



Analysis of Markets

CLASS NOTES

COMM 220

Gregory Lypny

About these notes

These are draft class notes for the topics that we will be discussing in Analysis of Markets. Neoclassical economics is emphasized. We will deal with behavioural economics in our class discussions. You can merge the two in your own notes.

If you intend to print these notes, I suggest printing one topic at a time because I sometimes post updates.

Gregory Lypny

Tuesday, August 26, 2014

NOTE 1

Building Economic Models

Twelve students take part in an experiment. Each is randomly assigned to be a buyer or a seller of an otherwise worthless metal token. Buyers are told that if they buy a token, they can turn it in at the end of the experiment and be paid its redemption value. The redemption value is different for each buyer, and is known only to the buyer. Sellers are given one token each, called an

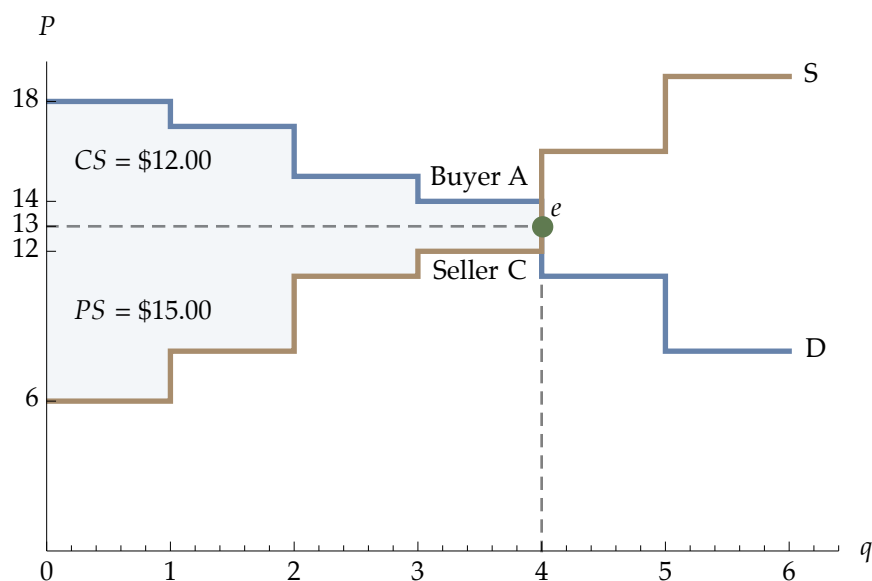
Redemption Value		Cost	
Buyer A	\$14	Seller A	\$16
Buyer B	\$17	Seller B	\$6
Buyer C	\$18	Seller C	\$12
Buyer D	\$11	Seller D	\$8
Buyer E	\$15	Seller E	\$11
Buyer F	\$8	Seller F	\$19

endowment, and told that if they sell their token, they can keep the sale price less the token's cost. The cost is different for each seller, and is known only to the seller.

The buyers and sellers trade by submitting their offers privately to Raymond, a skinny, bearded guy wearing Buddy Holly glasses and a Keep on Truckin' t-shirt that he bought in 1968. Raymond sorts the buyers' bids from highest to lowest and the sellers' asks from lowest to highest, and if there is a price for which the number of tokens that would be bought is equal to the number that would be sold, Raymond will call out that price as the *equilibrium* or market-clearing price. This type of market is called a *call market* because all of the offers are batched. All buyers who bid the market clearing at least the market-clearing price will receive a token supplied by the sellers who asked no more than the market-clearing price. What is your hypothesis for the equilibrium price and quantity in this market?

I think not more than four tokens will be traded at a price of about \$13. It may not happen the very first session of the experiment because the subjects may need some time to learn, especially since redemption values and costs are private information. I do know that it would be strange if the number of tokens traded were greater than four (can you say why?). Four tokens

at about \$13 should be what we observe on average. Draw supply and demand schedules like Raymond would.



If buyers want to earn some cash, they will bid less than their redemption values. I don't know how much lower; too low, and there is a risk of slipping to the right of the equilibrium and earning nothing. If sellers are greedy too, they will ask more than cost, but not too much more. So, it looks like greed is necessary to put us in the neighbourhood of e .

The \$13 equilibrium price needs a word or two because it is, after all, being announced by Raymond. It isn't the only possible equilibrium price. You can see from the graph that Buyer A and Seller C, at the margin, determine the equilibrium. Buyer A will buy a token as long as the price no higher than \$14, and Seller C will give one up as long as the price is no less than \$12. So that fourth token would trade at any price from \$12 to \$14. The equilibrium price is not unique. But Raymond's job as auctioneer is to call out one equilibrium price, if it exists, so he had to have a rule to deal with multiple equilibrium prices, and his rule, which he announced to everyone at the beginning of the experiment, is to use the mid-point of the marginal bid and ask. It didn't have to be split down the middle; that's just the way Raymond's head works.

Of models and assumptions

The world is a complicated place. There's a lot going on all of the time. Economic models, like all scientific theories, are simplified versions of some small bit of reality. A theory sets out the conditions or variables that are necessary to answer the question being asked. *What should be the equilibrium in a call market when information about cost and redemption value are private?* Simple is usually best: being able to explain or predict something with just two variables is better than needing three. A theory is not intended to duplicate reality but capture enough of we want to

explain or predict. When a variable important to a theory cannot be measured because it is unobservable, an assumption has to be made about it. What assumptions were made to arrive at the hypotheses in the tokens experiment? Greed is one: people prefer having more to less. Another is that people place offers without thinking that their offers could somehow influence the equilibrium. They behave as if they are price-takers. The tokens experiment has been done many times, and the hypothesis is strongly supported. That's a brownie point for economics because, in this simple market, goods flow from the people who value them least (the lowest cost sellers) to those who value them most (the buyers with the biggest redemption values). The value of trade to the people in this market is reflected in the consumers' and producers' surpluses shown as CS and PS in the graph. They are better off than they would be without trade.

Something to think about

What if each the sellers were endowed with an I-Love-Concordia coffee mug, and everyone knew that the mugs was selling for \$13.89 in the university bookstore?

Review

ask
assumption
bid
call market
consumer surplus
endowment
equilibrium
greed (self-interest, selfish, prefer more to less)
hypothesis
price-taker
private information
producer surplus
theory

NOTE 2


Utility Theory

Do I want to keep the I-Love-Concordia coffee mug that I had been given as a seller in the experiment in Note 1. Or should I sell it? Would I want to buy one if I were assigned the role as buyer? Economics is all about choices, and it assumes that we make choices that are best for us, that gives us the greatest satisfaction. The tokens version of the experiment was designed so that greed would motivate the decision to buy or sell and the offers to be made. The design, of course, stemmed from the assumption that people are greedy. Assuming greed seems like a pretty safe bet, but then there's philanthropy, altruism, and self-destructive behaviour, which do not fit so neatly in the best-for-me mould. The assumption of greed is not enough to form a hypothesis about equilibrium price and quantity, or whether there would any trade at all, in the mugs version of the experiment because everyone was told the price of the mug in the university bookstore, and everyone knew that the price was public information. A slightly broader assumption had to be made: people have different preferences or tastes for mugs. Some sellers might offer theirs for sale because they prefer cash, and some buyers might one. We can get a rough idea about someone's tastes and whether they are greedy after the fact, that is, by observing the actual choices they make. But it is much harder to do before the fact; tastes and greed are *unobservable* in that sense.¹ Utility theory uses mathematical functions represent tastes and greed in an economically meaningful way and to avoid having to measure them. Choices as well as responses to changes in prices and income can then be predicted.

¹ In science lingo, *ex ante* means before the fact or looking ahead, and *ex post* means after the fact or looking back at the past. You can also think of greed as the one aspect of our tastes that is the same for everyone.

Income, wealth, and budgets

Consuming the things that you have chosen is what gives you satisfaction. Call these things x and y . x and y are *flows* of goods, services, or any activity that is consumed.² Groceries per week, haircuts per year, and time spent listening to the blues are all consumption flows; combinations of these are sometimes called consumption bundles. Your *income* is a flow too. Income together with the prices of goods defines your budget, which determines how much you can consume. Suppose your weekly income is \$100, the price of x is \$5, the price of y is \$2, and you don't happen to own any x or y . You could consume 20 units of x per week if you spent all of your income on x , 50 units of y if you spent all of it on y , or any combination of the two that



doesn't cost more than \$100. All of the bundles whose cost equals your income lie on a budget line.³ Bundles a , b , and e are just some of the many bundles that are affordable, and you would choose the one that you preferred to all others on the line. The slope of the budget line is

$-\frac{1}{2} = -\frac{\$5}{\$2} = -\frac{p_x}{p_y}$ because x is two and half times more expensive than y . An x is worth $2\frac{1}{2}$ y 's.

The *price ratio*, $\frac{p_x}{p_y}$, and not the individual prices, is what is important in making choices. That

the price of x is \$5 or \$5,000 doesn't mean much unless you know the market prices of other

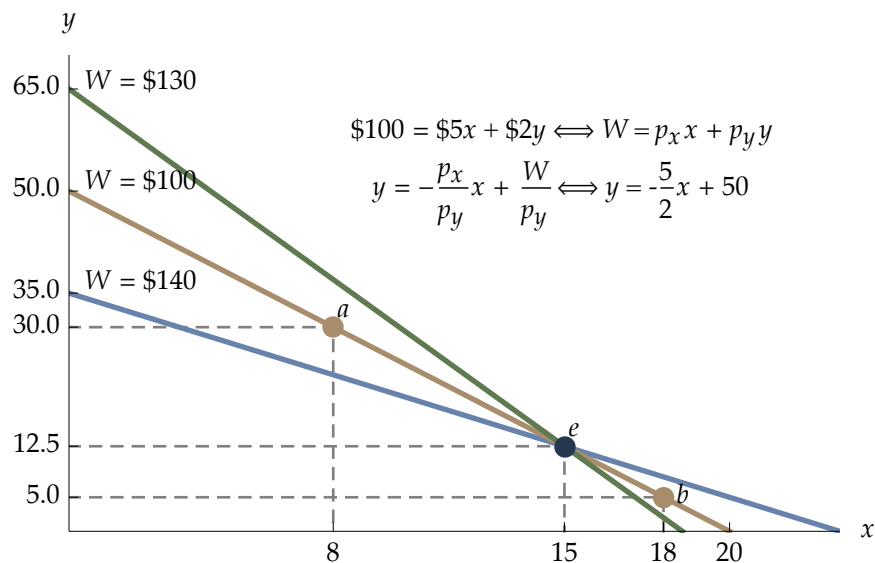
² To keep it short, I'll use goods to mean goods, services, or activities.

³ It is a line if there are only two goods or only two are being considered. It is more generally called a budget constraint or income constraint.

things that might be consumed along with x or instead of x . (I'll leave it to you to review how the budget line changes position if one or both prices or income changes. You've done it before.)

If you have no income but are endowed with 15 x and 12½ y then your *wealth* is the value of your *endowment* at market prices—whatever you own or whatever you're entitled to—and that's \$100 (\$75 worth of x and \$25 of y). Your wealth is \$100 everywhere along the budget line or wealth constraint, so long as prices stay the same. Wealth, unlike income, is a *stock*; it is a value at a point in time.⁴ A warehouse inventory, a firm's total assets, the number of bottles of beer in your fridge, my collection of Beatles records, and my daughter's 43 Beanie Babies are all stocks.

The budget line looks exactly the same whether you have an income (but no endowment) and are deciding how to spend it, or you have an endowment of goods (but no income) and are deciding whether to trade off some of one good to get more of the other. The only difference is that both the x and y intercepts of the wealth line change even when only one of the prices has




changed. This is because the wealth line always passes through the endowment point. This makes sense because, no matter how much the price of goods changes, you can always afford what is already yours. Can you say how prices must have changed for the initial wealth line (\$100) for it to become the wealth line at \$140 or \$130?

Utility functions and indifference curves

Your budget fixes the choices that you can *afford* to make. Your tastes set out the choices that you *want to* or are *willing* to make. Imagine listing all of the bundles of x and y that give you

⁴ I flunked accounting class seven times.

exactly the same satisfaction as your endowment of 15 x and $12\frac{1}{2} y$. This set of bundles is called an *indifference curve*, and it is how tastes are modelled in economics. Maybe your tastes for x and y can be represented by indifference curve shown in the figure. The value of the function at e is about 7.40852. This number is called *utility* and can be taken as a measure of how satisfied you



are with 15 x 's and $12\frac{1}{2} y$'s.⁵ All of the other points on the indifference curve are bundles that yield utility of 7.40852.


Bundles that a person prefers must lie on higher indifference curves. If someone offered you bundle g in return for e , you would take it. The utility of g is higher (I'll let you do the math), but it's not like you had a higher numerical value of utility in mind when you made the trade. You simply prefer g to e . Utility is not a real, physical, measurable thing like temperature, gravity, or the number of hairs on your head. It is just a ranking that is implied by the choices we make. But actual tastes are difficult to measure before a choice is made or if a choice is not observed. I don't think that many people believe that our tastes can truly be captured by mathematical functions, but if the functions are chosen carefully, they have plausible economic interpretations. A person like you with square root utility would choose g over e . Someone else with tastes represented by another function might not.

So what economic meaning is in the shape of indifference curves? They slope downward because they represent the tradeoff between two goods that a person is willing to make. Take away some of my x , and you'll have to give me a certain amount of y in return to leave me as well off as I was before. The tradeoff always exists if both goods are "good" in sense that more is preferred to less. It is not the case if one of the goods is a "bad" like secondhand cigarette smoke or construction noise. The steepness, ignoring the negative sign, is what really defines a per-

⁵ Utility means satisfaction, happiness, welfare, or wellbeing. Better off is an increase in utility and worse off, a decrease.

son's tastes because it is the most y that a person is willing to give up in order to gain one more x or the least y they demand in return for giving up one x . The absolute value of the slope is called the *marginal rate of substitution*. Think of it as a personal price ratio defined by personality or tastes. The marginal rate of substitution at e for the square root utility function is about 0.91: one more or one less x is worth 0.91 y you if you currently have 15 x and $12\frac{1}{2}$ y . The marginal rate of substitution at a is about 1.3 and at b , 2.3. Why does a unit of x increase in personal value moving from right to left along an indifference curve? Relative scarcity. You have less x compared to y at b than you have at a than at e . It makes sense that we value something more highly as it becomes scarcer relative to other stuff. Economics handles this by using convex functions for indifference curves. The functions chosen for indifference curves are also those that never flatten out completely because that would imply that there comes a point when you have so many x 's that you are unwilling to pay anything for one more (MRS is zero). That wouldn't sit well with greed assumption. If you prefer more to less, then one more slice of pizza always provides you with some positive utility, no matter how small, and it does not matter that you have just finished your 17th slice.


We have different tastes. I wish that high-waisted, pleated pants would come back in style. I bet that you do not. The marginal rate of substitution is the only way to tell one person from another. The next figure shows one of my indifference curves along with yours. I'm the log utili-



ty guy. We both have the same endowment. Who places a bigger value on the next unit of x ?

Consumption optimum

Utility theory brings together budgets and utility functions in a simple mathematical optimization: *people make choices that maximize their utility subject to their budgets*. English translation: people make the best choices for themselves that they can afford. As a person whose tastes are represented by a square root utility function, you would choose bundle C^* , about 5.7 x and



35.7 y , because this brings you the highest utility. I'll let you check that utility at the optimum is 8.3666. There is no other bundle on the budget line that gives higher utility because the indifference curve that passes through C^* is tangent to the budget line; it touches the line at just that one point. Moving the indifference curve any higher would place it above the budget line, and those bundles are not affordable. Because the optimum is a tangency, it is determined by equating the slope of the indifference curve to the slope of the budget line. Because the negative signs of the slopes cancel when they are set equal to one another, C^* is chosen so that the marginal rate of substitution is equal to the price ratio, which is constant at $2\frac{1}{2}$.

$$C^* : MRS = \frac{p_x}{p_y} = \frac{\$5}{\$2}$$

The marginal rate of substitution is the rate at which you are willing to trade x for y , and the price ratio is the rate at which the market is willing to trade them. You cannot make yourself better off when you reached the point where the value you place on the next unit of x is the same value that the market places on it. How did you get there? If utility theory says that people maximize utility, they are always at C^* . Why would anyone be anywhere else? If there was a change in prices or income or some other event that moved them away from C^* , they would quickly do whatever was necessary (like trade) to get back to C^* . That implies that we really

wouldn't find people at e for any appreciable length of time unless, of course, e and C^* were one and the same. But not many would understand utility theory if they were told that C^* simple is *because it is*, so the explanation almost always includes a story of how the person trades from their endowment to move up along the budget line until they reach C^* . At e , your marginal rate of substitution is less than the price ratio, which means that you value the next unit of x less than the market does. So you sell some of your x 's in return for $2\frac{1}{2}$ y 's each. As you climb up the budget line, your marginal rate of substitution increases because x is becoming relatively scarcer to you. The bliss point is reached at C^* when your marginal rate of substitution is exactly equal to the price ratio, and you stop selling.

