Chapter 6

Welfare

The preceding chapter presented examples where the competitive equilibrium allocation is the same as the solution to a social planner's problem. These examples illustrate the First Welfare Theorem of economics, which gives conditions under which a competitive equilibrium is Pareto optimal. The First Welfare Theorem is a remarkable result—arguably one of the most profound insights in economic theory. Market mechanisms can effectively coordinate the activities of large numbers of consumers and firms, despite being separated across space and time. Amazingly, under the conditions of the First Welfare Theorem (FWT), the actions of these disparate and heterogeneous people are coordinated in a way that is socially optimal even though each individual is simply pursuing their private interests. This is Adam Smith's invisible hand.[1]

Under the FWT, a competitive equilibrium is Pareto optimal. We focus on the concept of Pareto optimality because it is widely used as a "minimal" welfare criterion. It is, in particular, silent on the distributional implications of an allocation. For example, an allocation with extreme inequality can nevertheless be Pareto optimal. At times, macroeconomists assume more specific social welfare functions that express a preference for more even distributions of resources and we will comment upon this briefly.

While markets can work very well, they can also fail, which is to say that the competitive equilibrium may not be Pareto optimal. Such failures of the FWT can arise for a number of reasons, which we will discuss in detail. Our aim is to understand when market economies work well and when they do not. Indeed, much of modern macroeconomics is concerned with understanding the potential market failures that affect the economy and understanding the public policies that might lead to better outcomes. An understanding of the nature of market failure and how it can be corrected is an important part of evaluating the potential benefits of policy proposals.

After first reviewing the relationship between Pareto optimality and competitive equilibrium (the FWT) in an abstract, general setting, we map it into the more applied macroeconomic settings used in our text. We then go through the most common market failures that arise in macroeconomic analysis and indicate how each of these failures involve departures from the assumptions underlying the welfare theorem. We end the chapter with a short discussion of the role of government policy.

[1] An inspiring description of the power of markets can be found in the short video available at `https://www.youtube.com/watch?v=67tHtpac5ws`

6.1 The First Welfare Theorem

The First Welfare Theorem says that a competitive equilibrium is Pareto optimal. To make this statement more precise and to give a sketch of the proof we will consider an abstract economy with many different goods. These goods could be different commodities or goods at different dates. Later we will describe how different specific economic environments relate to this abstract economy.

There is a set $\mathcal{I}$ of different consumers, each indexed by $i \in \mathcal{I}$, and a set $\mathcal{J}$ of firms, each indexed by $j \in \mathcal{J}$. We assume $\mathcal{I}$ and $\mathcal{J}$ are finite, although we will later consider extensions in which they are infinite. Let x be a vector of consumption levels of different goods that are traded. The length of the vector specifies how many markets there are in the economy. For now, we will think of the vector length as finite, but it will be relevant to consider the possibility of infinite vectors later. Moreover, the vector is, at least for now, allowed to have negative elements. Similarly, y is a vector of production levels of the same goods and ω is a vector of exogenous endowment levels of the goods. Consumer i is endowed with ω_i and consumes x_i while firm j produces y_j. Thus, our resource constraint reads

$$\sum_{i \in \mathcal{I}} x_i \leq \sum_{j \in \mathcal{J}} y_j + \sum_{i \in \mathcal{I}} \omega_i. \tag{6.1}$$

This inequality applies to each element of the vector so, for every good, the total amount consumed cannot exceed the endowment plus the amount produced by all the firms.

Let X_i be a set of vectors that are feasible for agent i to consume. For example, if there are two commodities, then $X_i = \mathbb{R}^2_+$ would rule out consuming a negative amount of either good. Firms have production possibility sets denoted by Y_j. We assume that consumers maximize utility taking their budget constraint, prices, and consumption possibility sets as given. Utility maximization means they choose a consumption bundle that is not preference-dominated (as defined by a preference ordering $\succeq_i$) by any other choice available to them. Similarly, firms maximize profits, taking prices as given, by choosing a production plan in their production possibility set Y_j. Firm profits are given to the consumers that own the firms with $\theta_{i,j}$ denoting consumer i's share of firm j. Each firm is wholly owned by the consumers so $\sum_i \theta_{ij} = 1$. Finally, let p be the vector of prices for each good.

Key assumption We will use the weakest assumption under which the First Welfare Theorem holds: local non-satiation (LNS). LNS expresses that each consumer can always be made better off by an infinitesimally higher consumption of some good.[2]

A competitive equilibrium A competitive equilibrium is a consumption allocation $\{x_i^\star\}_{\forall i}$, a production allocation $\{y_j^\star\}_{\forall j}$, and a price system $p^\star$ such that

1. for each $i \in \mathcal{I}$, the consumption choice $x_i^\star$ is in X_i and there is no $x \in X_i$ such that $x \succ_i x_i^\star$ and $px \leq px_i^\star = p\omega_i + \sum_j \theta_{i,j} p y_j$;

2. for each $j \in \mathcal{J}$, $y_j^\star \in Y_j$ and there is no $y \in Y_j$ such that $py > py_j^\star$;

[2] This notion also presumes that the consumption possibility sets X_i allow small movements in at least one desirable direction.

3. and the market for each good clears (equation (6.1) holds with equality).

Theorem 6.1 (The First Welfare Theorem) *An allocation that is part of a competitive equilibrium is Pareto optimal.*

Proof. The proof is by contradiction. So suppose there exists an allocation $\{\tilde{x}_i\}_{\forall i}$, $\{\tilde{y}_j\}_{\forall j}$ that is feasible (these values are all in their possibility sets and the resource constraint is satisfied) and that Pareto dominates the allocation of the given equilibrium. Then, by definition,

$$\forall i \in \mathcal{I} : \tilde{x}_i \succeq_i x_i^\star,$$
$$\exists \tilde{i} \in \mathcal{I} : \tilde{x}_{\tilde{i}} \succ_{\tilde{i}} x_{\tilde{i}}^\star.$$

Now the first property of a competitive equilibrium combined with LNS can be used to conclude that

$$p\tilde{x}_i \geq px_i^\star$$

holds for all i. If it were strictly cheaper to buy the new allocation, the consumer could have spent more to improve on the existing choice and hence that choice could not have been optimal for the consumer. Moreover, for $\tilde{i}$ it must be that

$$p\tilde{x}_{\tilde{i}} > px_{\tilde{i}}^\star,$$

again because otherwise the original allocation must not have involved consumer optimization.

Summing the budget constraints of all consumers we have

$$p\sum_{i\in\mathcal{I}} x_i^\star = p\omega + p\sum_{j\in\mathcal{J}} y_j^\star,$$

where $\omega \equiv \sum_i \omega_i$ and we have used the fact that $\sum_i \theta_{ij} = 1\ \forall j$. By the arguments above, the alternative consumption allocation is more expensive

$$p\sum_{i\in\mathcal{I}} \tilde{x}_i > p\omega + p\sum_{j\in\mathcal{J}} y_j^\star.$$

Furthermore since $y_j^\star$ is profit maximizing for each j, it must be that $\sum_j p\tilde{y}_j \leq \sum_j py_j^\star$ so we have

$$p\sum_{i\in\mathcal{I}} \tilde{x}_i > p\omega + p\sum_{j\in\mathcal{J}} \tilde{y}_j.$$

In addition, the alternative allocation is resource feasible. Multiplying equation (6.1) by p we obtain

$$p\sum_{i\in\mathcal{I}} \tilde{x}_i \leq p\omega + p\sum_{j\in\mathcal{J}} \tilde{y}_j.$$

These last two inequalities contradict one another. ■

Mapping the setting into our macroeconomic models The power of the abstract proof above is that it can be applied to a large number of contexts. Let us first consider a static macroeconomic model with one agent with endowments of capital and labor equal to k and 1, and a neoclassical, constant-returns to scale production function that produces a consumption good. Preferences could be represented by a standard utility function, which satisfies LNS because it is strictly increasing. Then $x \in \mathbb{R}^3_+$ and a typical consumer choice would be $(c, 0, 0)$. Firms would have $(y, -k, -l) \in \mathbb{R}_+ \times \mathbb{R}^2_-$, the endowment vector would be $\omega = (0, k, 1)$, and the normalized price vector would be $(1, r, w)$.

The argument can easily be extended to include different types of goods. For example, if consumers value leisure we can simply treat them as "buying leisure" so that is just another good for them to consume. Other intermediate goods can be included too, as well as other resources, such as land. The theorem also applies to endowment economies by specifying the production possibility sets to only include the zero vector.

In a dynamic setting with $T < \infty$ time periods, goods and services at different dates are simply additional market goods and all the vectors become correspondingly longer. One can, in particular, define the vectors as simply containing T times the number of elements in the static model. As we will see below, the argument also applies to infinite horizon models but with some additional caveats. In the next chapter we will discuss models with uncertainty where goods are indexed not just by time, but also by the state of the world. The First Welfare Theorem applies to those settings as well.

Infinite horizon models We start by considering a case with a finite number of infinitely-lived consumers. Later in the chapter we will discuss an overlapping-generations economy where time is infinite but each consumer has a finite life.

All the vectors in the proof above are infinite-dimensional but the competitive equilibrium is defined as before. A key point here is that for an equilibrium to exist—especially for it to satisfy its first property—it would have to be that the value of total expenditures of the consumer is finite, i.e., the dot product of the infinite price vector and the infinite quantity vector is finite. For an infinite sum, this would require, unless quantities go to zero, that prices fall sufficiently fast. Typically in our models, prices fall asymptotically at a geometric rate (equal to $\beta < 1$ if there is no growth), which for constant quantities imply a finite total value. As we shall see below, although this property holds in dynastic models, it does not hold in some overlapping-generations models. This means that overlapping-generation models can have very different welfare properties.

