Chapter 5 Applications of Dynamic Programming

The versatility of the dynamic programming method is really only appreciated by exposure to a wide variety of applications. In this chapter we look at applications of the method organized under four distinct rubrics. In Section 1, we consider problems in which the choice is discrete, typically involving binary decisions at each stage. Section 2 considers a special class of discrete choice models called optimal stopping problems, that are central to models of search, entry and exit. Section 3 considers applications in which the choice variable is continuous. In this section we look at classic models of consumption and investment decision-making. Finally, Section 4 extends the continuous choice setting to one in which a fixed transaction cost must be borne each time the agent acts. The classic application in this area is to inventory problems, but we also study problems of capital replacement and durable goods.

Chapter 4 pressed the reader to think seriously on each occasion about whether conditions ensuring the validity of the dynamic programming approach are met. Specifically, boundedness of the state space and the value function were seen to be crucial elements in justifying the methodology. In this section, these concerns remain as important as ever. However, particularly where the necessary conditions can be easily confirmed, this chapter will frequently skip the essential steps of checking the validity of the key dynamic programming assumptions. In the interests of readability, we are sometimes going to be a little too casual here.

1. Discrete Choice Problems

The analysis in chapter 4 focused on dynamic programming problems where the choice variable was continuous – how much to invest, how much to consume, etc. But dynamic programming is very versatile, and the technique is also very useful for analyzing problems in which the choice variable consists of a small number of mutually exclusive options. In these cases, the Bellman equation may look a little different from those we have seen already. Consider a binary choice problem in which one of two actions, A and B, may be taken. Let x denote a random state variable, and let I be an indicator variable such that I=A indicates that choice A has been made and I=B indicates that choice Bhas been made. Then, the Bellman equation can be written as

$$V(x) = \max_{I = \{A,B\}} \left\{ u(x, x', I) + \beta \int V(x'; I) dF(x' \mid x, I) \right\}.$$
(1.1)

Equation (1.1) allows the choice of A or B to affect the current period's payoff, the future value function and the conditional distribution of x'. It is also common practice to split out the payoffs into separate parts according to the decision made. Thus, (1.1) may also be written as

$$V(x) = \max\left\{u(x, x', A) + \beta \int V(x'; A)dF(x' \mid x, A), \\ u(x, x', B) + \beta \int V(x'; B)dF(x' \mid x, B)\right\}$$

While this might not seem to be much of an improvement, most problems will not be so general and it usually be possible to simplify the notation considerably. For example, if choice A is to accept a payoff and end the dynamic problem and choice B is to postpone the payoff, we could write

$$V(x) = \max\left\{u(x), \beta \int V(x')dF(x'\mid x)\right\}$$

where the dependence of the first term on choosing A and the second term on choosing B will be apparent from the context.

Cake-Scoffing with Taste Shocks.

To begin, we consider yet another variation of the cake-eating problem already analyzed in various guises in Chapter 4 (see, especially, example 4.1 from that chapter). We assume now that the cake must be eaten in its entirety in one period. The taste shock, z, may take on only two values, $0 < z_l < z_h$. and let p_{lh} be the probability of tastes switching from z_l to z_h and let p_{hl} equal the probability of switching from z_h to z_l (this problem belongs to a class of stochastic one-shot decisions called *optimal stopping problems*, which are analyzed more generally in Section 2).

The value function can be written in two parts as

$$\begin{split} V(z_l) &= \max\left\{ z_l u(x), \, \beta \left[p_{lh} V(z_h) + (1 - p_{lh}) V(z_l) \right] \right\}, \\ V(z_h) &= \max\left\{ z_h u(x), \, \beta \left[p_{hl} V(z_l) + (1 - p_{hl}) V(z_h) \right] \right\}, \end{split}$$

where in each case the first term in braces is the payoff from immediate consumption and he second term is the value of postponing consumption.

Consider the choice when $z=z_h$. Immediate consumption implies $V(z_h) = z_h u(x)$. Clearly, if it is not optimal to eat the cake now it will never be optimal to eat the cake when $z=z_h$. Then, either the cake will never be eaten, which would give a return of zero, or it will be eaten when $z=z_h$, which would give a return of $0 < z_l u(x) < z_h u(x)$. Thus it must be optimal to eat the cake immediately when $z=z_h$, so that $V(z_h) = z_h u(x)$. The more substantive question rises when $z=z_l$. Substituting $V(z_h) = z_h u(x)$ into $V(z_l)$, we get

$$V(z_l) = \max\left\{z_l u(x), \beta\left[p_{lh} z_h u(x) + (1 - p_{lh}) V(z_l)\right]\right\}.$$

The payoff from not eating the cake is

$$V^{N}(z_{l}) = \beta \left[p_{lh} z_{h} u(x) + (1 - p_{lh}) V^{N}(z_{l}) \right],$$
(1.2)

which can be solved for

$$V^N(z_l) = \frac{\beta p_{lh} z_h u(x)}{1 - \beta (1 - p_{lh})} \,. \label{eq:VN}$$

The payoff from not waiting is simply $V^E(z_l) = z_l u(x)$, and so the agent will not eat the cake in state $z=z_l$ as long as $V^N(z_l) > V^E(z_l) = z_l u(x)$. This resolves to waiting as long as

APPLICATIONS OF DYNAMIC PROGRAMMING

$$\frac{z_l}{z_h - z_l} < \frac{\beta p_{lh}}{1 - \beta}.$$

For a given difference in the taste shocks, waiting is more likely the higher the discount factor and the greater the probability that the shock will switch from z_l to z_h .

A minor modification of the problem assumes that a fraction, δ , of the cake must be removed and thrown away each period, so that the size of the cake evolves according to $x' = (1 - \delta)x$. Then, the Bellman equations are

$$\begin{split} V(z_l, x) &= \max\left\{z_l u(x), \, \beta\left[p_{lh} V(z_h, (1-\delta)x) + (1-p_{lh}) V(z_l, (1-\delta)x)\right]\right\}, \\ V(z_h, x) &= \max\left\{z_h u(x), \, \beta\left[p_{hl} V(z_l, (1-\delta)x) + (1-p_{hl}) V(z_h, (1-\delta)x)\right]\right\}. \end{split}$$

Again, it is optimal to eat the cake immediately if $z=z_h$, so $V(z_h, x) = z_h u(x)$. The value of not eating the cake when $z=z_l$ is

$$V^{N}(z_{l},x) = \beta \left[p_{lh} z_{h} u(x) + (1 - p_{lh}) V^{N}(z_{l}, (1 - \delta)x) \right].$$
(1.3)

This is where we now hit a problem. When the cake does not decay, so that $\delta=0$, the argument of V^N is identical on both sides of (1.3) (see also eq. (1.2)), and solving for V^N for involved the trivial act of collecting terms. But with $\delta>0$, the arguments of V^N are no longer equal and this simple option is not available. To make progress, we need to try a different tack. One, which we follow here, is to assume a functional form for utility, and then solve by the method of undetermined coefficients.