Things to think about

Bundles that lie below the line. Affordable? Would a person choose one?

What is the difference between a utility function and an indifference curve?

What is the utility of bundle g ?

Why can't your indifference curves cross one another?

Wealth versus welfare?

There is something wrong about using the square root utility function for modelling utility. What is it?

How would you draw an indifference curve if one of the goods was a "bad," like second-hand cigarette smoke or construction noise?

Review

consumption optimum

endowment

flow vs. stock

income vs. wealth

income constraint vs. wealth constraint

marginal rate of substitution

price ratio

relative scarcity

utility (a.k.a welfare)

NOTE 3

Welfare and Exchange

What can be said about the welfare of society as a whole if, as economic theory assumes, each of us looks out for number 1? Go to school or not, get married or not, get a job or live in your parents' basement for the rest of your life. All of these me-first decisions inevitably involve us interacting with each other, affecting each other, and while not always apparent, take in place in countless different markets.

What then does it mean for a society to be as well off as can be? One gauge of best welfare is a Pareto improvement. An allocation is Pareto optimal if it is impossible to make someone better off without making others worse off. You can also say that an equilibrium is Pareto efficient or allocationally efficient. The table shows some other ways of identifying a Pareto optimal allocation. How do we get there?

An allocation is Pareto optimal if there is no way to...

make everyone better off.

to make one person better off without making at least one other person worse off

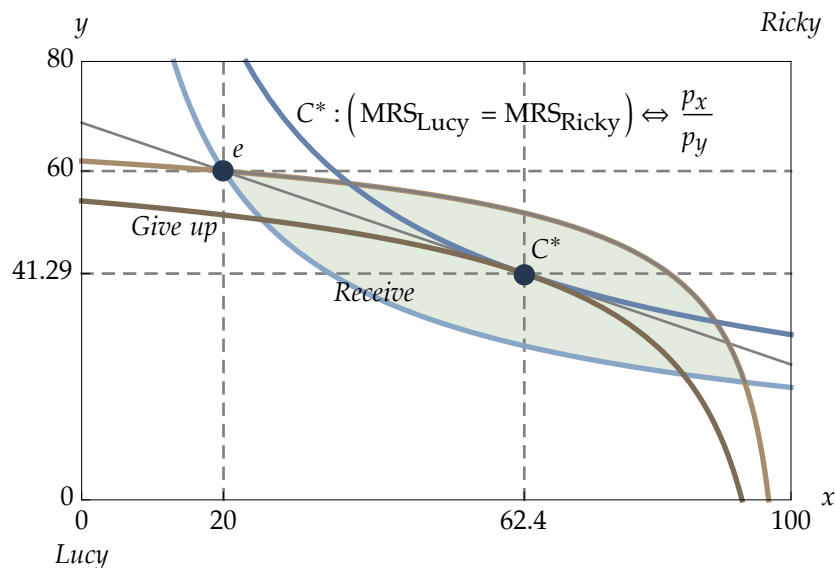
gain from trade without someone else losing from trade

Here is how to draw a Pareto optimum. Pretend that the economy is inhabited by just two people, Lucy and Ricky, and both have log utility functions of the form

$$u = a \ln(x) + (1 - a) \ln(y)$$

where $a = 0.4$ for Lucy and 0.6 for Ricky. Lucy is endowed with $20 x$ and $60 y$. Ricky is endowed with $80 x$ and $20 y$. This means that the aggregate endowment or supply of x and y in the economy is $100 x$ and $80 y$. The economy can be represented as a rectangle, called an Edgeworth exchange box, with the width equal to the aggregate endowment of x and the height equal to y . Lucy's origin is the lower-left corner and Ricky's the upper-right. Lucy has some level of utility associated with her endowment, with an indifference curve passing through it (the light blue one), and so must Ricky with his (the light brown one). Point e represents their endowments, read from Lucy's perspective. Lucy increases her utility if she can obtain a bundle that lies above her initial indifference curve. For Ricky, down is up because his origin is the top-right

corner of the Edgeworth box. He increases his utility by moving down towards Lucy's origin. The green lens-shaped area formed by the intersection of their initial indifference curves, in-



cluding the boundary, is the area of Pareto improvement. A reallocation to any point in the green area from e would make both better off, or at least one better off without hurting the other. Neither would object to such a reallocation.

	Lucy	Ricky
Taste parameter a	0.4	0.6
Endowment e	{20,60}	{80,20}
Optimum C^*	{62.4,41.2941}	{37.6,38.7059}
$MRS(e)$	2	0.107143
$MRS(C^*) = \text{price ratio}$	0.441176	0.441176
Utility(e)	3.6549	3.41162
Utility(C^*)	3.88586	3.6473

The allocation C^* is the Pareto optimum. At this point neither Lucy nor Ricky can be made better off without making the other worse off. If Lucy were to move to a higher indifference curve, Ricky would have to pull back to a lower one. The same goes for Ricky. To get to C^* , Lucy bought some x from Ricky in return for some y . It makes sense that Lucy would buy x from Ricky since her marginal rate of substitution at e is bigger than his (can you see that?). At the Pareto optimum, Lucy and Ricky place the same value on the next unit of x in terms of y , so there are no mutually advantageous trades left. Mathematically, their indifference curves are

tangent to one another (the grey line passing through e and C^*), which means that their marginal rates of substitution are equal. It is the condition that defines the equilibrium.

Their common marginal rate of substitution at the Pareto optimum must also be the equilibrium price ratio. If money were to take the place of barter, the dollar prices of x and y , whatever they might be, would have to be in the ratio 0.441176 (see table or, again, the tangent line passing through e and C^* in the graph). Notice that the price ratio is not given but is, in fact, determined jointly with the optimal allocation C^* , and is called a *general equilibrium*—the values of all important variables, such as consumption and the price ratio, are determined jointly. You are not expected to be able to compute the Pareto optimum in this course (take my FINA 385 for punishment), but you should be able to look at the figure or table above and interpret them.

When a market is allocatively efficient, the allocation of goods and services and the price ratio together reflect an aggregation of everyone's tastes. If production had also been included in the economy, the allocation and price ratio would also be driven by producers, which could be Lucy and Ricky, each choosing the mix of x and y to manufacture so as to maximize profits.⁶ You can think of an allocatively efficient market as one in which each person derives the most benefit—utility for consumers and profit, which in turn becomes utility too, for producers—at the same time that everyone else is doing the same. It is an equilibrium in the sense that there is no incentive for anyone to change: Ricky is happiest with his consumption bundle (37.6, 38.7) at C^* , and Lucy hers (62.4, 41.3), in light of the price ratio, 0.441176, and the price ratio is literally what it is because Lucy and Ricky are holding bundles for which their marginal rates of substitution are equal. I know what you are thinking: this whole general equilibrium looks like a tautology. Well, it is, but that is a problem of the economic paradigm that we'll have to leave for another day.

Arrow's Impossibility Theorem meets The First Welfare Theorem of Economics

Goodness, that's a long heading. It looks like trade is one way for an economy to reach a Pareto optimum, although it isn't clear what kind of trade, or whether the optimum is unique or exists at all. But are there other ways? Could a government or benevolent dictator come up with a formula to reallocate the endowments of its citizens to a Pareto optimum—doing what is best for the people? Does a democratic system of majority vote do the trick? Lotteries? Rankings? No such luck. Professor Kenneth Arrow figured out that there is no social decision rule that can guarantee a Pareto optimum (that's why it's called Arrow's Impossibility Theorem). The reason is pretty simple: governments, rulers, or anyone on the outside, cannot know people's tastes, so

⁶ The economy described here is sometimes called a *pure exchange economy* because there is trade but no production.

it would be a fluke if any social decision rule they created could pick off a Pareto optimum. Makes you wonder how the millions of shareholders in modern corporations agree to anything!

Trade saves the day though. But only one kind of trade or market guarantees a Pareto optimum, and that is pure competition. This is not competition as in destroy-my-competitors-no-matter-what kind of competition but the kind we read about in textbooks where everyone is taken to behave as if they were price-takers. This means that when Lucy is thinking about how many y 's she'd be willing to pay Ricky for so many x 's, she does not think that her demand for x could somehow influence its price in the bigger scheme of things. The same goes for Ricky. Of course, the demand and supply of the millions of people like Lucy and Ricky, taken all together, do influence prices and quantities. But individually, each person must act as if they were inconsequential, tiny fish in a big sea.

If Lucy and Ricky set out to maximize their utility and do so by trading as price takers, they will arrive at a competitive equilibrium. The *First Welfare Theorem of Economics* tells us (we won't prove it) that competitive equilibria are Pareto optimal. So, if you know that a market is in equilibrium in a supply-equals-demand sense *and* that it is a purely competitive—something harder to do—then you also can conclude that it is also Pareto optimal. You don't need to know anything about people's individual tastes or motives. It underlies Adam Smith's idea of the invisible hand at work: everyone acting in their own self interest makes for the best outcome for society as a whole. The implication of The First Welfare Theorem of Economics is profound because it suggests that free markets and economic freedom in general—private property and protection of property rights, reasonably low taxes, rights to trade and do business—should be the default path for improving societal welfare. What does that say about centralized economies?

Review

pure exchange economy

general equilibrium

Pareto improvement and optimum

Arrow's Impossibility Theorem

The First Welfare Theorem of Economics

price-taker

pure competition

NOTE 4

Information

In the previous note, one x is valued at 0.441176 y 's in (general) equilibrium. Everything about the economy is packed into that price ratio. When the price ratio changes, it signals change, although what change is not always clear: technology, weather, demography, tastes, health. Take a step back from the price ratio to the information about those things that might cause the price ratio to change. Firms have to buy or rent their factors of production. For this they need financing from investors in the form of equity and loans. While the relative prices of goods and services are a signal to firms to allocate their factors of production profitably, existing and potential investors also rely on signals from the prices of financial securities to help them know whether this is being done; otherwise, investment capital wouldn't necessarily flow to the firms earning the highest return for a given level of risk. If we are greedy, wouldn't we be monitoring all kinds of information, trying to figure out its effect on the prices of goods and services, and in turn, the effect on the prices of stocks and bonds of the companies that put out those goods and services? You bet. Good news about the future cash flow of a firm should translate into a rise in the price of its securities and bad news a fall because greedy investors act on the news. How big a rise or fall? The *present value* or *discounted value* of the change in future cash flow, such as dividends, that investors expect to receive.⁷ The level of and change in security prices presumably acts to discipline firms to produce efficiently.

How quickly should security prices adjust to new information? If enough of us jump up to buy or sell a financial security when there is news, its price should change quickly, so quickly in fact, that a stock that is underpriced at the moment of good news is never underpriced long enough for anyone to profit from buying it. It is never really underpriced to begin with because new information would be incorporated into the price instantly! (That's economics channeling physics.) That blindingly fast change in price should equal the present value of the expected changes in future cash flow. Some investment companies locate their operations as close as possible to the stock exchange so that they can minimize the length of wire runs from their computers to the stock exchange's servers. Enough said. The price is always right, in theory at least.

⁷ You'll get to do some present value calculations and other time value of money calculations in the quizzes.

A capital market is said to be *informationally efficient* if the price of its securities reflects all relevant information about future cash flow. This is also called the *efficient markets hypothesis*. Nowhere is there greater attention paid to the relationship between prices and information than in capital markets, where investors, analysts, alchemists, and clairvoyants try to decipher the effect of all sorts of information on the value of securities. (I remember watching a TV show on PBS about corporations and capitalism in America, and a commentator said something to the effect that if it is reported that a woman dies of breast cancer in America, the price of biotech stocks will go up. I thought it was Noam Chomsky who made the remark, but when I emailed Professor Chomsky about it, he replied that he did not recall.) Informational efficiency is to capital markets what allocational efficiency is to the markets for goods, services, and factors of production. In an informationally efficient capital market, you can't make yourself wealthier trading securities in the same way that Lucy and Ricky do not make themselves wealthier trading x and y in a pure exchange economy. To get rich trading financial securities—and that does not mean hitting the jackpot now and again because anyone can get lucky—you'd have to have valuable information that most others do not have or have a better understanding of the available information than everyone else.

Rational expectations

The assumption that people are soaking up and responding to information continuously also includes each of us taking into account how we think everyone else is motivated and responds to information. A long time ago, an insightful little story about this appeared on the back cover of *Journal of Political Economy*. A group of young people who, while taking a walk on a well-traveled country road, come across a peach tree laden with ripe fruit. When one of them suggests that they stop to pick some peaches, another quickly responds that they shouldn't bother because if the peaches were any good they would've already been picked. That's *rational expectations*. I like peaches, and all the better if they are free. But wait, don't most other people like free peaches? And wouldn't they pick them if they were free? Then why are those beautiful peaches still on the tree?

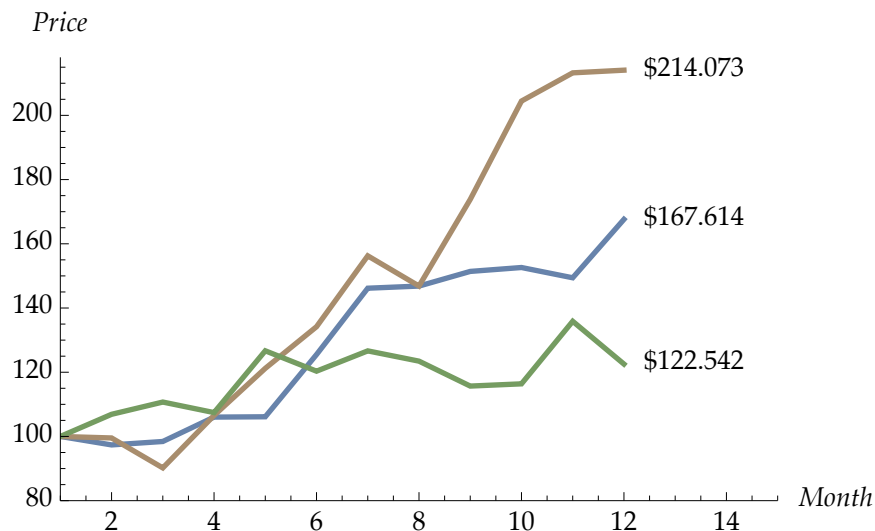
In utility theory, tastes are modelled by utility functions because tastes are unobservable. In the same way, individual decision processes are unobservable, so it is assumed that people behave *as if* they take into account every available scrap of information in making any decision, and that they are frighteningly logical about it. Rational expectations takes this individual logic one step further in assuming that we consider the effect of our choices on everyone else and everyone else's choices on us. The peaches never get picked. Assuming rational expectations to avoid having to include a specification of individual decision processes in economic models implies that people are assumed to behave as if they know the true model, that is, the economist's model of the economy (how convenient for economists), and that they behave as if everyone

else knows the true model of the economy (I know that you know that I know...). An implication of rational expectations is that monetary policy is ineffective. A government that pursues an inflationary monetary policy by printing money will be undone by unions who anticipate the inflation and, in turn, demand wage increases to compensate for the erosion of real wages.

Take a random walk

An implication of the efficient markets hypothesis is that you can't use the past to predict the future: markets have no memory. That's because news isn't really news unless it is, itself, a surprise. If news cannot be predicted, then our reaction to it must result in price changes that cannot be predicted. If you used past changes in price as your information or news, and identified a pattern in those prices that could predict the future ups and downs of prices, then there are probably other clever people who have found the same pattern.⁸ Everyone would exploit the pattern, buying when it signals a rise in price and selling when it signals a decline. The buying and selling would cause the pattern to self-destruct, taking down trading gains along with it. That's why it is sometimes said that price changes in an informationally efficient market should resemble a *random walk*, and markets where the past cannot be used to predict future is said to be *weak-form efficient*.

