which the Ramsey value v^R cannot be sustained by reverting to repetition of v^N as a consequence of deviation from the Ramsey y.

l. For the same parameter values as in part k, find another cut off value $\tilde{\delta}_c$ for δ below which Ramsey cannot be sustained by reverting after a deviation to an equilibrium attaining the worst subgame perfect equilibrium value. Compute the worst subgame perfect equilibrium value for $\tilde{\delta}_c$.

m. For $\delta = .08$, compute values associated with the best and worst subgame perfect equilibrium strategies.