Intuition One issue underlying the First Welfare Theorem is the fact that all consumers and firms face the same prices. For all consumers the ratio of the marginal utilities of any two goods will equal the ratio of prices and, as the prices are the same, they all have the same ratio of marginal utilities. It is therefore impossible to make a marginal change in the consumption allocation that results in a Pareto improvement. This point can be illustrated with an Edgeworth box as shown in the top panel of Figure 6.1.

The figure depicts an endowment economy with two consumers and two goods. Their endowments are indicated by the dot. Given relative prices, we can draw a budget line as shown by the dashed line. Notice that each consumer will choose a consumption bundle

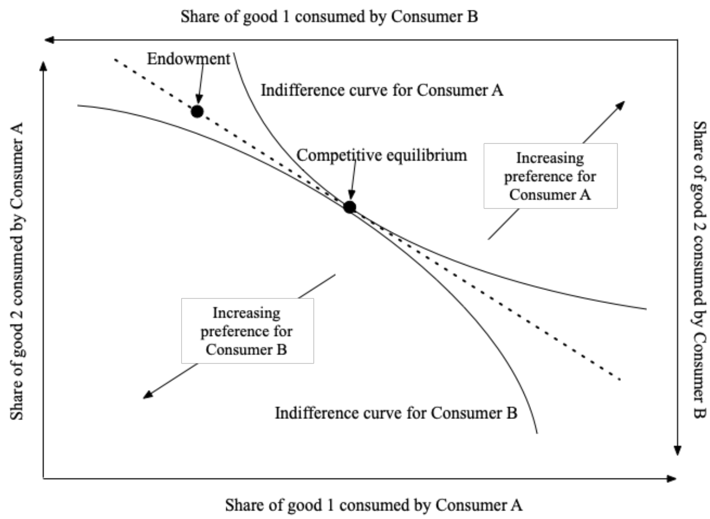

(a) Efficient equilibrium.

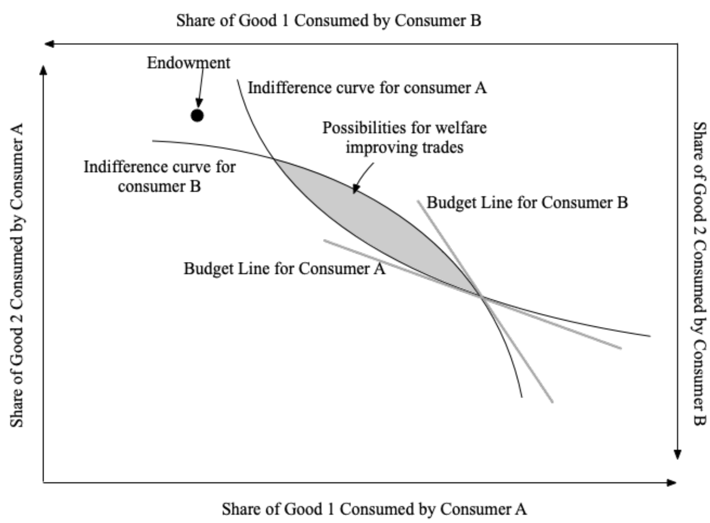

(b) Distorted market.

Figure 6.1: Edgeworth boxes of an endowment economy with two goods and two consumers.

Note: Panel (a) shows an efficient equilibrium in which both consumers face the same prices and have the same budget line. Panel (b) shows a distorted market where the two consumers face different prices.

that makes their indifference curve tangent to the budget line. When we impose market clearing, the points representing their consumption bundles in the Edgeworth box must coincide because what is not consumed by Consumer A must be consumed by Consumer B. Such a situation is marked as a competitive equilibrium in the figure. Notice that the indifference curves are tangent to the budget line and therefore also to each other. Because the indifference curves do not intersect, there is no way to move Consumer A to a higher indifference curve without harming Consumer B. Now consider the lower panel of the figure where we have imagined the two consumers face different prices. We will discuss several reasons this could occur, but a simple one is that Consumer A faces a tax on consuming one of the goods while Consumer B does not. If the consumers face different after-tax prices they will have different budget lines. As before each consumer chooses a consumption bundle where their indifference curve is tangent to their budget line. As the budget lines are different, the indifference curves now intersect and there are alternative allocations that yield higher welfare as shown by the shaded area in the figure.

Similar logic applies to firm profit maximization. Profit maximization leads firms to set the marginal rate of transformation equal to the ratio of prices. As firms face the same prices as consumers, the marginal rate of transformation between any two goods will therefore be equal to the marginal rate of substitution between any two goods. A marginal change in the production allocation therefore cannot transform goods in a way that improves consumer welfare. This point can be visualized as point of tangency between the indifference curves of the consumers (remember they all have the same slope in equilibrium) and the production possibility frontier. A marginal change along the production possibility frontier does not increase consumer welfare.

These intuitive arguments are not as powerful as the more abstract proof above because they apply to marginal changes in allocations while the abstract proof is global. Nevertheless, they can be helpful in developing an intuitive understanding for when the First Welfare Theorem will hold.

How do these ideas apply to macroeconomic models? In the dynastic growth model with optimal saving, the Euler equation sets the marginal rate of substitution between goods at t and $t+1$ equal to the relative price between them, i.e., the real interest rate. As all consumers face the same interest rate, they all have a common marginal rate of substitution between t and $t+1$ goods. Moreover, as firms rent capital from households at a rental rate that is equal to the real interest rate, the marginal rate of transformation between goods at t and $t+1$ is equal to the households' marginal rates of substitution. Similarly, if we consider the model with elastic labor supply, the firm will produce at a point where the marginal product of labor is equal to the wage. The consumer will supply labor such that the marginal rate of substitution between consumption and leisure is equal to the wage. Eliminating the wage from these two optimality conditions gives us that the marginal rate of substitution between goods and leisure is equal to the marginal rate of transformation between goods and leisure.

6.2 Tracing out the Pareto frontier

One useful way to construct a Pareto-optimal allocation is to solve a social planner's problem in which the planner maximizes the weighted sum of the agents' utilities. A solution to such a problem must be Pareto optimal because if it were not, the Pareto dominating allocation would increase the planner's objective function. As we will explain, the solution to this planner's problem is closely related to the concept of competitive equilibrium.

For the sake of simplicity, assume a set of infinitely-lived consumers, indexed by $i \in \mathcal{I}$, live in an exchange economy. Each period, there is a single consumption good and each consumer receives an endowment $\omega_{i,t}$. Each consumer has preferences given by

$$U_i \equiv \sum_{t=0}^{\infty} \beta^t u(c_{i,t})$$

with $u'(c) > 0$ and $u''(c) < 0$. Let p_t be the date-0 price of the date-t good. A competitive equilibrium is a consumption allocation $\{c_{i,t}\}_{\forall i,t}$ and a price system $\{p_t\}_{\forall t}$ such that all consumers are maximizing utility taking prices as given and markets clear. Market clearing requires that $\sum_i c_{i,t} = \sum_i \omega_{i,t}$ for all t.

In the competitive equilibrium with date-0 trading, the consumers maximize their utility subject to the date-0 budget constraint. The Lagrangian of this problem is

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t u(c_{i,t}) + \lambda_i \sum_{t=0}^{\infty} p_t(\omega_{i,t} - c_{i,t})$$

and the first-order condition for $c_{i,t}$ is

$$\beta^t u'(c_{i,t}) = \lambda_i p_t.$$

As the utility function is strictly concave, we can invert $u'(\cdot)$

$$c_{i,t} = (u')^{-1}\left(\beta^{-t}\lambda_i p_t\right), \tag{6.2}$$

which shows that knowledge of $\{\lambda_i\}_{\forall i}$ and $\{p_t\}_{\forall t}$ is enough to determine the entire consumption allocation. The Lagrange multiplier λ_i captures the shadow value of date-0 wealth and will be decreasing in date-0 wealth.

Now consider a social planner that seeks to maximize a weighted sum of the utilities of the consumers. The planner's objective is

$$\sum_{i \in \mathcal{I}} \mu_i U_i,$$

where μ_i is the weight on consumer i's utility. These weights are called Negishi weights or Pareto weights. The constraint on the planner is the resource constraint

$$\sum_{i \in \mathcal{I}} c_{i,t} \leq \sum_{i \in \mathcal{I}} \omega_{i,t},$$

which must hold at each date. The Lagrangian of this problem is

$$\mathcal{L} = \sum_{i \in \mathcal{I}} \mu_i \sum_{t=0}^{\infty} \beta^t u(c_{i,t}) + \sum_{t=0}^{\infty} \psi_t \sum_{i \in \mathcal{I}} (\omega_{i,t} - c_{i,t}),$$

where ψ_t is the Lagrange multiplier on the resource constraint at date t. The first-order condition of this problem with respect to $c_{i,t}$ is

$$\mu_i \beta^t u'(c_{i,t}) = \psi_t.$$

Inverting $u'(\cdot)$ as above yields

$$c_{i,t} = (u')^{-1} \left(\beta^{-t} \frac{1}{\mu_i} \psi_t \right), \tag{6.3}$$

which shows that knowledge of $\{\mu_i\}_{\forall i}$ and $\{\psi_t\}_{\forall t}$ is enough to determine the entire consumption allocation.