Here are three stocks, each of which happened to be worth \$100 at the beginning of month 1. Twelve months later there are big differences in their prices. Some people will look at a price chart like this and see patterns or trends. Some might say that the price of the brown stock has



reached its peak and will soon start to fall. Others might say that the blue stock is trending upwards. What would they think about the green stock? Those who didn't flunk their intro stats

⁸ Don't think that you are the only smart cookie out there.

course might compute the correlation of the month-to-month price changes of each stock, despite the fact that they have only 11 price changes to work with (or they might go further back in time to increase their sample size). They might find that one or two of them display *positive serial correlation*: a rise in price is more likely to be followed by another rise than a fall. They might find that one of the stock's price changes is not serially correlated or may be negatively serially correlated (what does that mean?).

If it were now the end of month 12, would you base your decision to buy one of the three stocks on a pattern that you see? It wouldn't matter because there are no patterns. They all follow random walks, where each month there is an equal chance that price rises by 20 per cent or falls by 10 per cent. The monthly changes are completely unpredictable. The only thing that can be predicted is that if you buy any one of them and hold it long enough, your average return will be five per cent over the long run.⁹

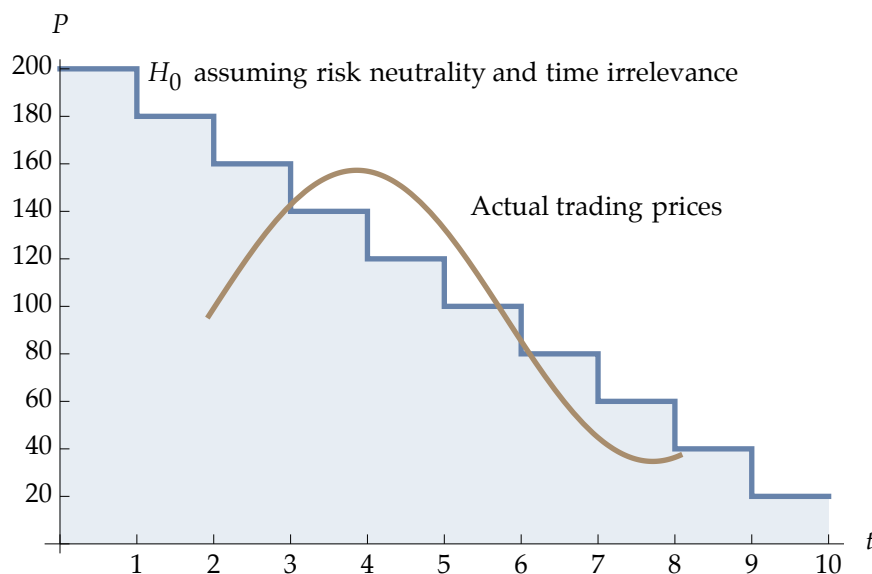
Bubbles and crashes

If people have rational expectations and information about a company's future dividends is widely available, the price of the company's stock should equal the present value of its expected (future) dividends. Professor Vernon Smith ran a series of clever experiments that demonstrated that a stock market bubble and crash could be created in a laboratory under controlled conditions, and that the rational expectations hypothesis is violated (for shame).

Subjects took part in an oral double auction market, trading a stock with a ten-period life. At the end of each period, a coin toss determined whether the stock paid a dividend of \$15 or \$25, and all participants knew this. The rational expectations price each period, that is, the equilibrium price implied by the efficient markets hypothesis, is simply the sum of expected dividends. Since the expected dividend each period is \$20, the predicted price for period 1 is \$200 (\$20 × 10 periods), \$180 for period 2 (\$20 × 9 periods), and so on down the blue steps in the figure. In real life, the predicted equilibrium prices would be less than the sum of expected dividends because the market would discount them at some positive interest rate to compensate for the time value of money and the fact that the dividends are risky. But in an experimental market it is reasonable to take the interest rate as zero because each period is short, say, 20 minutes, as opposed to months, quarters, or years, and subjects face little risk since they do not have to put up any of their own money to trade. Subjects are effectively assumed to treat the time between dividend payments as irrelevant and ignore risk. Someone who does not care about risk is said to be *risk neutral* as opposed to *risk averse*, which is natural for real life situations. Risk neutrality implies that subjects care only about the expected dividend and not how risky it is: they see no difference between a fifty-fifty chance of receiving \$15 or \$25 or a fifty-fifty chance of \$5 or \$35.

⁹ Since we are looking ahead, we can also say that the *expected* return is five per cent, a probability-weighted average, $0.5 \times 20\% + 0.5 \times (-10\%) = 5\%$. Probability-weighted averages are called *expected values*.

The actual average trading prices (brown line) sketch a classic stock market bubble and crash,



where at some point the price might be so high that even if all of the remaining dividends were \$25, it would not be enough to recoup the purchase price!

Professor Robert Shiller’s surveys of investors who experienced the stock market crash of 1987 and the boom and bust real estate markets of the mid to late 1980s are provocative. He finds there is a marked tendency for investors to focus most closely on recent price changes as their primary source of information about the future, resulting in herd behaviour. If everyone is watching price changes in order to guess what everyone else is thinking, who’s looking at economic fundamentals? This is hinted at in Professor Smith’s bubbles and crashes experiment in the comments of subjects who reported that they were aware that the stock was overpriced but bought anyway (or did not sell) because they were afraid of missing out on further possible price rises.

Anyone taken as an individual is tolerably sensible and reasonable — as a member of a crowd, he at once becomes a blockhead.

—Friedrich Von Schiller, as quoted by Bernard Baruch

Is private information reflected in prices?

Suppose you took part in an experiment that had you trading shares of a stock in a continuous, double oral auction market. The shares pay dividends in each of two periods, A and B. The dividends are certain, so you know exactly what you will earn, but yours won’t necessarily be the same as those paid to other investors (types I, II, or III). Information about dividends is also private. You do not know what others will earn, and they do not know what you earn, and everyone knows that.

Investor	Period A	Period B	Total
I	\$60	\$40	\$100
II	\$90	\$50	\$140
III	\$40	\$70	\$110

What does economic theory have to say about the equilibrium price in period A and B? If a price is an equilibrium price, there are no gains from trading at that price, and so there would be no point in trading. Start in period B and work back to period A. The equilibrium price in B must be \$70 because, at a price of \$70, type III investors are indifferent to buying and selling because that is exactly the dividend they will earn. At any price below \$70, type III would gladly buy from I or II and pocket the difference between the purchase price and \$70, and at any price above \$70, III would gladly sell, although neither I or II would buy. Because B is the last period, the equilibrium price is equal to the highest dividend *anyone* might receive, and that happens to be the dividend paid to type III investors. Now slide back in time to period A. The price in A must be at least \$140 because type II investors can earn \$140 simply by sitting back and collecting dividends. But all investors can trade, and that means that those earning lower dividends can gain by selling to those earning higher dividends, who can of course gain too. The equilibrium price in A must then be \$160 because the most that anyone (type II) can earn in A is \$90 and the most that anyone (type III) can earn in B is \$70. The right to trade means that price reflects the biggest benefit—cash flow in this case—that can be received at every point in time, but it does not matter who receives it. You can also think of the value of the right to trade as being \$20 in this experiment, the difference between the equilibrium price of \$160 and \$140, the biggest total cash flow that any one investor type could earn from dividends alone. That's neat because the right to trade is a legal construct or freedom, and we've just demonstrated that it has value; it results in a Pareto improvement. A right has non-negative value.

It may be hard at first to wrap your head around this result because you can't help but think of real people, such as yourself, trading. How can the price be \$160 in A if no one, not even investor III, knows that the biggest dividend to be paid in period B is \$70? That's where the finessing assumption of rational expectations comes in. Economic theory assumes that, even though dividends are private information, investors behave *as if* they do know each others' dividends (I act like I know even though I really don't know). Without the rational expectations assumption, the equilibrium price cannot be predicted. And with the assumption, the equilibrium price in each period appears instantaneously. The moment the bell rings to start trading in period A, the bid and ask prices must be \$160; the market is in equilibrium from the get-go and no trade occurs. The same goes for B. A market that is so informationally efficient that its prices even reflect private information is said to be strong-form informationally efficient. If real-world markets were strong-form efficient, insiders would not be able to enrich themselves at the ex-

pense of others, and investment advisors and hedge funds would not do better than the rest of us (which, in fact, they generally do not).

That is the theory. For experimental sessions, we wouldn't be quite so demanding. As a subject in the experiment, you are only human after all, and you cannot read the minds of other subjects. We'd expect trade to occur and hypothesize that the average trading price in A is \$160 and the average in B \$70, and that the trading prices should converge to these hypothesized values quickly rather than slowly. It turns out that in actual sessions the price in period B moves to \$70 pretty quickly in the first run, but the period A price hangs around \$140. If the experiment is repeated a number of times, subjects learn through the feedback of repetition that the stock is worth \$70 in B and then incorporate this information into the period A price in subsequent runs, driving up the price to \$160. It takes time to learn just as Bill Murray did in the 1993 movie Groundhog Day. In class, we'll discuss how the addition of a forward market can make the prices in this experimental market move more quickly to their efficient levels.

Review

informational efficiency or efficient markets hypothesis

rational expectations

weak-form informational efficiency

semi-strong form informational efficiency

strong-form informational efficiency


risk aversion and risk neutrality

Keynes beauty contest (look it up and see where it applies to this note)

NOTE 5

Production and Finance

Just as a person's decision to consume a certain mix of good can be modelled by a utility function, a firm's decision to produce a certain mix of goods, meat pies (x) and sausage (y), can be modelled using a function called a *production possibility curve*. For people it is a model of tastes; for firms it is a model of technology. A production possibility curve, like the one in the figure, sets out the most of one good that can be produced for a given amount of the other. If the firm



chooses to produce 364.577 meat pies then its technology permits it to produce up to 890.123 kg of sausage (point b). The shape of a firm's production possibility curve depends on the current state of technology and the quantity and quality (or productivity) of factors of production—land, labour, and capital. Another firm producing the same meat pies and sausage may have a production possibility curve that is flatter or more rounded or pushed out more or less. Production possibility curves can represent the production tradeoffs for a person, a firm, an industry, or a whole economy.

Like indifference curves, production possibility curves slope downward because they represent a tradeoff between two goods. This production tradeoff can be for a person, a firm, an industry, or a whole economy. Increasing production of meat pies to 489.577 means lowering

production of sausage to 790.88 kg.¹⁰ The absolute value of the slope of a production possibility curve is called the *marginal rate of transformation* (MRT). It is an opportunity cost because it tells us how many units of y must be given up to produce one more unit of x , or how many more units of y can be produced if one less unit of x is produced. At point c , the cost of producing one more meat pie on top of the 489.577 already produced is 0.967 kg of sausage, roughly one for one. But unlike indifference curves, production possibility curves are usually drawn strictly concave to the origin. This is so they can represent a diminishing marginal rate of transformation, which means that it becomes increasingly costly to produce each additional unit of x as x production is increased. That makes sense because producing more and more x means shifting factors out of y production and into x production, and at some point the firm will be tapping into factors that are not as well suited for making x as they are for making y . Some of the sausage makers are equally proficient at making meat pies, others less so.

As economics assumes that people are greedy, the people who own and run companies will produce a mix of x and y that maximizes profits. At a price of \$9 for meat pies and \$6 per kg of sausage, the firm will earn about \$7,881 if it produces 240 meat pies and 954 kg of sausage (point a). As a matter of fact, it will earn \$7,881 if it produces anywhere along the grey line,

whose slope (ignoring the minus sign) is $\frac{p_x}{p_y} = \frac{\$9}{\$6} = \frac{3}{2}$, and which passes through a . This is called

an *iso-profit* or *iso-income* line. But b is a better choice than a because profit with this mix is \$8,622.¹¹ Why the increase in profit? At a , producing one more x only costs four-tenths of a y . You earn \$9 for an extra x and \$2.40 for the four-tenths of a y which brought in \$6. As long as the price ratio is greater than the marginal rate of transformation, you can earn more by shifting production out of y and into x . Let's see the math.

$$\frac{p_x}{p_y} > MRT = \left| \frac{\Delta y}{\Delta x} \right|$$

Cross multiply to get

$$p_x |\Delta x| > p_y |\Delta y|$$

which says that the profit or cash flow gained from one more x is greater than the profit foregone from 0.392 fewer y 's.

$$\$9 \times 1 > \$6 \times 0.392$$

If the inequality was reversed, producing more y and less x would increase your profit (start at d , for example). Profit must be maximized when the mix is chosen so that net profit or net cash

¹⁰ Don't you just love Italian sausage on the barbeque? I barbeque all winter.

¹¹ Are we talking profit or revenue?

flow doesn't change at all if production of x is increased or decreased by one unit. That point is P^* , the tangency of the production possibility curve and the iso-profit line, where profit is \$9,372.30, right?

$$P^* : \frac{p_x}{p_y} = MRT$$
$$\therefore p_x |\Delta x| = p_y |\Delta y|$$

The second line of math says that the net change in cash flow is zero. For those who are interested in cranking out the numbers, are just plain curious, or even bored, the production possibility curve in the figure is part of an ellipse with equation

$$1 = \frac{x^2}{g^2} + \frac{y^2}{h^2}, \text{ where } g = 800, h = 1,000$$

Where's the connection to financial markets in all of this? Well, if profits each period are as big as they can be, then the market value of the company—its assets—are as big as they can be, and shareholders will thank you because the value of their shares will be as big as they can be (lots of bigs here). If not, say because the boss choose a production point other than P^* , cash flow would be smaller than it could be and share value lower. If investors have information about the company's technology—its production possibilities—the competence of the boss, and can make comparisons to similar companies, then they'd recognize that the company is not doing as well as it could be. Some will see the low share price as a bargain if they figure that profits can be improved by giving the boss the boot and installing one that will choose P^* . The gains go to the takeover artist. Welcome to the world of corporate finance.

Review

- production possibility curve
- marginal rate of transformation
- iso-profit line
- profit maximization

Something to think about

Why do firms exist? If markets are such good allocators of resources, why do we see so much economic activity within the confine of firms and other formal organizations?


NOTE 6

Government Intervention

Government intervention takes many forms but often involves manipulation of market prices. Taxes reduce incomes, subsidies lower prices, sometimes for select groups. This results in a misallocation of resources because prices will no longer signal the relative value that people place on different goods and services. But government intervention is not a bad thing so long as its costs do not outweigh its benefits (to state the obvious).

Taxing and subsidizing the market for business loans

Suppose that interest paid on business loans was not a tax-deductible expense, and that lenders were not taxed on interest earned from business loans. Suppose also that both borrowers and lenders are taxed at a rate of 50 per cent, and that this rate applies to all of the other usual types of income and expenses; it's just that interest business loans is left out. Point e in the graph is the equilibrium for business loans in this situation. The interest rate is assumed to be



eight per cent, and the quantity (dollars) of loans outstanding is 1,000. But then the government decides to change the tax rules. Business loan interest is now a deductible expense, and the banks will be taxed on that same income. To borrowers, interest expense deductibility is a subsidy because it reduces taxes payable. This type of subsidy is called an *ad-valorem subsidy* or *val-*

ue subsidy because the amount depends on the price (the interest rate). It is also called a *tax shield* in this case. The tax on interest earned by lenders is called an *ad-valorem tax* or *value tax* because, like the subsidy, the amount depends on the price.

How will the change in tax rules affect the market for business loans? The demand for loans will increase because interest deductibility makes the after-tax cost of a loan less than the before-tax cost. Borrowers can now afford to pay a higher before-tax rate of interest for business loans. You can find a point on the new demand schedule by asking how high the before-tax rate of interest—eight per cent before the change—can rise without costing borrowers more than eight per cent after tax. The answer is obviously 16 per cent because a borrower in a 50 per cent tax bracket who pays 16 cents in interest on every dollar borrowed gets a tax shield of eight cents, so they are really only paying eight cents in interest. Let x be the before-tax rate of interest and T_d (d for demand) the borrowers' tax rate, then

$$\begin{aligned} .08 &= x(1 - T_d) \\ x &= \frac{.08}{1 - .50} = .16 \end{aligned}$$

That makes sense because the subsidy or tax shield is eight cents on the dollar when the interest rate is 16 per cent.

$$\text{Tax shield} = T_d \times (r \times q) = .50 \times (.16 \times \$1) = \$0.08$$

The new demand curve must pass through (\$1,000, .16) no matter what the rest of the curve looks like (it doesn't have to be a straight line like the one shown). Substitute different before-tax interest rates in the numerator in the equation above, and you will see why the demand curve shifts up more (or forward if you like) at higher interest rates and less at lower ones. That's the nature of an ad-valorem subsidy.