Assume that $u(x) = \ln(x)$, so that we may write

$$V^{N}(z_{l},x) = \beta \Big[p_{lh} z_{h} \ln(1-\delta) + p_{lh} z_{h} \ln(x) + (1-p_{lh}) V^{N}(z_{l},(1-\delta)x) \Big].$$
(1.4)

Now, the within-period payoff takes the form $A+B\ln(x)$, so it is reasonable to assume that the value function takes the same form, for some unknown coefficients A and B. Substituting this guess into (1.4), we have

$$\begin{aligned} A + B \ln(x) &= \beta p_{lh} z_h \ln(1 - \delta) + \beta p_{lh} z_h \ln(x) \\ &+ \beta (1 - p_{lh}) \big(A + B \ln(1 - \delta) \big) + \beta (1 - p_{lh}) B \ln(x) \,, \end{aligned}$$

and matching coefficients gives

APPLICATIONS OF DYNAMIC PROGRAMMING

$$B = \frac{\beta p_{lh} z_h}{1 - \beta (1 - p_{lh})}$$

and

$$A = \frac{\beta \ln(1-\delta)}{1-\beta(1-p_{lh})} \left(p_{lh} z_h + \beta(1-p_{lh})B \right).$$

Thus, not eating the cake is optimal if

 $A + B\ln(x) > z_l \ln(x) \,,$

where A and B are constants just defined, and only A depends on δ . In fact, as A is decreasing in δ , an increase in the rate of decay makes immediate eating more likely. Note also that, A=0 when $\delta=1$, so the earlier result it immediately recovered. However, for $\delta \neq 0$, the term $\ln(x)$ is not eliminated from the inequality so the choice to eat or not eat the cake depends on the size of cake remaining.

EXERCISE 1.1 (Cake scoffing with random spoiling). Assume now that in each period, the cake goes bad and must be discarded with probability q. Derive the optimal policy.

A COUPLE MORE EXAMPLES WILL BE ADDED HERE

2. Optimal Stopping Problems

A special class of problems involving a discrete choice are those in which there is a single decision to put an end to an ongoing problem. The stock option and cake scoffing problems in Section 1 are examples of this type of problem. Other examples include

- A student must decide when to give up trying to solve a homework problem.
- A firm must decide when to exit an industry.

• A firm decides when to stop working on the development of a product and to launch it

• An unemployed worker decides when to accept a job from a sequence of offers made.

Such problems are collectively known as optimal stopping problems. We analyze some in this section. The first three problems, concerned with exercising a stock option, searching for the lowest price, and accepting the highest bid, will introduce some features common to most optimal stopping problems. We then turn to more substantive problems in labor markets and industry dynamics.

Stock Options (Ross [1983:4])

Let p_t denote the price of a given stock in period t. Assume that $p_{t+1} = p_t + x_{t+1}$, where x_t is an independently and identically distributed random variable with distribution F(x) and mean zero (i.e. the stock price follows a random walk). Suppose you have an option to buy the stock at price c and you have T periods to exercise the option.

The Bellman equation is

$$V_t(p) = \max\left\{ p - c, \, \beta \int V_{t+1}(p+x) dF(x) \right\},\tag{2.1}$$

with boundary condition

$$V_T(p) = \max\{p-c, 0\}.$$
 (2.2)

Equation (2.1) gives a choice between exercising the option this period, and receiving p-c, or waiting another time period to await a fresh realization for the stock price. Equation (2.2) states that in period T either the option is exercised or it becomes worthless. Note that, as this is a finite horizon problem, time subscripts on the value functions are important.

Although no explicit solution for the value function exists, we can characterize the solution strategy. In order to do so, we will make use of the following properties:

LEMMA 2.1. (a) $V_t(p) - p$ is decreasing in p. (b) $V_t(p)$ is increasing in p. (c) $V_t(p)$ is decreasing in t.

PROOF. We will prove part (a). Parts (b) and (c) are intuitive. In particular, having less time until the option expires cannot make one better off. The proof of part (a) is by induction. Assume $V_{t+1}(p) - p$ is decreasing in p. We can then show that $V_t(p) - p$ is also decreasing in p:

$$\begin{split} V_t(p) - p &= \max\left\{-c, \, \beta \int V_{t+1}(p+x) dF(x) - p\right\} \\ &= \max\left\{-c, \, \beta \int V_{t+1}(p+x) dF(x) - p - \int x dF(x)\right\} \\ &= \max\left\{-c, \, \beta \int \left[V_{t+1}(p+x) - (p+x)\right] dF(x)\right\}. \end{split}$$

In the second line, we subtracted the term $\int x dF(x)$. This is the mean of x and is zero by assumption. For each x, $V_{t+1}(p+x) - (p+x)$ is decreasing in p by assumption. Thus, $V_t(p) - p$ is decreasing in p if $V_{t+1}(p) - p$ is. The proof is then completed upon noting that

$$V_T(p) - p = \max\left\{p - c, 0\right\} - p$$
$$= \max\left\{-c, -p\right\}$$

is decreasing in p. •

Now, it is optimal to exercise the option in period t if $p-c \ge \beta \int V_{t+1}(p+x)dF(x)$. But if this is the case, then (2.1) also tells us that $V_t(p) = p-c$ or, equivalently, that $V_t(p) - p = -c$. The left hand side of this equation is decreasing in p by part (a) Lemma 2.1. Moreover, when $V_t(p) - p > -c$, equation (2.1) tells us that $\beta \int V_{t+1}(p+x)dF(x) > p-c$ and it is not optimal to exercise the option. Thus, there exists a strike price, p_t^* , such that for any $p < p_t^*$ the option is not exercised and for any $p \ge p_t^*$ the option is exercised. Part (c) of Lemma 2.1 shows that this strike price is decreasing in t.

One can also establish the classic result that the value of an option is increasing in the variance of the stock price, although we shall only illustrate the idea here and leave the formal proof as an exercise. The bold line in Figure 2.1 illustrates the value of an option on a stock with a constant price p. For $p \leq c$, the option has zero value, while it is equal to p-c for $p \ge c$. Thus, the value function is convex. Now imagine that x can take on the value +u with probability $\frac{1}{2}$ and -u with probability $\frac{1}{2}$, such that p+u>c. Then, as Figure 1 illustrates, V(p;u) is clearly going to be increasing in u. This is an illustration of the general result, from Jensen's inequality, that for convex functions $E[V(x)] \ge V(E[x])$.



