A Theory of Slack

How Economic Slack Shapes Markets, Business Cycles, and Policies

Pascal Michaillat

Draft version: January 2026

CHAPTER 9.

Market efficiency and inefficiency

We have now built a slackish model of markets, and studied its positive properties: what are the price, slack, and output prevailing in the market, and how these variables respond to supply and demand shocks. In this chapter we turn to the normative, welfare-related properties of the model: What is the efficient market allocation? Can we expect the market to operate efficiently? If not, how far from efficiency might the market be?

What do we mean by efficient? The efficient allocation is the allocation that maximizes the social welfare generated by the market. Knowing the efficient allocation is critical for designing the best possible stabilization policies (as discussed in part IV).

Before we begin we must decide which model to use for the welfare analysis. We have introduced a range of slackish models: a static model with exogenous capacity; a static model with endogenous capacity; a dynamic model with exogenous capacity, which itself could be extended with endogenous capacity. We have also seen that these models can be built around a number of matching functions—as long as they satisfy a few assumptions—and a wide array of price norms. Which one should we pick?

It turns out that we do not need to answer this question, because we adopt the sufficient-statistic approach to welfare analysis. We simply solve the problem of a social planner who is trying to maximize welfare in a more general framework. Typically, being more general adds complexity. Here, however, the generalization cleans things up and highlights the mechanisms and forces at play better.

Usually, welfare analysis is structural. This approach comes with two limitations that the sufficient-statistic approach resolves. First, the structural approach assumes the entire structure of the market and studies welfare and policy within this structure. The main limitation of this approach is that the market structure is often assumed religiously but largely determines the policy conclusions. The policy insights are in many ways baked into the initial assumptions about model structure. Second, the structural approach requires calibrating all the parameters in the model even though many of these parameters are generally not observable in the real world.

With the sufficient-statistic approach, we do not need to assume the entire structure of the market. Instead, we make a minimal set of assumptions on the welfare function and the structure of the market that allow for welfare analysis. This is beneficial because, unlike the structural approach, it allows the analysis to be applied to a range of models: any model that satisfies the minimal set of assumptions. Additionally, from a practical side, since we are making fewer assumptions, it is more likely that our analysis is valid in the real world. At the end of the chapter, we will apply the general analysis to each of the models that we have introduced in the book, and we will map the sufficient-statistic results into structural results. Another motto of the sufficient-statistic approach is to express the results in terms of statistics that can be estimated in the data, so theory and data are closely connected.

9.1. Market planner's problem

We consider a social planner who aims to allocate the goods supplied to the market efficiently. The social planner allocates some of the goods for consumption, some for matching, and leaves some of them unsold. The planner's objective is to maximize social welfare. What should the planner do?

9.1.1. Market slack

We assume that there is slack on the market, so a share $u \in (0,1)$ of all goods available on the market are unsold. The rest of the goods are sold to buyers.

9.1.2. Beveridge curve

We also assume that the market features a Beveridge curve. This means that the visit rate is a decreasing and convex function $v(u)$ of the slack rate. This is the sole structural assumption that we make in the welfare analysis. We do not need to assume a matching function or price norm; we only need the reduced-form Beveridge curve.

Why do visits enter the planner's problem? Because visits require resources, measured by the matching cost, $\kappa > 0$. The matching cost is the number of goods that have to be devoted to any one visit.

The Beveridge curve says that as the planner reduces the slack rate, the visit rate increases. So as the number of unsold goods falls, more goods must be allocated to visits and matching buyers and sellers. This is the central tradeoff that the planner must resolve: how to balance unsold goods against goods allocated to matching.

9.1.3. Market welfare

Market welfare measures the well-being generated by market activity. In our slackish market model, market welfare is not very complicated. It starts from the market capacity, which is k goods.

First, not all goods are sold. The slack rate is u, so only $(1-u)k$ goods are sold to buyers. The uk goods that remain unsold have no social value.

However, not all $(1-u)k$ sold goods are consumed. Some goods are instead devoted to matching and, as they are not consumed, they do not produce direct social value. Thus, we must net out the amount of goods devoted to matching: we need to subtract $\kappa v k$ goods from the $(1-u)k$ goods sold to buyers. (Since v is the visit rate, vk is the number of visits in the market.)

Overall, $(1-u)k - \kappa v k$ goods are consumed through market purchases, so market welfare is

$$[1 - u - \kappa v]\, k. \tag{9.1}$$

9.1.4. Market planner's problem

We are now in a position to state the market planner's problem. We begin by plugging the Beveridge curve into market welfare, to write market welfare as a function of the slack rate:

$$\mathcal{M}(u) = [1 - u - \kappa v(u)]\, k. \tag{9.2}$$

The market planner chooses a slack rate $u \in (0,1)$ to maximize market welfare $\mathcal{M}(u)$. The slack rate u^* that maximizes market welfare is the efficient slack rate. The efficient slack rate is the best slack rate from a social perspective.

9.2. Solution to the planner's problem

We now determine the efficient slack rate, u^*. Ideally the planner would prefer to have no unsold goods and no goods allocated to matching, but that is not feasible in a slackish market. The Beveridge curve imposes that some goods are always unsold, and some visits are always required so some buyers and sellers match.