For lenders, having their interest income taxed is an increase in the cost of lending because, after all, taxes are a cost of doing business. A bank facing a 50 per cent tax rate will now get to keep only \$50 of every \$100 in interest earned, so the supply of loans will decrease. You can find one point on the new supply curve by asking how high the interest rate would have to rise for lenders to continue earning eight per cent after tax. The answer is 16 per cent again because a lender in a 50 per cent tax bracket will have to pay eight cents in taxes to the government for every 16 cents of interest earned. If T_s is the lenders' tax rate (s for supply),


$$\begin{aligned} .08 &= x(1 - T_s) \\ x &= \frac{.08}{1 - .50} = .16 \end{aligned}$$

The new supply curve, like the new demand curve, passes through (\$1,000, 0.16), so the new equilibrium must be (\$1,000, .16). The market for business loans did not grow or shrink (it's still \$1,000) because the after-tax return to lenders and the after-tax cost to borrowers did not

change (still eight per cent). The change in tax policy is *neutral* because it does not affect consumers (borrowers) and producers (lenders). And if it is neutral for consumers and producers, it must be neutral for the government, that third player lurking in the background. The new tax revenue of \$80 is exactly offset by the new tax shield of \$80. The government's cash flow has not changed.

It should be obvious from the graph that the reason that the new tax rule is neutral is because borrowers and lenders face the same tax rate. When the two tax rates are the same, the decrease in supply is exactly offset by the increase in demand, and the equilibrium quantity of loans does not change. For the policy to be neutral, the two tax rates do not have to be 50 per cent; they just have to be the same. It follows that the policy will not be neutral if the tax rates of borrowers and lenders differ.

If lenders face a higher tax rate than borrowers, the supply of loans will decrease “more” than the demand for loans will increase, and the loan market will shrink. If lenders face a lower tax rate than borrowers, the opposite happens. Here is an example of the first case. Borrowers



are taxed at 35 per cent and lenders at 55 per cent. For the equilibrium quantity of loans to stay put at \$1,000, lenders require a interest rate of 17.78 per cent (point u), but borrowers will not borrow that amount if the new interest rate is anything higher than 12.3 per cent (point l). The new policy will cause supply to decrease “more” than demand increases, causing the new equilibrium, e' , to lie to the left of the old one. I calculated the new equilibrium quantity of loans to be \$818 and the new equilibrium interest rate to be 14.5%. But I could only do that knowing the equations for the two curves to find their intersection. In real life, we don't know what the supply and demand curves for loans looks like, and estimating them is a crap shoot. What we do know, however, is points l and u , and it is plain to see that the new interest rate must lie somewhere between the rates, 12.3 per cent and 17.78 per cent. The finance minister will not be impressed when you tell him that your best forecast of the effect of a policy change puts the inter-

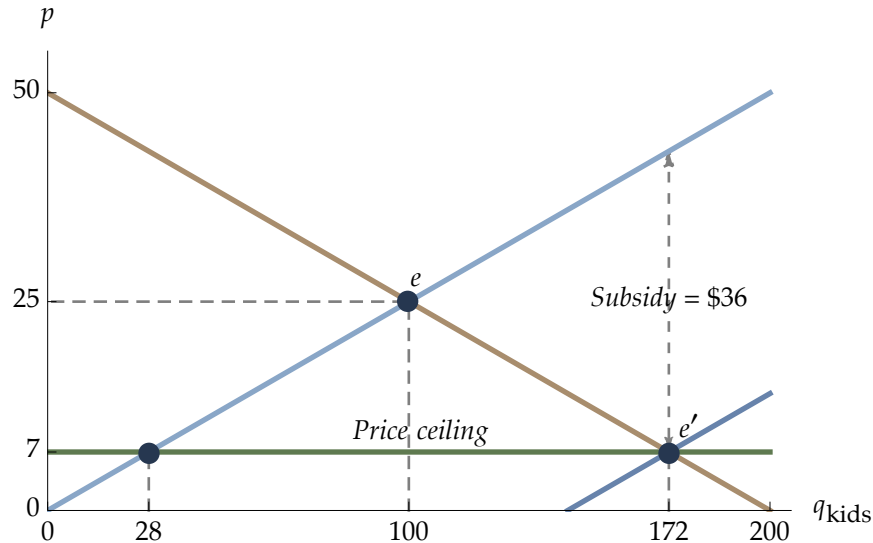
est rate in a five per cent range, but at least you're being honest about it. The minister will tell the public whatever he or she wants.

Because the loan market has shrunk, loans must have lost some of their appeal to borrowers and lenders compared to other forms of finance and investment. Look at points d and s , which lie directly below the new equilibrium. Those borrowers who continue to borrow after the policy change, now pay more for their loans after tax, 9.45 per cent (point d). And lenders who continue to lend earn less, 6.55 per cent (point s) down from eight. If borrowers pay more and lenders earn less, did the government generate a positive increase in its cash flow? Yes, it did, an increase of \$23.8017 from \$65.4548 in new tax revenue less \$41.653 in new tax shields. I'll leave it to you to do the arithmetic to confirm those numbers and identify the area in the graph that represents the \$23.8017.

Switch the tax rates of borrowers and lenders and rework the analysis. Explain it to a loved one. They'll thank you for it.

Seven-dollar a day daycare

When my daughter attended Concordia University's Garderie Les Petit Prof in the early 1990s, I think my wife and I paid about \$23 a day. Most registered daycares in Montreal at the



time were charging between \$20 and \$25. The picture shows a hypothetical equilibrium with a market-clearing price of \$25 and 100 children attending the various registered daycares in Montreal. Then in 1997 the Quebec government set a *price ceiling* of five dollars a day, to help parents balance work and childrearing and to make daycare more accessible to those with low incomes. The idea was that having access to daycare would allow single parents and stay-at-home parents of two-parent families to find jobs and better provide for their families. More

people working would boost the economy through increased spending. The ceiling was raised to seven dollars a day in 2004.

No daycare can survive on seven dollars a day per child, and if fixing the price is all that the government did, there would be a severe shortage of spots. The figure shows that at seven dollars, there would be demand for 172 spots but only 28 would be offered. If there was some way to accurately predict that 172 spots would be demanded at seven dollars a day, the government could give daycares a subsidy that would guarantee that everyone who wanted a spot would get it. A subsidy of \$36 per child per day does the trick in the example. This type of subsidy is called a *quantity subsidy* because the amount is fixed (\$36 per child per day). A smaller quantity subsidy would still leave a shortage, but the government may be satisfied as long as the number of spots increases above the previous level of 100.

One consequence of fixing the price of daycare is that it discourages variety. Some daycares are more costly to run, and so charge higher fees, because they offer better or more varied services. I don't know if the variety in daycare services runs from no-frills to Rolls Royce, but I am sure there are noticeable differences among them. The ones that are more costly to run and that fall under the government's jurisdiction had to cut costs by cutting back on services to be viable. Fewer field trips, that sort of thing. Daycares look more the alike, especially when it comes to "ratios," the number of children supervised by each caregiver. Daycares now maximize their revenue by taking in as many children allowable by law. As a matter of fact, the government requires it (so I'm told, but I have to fact-check this). Previously a selling point of some daycares was their low ratios; your child would receive more attention. Is there a benefit to this homogenization? It is estimated that in 2011, 215,000 kids were being cared for at an expected subsidy of \$2.215 billion or \$10,302 per child per year or \$28 per child per day, and that 70,000 more mothers were able to hold jobs, translating into an increase of 1.7 per cent in Quebec's GDP.¹²

Review

ad-valorem or value taxes and subsidies versus quantity taxes and subsidies
tax shield, before-tax return, after-tax return

Something to think about

Is the freeze on university tuition in Quebec like the daycare price ceiling? Are the consequences similar?

¹² Pierre Fortin, Luc Godbout, and Suzie St. Cerny, 2012, "Impact of Quebec's Universal Low-Fee Child-care Program on Female Labour Force Participation, Domestic Income, and Government Budgets," Working Paper, University of Sherbrooke.

NOTE 7

Market Failure

A market fails when price does not reflect true value in the sense that a price ratio formed with that price is not equal to peoples' marginal rates of substitution. In the extreme, the market may collapse or not exist to begin with because either there are no consumers willing buy at the "wrong" price or no sellers willing to sell. To understand why a market might fail, we need to look back at the First Welfare Theorem: A competitive equilibrium is Pareto optimal. People (and firms) behaving as price takers is a sufficient condition for a Pareto optimum. But there are a number of characteristics or defects that can cause an otherwise competitive market to be, well, less than competitive, scuttling the optimum.

What are these characteristics? One is economies of scale. If a firm can lower its costs without end simply by getting bigger, it will come to dominate a market and wield monopoly power. Too much market power is seen as a bad thing for this very reason, and that is why we have laws that protect competition. The second concerns who knows what. In real life, we do not all have access to the same information; there is *information asymmetry*. Sellers usually know more about the quality of their products and services than buyers do; people buying insurance know more about the risks they face than insurers do. What does this unequal distribution of information mean for the price of products and services? The third characteristic is the unavoidable by-products of our consumption and production. These are called *externalities*. Smoke a cigarette (that's consumption), and the poor fella next to you suffers, but it's not like he can ask you to compensate him for the hit to his utility. Well, he can ask, but he probably won't get very far. He may not have a right to clean air, and even if he did, how would he enforce the right? The question of compensation, of course, also applies to our primary production, not just inadvertent by-products. We want to be paid for what we produce. But what if the thing that you produce can be freely consumed by anyone who wants to—a book, a song, computer software—and you have no easy way to make them pay you or to stop consuming if they won't. Goods that are seemingly free to those who would consume them are called *public goods*. Public goods include infrastructure, such as roads, intellectual property, such as books, music, computer software, and scientific research, which may be produced publicly or privately. The fifth cause of market failure is extreme inequality in the distribution of endowments—income or wealth. It is possible for a Pareto optimum in the Edgeworth box to be one in which Lucy is

filthy rich, and Ricky is so poor that he can barely afford essentials like housing. No one said that a Pareto optimum has to be fair.

This note will take a brief look at three of the five causes of market failure: informational asymmetries, externalities, and public goods. Many examples will be discussed in class, but you can, of course, find examples for yourself. When markets are prone to failure, government intervention may be justified.

The lemons problem

Professor George Akerlof asked, “Why are all used cars lemons?” *Information asymmetry* was the answer. Sellers of used cars know more about the quality of what they are selling than buyers do. If buyers cannot easily tell the difference between a good used car and an otherwise identical lemon, the price buyers will offer will be somewhere between the unknown true values of the two. Anticipating a single price for all used cars, rather than a high price for good ones and a low price for lemons, sellers don’t bother putting good used cars up for sale. Buyers, then, anticipating that good used cars will be nowhere to be found, offer only the lower, lemon price. The middle price never materializes because rational expectations is at work here. Where there could have been a market for good used cars selling at the higher price and another market for lemons, selling at the lower price, there is only a market for lemons. Good used cars suffer *adverse selection*: the bad has driven out the good. The only way to establish a market for good used cars is for the owners to help potential buyers distinguish between their good used cars and lemons. That takes money and effort, both of which fall on sellers of good used cars.

Now, if you replace good used cars with bright students and lemons with, say, mediocre students, both looking for jobs, how would you retell the story in a world of grade inflation? Enjoy.

Moral hazard

You lock your garage every night so that your bicycle doesn’t get stolen. Would you be as cautious if you had insurance that covers theft? Economic theory says no. Taking care of your things requires the kind of effort you’d rather avoid. And why bother anyway, the bike is covered? If the insurance premium that you paid reflects the risk of your bike getting stolen before you decided to slack off, it is now too low. By being less diligent after being insured, you have tilted the deal in your favour, and ripped off the insurance company. That is *moral hazard*. But have you really ripped off the insurance company? Insurers are not dumb. They anticipate that you, and everyone else, may be less careful after buying insurance, so they jack up everyone’s premium in anticipation of moral hazard. So even if you are not *homo economicus* and would never dream of leaving your garage unlocked, you will pay for the possibility that you might.

That's rational expectations at work again in the face of asymmetric information. Good risks pay higher premiums than they should because insurance companies do not know as much about your risks as you do, and how you might purposely change those risks once you are insured. Does that mean that good risks are adversely selected? Yes, in that they are charged too much. But is the adverse selection so severe that many people decide to go uninsured, leaving a market of lemon risks? No, because the premium for a market of nothing but lemon risks might be astronomically high. There's a primitive and fairly effective fix for moral hazard, at least when it comes to insurance: *deductibles*. If your \$800 bike gets stolen, you'll only be covered for \$600. I wonder if something like a deductible could have worked in financial markets to prevent the mortgage brokers, bankers, and real estate agents from engaging in moral hazard and causing the sub-prime loan crisis.

Externalities

Pollution is a by-product of manufacturing cars, just as it is a by-product of manufacturing most things. Pollution hurts surrounding communities, sometimes even faraway ones. If car companies had the right to pollute by law—a property right—and people affected by the pollution could get together and negotiate with a car company to pollute less, say, by offering some amount of money, then all would be well economically. People wouldn't offer an amount that was greater than the value they place on the reduction in pollution, which presumably would come from a reduction in the number of cars produced or investment in cleaner manufacturing processes. And the car company wouldn't accept an amount that was less than the value of lost production or the cost of a cleaner way of doing things. Mind you, the government could have given the people the right to a minimum level of clean air rather than giving the car company the right to pollute. In that case, the people have the property right, and the car company might offer to pay the community something in return for allowing it to increase production, if it were worthwhile. In either case, a market is effectively putting a price on pollution. Professor Ronald Coase¹³ showed, in what is known as *The Coase Theorem*, that the "socially optimal" level of production of something like cars does not depend on who is allocated the property right for an externality, such as pollution, provided that the two sides can negotiate the outcome and compensate each other to their mutual satisfaction. Make the property rights clear in law, and then let people work it out for themselves.

In real life, property rights are not always clear, in which case pollution does not have a price, and car companies probably produce more cars than is socially optimal because the costs to society of pollution are not a cost of producing cars. But even if property rights are clear, negotiation between the two sides is tricky. One car company, many people in the community. How do the people in the community coordinate their effort to negotiate? If some people in the

¹³ Professor Coase died September, 2013, aged 104.

community choose not to participate, they still enjoy the benefits of a reduction in pollution negotiated by the neighbours. The incentive to *free ride* is strong. You probably come up against free riders every time you have to do a group project for one of your courses.

Public goods

What if you produce something for sale but people can consume it without paying. Your work is treated as if it were free. All sorts of *intellectual property* fits the bill: music, art, books and other writing, computer software, scientific research and inventions. All of these take a great deal of effort or cost a lot to produce but can be consumed easily because one person's consumption does not affect another person's consumption (my listening to a song never prevents you from listening to it), and it is difficult for the producer to exclude anyone from consuming. Without fair compensation, these goods will be underproduced because the incentive is not there, and those that are produced will be over-consumed. Even goods that are not free because they are paid for with our tax dollars have the appearance of being free will be over-consumed. Just look at the congestion on our roads.

Goods that are over-consumed and underproduced because they have the appearance of being free are called *public goods*. Copyright and patents are the imperfect fix for the problem that producers of intellectual property have in extracting fair compensation. Both provide legal protection of ownership in the form of a monopoly for a fixed period. But enforcement of the property right is still a problem. For public infrastructure such as roads, user-pay and pay-per-use are the remedies. Hey, suburbanites, ready for a toll on the next Champlain bridge?