Although we are not been able to explicitly solve the model, several key properties have been derived. First, there exists a minimum strike price for each period, and this strike price is decreasing with the passage of time. We have illustrated the idea that, because V(p) is a convex function, uncertainty raises the value of the stock option. This idea is of course no different from the fact that risk averse people with concave utility functions are willing to pay for insurance to reduce uncertainty, while risk takers with convex utility are willing to pay to gamble and increase uncertainty.

EXERCISE 2.1 We did not formally show that the value of the stock option is increasing in the variance. This exercise asks you to do so. Assume F(x) is the

normal distribution with mean zero and variance σ^2 . Prove that $V_t(p;\sigma^2)$ is increasing in σ^2 .

When $p=p^*$, the payoffs from exercising and not exercising the option are identical, and we can write

$$p_{t}^{*} - c = \beta \int V_{t+1} \left(p_{t}^{*} + x \right) dF(x) ,$$

$$p_{t}^{*} = c + \beta \int V_{t+1} \left(p_{t}^{*} + x \right) dF(x) .$$
(2.3)

The critical value, p_t^* , is defined by two components. The first is simply the cost of exercising the option and, if dynamics did not matter it would be worth exercising the option. The second term, therefore, is the option value. Equation (2.3) is a version of what is known generically as the **fundamental reservation price equation** for optimal stopping problems. In this case, the equation is not too useful because it contains the unknown value function. However, we already know that $V_{t+1} \ge 0$, so we can verify from (2.3) that $p_t^* > c$. In the next example, we will be able to derive a fundamental reservation price equation that does not involve the value function.

Searching for the Lowest Price

or

Consider an agent interested in purchasing a single unit of a good whose price varies from store to store. At each store visited, the individual is quoted a price, $p \ge 0$, a random draw form the distribution F(p). Sampling a price costs c. Stigler (1961) suggested that the individual should sample n stores and then buy from the lowest price quoted. After visiting n stores, the expected value of the minimum price is

$$m(n) = n \int_{0}^{\infty} pf(p) \left[1 - F(p) \right]^{n-1} dp .$$
(2.4)

Equation (2.4) gives the expected value of a price, $\int pf(p)dp$, conditional on all other prices exceeding this one (the term $[1 - F(p)]^{n-1}$. As any one of the *n* prices could be the largest, the whole expression is multiplied by *n*. An integration by parts (which we leave to the reader to do as an exercise) yields

$$m(n) = \int_{0}^{\infty} \left[1 - F(p)\right]^{n} dp .$$
(2.5)

The expected reduction in price from increasing n by one unit is therefore

$$g(n) = m(n) - m(n+1) = \int_{0}^{\infty} [1 - F(p)]^{n} dp - \int_{0}^{\infty} [1 - F(p)]^{n+1} dp$$
$$= \int_{0}^{\infty} [1 - F(p)]^{n} F(p) dp ,$$

which declines with n at a decreasing rate. Therefore n should be chosen so that

g(n+1) < c < g(n).

Stigler showed that if all customers follow this rule, each store faces a well-defined downward-sloping demand curve, the exact properties of which depend on the search cost, c, and the distribution F.

The sample-size rule proposed by St8gler is not very appealing. Even if the agent receives a price quote p < c, so that no further search could have positive value, he continues to search until n stores have been sampled. A more attractive rule would be one that indicates on the fly when search should stop. Let us imagine that the agent visits each store in succession at the rate of one per period. Then, given that the price quoted in the current period is p, the individual can choose either to stop now and purchase, or to sample again. If he stops now he receives u-p, where u is the utility of consumption. If he continues he enters the next period as an active searcher. This is now a dynamic programming problem with Bellman equation

$$V(p) = \max\left\{u - p, -\beta c + \beta \int_{0}^{\infty} V(p) dF(p)\right\}.$$

The second term in braces is a constant independent of the current quote, because prices are i.i.d. draws. The first term in braces is obviously declining in p. Because V(p) attains a maximum of u when p=0, there must exist a unique p^* such that the agent is indifferent between stopping and continuing. This is illustrated in Figure 2.2. Any price greater than p^* stimulates further search, while any price less than p^* induces a purchase.



FIGURE 2.2. The reservation price for the consumer search problem

The optimality condition implies that p^* satisfies

$$u - p^* = -\beta c + \beta \int_0^\infty V(p) dF(p)$$

= $-\beta c + \beta \int_0^{p^*} V(p) dF(p) + \beta \int_{p^*}^\infty V(p) dF(p).$ (2.6)

Now, any price quote under p^* is accepted, so V(p) = u - p for all $p \le p^*$. Any price quote over p^* yields a value that is independent of the quote. Let this value be v. As we have defined p^* as the maximum price that is accepted, it must be the case that $v = u - p^*$. Therefore, (2.6) can be written as

APPLICATIONS OF DYNAMIC PROGRAMMING

$$u - p^* = -\beta c + \beta \int_0^{p^*} (u - p) dF(p) + \beta (u - p^*) \int_{p^*}^{\infty} dF(p)$$

= $-\beta c + \beta \int_0^{p^*} (u - p) dF(p) + \beta (u - p^*) \left[1 - \int_0^{p^*} dF(p) \right]$
= $-\beta (u - p^*) + \beta \left[-c + \int_0^{p^*} (p^* - p) dF(p) \right].$

Therefore,

$$p^* = u - \frac{\beta}{1 - \beta} \left[-c + \int_0^{p^*} (p^* - p) dF(p) \right],$$
(2.7)

which is the fundamental reservation price equation for this problem. Even when a price quote less that u is received, the agent may continue searching in the hope that a lower quote arrives later. The second term in (2.7) provides the present value of maintaining the option to continue searching.

The reservation price principle of the optimal stopping problem remains true even if quotes previously rejected can be recalled. Let P denote the smallest quote received prior to the current period. Then, the Bellman equation with recall reads

$$V(p) = \max\left\{u - \min(P, p), -\beta c + \beta \int_{0}^{\infty} V(p) dF(p)\right\},\$$

so if $u-p \ge -\beta c + \beta \int V(p) dF(p)$ the current offer is accepted and if $u-P > -\beta c + \beta \int V(p) dF(p)$ a past offer is accepted. But these inequalities have the same value on the right and side and, as there is a discounting cost to waiting, the first time a quote provided a surplus exceeding $-\beta c + \beta \int V(p) dF(p)$ it would have been accepted immediately. Thus, the ability to recall earlier bids does not change the optimal solution at all.

