Unfortunately, Varian's coverage of production is one of the weakest links in a (frequently weak) book. You can profitably disregard the first 8 subsections of the chapter. In those subsections, Varian considers a technology in which there is just one input to production, labour, and just one output being produced, and that's just too trivial an economy to spend any time considering.

As in the chapter (32) on exchange, the analysis becomes much more realistic and useful if we just expand the economy a little, to allow for 2 goods [as Varian does from subsection 33.9 on] and 2 inputs [which Varian doesn't really consider, except behind the scenes].

So why is a "production economy" a little more realistic than an exchange economy. In an exchange economy, goods are not produced. We just have a fixed endowment of them, so that the question cannot be addressed of how an economy might change the quantities of goods and services which are produced.

How are goods and services produced? I'll assume that there is some given *technology* for producing the different goods, using the different inputs to production. Now I'm going to assume that this technology is given, rather than something which might be changed through innovation or research. That obviously ignores some of the most interesting and important economic questions we actually face. But Varian makes that same assumption of a given technology : the production possibilities set in his figure 33.7 is taken as given, rather than something we might change (with some exertion of time and effort).

So the technology is something like a recipe book. The technology of production of good 1 could be described by a *production function* for good 1, some

$$x_1 = F^1(l_1, k_1)$$

where x_1 is the quantity produced of good 1 by some firm, which is a function of the quantities the firm uses of different inputs to production. In my example, those inputs are quantities of labour and of capital, indicated (respectively) by l_1 and k_1.

This notion of a production function is something Varian doesn't really mention in chapter 33. But you may have seen it before — it's discussed in a fair bit of detail in chapter 19 of Varian.

So how much of good #1 could a country produce? If, in total, firms producing that good employ L_1 person–hours of labour, and K_1 machine–hours of capital, then the total quantity of good 1 which we could produce would be

$$X_1 = F^1(L_1, K_1)$$

Naturally, we assume [as Varian does in chapter 19] that $F^1(\cdot, \cdot)$ is an increasing function of L_1 and K_1 : employing more workers or more machinery leads to higher output. (So does employing more raw materials, land and other inputs, if we had a more realistic model with more than 2 inputs to production.) The derivatives F^1_L and F^1_K of this production function with respect to the quantities used of the inputs are the **marginal products** of the inputs.

But here is the trade–off a country faces in production : people like both goods, good 1 and good 2. More of either good enables people in the country to consume that added output, and makes them better off. The way we can make more of good 1 is [only] to employ more workers, or more machinery, or both, in the good–1 industry.

But now I assume that there is a fixed total supply of workers and machinery in the country. So instead of assuming fixed endowments of consumption goods, as in chapter 32, I am pushing the problem back a stage by assuming fixed quantities (in total) of inputs to production.

That assumption is not realistic. In the short run, the supply of labour can be varied, as people choose whether to work more or fewer hours, or whether or not to enter the formal labour force. In the long run, the supplies of many inputs can be varied, as we invest more in new equipment (increasing the supply of capital), or in training and education (raising the effective supply of labour). But there are limits to the availability of inputs — for example land, or non–renewable energy, so that assuming fixed total quantities of the inputs is a simple way of recognizing those limits.

So I'm assuming that we have fixed total quantities available of labour and capital, denoted $\bar{L}$ and $\bar{K}$. That means that production of consumption goods are constrained by the following algebra

$$X_1 = F^1(L_1, K_1)$$

$$X_2 = F^2(L_2, K_2)$$

$$L_1 + L_2 \leq \bar{L}$$

$$K_1 + K_2 \leq \bar{K}$$

The last two inequalities capture the resource scarcity constraints : we can't employ more inputs than we have.

It is inefficient to leave resources unemployed, so that we'll want those last two inequalities to be equalities. Why leave machinery unemployed, when we can increase output of good 1 or good 2 by employing a little more of it?

But now, with $L_1 + L_2 = \bar{L}$ and $K_1 + K_2 = \bar{K}$, we have this conflict between the 2 industries. Making more of good 1 requires increasing L_1 or K_1 (or both), and that means making less of good 2.