Review

- information asymmetry
- adverse selection
- moral hazard
- deductible (insurance)
- externality
- free ride
- public good
- The Coase Theorem

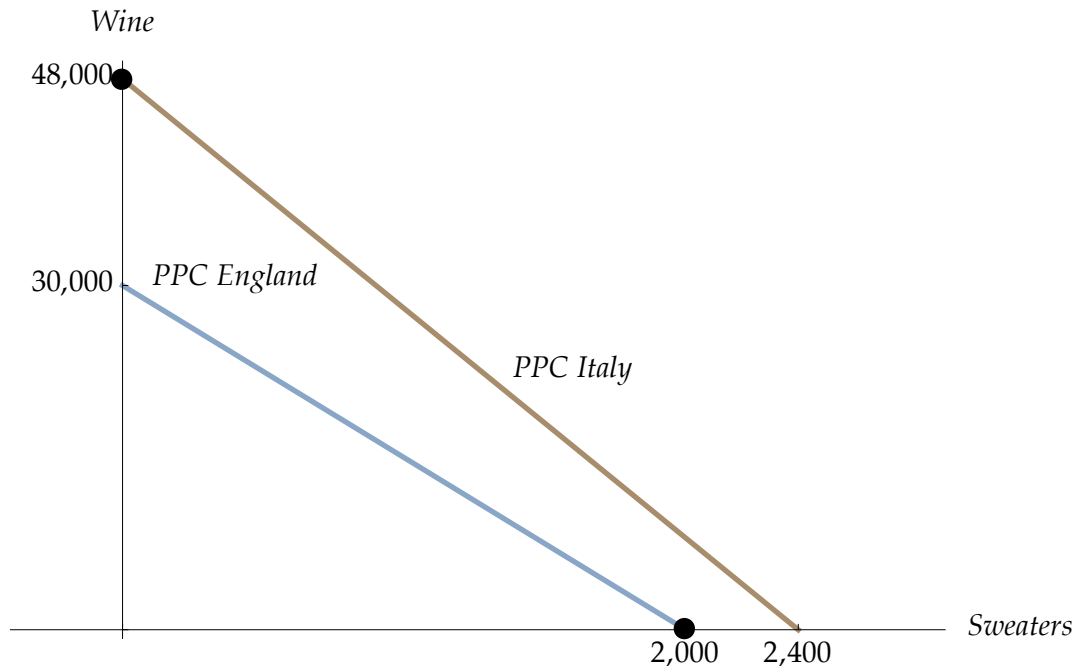
NOTE 8

International Markets

International trade can make people better off when their country manufactures goods whose costs are low compared to that of other countries. It's that simple. As for currencies, translating one to another isn't rocket science, but the difference in interest rates between two countries sends a message about how the market expects an exchange rate to change.

Gains from specialization and trade

Wine and wool sweaters are made in both England and Italy. Suppose that England can produce as much as 30,000 litres of wine or at most 2,000 sweaters in a year. Italy can produce 48,000 litres of wine or 2,400 sweaters. Both countries, of course, produce some mix of the two goods. Suppose also that the number of litres of wine that each country must give up to produce an extra sweater does not depend on the mix. That means that one more sweater always costs 15



litres of wine in England and 20 in Italy. The opportunity cost of a sweater in England is 15 litres

of wine, and the opportunity cost in Italy is 20. Because this about production, you could also say that the marginal rate of transformation of wine for sweaters is 15 in one country and 20 in the other. Constant marginal rates of transformation mean that all of the combinations of sweaters and wine that England and Italy can produce lie on straight lines or production possibility curves.

If sweaters can be made more cheaply in England than Italy, wouldn't there be a gain if England made more of them and Italy fewer? You bet. If England makes an extra 100 sweaters, it will cost 1,500 litres of wine. If Italy makes 100 fewer, it will be able to increase its output of wine by 2,000 litres. The number sweaters in the world is unchanged, but there are 500 more litres of wine to drink. The gain is equal to five litres of wine per sweater, the difference between England and Italy's marginal rates of transformation.

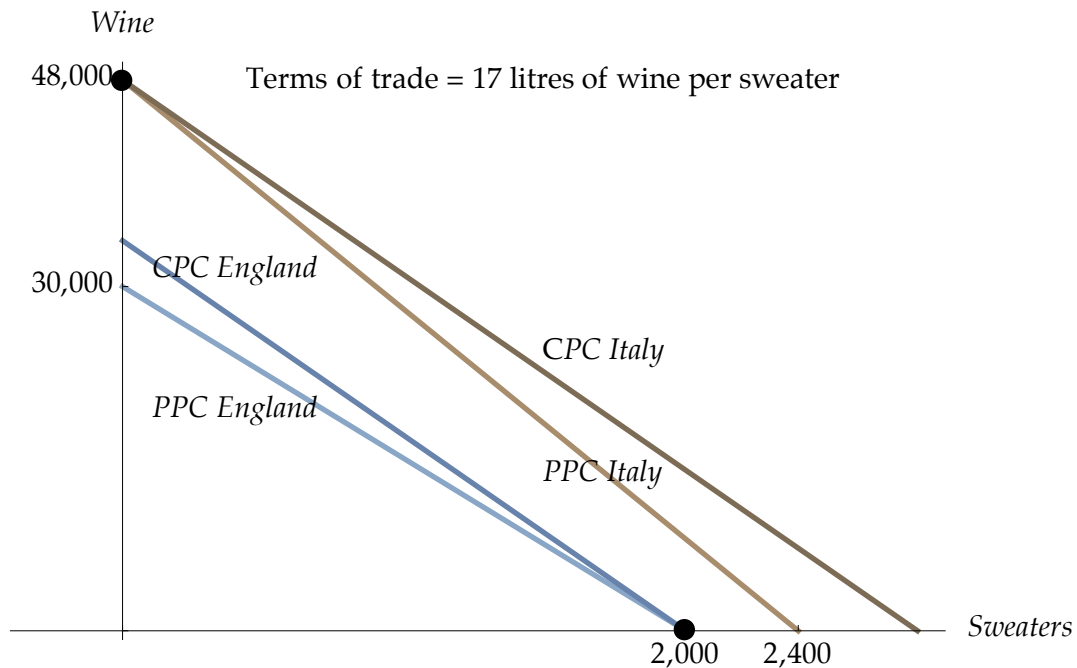
	UK	Italy	Gain
Δ Wine in litres	-1,500	2,000	500 litres
Δ Sweaters	100	-100	0
MRT litres per sweater	15	20	5 litres

If there is a gain of 500 litres of wine by shifting production of 100 sweaters from Italy to England, why not shift 200, 300, or more? The more England specializes in sweaters and Italy in wine, the greater the gain to the world. The biggest gain is had if England stops making wine altogether and specializes in nothing but sweaters, and Italy manufactures nothing but wine. Although, this is true only because we assumed that their marginal rates of transformation are constant. Normally as a country shifts more of its production to the good in which it has a comparative advantage, the cost of doing so will gradually increase, so partial specialization rules the day.

No one is going to enjoy a single drop of that 500 litres of found wine unless the two countries trade. England must export sweaters to Italy in return for vino. (Or Italy must export wine to England in return for sweaters. It doesn't matter how you say it.) Think about it. England is pretty dry sitting on nothing but a big pile of sweaters, and all of Italy is shivering during the winter. Brits gain as long as they can sell sweaters for more than it costs to make them, 15 litres of wine each; and Italians are happy to buy English sweaters if price is less than the 20 litres of wine it costs of make one at home.

If the two countries are friendly, they'll certainly explore the possibility of trading with one another. A world price ratio of sweaters in terms of wine, known as the *terms of trade* in international economics, will be determined by the tastes and incomes of the English and Italians and their production possibilities. Suppose the terms of trade are 17, and that England exports 1,000 sweaters to Italy. England will be ahead 2,000 litres of wine and Italy ahead 3,000 litres.

The gains to England and Italy have to add up to the gain to the world: 5,000 litres. It's not hard to see that the closer the terms of trade are to England's cost (MRT), the smaller the gain to England and the bigger the gain to Italy. And you would never expect to see a price ratio outside of the two MRTs; otherwise, one of the two countries would be losing wine for every sweater



bought or sold.

By specializing its production in the goods in which it has a comparative advantage and then trading internationally, each country can consume along its *consumption possibility curve* (darker blue and darker brown lines) resulting in a Pareto improvement. All free trade agreements are based implicitly on this idea.

Exchange rates

Suppose that the interest rate is three per cent in Canada and five per cent in France. What do you make of the difference? Both rates are for one year and are risk-free. They are rates that you would earn on safe government bonds.

One million Canadian dollars invested for a year here at home will return \$1,030,000 for sure. If instead you invest that \$1 million in France, you'll first have to buy Euros because French government bonds are priced in Euros. Say that the Euro spot price is \$1.3021. One million Canadian dollars will buy 767,990 Euros, rounded to the nearest Euro. Those Euros, invested at the five per cent French risk-free rate, will return €806,390 at the end of the year.

One million dollars is guaranteed to grow to \$1,030,000 or €806,390. What would you do? Invest in Canada or France? It's not clear because if you invest in France, you will have to sell the Euros for Canadian dollars so that you can buy smoked meat in Montreal. The decision would be easy if you knew today what the value of the Euro was going to be a year from now. But you do not, and that uncertainty about the exchange rate makes the French option risky even though the return on French government bonds is risk-free. If the Euro doesn't budge, you'll bring home \$1,050,000 (obviously, right?). If it falls by two cents, \$1,033,872, which is \$3,872 more than you'd earn in Canada. But if it falls by three cents, you'll be left with \$1,025,808, and it would have been better to have kept your money in Canada. Most people don't like risk; they need to be compensated to bear it. And if you're like most people, you might not be willing to invest in France even if you thought there was little chance of the Euro depreciating by three cents.

Ignore risk for a minute—I know, not realistic at all—and pretend that all you care about is earning as much as possible (what do we call that kind of attitude towards risk?). What would the spot price of the Euro have to be one year from now for you to break even? It would have to be \$1.2773 because selling €806,389.68 for \$1.2773 each gives \$1,030,000.

$$\$1,030,000 = \$1.2773 \times \text{€}806,389.68$$

It probably doesn't surprise you that \$1.2773 is about two per cent less than the current Euro spot price of \$1.3021. The interest rate in France is two percentage points higher than in Canada, so the Euro would have to depreciate by about two per cent to cancel France's interest rate advantage. Time for some notation.

$p_0 = \$1.3021$ is the spot price of the Euro today

$p_1 =$ unknown spot price of the Euro one year from now

$r_d = .03$ is Canada's interest rate (d for domestic)

$r_f = .05$ is France's interest rate (f for foreign)

The total return on \$1 million invested in Canada at three per cent is

$$\$1,000,000(1 + r_d) = \$1,000,000(1.03) = \$1,030,000$$

Or buy \$1 million worth of Euros

$$\frac{\$1,000,000}{p_0} = \frac{\$1,000,000}{\$1.3021} = \text{€}767,990.17$$

Invest them in France at five per cent for a total return of

$$\text{€}767,990.17(1 + r_f) = \text{€}767,990.17(1.05) = \text{€}806,389.68$$

and then sell your Euros in a year at the as yet unknown spot price

$$€767,990.17(1+r_f)p_1 = €767,990.17(1.05)p_1 = €806,389.68p_1$$

To calculate what the spot price a year from now would have to be to break even, equate the return in Canada with the return in France

$$\begin{aligned} \$1,000,000(1+r_d) &= \frac{\$1,000,000}{p_0}(1+r_f)p_1 \\ p_1 &= p_0 \frac{1+r_d}{1+r_f} \\ &= \$1.3021 \times \frac{1.03}{1.05} \\ &= \$1.2773 \end{aligned}$$

The calculation is done using the three known current prices—the two interest rates and the spot price of the Euro. If these three prices are equilibrium prices in the supply-equals-demand sense, then every investor must have already decided to keep their money here or invest in France. Equilibrium means that everyone is content keeping their money wherever they have decided to keep it, and because of that, it must also mean that the market *expects* to break-even, $E(p_1) = \$1.2773$. If this weren't true, people would still be moving their money around to exploit some perceived advantage and the three known prices would still be in flux; they would not be equilibrium prices. The equilibrium relationship between the interest rates in two countries and the exchange rate is known as *uncovered interest rate parity*.

$$\$1(1+r_d) = \frac{\$1}{p_0}(1+r_f)E(p_1)$$

The interpretation is neat. If the interest rate in France is higher than the interest rate in Canada, the market must be expecting the Euro to depreciate (\$1.2773 is less than \$1.3021). If France's interest rate is lower than Canada's, a Euro appreciation is expected. The two interest rates provide a market forecast of the future exchange rate! To see the relationship in percentage terms, put the two interest rates on one side of the equal sign and the spot prices on the other, and then subtract one from both sides

$$\begin{aligned} \frac{1+r_d}{1+r_f} - 1 &= \frac{E(p_1)}{p_0} - 1 \\ \frac{1.03}{1.05} - 1 &= -0.0190476 \approx -2\% \end{aligned}$$

The multiplicative difference in the interest rates, $\frac{1+r_d}{1+r_f} - 1$, predicts the expected percentage

change in the exchange rate, $\frac{E(p_1)}{p_0} - 1$, a depreciation of 1.9 per cent. The simple difference,

$r_d - r_f = -2$ per cent is a pretty good approximation when the two interest rates are not too far apart. You may earn two per cent more on French government bonds but on average you will lose that amount when you convert Euros back to Canadian dollars.

Greed keeps the parity in uncovered interest rate parity. You know that the market expects the Euro to depreciate to \$1.2773, but maybe you think you know better. You think the Euro will not fall by as much. Your forecast is \$1.28. Being the gutsy individualist that you are, you *speculate* by borrowing \$1 million in Canada at three per cent and use it to buy French government bonds paying five per cent.¹⁴ Your expected profit is \$2,178.79. Let \hat{p}_1 be your forecast for the Euro, which replaces the expected price implied by the market

$$\begin{aligned} \$1,000,000(1+r_d) &< \$1,000,000(1+r_f)\frac{\hat{p}_1}{p_0} \\ \$1,000,000(1.03) &< \$1,000,000(1.05)\frac{\$1.28}{\$1.3021} \\ \$1,030,000 &< \$1,032,178.7881 \end{aligned}$$

The left-hand side is what you owe on your loan and the right is the expected return on your investment. The difference is your expected profit. Hope things work out according to your forecast. There were other investors who speculate like you because they believe the Euro will end up above \$1,2773. They are on the upper part of the demand curve for Euros. Still others believe that the Euro will fall below \$1.2773, and they speculate by borrowing Euros at five per cent to invest in Canada at three. They are on the lower part of the supply curve for Euros. Together they create the equilibrium comprised of two interest rates and the current spot price.

If there is also a forward market for currencies, *covered interest rate parity* is possible. The equilibrium forward price of the Euro must be equal to the expected future spot price or else it is possible to make a completely riskless profit—just like printing money.

$$\$1(1+r_d) = \frac{\$1}{p_0}(1+r_f)f$$

$$f = E(p_1)$$

To see why the forward price must be equal to \$1.2773, assume that it is temporarily out of whack, say, \$1.25. A forward price of \$1.25 means that you can contract now to buy or sell Euros

¹⁴ By gutsy, I mean risk neutral. Gutsy sounds so much less clinical, don't you think?

in one year at \$1.25 no matter what the spot price turns out to be. I would borrow 767,990 Euros at five per cent to invest in Canada at three per cent and at the same time buy 806,390 Euros forward at \$1.25 because that is how many Euros I will need to pay off my loan. That leaves me a guaranteed profit of \$22,012.90. Follow the inequality. It's all there.

$$\begin{aligned} \$1,000,000(1.03) &> \frac{\$1,000,000}{\$1.3021} \times 1.05 \times \$1.25 \\ \$1,000,000(1.03) &> €767,990 \times 1.05 \times \$1.25 \\ \$1,030,000 &> €806,390 \times \$1.25 \\ \$1,030,000 &> \$1,007,987.0978 \end{aligned}$$

Earning a riskless profit like this without putting up any of your own money is, as you know, *arbitrage*. The two investments, Canada or France, should be perfect substitutes (for risk neutral investors) because they are both one-year riskless investments. Perfect substitutes must have the same price or, saying the same thing, yield the same return. If the forward exchange rate is \$1.25, they do not, and an attractive profit opportunity exists. Of course, everyone knows that the forward rate should be \$1.2773 just by looking at the interest rates and the spot price. So it should not be surprising that these sorts of opportunities are rare. The absence of arbitrage profit is another way of knowing that the markets are in equilibrium.

Review

- gains from trade
- terms of trade (price ratio)
- uncovered interest rate parity
- speculation in currency and interest rate markets
- covered interest rate parity
- arbitrage in currency and interest rate markets

NOTE 9

Intertemporal Markets

A capital market may be as simple as the opportunity to borrow and lend. Borrow to consume more now at the cost of having less later; lend to consume more later by giving up some now. Having the choice of shifting consumption across time leaves us better off. The price of consuming \$1 worth of anything today is the rate of interest. That is why the rate of interest is referred to as the time value of money. The opportunity to borrow can also leave us better off by making us wealthier when the money raised is used to finance real production that earns a higher return than the cost of borrowing.