Asset Selling

An even simpler problem concerns an agent with an asset he is trying to sell. Assume he receives offers at the rate of one per period. Denote these offers by p_0, p_1, \ldots , which are random i.i.d. draws from the closed interval $[p_L, p_H]$. If an offer is rejected, the agent must wait until the next period to get another offer.

The Bellman equation is

$$V(p) = \max\left\{p, \beta \int_{p_L}^{p_H} V(p) dF(p)\right\}.$$

The first term in braces is the value of accepting the current offer. The second term is the return from rejecting the offer and waiting for another draw next period. Clearly, the optimal policy is to accept p if $p \ge \beta E[V(p)]$. As E[V(p)] is a constant independent of p again, (to the i.i.d. assumption), this again implies there is a reservation price, p^* , below which the offer is rejected and above which it is accepted. If there is an interior solution, the reservation price satisfies $p^* = \beta E[V(p')]$. We will see in a moment that p^* is always less than p_H , so the only alternative for a non-interior solution is if $p^* = p_L$. Hence,

$$p^* = \max\left\{p_L, \beta \int_{p_L}^{p_H} V(p) dF(p)\right\}$$
$$= \max\left\{p_L, \beta \int_{p_L}^{p^*} V(p) dF(p) + \beta \int_{p^*}^{p_H} V(p) dF(p)\right\}.$$

Any offer over p^* is accepted, yielding V(p)=p. Any price below p^* yields a value independent and equal to p^* . Therefore,

$$p^{*} = \max\left\{p_{L}, \beta p^{*} \int_{p_{L}}^{p^{*}} dF(p) + \beta \int_{p^{*}}^{p_{H}} p dF(p)\right\}$$
$$= \max\left\{p_{L}, \beta p^{*} F(p^{*}) + \beta \int_{p^{*}}^{p_{H}} p dF(p)\right\}.$$
(2.8)

The solution to this equation solves the optimal stopping problem. We can easily verify that the solution is unique. Differentiating the second term on the right-hand side of (2.8) with respect to p^* yields $0 < \beta F(p^*) < \beta < 1$. As F is a cumulative distribution function, we also expect it to be a continuous function. Hence, (2.8) is a contraction, mapping values from the closed bounded interval $[p_L, p_H]$ back into itself. The solution is depicted in Figure 2.3 which plots the right hand side of (2.8) as the convex locus **aa**.¹ As $p^* \rightarrow p_H$, the right-hand side of (2.8) approaches βp_H . Thus, it will never be optimal to reject an offer greater than βp_H . As $p^* \rightarrow p_L$ it approaches $\beta E(p)$. If β is not too small, there is an interior solution, as indicated. For β sufficiently small, the locus **aa** lies below the 45° line for the entire range $[p_L, p_H]$. In that case $p^* = p_L$ and it is optimal to accept the first offer made. The locus shifts upward when β increases, so more patient people set a higher reservation price.



FIGURE 2.3 The reservation price for the asset selling problem

¹ The second derivative of (2.8) is $\beta F'(p^*) \ge 0$, so the function is convex.

APPLICATIONS OF DYNAMIC PROGRAMMING

Noting that $\int_{p_L}^{p^*} dF(p) = 1 - \int_{p^*}^{p_H} dF(p)$, a little rearrangement of (2.8) yields the fundamental reservation price equation

$$p^* = \frac{\beta}{1-\beta} \int_{p^*}^{p_H} (p-p^*) dF(p)$$

The right-hand side records the value of rejecting an offer: it is the value of maintaining the option to secure an improved offer in the future.

The simplicity of the solution technique for this and the previous problems serially rests in part on the assumption of i.i.d. draws. In many settings it will be reasonable to assume that bids and quotes are serially correlated, and these may quantitatively change the reservation price. However, even when bids are correlated so that p' is dependent on p, the optimal stopping policy is still to accept the first bid exceeding a constant threshold, p^* . To see this for the asset selling problem, note that the optimality condition can be written as

$$p^* = \beta \int_{p_L}^{P_H} V(p) dF(p \mid p^*), \qquad (2.9)$$

where the conditional distribution captures the serial correlation (for positive serial correlation, $F(p | p^*)$ is decreasing in p^* for any p). Clearly the solution to (2.9) does not involve any current or recent realizations of the sequence of bids. The only substantive change with serial correlation concerns the value function for offers below p^* . In the i.i.d. case, $V(p)=p^*$ for all $p \leq p^*$. In the case of positive serial correlation, a low current offer makes low offers more likely in the near future, and this makes the agent worse off. Thus, $V(p) < p^*$ for $p < p^*$, and the gap between p^* and V(p) gets larger the lower is p.

Commentary

Before turning to more substantive problems involving labor market applications and industry dynamics, it is useful at this stage to take stock of what we have learned. The optimal stopping problems we have studied have some common features, although we have drawn out different features for each one:

- There is a unique reservation price that triggers an end to the problem.
- The Bellman equation is a convex function of the state variable.

• The value of the Bellman equation is increasing in the variance of the state variable. That is, even though agents may be risk averse or risk neutral in general, risk is valuable in the context of optimal stopping. This result comes from applying Jensen's inequality to the convexity of the value function.

• Serial correlation in the state variable may have quantitative effects on the reservation price but it does not alter then reservation price principle of optimal stopping.

• The ability to recall previous value of the state variable has no quantitative effect on the reservation price.

At this point it is helpful to introduce an important caveat to the reservation price principle. In each problem studied so far, we have assumed that agents know the distribution form which the state variable is drawn. The reservation price principle does not generally survive an extension of these types of problems to situations where the agent must also learn the distribution. Rothschild (1974), who has studied this problem, provides a simple example. Imagine a consumer searching for the lowest price does not know the distribution of prices. His prior is that either the price is always \$3, or that it is \$2 with probability 0.01 and \$1 with probability 0.99. If the consumer receives an offer of \$3, he will accept because he now believes that all price quotes will be identical and further search is pointless. If he receives an offer of \$2 then (assuming search costs are not too high) he will not accept because the odds that the next offer is only \$1 is now perceived to be very high. Thus, the consumer accepts bids of \$1 and \$3, but not \$2, and the unique reservation price property has vanished.