So let me pick some $L_1 < \bar{L}$ and $K_1 < \bar{K}$ to be used in production of good 1. By allocating that much labour and capital to that industry, I am ensuring overall production of good 1 of $X_1 = F^1(L_1, K_1)$. But that allocation of labour and capital to industry 1 also determines how much labour and capital get allocated to industry 2, namely $L_2 = \bar{L} - L_1$ and $K_2 = \bar{K} - K_1$. And that determines how much we will get of good 2, $X_2 = F^2(L_2, K_2) = F^2(\bar{L} - L_1, \bar{K} - K_1)$.

So if I were in charge of the economy, and got to pick the input quantities (L_1, K_1) allocated to industry 1, then that choice will determine overall production (X_1, X_2) of each good. So any choice of (L_1, K_1), with $0 \leq L_1 \leq \bar{L}$ and $0 \leq K_1 \leq \bar{K}$ will determine what output production (X_1, X_2) we get.

And that's where the production possibility set in figure 33.7 comes from. Take all the possible input allocations, with $0 \leq L_1 \leq \bar{L}$ and $0 \leq K_1 \leq \bar{K}$ [and with $L_2 = \bar{L} - L_1$, $K_2 = \bar{K} - K_1$], and each of them will determine some overall quantities (X_1, X_2) produced. So we'll just put a dot in figure 33.7 for every possible input combination. And the set of all possible output combinations is the shaded area in figure 33.7.

Naturally, if I allocate more resources to industry 1, by increasing L_1 and K_1, then I will increase X_1, but I'll decrease X_2 since I am taking resources away from industry 2. And if I decrease L_1 and K_1 both, then I'll get less X_1 and more X_2. So in the first case I will move down and to the right in the shaded area. And in the second case I will move up and to the left.

What if I increase L_1 but decrease K_1? I don't know what happens to overall output of good 1. I am forcing industry 1 to substitute labour for capital, and whether that increases or decreases overall output depends on how good of a substitute one input is for the other — and on how much I added of one input and took away of the other. And this increase in L_1 and decrease in K_1 means that I'm decreasing L_2 and increasing K_2, forcing industry 2 to substitute capital for labour, a substitution in the opposite direction from that in industry 1.

Overall, this change might actually enable me to increase production of both goods, to move up and to the right in the shaded area in figure 33.7. That would happen if I arranged the input shifts relatively cleverly. If labour is much more productive at the margin in industry 1 than it is in industry 2, at the margin, then my shift is a smart idea. By increasing L_1 and decreasing K_1 I am moving labour from the industry in which it is relatively less productive (industry 2) to the industry in which it is relatively more productive (industry 1). And I am shifting capital from the industry in which it is relatively less productive (industry 1) to the industry in which it is relatively more productive (industry 2).

It runs out that such a "win–win" trade of inputs across industry is possible — whenever the relative productivities of the two inputs differ across industries. That will be the case whenever

$$\frac{F_K^1}{F_L^1} \neq \frac{F_K^2}{F_L^2}$$

where

$$F_K^1 \equiv \frac{\partial F^1(K_1, L_1)}{\partial K_1} \qquad a$$

and so on.

That is, whenever equation (a) holds, the resulting (X_1, X_2) combination is **inside** the production possibility set in figure 33.7 : there is some other feasible input allocation which leads to an outcome to the northeast, meaning more output of each good.

So there is a condition for efficiency in production — just as there was a condition for Pareto efficiency in exchange in chapter 32. In chapter 32, an allocation of goods and services among consumers was efficient if and only if the marginal rates of substitution between any two goods were the same for all consumers. That is what the tangency of the indifference curves in the Edgeworth box diagram meant.