Bonds

A bond is a loan to a company. A financial institution, often an investment bank, lends money to a company, which in turn promises to make periodic interest payments, usually semi-annually, and a lump sum repayment of the principal at maturity. The investment bank chops up the loan into many smaller loans, which it offers to sell to its clients, so that they can collect

Par or face value	1,000
Coupon rate (c)	0.05
Coupon (C)	\$50
Interest rate (r)	6.5%
Maturity (T years)	25
Compounding (m times a year)	2
Present value of coupons	\$613.79
Present value of par	\$202.07
Bond value (P)	\$815.86

interest payments and a proportional lump-sum payment at maturity. A \$50 million loan to a company might be distributed by the investment bank to its clients as 50,000 bonds, each a loan of \$1,000 (the *par* value or *face value* of the bond), with interest payments, known as *coupon pay-*

ments, computed on the face value. Some bonds pay only interest and others only a lump sum at maturity. Bonds are traded. Retail investors like you and me can buy bonds in the secondary market. Bonds are risky. The price of a bond at any time depends on the market's expectation that the issuing company can make the interest payments as they come due and the repay of the face value at maturity. Changes in the price of a bond caused by changes in a company's prospects is called *default risk*. The price of a bond also changes when the interest rate or the yield on bonds of similar risk changes. This is because alternative investments may have become more or less attractive. Changes in the price of a bond caused by changes in interest rates is called *interest rate risk*.

The bond described in the table matures in 25 years and will pay interest at a rate of five per cent on a par value of \$1,000. That's 50 payments of \$25 because the coupons are paid every semi-annually. At the end of 25 years, a lump sum payment of \$1,000 is made in addition to the last coupon payment. The price of the bond today is \$815.86, which is the sum of the present value of the coupons, \$613.79, and the present value of the par, \$202.07, both cash flows having been discounted at 6.5 per cent. The coupons are an annuity, and that's how they are treated in the first term of the bond valuation equation. The face value is a straight forward discounting of a single amount. Everything rides on that 6.5 per cent discount rate. It is an opportunity cost, which in bond circles is called the required *yield* or *yield-to-maturity*. Where does it come from? You can't look it up. As an opportunity cost, it must be the rate that could be earned on an investment of the same risk, say, the bonds of another company in a similar line of business. But how was the yield-to-maturity of the bonds of the other company determined? Going around in a circle.

$$P = \frac{C}{m} \left(\frac{1 - \left(1 + \frac{r}{m}\right)^{-mT}}{\frac{r}{m}} \right) + \frac{F}{\left(1 + \frac{r}{m}\right)^{mT}} = \frac{\$50}{2} \left(\frac{1 - \left(1 + \frac{.065}{2}\right)^{-2 \times 25}}{\frac{.065}{2}} \right) + \frac{\$1,000}{\left(1 + \frac{.065}{2}\right)^{2 \times 25}}$$

$$= \$613.79 + \$202.07 = \$815.86$$

In fact, 6.5 per cent is the bond's price just as \$815.86 is the bond's price. There's a dollar price on the left-hand side of the equation and a price in per cent on the right, which means that the equation is saying that \$815.86 today is worth the same as a series of \$25 dollar payments and a lump-sum payment of \$1,000, 25 years down the road, because that cash flow has a price of 6.5 per cent. So take your pick: supply and demand for bonds with quantity on the horizontal axis and dollar price on the vertical axis or per cent yield on the vertical axis. It's all the same. But lurking within the supply and demand schedules are investors' tastes for risk and return, and their expectations regarding the company's future prospects and ability to make the payments.

Does buying a bond make you wealthier? You already know the answer to that: no if the market is informationally efficient. In buying the bond, you are giving up \$815.86 now for a future cash flow whose value today is, to beat the point to death, \$815.86. And of course, the same goes for the seller. It makes sense. In buying and selling financial securities, we are not producing anything of value; we are only moving money across time. Then why do it? See section 3.

The term structure of interest

Suppose the risk-free, four-year spot rate of interest is 9.75 per cent. It is the rate that you might earn (might have earned many years ago) on an insured deposit, such as a guaranteed investment certificate or a treasury bond, if you were willing to tie up your money for the full four years. One dollar will grow to \$1.46344 in four years if interest is compounded semi-annually.

$$\$1 \left(1 + \frac{0.0975}{2} \right)^{2 \times 4} = \$1.46344$$

Now, what if the risk-free, two-year spot rate of interest was 9.5 per cent? You could invest or lend one dollar for two years and earn \$1.20397.

$$\$1 \left(1 + \frac{0.095}{2} \right)^{2 \times 2} = \$1.20397$$

Which would you choose: \$1.46344 in four years or \$1.20397 in two? The two investments can't be compared as is because their maturities differ. Anyone would naturally say, "It depends on what I think I could earn in years 3 and 4 if I were to invest for two years and then reinvest my principal and interest for another two years." But you can't know now what the two-year spot rate of interest will be in two year's time. The two-year investment rolled over for another two is risky, while the four-year investment is not. Cheat by assuming that you and all other investors are risk neutral, so that all that anyone cares about is earning as much money as possible at the end of four years, regardless of the risk. Now the investments can be compared.

If you thought that the two-year spot rate in years 3 and 4 would climb to 10.25 per cent from its current level of 9.5 per cent, then you would take the chance by investing short-term and rolling over the investment for another term because the expected total return is 1.47042, and that's more than the \$1.46344 you get by investing long term.

$$\$1 \left(1 + \frac{0.095}{2} \right)^{2 \times 2} \left(1 + \frac{0.1025}{2} \right)^{2 \times 2} = \$1.47042$$

What if you don't have a dollar to lend (darn those video lottery terminals)? No problem, borrow it at the four-year rate. At the end of four years you will owe \$1.46344 which you are expecting to be able to repay from a return of \$1.47042 for an expected profit of \$0.00697287 or

\$6.97 per \$1,000. If the future two-year interest rate ends up even higher, your smile will be bigger; on the other hand, you'll be in trouble if it doesn't climb high enough. What would your speculative strategy be if you expected the future two-year spot rate to be 9.8 per cent?

Every person acts on their expectations, but the fact remains that the short-term interest rate is 9.5 per cent and the long-term rate is 9.75 per cent. If these two interest rates are equilibrium interest rates, in the supply-equals-demand sense as always, then money is not being shifted from short-term to long-term or vice versa. Those who expect the future two-year spot rate to be "high" have invested short term; those who expect it to be "low" have invested or lent long term. And those dollars that lie at the two intersection of the supply and demand schedules represent investors who are indifferent between investing short term or long term. They must expect the short-term investment rolled over to yield the same return as the long-term investment. That break-even rate, the equilibrium expected future spot interest rate (the word *future* is redundant I think), is 10.0003 per cent, yes, just a hair over 10 per cent.

$$\left(1 + \frac{0.095}{2}\right)^{2 \times 2} \left(1 + \frac{x}{2}\right)^{2 \times 2} = \left(1 + \frac{0.0975}{2}\right)^{2 \times 4} \Rightarrow x = 0.100003$$

This is neat because the two spot interest rates, which exist and which we can transact with, tell us what the market expects will be the short-term spot rate in two year's time, in this case, 10.0003 per cent. You can now say why the long-term interest rate is higher than the short-term rate: the market expects the short-term rate to rise in the future. If the four-year rate was lower than the two-year rate, the market must be expecting the short-term rate to fall. This is known as the *Expectations Theorem of the Term Structure*. The term structure or yield curve refers to the shape of the graph of interest rates y-axis against their terms, two years or four. In our example, the yield curve is upward sloping, which is usual, but there is nothing unusual about a flat or even a downward-sloping term since the difference between the short and long rates is determined by expectations.

If a forward market for borrowing and lending exists, then the forward rate of interest must be equal to the equilibrium expected future spot rate implied by the expectations theorem

$$f^* = x = 0.100003$$

because if it wasn't, you and everyone else could earn a riskless profit using a strategy called arbitrage. Suppose, by some glitch, the actual two-year forward interest rate is 9.9 per cent instead of its equilibrium value 10.0003 per cent. You would borrow one dollar for two years at 9.5 per cent and invest it for four years at 9.75 per cent. You already know that at the end of four years you will receive \$1.46344. But how will you repay the loan after just two years? That's \$1.20397 you need to come up with. The answer is to borrow forward. This means that *today* you arrange a two-year loan of \$1.20397 to start in year 3, at the forward rate of 9.9 per cent.

What you have done is rolled over or refinanced your original loan at a rate that is fixed in advance. At the end of year 4, you will owe \$1.46065 for sure.

$$\$1 \left(1 + \frac{0.095}{2}\right)^{2 \times 2} \left(1 + \frac{0.099}{2}\right)^{2 \times 2} = \$1.46065$$

Earn \$1.4634 and repay \$1.46065 for an arbitrage profit of \$0.00279382 per dollar played or \$2,793.82 per \$1 million. Too good to be true? Pretty much, or at least too fleeting to be captured, because if you can do it, so can everyone else. The efficient markets hypothesis tells us that arbitrage profits should not exist! That is why the forward price is always equal the equilibrium expected future spot price: $f^* = E(r^*)$.

Intertemporal consumption optimum

Like most economic theories, utility theory tries to be one-size-fits-all. It can be used to describe someone's choice of how much to spend now and in the future just as easily as it can be used to describe their choice of goods, at any given time, as we did in chapter 2. Good x becomes C_0 , for current consumption or the dollars that you have today, and good y becomes C_1 , for future consumption or the dollars you will have next year. You can't consume dollars directly, of course, but the more you have, the more real stuff you can consume, and it really doesn't matter what that stuff is. The graph shows Lucy's intertemporal consumption optimum. Her tastes are represented by a natural log utility function, $u = \ln(C_0) + \ln(C_1)$. Dollars plotted against dollars may seem strange at first, but they are different things because one is dollars today and the other is dollars next year. (And speaking of one-size-fits all, there is nothing stopping us from extending the dates to year 2, 3, 4, right up to when Lucy kicks the proverbial bucket, and even beyond if she intends to leave an inheritance to her pet hamster, Francesca.) We'll assume that the future is certain despite the fact that it is not. This allows us to consider Lucy's choice as involving only her preference for consumption over time independent of her preferences for risk.

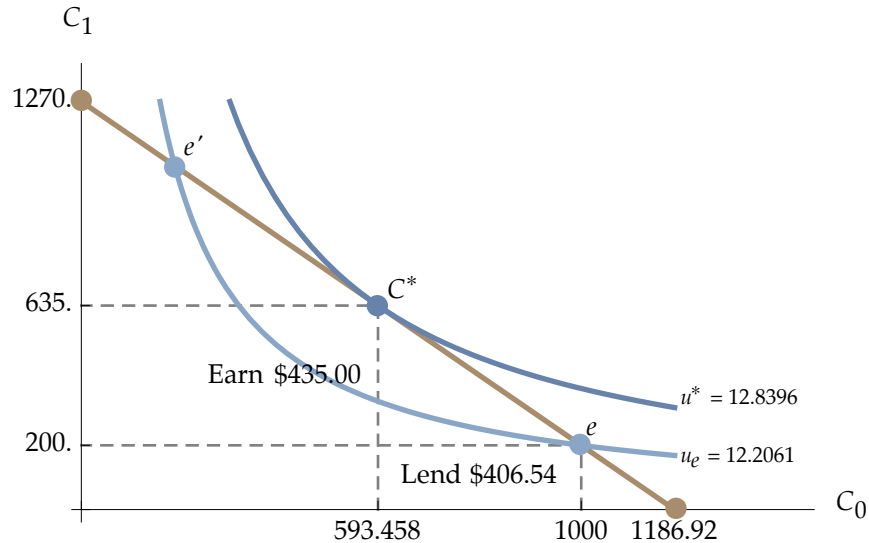
Lucy has an endowment of \$1,000 today and \$200 next year. At a seven per cent rate of interest, her wealth is \$1,186.92, which is where his budget line crosses the horizontal axis. The budget line, or wealth constraint since it is based on the market value of an endowment, is exactly the same as it was in note 2 except that it is now represents choices for consumption dollars rather than goods x and y , so

$$W_0 = p_x x + p_y y$$

becomes

$$W_0 = p_0 C_0 + p_1 C_1$$

and the only thing that needs to be done is to interpret p_0 and p_1 . If C_0 is current consumption measured in dollars, then p_0 must be 1 because the value of one dollar today is, well, one dollar. And if C_1 is consumption next year but wealth is being measured in dollars today, then p_1 must be the present value of one dollar, $\frac{1}{1+r}$, where r is the risk-free rate of interest. Substituting the



interpreted prices into the wealth constraint gives

$$W_0 = C_0 + \frac{C_1}{1+r} = \$1,000 + \frac{\$200}{1.07} = \$1,186.92$$

which says, very neatly, that your wealth is the present value of the cash flow stream to which you are entitled over the course of your life. Now rearrange the equation so that it is in the standard format for an equation for a line

$$C_1 = -(1+r)C_0 + (1+r)W_0 = -(1+r)C_0 + W_1$$

The price ratio is $1+r$, and it is the absolute value of the slope of the budget line as before. Consuming one more dollar today means foregoing a dollar plus interest later and vice versa. Another way to say this is that $1+r$ is the price of current consumption. $(1+r)W_0$ is the intercept on the C_1 or vertical axis, which is simply wealth measured in one year's time or W_1 .

Lucy maximizes her utility by choosing to consume \$593.458 today and \$635 next year. She gave up about \$406.55 this year, which at seven per cent, gives her an extra \$435 next year for a total of \$635. It's easy to see why Lucy is a lender and not a borrower. Her marginal rate of substitution at his endowment point is less than the price ratio

$$MRS_e < 1+r$$

With her endowment Lucy places a lower value on consuming one more dollar today than the market does. So she lends, moving up her budget line, and with each dollar lent, the dollars she has to consume now are becoming relatively scarcer, pushing up her marginal rate of substitution. She stops lending, having maximized her utility, when her marginal rate of substitution is exactly equal to $1+r$ because both she and everyone else values the next current dollar at exactly seven per cent. That is why the optimum is defined by

$$C^* : MRS = 1+r$$

A few things to note about the intertemporal consumption optimum. One is that the utility-maximizing choice is less extreme than the endowment in that \$593.458 and \$635 are closer to one another than are \$1,000 and \$200. This is called *consumption smoothing*, and it makes sense: most people prefer to have an income stream that is smooth (or, better still, smooth and increasing over time) than one that jumps around. You can see that consumption smoothing is a direct result of modelling with convex indifference curves. If Lucy's endowment was e' instead of e , she'd borrow to reach C^* . No matter what specific mathematical function is used to represent a person's tastes, if it is convex, the consumption optimum will always lie between e and e' ; the optimum is a weighted average of the two and therefore "smoother." The other thing to note is that Lucy is not wealthier for having chosen C^* , and she is not wealthier for having lent. She is better off because she prefers C^* to e . But C^* is worth \$1,186.92 just as e and every other point on the budget line is. Despite Scotiabank's slogan, *you're richer than you think*, a capital market by itself does not create wealth. It gives us choice, and that is a good thing. For wealth creation, something else is needed.

The last thing to note is subtle. Everyone else has presumably borrowed or lent to reach their own consumption optimums, so everyone must be evaluating their next dollar to spend (borrow to take that vacation?) at the market rate of interest. Utility theory implies that all of us implicitly agree what the time value of money is. If that is true in real life—an empirical question—then governments can justify using market interest rates as discount factors in cost-benefit analysis. If not, those analyses don't mean all that much.

Things to think about

Can you explain why and how Lucy would choose the same consumption optimum if her endowment was e' rather than e ?

Real investment and exchange

To make herself wealthier, Lucy must produce something that is worth more than it costs. Pretty basic, eh? Suppose that she can invest in any of the five projects shown in the table. Each is a real investment. The projects are risk-free in keeping with framework of looking at choices

over time in the absence of risk. If Lucy takes project C, for example, she would invest \$200 out of her current endowment of \$1,000 to earn a sure \$202 in one year. That's a return on investment of one per cent. But she'd never invest in C because she can always earn seven per cent by lending her money. The value today at seven per cent of receiving \$202 in a year is only about \$189. That means that the net present value of project C is -\$11. By taking project C, Lucy make herself poorer by \$11 because she'd be giving up \$200 for something that is worth only \$189.

$$NPV = -Cost + \frac{Income}{1+r}$$


$$NPV(C) = -\$200 + \frac{\$202}{1.07} = -\$11.2150$$

Being greedy, Lucy will take only those projects that increase her wealth, and that is the same as taking projects whose cash flows have positive net present value. She will invest \$400, \$190 in project A and and \$210 in B—both of which have a crazy high return, but the point is

Project	Cost	Income	ROI	NPV at 7%
A	\$190	\$305.90	61%	\$95.89
B	\$210	\$390.60	86%	\$155.05
C	\$200	\$202.00	1%	-\$11.22
D	\$180	\$185.40	3%	-\$6.73
E	\$220	\$178.20	-19%	-\$53.46

that their returns are greater than seven per cent—for a combined net present value of \$250.935. Her wealth increases by \$250.935 to \$1,437.85 from \$1,186.92 if she did nothing.