It is also no longer true that the presence or absence of recall is unimportant. If no recall is possible, the last offer observed will always be the one accepted. However, if recall is possible, then exploration might be valuable. For example, imagine you are searching for an apartment along a long road. You do not know the location of good and bad apartments, so you drive along the road observing from your car. You keep driving after seeing some acceptable places until you see you have entered a bad neighborhood. You then backtrack to select an apartment you saw earlier.

However, you may prefer not to backtrack if evaluating the quality of an apartment is very costly. imagine you have to make an appointment to see each apartment and this is the only way to evaluate its quality. Although you may remain unaware of the distribution of apartments, the high cost of exploration may induce you to accept the first apartment that meets your minimum res4rvation quality. This reservation quality will depend, of course, on your prior about the distribution; but the point is that you may not explore even though you know that your prior could be wrong. This example is consistent with formal results obtained by Rothschild, who shows that the reservation price property of the optimal stopping problem survives only if search costs are large enough; if they are small, the agent will undertake active exploration to learn about the true distribution.

Further exploration of problems with an unknown distribution takes us further into the theory of search and learning than is merited at this point. In the remainder of this section, therefore, we look at two substantive applications of optimal stopping. The first is concerned with labor market job search. The second with firms' decisions about industry exit. In both applications, we will maintain the assumption that the population distribution of the state variable is known.

Job Search

A particularly well-mined application of optimal stopping problems concerns search in the labor market. We consider some simple examples here. An infinitely-lived individual maximizes

$$E_0 \sum_{t=0}^{\infty} \beta^t y_t \; ,$$

where $y_i = w$ if employed at job w and $y_i = c$ if unemployed. The agent received one job offer with probability ϕ in each period, and no offers with probability $1-\phi$. The offer consists of a wage, $w \in [0, \overline{w}]$, which is a random draw from the distribution F(w). If the offer is accepted, the job is assumed to last forever. If the offer is rejected, the agent earns c for that period, and must wait until the next period to have a chance p of receiving another offer.

Let V(w) be the value of having an offer with wage w, let v denote the value of being unemployed without an offer. Then, the Bellman equation is

$$V(w) = \max\left\{\frac{w}{1-\beta}, c+\beta\left[\phi EV(w') + (1-\phi)v\right]\right\},$$
(2.10)

where

$$v = c + \beta \left[\phi E \left[V(w') \right] + (1 - \phi) v \right].$$
(2.11)

The first term on the right-hand side of (2.10) is the discounted present value of earning w in each period forever, and it represents the value to the worker of accepting the current offer. The second term is the value to the worker of rejecting the offer. The worker immediately receives the unemployment benefit c. In the subsequent period, he receives with probability ϕ a random offer, w', yielding expected value E[V(w')].² With probability $1-\phi$, no offer is received and this has the value v. Equation (2.11) clarifies what v is. It consists of earning c as an unemployed worker, and then in the next period either receiving a random offer or not receiving a random offer. That is, the value of not receiving an offer is the same as the value of turning down an offer.

No choice is involved with (2.11), so we can simply rearrange it to obtain

$$v = \frac{c}{1 - \beta(1 - \phi)} + \frac{\beta\phi}{1 - \beta(1 - \phi)} E\left[V(w')\right],$$

and substitute back into (2.10):

$$V(w) = \max\left\{\frac{w}{1-\beta}, \frac{c}{1-\beta(1-\phi)} + \frac{\beta\phi E[V(w')]}{1-\beta(1-\phi)}\right\}.$$
(2.12)

As before, we anticipate a unique reservation wage, w^* , defining the minimum wage necessary for the agent to accept an offer. Assuming an interior solution, the reservation wage satisfies

² Because this is a stationary problem, V(w) on the left hand side and V(w') on the right hand side must be the same function.

APPLICATIONS OF DYNAMIC PROGRAMMING

$$\frac{w}{1-\beta} = \frac{c}{1-\beta(1-\phi)} + \frac{\beta\phi E\left[V(w^{\,\prime})\right]}{1-\beta(1-\phi)}.$$

Clearly, for $w \le w^*$, V(w) is constant because the size of a rejected offer has no bearing on future returns. But, we know that $V(w^*) = w^*/(1-\beta)$, so $V(w) = w^*/(1-\beta)$ for all $w \le w^*$. For $w \ge w^*$, of course, the offer is accepted and $V(w) = w/(1-\beta)$. Thus, we have

$$\frac{w^*}{1-\beta} = \frac{c}{1-\beta(1-\phi)} + \frac{\beta\phi}{1-\beta(1-\phi)} \left(\int_0^{w^*} \frac{w^*}{1-\beta} \, dF(w) + \int_{w^*}^{\overline{w}} \frac{w}{1-\beta} \, dF(w) \right).$$
(2.13)

Noting that $\int_0^{w^*} w^* dF(w) = w^* - \int_{w^*}^{\overline{w}} w^* dF(w)$, (2.7) can be written as

$$\frac{w^*}{1-\beta} = \frac{c}{1-\beta(1-\phi)} + \frac{\beta\phi}{(1-\beta)(1-\beta(1-\phi))} \left(w^* + \int_{w^*}^{\overline{w}} (w-w^*) dF(w) \right)$$

or

$$w^* = c + \frac{\beta\phi}{(1-\beta)} \int_{w^*}^{\overline{w}} (w - w^*) dF(w) .$$
(2.14)

This is the fundamental reservation price equation. The right-hand side again records the value of rejecting an offer. It is the sum of the compensation, c, received while unemployed and the option value of staying unemployed to secure an improved offer. It is easy to calculate that the right hand side of (2.14) is decreasing in w^* , reaching a minimum of c when $w^* = \overline{w}$. The unique fixed point is shown in Figure 2.4. Note also that a reduction in ϕ rotates the curve downward, reducing w^* . That is, when offers are secured less frequently, the rational job hunter accepts worse offers.