The production problem could be described as follows

production problem : given the technology, and given the available input supplies, a production plan is efficient if and only if there is no other feasible production plan which yields more of every good

By "production plan" here, I mean a division of the available resources of each input among the different industries. An efficient production plan solves the following mathematical problem

production maximization problem : for a *given* quantity of good 1 being produced, try and produce as much as possible of good 2, that is
choose L_1 and K_1 to
maximize $X_2 = F^2(\bar{L} - L_1, \bar{K} - K_1)$
subject to the constraint $X_1 = F^1(L_1, K_1) \geq \bar{X}_1$
where $\bar{X}_1$ is the given required quantity of good 1

The production maximization problem can be solved using the method of Lagrange, by maximizing the Lagrangean function

$$\mathcal{L}(L_1, K_1, \lambda) \equiv F^2(\bar{L} - L_1, \bar{K} - K_1) + \lambda[F^1(L_1, K_1) - \bar{X}_1]$$

Maximizing this function with respect to L_1 and K_1 and λ means taking the derivatives of $\mathcal{L}(L_1, K_1, \lambda)$ with respect to L_1, K_1 and λ, and setting those derivatives equal to 0 :

$$-F^2_L + \lambda F^1_L = 0 \tag{L}$$

$$-F^2_K + \lambda F^1_K = 0 \tag{K}$$

$$F(L_1, K_1) - \bar{X}_1 = 0 \tag{\lambda}$$

Equations (L) and (K) imply the following required condition for efficiency in production :

$$\frac{F^1_L}{F^1_K} = \frac{F^2_L}{F^2_K} \tag{prodeff}$$

which is the exact opposite of my earlier condition (a) for production to be *inefficient*.

An allocation of inputs to industries will be efficient in production if and only if condition $(prodeff)$ holds if and only if it leads to an output combination on the production possibility frontier, the upper boundary of the production possibility set in figure 33.7.

And the nice thing about looking at the production technology, which Varian doesn't want to do, is it relates efficiency in production to the fundamental theorems of welfare economics in the previous chapter.

The constrained optimization problem I described above is one that a central planner would want to solve if she controlled directly the whole economy : "how can I maximize production of military weapons subject to allocating enough resources to food production that people don't starve". But we know that there are a lot of problems with economies run by a central planner.

What would happen if, instead, the economy were not planned, but consisted of a lot of competitive profit–maximizing firms, many in each industry. If I run a firm which produces good 1, I have to hire labour and capital, and pay them for their services. In a competitive economy, I have to pay workers a wage of w per hour, and rent machinery for r per hour, where w and r are the market prices of labour and machinery respectively. Then I take the output that these worker and machines produce (together) and I sell it on output markets. The market price of what I produce is p_1, so I am producing good 1. In a competitive economy, I take the market prices of both my output of good 1, and my inputs of capital and labour as given.

So my problem, as owner of a competitive good–1–producing firm, is to choose labour and capital inputs l_1 and k_1 to maximize my profits

$$\pi_1 = p_1 F^1(l_1, k_1) - wl_1 - rk_1$$

The first term in the expression above is my sales revenue, the price per kilo of good 1, times the quantity $F^1(l_1, k_1)$ which my firm produces. The second and third terms are my labour and capital expenses. Choosing l_1 and k_1 to maximize the above expression means setting the derivatives equal to 0 :

$$p_1 F^1_L(l_1, k_1) - w = 0 \qquad (L1)$$

$$p_1 F^1_K(l_1, k_1) - r = 0 \qquad (K1)$$

Similarly, a firm in the other industry, producing good 2, chooses its inputs l_2 and k_2 to maximize its profits $p_2 F^2(l_2, k_2) - wl_2 - rk_2$, resulting in first–order conditions

$$p_2 F^2_L(l_2, k_2) - w = 0 \qquad (L2)$$

$$p_2 F^2_K(l_2, k_2) - r = 0 \qquad (K2)$$

Now if I combine equations $(L1)$ and $(K1)$ I get

$$\frac{F^1_L}{F^1_K} = \frac{w}{r} \qquad (e1)$$

and if I combine equations $(L2)$ and $(K2)$ I get

$$\frac{F^2_L}{F^2_K} = \frac{w}{r} \qquad (e2)$$

So if the firms in both industries have to pay the same wages for labour, and the same price for capital, then equations $(e1)$ and $(e2)$ mean that

$$\frac{F^1_L}{F^1_K} = \frac{F^2_L}{F^2_K} \qquad (eq)$$

which is **exactly** my condition $(prodeff)$ above for efficiency in production.