Lucy will also be better off if she increases her wealth this way. That seems obvious but best to check. You can do that by adding the five projects to Lucy's consumption diagram. The



projects are shown as green line segments, ordered from highest to lowest rate of return, starting at her endowment. The first is project B with a rate of return of 86 per cent. The cost of \$210 (\$1,000 left to \$790 on the horizontal axis) and income of \$390.60 (\$200 up to \$591) on the vertical axis gives the line segment a absolute slope of 0.86. Then comes A with a return of 61 per cent, D at three percent, and so on. Connecting the segments end to end creates a piece-wise curve that is concave to the origin, and is called an *investment opportunity schedule*. If the projects were infinitely divisible, the investment opportunity schedule would be smooth and look just like a production possibility curve. And that's because it is a production possibility curve, except that it represents dollars today being transformed into dollars tomorrow rather than combinations of goods.

To spot which projects have positive net present values, shift the budget line parallel and upwards until it touches the investment opportunity schedule at just one point. This happens at point P^* (which would be a tangency if the investment opportunity schedule were smooth). Projects to the right of P^* have positive net present values and should be undertaken, and those to the left should be passed over because they do not. That makes sense because the budget line has a slope equal to one plus the rate of interest (ignoring the negative sign), and the slope of each line segment representing each project has a slope equal to one plus the project's return on investment. The projects to the right of P^* earn more than the rate of interest, which in a two-period world means they have positive net present values. The horizontal intercept of the pushed-out budget line, \$1,438, is Lucy's maximized wealth. The distance, \$251, from the old intercept is the combined net present values of projects A and B.

Lucy's investment choice has nothing to do with her taste for consumption over time. It is driven by greed, and anyone who is greedy and who has the same opportunities would also choose A and B. That's a good thing because if these investment opportunities belonged to a corporation, there would be no disagreement among shareholders as to which projects should be undertaken. Greed and a shared gauge for profitability—the market rate of interest—gets us around Arrow's impossibility result. It allows shareholders to hire professional managers to do the project picking for them, directing them with one simple rule: make us richer. (In the uncertainty of the real world, making sure that managers actually do their job, and that they do not rob shareholders blind is another story.)

Having made herself as wealthy as can be, Lucy then expresses her preference by choosing the cash flow stream she wants to consume. That is \$719 today and \$769 next year, shown at C^* . After investing \$400 in projects A and B, Lucy has only \$600 remaining out of her current endowment of \$1,000. Not a problem. With a capital market, she simply borrows \$119 at seven per cent to bring her cash on hand to \$719. So there are two reasons that Lucy's utility is higher than it would be if she consumed her endowment: one is the increase in wealth from profitable investments; the other is the choice cash flow streams to consume provided by the ability to borrow or lend (otherwise she'd have to consume P^* rather than C^*). It is the existence of a market

rate of interest, at which everyone can borrow or lend, that allows Lucy to separate her investment decision and your consumption choice. One does not interfere with the other, so Lucy can maximize her wealth, without concern about what she wants to consume, and in that way, boost her utility. The separation of consumption and investment is known as *Fisher Separation* after economist Irving Fisher. Fisher Separation can be summarized like this:

$$P^* : MRT = 1 + r$$

$$C^* : MRS = 1 + r$$

MRT is the marginal rate of transformation or the absolute value of the slope of the investment opportunity schedule, just as it is for production possibility curves. But here *MRT* is interpreted as one plus the marginal return on investment. The condition $MRT = 1 + r$, says to keep investing until the last penny invested earns exactly the rate of interest; in other words, chose all projects with positive net present values.¹⁵ You are already familiar with the second condition, $MRS = 1 + r$.

Things to think about

1. Point C is the cash flow from section 3 that Lucy would consume if she did not have investment opportunities. Notice that Lucy consumes both more today and next year at C^* than she does at C. What does this tell you about current and future consumption dollars?
2. Is Lucy smoothing her consumption by choosing C^* ?
3. Can you eyeball where Lucy would consume if she could invest but there was no capital market?
4. Lucy makes herself richer with positive net present value projects, but where do positive net present values come from?

The interest rate

Up to this point we have been looking at intertemporal consumption in partial equilibrium. In section 3, Lucy achieved her preferred consumption stream by lending part of her current endowment given that the interest rate was seven per cent. In section 4, she achieved her preferred consumption stream by investing in projects with positive net present values and borrowing, again, given a seven per cent interest rate.

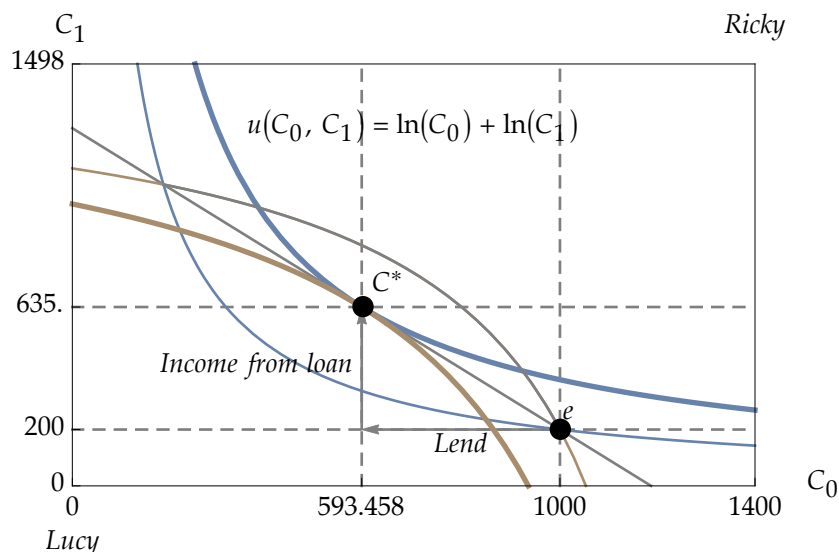
How does the interest rate come to be seven per cent? Equilibrium in the market for money. The equilibrium interest rate is the one that equates the number of dollars potential lenders are willing to lend with the number of dollars that potential borrowers are willing to borrow. You can guess that the supply and demand for dollars today depends upon everyone's endow-

¹⁵ Zero present values are okay too.

ment (think of this as the distribution of income both now and in the future), their tastes for consumption ('high' or 'low' marginal rate of substitution at the endowment), and their individual real investment opportunities. And all three of these variables interact.

People whose current endowments are small compared to their future endowments will have biggish marginal rates of substitution, and they will generally want to increase their current consumption. If these same people comprise a big share of the aggregate current endowment of the economy, the borrowing that they do from those whose current endowments are so comparatively small will occur at a higher interest rate than it otherwise would. One example of this is an economy dominated by people facing lean times now compared to how they'll be doing next year. But we can't talk about endowments without talking about tastes. You know that two people with the same endowment will have different marginal rates of substitution if their tastes differ. That is, after all, what it means to have different tastes. It is also possible that someone with a small current endowment relative to their future endowment has a smaller marginal rate of substitution than someone with a bigger current endowment relative to their future consumption. (Confused? Draw some indifference curves.) The person with the smaller relative current endowment values more current consumption less than the person with the bigger relative current endowment because differences in their preferences (I don't have much now, but it doesn't matter because I really don't value consuming more now). So, endowments and tastes necessarily interact.

Here is an example of the determination of a seven per cent rate of interest in general equilibrium. Including real investment opportunities in the Edgeworth box is a little advanced for this course, so the equilibrium is presented for a pure exchange economy. This would then correspond to Lucy's partial equilibrium in section 3. The aggregate current endowment is \$1,400, of which \$1,000 belongs to Lucy and \$400 to Ricky. The aggregate future endowment is \$1,498 with \$200 going to Lucy and \$1,298 to Ricky. Both Lucy and Ricky have log utility functions. At



the endowment point, e , Lucy's marginal rate of substitution is smaller than Ricky's, so Lucy lends to Ricky. The loan that is Pareto optimal is \$406.54 with a repayment of \$435 in a year. At the optimum, C^* , Lucy's marginal rate of substitution is equal to Ricky's is equal to 1.07. That implies an interest rate of seven per cent.

Notice that the growth rate of the economy is also seven per cent ($\$1,498 \div \$1,400 = 1.07$)? This is the exception and not the rule. It happens in this example because both Lucy and Ricky are assumed to have log tastes

$$u = \ln(C_0) + \ln(C_1)$$

which is a function with special properties (I'm not allowed to tell you exactly what they are). The same individual endowments with any other utility function would result in a different interest rate. For example, using this variation of the log function

$$u = \frac{4}{10} \ln(C_0) + \frac{6}{10} \ln(C_1) \text{ for Lucy}$$

and

$$u = \frac{6}{10} \ln(C_0) + \frac{4}{10} \ln(C_1) \text{ for Ricky}$$

would result in an interest rate of 13 per cent. But if Lucy's endowment were changed to (\$900, \$600) and Ricky's to (\$500, \$848), keeping the aggregate endowment the same, the interest rate would be about five per cent. Can you relate these different rates to the earlier discussion on the interplay of endowments and tastes?

Things to think about

Are Lucy and Ricky's consumption streams smoother at the Pareto optimum or at the endowment? In equilibrium, are both of them consuming more next year than this year?

Can the interest rate be negative? If so, roughly what would have to be the conditions in the economy? Hint: draw an Edgeworth box and a Pareto optimum.

The Life Cycle Theory of Consumption and Savings

Consumption smoothing is one result of utility maximization. Both Lucy and Ricky smoothed their consumption. Life has its ups and downs but we prefer to take the edge off, and we can do that if there is a capital market where we can borrow and lend. Does utility theory predict that you will smooth your consumption if you live to be 88? Yes, all of the results for the two-period model still hold. This means that when you are just starting out you will probably be a borrower, or net dissaver to use the jargon, because your income will be less than the smoothed amount you prefer. By mid life, your income will have grown, and you will hopefully

be a net saver (lender). Income stops coming in or shrinks to a pension when you retire, but utility theory predicts that you will still want to maintain the lifestyle to which you are accustomed, so you will likely become a net dissaver again and stay that way until you kick the proverbial bucket. This is the essence of the *Life Cycle Theory of Consumption and Savings*: people want to be able to spend smoothly over their lifetimes—think of it as maintaining a certain standard of living—and they adjust their savings, cycling through borrowing and lending, in order to achieve this. Let’s look at the special case of perfect consumption smoothing, where the optimal amount consumed each year is constant, and then see how it relates to consumption smoothing in general, that is, smooth by increasing every year.

A guy named Nigel will be the player in this example. Nigel is 20 years old. He just landed his first job, with a starting salary of \$100,000, which will grow by 1.5 per cent every year right up until he retires at age 70. He owns a \$300,000 house, a \$75,000 investment portfolio, and a

Nigel in a nutshell

Preferences	homothetic
Age now	20
Age at retirement	70
Age at death	88
Interest rate	0%
Annual income	\$100,000
Growth rate of income	1.5%
<i>Assets</i>	
House	\$300,000
Portfolio	\$75,000
Beatles records	\$5,000

collection of Beatles records worth \$5,000. Not bad for a 20-year old, eh? One other thing about Nigel, and he knows this himself, is that he will live to be 88. Not the kind of thing most of us would want to know, but we have assumed a world of perfect certainty.

What amount will Nigel consume each year if he maximizes his utility? This optimal consumption is also known as *permanent income* because, by spending that amount each year, Nigel will have drawn his wealth down to zero by the time he dies. The present value of the optimal consumption stream is exactly equal to his wealth, as it must be because you saw in the two-period case that Lucy’s optimal consumption lies on the budget line; it satisfies the wealth con-

straint. So, to compute permanent income, we need an interest rate in order to figure out Nigel's wealth, and we need to know Nigel's utility function. There are two conditions for permanent income to be constant, that is, perfect smoothing or exactly the same from year to year: one is that the interest rate must be zero, so that there is no incentive to push a few extra dollars into the future; and the other is that Nigel's utility function must be such that the enjoyment he gets from an extra dollar doesn't depend on when he gets it. The trusty log utility function does the trick as long as the coefficients a and b are equal.

$$u = a \ln(C_0) + b \ln(C_1), \quad a = b$$

Any utility function that is symmetric in this way would also work.

Now for Nigel's wealth. His wealth is the market value of his assets plus the present value of the income he will earn over his lifetime.

$$\text{Wealth} = \text{Market value of assets plus the present value of income stream}$$

The combined value of his house, portfolio, and record collection is \$380,000, so his assets are taken care of. His income is an annuity that will grow 1.5 per cent a year. He will receive \$100,000 in one year, \$101,500 in two, \$103,022 a year after that, and when he turns 70, his 50th and last pay cheque will be \$207,413. The present value of his income stream is

$$PV(\text{income}) = \$100,000 \left[\frac{(1.015)^0}{(1+r)^1} + \frac{(1.015)^1}{(1+r)^2} + \dots + \frac{(1.015)^{N-1}}{(1+r)^N} \right]$$

Let m be Nigel's \$100,000 starting salary, g be the 1.5 per cent the growth rate, and N be the number of years he will work. The present value can then be written as

$$PV(\text{income}) = m \left[\frac{(1+g)^0}{(1+r)^1} + \frac{(1+g)^1}{(1+r)^2} + \dots + \frac{(1+g)^{N-1}}{(1+r)^N} \right]$$

You can show that the series simplifies to

$$PV(\text{income}) = m \left[\frac{1 - \left(\frac{1+g}{1+r} \right)^N}{r-g} \right]$$

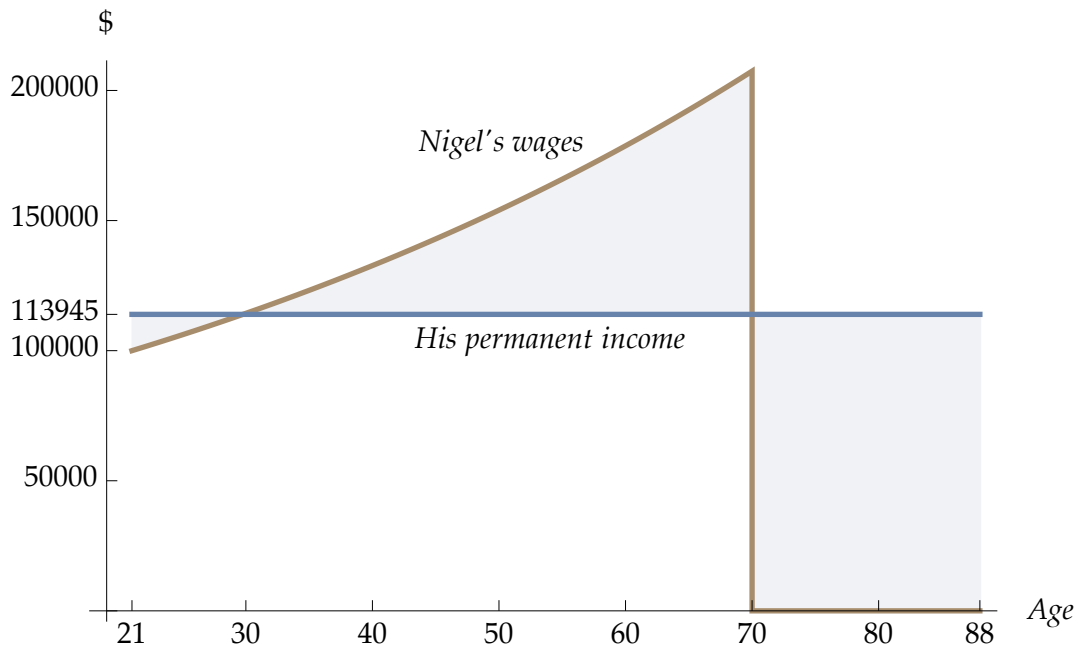
which can be used to compute the present value of any annuity that grows at a fixed rate.