Firms may also conduct searches for workers and in the interest of symmetry we consider a simple example here. Imagine a firm can locate at most one potential worker each period, who demands a wage W from the distribution G(W). If hired, the worker produces one unit of output at a price p, forever. In each period, there is a probability π that no potential worker is found. The Bellman equation for the firm is

$$J(W) = \max\left\{\frac{p - W}{1 - \beta}, \beta\left[\pi E\left[J(W')\right] + (1 - \pi)j\right]\right\},$$
(2.15)



FIGURE 2.2 The Reservation Wage for Job Hunters $% \left({{{\rm{A}}_{{\rm{B}}}} \right)$

where J(W) is the value of having a potential employee demanding W, and

$$j = \beta \left[\pi E \left[J(W') \right] + (1 - \pi) j \right]$$
(2.16)

is the value of not having found a potential employee. Solving (2.16) for j and substituting into (2.15) yields

$$J(W) = \max\left\{\frac{p-W}{1-\beta}, \frac{\beta\pi E\left[J(W^{\,\prime})\right]}{1-\beta(1-\pi)}\right\}.$$

Again assuming an interior solution, the reservation wage (this time a maximum wage)

satisfies

$$\frac{p - W^*}{1 - \beta} = \frac{\beta \pi E \left[J(W') \right]}{1 - \beta (1 - \pi)} \,.$$

Following by now familiar arguments, $J(W) = (p - W^*)/(1 - \beta)$ for any $W \ge W^*$, and $J(W) = (p - W)/(1 - \beta)$ for any $W \le W^*$. Thus,

$$\frac{p - W^*}{1 - \beta} = \frac{\beta \pi}{1 - \beta(1 - \pi)} \left[\int_0^{W^*} \frac{p - W}{1 - \beta} \, dG(W) + \int_{W^*}^{\infty} \frac{p - W^*}{1 - \beta} \, dG(W) \right]$$
$$= \frac{\beta \pi}{(1 - \beta) \left(1 - \beta(1 - \pi)\right)} \left[p - W^* + \int_{W^*}^{\infty} \left(W^* - W\right) \, dG(W) \right].$$

Some rearrangement gives

$$W^* = p - \frac{\beta \pi}{(1-\beta)} \int_{W^*}^{\infty} (W^* - W) \, dG(W) \,, \tag{2.17}$$

which can be compared with the fundamental reservation wage equation (2.14) for workers. Notice here that the option value of continued search implies that firms offer a wage below the marginal product of labor. The tighter the job market, in the sense that potential employees are harder to find, the greater the wage the firm is willing to offer.

Optimal stopping models of job search have dominated equilibrium models of unemployment and wage determination over the last 20 years. We have developed here a labor supply curve that provides a reservation wage w^* below which workers will remain unemployed, and a reservation wage W^* above which firms will not hire workers. Only if $w^* \leq W^*$ is there an opportunity to match unemployed workers with vacancies. However, it is not obvious that this inequality will hold. In fact, much of the literature on job search has been concerned with how variations in job turnover rates, bargaining institutions, and numerous other aspects of the labor market affect the joint determination of equilibrium wages and unemployment. Studying this literature further takes us too far afield from the main task at hand in this chapter. Mortensen and Pissarides (1999) provide an excellent review of the theory.

Firm Exit

Among a short list of important papers on the evolution of industrial structure are two outstanding applications of optimal stopping by Jovanovic (1982) and Hopenhayn (1992). In Jovanovic's paper, firms are not sure of their abilities upon entering. Time slowly and noisily reveals to the firms just how good they are, and the least-able firms eventually discover that their expected costs are too high and decide to exit. In Hopenhayn's model, all firms are equally able, but each firm experiences a sequence of serially correlated shocks to productivity that affect their expected costs in subsequent periods. Firms suffering a sequence of sufficiently undesirable shocks to productivity will choose to exit. In both papers, the exit decision is embodied in a fully-specified equilibrium model of the market. For the purposes of the example here, however, we will concentrate on the optimal stopping problem. The language used in describing the firm's problem borrows from Jovanovic rather than Hopenhayn. However, as will become apparent, the transition between one model and the other can be undertaken in a few short steps.

Production costs for output level q are given by c(q)x, where c is a strictly increasing, convex, differentiable function satisfying $\lim_{q\to\infty} c'(q) = \infty$. The variable x measures firm efficiency and is a strictly positive random variable, fluctuating from period to period. Decisions will be made on the basis of the conditional expectation of x, which we denote by y. That is, $y_t = E[x_t | I_{t-1}]$, and as x is strictly positive so is y. In both models, y is a serially correlated random variable, although for different reasons, and in both models we will call y expected efficiency.

Firms are price-takers. Given a constant price p, which we normalize to p=1, the firm chooses output to maximize expected profits in any given period:

$$\pi(y) = \max_{q} q - c(q) y , \qquad (2.18)$$

which yields the first-order condition

$$1 = c'(q)y. (2.19)$$

Let the solution to (2.19) be denoted by q(y). One can easily calculate from (2.19) that q'(y) = -c' / yc'' < 0, so optimal output is decreasing in expected efficiency. Substituting the solution into (2.18), we have

$$\pi(y) = q(y) - c(q(y))y$$

yielding $\pi^{/}(y) = -c < 0$ and $\pi^{//}(y) = -c^{/}q^{/} > 0$, where the envelope theorem was used in evaluating the first derivative. Thus, expected profits are a decreasing convex function of expected efficiency.³

It should be noted that we have treated profit maximization as a static problem. Because profits depend only on the value of the single state variable, y, this treatment is valid only if the choice variable does not affect the evolution of the state variable. In both models, as we will later see, y turns out to be a purely exogenous random variable, so the static profit maximization approach is valid.

Let $F(y \mid y)$ denote the distribution of next period's expected efficiency given this period's expected efficiency. It is assumed that $\partial F(y \mid y) / \partial y \leq 0$ for all y. The dynamic problem facing the firm is to decide if and when to exit the industry. If the firm stays in the industry through the next period it will earn profits $\pi(y')$ and maintain the option to continue thereafter. If it decides to leave the industry at the end of the current period, it will earn W, where W represents the sum of the value of selling the firm's assets on the secondary market and the entrepreneur's value in an outside activity. Thus, with a discount factor of β , $(1-\beta)W$ is the opportunity cost per period of remaining in the industry.

The Bellman equation for the firm is

$$V(y) = \pi(y) + \beta \int \max[W, V(y')] dF(y' | y), \qquad (2.20)$$

where the conditional distribution is assumed to have the Feller property. The maximization problem is only with respect to the binary choice of continuation versus exit because expected profits have already been maximized and there is no linkage between $\pi(y)$ and V(y'). The problem is therefore a pure optimal stopping problem.

We can analyze the key characteristics of this optimal stopping problem exploiting the theorems of Chapter 4. First, we show that there is a unique solution to (2.20). Note

³ The explanation for convexity is exactly the same as the reason why profit functions are convex in price, which you will recall from microeconomic theory.

first that $\pi(y)$ is bounded because p=1 is bounded and y is strictly positive.⁴ As $\pi(y)$ is continuous, then so is $\sum_{t=0}^{T} \beta^t \pi(y_t)$ for any T and any feasible sequence $\{y_t\}_{t=0}^{T}$, and so, in turn, are $V(y_0) = \sum_{t=0}^{T} \beta^t \pi(y_t) + \beta^T W$ and $\max[W, V(y_0)]$ Finally, as F has the Feller property, continuity is preserved in the integration. Thus, the operator T defined in (2.20) transforms bounded continuous functions into bounded continuous functions.