9.2.1. Analytical solution

We start by analytically determining the efficient slack rate, u^*. Because the Beveridge curve $v(u)$ is convex, the welfare function $\mathcal{M}(u)$ is concave.[1] Hence, the first-order condition is sufficient to find the unique maximum of the welfare function.

To find the efficient slack rate, we simply set the derivative of the welfare function to zero: $\mathcal{M}'(u) = 0$. Thus, the efficient slack rate satisfies

$$0 = -\left[1 + \kappa v'(u)\right] k.$$

This means that the efficient slack rate is implicitly defined by

$$v'(u) = -\frac{1}{\kappa}. \tag{9.3}$$

Efficiency requires the slope of the Beveridge curve equals $-1/\kappa$. There is a tradeoff between too much slack and too little slack, both of which are undesirable, modulated by the Beveridge curve. The efficient slack rate guarantees that there is some slack to keep resources devoted to matching low, but not too much to keep the number of unsold goods low.

Formally, formula (9.3) says that when the market operates efficiently, the welfare costs and benefits from moving one good from sold to unsold are equalized. When one good is not sold anymore, consumption drops by 1 good. Having one more good unsold also means having $-v'(u) > 0$ fewer visits. Each visit reduces consumption by the matching cost, κ, so consumption increases by $-v'(u)\kappa > 0$ goods. When costs and benefits are equalized, we have $1 = -v'(u)\kappa$, which is equivalent to (9.3).

9.2.2. Graphical solution

Next, we graphically determine the efficient slack rate (figure 9.1A). To do that, we introduce isowelfare curves:

$$v = W - \frac{1}{\kappa} \cdot u, \tag{9.4}$$

where $W > 0$ governs the welfare level along the curve. On any isowelfare curve, the visit rate and slack rate are such that social welfare does not change.

Let's imagine that the market planner could place the market on any isowelfare curve. Which one would they choose? Well, the planner would want to be on an isowelfare curve as close to the origin as possible. Points closer to the origin are characterized by

[1] We see in (9.2) that $\mathcal{M}(u)$ is the sum of two concave functions and therefore concave. The first function is $[1 - u]\, k$, which is linear and thus concave. The second function is $-v(u)\kappa k$, which is the opposite of a convex function and thus concave.

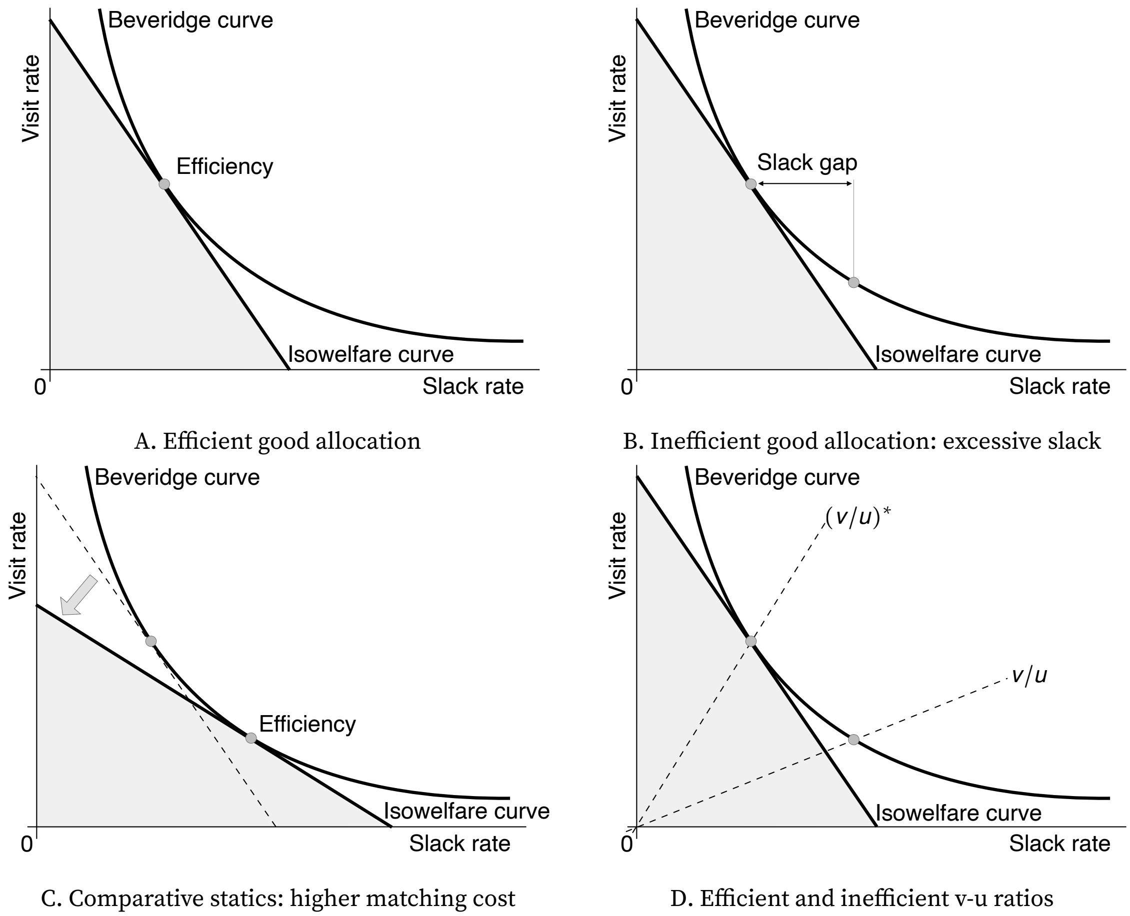

FIGURE 9.1. Efficiency and inefficiency in slackish market model

The isowelfare line is given by (9.4). The shaded area indicates all good allocations with higher welfare than the isowelfare line.

fewer visits and fewer unsold goods. Visits and unsold goods are costly, so the points on isowelfare curves that lie closer to the origin have higher welfare than the points that are further out.