There is a clear symmetry between equations (6.2) and (6.3). If $\mu_i = 1/\lambda_i$ and $\psi_t = p_t$, the two equations will give rise to the same consumption allocation. Let's assume that $\mu_i = 1/\lambda_i$. It must then be the case that $\psi_t = p_t$ because in the competitive equilibrium markets must clear and in the solution to the planner's problem the aggregate resource constraint must hold, which means in both cases the total consumption must equal the total endowment. By using equations (6.2) and (6.3) we can see that the only way we achieve the same total consumption at date t is if $\psi_t = p_t$. It follows that if we have the right Negishi weights, the competitive equilibrium is optimal in the eyes of a planner that attaches more weight to those with higher date-0 wealth. This argument is closely related to the First Welfare Theorem: there are some Negishi weights that make the competitive equilibrium optimal for the planner, which implies the competitive equilibrium is Pareto optimal.

Now let's flip the argument in reverse. As we vary the Negishi weights, we will arrive at different consumption allocations as solutions to the planner's problem. Each one of these allocations is Pareto optimal because it solves a planner's problem—again, if it were not Pareto optimal the planner would not choose it. By varying these weights, we can therefore map out many different Pareto-optimal allocations or in other words we can trace out the Pareto frontier. Similarly, by varying date-0 wealth in the market economy we can generate different competitive equilibria with different sets of Lagrange multipliers λ_i. By choosing the appropriate distribution of date-0 wealth we can engineer a competitive equilibrium that mimics the solution to the planner's problem with particular Negishi weights and therefore gives rise to a particular Pareto-optimal consumption allocation. In this case, the consumers still trade with one another, but we redistribute resources between them to change the distribution of consumption.

This procedure of constructing a competitive equilibrium to deliver a particular Pareto-optimal consumption allocation is closely related to the Second Welfare Theorem. The Second Welfare Theorem begins with a Pareto-optimal allocation and then gives conditions under which there exists a competitive equilibrium delivering this allocation as an outcome. The conditions involve ensuring that consumers' and firms' maximization problems have well-defined solutions, which in turn in general necessitates assumptions of convexity (e.g., the consumer's utility function is globally concave). These additional assumptions are typically met in our macroeconomic applications so they are not a problem per se, but the statement of a general theorem is more cumbersome so we will not present it here.

Market equilibria, as observed in actual economies, have radically different consumption levels across agents, so viewed from the perspective of frictionless competitive equilibria and an additive social welfare function of the sort just described, they correspond to points on the Pareto frontier with high Negishi weights on the high-consumption agents. The social welfare function with Negishi weights is just an analytical tool that we can use to construct different Pareto-optimal allocations. To be clear, the argument we have made here is not saying that this distribution of consumption is desirable or just but simply that there are some Pareto weights that could make the planner choose that distribution of consumption.

6.3 Inefficient market outcomes

We will now very briefly, by means of simple examples, cover a number of commonly considered departures from the abstract frictionless economy considered above. For each case, we will comment on efficiency properties by describing how the proof of the First Welfare Theorem 6.1 may or may not go through as well as explaining how the more intuitive marginal efficiency conditions may or may not hold. We begin with a case where the culprit is not the market but distortionary taxation and then look at externalities, missing markets, and monopoly power.

6.3.1 Taxes

First, consider lump-sum taxes (in a frictionless economy). Lump-sum taxes do not have to be equal across agents; the key is that the amount given to agent i does not depend on the behavior of agent i. From the perspective of the marginal conditions characterizing optima, since the marginal conditions of competitive equilibria do not involve lump-sum taxes or transfers, equilibria remain optimal. From the perspective of our abstract proof, all lump-sum taxes do is redistribute wealth (and thus "utils") across agents, i.e., move us along the Pareto frontier.[3]

Second, consider the two most commonly studied taxes in macroeconomic applications: taxes on labor earnings and taxes on capital income. Beginning with taxes on labor earnings, consider a tax that is proportional to earnings. Hence, instead of w_t an agent receives $w_t(1-\tau_\ell)$ for each unit of hours worked. In an economy where consumers do not value leisure, this tax does not affect any first-order conditions; hence, it acts as a lump-sum tax and does not disturb the efficiency properties of equilibrium. However, if leisure is valued, the marginal rate of substitution between consumption and leisure will differ from the marginal rate of transformation between consumption and leisure by a factor $1-\tau_\ell$. Similarly, if capital income is taxed, the marginal rate of substitution between goods at t and $t+1$ will differ from the marginal rate of transformation due to the tax.

What goes wrong? First, intuitively the firms and the households face different (after-tax) prices and therefore they may not exploit all of the possible trades that they should. Now looking at our abstract proof of equilibrium efficiency, what goes wrong in trying to use it? The key is that the tax appears in the equations. Suppose that the p in the proof is the

[3]To see this formally, one can consider the government to be one of the consumers in the economy. When we sum across the budget constraints of the consumers in the proof of the First Welfare Theorem, the lump-sum taxes will cancel out as the government's revenue equals the other consumer's tax payments.

pre-tax prices. What would not hold is the point where we say the alternative consumption allocation must cost more: $p\sum_{i\in\mathcal{I}}\tilde{x}_i > p\omega + p\sum_{j\in\mathcal{J}} y_j^\star$. This inequality may not hold because the consumer is maximizing utility with respect to the after-tax prices.[4] Hence the proof does not in general go through. However, the proof is valid when leisure is not valued, because then the leisure chosen, which is an element in x, is zero, so the fact that p is different for this good does not matter. This confirms the intuition that a proportional tax on an inelastic labor supply is non-distortionary.

Similarly, a proportional tax on capital income, e.g., with r_t replaced by $r_t(1-\tau_k)$, will appear in the Euler equation and therefore make the marginal rate of substitution between goods at $t-1$ and t differ from the corresponding marginal rate of transformation. In the proof of the First Welfare Theorem, the relative price of the time t good is higher for consumers than for producers and hence the proof cannot be completed.

6.3.2 Externalities

An externality arises when one agent's activity has payoff relevance to other agents that is not reflected in the market price associated with this activity. Let us use an example with a negative TFP externality: production by one firm damages the production carried out in other firms. Consider a static economy with a representative consumer who values leisure. In particular, the consumer maximizes $u(c,\ell) = \log c - B\ell^{1+1/\theta}/(1+1/\theta)$ subject to

$$c = rk + w\ell$$

by choice of (c,ℓ). Output of a typical firm j equals

$$A(\bar{y})k_j^\alpha \ell_j^{1-\alpha},$$

with $\bar{y} = (\sum_j y_j)/\mu$, where μ is the number of firms. Here, A is a decreasing function, expressing a negative production externality: the higher is total production in the economy, the lower is the productivity of each firm. Let us also assume that μ is large enough that each firm j ignores its own impact on $\bar{y}$; for simplicity, think of the set of firms as a continuum on $[0,1]$, with $\mu = 1$, so that there is a notion of a representative firm. Hence, we can drop the subscript j and the representative firm solves

$$\max_{k,\ell} \; A(\bar{y})k^\alpha \ell^{1-\alpha} - rk - w\ell,$$

taking $\bar{y}$ as given. In this static economy, output is determined by equilibrium in the labor market. From the consumer, we obtain

$$\frac{w}{rk + w\ell} = B\ell^{\frac{1}{\theta}}$$

and the firm's first-order conditions imply that

$$\frac{r}{w} = \frac{\alpha}{1-\alpha}\frac{\ell}{k}.$$

[4] We cannot just interpret p as the after-tax price because then when we say firms are profit maximizing, they are not maximizing profits with respect to p.

Combining the two equations, we obtain (with a small amount of algebra) that the equilibrium outcome for hours worked is given by the unique solution to

$$1 - \alpha = B\ell^{1+\frac{1}{\theta}}.$$

A striking feature of this solution is that it does not depend on the strength of the externality as captured by the function A: A does not appear in the equation.[5]

The efficient allocation, on the other hand, is given by the solution to

$$\max_{\ell, y} \; \log y - B\frac{\ell^{1+\frac{1}{\theta}}}{1+\frac{1}{\theta}} \quad \text{subject to} \quad y = A(y)k^{\alpha}\ell^{1-\alpha}.$$

Here, the first-order conditions (with λ denoting the multiplier on the constraint) are

$$B\ell^{\frac{1}{\theta}} = \lambda(1-\alpha)A(y)k^{\alpha}\ell^{-\alpha} \quad \text{and} \quad \frac{1}{y} = \lambda\left(1 - A'(y)k^{\alpha}\ell^{1-\alpha}\right),$$

which delivers

$$\frac{1-\alpha}{1 - A'(y)k^{\alpha}\ell^{1-\alpha}} = B\ell^{1+\frac{1}{\theta}} \quad \text{and} \quad y = A(y)k^{\alpha}\ell^{1-\alpha}$$

as two equations in the two unknowns ℓ and y. The first of these equations can be compared to the equilibrium outcome: the equilibrium is not optimal, unless $A'(y) = 0$, i.e., unless there is no externality. Intuitively, in their input choices, firms do not take into account how their production hurts others, and they should.

What goes wrong? To see how our abstract proof of the First Welfare Theorem would go wrong in this case, first note that an equilibrium with externalities would be defined by letting the Y_js—each firm's production possibility set—be endogenous and interdependent: $Y_j = Y_j((y_{j'})_{j' \neq j})$, where $(y_{j'})_{j' \neq j}$ is the vector of choices of other firms.[6] A key step in our abstract proof was that for each j, $p\tilde{y}_j \leq py_j^{\star}$, from profit maximization. This no longer follows, since the alternative allocation implies a different choice set for firm j, as given by $Y_j = Y_j((y_{j'})_{j' \neq j})$. In concrete terms, if other firms scale down their production relative to that in equilibrium, your choice set improves—your TFP increases—and your original choice was not optimal.[7]

The negative externality considered here is closely related to how the externalities due to climate change are usually modeled. There, the externality is usually assumed to affect TFP, as in our example here. Rather than occurring through overall production, however, the externality occurs through carbon emissions, which result from production of a specific good—energy derived from fossil fuels. Climate change is covered in Chapter 25.