welfare theorem for production : if output is produced by profit–maximizing price–taking firms, then the overall production outcome will be efficient, lying on the production possibility frontier

Just as there are lots of Pareto efficient allocations of a given endowment of goods, in chapter 32, there are many production plans which are efficient. Any allocation of inputs which results in a production plan on the production possibility frontier is efficient. And the output combinations in the frontier depicted in figure 33.7 cover quite a range : from none of good 1, and the maximum possible amount of good 2 on the top left to the bottom right point, yielding none of good 2 and the maximum possible quantity of good 1.

Which efficient production plan is best depends on what the consumers want to consume. There has to be some link between the production decisions in chapter 33, determining the available quantities of goods and services, and consumers tastes for those goods and services, determining how they are allocated among consumers.

Figure 33.9 illustrates the relationship between the two sides of the economy. What it shows is a production plan in which the slope of the production possibility frontier equals the slope of the consumers' indifference curves.

Some terminology here. Recall that the (absolute value of the) slope of a consumer's indifference curve is called the marginal rate of substitution, since it measures the rate at which she is just willing to substitute one good for the other. The (absolute value of the) slope of the production possibility frontier is called the *marginal rate of transformation*, since it measures the rate at which one good can be transformed into another.

That is, suppose we are at some point on the production possibility frontier, producing X_1 kilos of good 1 and X_2 litres of good 2. We know that there is no feasible way of producing more of both goods. If we want more production of good 2, we will have to produce less of good 1 : we'll have to move up and to the left along the production frontier. How much more will we get of good 2, if we reallocate labour and capital so as to produce 1 kilo less of good 1? That depends on the technology of course. But the slope of the possibility frontier is exactly the answer to that question. If the equation of the production possibility frontier is $X_2 = G(X_1)$, then its slope $G'(X_1)$ is dX_2/dX_1, which is the change in production of good 2 associated with a 1 kilo change in the production of good 1.

A person's $|MRS|$ is how much she wants to get of good 2, in order to give up 1 kilo of good 1.

Now suppose we had a production combination (X_1, X_2) which was on the production frontier, and suppose that the MRT was greater in absolute value then every person's MRS. For example, suppose that the MRT was 4, and each person's MRS was 2. Consider the following change in production : reallocate labour and capital from good 1 production to good 2 production. Reallocate just enough capital and labour so that $\Delta X_1 = -2$ and $\Delta X_2 = 8$. How do we know that we can do this? If the $|MRT|$, that means exactly that every 1 kilo reduction in production of good 1 results in a 4 litre increase in good 2 production. So the fact that, at the current production plan, $|MRT| = 4$ means that a 2–kilo reduction in good 1 production results in an 8 litre increase in good 2 production, if we allocate the inputs efficiently.

Now take those additional 8 litres of good 2, and give each of the two people 4 litres. Of course, at the same time, we have reduced overall production of good 1 by 2 kilos, so we'll have to reduce people's consumption. So reduce each person's consumption of good 1 by 1 kilo.

So, with $|MRT| = 4$, we have chosen to reduce overall production of good 1 by 2 kilos, and increase overall good 2 production by 8 litres, and then we have divided these changes evenly between the two people. How do the people feel. Person A has seen her consumption of good 1 fall by 1 kilo and her consumption of good 2 increase by 4 litres. But I said earlier that her $|MRS|$ was 2 : she was just willing to give up 1 kilo of good 1 in order to get 2 litres more of good 2. Making her give up 1 kilo of good 1, and giving her 4 litres (instead of 2) of good 2, makes her better off, since she's getting more than she needed to compensate for the reduction in her consumption of good 1. The same improvement in well–being happens to person B, since he got the same deal, and his $|MRS|$ also was 2.