We can next readily verify that T is a contraction mapping. We do this directly by means of the supremum norm, rather than by Blackwell's theorem. Let $f(y) \neq g(y)$ denote any two bounded continuous functions. Then

$$\begin{aligned} \|Tf(y') - Tg(y')\| &= \beta \left\| \int \max[W, f(y')] dF(y' \mid y) - \int \max[W, g(y')] dF(y' \mid y) \right\| \\ &= \beta \left\| \int \max[W, f(y')] - \max[W, g(y')] dF(y' \mid y) \right\| \\ &\leq \beta \left\| \max[W, f(y')] - \max[W, g(y')] \right\| \\ &\leq \beta \left\| f(y') - g(y') \right\| \\ &< \left\| f(y') - g(y') \right\|. \end{aligned}$$

The first inequality is because $F(y \mid y)$ has the Feller property and maps the functions fand g into the same bounded intervals. The second inequality comes from the following argument. Consider any y' such that $f(y') \ge g(y')$. Then: $W \ge f(y') \Rightarrow W \ge g(y') \Rightarrow$ $\max[W, f(y')] - \max[W, g(y')] = 0 \le f(y') - g(y'); \quad f(y') \ge W \ge g(y') \Rightarrow \max[W, f(y')]$ $-\max[W, g(y')] = f(y') - W \le f(y') - g(y'); \quad (\text{iii}) \quad f(y') \ge g(y') \ge W \Rightarrow \max[W, f(y')] - \max[W, g(y')] = f(y') - g(y')$. Thus, if the supremum ||f(y') - g(y')|| is found at a point where $f(y') \ge g(y')$, the inequality is proved. But for any two functions h_1 and h_2 ,

⁴ These limits on p and y, in conjunction with the continuity of c' and the assumption that $\lim_{q\to\infty} c'(q) = \infty$, imply that there exists a finite q satisfying c'(q) = 1/y. As q is finite, so are revenues.

 $\|h_1 - h_2\| = \|h_2 - h_1\|$ by the symmetry of distance functions, and the same arguments can therefore be applied for any $g(y') \ge f(y')$.⁵

So now we have established that T is a contraction mapping and that there is a unique solution to (3), we can explore some of the properties of this solution. It turns out that these depend critically on the properties of the conditional distribution $F(y \mid y)$. We have already mentioned that both Hopenhayn and Jovanovic introduce some persistence in y by assuming that $F(y \mid y)$ is weakly decreasing in y. Given this assumption, we can now show that V(y) is strictly decreasing in y. Recall that, as T is a contraction mapping, $V(y) = \lim_{n\to\infty} T^n g(y)$ for any bounded continuous function g(y). Well, let's assume that g(y) is decreasing in y.⁶ Then it will be the case that $\max[W, g(y)]$ is weakly decreasing in y. To see this last one, let a be the minimum value for y, let y^* be the value of y such that g(y) = W. Then, as g(y) is decreasing in y, $\max[W, g(y)] = W$ for any $y \ge y^*$, and $\max[W, g(y)] = g(y)$ for any $y \le y^*$. That is,

$$\int_{a}^{\infty} \max[W, g(y')] dF(y' \mid y) = W \int_{y^{*}}^{\infty} dF(y' \mid y) + \int_{a}^{y^{*}} g(y) dF(y' \mid y)$$

$$= W \left(1 - F(y^{*} \mid y) \right) + g(y) F(y' \mid y) \Big|_{a}^{y^{*}} - \int_{a}^{y^{*}} g_{y'}(y') F(y' \mid y) dy'$$

$$= W \left(1 - F(y^{*} \mid y) \right) + g(y^{*}) F(y^{*} \mid y) - \int_{a}^{y^{*}} g_{y'}(y') F(y' \mid y) dy'$$

$$= W - \int_{a}^{y^{*}} g_{y'}(y') F(y' \mid y) dy'. \qquad (2.21)$$

 $^{^{5}}$ You do not need to go through this argument each time. It is well known that the max operator drops out of such supremum functions so you can go from the third to the fourth line without comment.

⁶ If we chose a function g(y) that is increasing in y, we would still end up with a function V(y) that is decreasing in y, but we would just have a hard time proving it.

Differentiating with respect to y yields

$$\frac{d}{dy} \int_{a}^{\infty} \max[W, g(y')] dF(y' \mid y) = -\int_{a}^{y^{*}} g_{y'} F_{y}(y' \mid y) dy' \ge 0.$$

Finally, as $\pi(y)$ is strictly decreasing in y, Tg(y) is strictly decreasing. We can repeat this to see that $T^ng(y)$ is strictly decreasing in y for any n. But as letting $n \to \infty$ yields the fixed point V(y), we have proved that V(y) is decreasing in y.⁷

So the value of being an active firm is strictly decreasing in y. But as exit is preferable when $W \ge V(y)$, there is a unique y, say y^* , above which exit is chosen. Now, as output is decreasing in y, we have proved that there is a minimum firm size, say q^* , below which a firm chooses to exit.

Let us take a brief digression here. Although it is perhaps less useful in this case, we can construct the fundamental reservation equations for this model. Replace the arbitrary function g(y') in (2.21) with V(y') so we can write

$$\begin{split} V(y) &= \pi(y) + \beta \int_{a}^{\infty} \max\left[W, V(y^{\,\prime})\right] dF(y^{\,\prime} \mid y) \\ &= \pi(y) + \beta \left[W - \int_{a}^{y^{*}} V_{y^{\,\prime}}(y^{\,\prime})F(y^{\,\prime} \mid y) dy^{\,\prime}\right], \end{split}$$

using (2.21) for the second line. At y^* , $V(y^*) = W$, so we may write

$$W = \pi(y^*) + \beta \left[W - \int_a^{y^*} V_{y'}(y') F(y' \mid y^*) dy' \right]$$

or

⁷ It is worth drawing attention to this remarkable proof, so the central trick is not overlooked. The contraction mapping theorem tells us that $V(x) = \lim_{n\to\infty} T^n g(x)$ for any appropriately bounded and continuous function g(x). So the trick is to choose a function that has the right properties to help pin down the properties of V(x). We saw in Chapter 4 an example in which we used this result to explicitly solve a model. Here the procedure is no help in solving the model, but it has turned out to be helpful in establishing a central property of the model.