[5]This particular feature, which follows because income and substitution effects cancel in the labor-supply specification, makes the example stark but is not necessary for the arguments here.

[6]For simplicity here we abstract from your own production lowering your own TFP.

[7]In the present example, equilibrium profits are zero, and if all other firms decrease their output below the equilibrium level, your profits can be made positive (and unboundedly large) by simply scaling up your input choices.

6.3.3 Missing markets: an example with constraints on borrowing

A friction that is commonly studied in macroeconomic applications involves "missing markets": restrictions on trade in one way or another. For example, there is no market that allows workers to buy insurance against unfavorable changes in their salaries because of the moral hazard that workers will have incentives to collect the insurance payment rather than work hard. Another example is that households and firms may not be able to borrow much as they would like. Borrowing constraints are thought to be important constraints both for firms (for funding their day-to-day operations as well as long-term investments) and consumers (for purchasing homes and durable goods more generally). These borrowing constraints ultimately stem from features of the economic environment, such as the difficulty of enforcing repayment, but are sometimes modeled as simple constraints. We focus on the case of borrowing constraints here, but the logic applies more generally to settings where certain goods or agreements for borrowing or insurance are not traded.

We begin with a two-agent dynastic endowment economy: a two-type special case of that presented in Section 5.3.1. The total endowment each period is constant at ω, but let us now assume that agent 1 is endowed with $2\omega/3$ in odd periods and $\omega/3$ in even periods; agent 2 has $\omega/3$ and $2\omega/3$ in odd and even periods, respectively (the odd-numbered agent is rich in odd periods). Let us also assume that both agents have logarithmic utility. The economy starts in period 0. In a competitive equilibrium with unrestricted markets we have (i) full consumption smoothing, so that $c_{1,t} = c_1$ for all t and $c_{2,t} = c_2$ for all t; prices for goods at different dates satisfy $p_t = \beta^t$, i.e., the gross real interest rate is $1/\beta$; and (iii)

$$c_1 = \left(\frac{1}{3} + \frac{2\beta}{3}\right)\frac{\omega}{1+\beta}$$

and

$$c_2 = \left(\frac{2}{3} + \frac{\beta}{3}\right)\frac{\omega}{1+\beta}.$$

Consumer 2 is richer, and hence consumes more, due to an endowment stream that is higher in present value because consumer 2 receives the larger endowment one period before consumer 1 does.

Note that the unrestricted competitive equilibrium has active borrowing and lending: agent 1 borrows in even periods and repays in odd periods. Suppose, then, that borrowing is simply not allowed. In terms of a sequence of budget constraints $c_t + q_t a_{t+1} = \omega_t + a_t$, where a is asset holdings and q is price of a one-period real bond delivering 1 unit of consumption next period, no borrowing means that the consumer is facing an additional constraint: $a_{t+1} \geq 0$ for all t. The consumer's maximization problem then reads

$$\max_{\{a_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log\left(\omega_t + a_t - q_t a_{t+1}\right) \quad \text{subject to} \quad a_{t+1} \geq 0 \quad \forall t.$$

Differentiation with respect to a_{t+1} delivers

$$\frac{q_t}{c_t} = \beta\frac{1}{c_{t+1}} + \mu_t,$$

where μ_t is the Kuhn-Tucker multiplier on the borrowing constraint. This multiplier is non-negative: positive when the constraint binds and zero otherwise. Hence, we conclude that under a borrowing constraint, the Euler equation generally is an inequality constraint: $q_t u'(c_t) \geq \beta u'(c_{t+1})$. Intuitively, the marginal value of consumption today can be strictly higher than the consumer could obtain by consumption tomorrow, using market prices, but the consumer cannot increase consumption today further due to the borrowing constraint.

In our simple example, no borrowing means autarky: consumer 1 always consumes $2\omega/3$ in odd periods and $\omega/3$ in even periods, with the remainder of the total endowment consumed by agent 2.[8] The allocation is clearly not optimal, since consumption is not smoothed. What is the prediction for interest rates? In even periods, agent 2 is not constrained so we know that

$$\frac{q_t}{2\omega/3} = \beta \frac{1}{\omega/3}$$

when t is even. Hence, $q_t = 2\beta$: the real interest rate $1/q_t - 1$ is lower than under unrestricted borrowing, since the lender needs to be discouraged from lending.[9] The same logic applies in odd periods: $q_t = 2\beta$. We even see that the real interest rate will be negative provided that $\beta > 0.5$.

What goes wrong? Clearly, marginal rates of substitution are not equalized across agents here. From the perspective of our abstract proof, what goes wrong? As in the case with externalities, the sets from which agents are choosing will change: they will contain more restrictions and they will be endogenous. For example, without borrowing constraints, the main restriction on the consumption possibility sets is non-negativity. With borrowing constraints, it will include an additional constraint for each good: $c_0 \leq a_0 + \omega_0$, $q_0 c_1 \leq \omega_0 + a_0 - c_0 - q_0 \omega_1$, etc. The equilibrium allocation can be dominated by an alternative allocation—say, one with full smoothing—simply because this alternative allocation is not within the agent's consumption possibility set.

Moreover, as we see here, the consumption possibility sets will now also depend on prices, departing conceptually from the abstract setup where the consumption possibility sets are exogenous. As in the externality example, there is an interdependence between the possibility sets the agents face. And just like the externality case, the agents will not take into account how their actions affect the possibility sets that other face. These are *pecuniary externalities*—one agent's choices affect the prices that appear in other agents' constraints. As a result of these pecuniary externalities, if a planner were to select the consumption allocation subject to the budget constraints and borrowing constraints, the planner may choose a different allocation than the competitive equilibrium because the planner would take account of the pecuniary externalities.

[8] As agent 1 is unable to borrow in period 0, agent 2 has nobody to lend to and cannot save. In period 1, the two agents face the same situation as period 0 with their roles reversed.

[9] Values of q_t above 2β will also constitute equilibria, since all agents are then strictly constrained. However, we view the case with an interior solution as the interesting one, since it is the limiting equilibrium for economies with borrowing constraints $a_{t+1} \geq \underline{a}$ where $\underline{a} < 0$, with $\underline{a} \to 0$.

6.3.4 Lack of commitment

In the dynamic equilibrium models we looked at above, we always assumed that all agents engaged in intertemporal trade could commit to delivering on their promises. When such commitment is not available, the allocations are of course affected. One could, for example, imagine that it is possible to default on debt and that such defaults are not punished in any way. This would, given that lenders are rational, lead to de-facto borrowing constraints, as in Section 6.3.3 above. One could also imagine that default can occur but that it is punished, in which case intertemporal trade generally will occur. Note, however, that the punishment mechanism per se needs to be committed to and that it is not clear that punishment will be rational ex post. When punishments are assumed in economic models, there is therefore a presumption that some agent (such as the government) has an ability to commit to it.

Macroeconomic models where consumers cannot fully commit, including where they do default in equilibrium, exist and are interesting cases of "endogenously incomplete markets," but they are not discussed in this text. A more common example involves government policy, such as the case where governments cannot fully commit to future tax policy; this case is discussed in Chapter 15 below. In the context of international economics, another important example is the case where default on sovereign debt (involving lending between countries) can occur; this is a central example in Chapter 24 on emerging markets.

6.3.5 Market power

A well-known case, and also one of great relevance for macroeconomics, is where some agents in the economy have market power. That is, they can, through their own behavior, affect the prices they are transacting at. This may occur on the individual level—you may be able to bargain for a higher salary—or at the firm level—the firm can set the price of its product, or try to bargain with their input suppliers over those prices. We will discuss examples of both of these occurrences in the book. When we study labor-market frictions, it will be natural to discuss bargaining. When we study growth and business cycles, we will study markups, i.e., how firms with market power set prices above marginal cost.

In macroeconomic models with market power, we will almost always assume that this power is quite limited: it is limited to an individual transaction and not affecting aggregates. Thus, for monopoly settings we most often study monopolistic competition, where each firm holds a monopoly over a specific good but faces competition with a continuum of imperfect substitutes, with each one playing a negligible role in the aggregate. Similarly, a worker will be assumed not to have any influence on aggregates in the wage negotiations. Clearly, there are examples where these assumptions are not appropriate (Apple can likely affect the entire market for smartphones, and perhaps even world GDP, in their pricing; particularly gifted vaccine researchers could be seen to have similar powers). But they are likely rare. So let us now briefly and compactly describe the basic, static model of monopolistic competition building on Dixit and Stiglitz (1977).

A model with monopolistic competition

Let us assume a representative consumer with preferences defined over a continuum of imperfectly substitutable goods and labor effort, L:

$$U\left((c(i))_{i=0}^{1}, L\right) = u\left(\left(\int_0^1 c\left(i\right)^{1-\frac{1}{\varepsilon}}\, di\right)^{\frac{\varepsilon}{\varepsilon-1}}\right) - v(L) \equiv u(C) - v(L).$$

This function implies a constant elasticity of substitution $\varepsilon \geq 0$ across different consumption goods (a CES function). We will specialize u to be logarithmic below in an example.