So I started from a situation in which $|MRT| > |MRS^A| = |MRS^B|$, and I found a feasible way to make both people better off, without making any person worse off. That means that the original allocation could not have been Pareto efficient. On the other hand, if both people's $|MRS|$'s exceeded the $|MRT|$ then I could have made them both better off by increasing production of good 1, reducing production of good 2, and splitting the changes in production between the two people. Therefore

overall efficiency : An production plan and allocation of the goods produced will be Pareto efficient only if $MRT = MRS^A = MRS^B$.

That means that, in an economy with production, there are three sets of conditions which must be satisfied for the overall allocation to be efficient. The first set of conditions was what was derived in chapter 32 : they hold for an exchange economy, and they still must hold for an economy with production. The second set of conditions came in this section (but were ignored by Varian) : they ensure that there is no better way of reallocating inputs to production so as to get more of every good produced. And the third set of conditions was the one I just discussed, relating the production and consumption sides of the economy.

For an allocation of inputs among industries, and of goods among people, to be Pareto efficient, the following three sets of conditions must hold :

(*exch*) for any pair of consumption goods, any two people must have the exact same marginal rate of substitution between these two goods

(*prod*) for any pair of inputs, and for any two firms [maybe in different industries], the *technical rate of substitution* between the inputs must be the same for both goods, where the technical rate of substitution is defined [in chapter 19] as the ratio of the marginal products of the two inputs (the slope of the firm's *isoquant*)

$$MP_L^i/MP_K^i = MP_L^j/MP_K^j$$

for any two firms i and j, and any two inputs K and L

(*overall*) the marginal rate of transformation between any 2 goods must equal the marginal rate of substitution of any consumer between those goods

What is the marginal rate of transformation? It's the slope of the production possibility curve. But what does it equal?

Suppose that we want to reduce production of good 1 and increase production of good 2. One way of doing that would be to reallocate 1 hour of 1 worker's labour from the first industry to the second. That would increase production of good 2 by MP_L^2, and decrease production of good 1 by MP_L^1. The ratio of the changes, the increase in X_2 production divided by the decrease in good 1 production, would be

$$|\frac{dX_2}{dX_1}| = \frac{MP_L^2}{MP_L^1}$$

On the other hand, we could accomplish these changes in output by reallocating capital instead of labour, so that

$$|\frac{dX_2}{dX_1}| = \frac{MP_K^2}{MP_K^1}$$

But notice that, if we are efficient in production — which we must be if we are on the production possibility frontier — condition (*prod*) for efficiency says that

$$\frac{MP_L^1}{MP_K^1} = \frac{MP_L^2}{MP_K^2}$$

so that both ways of changing production give the same answer. However we reallocate inputs at the margin, it must be true that

$$|MRT| = \frac{MP_L^2}{MP_L^1} = \frac{MP_K^2}{MP_K^1}$$

Now what has happened to my theorems of welfare economics. In chapter 32, Varian showed that a price–taking market economy would be efficient in exchange — condition (*exch*) must hold in any competitive equilibrium. A few pages ago, I showed that if firms were price–taking profit maximizers, then condition (*prod*) must hold in any competitive equilibrium. In showing that, I used the firms' conditions for profit maximization. Each firm's optimal choice of labour to hire obeyed the equations

$$p_1 F_L^1(l_1, k_1) - w = 0 \qquad (L1)$$

$$p_2 F_L^2(l_2, k_2) - w = 0 \qquad (L2)$$

Those equations imply that

$$p_1 F_L^1(l_1, k_1) = p_2 F_L^2(l_2, k_2)$$

or

$$|MRT| = \frac{F_L^2}{F_L^1} = \frac{p_1}{p_2}$$

So if production is done by competitive profit–maximizing firms, we must be at a point on the production possibility frontier at which $|MRT| = p_1/p_2$.