However, the isowelfare curves that are deep inside, close to the origin, do not cross the Beveridge curve: they are not feasible. The market planner must respect the Beveridge curve so must choose a point that is also on the Beveridge curve. Accordingly, the optimal choice for the planner is to pick the isowelfare curve that is as far in as possible while still touching the Beveridge curve: the isowelfare curve that is tangent to the Beveridge curve. The tangent isowelfare curve has the highest welfare while respecting the feasibility constraint represented by the Beveridge curve.

Then, the efficient allocation is the point of tangency between the Beveridge curve and isowelfare curve. Since the slope of the Beveridge curve is $v'(u)$ and the slope of the isowelfare curve is $-1/\kappa$, we recover the efficiency condition (9.3) with our graphical approach.

At the point of tangency, the market operates efficiently; at other points on the Beveridge curve, it operates inefficiently. At any point below the efficiency point, for instance, slack is inefficiently high (figure 9.1B). This situation generates a positive slack gap, $u - u^* > 0$. As the market becomes tighter and tighter, at some point it crosses efficiency. If the market keeps tightening, the slack gap becomes negative, that is, $u - u^* < 0$. In this case, the slack rate is below the efficient slack rate.

9.2.3. Comparative statics

From equation (9.3), we see what happens if the matching cost changes. Given that the Beveridge curve $v(u)$ is convex, $v'(u)$ is increasing in u. If the matching cost κ goes up, then $-1/\kappa$ falls, meaning that the Beveridge curve must be flatter at efficiency. As a result, the efficient slack rate u^* increases. If the matching cost goes up, that makes the concern about too many goods devoted to matching more important, which tilts the balance in the tradeoff towards allowing a bit more slack.

We can also perform the comparative statics on the efficiency diagram (figure 9.1C). If the matching cost κ goes up, the isowelfare curve (9.4) becomes flatter. A flatter line means that the new tangency point occurs further out. Thus, the new efficient slack rate is now higher.

9.3. Sufficient-statistic formula for the efficient v-u ratio

There are two potential issues with the efficiency formula (9.3) that make it hard to implement in practice. First, the formula defines the efficient slack rate or visit rate only implicitly. If the formula holds, then the slack and visit rates are efficient, but their values are not given. Second, both in theory and in practice, the Beveridge curve is convex, so the slope of the Beveridge curve constantly changes with the tightness of the market. Measuring a slope in the data is hard enough: measuring the slope at any point in time to know if the condition is satisfied would be so hard as to be impractical.

To alleviate these issues, we reformulate the efficiency condition so that it involves model variables explicitly and so that it only involves stable statistics (unlike the Beveridge slope). We are looking for a formula in which we can plug market statistics that give us the efficient level of a market variable directly. Computing such formula is actually simple.

9.3.1. V-u ratio

Our first step in improving the efficiency formula is to introduce the v-u ratio, v/u. This is the ratio of the visit rate to the slack rate, which is also the ratio of the number of visits to the number of unsold goods. Furthermore, in the dynamic model of chapter 8, the v-u ratio coincides with the market tightness, $v/u = \theta$.

In general the market tightness is the ratio of the number of visits to the number of goods for sale in the matching function. In static models the number of goods for sale is the market capacity, while in dynamic models the number of goods for sale is the number of goods currently unsold, which is why the v-u ratio corresponds to the market tightness only in dynamic models. In any case, the rest of the section is about the v-u ratio.

9.3.2. Beveridge elasticity

Our second step is to introduce the Beveridge elasticity, which we denote β. The Beveridge elasticity is defined by

$$\beta = -\frac{u}{v} \cdot v'(u).$$

Hence the Beveridge elasticity is the elasticity of the Beveridge curve, normalized to be positive by multiplying it by -1. (Recall that the Beveridge curve is downward sloping so $v'(u) < 0$.)

9.3.3. Sufficient-statistic formula

To obtain a sufficient-statistic formula, we rework the efficiency condition (9.3). We start by multiplying both sides of (9.3) by $-(v/u)(u/v) = -1$. We obtain:

$$\frac{v}{u} \cdot \left[-\frac{u}{v} \cdot v'(u)\right] = \frac{1}{\kappa}.$$

The left-hand side is just the v-u ratio times the Beveridge elasticity, β. Dividing both sides by β, we therefore get:

$$\left(\frac{v}{u}\right)^* = \frac{1}{\beta\kappa}. \tag{9.5}$$

Equation (9.5) is our sufficient-statistic formula for market efficiency. It says that the efficient v-u ratio is determined by two sufficient statistics. The first sufficient statistic is the matching cost, κ. The key comparative static is that if the matching cost increases, then $(v/u)^*$ falls and u^* increases. If you have a higher matching cost, it is much more costly to visit shops, so the planner prefers to leave more goods unsold so fewer goods are devoted to matching and in the end more goods are consumed.