The consumer's budget is

$$\int_0^1 p\left(i\right) c\left(i\right) di = y,$$

where y is income. We will now derive a price index by considering how the consumer should allocate goods in the cheapest way in order to reach a given level of C . So consider

$$\min_{(c(i))_{i=0}^{1}} \int_0^1 p\left(i\right) c\left(i\right) di \quad \text{subject to} \quad \left(\int_0^1 c\left(i\right)^{1-\frac{1}{\varepsilon}}\, di\right)^{\frac{\varepsilon}{\varepsilon-1}} \geq C.$$

We obtain, with λ denoting the multiplier for the constraint,

$$\begin{aligned} p\left(i\right) &= \lambda \frac{\varepsilon}{\varepsilon-1} \left(\int_0^1 c\left(i\right)^{1-\frac{1}{\varepsilon}}\, di\right)^{\frac{\varepsilon}{\varepsilon-1}-1} \left(1-\frac{1}{\varepsilon}\right) c\left(i\right)^{-\frac{1}{\varepsilon}} \\ &= \lambda \left(\int_0^1 c\left(i\right)^{1-\frac{1}{\varepsilon}}\, di\right)^{\frac{\varepsilon}{\varepsilon-1}-1} c\left(i\right)^{-\frac{1}{\varepsilon}} = \lambda \left(\frac{c\left(i\right)}{C}\right)^{-\frac{1}{\varepsilon}}. \end{aligned} \tag{6.4}$$

Multiply by $c(i)$ on both sides, sum across goods, and simplify to obtain

$$\int_0^1 p\left(i\right) c\left(i\right) di = \lambda C.$$

That is, λ can be interpreted as a "unit price of C," the whole basket. What is λ in terms of primitives? Often the multiplier is merely used to set up and solve a maximization problem but sometimes it carries important content, such as here, and then it is relevant to compute its value. To derive a formula for the unit price, use the expression for $p(i)$ again: raise it to $1-\varepsilon$ and sum across i. This delivers

$$P \equiv \lambda = \left(\int_0^1 p\left(i\right)^{1-\varepsilon}\, di\right)^{\frac{1}{1-\varepsilon}}.$$

So P, the unit price, is itself a CES function that is increasing in all prices and homogeneous of degree one.

A key other implication of equation (6.4) is that it expresses an *inverse demand* function: it expresses a relation between the price of good i and the demand for it, given an overall level of spending PC. The demand function itself becomes, after solving for $c(i)$ and using $\lambda = P$,

$$c(i) = C\left(\frac{p(i)}{P}\right)^{-\epsilon}.$$

Here we see that the price elasticity of demand is constant and equal to ϵ. (One can replace C by y/P, where y is the consumer's income.)

The demand function is used as a central object in the definition of a monopolistically competitive equilibrium. We will now state it. We assume that one firm produces each kind of consumption good and that all production functions are identical and linear in labor: $c(i) = A\ell(i)$ for all i. We assume that labor supply is exogenous and equal to 1.

A monopolistically competitive equilibrium For the economy described above, a monopolistically competitive equilibrium is a consumption allocation $\{c(i)^\star\}_{\forall i}$, a labor allocation $\{\ell(i)^\star\}_{\forall i}$, a price vector $\{p(i)^\star\}_{\forall i}$, a profit vector $\{\pi(i)^\star\}_{\forall i}$, and a wage $w^\star$ such that

1. $(\{c(i)^\star\}_{\forall i}, L^\star)$, where $L^\star \equiv \int_0^1 \ell^\star(i)di$, solves

$$\max_{(c(i))_{i=0}^1, L} u\left(\left(\int_0^1 c\,(i)^{1-\frac{1}{\varepsilon}}\,di\right)^{\frac{\varepsilon}{\varepsilon-1}}\right) - v(L)$$

   subject to $\int_0^1 p^\star(i)c(i)di = w^\star L + \int_0^1 \pi^\star(i)di$.

2. for each i, $(p^\star(i), c(i)^\star, \ell(i)^\star)$ solves the maximization problem

$$\max_{p,c,\ell}\; pc - w^\star\ell \quad \text{subject to} \quad c = A\ell \quad \text{and} \quad p = P^\star\left(\frac{c}{C^\star}\right)^{-\frac{1}{\epsilon}},$$

   where

$$C^\star = \left(\int_0^1 c^\star(i)^{1-\frac{1}{\varepsilon}}di\right)^{\frac{\varepsilon}{\varepsilon-1}} \text{ and } P^\star = \left(\int_0^1 p^\star\,(i)^{1-\varepsilon}\,di\right)^{\frac{1}{1-\varepsilon}}$$

   and $\pi^\star(i)$ defines the maximum obtained.

Notice (i) that no market-clearing condition is needed as it is satisfied immediately as part of the firms' problems; (ii) that firms make profits, which accrue to the consumer; (iii) that firm i solves a problem that does not depend on i, since all goods are symmetric here (still, we label the solutions with i); and (iv) that a key input into the firm's problem is the inverse demand function, $p^\star(c)$. We stated this last condition with knowledge of form for the inverse demand function for each good; this function, of course, has to be consistent with condition 1 of the definition and this was ensured in our derivations leading up to the definition. A monopolistically competitive equilibrium can alternatively be defined to include this demand function explicitly as an equilibrium object, with the added condition that it is consistent with consumer maximization.

To see what implications follow from this setup, let us solve the firm's problem. Substitute the constraints into the objective, so that it reads

$$\max_c\; PC^{\frac{1}{\varepsilon}}c^{1-\frac{1}{\varepsilon}} - \frac{w}{A}c,$$

where we have dropped the i due to symmetry and *s for notational convenience. Clearly, this is a well-defined problem if $\varepsilon > 1$; if $\varepsilon \leq 1$, goods are not sufficiently substitutable, and the monopolist's problem does not have a solution.[10]

$$c = C \cdot \left(\frac{w}{A} \frac{1}{(1 - 1/\varepsilon)\, P} \right)^{-\varepsilon} \qquad \text{and} \qquad p = \frac{\varepsilon}{\varepsilon - 1} \frac{w}{A};$$

thus, $\mu \equiv \varepsilon/(\varepsilon - 1) > 1$ expresses that the firm charges a *markup* over marginal cost, w/A, that is constant in percent. Profits π are given by $(\mu - 1)(w/A)c$.

The equilibrium is symmetric, so $c = C$. We can normalize one price, or a combination of prices, so we set $P = 1$. We then see that the equilibrium wage has to equal $w = A(1-1/\varepsilon)$: the wage is below marginal productivity (as $\varepsilon > 1$). Equilibrium work effort and consumption are then solved from $u'(C)w = v'(L)$. With a logarithmic u we obtain

$$A(1 - 1/\varepsilon) = Cv'(C/A),$$

which has a unique solution for C assuming $v(L)$ is convex and, hence, the right-hand side is strictly increasing in C.

It is easy to see that the outcome is inefficient. The planner will choose symmetry across goods and hence simply maximizes $u(C) - v(C/A)$. The outcome is the first-order condition

$$A = Cv'(C/A),$$

whose left-hand side is larger than in the monopolistic case and, hence, output and hours worked are too low in equilibrium.

What goes wrong with market power? Firms choose inputs taking into account how their demand is affected. As a result, firms tend to under-produce relative to the optimum, since that gives them a higher price. Thus, in equilibrium, firms transform labor into consumption goods at a marginal rate that is higher than the rate at which consumers value leisure relative to the consumption good. Our abstract proof of the First Welfare Theorem, moreover, cannot be used directly since an equilibrium with monopoly power is defined quite differently, as we have seen: some agents do not take p as given.

6.3.6 Quantifying welfare losses

In applied macroeconomic analysis, the aim is most often quantitative. The researcher wants to go beyond qualitative statements like "the equilibrium is not optimal" to a quantitative one indicating *how much* worse, or better, one allocation is than another. There are numerous ways to do this and we will focus on the most common one here. We will use distortionary taxes in the context of a representative-agent economy as an example.

So suppose the government has an amount of expenditures—e.g., military purchases—that it needs to finance with taxes. One option, at least in theory, is to use lump-sum taxes.

[10] When $\varepsilon < 1$, infinite profits can be obtained by raising prices toward infinity; the reader is invited to show this by considering the implied maximization problem. The case $\varepsilon = 1$ is special: since revenues are then independent of the price, higher prices, and consequently lower quantities sold and therefore production costs, are always better and no profit maximum exists. Its supremum equals the revenue.

Another one is to tax labor earnings at a proportional rate. How much worse is it to use distortionary taxes? Let us focus on a static model because it is simple; the principles are the same in a dynamic context.

The procedure is simple. First compute an equilibrium for each of the two tax systems. Denote the resulting allocations of consumption and hours worked (c_l, ℓ_l) and (c_d, ℓ_d) (l for lump-sum and d for distortionary). Clearly, $u(c_l, \ell_l) > u(c_d, \ell_d)$, where the utility is defined over consumption and hours (hours affect utility negatively). Now it is always possible, with standard utility functions, to find a $\Delta > 0$ such that $u(c_l(1-\Delta), \ell_l) = u(c_d, \ell_d)$. The Δ expresses how much, in percent, consumption needs to be decreased in the better allocation to generate the same utility as in the worse allocation, while maintaining the same hours choice.

The key in the example is that Δ has a real interpretation; while one could compute the difference in consumer "utils" between two allocations, such a measure would not have an interpretation as "utils" have no meaning per se. There are, of course, alternative ways to define a Δ. One could, for example, imagine both reducing consumption and raising hours worked at the same time. One could also define a Δ as the percentage increase in consumption in the worse allocation that would make it as good as the better allocation. This would deliver a different Δ. One can, finally, define Δ as an amount of wealth that the consumer would need to be given as a lump-sum transfer in order to be indifferent between two allocations.

In dynamic models, the Δ is defined as a percentage change in consumption in all time periods. Consumption need not be constant over time in either of the allocations considered, but a unique Δ can still be defined: it is the percentage decrease in consumption applied in every period that would make the consumer indifferent with the worse allocation.