APPLICATIONS OF DYNAMIC PROGRAMMING

$$W = \frac{\pi(y^*)}{1-\beta} - \frac{\beta}{1-\beta} \int_a^{y^*} V_{y'}(y') F(y' \mid y^*) dy', \qquad (2.22)$$

which is the fundamental reservation price equation. We have already seen that $V_{y'}(y') < 0$. Thus, the reservation value required for continued activity is *greater* than the discounted present value of receiving $\pi(y^*)$ forever. However, in this case, the fundamental equation does not yield an obviously intuitive result until we rearrange it slightly to get

$$\pi(y^*) - (1 - \beta)W = \beta \int_a^{y^*} V_{y'}(y')F(y' \mid y^*)dy' \le 0$$
(2.23)

Here $(1-\beta)W$ can be interpreted as the per-period fixed cost of being active in the industry. Equation (2.23) states that the firm does not exit until it is making possibly substantial losses. The firm does not exit as soon as losses fall to zero because it wants to preserve the option value of receiving the positive profits that would be secured by a run of good draws for y in the future.

We mentioned at the beginning of this subsection that there were important differences between the Hopenhayn and Jovanovic versions of the optimal stopping problem. We turn to that distinction now. Hopenhayn assumes that efficiency, x, is a random variable that fluctuates from period to period but exhibits persistence. That is, if $G(x \mid x)$ is the conditional distribution of x', then G is decreasing in x. Clearly if we expect x to be low today, we will also expect it to be low tomorrow. But this is no more than saying that the conditional distribution of expected efficiency, $F(y \mid y)$, is decreasing in y. Thus, we have already analyzed the heart of Hopenhayn's optimal stopping problem, and we can use this to think about survival probabilities. Consider a firm of current size $q > q^*$, where q depends negatively on y. The probability that $y' > y^*$ next period is increasing in yand therefore decreasing in q. Hence, the smaller a firm is today, the more likely it is to exit tomorrow. Equivalently, the probability of survival is increasing in current size and, moreover, size is a sufficient statistic for survival.

The difficulty with Hopenhayn's result is that the empirical evidence suggests a more complex empirical relationship between survival and observable firm characteristics. Dunne, Robert and Samuelson (1989), in particular, have shown that firm age also matters: conditional on size, younger firm are more likely to exit and conditional on age smaller firms are more likely to exit.. But in Hopenhayn, once one conditions on size, age does not matter.

In contrast, Jovanovic's model creates a role for age as well as size. Jovanovic also assumes that x is a random variable but, unlike Hopenhayn, there is no persistence in the shocks; x is subject an to i.i.d. shock in each period. However, the mean of x varies from firm to firm, and firms do not know their own mean. Put another way, they do not know how efficient they are on average, but must learn it from a sequence of noisy signals. Specifically, assume that $x=g(\eta)$, where g is a strictly increasing, strictly positive function. The parameter η is given by $\eta_t = \theta + \varepsilon_t$, where ε_t is a draw from a normal distribution with mean zero and variance σ_{ε}^2 . The firm does not know its θ . Upon entry it only knows that it will be a draw from a normal distribution with mean $\overline{\theta}$ and variance σ_{θ}^2 . The firm observes its costs each period, and this allows it to update its beliefs about what θ is. The calculations will not be shown here; instead we shall just note here that after T periods the firm's beliefs about its θ are that it is a draw from a normal distribution with mean $\overline{\eta} = T^{-1} \sum_{t=1}^{T} \eta_t$ and variance $\hat{\sigma}^2 = \left(\sigma_{\varepsilon}^2 \sigma_{\theta}^2\right) / \left(T \sigma_{\theta}^2 + \sigma_{\varepsilon}^2\right)$. That is, the age of the firm and its past mean efficiency are all we need to know to describe the firm's beliefs about its average efficiency. We also see that the variance declines with age, as the firm becomes more confident about what its true efficiency is. The conditional distribution for expected efficiency can therefore be written as $F(y \mid y, T)$. Intuitively, if the firm has received a lot of signals causing it to believe x is high, it will also believe that x' will be high. Thus $F(y \mid y, T)$ is decreasing in y, as maintained up to now. But the variance of the subjective beliefs about θ , and thus the variance of y, is greater for younger firms.



Figure 2.5.

This says nothing more than the fact that a firm with little information is likely to revise its beliefs significantly when it receives more information.

So now consider a firm with expected efficiency $y > y^*$. The younger firm is more likely to revise y drastically, and hence is more likely to draw a $y' > y^*$. However, what this means for the relationship between survival and age turns out to be a more complicated story. We have already established that $\pi(y)$ is a convex, decreasing function. But then V(y) is also convex, as drawn in Figure 2.5, and this implies that the stopping point y^* depends monotonically on age. In Figure 2.5, V(y) corresponds to the value of continuing in the industry for one more period regardless of the optimal decision. An older firm, with a relatively small variance on expected efficiency has a stopping efficiency of y_1^* . A younger firm with a greater variance has a stopping efficiency $y_2 > y_1^*$. Thus, given a current expected efficiency y, it is true that any y' > y is more likely to be attained by a younger firm. However, offsetting this is the fact that the y' that must be attained to induce exit is further away for the younger firm. That two effects of age offset each other implies that one cannot make general statements about *how* age affects survival (See Fig-



FIGURE 2.6. Density of y/ condition on y. Shaded area gives probability that the firm will exit next period. The probability may be larger or smaller for the young firm.

ure 2.6). However, one *can* make the claim that among firms that exit, younger firms will be smaller.

3. Continuous Choice Models

 $Consumption \ Problems$

Investment Problems

4. Transaction Costs

(S,s) Inventory Problems

Vintage Capital

The Used Car Market

Empty Factories

5. Time Inconsistency

Notes on Further Reading

References

- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson (1989): "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4):671-698.
- Hopenhayn, Hugo A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 60(5):1127-50.
- Jovanovic, Boyan (1982): "Selection and the Evolution of Industry." *Econometrica*, **50**(7):649-670.
- Mortensen, D.T. and C.A. Pissarides (1999): "New developments in models of search in the labor market." In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 3B, North-Holland, Amsterdam.
- Ross, Sheldon M. (1983): Introduction to Stochastic Dynamic Programming. San Diego: Academic Press.
- Rothschild, Michael (1974): "Searching for the lowest price when the distribution of prices is unknown." *Journal of Political Economy*, **82**(4):689-711.
- Stigler, George J. (1961): "The economics of information." Journal of Political Economy, 69:213-25.

Applications of Dynamic Programming

203