The second sufficient statistic is the Beveridge elasticity, β. This statistic matters because the social planner is trying to solve the tradeoff between visits and slack. When the elasticity of the Beveridge curve is high, it is costly to cut slack a bit, in the sense that a small reduction in slack requires many visits. Conversely, by increasing just a bit the slack rate, you can reduce the number of visits considerably. In other words, slack is more favorable when the Beveridge curve is steeper: a higher Beveridge elasticity leads to a lower $(v/u)^*$, so a slacker market.

This formula for market efficiency is quite simple given that it only involves two statistics. If a government wants to know the efficient v-u ratio in any market, they just need to know the Beveridge elasticity and matching cost in that market to determine the efficient market allocation.

The efficient v-u ratio is depicted in figure 9.1D, together with a v-u ratio that is inefficient. In the diagram, the v-u ratio is represented by the slope of any ray through the origin. The efficient slack and visit rates are then given by the intersection of the Beveridge curve with the ray whose slope is given by $(v/u)^*$.

9.4. Applying the formula to the static market

Formula (9.5) expresses the efficient v-u ratio as a function of two sufficient statistics. It is natural to ask how the efficiency condition would look like in specific models. The sufficient statistics do not necessarily correspond to parameters from the models, so it is unclear how efficiency might look like in the models. In the rest of the chapter we show how the formula applies to all the models presented so far in the book. In this section we apply the formula to the basic static model of chapter 5 .

9.4.1. Verifying the sufficient-statistic assumptions

We need to check that the static model satisfies all the assumptions made in the sufficient-statistic analysis.

First, welfare is solely determined by the consumption of buyers. Assuming that all buyers are the same to avoid distributional issues, welfare is solely determined by aggregate consumption. This in turn is just the number of goods for sale k minus unsold goods uk minus goods used for matching $\kappa v k$. Hence, welfare is proportional to $1 - u - \kappa v$, just as in the sufficient-statistic analysis.

Second, the model features a Beveridge curve. The slack rate is $u = 1 - f(\theta)$, where θ is the market tightness, which is just equal to the visit rate: $\theta = v$. So the definition of the slack rate implicitly defines a Beveridge curve:

$$u = 1 - f(v). \tag{9.6}$$

Given that the selling probability f is strictly increasing and concave, we learn that the slack rate u is a strictly decreasing and convex function of the visit rate v. Equivalently, the visit rate v is a strictly decreasing and convex function of the slack rate u—which means that the model produces a Beveridge curve. This can be seen graphically. The inverse of a function is obtained by reflecting its graph across the 45° line; this reflection preserves both monotonicity and convexity for strictly decreasing functions. Hence any strictly convex and decreasing function, once inverted, remains strictly decreasing and convex.[2] The bottom line is that the model admits a Beveridge curve: $v = v(u)$, where the function is strictly decreasing and convex. As an illustration, figure 9.2A displays the Beveridge curve obtained in the static model.

9.4.2. Applying the formula

All the sufficient-statistic conditions are satisfied, which tells us that formula (9.5) holds. Applying it to the static model, we get $(v/u)^* = 1/(\beta\kappa)$.

The final step is to express the Beveridge elasticity as a function of model parameters. Equation (9.6) implies that the elasticity of u with respect to v is $\epsilon^u_v = -f(v)/u \cdot (1-\eta)$, where η is the matching elasticity, and thus $1-\eta$ is the elasticity of f (see equation (4.7). Since $v = \theta$ and $u = 1 - f(\theta)$, the elasticity is $\epsilon^u_v = -(1-\eta) f(\theta)/[1 - f(\theta)]$. The Beveridge elasticity is defined by $\beta = -\epsilon^v_u = -1/\epsilon^u_v$. Therefore, the Beveridge elasticity relates to the

[2] This can also be shown analytically. Consider a twice-differentiable function $g(x)$ that is strictly decreasing and convex, so that $g' < 0$ and $g'' > 0$. By the inverse function theorem, its inverse has first derivative $(g^{-1})'(y) = 1/g'(g^{-1}(y))$ and second derivative $(g^{-1})''(y) = -1/[g'(g^{-1}(y))]^2 \cdot g''(g^{-1}(y)) \cdot (g^{-1})'(y) = -g''(g^{-1}(y))/[g'(g^{-1}(y))]^3$. From these expressions we see that $(g^{-1})' < 0$, so the inverse is strictly decreasing, and $(g^{-1})'' > 0$, so the inverse is strictly convex.