6.4 Overlapping generations

The efficiency properties of overlapping-generations (OG) models require separate discussion and have been thoroughly studied in the literature. Here, we highlight only the main results and insights. The key takeaway is that even in the absence of frictions—no missing markets, distortionary taxes, monopoly power, or externalities—equilibria can be, though need not be, Pareto inefficient. Moreover, a simple litmus test exists to determine whether an equilibrium is efficient. We will provide this test, but not prove it.

Before proceeding, let us note that there are models that share properties with overlapping generations models, such as those where people die randomly and new people appear. Here, equilibria are not necessarily efficient either.

6.4.1 The endowment case

Consider a simple example for illustration: there is a representative consumer in each cohort who lives for two periods. Let us also use a concrete, very simple utility function: for every generation $t \geq 0$ preferences are represented by

$$u_t(c_y, c_o) = \log c_y + \log c_o$$

where we evaluate at two arbitrary consumption levels (c_y, c_o). The preferences of generation $t = -1$, who are old as time begins, are similarly represented by $u_{-1}(c) = \log c$.

Suppose stationary endowment sequences are given by

$$\omega_{y,t} = \omega_y$$

and

$$\omega_{o,t} = \omega_o$$

for all t, where t represents the time period and $\omega_{o,t}$ is the endowment of the old (who were born at $t-1$) at time t.

In the competitive equilibrium we consider, trading is sequential and there are no borrowing constraints. With q_t being the price of a bond at t, the agent born at $t \geq 0$ solves

$$\max_{c_y, c_o} \log c_y + \log c_o$$

subject to

$$c_y + q_t c_o = \omega_y + q_t \omega_o.$$

It is straightforward to solve this maximization problem. It delivers

$$c_{y,t} = \frac{1}{2}(\omega_y + q_t \omega_o) \tag{6.5}$$

and

$$c_{o,t+1} = \frac{1}{2}\left(\frac{\omega_y}{q_t} + \omega_o\right), \tag{6.6}$$

where we have now given the consumption choices sub-indexes for the time period in which they occur. Note that the consumer's saving when young is $\omega_y - c_{y,t} = (\omega_y - q_t\omega_o)/2$.

The old agent at time zero maximizes utility subject to the budget $c_{o,0} = \omega_o$ and hence the choice is trivially given by the budget.

Market clearing in this overlapping-generations economy for period $t = 0$ reads

$$c_{o,0} + c_{y,0} = \omega_y + \omega_o.$$

Since the old's choice is given, we conclude that $c_{y,0} = \omega_y$: the young also consumes the endowment. The bond price that clears the market is $q_0 = \omega_y/\omega_o$. It follows that $c_{o,1} = \omega_o$: the old at 1 will consume the endowment in the second period of life as well.

The argument is then repeated and we obtain, period by period, that

$$c_{y,t} = \omega_y,$$

$$c_{o,t} = \omega_o,$$

and

$$q_t = \frac{\omega_y}{\omega_o}.$$

This constant sequence supports the equilibrium where agents do not trade: the prices induce people to consume their initial endowments.

Let us now plug in specific numbers. Let $\omega_y = 3$ and $\omega_o = 1$. It follows that $q_t = 3$. Thus, the gross real interest rate is 1/3; the net interest rate is - 67%.

Is this allocation Pareto efficient? Consider the following alternative feasible allocation: for all t,

$$\tilde{c}_{y,t} = 2$$

and

$$\tilde{c}_{o,t} = 2.$$

That is, the alternative allocation $\tilde{c}$ is obtained from a chain of intergenerational goods transfers that consists of the young in every period giving a unit of their endowment to the old in that period. Notice that for all generations $t \geq 0$, this is just a modification of the timing in their consumption, since total goods consumed throughout their lifetime remain at 4. For the initial old, this is an increase from 1 to 2 units of consumption when old. It is clear, then, that the initial old strictly prefer the new allocation. We need to check what the remaining generations think about the change. It is clear that since utility is concave (the log function is concave), this even split of the same total amount will yield a higher utility value: $\log 2 + \log 2 = 2 \cdot \log 2 = \log 4 > \log 3 + \log 1 = \log 3$.

We conclude that the competitive equilibrium, which we solved for uniquely, is not Pareto optimal: it is dominated by an allocation where each cohort gives a transfer when young and receives one when old. But why is the equilibrium not Pareto optimal? There is no friction: no agent is prevented from trading, there are no externalities or elements of market power. We will return to this question shortly, but first consider the reverse case: $\omega_y = 1$ and $\omega_o = 3$. Now, q_t becomes 1/3 each period; the net real interest rate is +67%. Again, consider the alternative $(2, 2)$ allocation: is it a Pareto improvement? For all generations born at $t \geq 0$, the answer is yes, as in the previous example. However, the old at 0 will be made worse off. Hence, the proposed alternative is not a Pareto improvement. Is there some other, smarter alternative allocation that does the job? The answer is, in fact, no.

In the overlapping generations model, equilibria are sometimes Pareto optimal and sometimes not. We will elaborate on the intuition later, but let us instead go back and revisit our abstract proof of the First Welfare Theorem in Section 6.1. The notation there can, as in our application, include infinite sequences, and the proof goes through in all parts, except possibly in one: when summing the budget over all agents, $\mathcal{I}$ is now infinite. Prices should now be viewed as given by the sequence $p = \{p_0, p_1, p_2, \dots\}$ where $p_t = \prod_{\tau=0}^{t-1} q_\tau$. In our given equilibrium where $q_t = 3$ for all t, then, the present value of the equilibrium allocation is $\omega_o + (\omega_y + \omega_o) + 3(\omega_y + \omega_o) + 3^2(\omega_y + \omega_o) + \dots$ and this sum is infinite. Hence, this proof strategy cannot be used. However, in the equilibrium where $q_t = 1/3$, the present value is finite, and the proof does go through. Thus, we in fact have a proof that there is no allocation that can Pareto improve on the autarkic allocation $(c_{y,t}, c_{o,t}) = (1, 3)$ for all t!

When can equilibria where the present-value budget sum across all agents is infinite be Pareto improved upon? We cannot rely on the standard proof of the First Welfare Theorem, but it turns out that there is a general theorem—one provided in Balasko and Shell (1980)—that gives us the answer. The theorem relies on some assumptions, but these assumptions are rather weak: aside from regularity conditions and a bounded sequence for total endowments, they mainly restrict the curvature of consumers' indifference curves away from the two extreme cases (linear and kinked). The Balasko-Shell result is: A competitive equilibrium in an endowment economy populated by overlapping generations of agents is Pareto optimal if and only if

$$\sum_{t=0}^{\infty} \frac{1}{p_t} = \infty.$$

The proof is quite involved and we refer the reader to the source. The two special cases we have looked at are consistent with the theorem: $q_t = 3$ for all t implies that $\sum_{t=0}^{\infty}(1/p_t) = \sum_{t=0}^{\infty} 3^{-t} = 3/2$ is finite, and hence the equilibrium is not optimal; $q_t = 1/3$ for all t implies that $\sum_{t=0}^{\infty}(1/p_t) = \sum_{t=0}^{\infty} 3^t$ is infinite, and hence the equilibrium is optimal. Now, however, we can also evaluate the middle case where $(\omega_y, \omega_o) = (2, 2)$. Here, since $q_t = 1$ implies an infinite sum, our abstract proof cannot be used—as in the $(\omega_y, \omega_o) = (3, 1)$ case—but now the theorem tells us that equilibrium is optimal (a infinite sum of 1s is infinite).

The Balasko-Shell theorem can be applied also to non-constant sequences; indeed, it applies also for non-constant endowment sequences. An important observation, however, is that whether $\sum_{t=0}^{\infty}(1/p_t)$ is finite or not does not depend on anything but how p_t behaves as t literally goes to infinity. Thus, whenever the gross real interest rate p_t/p_{t+1} converges, we know that the equilibrium is optimal if and only if the limit net real interest rate is equal to or above 0.[11] Relatedly, if the economy has a finite time horizon, no inefficiency can occur, no matter how the real interest rate evolves: the present value of the sum of all budgets is finite in this case, and the standard proof can be used. It is the combination of (i) infinite time and (ii) a corresponding infinite set of cohorts of consumers—each of which has a finite budget—that can make markets fail.

Third and finally, whenever the equilibrium is sub-optimal, a straightforward government policy of transferring resources from the young to the old will allow all generations to be made better off. This is an argument for the introduction of social security as a pure government-mediated transfer scheme: a "pay-as-you-go" system. The government works like a bank here: you give it money when young and get it back when old. What allows the government to achieve a better allocation than markets can deliver? On the one hand, the inefficiency of the overlapping generations model is a pure market failure, and in that sense a government can improve on it. But it does require an ability of the government to implement a sequence of transfers that stretches into eternity. If, in our simple $(\omega_y, \omega_o) = (3, 1)$ case transfers stopped at some point in time, the young in the very last period of transfers will be worse off by having given something away, with nothing in return.

Let us now provide some intuition for the Balasko-Shell result, and let us continue with our example and focus on the "toughest" case: $\omega_y = \omega_o = 2$, where a First Welfare Theorem cannot be proved the standard way and yet applies. First, we restrict attention to *stationary* allocations, i.e., allocations such that $c_{y,t} = c_y$ for all t and $c_{o,t} = c_o$ for all t. So is there a stationary allocation that Pareto dominates (2, 2)? Figure 6.2 shows the resource constraint of the economy, plotted together with the utility level curve corresponding to the allocation (2, 2).