If consumers are price–takers, and choose consumption levels to get to the highest indifference curve on their budget line, then chapters 2 – 9 show that they'll pick a consumption bundle at which their marginal rate of substitution equals the slope of their budget line. That means that, in a competitive economy

$$|MRS^A| = |MRS^B = p_1/p_2$$

So the third set of efficiency conditions, the conditions for overall efficiency, must be satisfied in equilibrium in a perfectly competitive market economy. The theorems of welfare economics continue to hold when we consider a more realistic environment in which goods and services are produced, using various inputs.

The only form of market organization which I (or Varian) have mentioned is a perfectly competitive market economy. And the first theorem of welfare economics says that the allocation resulting from this form of organization is Pareto efficient : no other form of economy organization can give rise to an allocation which is preferred by everyone.

So are there other forms of market organization? Are there forms of organization which lead to inefficient outcomes?

Here I will just provide two examples of inefficient policies, just to show that the efficiency of some form of organization is not a guaranteed outcome. Both examples are taxes, although they're certainly not the only examples of policies which lead to inefficient outcomes.

First, consider a tax on the use of capital by firms producing good 1. That is, firms producing good 1 must pay a price of $r + t_K^1$ for every unit of capital they use, r to the owner of the equipment and t_K^1 in taxes to the government. This tax is levied only on capital employed, not on labour. And it's levied on on firms producing good 1, not firms producing good 2.

In that case, firm 1's profit is $p_1 F^1(l_1, k_1) - wl_1 - (r + t_K^1)k_1$, and when it chooses how much labour and capital to hire, so as to maximize profits, it's first–order conditions for profit maximization are

$$p_1 F_L^1(l_1, k_1) - w = 0 \tag{L1}$$

$$p_1 F_K^1(l_1, k_1) - (r + t_K^1) = 0 \tag{K1'}$$

A firm in the other industry, producing good 2, chooses its inputs l_2 and k_2 to maximize its profits $p_2 F^2(l_2, k_2) - wl_2 - rk_2$, resulting in the same first–order conditions as before

$$p_2 F_L^2(l_2, k_2) - w = 0 \tag{L2}$$

$$p_2 F_K^2(l_2, k_2) - r = 0 \tag{K2}$$

From equations $(L1)$ and $(K1')$

$$\frac{F_L^1}{F_K^1} = \frac{w}{r + t_K^1}$$

and (as before) equations $(L2)$ and $(K2)$ imply that

$$\frac{F_L^2}{F_K^2} = \frac{w}{r}$$

That means that — when the tax rate does not equal zero —

$$\frac{F_L^1}{F_K^1} \neq \frac{F_L^2}{F_K^2}$$

So putting a tax on a single input — but only in a particular industry — results in choices by firms which do **not** satisfy the condition for efficiency in production. It results in an output combination (X_1, X_2) which is strictly inside the production possibility frontier (when firms are price–taking profit maximizers).

So, in a competitive economy, taxing the use of one input in one industry results in an inefficient outcome. Slightly modifying the algebra, so does taxing the use of that input at different rates in different industries.

Who would be so stupid as to set tax rates in this fashion? The government of Canada — and of Canadian provinces, and of many many jurisdictions. Canada levies an income tax on the income of corporations, as do 9 of the 10 Canadian provinces. The rules of that tax are quite complicated, but they determine the effective cost to a firm of using more capital, after all the tax consequences are taken into account. (And firms should, and do, take those tax consequences into account. It's their net–of–tax profits which matter tot he firms' owners.) And the way our corporate income tax rules are structured, the tax rate on the use of capital and equipment varies a lot across industries, and it varies a lot across provinces. That means that the tax rate on the use of capital by an oil exploration firm in Alberta is much much lower than the tax rate on the use of capital by a grocery chain in Ontario. These variations across industries and inputs in the effective tax rate lead to output choices in Canada which are inefficient — we could increase production of both oil and groceries by reallocating inputs among industries.

My second example also involves taxes. But this time it's on the consumption side. Suppose that some government levies a 13% tax on the consumption of good 2 – but **not** on the consumption of good 1. That means that any consumer faces a price of p_1 on good 1, and a price of $P_2 \equiv p_2(1+\tau_2)$ on good 2, where τ is the proportional tax rate of the consumption tax on good 2, 13% in this case.