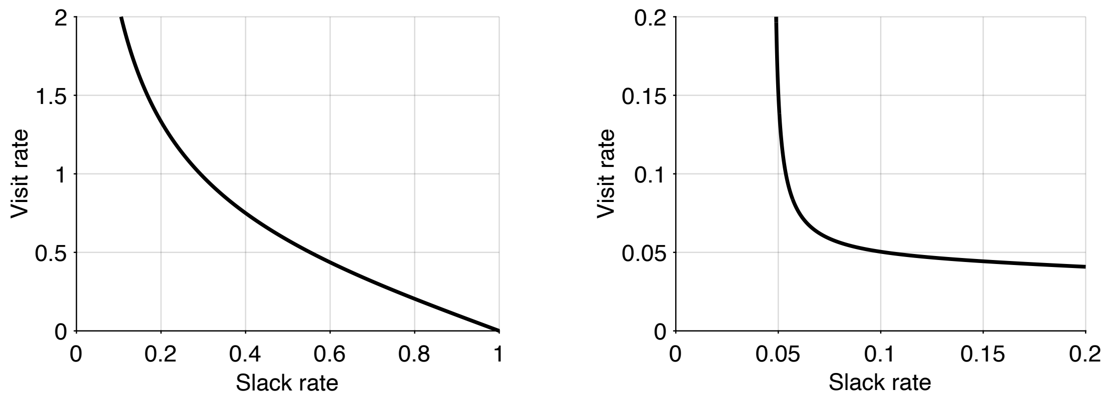

A. Static model

B. Dynamic model

FIGURE 9.2. Beveridge curve in slackish market models

In the static model, the Beveridge curve is given by (9.6). In the dynamic model, the Beveridge curve is given by (8.5). The matching function is CES, given by (4.10). Parameters are set to $\sigma = 2$ and $\lambda = 5\%$.

matching elasticity as follows:

$$\beta = \frac{1}{1-\eta} \cdot \frac{1-f(\theta)}{f(\theta)}.$$

Given the expression for the Beveridge elasticity, we obtain the efficiency condition in the static model. The sufficient-statistic formula becomes $(v/u)^* = f(\theta^*)/[1-f(\theta^*)] \cdot (1-\eta)/\kappa$, or since $v = \theta$ and $u = 1 - f(\theta)$, $\theta^*/f(\theta^*) = (1-\eta)/\kappa$, or using (4.5),

$$(1-\eta)q(\theta^*) = \kappa. \tag{9.7}$$

If the matching function is of Cobb-Douglas form, η is constant. For other matching functions, the matching elasticity is a function of tightness, as we saw in chapter 4. Then the left-hand side of the equation is a function of tightness, $[1-\eta(\theta)]q(\theta)$. In all the matching functions considered in chapter 4, $\eta(\theta)$ is an increasing function of θ, so $1-\eta(\theta)$ is a decreasing function of θ, just like $q(\theta)$, so that the left-hand side is strictly decreasing in tightness, implying that the efficiency condition implicitly determines a unique tightness.

9.4.3. Hosios condition

By looking closely, we see a final interesting property of formula (9.7): it implies the famous Hosios (1990) condition.

The Hosios condition says that in a market organized around a matching function, and with bargained prices, efficiency is achieved when the seller's bargaining power equals the matching elasticity:

$$\chi = \eta.$$

It turns out that this is just what formula (9.7) implies.

To see the connection with the Hosios condition, let's compute the matching wedge given by (9.7):

$$q(\theta^*) = \frac{\kappa}{1-\eta} \quad \text{so} \quad \tau(\theta^*) = \frac{\kappa}{\frac{\kappa}{1-\eta} - \kappa}.$$

The expression for $\tau(\theta^*)$ given by (9.7) simplifies to

$$\tau(\theta^*) = \frac{1-\eta}{\eta}. \tag{9.8}$$

This is the value that the matching wedge must take at efficiency.

But under bargaining, we saw in equation (6.5) that the matching wedge is equal to $\tau(\theta) = (1-\chi)/\chi$. Hence, for the market to operate efficiently under bargaining, it must be that $\chi = \eta$!

9.5. Applying the formula to the market with endogenous capacity

The sufficient-statistic analysis did not mention endogenous capacity, so it is natural to worry that the formula might not apply with endogenous market participation. In this section, we show that this worry is not warranted: by an envelope argument, the sufficient-statistic formula continues to apply even when market participation is endogenous, as in the model of chapter 7.

9.5.1. Validity of the formula with endogenous capacity

Is (9.7) still a valid efficiency condition when sellers enter endogenously? We now turn to the welfare function in this generalized framework. Note first that we do not need to take the utility from money into account since the stock of money is fixed and money just changes hands between buyers and sellers, who value it similarly. So we only need to take into account utility from consumption and disutility from participation.

Social welfare is the sum of individual utilities for all buyers and sellers. All buyers are the same by assumption (to avoid distributional issues) so their utility is just the utility from aggregate consumption: $\mathcal{B}(c(u)) = ac(u)^{1-\alpha}$. Aggregate consumption c is a function of the unemployment rate u given by

$$c(u) = [1 - u - \kappa v(u)]\, kh(u),$$

where $v(u)$ is the Beveridge curve. Aggregate consumption is obtained by subtracting matching costs $\kappa v k h(u)$ from output $(1-u)kh(u)$.

All sellers have different utility from nonparticipation, so the aggregate utility is obtained by integrating the individual utilities of all potential sellers who remain outside of the market: $\mathcal{S}(h(u)) = \int_{h(u)}^{1} \xi i^{\phi} di$, where the participation rate $h(u)$ comes from (7.2).

Overall, then, market welfare is a function of the slack rate: $\mathcal{M}(u) = \mathcal{B}(c(u)) + \mathcal{S}(h(u))$. The planner cannot do anything about matching on the market, so it takes the Beveridge curve as given. It cannot force people to enter the market or leave it, so it must respect people's voluntary entry or exit of the market, described by (7.2). What the planner can do is influence prices, for instance through taxes or subsidies or price caps, or purchase goods directly on the market. Through these actions it can determine the slack rate u.