The shaded area is the feasible set; its frontier given by the line $c_y + c_o = 4$. It is clear from the picture with a tangency at (2, 2) (recall that the utility function is $\log c_y + \log c_o$) that it is not possible to find an alternative allocation that Pareto dominates this one. Now,

[11] For an economy where endowments grow at some net rate g in the limit, a similar theorem applies: the equilibrium is optimal if and only if the limit interest rate is equal to or above g.

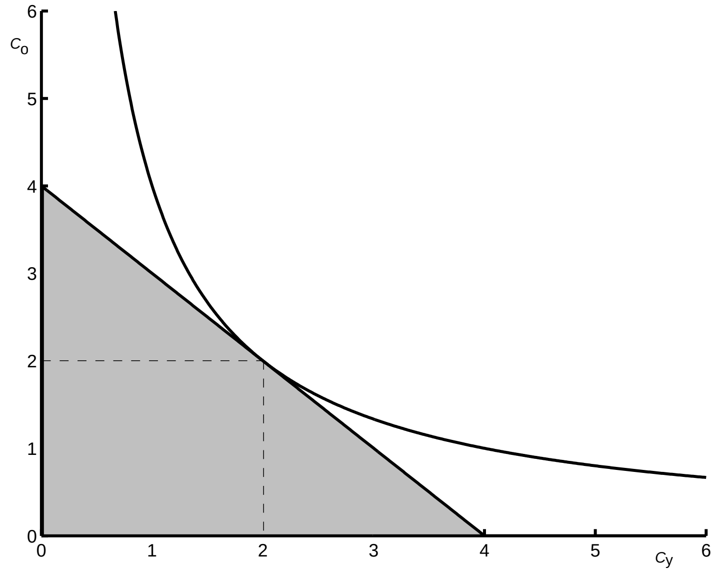

Figure 6.2: Pareto optimality of (2, 2) allocation.

let us however admit non-stationary allocations: could there be a non-stationary allocation that dominates (2, 2)? In order to implement such a non-stationary allocation, a chain of inter-generational transfers would require a transfer from young to old at some arbitrary point in time t. The agents giving away endowment units in their youth would have to be compensated when old. The question is how many units of goods would be required for this compensation.

Figure 6.3 illustrates that, given an initial transfer ε_1 from young to old at t, the transfer ε_2 required to compensate generation t must be larger than ε_1, given the convexity of the indifference curves. This in turn will command a still larger ε_3, and so on. Is the sequence $\{\varepsilon_t\}_{t=0}^{\infty}$ formed feasible? No: eventually the transfer will exceed the young agent's endowment.

In the "simpler" case, where the equilibrium involves a gross real interest rate less than one, a constant transfer sequence is always possible: one can select another stationary allocation and it is better for everybody. In the case where the gross real interest rate is above one, the argument is as illustrated in Figure 6.3, except even harder, because now the young needs to be compensated even more when old, given that their indifference curves have a higher slope (in absolute value). So no feasible better path exists here either.

6.4.2 Intertemporal production

Intertemporal production in an overlapping generations economy raises further issues. In some cases, the introduction of the possibility to save—in the form of "capital"—can help overcome an inefficiency; in others, it can lead to new market failures. We will keep the discussion brief and merely provide some illustrations.

Let us start with the endowment economy $(\omega_y, \omega_o) = (3, 1)$, where we know that the equilibrium is not optimal. Suppose that there is a simple storage technology allowing the consumers to save $k \geq 0$ units today and receive χk units in return tomorrow. Clearly,

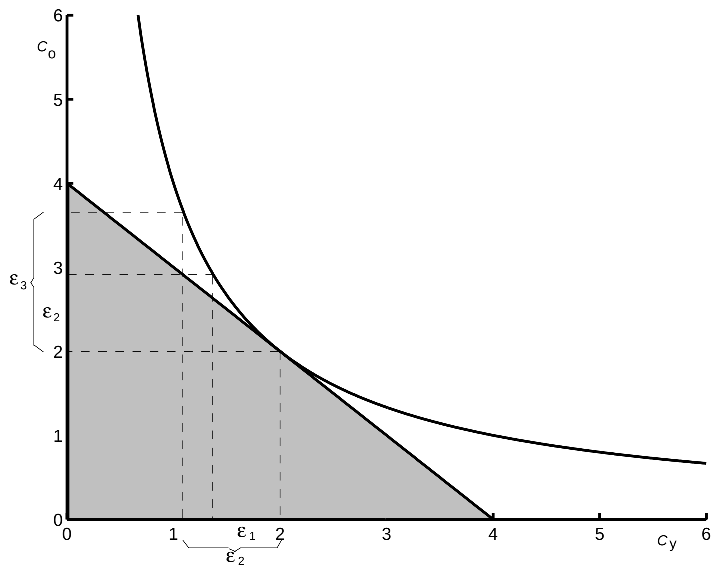

Figure 6.3: (2, 2) cannot be improved upon.

if $\chi > 1/3$ they will: it gives them a higher return than in the original equilibrium. Is the resulting equilibrium allocation optimal? We can solve for individuals' optimal saving choices by simply setting $q_t = \chi > 1/3$ in equations (6.5)–(6.6). This is our equilibrium, which again is autarky, but with active individual storage.

If we regard the equilibrium allocation as a new endowment point and ask if Pareto improvements are possible, the answer is given by the Balasko-Shell theorem: improvements are possible if and only if $\chi < 1$. However, this is not the final answer: here, the answer is that the equilibrium is actually optimal when $\chi > 1$ but not when $\chi \leq 1$. In particular, $\chi = 1$ is not optimal. The reason is shockingly simple. In the equilibrium, all young agents store 1 unit and, hence, achieve the consumption allocation $(2, 2)$. So far so good; this is in fact the same allocation for people born at $t = 0$ and later as that with our social-security scheme. However, the old at 0 are still consuming only 1 unit, so a Pareto improvement can be obtained by instead carrying out the social-security scheme: the old at zero are strictly better off and no-one else worse off (everybody else is indifferent).

The market failure with storage here, and $\chi = 1$, is an example of *dynamic inefficiency*: it is possible, by means of different saving choices, to create more resources at at least one point in time without forsaking resources at any other point in time. Thus, there is even a "free lunch" in equilibrium! The key insight here is that *oversaving* can occur in equilibrium. This is different than the market failure with fixed resources that we looked at before in our overlapping generations model: there, it was a matter of redistribution, and here, it is intertemporal production that is inefficient, again despite the absence of frictions.

As we shall see soon, the inefficiency in the special case with a one-for-one storage technology relies on the linearity of the production technology. If the intertemporal production technology is instead neoclassical, i.e., if it has decreasing returns to scale, the Balasko-Shell condition that the equilibrium is efficient if and only if, in the limit, the (gross) real interest rate is greater than or equal to 1, will be recovered. Let us therefore briefly revisit the neoclassical model here. Now let us interpret (ω_y, ω_o) as the endowments of a cohort in labor efficiency units. An agent's budget sets in period t and $t+1$ are then

$$c_y + s = \omega_y w_t \quad \text{and} \quad c_o = s(1-\delta+r_{t+1}) + \omega_o w_{t+1},$$

with s denoting saving. If the utility function is strictly quasiconcave, the saving choice at t is given uniquely by some function h:

$$s_t = h\left(w_t,\, r_{t+1},\, w_{t+1}\right)$$

Asset markets clear when $s_t = k_{t+1}$, where k denotes the economy's total capital stock: the young buy the entire capital stock for the next period. Competitive pricing of inputs as usual implies $r_t = F_1\left(k_t,\, \omega_y + \omega_o\right)$ and $w_t = F_2\left(k_t,\, \omega_y + \omega_o\right)$. Thus, our equilibrium can be computed as the solution to the non-linear first-order difference equation

$$k_{t+1} = h\left(F_2(k_t, \omega_y + \omega_o), F_1(k_{t+1}, \omega_y + \omega_o), F_2(k_{t+1}, \omega_y + \omega_o)\right).$$

This equation implicitly determines k_{t+1} as a function of k_t. We then have the following result that allows us to check whether the equilibrium savings choices are dynamically efficient.

Theorem 6.2 *Define $R_t = 1 - \delta + F_1(k_t, \omega_t)$, where ω_t is the total labor endowment at t. Then $\{k_t\}_{t=0}^{\infty}$ is dynamically efficient if and only if*

$$\sum_{t=0}^{\infty}\left[\prod_{s=1}^{t} R_s\left(k_s\right)\right] = \infty.$$

The theorem, whose assumptions are suppressed in the statement for brevity, relies in part on a proof of production efficiency that is similar in spirit to the proof of Pareto efficiency of an overlapping generations allocation with fixed resources.[12] The key point, however, is that it is the same condition as under fixed endowments: $\prod_{s=1}^{t} R_s\left(k_s\right) = 1/p_t$, where p_t is our Arrow-Debreu price of consumption good at t in terms of consumption good at 0.

To obtain intuition, it is instructive to restrict attention to steady states, i.e., solutions to

$$k_{ss} = h\left(F_2(k_{ss}, \omega_y + \omega_o), F_1(k_{ss}, \omega_y + \omega_o), F_2(k_{ss}, \omega_y + \omega_o)\right).$$

Clearly, from the theorem, a steady state is efficient if and only if $R = F_1(k_{ss}, \omega_y + \omega_o) + 1 - \delta \geq 1$, that is, if the net interest rate is non-negative.[13] Let us start with the case of inefficiency. When $F_1(k_{ss}, \omega_y + \omega_o) + 1 - \delta < 1$, an alternative saving plan is possible where savings at time t are reduced by a small amount $\epsilon > 0$. This frees up resources for consumption at t. At $t+1$, there are now fewer resources available, but the reduction is less than ϵ, given that the marginal return to saving was strictly less than one. However, suppose that at $t+1$, saving is again decreased by ϵ relative to the given steady state: then resources are freed up this period as well on net (there is less production, by less than ϵ, but also less saving, by

[12] The assumptions underlying it are that the sequence $\{R_t\}_{t=0}^{\infty}$ be uniformly bounded above and below away from zero and on the production-function curvature being bounded as follows: $0 < a \leq -f_t''\left(k_t\right) \leq M < \infty \;\; \forall t, \forall k_t$, where $f_t(k)$ is defined as $F(k/\omega_t, 1)$ and F is assumed to have CRS.