Now if a consumer is buying good 2, she faces a price of P_2 for good 2, and p_1 for good 1. The slope of her budget line is p_1/P_2, not p_1/p_2. And

$$p_1/P_2 = \frac{p_1}{p_2}\frac{1}{1+\tau}$$

So a consumer who chooses a consumption bundle so as to reach the highest indifference curve on her budget set will choose a consumption bundle where her indifference curve is tangent to a line with slope p_1/P_2, so that, if person A is such a person,

$$|MRS^A| = \frac{p_1}{P_2} = \frac{p_1}{p_2}\frac{1}{1+\tau}$$

Now private firms make their profit–maximizing decisions using the net–of–tax prices p_1 and p_2 : p_2 is what a firm gets for each litre of good 2 it sells, since the rest of the consumer's payment, the taxes of τp_2 don't go to the seller, but to the government. So profit–maximizing behaviour by firms results in a production plan on the production possibility frontier, and we have

$$|MRT| = \frac{p_1}{p_2}$$

The condition for overall efficiency does not hold when $\tau \neq 0$, so that the allocation is not Pareto efficient. It is possible to reallocate inputs from one industry to another in such a way as to make all consumers better off.

Of course, the example I have chosen is an actual Canadian example. Consumer purchases of most goods and services are subject to HST, and the rate of 13%. But notice that the inefficiency in my calculation resulted not from the assumption that consumer purchases were taxed, but from the assumption that **only some** consumer purchases were taxed. The introduction of the current HST in Ontario, about 5 year ago, involved an expansion of HST coverage. And the reason for this expansion of coverage was to address exactly the inefficiency in my example. If all goods and services were subject to HST at the same 13% rate, then consumers would pay a price of $P_1 = p_1(1+\tau)$ on purchases of good 1 as well, so that $P_1/P_2 = \frac{p_1(1+\tau)}{p_2(1+\tau)} = p_1/p_2$, so that we would have $MRS^A = MRT$ in equilibrium, and overall efficiency. The continued inefficiency of Canada's HST stems from the fact that there still are quite a few goods which are exempt from taxation : groceries, children's clothing, repairs which are done at home by the owner–occupier. As long as we have significant exemptions from HST, the tax gives rise to an inefficient allocation.

So there are two examples of inefficient policies. Does that mean that we would all be better off if our governments got rid of the corporate income tax and the HST? As a practical matter, not really. First of all, we can't get rid of all taxes, unless we want to get rid of all expenditures by all levels of government. So, if anyone who believes that some significant government expenditure is needed had better find a good tax system. [And you do believe in government expenditure if you want to keep any of : subways, snow plowing,public education, police services, unemployment insurance, hospitals.]

The analysis of this chapter, and of chapter 32, suggest that a tax system leads to an efficient allocation only if it does not distort any of the three sets of efficiency condition. I have described briefly how sales taxation, such as the HST, or corporate income taxation, both lead to inefficiency as currently structured. But if I had the time to look into people's job choice decisions, and people's education decisions, I could show that most realistic taxes give rise to significant inefficiencies. That includes a "more efficient" HST, even if it did cover all market purchases by consumers. And it includes a personal income tax, which is quantitatively the most important government revenue source in Canada.

The only taxes which do not lead to inefficient outcomes are what Varian calls **lump–sum** taxes in subsection 2.6. Those are taxes in which your own tax liability is unaffected by any decisions you might make : about what to buy, where to live, where and when to work. And it is very hard to think of examples of taxes with that property. It's even harder to think of taxes which have that property and which seem reasonably fair, in that they don't put much of the tax burden on those least able to bear it.

So the disappointing policy conclusion is that, realistically, we have to live with distortions from Pareto efficiency. How big those distortions should be, and what we can do about them, are very important application of microeconomics. But I'll have to defer those topics to other courses (particularly econ 4070 and 4080).