The planner chooses the slack rate u to maximize market welfare $\mathcal{M}(u)$. The first-order condition for the maximization problem is $0 = \mathcal{M}'(u)$, which becomes:

$$0 = \mathcal{B}'(c) \cdot \left\{\left[-1 - \kappa v'(u)\right] kh(u) + \left[1 - u - \kappa v(u)\right] kh'(u)\right\} + \mathcal{S}'(h)h'(u).$$

Rearranging terms, we get

$$0 = \mathcal{B}'(c)kh(u) \cdot \left[-1 - \kappa v'(u)\right] + \frac{h'(u)}{h(u)} \cdot \left[\mathcal{B}'(c)c(u) + \mathcal{S}'(h)h(u)\right].$$

We now need to rework the second term in the first-order condition. Equation (5.14) tells us that

$$\mathcal{B}'(c) = a(1-\alpha)c^{-\alpha} = p\left[1 + \tau(\theta)\right] \quad \text{so} \quad \mathcal{B}'(c)c = p\left[1 + \tau(\theta)\right]c = py.$$

By Leibniz rule, we see that $\mathcal{S}'(h) = -\xi h^{\phi}$. And then using equation (7.2), we learn that

$$\mathcal{S}'(h) = -p(1-u)k \quad \text{so} \quad \mathcal{S}'(h)h = -p(1-u)kh = -py.$$

Hence, we have just established that

$$\mathcal{B}'(c)c(u) + \mathcal{S}'(h)h(u) = 0.$$

This result is key, and it hinges on an envelope-theorem logic. Because market participants behave optimally, they enter until the marginal cost from entering (foregone leisure) is just equal to the marginal benefit from entering, which is determined by the price of goods. But the price of goods also determines the marginal utility of consumption given that buyers behave optimally. All in all, when the planner changes the slack rate, it does affect participation, but this change in participation has no effect on welfare because participation has been chosen optimally by sellers.

Critically, we have just learned that the second term in the first-order equation is just

0. Therefore, the first-order condition reduces to $v'(u) = -1/\kappa$, just as in the baseline case with fixed participation. We therefore recover the result that the efficient v-u ratio is given by $(v/u)^* = 1/(\beta\kappa)$.

In sum, endogenizing market participation does not change the welfare analysis. This result stems from an envelope-theorem logic that is classic in public economics. Even if the social planner alters market participation by changing the slack rate, welfare is unaffected because the sellers who move in or out of the market are indifferent between participating or not. Indeed, if sellers strictly preferred participation to nonparticipation, they would move into the market. Conversely, if they strictly preferred nonparticipation, they would move out of the market.

9.5.2. Applying the formula

In the static market with endogenous capacity, the welfare function is different because it also includes the welfare of sellers, who decide whether to enter the market or not. What makes this model all the more intriguing is that the planner takes into account the fact that it can affect the number of sellers in the market by picking the tightness of the market, as equation (7.2) shows. Despite these differences, we have seen that the v-u ratio remains given by the sufficient-statistic formula (9.5) in this model.

How does the sufficient-statistic formula (9.5) look like in the model? Since the Beveridge curve is the same as in the basic market—obtained from the definition of the slack rate, $u = 1 - f(v)$—the sufficient-statistic formula also translates to (9.7).

9.6. Applying the formula to the dynamic market

Finally, we apply the sufficient-statistic formula to the dynamic model of chapter 8. Although we did not mention any dynamics in the welfare analysis, the formula remains valid here because the dynamic model features a Beveridge curve too.

9.6.1. Sufficient-statistic formula for efficient market tightness

In a dynamic model, the v-u ratio is the market tightness: $\theta = v/u$. Thus formula (9.5) gives the efficient market tightness as a function of two sufficient statistics:

$$\theta^* = \frac{1}{\beta\kappa}. \tag{9.9}$$

Most quantitative applications should use a dynamic model—as the real world is dynamic, with long-term relationships. This efficient-tightness formula is useful in that context.

9.6.2. Verifying the sufficient-statistic assumptions

We now check that our dynamic model satisfies all the assumptions made in this chapter's welfare analysis. First, welfare is solely determined by the consumption of buyers, just as in the static case.

Second, the model features a Beveridge curve, given implicitly by (8.5). In chapter 8 we established that the Beveridge curve $v(u)$ is downward sloping but we did not check its convexity, which is key for the welfare analysis. We can do that here. We need to differentiate twice (8.5). The first implicit differentiation gives (8.6). We now implicitly differentiate that equation again with respect to u and obtain

$$\frac{\partial^2 m}{\partial u^2} + \frac{\partial^2 m}{\partial v^2} \cdot \left[v'(u)\right]^2 + \frac{\partial m}{\partial v} \cdot v''(u) = 0.$$

Reshuffling the terms to isolate the second derivative of the Beveridge curve, we obtain:

$$v''(u) = -\frac{\frac{\partial^2 m}{\partial u^2} + \frac{\partial^2 m}{\partial v^2} \cdot \left[v'(u)\right]^2}{\frac{\partial m}{\partial v}}.$$

Since the matching function is increasing in both arguments, $\partial m/\partial v > 0$. Since the matching function is concave in both arguments, $\partial^2 m/\partial u^2 < 0$ and $\partial^2 m/\partial v^2 < 0$. Of course, $\left[v'(u)\right]^2 > 0$. Hence the numerator of the fraction is negative, its denominator is positive, so with the minus sign, the whole expression is positive, which tells us what we wanted to know: $v''(u) > 0$. Thus, the Beveridge curve in the dynamic model is convex.