[13] In a balanced-growth version of this economy, the efficiency requires the interest rate not to be below the growth rate.

ϵ). This procedure is repeated and, as a result, there are strictly more resources available to consume at all points in time beginning with period t. Now suppose there is a steady state with a gross return one or greater than one. If one were to try to create more resources at some point t by reducing saving by ϵ, the resources available at $t+1$ would be reduced by more than ϵ.[14] The future reductions in saving required to keep resources at least as high as before will have to grow over time and will finally, become infeasible. The proofs of these statements, which are available Appendix 6.A, are the key behind understanding the logic of the theorem above.

In conclusion: in the overlapping generations model, equilibria—with or without production—can be efficient or inefficient. The key condition for efficiency, which holds rather generally, is that the asymptotic net real interest rate be non-negative or, in the case of growth, no less than the rate of growth.

6.4.3 Dynamic inefficiency in the warm glow model

The warm glow model, discussed briefly in Chapter 5, can also deliver outcomes with dynamic inefficiency properties. Let us illustrate this point with an example. Assume individuals live for one period only but give bequests according to the warm-glow model, with utility function $u(c_t)+v(b_{t+1})$ for agents alive at time t, and that they are in a standard neoclassical environment; capital accumulation comes from bequests. Thus, the agent's budget reads $c_t+b_{t+1}=(1+r_t-\delta)b_t+w_t$. Clearly, the offspring (at $t+1$) obtains b_{t+1} plus the net return.

Assume now that v is linear, $v(b)=Ab$ for some $A>0$, and that u is continuously differentiable and strictly concave. Then the optimization problem gives a solution for c_t from $u'(c_t)=A$; denote this solution $\bar{c}>0$. Bequests will then satisfy $b_{t+1}=(1+r_t-\delta)b_t+w_t-\bar{c}$. In equilibrium, $b_{t+1}=k_{t+1}$ so capital accumulation will be given by $k_{t+1}=(1+r_t-\delta)k_t+w_t-\bar{c}=(1-\delta)k_t+F(k_t,1)-\bar{c}$. This equation is a first-order difference equation—a convenience compared to the standard dynastic model—and it is easy to see from simple inspection of the function in (k_t,k_{t+1}) space, that there are two positive steady states. The lowest of these is unstable whereas the highest steady state, $\bar{k}$, is stable. Therefore, for a high enough initial capital stock, there is convergence to the higher steady state. At this steady state, moreover, the slope of the function, $1+F_1(\bar{k},1)-\delta$, is less than 1, but this also means that the long-run real interest rate $r-\delta$, is negative. Hence the economy is dynamically inefficient: there is "too much" capital accumulation, so that output could be increased at all times by lowering the capital stock, along the lines of the previous section. Intuitively, in this model people care about capital savings per se, leading to there being a free lunch in equilibrium.[15]

---

[14] If the net return is initially zero, it will raise above zero when we reduce saving.

[15] This can occur since the model has a non-standard feature. Note that for a statement about Pareto (in-)efficiency, one would need to think about which objective function to use. Lowering capital below the dynamically inefficient steady state could give all consumers more resources to consume, but they might not be happier given the warm glow from capital itself.

6.5 Optimal government policy

We have now seen a number of examples of how markets may deliver inefficient outcomes. A natural suggestion in each of these cases is to propose a government policy to improve on the allocation. In the case of distortionary taxes, the origin of the inefficiency is in government policy itself, but in the other cases, what would we, as macroeconomic analysts, tell the government?

But let us start with the tax case, because it is still interesting. In particular, it is often argued, and for good reasons, that lump-sum taxes are hard to implement in practice. Thus, tax analysis could be carried out by comparing tax policies that are deemed feasible to implement. For example, in a dynamic model one can compare proportional taxes on labor earnings to proportional taxes on capital income: which is the better system from a welfare perspective, and by how much? Such an approach is referred to as Ramsey analysis, after Frank Ramsey's early work (see Ramsey, 1927). The approach is, however, somewhat problematic since it is often not clear which policies are feasible and which are not. Ramsey taxation will be studied in Chapter 15.

The case of externalities is well understood; here, Pigou (1920) suggested a tax, or transfer, that would counteract the distortion and cancel its effects exactly. In the example of pollution, the externality could be corrected by charging a per-unit output tax on firms equal to what the externality would be, evaluated at the optimal allocation, and it would deliver the optimal allocation as an equilibrium outcome. Similarly, positive externalities can be encouraged with per-unit subsidies where they occur. A different path here would be to follow Coase: assign property rights so that the side effects of agents' actions can be incorporated as market transactions. This path can, potentially, be easier to take, since there is no need for the government to compute what an appropriate tax rate would be: once property rights are assigned and enforced, the property value and the price its owner will charge for its use will be determined by markets and, in the absence of further frictions, lead to an optimal allocation. Sometimes, however, as in the case of damages due to climate change, ownership cannot be assigned: the earth's atmosphere is not possible to own.

Monopoly distortions that result in inefficiently low production are easy to handle in principle: one can, for example, subsidize production at a per-unit rate. Again, this requires a calculation of the appropriate tax rate, which requires much detailed knowledge. Therefore, whenever it is possible, anti-trust regulation can be used to minimize the presence of monopoly pricing. Monopoly power, finally, can play another role in the economy: it can provide incentives to invent new products if, namely, there is patent protection (or it is difficult for other reasons to imitate the product). Invention is typically costly so to incentivize inventors one might wish to exclude others from using an invention. Thus, regulating against monopoly has drawbacks too; we will study this issue in Chapter 13.

6.5.1 Missing markets and the "chicken model"

What about the missing markets case? Recall the endowment economy where endowments alternated between agents and no borrowing was allowed. Here, a reasonably simple policy would seem to be available to the government: each period, tax the agent with a high endowment and transfer the proceeds to the other agent, so as to achieve full consumption smoothing. Both agents would then be better off, at least if the transfer is small enough. Similarly, in the context of missing insurance markets, the government could simply compensate people who received bad shocks and tax the luckier ones to finance the transfers. Is it reasonable to propose such a policy? This is not so clear. In reality, when a market is missing, it is usually missing for a reason. That reason could, for example, be problems of private information or moral hazard. If borrowing/lending and insurance are so beneficial, why do they not materialize, in the cases where they appear not to be present?[16] A consequence of this point is that it is entirely conceivable that there would be negative side effects of the governments intervention: those side effects that made markets missing in the first place.

The above considerations have given rise to the concept of *a "chicken model" of government*. The model here refers to an argument for government intervention in markets and goes as follows. Assume that (i) people like chicken; (ii) the market economy cannot produce them; and (iii) the government can produce this good. The result then follows: the government should produce chicken. From our perspective, the lesson should be: in cases where government intervention is proposed, think about which friction is at work, and whether it is one that the government is likely to be able to deal with well, or better than the market. Sometimes the answer is likely yes, and other times no.

Clearly one kind of approach available here would be to try to model the causes of frictions, such as private information, explicitly. Thus, analysis following the work of Mirrlees (1971) has been used to study optimal taxes and transfers when markets appear to be functioning imperfectly. Another reason why some markets may not exist is a lack of commitment, as discussed in Section 6.3.4 above.

6.5.2 Redistribution policy

The discussion so far centered around minimizing frictions. In practice, a separate aim is often redistribution, i.e., the idea is not to Pareto improve on the given allocation but rather to achieve a more equitable distribution of consumption even if some agents are made worse off. In macroeconomic models with heterogeneity, some of which will be studied later in this book, researchers often adopt a social welfare function to guide policy choices. In such cases, the welfare weights on different agents would represent the policymaker's preferences but, of course, not necessarily those of the researcher.

An often used social welfare function is an additive ("utilitarian") formulation. The most common assumption then is that the utilities of the agents are weighted equally. Clearly, equal weights would amount to equal consumption in an economy where direct, non-distortionary redistribution is available. Hence, equal weights embody a strong desire for equality. If taxes are distortionary, then equal weights still express the same desire but the optimal level of redistribution will typically not fully eliminate consumption inequality.

An argument for equally-weighted utilitarian social welfare functions that have been used is the "behind-the-veil-of-ignorance" notion. So imagine that a person, before they are born into a household somewhere in the world, possibly also without knowing what

[16] Think about whether you should write an insurance contract with your fellow graduate students, making sure that your post-graduation salaries are all the same, after taking transfers between you into account.

genetic skills they will have once born, is asked to consider potential distributional policies. Then redistribution could potentially be viewed as an optimal insurance scheme; in concrete terms, maximize $\pi u(c_A) + (1-\pi)u(c_B)$ subject to a resource constraint $\pi c_A + (1-\pi)c_B = \pi\omega_A + (1-\pi)\omega_B$, where $\omega_A > \omega_B$ would be the market incomes of the two types of agents; π is the fraction of people who will be born as A types. Clearly, this optimization problem embodies equal weights and delivers $c_A = c_B$. The insurance solution cannot be offered by markets, since the agents are not around to sign the contract before they are born, but governments can nevertheless carry out a policy which achieves the insurance outcome. Is this therefore a chicken model of government? Not so much; it is more seen a potential guiding philosophical principle with which you may agree or disagree.