If the interest rate is zero, the equation for the present value of his income is even simpler

$$PV(\text{income}) = m \frac{(1+g)^N - 1}{g} = \$100,000 \frac{1.015^{50} - 1}{0.015} = \$7,368,280$$

Nigel's wealth is \$7,748,280, \$380,000 from his assets and the bulk, almost \$7.4 million, from his salary. He will spread this evenly over the 68 years he has left to live. That means he will consume \$113,945 each year. If the interest rate was positive, Nigel's permanent income would be increasing over the years, although he'd still be smoothing his consumption.¹⁶

The first thing to note about Nigel's permanent income of \$113,945 is that it is almost \$14 thousand bigger than his first few pay cheques. As a matter of fact, he will have to borrow or sell some of his assets to make up the difference until he is 29. His life cycle of consumption and savings is shown in the chart. The timeline starts when Nigel is 21 ($t = 1$) rather than 20 ($t = 0$)



because we assumed that cash flow occurs at the end of the year and that is when it is consumed. Can you spot the times when he is a net saver and dissaver?

The other, and more important, thing to note about the permanent income hypothesis is that if Nigel's permanent income changes, he will increase or decrease his consumption by exactly the change in permanent income. If he finds ten dollars lying on the sidewalk right now (happens to me all the time), he will treat himself by spending about 15 cents more each year (can you figure that out?). He will not run down the street and spend the \$10 at Double Pizza.

¹⁶ That calculation is left for a more advanced course.

NOTE 10

Markets for Risk

A capital market not only allows us to shift consumption over time, it allows us to shift risk. Those who are willing to take on more risk buy it from those who want to get rid of some. Risk-averse buyers will demand a higher rate of return for taking bigger risks, and risk-averse sellers decide how much return they are willing to give up to shed risk. Markets for risk, such as those for bonds, stocks, and insurance, are no different from any other in that they benefit society by providing choice.

States and dates

Suppose that there is no risk now. Lucy and Ricky know everything about the present moment: their bank balances, pairs of clean socks in their dresser drawers, whether the salami in their refrigerators is still fresh, and the condition of the shingles on their roofs. You know that you can't know everything even if there is no uncertainty, but there is an awful lot that is knowable. That is period 0 from the intertemporal choice model in chapter 9. The future, however, is uncertain, so consumption in period 1 must now be treated differently (likewise, periods 2, 3, 4 and so on, if they are part of the model). So, risk and time are intertwined. In a world without risk, Lucy's period 0 endowment might be \$168 and her period 1 endowment \$208.80. If she chooses to consume her endowment, she knows right now that \$208.80 will be waiting for her next year for sure.

Not so in a risky world. Lucy's current endowment is still a sure \$168, but her future endowment might be any amount. Probably not \$1 billion and hopefully more than \$3. How is risk to be brought into the model? One way is to assume that future, risky consumption is a random variable that behaves according to a known probability distribution. Probability distributions are made up of two things: outcomes (consumption or income in our case) and the probabilities of those outcomes occurring. Lucy, for example, might face a 0.4 chance of ending up with \$72 next year and a 0.6 chance of \$300. There are only two outcomes, but that's arbitrary. We could have 11 or 546. In real life I don't think we can truly know what the outcomes are or how many are possible, and we would struggle to attach probabilities to them. Economists refer to the circumstances in which particular outcomes occur as *states* or *states of nature*. If

there are two, then they might be called state a and state b (like Thing 1 and Thing 2 from Dr. Seuss).

States are mutually exclusive: rain or shine, war or peace, boom or bust. State a has a 0.4 chance of occurring, and if it does, Lucy will have \$72; state b has a 0.6 chance of occurring, in which case Lucy will have \$300. In a riskless model, Lucy's endowment is written as

$$e = (C_0, C_1) = (168, 208.80)$$

but in a *time-state preference model*, her endowment or any consumption bundle is written as

$$e = (C_0, C_1) = (C_0, (\pi_a, C_a; \pi_b, C_b)) = (168, (0.4, 72; 0.6, 300))$$

where C_1 is now a random variable (sometimes this is indicated by putting a tilde over the variable, as in \tilde{C}_1); π_a and π_b are the probabilities of having consumption C_a in state a or C_b in b . It is important to remember that Lucy will never have \$72 *and* \$300 next year. She knows now that she will have either \$72 *or* \$300 depending on what happens. *Que sera sera*.

Lucy's wealth in a riskless world is the present value of her endowment computed at whatever the risk-free rate of interest happens to be

$$W_0 = p_0 C_0 + p_1 C_1 = C_0 + \frac{C_1}{1+r} = \$168 + \frac{\$208.80}{1+r}$$

In a risky world, her wealth is still the present value of her endowment but it is expected future consumption, $E(C_1)$, that must be discounted at a higher rate, k , to take account of the risk

$$W_0 = C_0 + \frac{E(C_1)}{1+k} = \$168 + \frac{0.4 \times \$72 + 0.6 \times \$300}{1+k} = \$168 + \frac{\$208.80}{1+k}, \quad k > r$$

In the *time-state preference model*, it is convenient to write wealth in terms of prices rather than interest rates or discount factors. We can figure out the rates later.

$$W_0 = p_0 C_0 + p_1 C_1 = p_0 C_0 + p_a C_a + p_b C_b$$

Instead of having one price for future consumption, p_1 , there is a price today for consumption in each future state of nature, p_a and p_b . These are called *pure state prices*, *primitive security prices*, or *Arrow-Debreu prices*, after the two economists who, independently, developed the model. Take a moment to think about what a pure state price is. It is the price or value today of having exactly one dollar in a particular future state and nothing in all of the other states. That means that you can think of a pure state price as a present value factor that is attached to a particular future event or, even easier, you can think of it as the market price of a gamble that pays one dollar if a particular event occurs and nothing if anything else occurs.

p_a is the price today of a gamble that pays \$1 in state a and \$0 in state b

p_b is the price today of a gamble that pays \$1 in state b and \$0 in state a

While pure state prices are theoretical—you can't look them up anywhere—they do exist in real markets various guises, some of which we will discuss in class. For now, here's an example: the price of a lottery ticket that pays \$1 if you earn an A or better in my course and \$0 if you don't.

In the economy we'll be discussing in the rest of this note, $p_a = 0.374166$ and $p_b = 0.614817$. I'm giving you the values because calculating them is a bit too advanced for this course. But their interpretation is more important than their computation: A dollar next year in state a is worth about 37 cents now and a dollar in state b is about 61 cents. You'll be able to use them to calculate Lucy and Ricky's wealth, as well as the equilibrium risk-free rate of interest and the expected return on the stock market index.

Description of the economy


Information about a two-period Canadian economy is given in the table. The states and dates are in the first panel. The state probabilities and pure state prices are carried over from the previous section. The state prices are in the rows labeled time period 1 to associate them with states a and b , but they are time 0 or *now* prices. The middle panel shows Lucy's endowment, and Ricky has been brought into the story. Since Lucy and Ricky are Canada's only inhabitants, the country is endowed with \$420 today and will grow to \$480 next year if state a occurs or decline to

Contingencies				Endowments			Securities	
Date (t)	State (s)	Prob. (π)	P	Lucy	Ricky	Canada	M	F
0	-	1	1	168	252	420	P	P
1	a	0.4	0.374166	72	408	480	4.8	1
1	b	0.6	0.614817	300	100	400	4	1

\$400 if it's b . You can think of the aggregate endowment as Canada's gross domestic product. The last panel introduces two financial securities. M is a risky stock that pays a dividend of \$4.80 in state a and \$4 in state b . M 's price is P_M . M is like a stock market index or super stock that includes all firms in the economy, so it pays out the aggregate endowment of the economy. M 's payoff is one one-hundredth the aggregate endowment, so the supply of shares of M must be 100. If M 's payoffs were \$2.40 and \$2, there would be 200 outstanding because the payoffs have to be proportional to the aggregate endowment. Anyone holding nothing but shares in M is bearing the same risk as the economy as a whole. F is a riskless security because it pays \$1 no matter what happens. It is like a pure discount bond with price P_F because it does not pay interest, and its payoff is not related to the aggregate income of the economy. You can think of F as a Treasury bill with a face value of \$1. Since F is debt security, its net supply must be zero—for

every dollar borrowed there is a dollar lent. By figuring out the equilibrium price of M and F , we will know the expected return on risky stock and the risk-free rate of interest.

The figure shows the economy represented in an Edgeworth box. It includes the Pareto optimum, $C^* = (C_a, C_b) = (238.445, 198.705)$. Notice there is no axis for C_0 . If it was drawn, the C_0 axis would come straight out into your face from Lucy's origin, and the Edgeworth box would be three dimensional. So $e_0 = 168 = C_0^*$ could be shown on a third axis but it isn't necessary because the time-state preference model describes how Lucy and Ricky trade their future endow-



ment now, e , to end up at the Pareto optimum, C^* , and that is also why optimal current consumption is fixed at 168. What you are looking at in two dimensions is next year or period 1, and the axes are the two outcomes. Even though the probabilities are not shown anywhere (you'll have to remember them or look back at the table), every point in the Edgeworth box is a distribution of future consumption or income. Endowment point e is the distribution (0.4, \$72, 0.6, \$300) for Lucy, and by subtracting from the aggregate endowment of the economy, (0.4, \$408; 0.6, \$100) for Ricky. Canada is expected to grow by 2.9 per cent

$$\text{Expected rate of growth} = \frac{E(C_1)}{C_0} - 1 = \frac{0.4 \times 480 + 0.6 \times 400}{420} - 1 = 0.0285714$$

Pareto optimum

The Pareto optimum is a general equilibrium and is determined in the usual way: Lucy and Ricky's marginal rates of substitution over uncertain state consumption are equal, and this

happens at a value of 0.608581, which is the implied equilibrium price ratio. One dollar of consumption in state a is worth 0.608581 dollars of consumption in state b .

$$C^* : MRS_{Lucy} = MRS_{Ricky} = \frac{p_a}{p_b} = \frac{0.374166}{0.614817} = 0.608581$$

Notice that you cannot compute the pure state prices just from the implied price ratio. But having given you those values, $p_a = 0.374166$ and $p_b = 0.614817$, you can compute Lucy's wealth to be \$379.385, Ricky's as \$466.141, making \$845.526 the wealth of the nation. Why is a dollar of consumption worth so much less in state a than in b ? You can find two reasons in the table.

If wealth is the present value of a person's income, then the price or market value of a security, in the same way, is the present value of the income it provides. The equilibrium price of risky stock M is about \$4.26 for an expected return of 1.52 per cent. The price of safe bond F is \$0.9889, making the risk-free rate of interest 1.11 per cent. Letting X_s be the state income or payoff of any security, then

$$\begin{aligned} P_M &= p_a X_a + p_b X_b = 0.374166 \times \$4.80 + 0.614817 \times \$4 = \$4.25526 \\ E(R_M) &= \frac{E(X)}{P_M} - 1 = \frac{\pi_a X_a + \pi_b X_b}{P_M} = \frac{0.4 \times \$4.80 + 0.6 \times \$4}{\$4.25526} = 0.0152132 \\ P_F &= p_a X_a + p_b X_b = (p_a + p_b) X = 0.374166 \times \$1 + 0.614817 \times \$1 = \$0.988983 \\ r_f &= \frac{X}{P_f} - 1 = \frac{\$1}{\$0.988983} - 1 = 0.0111399 \end{aligned}$$

The difference between the expected return on the stock market index and risk-free bonds is called the *market risk premium*. It is the extra return that the market compensates you with for bearing risk. The risk premium is only 0.4 per cent in this economy. That's tiny. A market risk premium in the range of three to six per cent was typical during the 20th century. Tweak the parameters of the time-state preference model all you want and you won't be able to crank out realistic risk premiums unless you make Lucy and Ricky pathologically risk averse. This anomaly, that real life risk premiums are so different from what theory predicts, is known as *the equity premium puzzle*.

Portfolios and trade

C^* is an equilibrium because, mutually, Lucy and Ricky prefer it to any other distribution of future consumption. At any other point in the Edgeworth box, their marginal rates of substitution would differ from one another, and they would then do something to get to C^* . How do they get to C^* if, for whatever reason, they are not there? The government could try to reallocate their endowments. So could a benevolent monarch or dictator. But we haven't introduced the state (as in government) in this model. Moving around the Edgeworth box means that Lucy and

Ricky must be trading future distributions of consumption. Since the only thing to trade in this economy is the two securities, and there is no production, every distribution must be a *portfolio* of shares of risky stock M and safe bond F . You can see this by considering two references. One is the main diagonal, drawn as a solid green line, and the other is the two dashed green lines, one leaving Lucy's origin at 45 degrees and the other leaving Ricky's at the same angle. The main diagonal is the set of all portfolios that are comprised of nothing but shares of M because every distribution along the line pays proportionally to the economy as a whole, \$4 in state b for every \$4.80 in a . At point g , for example, Lucy's portfolio pays \$120 in a and \$100 in b , so she must be holding 25 shares of M and no bonds. Ricky owns the other 75 because all 100 shares must be accounted for. Neither owns bonds (is a lender) or has issued bonds (is a borrower) because receiving bond income or paying down bond debt would throw the portfolio payoffs off the \$4 to \$4.80 ratio, and that is not possible if g is on the main diagonal. Both are invested solely in the market index and are bearing the same risk as Canada as a whole. Now consider point h , which is on the 45-degree line leaving Ricky's origin. A 45-degree line has a slope of 1 and because it passes through $(0,0)$, all of the state a consumption payoffs are equal to all the state b payoffs along the line. It is a *certainty line* for the person situated at its origin. Ricky's portfolio h is riskless because it pays him \$160 in either state. He must own nothing but bonds, 160 F bonds to be exact because the F bonds have a one dollar face value. Lucy's portfolio h cannot be riskless because if Canada's future income is risky and Ricky is not bearing any of that risk, someone else must be. A less slick way of seeing that Lucy's portfolio h is risky is that it pays off different amounts in the two states: \$240 in b to \$320 in a , so it cannot be comprised of just safe bonds. And it cannot be comprised of only M shares because the payoffs are not in the ratio of \$4 to \$4.80 (h is not on the main diagonal). This means that Lucy's portfolio h must be comprised of both bonds and stocks. And if Ricky is a lender with portfolio h , Lucy must have borrowed from him because the net supply of debt is zero. Lucy owns -160 bonds.¹⁷ She owns all 100 shares of M but had to borrow to get them.

Points e and C^* are the two distributions that are important to the time-state preference model. Working out the composition of these portfolios is as simple as solving two equations in two unknowns. At e , Lucy receives \$72 in state a or \$300 in state b .

$$e_{Lucy} \begin{cases} C_a = 72 = 4.80 \cdot q_M + 1 \cdot q_F \\ C_b = 300 = 4 \cdot q_M + 1 \cdot q_F \end{cases} \rightarrow q_M = -285, q_F = 1,440$$

Those amounts, \$72 and \$300, have to come from the payoffs of M and F . Lucy is short 285 shares. She borrowed those from Ricky. This is a risky loan because the amount she'll have to repay Ricky depends on which state prevails. If a then she has to pay him \$1,368 but if b , only

¹⁷ Owning a positive quantity of a security is sometimes called a *long position*. Owning a negative number is called a *short position* or *short selling*, and it is the same as having borrowed.

\$1,140. She used the loan to buy safe 1,440 bonds. Kind of strange, isn't it? Taking out a risky loan to buy something safe? Strange but not impossible; this is a theoretical model, and e was arbitrary. Ricky is long 385 risky shares and short 1,440 safe bonds. No need to solve equations; it's just bookkeeping. I'll let you confirm that at the optimum Lucy holds about 49.68 shares, Ricky holds 50.32, and there is no debt outstanding (C^* is on the Market line in the example). Could you have gotten a sense of the composition of Lucy's e portfolio just by looking at its location in the Edgeworth box? Move e around. What happens to its composition if it is moved closed to Lucy's certainty line? Beyond that, as it is moved closer to the Market line? Beyond the Market line to the lower-right of the lower-right?

Here's some homework for you. Pick any point in the Edgeworth box and see if you can figure out the composition of Lucy and Ricky's portfolios at that point using the payoffs of securities M and F . You'll find that you can because the following conditions are met: there are at least as many securities as there are states of nature (M and F , a and b); the securities are not perfect substitutes (you cannot create the payoffs of M by multiplying the payoffs of F by some constant); and negative quantities are allowed (short selling). These three conditions are necessary for a *complete market*, which is one where any distribution of payoffs can be created by combining existing securities into portfolios. The upshot of this is that Lucy and Ricky can trade to any point in the Edgeworth box that they prefer, and that includes C^* . Choice is good.

Review

states of nature

pure state price, primitive security price or Arrow-Debreu price

time-state preference model

complete market