As an illustration, figure 9.2B displays the Beveridge curve obtained in the dynamic model under our usual calibration. Just like the market supply, the Beveridge curve is much more convex in the dynamic model than in the static model.

9.6.3. Applying the formula

All the sufficient-statistic conditions are satisfied, so formula (9.5) holds. Applying it to the dynamic model, we get $\theta^* = 1/(\beta\kappa)$. The final step is to link the Beveridge elasticity β to model parameters.

From equation (8.5), which implicitly defines the Beveridge curve, we computed the slope of the Beveridge curve, and found that it was given by equation (8.7). We now rework this equation to compute the Beveridge elasticity in the dynamic model. We multiply both sides of equation (8.7) by $-u/v$, and we obtain:

$$-\frac{u}{v}v'(u) = \frac{\lambda \cdot \frac{u}{m} + \frac{u}{m} \cdot \frac{\partial m}{\partial u}}{\frac{v}{m} \cdot \frac{\partial m}{\partial v}}.$$

We divided numerator and denominator of the right-hand side fraction by the number of matches m to make the elasticities of the matching function appear.

We now simplify this expression. First, the left-hand side is just the Beveridge elasticity β, by definition. The denominator of the right-hand side fraction is just the elasticity of the matching function with respect to the number of visits. This is just $1-\eta$, where η is the matching elasticity, as we established in (4.2). The second term in the numerator is just η, the matching elasticity. Finally, by definition, $m/u = f(\theta)$ is just the selling rate, so the first term in the right-hand numerator boils down to $\lambda/f(\theta)$.

All in all, the equation gives the following expression for the Beveridge elasticity:

$$\beta = \frac{1}{1-\eta}\left(\eta + \frac{\lambda}{f(\theta)}\right).$$

Just as in the static case, the Beveridge elasticity is closely related to the matching elasticity, η. However, because the slack rate is generally small, the term $\lambda/f(\theta)$ is much smaller than η, so the Beveridge elasticity given by the dynamic model is likely to be stable even if slack varies over time. This is different from what happens in the static model, where the Beveridge elasticity varies strongly with slack. We therefore see in the dynamic model why formula (9.5) is more usable than the original formula (9.3): because the statistics that it involves are stable over time, so they do not need to be constantly re-estimated.

Finally, we express efficiency condition (9.5) in terms of model parameters:

$$\theta^* = \frac{1-\eta}{\eta + \lambda/f(\theta)} \cdot \frac{1}{\kappa}.$$

Accordingly, the efficient tightness θ^* is implicitly defined by

$$\eta\theta^* + \frac{\lambda}{q(\theta^*)} = \frac{1-\eta}{\kappa},$$

where $q(\theta) = f(\theta)/\theta$ is the buying rate. The left-hand side of the equation is continuous and strictly increasing from 0 to ∞ when θ goes from 0 to ∞. Since the right-hand side is a positive number, the equation admits a unique solution.

This equation links the efficient tightness to model parameters. It is similar, albeit more complicated, than equation (9.7). It depends on the matching elasticity η and matching cost κ, just like in the static case, but also involves the separation rate λ.

9.7. Modulating the social cost of slack

So far we have assumed that unsold goods are pure waste: they have no social value or cost besides their wastefulness.

This is a valid assumption for services that are not rendered or goods that are perishable. But if, for instance, unsold goods can be donated, these goods would have some social value.[3] On the other hand, if unsold goods are sent to landfills where they produce noxious gases while decomposing, these goods impose a social cost.[4]

Similarly, the assumption seems appropriate for employed workers that are idle. But if an unemployed worker engages in home production, their time might have social value. If on the other hand the time spent unemployed is painful psychologically, their time generates a social cost.

To allow for the possibility that unsold goods have some social value or cost beyond the waste they impose, we introduce a statistic $\zeta < 1$, which captures the social value of unsold goods compared to sold goods. The case $\zeta = 0$, which we have considered so far, is when unsold goods are pure waste. The case $\zeta < 0$ is when unsold goods generate a cost on top of being wasteful. And the case $\zeta \in (0, 1)$ is when unsold goods have some social value that alleviates the wastefulness of unsold goods.

In this extension, the market welfare has an additional element. Since each of the uk unsold goods have a social value ζ relative to sold goods, an amount ζuk must be added to market welfare (9.1). In total, market welfare becomes

$$\mathcal{M}(u) = [1 - \kappa v(u) - (1 - \zeta)u]\, k.$$

Given that the welfare cost of a good failing to sell is reduced by $1 - \zeta$, the first-order condition (9.3) generalizes to

$$v'(u) = -\frac{1 - \zeta}{\kappa}. \tag{9.10}$$

If slack has some social value ($\zeta > 0$), the isowelfare curve is flatter in the efficiency diagram, so the tangency point occurs further out, at a higher efficient slack rate (just like in figure 9.1C). Along the Beveridge curve, the efficient slack rate u^* increases. Intuitively, when the cost of having unsold goods is a bit less, the efficient allocation of goods tolerates a bit more slack. Conversely, if slack imposes a social cost beyond its wastefulness ($\zeta < 0$), the isowelfare curve is steeper, so the efficient slack rate is lower.

Finally, the sufficient-statistic formula (9.5) generalizes to

$$\left(\frac{v}{u}\right)^* = \frac{1 - \zeta}{\beta \kappa}.$$

In this generalized formula, a third sufficient statistic appears: the social value of unsold goods, ζ. The efficient v-u ratio responds to changes in the social value of unsold goods just like to changes in the matching cost κ: when the matching cost increases or when the

[3] In the United States in 2023, retailers donate 12% of donatable unsold food to food banks (ReFED 2025).

[4] In the United States in 2023, about 25 million tons of surplus food end up in landfill, emitting a sizable amount of methane (ReFED 2025).

social value increases, the efficient v-u ratio decreases, so the efficient slack rate increases.

9.8. Are markets efficient?

The last question we ask in this chapter is: Can we expect markets to operate efficiently? One of the key theorems of neoclassical economics is that Walrasian markets operate efficiently. But slackish markets are different: there is absolutely no guarantee that a slackish market operates efficiently.

From a theoretical perspective, the key reason why Walrasian markets are efficient while slackish markets are not is that prices are set differently. In a slackish market, there exists a price norm that guarantees efficiency, but this price norm might not be adopted in the real world. Under all other price norms, efficiency is not guaranteed. In effect, there is no guarantee that the price norm that prevails in the real world leads to efficiency.

This is maybe easiest to see in the basic model of chapter 5. The market tightness is given by the supply-equals-demand condition, (5.22). When we solve the model, we use the price given by the price norm, plug it into the market demand, and determine the market tightness that equalizes supply and demand. Here we use the condition in the other way: we set tightness to its efficient level θ^*, equivalently given by (9.7) or (9.8). Then we use the supply-equals-demand condition to find the efficient price p^*, defined by

$$y^d(\theta^*, p^*) = y^s(\theta^*).$$

We know that the market demand is strictly decreasing in the price, so that equation defines a unique efficient price p^*. Furthermore, by using the functional forms for market supply and demand, given by (5.16) and (5.6), we easily obtain an expression for the efficient price. The supply-equals-demand condition requires:

$$\left[\frac{(1-\alpha)a}{p^*}\right]^{1/\alpha}\left[1+\tau(\theta^*)\right]^{1-1/\alpha} = f(\theta^*)k.$$

From (9.8) we see that $1+\tau(\theta^*) = 1/\eta$. By reshuffling terms, we obtain the following expression for the efficient price:

$$p^* = \frac{a}{k^\alpha}\cdot\frac{(1-\alpha)\eta^{1-\alpha}}{f(\theta^*)^\alpha}.$$

It is also possible to express the efficient price solely as a function of model parameters. We do that by removing the efficient tightness θ^* from the equation, using (9.7):

$$p^* = \frac{a}{k^\alpha}\cdot\frac{(1-\alpha)\eta^{1-\alpha}}{f\left(q^{-1}\left(\frac{\kappa}{1-\eta}\right)\right)^\alpha}.$$

An interesting aspect of the efficient price is that it is a flexible case of the price norm (6.1) introduced in chapter 6. Indeed, we obtain the price p^* from the price norm by setting the price-rigidity parameter to $\gamma = 0$ and the price level to

$$\rho = \frac{(1-\alpha)\eta^{1-\alpha}}{f\left(q^{-1}\left(\frac{\kappa}{1-\eta}\right)\right)^{\alpha}}.$$

We argued earlier in the book that bargaining was a possible mechanism pushing prices towards flexibility. We see here that efficiency is another one. Any social or economic mechanism pushing markets towards efficiency makes prices more flexible. In chapter 14 we will introduce such a market mechanism: Moen (1997)'s directed search.

Because market inefficiency is generic, government interventions might be needed to bring markets and the economy closer to efficiency—as we discuss in part IV. Here there is no guarantee that the invisible hand maintains markets at efficiency. Instead, the hand is generally inefficient: it does not ensure that markets operate efficiently on their own. The reason is that with slack, there is no mechanism that keeps prices at their efficient level, and so there is generally too much or too little slack on any market.

9.9. Summary

In this chapter we define efficiency as the level of slack that maximizes social welfare. Rather than specifying a complete structural model, we adopted a sufficient-statistic approach: efficiency can be analyzed in any market that admits a Beveridge curve, and results can be expressed in terms of observable statistics.

We found that the efficient slack rate equates the slope of the Beveridge curve to the inverse of the matching cost. From this, we obtained a sufficient-statistic formula that links the efficient v-u ratio to the matching cost κ and Beveridge elasticity β: $(v/u)^* = 1/(\beta\kappa)$.

Finally, we argue that slackish markets are generically inefficient, since no market mechanism ensures that prices adjust to their efficient level. In practice, price norms might generate either excessive or insufficient slack, justifying policy interventions that move markets toward their efficient operating points.

Bibliography

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies* 57 (2): 279–298. https://doi.org/10.2307/2297382.

Moen, Espen R. 1997. "Competitive Search Equilibrium." *Journal of Political Economy* 105 (2): 385–411. https://doi.org/10.1086/262077.

ReFED. 2025. "From Surplus to Solutions." US Food Waste Report. https://perma.cc/3ZLP-QKHN.